arXiv:submit/8429108 [cs.LG] 25 Nov 2026

Guiding LLMs at Inference Time via Bregman Tilting

Jacob L. Block* Aryan Mokhtari*† Sanjay Shakkottai*

*UT Austin †Google Research

Abstract

We study inference-time guidance for large language models (LLMs) in iterative problem solving, where the LLM generates an attempt, receives text feedback from the environment, and tries again. While input prompting, which puts prior attempts and feedback in context, is a standard approach, it gives no control over how the LLM uses the feedback in its next attempt. Thus, we introduce *Bregman tilting*, which complements any input-prompting approach by directly modifying the LLM's next-token distribution. At each decoding step, it scores each token by its change in log-probability when the feedback is added to a prompt that contains the prior attempt, and updates the distribution via a mirror ascent step on the expected score, regularized by a Tsallis-entropy-based family of Bregman divergences indexed by $\alpha \in [0, 1]$. At $\alpha = 1$ this recovers the exponential tilt of prior methods, while $\alpha < 1$ changes the divergence geometry to discount the influence of unlikely tokens under input prompting. The resulting update damps each token's tilt by its input-prompting likelihood, and it locally improves the true correctness probability whenever the score positively covaries with correctness under the input-prompting distribution sharpened to a lower temperature. Experiments on coding and math show that Bregman tilting outperforms input prompting in every setting, while $\alpha = 0$ performs best, beating the exponential tilt at $\alpha = 1$ and highlighting the value of our framework as a whole.

Corresponding author: jblock@utexas.edu

1 Introduction

A common paradigm for solving reasoning problems with large language models (LLMs) is iterative generation at test time: the model generates an attempt, receives feedback on it in the form of text, and generates again. The feedback can come from the environment, such as a compiler and unit tests [CLSZ24] or a proof verifier [Lin+26], from a human user [BHPSRK26], from another LLM [OIWGS24], or from the solver itself [Mad+23]. In each case, the goal is for the LLM to produce a better attempt at each round given the text history of its previous attempts and their feedback.

We study a training-free, inference-time setting, where the underlying model parameters are kept frozen. In this setting, the standard approach is *input prompting*, which places the previous attempts and their feedback in the prompt given to the LLM at the next round and resamples [Mad+23; SCGNY23; Ven+26; PWSNHW25]. While this is simple to implement and can improve performance round over round, input prompting updates the distribution over rollouts only through the LLM's context. This treats the LLM as a black-box sampler conditioned on its prompt and gives no control over how strongly or faithfully the LLM incorporates the feedback into its next attempt. We address this issue with *Bregman tilting*, a principled inference-time framework that complements a given input-prompting approach by using text feedback to explicitly modify the LLM's output distribution.

Contributions. After each round, we search for a distribution over the next round's rollouts which improves the probability of producing an acceptable response (Section 2). We propose a per-token feedback score that serves as a proxy for this intractable objective, and a novel guidance method that updates the LLM's next-token distribution to increase the expected score. We detail these contributions below.

At any prefix, the optimal sampler maximizes the probability of generating a token which leads to an acceptable rollout. However, this objective is infeasible without access to a verifier of correctness. Instead, we define a per-token feedback score as a surrogate, computed as the change in the log-probability of each candidate next token when the feedback is added to the prompt (7). It compares the LLM's next-token distribution under the prompt that shows the previous attempts and their feedback to its distribution under a reference prompt that shows the same previous attempts but without the feedback. Similar contrastive scores appear in contrastive decoding [LHFLEHZL23], context-aware decoding [SHLTZY24], classifier-free guidance for LLMs [SSFLB24], and self-distillation [Hüb+26], which contrast against a weaker LLM or the same LLM prompted without some added context. In iterative refinement, however, a reference prompt that omits the previous attempt can reward tokens for agreeing with that attempt regardless of correctness. Because the two prompts differ only in the feedback, our score isolates its effect, while the question-only reference prompt construction of prior work can fail to outperform vanilla input prompting (Table 2).

We then guide the LLM's output distribution at inference time with our proposed Bregman tilting framework. At each decoding step, Bregman tilting takes a single mirror ascent step on the objective of maximizing the expected score over next tokens (11). This regularizes the update using the Bregman divergence to the input-prompting LLM's next-token distribution, penalizing deviations in terms of the geometry induced by the mirror map.

We consider a family of mirror maps parameterized by $\alpha \in [0, 1]$ based on the Tsallis entropy which smoothly interpolates between the log-barrier at $\alpha = 0$ and the negative Shannon entropy at $\alpha = 1$.

At $\alpha = 1$, the Bregman divergence is the KL divergence, and the update exponentially tilts the input-prompting LLM with the feedback score, a common update form for inference-time guidance [SHLTZY24; SSFLB24; HS22]. However, setting $\alpha < 1$ leads to a generalized tilting update which discounts the effects of tokens which the input-prompting LLM considers unlikely.

We show that, in the limit as a token's probability under the input-prompting distribution goes to zero, Bregman tilting changes its log-probability by the exponential tilt's amount, damped by a factor of its input-prompting probability raised to the power $1 - \alpha$ (Proposition 4.2). We further analyze the non-asymptotic case, upper-bounding the factor by which the update can increase the probability of any token whose input-prompting probability lies below a threshold set by the step size and the token's score. For $\alpha < 1$, this factor approaches one as the token's input-prompting probability decreases, whereas at $\alpha = 1$ it grows exponentially in the score regardless of the token's probability (Proposition 4.3). Thus, Bregman tilting incorporates the score in the updated distribution while decaying the tilt strength in terms of each token's probability under input prompting, providing effective guidance while avoiding extreme jumps in the probabilities of tokens which are disfavored by input prompting. Further, we derive a sufficient condition for when the update, which optimizes a surrogate, also increases the true per-token objective of maximizing the probability of leading to an acceptable rollout. We show that if the feedback score positively covaries with the probability of acceptance over tokens drawn from the input-prompting distribution sharpened with temperature $1/(2 - \alpha) \leq 1$, then the Bregman tilting update with small step size provably increases this objective (Theorem 4.5). Smaller α amplifies this sharpening, concentrating the PMF on the likely tokens and lessening the dependence of this condition on potentially erroneous scores on unlikely tokens.

Experiments. We apply Bregman tilting at $\alpha \in \{0, 1\}$ with two input-prompting methods, Self-Refine [Mad+23] and Recursive Self-Aggregation (RSA) [Ven+26], with Qwen3-4B-Instruct-2507 [Yan+25]. We use two code benchmarks, LiveCodeBench and MBPP+, with feedback from running each program on tests, and two math benchmarks, AIME 2025 and HMMT February 2025, with feedback from a stronger reviewer LLM. In all eight settings, Bregman tilting with either value of α improves performance over input prompting, and $\alpha = 0$ achieves the best mean pass@1 and pass@16 across all benchmarks and input-prompting methods, showing the key role played by α in shaping the Bregman tilting update geometry beyond the exponential tilting form of $\alpha = 1$.

1.1 Related Work

A large body of work improves LLM reasoning at inference time by alternating between generation and feedback. Self-Refine [Mad+23] has the LLM critique and revise its own output, Reflexion [SCGNY23] keeps verbal reflections on past failures in context, and CRITIC [GSGSYDC24] and Self-Debugging [CLSZ24] incorporate feedback from external tools such as interpreters and unit tests. Recursive Self-Aggregation [Ven+26] refines a population of attempts by aggregating subsets of prior attempts to show to the model at each round. ProRefine [PWSNHW25] considers a prompt-optimization approach, using one LLM to generate natural language feedback and another to refine the prompt given to the solver at the next round. A related line of work optimizes the prompt itself through evolutionary search [FBMOR24; Agr+26]. All of these methods act on the LLM only through its context and treat it as a black-box sampler. In contrast, Bregman tilting directly modifies the LLM's next-token distribution and complements any input-prompting method, which we demonstrate on top of Self-Refine and RSA.

Prior inference-time guidance methods modify the LLM's next-token distribution directly, most commonly via exponential tilting. This form arises in posterior sampling, where Bayes' rule expresses the target as the original, or prior, distribution tilted by a learned classifier [DN21; YK21; BMMS26], and in KL-regularized RL, where the optimal policy tilts the original distribution by a per-token reward or value function [Mud+24; XSKZAHG25; RWGKP25]. Related reward-guided decoding methods apply the same tilt, adding a reward model's score of each candidate token to its log-probability [KBL24]. Other guidance methods take the score to be the log-ratio of two LLM next-token distributions. These can come from a strong and a weak LLM, as in contrastive decoding [LHFLEHZL23], or from the same LLM under two prompts: with and without retrieved context [SHLTZY24], with and without an attribute to condition on [HS22; SSFLB24], or against a deliberately corrupted instruction in instructive decoding [KKLY24]. Bregman tilting differs from these methods in two ways. First, its reference prompt contains the question and the previous rollout but not the feedback, so the score isolates the effect of the feedback rather than rewarding tokens which simply copy the previous rollout shown to the LLM. Second, it replaces the exponential tilt with the solution of a mirror ascent step under a family of Bregman divergences, which recovers the exponential tilt as a special case while inducing a family of more performant updates.

Bregman tilting performs a single step of mirror ascent [NY83; BT03], which generalizes the gradient ascent update to non-Euclidean geometries via the Bregman divergence of a strictly convex function, called the mirror map. We define a family of mirror maps ψ_α for $\alpha \in [0, 1]$. On the interior $\alpha \in (0, 1)$, ψ_α is, up to an additive constant, the negative Tsallis entropy of index α [Tsa88], scaled by $1/\alpha$, while at the endpoints we define ψ_0 as the log barrier and ψ_1 as the negative Shannon entropy respectively. This is a natural extension of the Tsallis family: up to additive constants, ψ_0 and ψ_1 are the limits of ψ_α as $\alpha \to 0$ and $\alpha \to 1$, and the induced Bregman divergence of ψ_α varies continuously in $\alpha \in [0, 1]$ (Proposition 3.2). This family of mirror maps is standard in online learning over the simplex, notably for multi-armed bandits [WL18; ZS21]. We bring this geometry to inference-time guidance, leading to our generalized tilting update.

Lastly, while we focus on inference-time methods, there is substantial prior work on training-based methods which update the LLM parameters. Reinforcement learning with a KL penalty to the original LLM, whose optimal policy is an exponential tilt of the original LLM by the reward, maximizes a reward that is learned, implicit in preference data, or given by a verifier [SWDRK17; RSMMEF23; Sha+24]. When text feedback or reference solutions are available, self-distillation methods instead train the LLM to match its own next-token distribution when it is additionally given this privileged information in context [ZXLHPCG26; PVGLCC26; Hüb+26], extending context distillation [SKZ22; YDWHW26].

With the reverse KL divergence as its objective, self-distillation is a policy gradient method whose per-token advantage is the log-ratio of the LLM's next-token distributions with and without the feedback in its prompt [Hüb+26], a contrastive signal of the same form as our score. However, self-distillation uses this signal to move the effect of the privileged information into the weights, so that the LLM prompted with the question alone matches the LLM when given the feedback in-context. Its target is therefore the feedback-prompted LLM itself, and its goal is to learn weights that generalize to new questions. Test-time variants apply the same update repeatedly to a single question [Hüb+26], but they still aim to distill the feedback into the weights and sample rollouts by prompting the LLM with the question alone. In contrast, our setting keeps the weights frozen, places the history of attempts and feedback in the prompt, and focuses on how to utilize this history to further update the feedback-prompted LLM distribution when generating the next attempt.

Bregman tilting therefore starts from the feedback-prompted LLM, which self-distillation uses as its training target, and tilts it further along the score without any backpropagation through the LLM.

Notation. Let $\mathcal{V}$ be the vocabulary and M the maximum sequence length, and write $\mathcal{V}^M$ for the set of token sequences of length M.[1] For a sequence x, x_m is its mth token and $x_{<m} \coloneqq (x_1, \dots, x_{m-1})$ is its prefix of length $m-1$ (empty for $m = 1$); (x, y) is the concatenation of two sequences, and $[n] \coloneqq \{1, \dots, n\}$.

We write $\Delta^+(\mathcal{V}) \coloneqq \{p \in \mathbb{R}^{\mathcal{V}} : \sum_{v \in \mathcal{V}} p(v) = 1,\ p(v) > 0 \text{ for every } v \in \mathcal{V}\}$ for the set of probability mass functions (PMFs) over the vocabulary with full support. Every next-token PMF of an LLM lies in $\Delta^+(\mathcal{V})$, since it is a softmax over the LLM's logits. An *autoregressive* PMF ρ assigns a next-token PMF $\rho(\cdot \mid h) \in \Delta^+(\mathcal{V})$ to every prefix of tokens h. Sampling n tokens one at a time from ρ after the prefix h induces a PMF on continuations $r \in \mathcal{V}^n$, $\rho(r \mid h) \coloneqq \prod_{m=1}^{n} \rho(r_m \mid h, r_{<m})$. We overload notation and write both the PMF over next tokens and over continuations as $\rho(\cdot \mid h)$, where the surrounding context makes the distinction clear. We write $\rho(r)$ for the PMF over rollouts with the empty prefix. We use π to denote the specific autoregressive PMF induced by the LLM; ρ refers to an arbitrary autoregressive PMF that need not be the output of any LLM. In contrast, p and q denote standalone PMFs in $\Delta^+(\mathcal{V})$.

We keep the following conventions throughout. Lower-case letters denote realized quantities: h is a prompt, r a rollout, f a feedback text, and v a token. The sans-serif capitals R and V denote a random rollout and a random token. $\mathbb{P}$ and $\mathbb{E}$ are the probability and expectation operators, always with the sampling law in the subscript, as in $\mathbb{P}_{\mathsf{R} \sim \rho}(\mathsf{R} \in \mathcal{A})$ or $\mathbb{E}_{\mathsf{V} \sim q}[\,\cdot\,]$. A subscript on a sequence or a PMF indexes a token position, as in r_m or $r_{<m}$; the round t and, when several chains run in parallel, the chain j are superscripts in parentheses, as in $h^{(t)}$ and $r^{(t,j)}$. Vectors indexed by $\mathcal{V}$, such as PMFs, are written as functions of v, so every operation on them is componentwise, and $\mathbf{1} \in \mathbb{R}^{\mathcal{V}}$ is the all-ones vector.

2 Inference-Time Pipeline

We first describe the inference-time pipeline which encompasses our method. Given an LLM π, we aim to solve a fixed question by iteratively generating answer attempts, which we refer to as rollouts, and receiving feedback on them. We assume access to a feedback mechanism Eval that takes a rollout r and returns a text $\mathrm{Eval}(r)$ about its quality. The mechanism can be a system such as a compiler or a test suite, another LLM, or a human.

At each round of the iterative pipeline, we prompt the LLM to generate rollouts, observe the feedback, and then update the distribution used for the next round's rollouts. We focus on tasks such as math and code, where correctness of each rollout is deterministically verifiable, for example by comparing its final answer to a reference or running its program on a set of tests, even though the pipeline itself cannot access this verifier. Our goal is simply to produce acceptable rollouts for the task.

[1] A sequence shorter than M is represented with an end-of-sequence token followed by padding, and we implicitly restrict our focus to autoregressive PMFs which generate only padding after this token.

2.1 Setup

We first define the pipeline mechanics. The pipeline runs for T rounds with Q chains per question. A chain is one sequence of attempts at the question, in which each round produces one rollout and its feedback. We write $r^{(t,j)}$ and $f^{(t,j)} = \mathrm{Eval}(r^{(t,j)})$ for the rollout and feedback of chain $j \in [Q]$ at round t. For brevity, we describe the pipeline for a single chain and drop the chain index j.

After round t, we have available the LLM π, the zero-shot prompt $h^{(1)}$, which contains the question and any instructions on the format of the answer, and the rollouts $r^{(1)}, \dots, r^{(t)}$ with their feedback $f^{(1)}, \dots, f^{(t)}$. We use this information to construct an autoregressive PMF $\rho^{(t+1)}$ for the next round, with the goal of increasing the probability of generating an acceptable rollout.

Formally, we define $\mathcal{A} \subseteq \mathcal{V}^M$ as the fixed set of acceptable rollouts. For an autoregressive PMF ρ, we define the acceptance probability as

$$A(\rho) \coloneqq \mathbb{P}_{\mathsf{R}\sim\rho}\left(\mathsf{R} \in \mathcal{A}\right), \tag{1}$$

and more generally, for a prefix x and autoregressive PMF ρ, we write the prefix-conditioned version

$$A(\rho\,; x) \coloneqq \mathbb{P}_{\mathsf{R}\sim\rho(\,\cdot\,|x)}\left((x, \mathsf{R}) \in \mathcal{A}\right).$$

For the empty prefix x, $A(\rho) = A(\rho\,; x)$. Our goal is then to construct the PMF $\rho^{(t+1)}$ for the next round so that $A(\rho^{(t)}) \le A(\rho^{(t+1)})$. However, since $\mathcal{A}$ is unknown, we cannot optimize (1) directly. Instead, we use the rollouts and feedback observed so far to update the LLM's rollout distribution with the goal of increasing the value of $A(\,\cdot\,)$ across rounds.

2.2 Input Prompting

We describe the standard input-prompting approach which our Bregman tilting method is designed to complement. The simplest way to construct $\rho^{(t+1)}$ is input prompting, which places the observed rollouts and their feedback in the prompt of the next round. At round 1 there is no feedback yet, and prompting the LLM with the zero-shot prompt gives the autoregressive PMF

$$\rho^{(1)}(r) = \pi(r \mid h^{(1)}). \tag{2}$$

An input-prompting method is defined by its prompt update, which constructs the next prompt $h^{(t+1)}$ from the previous prompts, rollouts, and feedback [Mad+23; SCGNY23; Ven+26]. Formally, the method defines a function $\mathrm{PromptUpdate}\,(\,\cdot\,)$ and sets

$$h^{(t+1)} = \mathrm{PromptUpdate}\left(h^{(1:t)}, r^{(1:t)}, f^{(1:t)}\right), \qquad \rho^{(t+1)}(r) = \pi(r \mid h^{(t+1)}).$$

In RSA [Ven+26], the update can also depend on the rollouts of the other chains of the question, but we describe the single-chain case here for brevity. We write $\pi^{(t)}$ for the *input-prompting LLM* at round t, the autoregressive PMF obtained from the same LLM π by fixing the prompt $h^{(t)}$. Formally, for every prefix x of a rollout, we write

$$\pi^{(t)}(\,\cdot \mid x) \coloneqq \pi(\,\cdot \mid h^{(t)}, x). \tag{3}$$

Thus, input prompting sets $\rho^{(t)} = \pi^{(t)}$ for each round t. This updates each $\rho^{(t)}$ only through the context of the LLM, allowing the model to see its past mistakes and reason about a better solution. While we cannot guarantee that the LLM improves its rollouts based on implicit reasoning given prior attempts, prior work on iterative prompting [Mad+23; SCGNY23; Ven+26] as well as our experiments in Section 5 demonstrate that existing iterative prompting methods improve performance round over round and significantly outperform the round-1 zero-shot model $\pi^{(1)}$. Thus, it serves as a sound starting point for our proposed framework.

3 Our Proposed Framework: Bregman Tilting

Although input prompting lets the LLM see its prior mistakes in the updated context $h^{(t+1)}$, it gives no control over how the LLM uses this information in its next rollout. Therefore, we complement a given input-prompting method by also directly modifying the next-token PMF of the LLM under the prompt $h^{(t+1)}$. To do so, we define a feedback score that measures how the feedback changes the likelihood of each candidate token, acting as a surrogate for the true acceptance probability. We then propose *Bregman tilting*, an inference-time guidance framework that takes a mirror ascent step from the input-prompting PMF, with the expected feedback score as the objective and a Tsallis-entropy-based mirror map (Section 3.2).

3.1 Per-Token Surrogate Objective

We first define the per-token counterpart $\bar{a}$ of the sequence-level objective A in (1). Fix the prompt $h^{(t+1)}$, and consider a prefix x of the rollout being generated. For each token v, define the probability that the input-prompting LLM $\pi^{(t+1)}$ defined in (3) completes the prefix (x, v) with a rollout in $\mathcal{A}$ as

$$a(v\,;x) \coloneqq A\left(\pi^{(t+1)}\,;(x,v)\right). \tag{4}$$

For any autoregressive PMF ρ, we define the per-token objective $\bar{a}(\rho\,;x)$ as the expected value of $a(\mathsf{V}\,;x)$ under its next-token PMF $\rho(\cdot \mid x)$,

$$\bar{a}(\rho\,;x) \coloneqq \mathbb{E}_{\mathsf{V}\sim\rho(\,\cdot\,|x)}\left[a(\mathsf{V}\,;x)\right]. \tag{5}$$

Importantly, the gain in the sequence-level objective A of (1) over the input-prompting LLM $\pi^{(t+1)}$ can be expressed as the sum of expected per-token gains in $\bar{a}$.

Proposition 3.1. *For every autoregressive PMF ρ,*

$$A(\rho) - A(\pi^{(t+1)}) = \sum_{m=1}^{M} \mathbb{E}_{\mathsf{R}\sim\rho}\left[\bar{a}\left(\rho\,;\mathsf{R}_{<m}\right) - \bar{a}(\pi^{(t+1)}\,;\mathsf{R}_{<m})\right]. \tag{6}$$

It therefore suffices to study the update at a single position m with a fixed prefix $r_{<m}$. At this prefix, $\bar{a}(\rho\,;r_{<m})$ depends on ρ only through its next-token PMF $\rho(\cdot \mid r_{<m})$, so we treat this PMF as a free variable $p \in \Delta^{+}(\mathcal{V})$ and define

$$A_m(p) \coloneqq \mathbb{E}_{\mathsf{V}\sim p}\left[a(\mathsf{V}\,;r_{<m})\right] = \sum_{v\in\mathcal{V}} p(v)\, a(v\,;r_{<m}).$$

The subscript m marks the dependence on the fixed prefix $r_{<m}$, so that $\bar{a}(\rho\,; r_{<m}) = A_m\left(\rho(\cdot \mid r_{<m})\right)$.

While A_m is the ideal per-token objective, like A it is unobserved. We thus replace it with a tractable surrogate. We define a per-token feedback score $s^{(t+1)}(v \mid r_{<m})$ which computes the log-probability gain for the token v when the LLM π is given the updated prompt $h^{(t+1)}$, which contains the previous rollout $r^{(t)}$ and its feedback $f^{(t)}$, compared to when it is given a modified reference prompt $h^{(t+1)}_{\setminus f}$ containing $r^{(t)}$ *but not the feedback* $f^{(t)}$. Formally, for a token v and prefix $r_{<m}$,

$$s^{(t+1)}(v \mid r_{<m}) \coloneqq \ln \pi(v \mid h^{(t+1)}, r_{<m}) - \ln \pi(v \mid h^{(t+1)}_{\setminus f}, r_{<m}). \tag{7}$$

See Appendix B.5 for the exact prompts we use. Going forward, we abbreviate the input-prompting PMF and the score given the fixed prefix $r_{<m}$ as

$$\pi_m \coloneqq \pi^{(t+1)}(\cdot \mid r_{<m}) \in \Delta^+(\mathcal{V}), \qquad s_m \coloneqq s^{(t+1)}(\cdot \mid r_{<m}) \in \mathbb{R}^{\mathcal{V}}. \tag{8}$$

This score $s_m(v)$ is positive when the feedback makes v more likely and negative when it makes v less likely at the prefix $r_{<m}$. For each next token of the round-$(t+1)$ rollout, we use this score in place of a in (5), replacing the intractable objective with the following surrogate.

$$\tilde{A}_m(p) \coloneqq \mathbb{E}_{\mathsf{V}\sim p}\left[s_m(\mathsf{V})\right] = \sum_{v\in\mathcal{V}} p(v)\, s_m(v). \tag{9}$$

Since $\tilde{A}_m$ is a surrogate, we do not maximize it outright, as this would simply place all probability mass on the highest scoring token. Instead, we search for a distribution which remains close to the input-prompting PMF π_m while increasing the value of $\tilde{A}_m$, leading to our mirror-ascent-based Bregman tilting update.

3.2 Bregman Tilting Update

The standard way to perturb π_m toward the tokens that the feedback favors is through exponential tilting, which returns the updated next-token PMF $\hat{p}_m(v) \propto \pi_m(v)\, e^{\eta\, s_m(v)}$ [SHLTZY24; HS22; SSFLB24]. It reweights the probability assigned to each token by the factor $e^{\eta\, s_m(v)}$, potentially increasing the probability of a token that π_m considers unlikely by a large factor. We instead propose *Bregman tilting*, which returns the output of a single mirror ascent step on the surrogate objective $\tilde{A}_m$. We consider a family of mirror maps ψ_α for $\alpha \in [0,1]$ which strictly generalize the exponential tilting update and result in a π_m-*dependent tilt* which adapts the update geometry to better resist changes to low-probability tokens under π_m.

Treating the optimization variable p as a vector of next-token probabilities, we can write

$$\tilde{A}_m(p) = \langle s_m, p\rangle.$$

We maximize this objective over the simplex $\Delta^+(\mathcal{V})$, regularized by the Bregman divergence $D_\psi(p,q)$ of a strictly convex, differentiable function ψ, defined by

$$D_\psi(p,q) \coloneqq \psi(p) - \psi(q) - \langle \nabla\psi(q), p - q\rangle. \tag{10}$$

The *Bregman tilting* update is then

$$\hat{p}_m \coloneqq \operatorname*{argmax}_{p \in \Delta^+(\mathcal{V})} \langle s_m, p \rangle - \frac{1}{\eta} D_\psi(p, \pi_m). \tag{11}$$

Since $s_m = \nabla \tilde{A}_m(p)$, Equation (11) can be interpreted as a single step of mirror ascent [NY83] applied to the objective $\tilde{A}_m(p)$. We consider a particular family of mirror maps ψ_α parameterized by $\alpha \in [0, 1]$. On the interior $\alpha \in (0, 1)$ we define

$$\psi_\alpha(p) \coloneqq -\frac{1}{\alpha(1-\alpha)} \sum_{v \in \mathcal{V}} p(v)^\alpha. \tag{12}$$

This is the negative Tsallis entropy of index α [Tsa88], scaled by $1/\alpha$ and shifted by a constant. We then define the endpoints ψ_0 and ψ_1 as the log-barrier and negative Shannon entropy respectively,

$$\psi_0(p) \coloneqq -\sum_{v \in \mathcal{V}} \ln p(v), \qquad \psi_1(p) \coloneqq \sum_{v \in \mathcal{V}} p(v) \ln p(v). \tag{13}$$

By design, this family of strictly convex functions ψ_α for $\alpha \in [0, 1]$ is continuous in terms of the Bregman divergence, the key operation which defines the Bregman tilting update in Equation (11).

Proposition 3.2. *Consider the family of mirror maps ψ_α defined by Equations (12), (13), and arbitrary PMFs $p, q \in \Delta^+(\mathcal{V})$.*

(i) On the interior $\alpha \in (0, 1)$, the Bregman divergence of ψ_α is

$$D_{\psi_\alpha}(p, q) = \frac{1}{\alpha(1-\alpha)} \sum_{v \in \mathcal{V}} \left(q(v)^\alpha - p(v)^\alpha + \alpha\, q(v)^{\alpha-1} \left(p(v) - q(v) \right) \right). \tag{14}$$

(ii) The Bregman divergence D_{ψ_α} is a continuous function of α at the endpoints $\{0, 1\}$, with limits

$$\lim_{\alpha \to 0} D_{\psi_\alpha}(p, q) = D_{\psi_0}(p, q) = \sum_{v \in \mathcal{V}} \left(\frac{p(v)}{q(v)} - \ln \frac{p(v)}{q(v)} - 1 \right), \tag{15}$$

$$\lim_{\alpha \to 1} D_{\psi_\alpha}(p, q) = D_{\psi_1}(p, q) = \sum_{v \in \mathcal{V}} p(v) \ln \frac{p(v)}{q(v)} = \mathrm{KL}\left(p \,\|\, q \right). \tag{16}$$

For this family of Bregman divergences, we compute the update rule below.

Proposition 3.3. *For $\alpha \in [0, 1]$ and $\eta > 0$, Equation (11) with the mirror map ψ_α has a unique solution, which we denote $\hat{p}_m(\,\cdot\,; \alpha)$. Moreover, there exists a scalar $\hat{\lambda}$ such that $(p, \lambda) = (\hat{p}_m(\,\cdot\,; \alpha), \hat{\lambda})$ is the unique pair in $\Delta^+(\mathcal{V}) \times \mathbb{R}$ which satisfies*

$$\nabla \psi_\alpha(p) = \nabla \psi_\alpha(\pi_m) + \eta\, s_m - \lambda \mathbf{1}. \tag{17}$$

While Proposition 3.3 implicitly defines the update, we further compute it in closed form. For $\alpha = 1$, the Bregman divergence $D_{\psi_1}(p, q)$ is the KL divergence $\mathrm{KL}\left(p \,\|\, q \right)$ (Equation (16)), so Equation (11) is the familiar KL-regularized reward maximization problem with the score s_m as the reward [RSMMEF23]. Substituting $[\nabla \psi_1(p)](v) = \ln p(v) + 1$ into Equation (17) gives

$$\hat{p}_m(v\,; 1) = \pi_m(v)\, e^{\eta\, s_m(v) - \hat{\lambda}}. \tag{18}$$

Thus, $\alpha = 1$ recovers standard exponential tilting, reweighting the probability of each token under π_m by $e^{\eta\, s_m(v) - \hat{\lambda}}$. For $\alpha \in [0, 1)$, substituting $[\nabla\psi_\alpha(p)](v) = -\frac{1}{1-\alpha} p(v)^{\alpha-1}$ into Equation (17) gives

$$\hat{p}_m(v\,;\alpha) = \frac{\pi_m(v)}{\Big(1 + (1-\alpha)\pi_m(v)^{1-\alpha}\Big(\hat{\lambda} - \eta\, s_m(v)\Big)\Big)^{\frac{1}{1-\alpha}}}, \tag{19}$$

which recovers Equation (18) in the limit $\alpha \to 1$. This update for $\alpha < 1$ reweights the probability $\pi_m(v)$ by the π_m-dependent tilt $\Big(1 + (1-\alpha)\pi_m(v)^{1-\alpha}(\hat{\lambda} - \eta\, s_m(v))\Big)^{-\frac{1}{1-\alpha}}$. Similar to the $\alpha = 1$ case, the tilt factor is increasing in the score $s_m(v)$, thereby boosting the probability of tokens with higher scores. However, the tilt factor is independent of $\pi_m(v)$ only in the special case $\alpha = 1$; for $\alpha < 1$, the update modulates the tilt strength in terms of the probability assigned by the input-prompting PMF $\pi_m(v)$. We quantify the consequences of the choice of α in the next section.

Remark 3.4 (Solving for $\hat{\lambda}$). Computing the update requires the multiplier $\hat{\lambda}$ that ensures $\hat{p}_m(\,\cdot\,;\alpha)$ sums to one. For $\alpha = 1$ it is the log-normalizer $\hat{\lambda} = \ln \sum_v \pi_m(v)\, e^{\eta\, s_m(v)}$. For $\alpha < 1$ it has no closed form, so we resort to numerical approximation. As a function of λ, the total mass of the right-hand side of Equation (19) is strictly decreasing and convex, and $\hat{\lambda}$ lies in a known interval. Thus, a few steps of bisection and Newton's method can efficiently approximate $\hat{\lambda}$, while adding negligible cost to the LLM decoding step. See Appendix A.3 for details.

4 Theoretical Analysis

We analyze the Bregman tilting update in two parts. First, we show how α shapes the geometry of the Bregman divergence regularizer, and how this results in a damped exponential tilt update that limits the change in probabilities of tokens that are unlikely under the input-prompting PMF π_m. Second, we connect the update, which ascends the surrogate objective $\tilde{A}_m$, to the true per-token objective A_m, deriving a sufficient condition for ascending the true objective. Below, we write $\hat{p}_m(\,\cdot\,;\alpha,\eta)$ for the solution of Equation (11) with the mirror map ψ_α and step size $\eta \geq 0$, and $a_m(v) \coloneqq a(v\,; r_{<m})$, so that $A_m(p) = \mathbb{E}_{\mathsf{V}\sim p}\,[a_m(\mathsf{V})]$. All other notation follows Section 3.2.

4.1 The Geometry Induced by α

We study how the index α of the mirror maps (12) and (13) shapes the Bregman tilting update. Importantly, we highlight how smaller values of α smoothly damp the update strength for each token in terms of its probability under the input-prompting PMF π_m, incorporating the feedback score into the update without drastically increasing the probability of unlikely tokens.

We first analyze the local behavior of the divergence around π_m. For small perturbations, the divergence acts as the squared weighted Euclidean norm induced by the positive definite Hessian of the mirror map ψ_α.

Proposition 4.1. *Let $\alpha \in [0, 1]$. For every $p \in \Delta^+(\mathcal{V})$, $\nabla^2\psi_\alpha(p) = \operatorname{diag}\big(p(v)^{\alpha-2}\big)_{v\in\mathcal{V}}$, and for every*

perturbation $u \in \mathbb{R}^{\mathcal{V}}$ *with* $\sum_v u(v) = 0$,

$$D_{\psi_\alpha}(\pi_m + u, \pi_m) = \frac{1}{2} \sum_{v \in \mathcal{V}} \pi_m(v)^\alpha \left(\frac{u(v)}{\pi_m(v)} \right)^2 + \mathcal{O}\left(\|u\|^3 \right) \quad \text{as } u \to 0. \tag{20}$$

The divergence locally penalizes the perturbation $u(v)$ of each token by its squared relative change $(u(v)/\pi_m(v))^2$, weighted by $\pi_m(v)^\alpha$. At $\alpha = 1$, the weight is $\pi_m(v)$ itself, so changing the probability of an unlikely token by a large factor costs almost nothing, but as α decreases to 0 the weights approach one, and at $\alpha = 0$ every weight is exactly one. Since an unlikely token can only gain significant probability by a large relative change, a smaller α makes such a change much more costly.

While this result concerns the regularizer, we show that the update inherits this behavior by analyzing how it changes the probability of a single token v asymptotically as $\pi_m(v)$ goes to 0.

Proposition 4.2. *Let* $\alpha \in [0, 1]$ *and* $\eta > 0$*, and fix the score* s_m*. For every token* $v \in \mathcal{V}$*, the log-ratio of the mass assigned to* v *by the update* $\hat{p}_m(v\,; \alpha)$ *relative to* $\pi_m(v)$ *satisfies*

$$\ln \frac{\hat{p}_m(v\,; \alpha)}{\pi_m(v)} = \pi_m(v)^{1-\alpha} \left(\eta\, s_m(v) - \hat{\lambda} \right) + \mathcal{O}\left(\pi_m(v)^{2(1-\alpha)} \right) \quad \text{as } \pi_m(v) \to 0. \tag{21}$$

Thus, to leading order, Bregman tilting changes the log-probability of an unlikely token by $\eta\, s_m(v) - \hat{\lambda}$, the exact amount of the exponential tilt (18), but *damped* by the factor $\pi_m(v)^{1-\alpha}$. At $\alpha = 1$ the damping factor is one, and every token is tilted with the same step size η no matter how unlikely it is. For $\alpha < 1$, the damping factor decreases smoothly to zero as $\pi_m(v)$ decreases, and at $\alpha = 0$ it is $\pi_m(v)$ itself. Thus, the step size η sets the strength of the tilt, while α controls how quickly this strength decays on the tokens that the input-prompting PMF π_m considers unlikely.

While Proposition 4.2 describes the limit $\pi_m(v) \to 0$, we lastly give a non-asymptotic bound on how far the update can promote a token, which holds at every step size η for every token whose probability under π_m is small relative to its score. To state it, we define the sharpened PMF $\pi_m^{[\alpha]} \in \Delta^+(\mathcal{V})$ and the expected score $\bar{s}_m^{[\alpha]}$ under it,

$$\pi_m^{[\alpha]}(v) \coloneqq \frac{\pi_m(v)^{2-\alpha}}{\sum_{u \in \mathcal{V}} \pi_m(u)^{2-\alpha}}, \qquad \bar{s}_m^{[\alpha]} \coloneqq \mathbb{E}_{\mathsf{V} \sim \pi_m^{[\alpha]}} \left[s_m(\mathsf{V}) \right]. \tag{22}$$

The sharpened PMF $\pi_m^{[\alpha]}$ raises each probability to the power $2 - \alpha \geq 1$ and renormalizes, which is equivalent to applying the temperature $1/(2-\alpha) \leq 1$ to π_m. At $\alpha = 1$ it is π_m itself, and for $\alpha < 1$ it moves probability mass from unlikely tokens to likely ones.

Proposition 4.3. *Let* $\alpha \in [0, 1]$ *and* $\eta > 0$*. If a token* $v \in \mathcal{V}$ *satisfies*

$$(1 - \alpha)\, \eta\, \pi_m(v)^{1-\alpha} \left(s_m(v) - \bar{s}_m^{[\alpha]} \right) < 1, \tag{23}$$

then

$$\frac{\hat{p}_m(v\,; \alpha)}{\pi_m(v)} \leq \begin{cases} \left(1 - (1 - \alpha)\, \eta\, \pi_m(v)^{1-\alpha} \left(s_m(v) - \bar{s}_m^{[\alpha]} \right) \right)^{-\frac{1}{1-\alpha}} & \text{if } \alpha \in [0, 1), \\ e^{\eta \left(s_m(v) - \bar{s}_m^{[1]} \right)} & \text{if } \alpha = 1. \end{cases} \tag{24}$$

This bounds the factor by which the update can increase the probability of a token which has small probability relative to its score as defined by the condition in (23). At $\alpha = 1$ the cap grows exponentially in the score and is independent of $\pi_m(v)$, so an unlikely token with a large score can gain probability by a large factor. In contrast, for $\alpha < 1$ the score enters the bound only through the damped quantity $\eta\,\pi_m(v)^{1-\alpha}(s_m(v) - \bar{s}_m^{[\alpha]})$, making the bound close to one when this quantity is small. For example, at $\alpha = 0$ the bound is $1/\Big(1 - \eta\,\pi_m(v)(s_m(v) - \bar{s}_m^{[0]})\Big)$, so even a large step size increases the probability of an unlikely token by only a small factor unless its score is extreme.

Notably, the condition (23) always holds at $\alpha = 1$, where its left-hand side is zero, as well as when $\alpha < 1$ for every token v whose score is at most the sharpened mean, $s_m(v) \le \bar{s}_m^{[\alpha]}$. In general, it holds for every token whose probability under π_m lies below a constant determined by η and its centered score $s_m(v) - \bar{s}_m^{[\alpha]}$, and it fails only for tokens that are both relatively likely and favored by the feedback, which are the tokens the tilt is meant to move.

Together, these results show that α smoothly controls the strength of the update for each token v in terms of the input-prompting probability $\pi_m(v)$. At $\alpha = 1$, every token is tilted by the same rule regardless of its probability, while a smaller α concentrates the tilt on the likely tokens and restricts it on the unlikely ones. This is preferable when the score is reliable on the tokens that the input-prompting LLM considers likely but possibly unreliable on the rest, since a smaller α prevents a large, possibly erroneous score on an unlikely token from greatly increasing its probability. Section 4.2 makes this interpretation precise, showing that the same discounting appears in a sufficient condition for guaranteeing that the update ascends the true objective.

4.2 Ascending the True Objective

We connect the Bregman tilting update on the surrogate $\tilde{A}_m$ to the true objective A_m. We first show that Bregman tilting is guaranteed to ascend the surrogate, a standard result for mirror ascent.

Proposition 4.4. *Let $\alpha \in [0, 1]$ and $\eta > 0$, and write $\hat{p}_m = \hat{p}_m(\,\cdot\,; \alpha, \eta)$. Then*

$$\tilde{A}_m(\hat{p}_m) - \tilde{A}_m(\pi_m) = \frac{1}{\eta}\Big(D_{\psi_\alpha}(\hat{p}_m, \pi_m) + D_{\psi_\alpha}(\pi_m, \hat{p}_m)\Big) \ge 0. \tag{25}$$

Further, if the score s_m is non-constant over $\mathcal{V}$, the gain is strictly positive, $\tilde{A}_m(\hat{p}_m) - \tilde{A}_m(\pi_m) > 0$.

To relate this gain to A_m, we compare the rates at which the two objectives increase as the step size η grows from zero. For $\alpha \in [0, 1]$, we now define $\Gamma_m(\alpha)$ and $\tilde{\Gamma}_m(\alpha)$ as the derivatives with respect to η of A_m and $\tilde{A}_m$ along the update, evaluated at $\eta = 0$,

$$\Gamma_m(\alpha) \coloneqq \frac{d}{d\eta}\, A_m\left(\hat{p}_m(\,\cdot\,; \alpha, \eta)\right)\Big|_{\eta=0}, \qquad \tilde{\Gamma}_m(\alpha) \coloneqq \frac{d}{d\eta}\, \tilde{A}_m\left(\hat{p}_m(\,\cdot\,; \alpha, \eta)\right)\Big|_{\eta=0}. \tag{26}$$

We show that the ratio of the rates $\Gamma_m(\alpha)$ and $\tilde{\Gamma}_m(\alpha)$ is determined by the covariance and variance of a token drawn from the sharpened input-prompting PMF $\pi_m^{[\alpha]}$ defined in (22).

Theorem 4.5. *Let $\alpha \in [0, 1]$, and let the score s_m be non-constant over the vocabulary $\mathcal{V}$. Then the derivatives in Equation (26) exist, $\tilde{\Gamma}_m(\alpha) > 0$, and*

$$\frac{\Gamma_m(\alpha)}{\tilde{\Gamma}_m(\alpha)} = \frac{\mathrm{Cov}_{\mathsf{V}\sim\pi_m^{[\alpha]}}\left(s_m(\mathsf{V}),\, a_m(\mathsf{V})\right)}{\mathrm{Var}_{\mathsf{V}\sim\pi_m^{[\alpha]}}\left(s_m(\mathsf{V})\right)}. \tag{27}$$

Algorithm 1 Bregman Tilting

Require: LLM π, feedback mechanism Eval, input-prompting method, number of chains Q, number of rounds T, index $\alpha \in [0,1]$, step size $\eta \geq 0$

for $t = 1, \ldots, T$ **do**
  Build $h^{(t,j)}$ and $h^{(t,j)}_{\setminus \hat{f}}$ for every $j \in [Q]$ with the input-prompting method
  for $j \in [Q]$ **do**
    for $m = 1, 2, \ldots$ until the end of the sequence **do**
      $\pi_m \leftarrow \pi(\cdot \mid h^{(t,j)}, r_{<m}), \quad s_m \leftarrow \ln \pi_m - \ln \pi(\cdot \mid h^{(t,j)}_{\setminus \hat{f}}, r_{<m})$
      $\hat{p}_m(\cdot\,; \alpha) \leftarrow$ Bregman tilting update of π_m with score s_m ▷ (18), (19), Remark 3.4
      Draw $r_m \sim \hat{p}_m(\cdot\,; \alpha)$
    end for
    $r^{(t,j)} \leftarrow (r_1, r_2, \ldots), \quad f^{(t,j)} \leftarrow \mathrm{Eval}(r^{(t,j)})$
  end for
end for
return $r^{(T,j)}$, $j \in [Q]$

The right-hand side of Equation (27) is the coefficient of the least-squares regression of $a_m(\mathsf{V})$ on $s_m(\mathsf{V})$ under $\mathsf{V} \sim \pi_m^{[\alpha]}$. Thus, Theorem 4.5 shows that, per unit of gain in the surrogate, the gain in the true objective is determined by how well the score vector s_m predicts the true vector of acceptance probabilities a_m. Since the denominators on both sides of Equation (27) are strictly positive, this gives the following condition for ascent of the true objective.

Corollary 4.5.1. *Let $\alpha \in [0,1]$, and suppose $\mathrm{Cov}_{\mathsf{V} \sim \pi_m^{[\alpha]}}(s_m(\mathsf{V}), a_m(\mathsf{V})) > 0$. Then there exists $\eta_0 > 0$ such that $A_m(\hat{p}_m(\cdot\,; \alpha, \eta)) > A_m(\pi_m)$ for every $\eta \in (0, \eta_0)$.*

The Bregman tilting update ascends the true per-token objective A_m at small step sizes whenever the feedback score $s_m(\mathsf{V})$ and the acceptance probability $a_m(\mathsf{V})$ positively covary under $\mathsf{V} \sim \pi_m^{[\alpha]}$. Since $\pi_m^{[\alpha]}(v) \propto \pi_m(v)^{1-\alpha}\, \pi_m(v)$, the weight of each token relative to $\alpha = 1$ is multiplied by $\pi_m(v)^{1-\alpha}$, the same damping factor as in Equation (21). Thus, in this condition as well, a smaller α discounts the score's errors on the tokens that input prompting considers unlikely when determining whether ascent of the surrogate carries over to the true objective.

Overall, α controls how far the score can move the tokens that input prompting considers unlikely (Propositions 4.2 and 4.3), and Corollary 4.5.1 gives a sufficient condition, at small step sizes, for the update to improve the true objective. Since this condition depends on the unknown acceptance probabilities a_m, the theory does not determine the best α. In Section 5, we compare $\alpha = 1$ and $\alpha = 0$ empirically and find that $\alpha = 0$ performs best, highlighting both the value of the framework and the robustness to unlikely tokens that our analysis identifies.

5 Experiments

We apply Bregman tilting with $\alpha = 1$ and $\alpha = 0$ on top of two input-prompting methods using two code and two math benchmarks. Appendix B gives the full experimental details.

Table 1: Mean pass@1 and pass@16 (%) at round 5, averaged over four seeds, with $Q = 16$ rollouts per question, along with the standard error. Zero-shot is the round-1 solver which all methods share within a seed. The shaded columns correspond to our Bregman tilting. The best score in each row and metric is highlighted in **green**.

		Pass@1				Pass@16			
				Bregman tilting				Bregman tilting	
Benchmark	Method	Zero-shot	Prompting	$\alpha = 1$	$\alpha = 0$	Zero-shot	Prompting	$\alpha = 1$	$\alpha = 0$
LiveCodeBench	Self-Refine	34.1 ±0.2	38.6 ±0.2	39.9 ±0.2	**41.0** ±0.3	47.5 ±0.7	52.1 ±0.7	55.5 ±0.8	**59.0** ±0.7
	RSA	34.1 ±0.2	43.6 ±0.5	45.9 ±0.6	**47.0** ±0.6	47.5 ±0.7	47.9 ±0.8	49.6 ±0.8	**51.5** ±0.8
MBPP+ (hard)	Self-Refine	40.8 ±0.3	64.3 ±0.3	67.1 ±0.3	**67.8** ±0.3	62.3 ±0.7	83.4 ±0.5	83.5 ±0.5	**85.0** ±0.5
	RSA	40.8 ±0.3	72.9 ±0.6	77.5 ±0.5	**78.9** ±0.4	62.3 ±0.7	77.7 ±0.5	81.1 ±0.6	**82.6** ±0.4
AIME 2025	Self-Refine	41.4 ±0.6	70.3 ±0.6	71.8 ±0.6	**73.2** ±0.6	67.5 ±2.0	87.5 ±1.6	90.0 ±1.4	**92.5** ±0.8
	RSA	41.4 ±0.6	80.7 ±0.9	84.5 ±1.0	**85.9** ±1.2	67.5 ±2.0	85.8 ±1.6	89.2 ±1.4	**90.0** ±1.5
HMMT Feb 2025	Self-Refine	29.8 ±0.5	55.3 ±0.7	57.8 ±0.7	**58.2** ±0.7	52.5 ±2.3	78.3 ±1.5	**83.3** ±1.5	**83.3** ±1.5
	RSA	29.8 ±0.5	68.7 ±1.6	74.5 ±1.4	**74.6** ±1.0	52.5 ±2.3	75.0 ±1.9	80.8 ±1.6	**83.3** ±1.5

Algorithm. We present the formal Bregman tilting algorithm in Algorithm 1. In each round t, the input-prompting method builds a prompt $h^{(t,j)}$ for each chain j from the previous rollouts and their feedback (Section 2.2). We pair it with its no-feedback counterpart $h^{(t,j)}_{\setminus f}$, which just removes the feedback and the instructions that refer to it. The rollout is generated one token at a time: at each position m, we evaluate π on both prompts after the same prefix $r_{<m}$, form the score s_m of Equation (7), solve for $\hat{\lambda}$, and sample the next token from $\hat{p}_m(\cdot\,; \alpha)$. In round 1 both prompts are the zero-shot prompt $h^{(1)}$, so the score is zero and the chain samples from the zero-shot PMF (2). See Appendix B.5 for the exact prompts.

Input-Prompting Methods. We use *Self-Refine* [Mad+23] and *Recursive Self-Aggregation* (RSA) [Ven+26], modified to incorporate feedback. Self-Refine has the model first critique its own previous output, then generate a refined response using this critique. We adapt this to our setting, where the prompt $h^{(t+1)}$ shows the question, the chain's most recent rollout $r^{(t)}$, and its feedback $\mathrm{Eval}(r^{(t)})$, and asks for a refined solution that addresses the feedback. For Self-Refine, the Q chains of a question run independently. In contrast, RSA lets the Q chains of a question interact. After each round, the RSA prompt of every chain shows $K = 4$ rollout-feedback pairs drawn at random from the previous round's population and asks for a single improved solution. This allows a correct step found by one chain to spread to the others. For both methods, the no-feedback prompt $h^{(t+1)}_{\setminus f}$ is identical to $h^{(t+1)}$, except the feedback and corresponding instructions are removed (Appendix B.5).

Model and Sampling. We use Qwen3-4B-Instruct-2507 as the solver, sampled at temperature 1. Every method runs for $T = 5$ rounds with $Q = 16$ chains per question, across four seeds. Within a seed, all methods start from the same zero-shot round-1 rollouts and feedback, so the comparisons are paired by question. Bregman tilting uses $\eta = 1$ at $\alpha = 1$ and $\eta = 48$ at $\alpha = 0$.

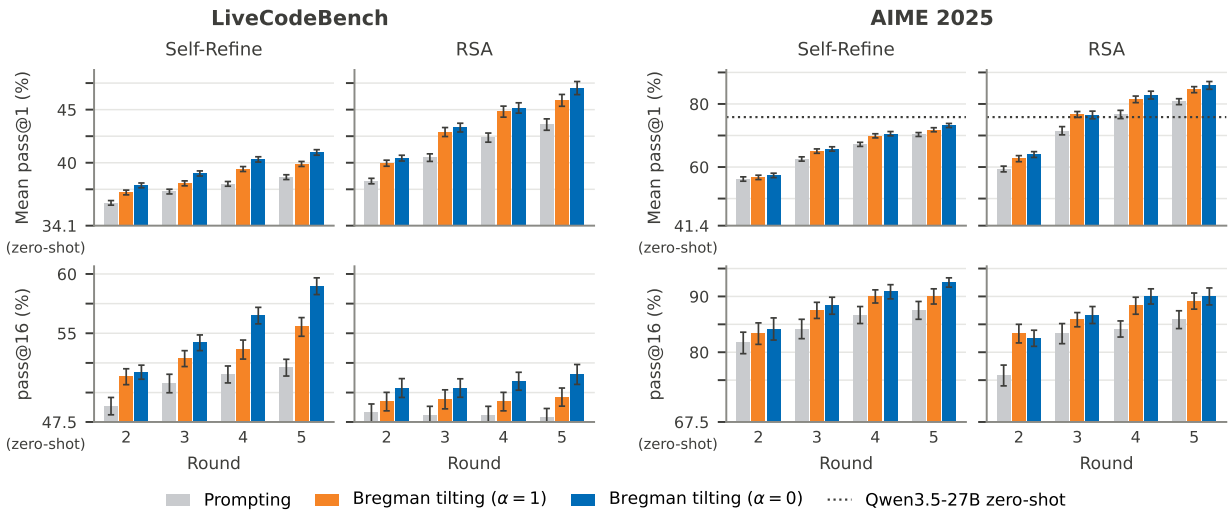

Figure 1: Mean pass@1 (top) and pass@16 (bottom) by round on LiveCodeBench (left) and AIME 2025 (right) for Self-Refine and RSA. Prompting refers to input prompting only, and we compare to our Bregman tilting at $\alpha = 0$ and $\alpha = 1$. On LiveCodeBench the feedback comes from running each program on a test suite, and on AIME 2025 from Qwen3.5-27B as a reviewer. Bars are means over four seeds, with standard error bars. Within each seed, all methods use the same round-1 rollouts, drawn from the zero-shot Qwen3-4B solver, so each row's y-axis starts at the round-1 score, marked "(zero-shot)", and the height of a bar above it is the method's gain over zero-shot. For AIME, the dotted line is the zero-shot mean pass@1 of the Qwen3.5-27B feedback model.

Benchmarks. For code, we use LiveCodeBench-v6 [JHGLYZWSSS25] and MBPP+ [LXWZ23]. We consider a generated program correct if it passes all exposed and hidden tests. For LiveCodeBench we use the problems published from February to April 2025, and for MBPP+ we filter out the 223 of the 378 which Qwen3-4B-Instruct-2507 already solves in all 16 of its seed-0 zero-shot rollouts. We report results on the remaining *hard set* in the main body (see Appendix B.8 for full-set performance). For math, we use AIME 2025 and HMMT February 2025 via MathArena [BDPJV25].

Feedback. On code, the feedback mechanism Eval(r) runs the program of r on a set of exposed tests and reports the first test it fails, with the input, the program's output, and the expected output, or the error it raises. The benchmarks' public tests are few and miss many errors, so we expose the public tests together with a fixed random half of the benchmarks' remaining private tests for each problem as in prior work [Hüb+26]. On LiveCodeBench, this raises the share of wrong programs that the feedback flags from about 60% to over 90%. On math, we simulate the feedback of a human expert by prompting Qwen3.5-27B [Qwe26] to review the rollout r without seeing the reference answer and to point to the first step that fails and why, without giving an answer of its own.

Metrics. We report mean pass@1, the fraction of a round's Q rollouts that are correct, and pass@16, whether any of them is correct, each averaged over the questions. Under Self-Refine the Q chains of a question are independent, so the rollouts of a round are independent draws from that round's distribution and pass@16 is its usual pass@k with $k = Q$. Under RSA the chains share rollouts through their prompts, so pass@16 instead measures how many questions the population

Table 2: Necessity of feedback isolation in the reference prompt, on AIME 2025 with Self-Refine. We compare Context-Aware Decoding (CAD) [SHLTZY24], which is Bregman tilting at $\alpha = 1$ with the zero-shot reference prompt $h^{(1)}$, to Bregman tilting at $\alpha = 1$ and $\alpha = 0$ with our no-feedback reference prompt $h_{\setminus f}^{(t+1)}$. We report mean pass@1 and pass@16 (%) at round 5, averaged over four seeds, as well as the percentage of rollouts with a reviewer-flagged wrong final answer whose next-round rollout repeats the same answer, pooled over rounds 1 to 4 and all seeds.

Method	Reference prompt	Mean pass@1 (↑)	Pass@16 (↑)	Flagged Wrong Answer Repetition Rate (%) (↓)
Prompting	—	70.3	87.5	44.3
CAD	Zero-shot $h^{(1)}$	69.8	86.7	43.4
Bregman tilting, $\alpha = 1$	No feedback $h_{\setminus f}^{(t+1)}$ (ours)	71.8	90.0	29.9
Bregman tilting, $\alpha = 0$	No feedback $h_{\setminus f}^{(t+1)}$ (ours)	**73.2**	**92.5**	**21.7**

that RSA produces covers. In both cases, we compute pass@16 from the rollouts directly. Main results are shown with standard errors. See Appendix B.4 for details.

5.1 Results

Main Results. Table 1 reports mean pass@1 and pass@16 after five rounds. Input prompting alone improves substantially over the zero-shot solver on every benchmark, while Bregman tilting further improves on top of input prompting for both values of α. Crucially, $\alpha = 0$ achieves the best mean pass@1 and pass@16 in every setting, strictly improving over $\alpha = 1$ everywhere except HMMT Self-Refine pass@16, where they tie. Figure 1 highlights the round-by-round performance on LiveCodeBench and AIME 2025. Throughout the rounds, both settings of Bregman tilting lead prompting from the first tilted round onward in both metrics, and $\alpha = 0$ leads $\alpha = 1$ in all but two of the 32 round-and-metric comparisons, both in intermediate rounds with RSA on AIME 2025. Since $\alpha = 1$ reduces to the exponential tilt used by prior guidance methods, applied here with our feedback score, these results demonstrate the value of the framework as a whole; the feedback score alone gives consistent gains over input prompting, while moving beyond the exponential tilt to $\alpha = 0$ improves even further. See Appendix B.6 for the round-by-round results on all benchmarks.

Necessity of Feedback Isolation. Our feedback score (7) contrasts the LLM under the prompt $h^{(t+1)}$, containing the previous rollout and its feedback, against the reference prompt $h_{\setminus f}^{(t+1)}$, which contains the same rollout but no feedback. Prior contrastive scores [SHLTZY24; Hüb+26] instead contrast against a reference prompt without any added context, which in our pipeline is the zero-shot prompt $h^{(1)}$. In this case, Bregman tilting at $\alpha = 1$ reduces to Context-Aware Decoding (CAD) [SHLTZY24] with score $\ln \pi(v \mid h^{(t+1)}, r_{<m}) - \ln \pi(v \mid h^{(1)}, r_{<m})$. Since $h^{(1)}$ omits both the previous rollout and its feedback, this score muddles the effect of showing the previous rollout with the effect of the feedback, rewarding tokens that agree with the previous rollout regardless of correctness.

To demonstrate this effect, we run Bregman tilting at $\alpha = 1$ with the zero-shot reference on AIME 2025 with Self-Refine, keeping each seed's round-1 rollouts and feedback unchanged. Along with mean pass@1 and pass@16, we measure how often the solver repeats a wrong answer when the reviewer model has flagged it as incorrect, as the reviewer is prompted to end each reply with a

Table 3: Runtime comparison with Qwen3-4B-Instruct-2507. Every method decodes the same AIME 2025 round-2 Self-Refine rollouts, with 1 or 8 chains decoded in parallel. We report the decoding time per generated token and its ratio to prompting, along with the time to process the prompts of a rollout before decoding and the peak GPU memory. CAD [SHLTZY24] is Bregman tilting at $\alpha = 1$ with the zero-shot reference prompt $h^{(1)}$.

Chains	Method	ms / token	× Prompting	Prompt processing (s)	Peak memory (GiB)
1	Prompting	5.46	1.00	0.20	12.4
	CAD ($\alpha = 1$, reference $h^{(1)}$)	5.96	1.09	0.22	15.2
	Bregman tilting, $\alpha = 1$	5.96	1.09	0.39	15.2
	Bregman tilting, $\alpha = 0$	6.07	1.11	0.39	15.2
8	Prompting	1.58	1.00	0.21	32.1
	CAD ($\alpha = 1$, reference $h^{(1)}$)	2.04	1.30	0.23	54.6
	Bregman tilting, $\alpha = 1$	2.09	1.32	0.41	54.6
	Bregman tilting, $\alpha = 0$	2.11	1.34	0.41	54.6

correctness verdict which is stripped from the returned feedback (Appendix B.5.2). Table 2 shows that with the zero-shot reference, Bregman tilting at $\alpha = 1$ is no better than input prompting in either metric, and it trails our reference prompt at every round. Interestingly, the zero-shot reference repeats flagged wrong answers as often as input prompting. Our reference prompt reduces this repetition pitfall at $\alpha = 1$, with the most reduction at $\alpha = 0$. Thus, the zero-shot reference biases the tilt towards copying the previous rollout, while our reference lets the tilt incorporate the feedback more effectively. See Figure 2 for round-by-round results.

Runtime. We compare the runtime of Bregman tilting, input prompting, and CAD, which is Bregman tilting with $\alpha = 1$ and the zero-shot reference prompt $h^{(1)}$. At each step, both CAD and Bregman tilting additionally evaluate the LLM on a reference prompt in the same batch to compute the guidance score. Bregman tilting at $\alpha = 0$ also solves for $\hat{\lambda}$ to construct the updated distribution (Remark 3.4). We use a set of fixed Self-Refine rollouts on AIME 2025, where at each position, each method computes its next-token PMF and samples a token, but decoding then continues from the recorded token rather than the sample. This makes the comparison fair, since every method decodes from the same prefixes. Table 3 reports the time per generated token on a single NVIDIA H100 GPU, for both 1 chain and 8 chains decoded in parallel. We observe that solving for $\hat{\lambda}$ is nearly free, as moving from $\alpha = 1$ to $\alpha = 0$ adds only 1–2%. The rest of the overhead comes from the reference prompt. Each decode has a fixed cost shared by all sequences in the batch, so the time per token of input prompting drops from 5.46 ms with 1 chain to 1.58 ms with 8 chains. The reference prompt doubles the batch size, adding about 0.5 ms per token in both cases, so the relative overhead for the guidance methods grows from about 1.1× to 1.3×. See Appendix B.9 for details.

6 Conclusion

We introduced Bregman tilting, an inference-time framework for guiding an LLM with text feedback in iterative problem solving. We defined a per-token feedback score as the change in its log-probability when the feedback was added to a prompt containing the prior attempt, isolating the effect of the

feedback, and incorporated the score into the updated next-token distribution via a mirror ascent step regularized by a Bregman divergence based on the Tsallis entropy parameterized by $\alpha \in [0, 1]$. Analysis showed that this update resulted in a generalized exponential tilt that damped the tilt strength on unlikely tokens under input prompting, and that at small step sizes it provably increased the true correctness probability whenever the feedback score covaried positively with correctness under a sharpened input-prompting distribution. Empirically, Bregman tilting outperformed input prompting in every setting, and $\alpha = 0$, which damped the most, performed best throughout.

Acknowledgments

This work was supported in part by NSF Grants 2019844, 2107037, 2112471, and 2505865, ONR Grant N00014-19-1-2566, the Machine Learning Lab (MLL) at UT Austin, the NSF AI Institute for Foundations of Machine Learning (IFML), the NSF AI-EDGE Institute, and the Wireless Networking and Communications Group (WNCG) Industrial Affiliates Program. We are grateful for computing support on the Vista GPU Cluster through the Center for Generative AI (CGAI) and the Texas Advanced Computing Center (TACC) at the University of Texas at Austin.

Contents of Appendix

A Theoretical Analysis **20**

A.1 Preliminaries 20

A.2 Proofs of Main Results 22

A.2.1 Section 3.1 22

A.2.2 Section 3.2 24

A.2.3 Section 4.1 29

A.2.4 Section 4.2 33

A.3 Solving for $\hat{\lambda}$ 36

B Experimental Details **38**

B.1 Recursive Self-Aggregation 38

B.2 Benchmarks and Grading 38

B.3 Solver and Feedback 39

B.4 Protocol, Metrics, and Statistics 40

B.5 Prompts 41

B.5.1 Code 41

B.5.2 Math 44

B.6 Results by Round 48

B.7 Reference Prompt by Round 50

B.8 MBPP+ Full Set 51

B.9 Runtime Experiment Details 51

A Theoretical Analysis

In this section, we prove every formal statement from the main body. The proofs are organized by the sections in which they appear. Appendix A.1 establishes general properties of Bregman divergences and of the mirror-ascent update (11), along with basic facts about the specific family of mirror maps ψ_α. We prove the main results in the following sections.

A.1 Preliminaries

Lemma A.1. *Let ψ be strictly convex and differentiable on a convex open set $U \subseteq \mathbb{R}^{\mathcal{V}}$, and let D_ψ be its Bregman divergence (10). Then the following hold for every $p, q, z \in U$.*

(i) $D_\psi(p, q) \geq 0$, with equality if and only if $p = q$.

(ii) The map $p \mapsto D_\psi(p, q)$ is strictly convex on U.

(iii) If $\tilde{\psi}(p) = \psi(p) + \langle w, p\rangle + c$ for a vector $w \in \mathbb{R}^{\mathcal{V}}$ and a scalar $c \in \mathbb{R}$, then $D_{\tilde{\psi}} = D_\psi$.

(iv) $D_\psi(p, q) + D_\psi(q, z) - D_\psi(p, z) = \langle \nabla\psi(z) - \nabla\psi(q), p - q\rangle$.

Proof of Lemma A.1. (i) Since ψ is strictly convex and differentiable, it lies strictly above each of its tangent lines, except at the point of reference [BV04], i.e., for every $p \neq q$,

$$\psi(p) > \psi(q) + \langle \nabla\psi(q), p - q\rangle.$$

Thus for $p \neq q$,

$$D_\psi(p, q) = \psi(p) - (\psi(q) + \langle \nabla\psi(q), p - q\rangle) > 0.$$

Further by definition,

$$D_\psi(q, q) = \psi(q) - (\psi(q) + \langle \nabla\psi(q), q - q\rangle) = 0.$$

(ii) As a function of p, $D_\psi(p, q)$ is the sum of the strictly convex function $\psi(p)$ and the affine function $-\psi(q) - \langle \nabla\psi(q), p - q\rangle$, so it is strictly convex.

(iii) Using the assumption, $\nabla\tilde{\psi}(q) = \nabla\psi(q) + w$. Thus,

$$\begin{aligned} D_{\tilde{\psi}}(p, q) &= \tilde{\psi}(p) - \tilde{\psi}(q) - \langle \nabla\tilde{\psi}(q), p - q\rangle \\ &= \psi(p) + \langle w, p\rangle - \psi(q) - \langle w, q\rangle - \langle \nabla\psi(q) + w, p - q\rangle \\ &= \psi(p) - \psi(q) + \langle w, p - q\rangle - \langle \nabla\psi(q), p - q\rangle - \langle w, p - q\rangle \\ &= \psi(p) - \psi(q) - \langle \nabla\psi(q), p - q\rangle \\ &= D_\psi(p, q). \end{aligned}$$

(iv) By definition of D_ψ, each $\psi(p), \psi(q), \psi(z)$ cancels out, giving

$$\begin{aligned} D_\psi(p, q) + D_\psi(q, z) - D_\psi(p, z) &= -\langle \nabla\psi(q), p - q\rangle - \langle \nabla\psi(z), q - z\rangle + \langle \nabla\psi(z), p - z\rangle \\ &= -\langle \nabla\psi(q), p - q\rangle + \langle \nabla\psi(z), p - q\rangle \\ &= \langle \nabla\psi(z) - \nabla\psi(q), p - q\rangle. \end{aligned}$$

□

Lemma A.2. *Let $\alpha \in [0,1]$. The mirror map ψ_α of Equations (12) and (13) is infinitely differentiable on $(0,\infty)^{\mathcal{V}}$, and for every $p \in (0,\infty)^{\mathcal{V}}$ and $v \in \mathcal{V}$,*

$$[\nabla \psi_\alpha(p)](v) = \begin{cases} -\dfrac{p(v)^{\alpha-1}}{1-\alpha} & \text{if } \alpha \in [0,1), \\ \ln p(v) + 1 & \text{if } \alpha = 1, \end{cases} \qquad \nabla^2 \psi_\alpha(p) = \operatorname{diag}\left(p(v)^{\alpha-2}\right)_{v\in\mathcal{V}}.$$

In particular, ψ_α is strictly convex on $(0,\infty)^{\mathcal{V}}$, and $\lim_{p(v)\to 0^+}[\nabla\psi_\alpha(p)](v) = -\infty$.

Proof of Lemma A.2. By Equations (12) and (13), $\psi_\alpha(p) = \sum_{v\in\mathcal{V}} \phi_\alpha(p(v))$, where

$$\phi_\alpha(x) \coloneqq -\frac{x^\alpha}{\alpha(1-\alpha)} \quad \text{for } \alpha \in (0,1), \qquad \phi_0(x) \coloneqq -\ln x, \qquad \phi_1(x) \coloneqq x \ln x.$$

Each ϕ_α is infinitely differentiable on $(0,\infty)$, so ψ_α is infinitely differentiable on $(0,\infty)^{\mathcal{V}}$, with $[\nabla\psi_\alpha(p)](v) = \phi_\alpha'(p(v))$ and $\nabla^2\psi_\alpha(p) = \operatorname{diag}\left(\phi_\alpha''(p(v))\right)_{v\in\mathcal{V}}$.

For $\alpha \in (0,1)$,

$$\phi_\alpha'(x) = -\frac{\alpha\, x^{\alpha-1}}{\alpha(1-\alpha)} = -\frac{x^{\alpha-1}}{1-\alpha}, \qquad \phi_\alpha''(x) = -\frac{(\alpha-1)\, x^{\alpha-2}}{1-\alpha} = x^{\alpha-2}. \tag{28}$$

At $\alpha = 0$,

$$\phi_0'(x) = -\frac{1}{x}, \qquad \phi_0''(x) = \frac{1}{x^2},$$

which is equivalent to (28) evaluated at $\alpha = 0$.

At $\alpha = 1$,

$$\phi_1'(x) = 1 + \ln(x), \qquad \phi_1''(x) = \frac{1}{x}.$$

This proves the form of the gradient and Hessian. Since the Hessian is diagonal with strictly positive entries, it is positive definite, implying strict convexity of ψ_α [BV04]. Lastly,

$$\lim_{p(v)\to 0^+} [\nabla\psi_\alpha(p)](v) = \begin{cases} \lim_{x\to 0^+} -\dfrac{x^{\alpha-1}}{1-\alpha} = -\infty & \text{if } \alpha \in [0,1), \\ \lim_{x\to 0^+} 1 + \ln(x) = -\infty & \text{if } \alpha = 1. \end{cases}$$

□

Lemma A.3. *For the family of mirror maps ψ_α for $\alpha \in [0,1]$ defined in Equations (12) and (13), ψ_0 and ψ_1 are the limits of ψ_α up to additive constants. Precisely, for every $p \in \Delta^+(\mathcal{V})$,*

$$\lim_{\alpha\to 0}\left(\psi_\alpha(p) + |\mathcal{V}|\left(\frac{1}{\alpha}+1\right)\right) = \psi_0(p), \qquad \lim_{\alpha\to 1}\left(\psi_\alpha(p) + \frac{1}{\alpha(1-\alpha)}\right) = \psi_1(p).$$

Proof of Lemma A.3. For the first limit, for every $v \in \mathcal{V}$, we have that

$$p(v)^\alpha = e^{\alpha \ln p(v)} = 1 + \alpha \ln p(v) + \mathcal{O}(\alpha^2).$$

Then, using the identity $\frac{1}{\alpha(1-\alpha)} = \frac{1}{\alpha} + \frac{1}{1-\alpha}$ and substituting these into Equation (12),

$$\begin{aligned}
\lim_{\alpha\to 0}\left(\psi_\alpha(p) + |\mathcal{V}|\left(\frac{1}{\alpha}+1\right)\right) &= \lim_{\alpha\to 0}\left(|\mathcal{V}|\left(\frac{1}{\alpha}+1\right) - \left(\frac{1}{\alpha}+\frac{1}{1-\alpha}\right)\sum_{v\in\mathcal{V}}\left(1+\alpha\ln p(v)+\mathcal{O}(\alpha^2)\right)\right) \\
&= \lim_{\alpha\to 0}\left(|\mathcal{V}|\left(\frac{1}{\alpha}+1\right) - |\mathcal{V}|\left(\frac{1}{\alpha}+\frac{1}{1-\alpha}\right) - \frac{1}{1-\alpha}\sum_{v\in\mathcal{V}}\ln p(v) + \mathcal{O}(\alpha)\right) \\
&= \lim_{\alpha\to 0}\left(|\mathcal{V}|\left(1-\frac{1}{1-\alpha}\right) - \sum_{v\in\mathcal{V}}\ln p(v) + \mathcal{O}(\alpha)\right) \\
&= -\sum_{v\in\mathcal{V}}\ln p(v).
\end{aligned}$$

For the second limit, let $\epsilon \coloneqq 1-\alpha$. For every $v\in\mathcal{V}$,

$$p(v)^{1-\epsilon} = p(v)\,e^{-\epsilon\ln p(v)} = p(v) - \epsilon\,p(v)\ln p(v) + \mathcal{O}(\epsilon^2).$$

Substituting this into Equation (12), and using $\sum_v p(v) = 1$,

$$\begin{aligned}
\lim_{\alpha\to 1}\left(\psi_\alpha(p) + \frac{1}{\alpha(1-\alpha)}\right) &= \lim_{\alpha\to 1}\left(\frac{1}{\alpha(1-\alpha)} - \frac{1}{\alpha(1-\alpha)}\sum_{v\in\mathcal{V}}p(v)^\alpha\right) \\
&= \lim_{\epsilon\to 0}\left(\frac{1}{(1-\epsilon)\epsilon}\left(1-\sum_{v\in\mathcal{V}}p(v)^{1-\epsilon}\right)\right) \\
&= \lim_{\epsilon\to 0}\left(\frac{1}{(1-\epsilon)\epsilon}\left[1-\left(\sum_{v\in\mathcal{V}}p(v)\right)+\epsilon\left(\sum_{v\in\mathcal{V}}p(v)\ln p(v)\right)+\mathcal{O}(\epsilon^2)\right]\right) \\
&= \lim_{\epsilon\to 0}\left(\frac{1}{1-\epsilon}\sum_{v\in\mathcal{V}}p(v)\ln p(v) + \mathcal{O}(\epsilon)\right) \\
&= \sum_{v\in\mathcal{V}}p(v)\ln p(v).
\end{aligned}$$

□

A.2 Proofs of Main Results

A.2.1 Section 3.1

Proof of Proposition 3.1

Proposition. *For every autoregressive PMF ρ,*

$$A(\rho) - A(\pi^{(t+1)}) = \sum_{m=1}^{M}\mathbb{E}_{\mathsf{R}\sim\rho}\left[\bar{a}\left(\rho\,;\mathsf{R}_{<m}\right) - \bar{a}(\pi^{(t+1)}\,;\mathsf{R}_{<m})\right].$$

Proof of Proposition 3.1. Fix a prefix x of length less than M. Conditioning on the first token v that $\pi^{(t+1)}$ generates after x,

$$A(\pi^{(t+1)}\,;x) = \sum_{v\in\mathcal{V}} \pi^{(t+1)}(v \mid x)\, A\left(\pi^{(t+1)}\,;(x,v)\right).$$

By the definition (4), $A\left(\pi^{(t+1)}\,;(x,v)\right) = a(v\,;x)$, so by the definition (5),

$$A(\pi^{(t+1)}\,;x) = \sum_{v\in\mathcal{V}} \pi^{(t+1)}(v \mid x)\, a(v\,;x) = \bar{a}(\pi^{(t+1)}\,;x). \tag{29}$$

We now show the result.

$$\begin{aligned} A(\rho) - A(\pi^{(t+1)}) &= \mathbb{E}_{\mathsf{R}\sim\rho}\left[\mathbb{1}\{\mathsf{R}\in\mathcal{A}\}\right] - A(\pi^{(t+1)}) \\ &= \mathbb{E}_{\mathsf{R}\sim\rho}\left[\mathbb{1}\{\mathsf{R}\in\mathcal{A}\} - A(\pi^{(t+1)})\right], \end{aligned} \tag{30}$$

since $A(\pi^{(t+1)})$ is a constant.

Since $\mathsf{R} = \mathsf{R}_{<M+1}$ is a complete rollout, we can write $\mathbb{1}\{\mathsf{R}\in\mathcal{A}\} = A(\pi^{(t+1)}\,;\mathsf{R}_{<M+1})$. Similarly, since $\mathsf{R}_{<1}$ is empty, we have that $A(\pi^{(t+1)}) = A(\pi^{(t+1)}\,;\mathsf{R}_{<1})$. Thus,

$$\mathbb{E}_{\mathsf{R}\sim\rho}\left[\mathbb{1}\{\mathsf{R}\in\mathcal{A}\} - A(\pi^{(t+1)})\right] = \mathbb{E}_{\mathsf{R}\sim\rho}\left[A(\pi^{(t+1)}\,;\mathsf{R}_{<M+1}) - A(\pi^{(t+1)}\,;\mathsf{R}_{<1})\right]. \tag{31}$$

We then telescope over the prefixes of R, giving that

$$\begin{aligned} &\mathbb{E}_{\mathsf{R}\sim\rho}\left[A(\pi^{(t+1)}\,;\mathsf{R}_{<M+1}) - A(\pi^{(t+1)}\,;\mathsf{R}_{<1})\right] \\ &\quad= \mathbb{E}_{\mathsf{R}\sim\rho}\left[\sum_{m=1}^{M} A(\pi^{(t+1)}\,;\mathsf{R}_{<m+1}) - A(\pi^{(t+1)}\,;\mathsf{R}_{<m})\right] \\ &\quad= \mathbb{E}_{\mathsf{R}\sim\rho}\left[\sum_{m=1}^{M} A\left(\pi^{(t+1)}\,;(\mathsf{R}_{<m},\mathsf{R}_m)\right) - A(\pi^{(t+1)}\,;\mathsf{R}_{<m})\right] \\ &\quad= \sum_{m=1}^{M} \mathbb{E}_{\mathsf{R}\sim\rho}\left[A\left(\pi^{(t+1)}\,;(\mathsf{R}_{<m},\mathsf{R}_m)\right) - A(\pi^{(t+1)}\,;\mathsf{R}_{<m})\right] \\ &\quad= \sum_{m=1}^{M} \mathbb{E}_{\mathsf{R}\sim\rho}\left[a\left(\mathsf{R}_m\,;\mathsf{R}_{<m}\right) - A(\pi^{(t+1)}\,;\mathsf{R}_{<m})\right] \\ &\quad= \sum_{m=1}^{M} \mathbb{E}_{\mathsf{R}\sim\rho}\left[\mathbb{E}_{\mathsf{R}|\mathsf{R}_{<m}}\left[a\left(\mathsf{R}_m\,;\mathsf{R}_{<m}\right) - A(\pi^{(t+1)}\,;\mathsf{R}_{<m})\right]\right], \end{aligned} \tag{32}$$

where the last line uses the law of iterated expectation, conditioning on the prefix $\mathsf{R}_{<m}$ of $\mathsf{R}\sim\rho$. Since $A(\pi^{(t+1)}\,;\mathsf{R}_{<m})$ depends only on $\mathsf{R}_{<m}$, which has length $m-1<M$, Equation (29) gives that

$$\mathbb{E}_{\mathsf{R}|\mathsf{R}_{<m}}\left[A(\pi^{(t+1)}\,;\mathsf{R}_{<m})\right] = A(\pi^{(t+1)}\,;\mathsf{R}_{<m}) = \bar{a}(\pi^{(t+1)}\,;\mathsf{R}_{<m}).$$

Further, since $a\left(\mathsf{R}_m\,;\mathsf{R}_{<m}\right)$ only depends on $\mathsf{R}_{<m}$ and R_m, and since $\mathsf{R}_m \mid \mathsf{R}_{<m}$ is distributed according to $\rho(\,\cdot \mid \mathsf{R}_{<m})$ for $\mathsf{R}\sim\rho$ by definition,

$$\mathbb{E}_{\mathsf{R}|\mathsf{R}_{<m}}\left[a\left(\mathsf{R}_m\,;\mathsf{R}_{<m}\right)\right] = \mathbb{E}_{\mathsf{R}_m\sim\rho(\,\cdot\,|\mathsf{R}_{<m})}\left[a\left(\mathsf{R}_m\,;\mathsf{R}_{<m}\right)\right] = \bar{a}(\rho\,;\mathsf{R}_{<m}).$$

Applying these two simplifications to the chain of equalities above,

$$\sum_{m=1}^{M} \mathbb{E}_{\mathsf{R}\sim\rho}\left[\mathbb{E}_{\mathsf{R}|\mathsf{R}_{<m}}\left[a\left(\mathsf{R}_m\,;\mathsf{R}_{<m}\right) - A(\pi^{(t+1)}\,;\mathsf{R}_{<m})\right]\right] = \sum_{m=1}^{M} \mathbb{E}_{\mathsf{R}\sim\rho}\left[\bar{a}(\rho\,;\mathsf{R}_{<m}) - \bar{a}(\pi^{(t+1)}\,;\mathsf{R}_{<m})\right]. \tag{33}$$

Then chaining the string of equalities in (30), (31), (32), (33) gives the result.

□

A.2.2 Section 3.2

Proof of Proposition 3.2

Proposition. *Consider the family of mirror maps ψ_α defined by Equations (12), (13), and arbitrary PMFs $p, q \in \Delta^+(\mathcal{V})$.*

(i) On the interior $\alpha \in (0,1)$, the Bregman divergence of ψ_α is

$$D_{\psi_\alpha}(p,q) = \frac{1}{\alpha(1-\alpha)} \sum_{v\in\mathcal{V}} \left(q(v)^\alpha - p(v)^\alpha + \alpha\, q(v)^{\alpha-1}\left(p(v)-q(v)\right)\right).$$

(ii) The Bregman divergence D_{ψ_α} is a continuous function of α at the endpoints $\{0,1\}$, with limits

$$\lim_{\alpha\to 0} D_{\psi_\alpha}(p,q) = D_{\psi_0}(p,q) = \sum_{v\in\mathcal{V}} \left(\frac{p(v)}{q(v)} - \ln\frac{p(v)}{q(v)} - 1\right),$$

$$\lim_{\alpha\to 1} D_{\psi_\alpha}(p,q) = D_{\psi_1}(p,q) = \sum_{v\in\mathcal{V}} p(v)\ln\frac{p(v)}{q(v)} = \mathrm{KL}\left(p \,\|\, q\right).$$

Proof of Proposition 3.2. Since $p, q \in \Delta^+(\mathcal{V})$ have full support, ψ_α is differentiable at both for every $\alpha \in [0,1]$ (Lemma A.2).

(i) Let $\alpha \in (0,1)$. By Lemma A.2,

$$[\nabla\psi_\alpha(q)](v) = -\frac{q(v)^{\alpha-1}}{1-\alpha}.$$

Substituting this and Equation (12) into the definition (10),

$$\begin{aligned} D_{\psi_\alpha}(p,q) &= \psi_\alpha(p) - \psi_\alpha(q) - \langle\nabla\psi_\alpha(q), p-q\rangle \\ &= -\frac{1}{\alpha(1-\alpha)}\sum_{v\in\mathcal{V}} p(v)^\alpha + \frac{1}{\alpha(1-\alpha)}\sum_{v\in\mathcal{V}} q(v)^\alpha + \frac{1}{1-\alpha}\sum_{v\in\mathcal{V}} q(v)^{\alpha-1}\left(p(v)-q(v)\right) \\ &= \frac{1}{\alpha(1-\alpha)}\sum_{v\in\mathcal{V}}\left(q(v)^\alpha - p(v)^\alpha + \alpha\, q(v)^{\alpha-1}\left(p(v)-q(v)\right)\right). \end{aligned}$$

(ii) We first compute the Bregman divergences of the endpoints. By Lemma A.2, $[\nabla\psi_0(q)](v) = -1/q(v)$. Substituting this and ψ_0 from Equation (13) into the definition (10),

$$\begin{aligned} D_{\psi_0}(p,q) &= -\sum_{v\in\mathcal{V}} \ln p(v) + \sum_{v\in\mathcal{V}} \ln q(v) + \sum_{v\in\mathcal{V}} \frac{p(v)-q(v)}{q(v)} \\ &= \sum_{v\in\mathcal{V}} \left(\frac{p(v)}{q(v)} - \ln \frac{p(v)}{q(v)} - 1 \right). \end{aligned}$$

By Lemma A.2, $[\nabla\psi_1(q)](v) = \ln q(v) + 1$. Substituting this and ψ_1 from Equation (13) into the definition (10),

$$\begin{aligned} D_{\psi_1}(p,q) &= \sum_{v\in\mathcal{V}} p(v)\ln p(v) - \sum_{v\in\mathcal{V}} q(v)\ln q(v) - \sum_{v\in\mathcal{V}} (\ln q(v)+1)\,(p(v)-q(v)) \\ &= \sum_{v\in\mathcal{V}} p(v)\ln p(v) - \sum_{v\in\mathcal{V}} p(v)\ln q(v) - \sum_{v\in\mathcal{V}} (p(v)-q(v)) \\ &= \sum_{v\in\mathcal{V}} p(v)\ln \frac{p(v)}{q(v)}, \end{aligned}$$

where the last step uses $\sum_v p(v) = \sum_v q(v) = 1$.

It remains to compute the limits. Define the two functions of α

$$w_0(\alpha) = |\mathcal{V}| \left(\frac{1}{\alpha} + 1 \right) \quad \text{and} \quad w_1(\alpha) = \frac{1}{\alpha(1-\alpha)}.$$

Fix $k \in \{0,1\}$. By the definition of the Bregman divergence (10),

$$D_{\psi_\alpha}(p,q) = (\psi_\alpha(p) + w_k(\alpha)) - (\psi_\alpha(q) + w_k(\alpha)) - \langle \nabla\psi_\alpha(q), p-q\rangle,$$

where we add and subtract $w_k(\alpha)$. We aim to show $D_{\psi_\alpha}(p,q)$ converges to $D_{\psi_k}(p,q)$ as $\alpha \to k$.

By Lemma A.3, the first two terms converge to the value at the endpoint,

$$\begin{aligned} \lim_{\alpha\to k} (\psi_\alpha(p) + w_k(\alpha) - \psi_\alpha(q) - w_k(\alpha)) &= \lim_{\alpha\to k} (\psi_\alpha(p) + w_k(\alpha)) - \lim_{\alpha\to k} (\psi_\alpha(q) + w_k(\alpha)) \\ &= \psi_k(p) - \psi_k(q). \end{aligned} \tag{34}$$

It then suffices to show that $\lim_{\alpha\to k}\langle \nabla\psi_\alpha(q), p-q\rangle = \langle \nabla\psi_k(q), p-q\rangle$. We consider both cases.

Case 1: $k = 0$.

By Lemma A.2,

$$\lim_{\alpha\to 0} [\nabla\psi_\alpha(q)](v) = \lim_{\alpha\to 0} -\frac{q(v)^{\alpha-1}}{1-\alpha} = -\frac{1}{q(v)} = [\nabla\psi_0(q)](v).$$

Thus,

$$\lim_{\alpha\to 0} \langle \nabla\psi_\alpha(q), p-q\rangle = \langle \nabla\psi_0(q), p-q\rangle.$$

Case 2: $k = 1$.

First, since $\langle \mathbf{1}, p - q \rangle = 0$, we have that for every $c \in \mathbb{R}$,

$$\langle \nabla \psi_\alpha(q), p - q \rangle = \langle \nabla \psi_\alpha(q) + c\mathbf{1}, p - q \rangle.$$

Take $c = \frac{1}{1-\alpha}$ and let $\epsilon \coloneqq 1 - \alpha$. Since $q(v)^{-\epsilon} = 1 - \epsilon \ln q(v) + \mathcal{O}(\epsilon^2)$,

$$\begin{aligned}
\lim_{\alpha \to 1} \left([\nabla \psi_\alpha(q)](v) + \frac{1}{1-\alpha} \right) &= \lim_{\alpha \to 1} \left(-\frac{q(v)^{\alpha-1}}{1-\alpha} + \frac{1}{1-\alpha} \right) \\
&= \lim_{\epsilon \to 0} \frac{1 - q(v)^{-\epsilon}}{\epsilon} \\
&= \ln q(v).
\end{aligned}$$

Then since $\nabla \psi_1(q) = \ln q + \mathbf{1}$ by Lemma A.2,

$$\begin{aligned}
\lim_{\alpha \to 1} \langle \nabla \psi_\alpha(q), p - q \rangle &= \lim_{\alpha \to 1} \langle \nabla \psi_\alpha(q) + \frac{1}{1-\alpha} \cdot \mathbf{1}, p - q \rangle \\
&= \langle \ln q, p - q \rangle \\
&= \langle \ln q + \mathbf{1}, p - q \rangle \\
&= \langle \nabla \psi_1(q), p - q \rangle.
\end{aligned}$$

Combining the result of these cases with (34), for $k \in \{0, 1\}$,

$$\begin{aligned}
\lim_{\alpha \to k} D_{\psi_\alpha}(p, q) &= \lim_{\alpha \to k} \left(\psi_\alpha(p) + w_k(\alpha) - \psi_\alpha(q) - w_k(\alpha) - \langle \nabla \psi_\alpha(q), p - q \rangle \right) \\
&= \lim_{\alpha \to k} \left(\psi_\alpha(p) + w_k(\alpha) \right) - \lim_{\alpha \to k} \left(\psi_\alpha(q) + w_k(\alpha) \right) - \lim_{\alpha \to k} \left(\langle \nabla \psi_\alpha(q), p - q \rangle \right) \\
&= \psi_k(p) - \psi_k(q) - \langle \nabla \psi_k(q), p - q \rangle \\
&= D_{\psi_k}(p, q).
\end{aligned}$$

□

Proof of Proposition 3.3

Proposition. *For $\alpha \in [0, 1]$ and $\eta > 0$, Equation (11) with the mirror map ψ_α has a unique solution, which we denote $\hat{p}_m(\,\cdot\,; \alpha)$. Moreover, there exists a scalar $\hat{\lambda}$ such that $(p, \lambda) = (\hat{p}_m(\,\cdot\,; \alpha), \hat{\lambda})$ is the unique pair in $\Delta^+(\mathcal{V}) \times \mathbb{R}$ which satisfies*

$$\nabla \psi_\alpha(p) = \nabla \psi_\alpha(\pi_m) + \eta\, s_m - \lambda \mathbf{1} \quad \text{(Equation 17)}.$$

Proof of Proposition 3.3. Fix $\alpha \in [0, 1]$ and $\eta > 0$. We first solve (17) for p at a fixed λ, which reduces (17) to a scalar equation in λ. We then show that this scalar equation has a unique solution $\hat{\lambda}$, so that (17) has a unique solution in $\Delta^+(\mathcal{V}) \times \mathbb{R}$. Finally, we show that the first entry of this solution coincides with the unique maximizer of the objective in (11) over $\Delta^+(\mathcal{V})$.

Let $(p, \lambda) \in \Delta^+(\mathcal{V}) \times \mathbb{R}$. We solve (17) for p in terms of λ. For $\alpha = 1$, the condition in (17) can be written as

$$p(v) = \pi_m(v)\, e^{\eta\, s_m(v) - \lambda} \qquad \text{(cf. Equation 18)}.$$

For $\alpha \in [0,1)$, the condition in (17) can be written as

$$p(v)^{\alpha-1} = (1-\alpha)\,(\lambda - b(v))\,, \qquad \text{where} \quad b(v) \coloneqq \eta\, s_m(v) - \frac{\pi_m(v)^{\alpha-1}}{1-\alpha}.$$

Since the left-hand side is positive, this can hold only if $\lambda > B \coloneqq \max_{v\in\mathcal{V}} b(v)$, in which case it is equivalent to

$$p(v) = \big((1-\alpha)\,(\lambda - b(v))\,\big)^{-\frac{1}{1-\alpha}} = \frac{\pi_m(v)}{(1+(1-\alpha)\pi_m(v)^{1-\alpha}\,(\lambda - \eta\, s_m(v)))^{\frac{1}{1-\alpha}}} \qquad \text{(cf. Equation 19)}.$$

To treat both cases at once, we write the set of feasible λ as $\Lambda_1 \coloneqq \mathbb{R}$ and $\Lambda_\alpha \coloneqq (B,\infty)$ for $\alpha \in [0,1)$. Then, for any $\lambda \in \Lambda_\alpha$, we write the corresponding solution to (17) over the set of strictly positive vectors as $\nu_\alpha(\,\cdot\,;\lambda) \in (0,\infty)^{\mathcal{V}}$, with total mass $Z_\alpha(\lambda)$, defined by

$$\nu_\alpha(v\,;\lambda) \coloneqq \begin{cases} \pi_m(v)\, e^{\eta\, s_m(v) - \lambda} & \text{if } \alpha = 1, \\ \big((1-\alpha)\,(\lambda - b(v))\,\big)^{-\frac{1}{1-\alpha}} & \text{if } \alpha \in [0,1), \end{cases} \qquad Z_\alpha(\lambda) \coloneqq \sum_{v\in\mathcal{V}} \nu_\alpha(v\,;\lambda). \tag{35}$$

Thus, $(p,\lambda) \in \Delta^+(\mathcal{V}) \times \mathbb{R}$ satisfies (17) if and only if $\lambda \in \Lambda_\alpha$ and $p = \nu_\alpha(\,\cdot\,;\lambda)$. Since p is a valid PMF, this requires $Z_\alpha(\lambda) = 1$. Conversely, if $\lambda \in \Lambda_\alpha$ and $Z_\alpha(\lambda) = 1$, then $\nu_\alpha(\,\cdot\,;\lambda)$ is a PMF in $\Delta^+(\mathcal{V})$ and the pair $(\nu_\alpha(\,\cdot\,;\lambda),\lambda)$ satisfies (17). Thus, the solutions of (17) in $\Delta^+(\mathcal{V}) \times \mathbb{R}$ are exactly the pairs

$$(p,\lambda) = \big(\nu_\alpha(\,\cdot\,;\lambda),\ \lambda\big) \quad \text{with} \quad \lambda \in \Lambda_\alpha \ \text{ and } \ Z_\alpha(\lambda) = 1.$$

We next show that the equation $Z_\alpha(\lambda) = 1$ has a unique solution $\hat{\lambda} \in \Lambda_\alpha$. We proceed in two cases.

Case 1: $\alpha = 1$.

By (35), for all $\lambda \in \mathbb{R}$,

$$Z_1(\lambda) = e^{-\lambda} \left(\sum_{v\in\mathcal{V}} \pi_m(v)\, e^{\eta\, s_m(v)} \right).$$

Thus, $\hat{\lambda} = \ln \sum_{v\in\mathcal{V}} \pi_m(v) e^{\eta\, s_m(v)}$ satisfies $Z_1(\hat{\lambda}) = 1$. Since Z_1 is strictly decreasing over $\mathbb{R}$, this solution is unique.

Case 2: $\alpha \in [0,1)$.

By (35), Z_α is a continuous function of λ on $\Lambda_\alpha = (B,\infty)$. Consider the following two values of λ:

$$\lambda_{\text{lo}} \coloneqq B + \frac{1}{1-\alpha}, \qquad \lambda_{\text{hi}} \coloneqq \max_{v\in\mathcal{V}}\ \eta\, s_m(v). \tag{36}$$

Since $\pi_m(v) \le 1$, we have $b(v) \le \eta\, s_m(v) - \frac{1}{1-\alpha}$ for all $v \in \mathcal{V}$, and thus $B = \max_{v\in\mathcal{V}} b(v) < \lambda_{\text{lo}} \le \lambda_{\text{hi}}$. Further,

$$Z_\alpha(\lambda_{\text{hi}}) \le 1 \le Z_\alpha(\lambda_{\text{lo}}).$$

The first inequality follows since

$$\begin{aligned} Z_\alpha(\lambda_{\mathrm{hi}}) &= \sum_{v\in\mathcal{V}} \left((1-\alpha)\left(\lambda_{\mathrm{hi}} - b(v)\right)\right)^{-\frac{1}{1-\alpha}} \\ &\le \sum_{v\in\mathcal{V}} \left((1-\alpha)\left(\eta s_m(v) - b(v)\right)\right)^{-\frac{1}{1-\alpha}} \\ &= \sum_{v\in\mathcal{V}} \pi_m(v) = 1, \end{aligned}$$

and the second since keeping only the summand of $Z_\alpha(\lambda_{\mathrm{lo}})$ with $b(v) = B$ gives

$$Z_\alpha(\lambda_{\mathrm{lo}}) \ge \left((1-\alpha)\left(B + \frac{1}{1-\alpha} - B\right)\right)^{-\frac{1}{1-\alpha}} = 1.$$

Thus, by the intermediate value theorem, there must exist $\hat{\lambda} \in [\lambda_{\mathrm{lo}}, \lambda_{\mathrm{hi}}] \subset \Lambda_\alpha$ such that $Z_\alpha(\hat{\lambda}) = 1$. Lastly, for all $\lambda \in \Lambda_\alpha$, differentiating (35) with respect to λ gives

$$\frac{d}{d\lambda} Z_\alpha(\lambda) = -\sum_{v\in\mathcal{V}} \nu_\alpha(v\,;\lambda)^{2-\alpha} < 0. \tag{37}$$

Thus, Z_α is strictly decreasing on Λ_α, so this solution to $Z_\alpha(\lambda) = 1$ must be unique.

In both cases, $Z_\alpha(\lambda) = 1$ has a unique solution $\hat{\lambda} \in \Lambda_\alpha$. Thus, by the above, $(p, \lambda) = (\nu_\alpha(\,\cdot\,;\hat{\lambda}), \hat{\lambda})$ is the unique pair in $\Delta^+(\mathcal{V}) \times \mathbb{R}$ which satisfies (17).

Going forward, we define the first entry of the solution to (17) as

$$\hat{p} \coloneqq \nu_\alpha(\,\cdot\,;\hat{\lambda}).$$

We lastly show that $\hat{p}$ is the unique solution of (11). For $p \in \Delta^+(\mathcal{V})$, we define $G(p)$ as the objective in (11) with mirror map ψ_α,

$$G(p) \coloneqq \langle s_m, p\rangle - \frac{1}{\eta} D_{\psi_\alpha}(p, \pi_m).$$

We must show that $\hat{p}$ is the unique maximizer of G over $\Delta^+(\mathcal{V})$.

By Lemma A.2, ψ_α is differentiable and strictly convex on the convex open set $(0,\infty)^{\mathcal{V}}$, which contains $\Delta^+(\mathcal{V})$. Then, by the three-point identity in Lemma A.1 (iv), for any $q \in \Delta^+(\mathcal{V})$,

$$D_{\psi_\alpha}(q, \hat{p}) + D_{\psi_\alpha}(\hat{p}, \pi_m) - D_{\psi_\alpha}(q, \pi_m) = \langle \nabla\psi_\alpha(\pi_m) - \nabla\psi_\alpha(\hat{p}), q - \hat{p}\rangle.$$

Rearranging,

$$D_{\psi_\alpha}(q, \pi_m) - D_{\psi_\alpha}(\hat{p}, \pi_m) = D_{\psi_\alpha}(q, \hat{p}) + \langle \nabla\psi_\alpha(\hat{p}) - \nabla\psi_\alpha(\pi_m), q - \hat{p}\rangle.$$

Then since $(\hat{p}, \hat{\lambda})$ satisfies (17), applying $\nabla\psi_\alpha(\hat{p}) - \nabla\psi_\alpha(\pi_m) = \eta\, s_m - \hat{\lambda}\mathbf{1}$ gives that

$$\begin{aligned} D_{\psi_\alpha}(q, \pi_m) - D_{\psi_\alpha}(\hat{p}, \pi_m) &= D_{\psi_\alpha}(q, \hat{p}) + \langle \eta\, s_m - \hat{\lambda}\mathbf{1}, q - \hat{p}\rangle \\ &= D_{\psi_\alpha}(q, \hat{p}) + \langle \eta\, s_m, q - \hat{p}\rangle, \end{aligned}$$

since $\langle \hat{\lambda}\mathbf{1}, q - \hat{p}\rangle = \hat{\lambda}\big(\langle \mathbf{1}, q\rangle - \langle \mathbf{1}, \hat{p}\rangle\big) = 0$, as $q, \hat{p} \in \Delta^+(\mathcal{V})$ are both valid PMFs which sum to one. Further rearranging gives that

$$\langle s_m, \hat{p}\rangle - \frac{1}{\eta} D_{\psi_\alpha}(\hat{p}, \pi_m) - \left(\langle s_m, q\rangle - \frac{1}{\eta} D_{\psi_\alpha}(q, \pi_m)\right) = \frac{1}{\eta} D_{\psi_\alpha}(q, \hat{p}).$$

Applying the definition of G then gives

$$G(\hat{p}) - G(q) = \frac{1}{\eta} D_{\psi_\alpha}(q, \hat{p}).$$

By Lemma A.1 (i), $D_{\psi_\alpha}(q, \hat{p}) > 0$ for all $q \neq \hat{p}$, so

$$G(\hat{p}) > G(q) \quad \forall q \in \Delta^+(\mathcal{V}),\ q \neq \hat{p}.$$

Thus, $\hat{p}$ uniquely maximizes G over $\Delta^+(\mathcal{V})$, i.e., (11) has the unique solution $\hat{p}_m(\,\cdot\,;\alpha) = \hat{p} = \nu_\alpha(\,\cdot\,;\hat{\lambda})$, and using the results above, $(p, \lambda) = (\hat{p}_m(\,\cdot\,;\alpha), \hat{\lambda})$ is the unique solution to (17) over $\Delta^+(\mathcal{V}) \times \mathbb{R}$. □

A.2.3 Section 4.1

Proof of Proposition 4.1

Proposition. *Let $\alpha \in [0, 1]$. For every $p \in \Delta^+(\mathcal{V})$, $\nabla^2 \psi_\alpha(p) = \operatorname{diag}\big(p(v)^{\alpha-2}\big)_{v\in\mathcal{V}}$, and for every perturbation $u \in \mathbb{R}^{\mathcal{V}}$ with $\sum_v u(v) = 0$,*

$$D_{\psi_\alpha}(\pi_m + u, \pi_m) = \frac{1}{2} \sum_{v\in\mathcal{V}} \pi_m(v)^\alpha \left(\frac{u(v)}{\pi_m(v)}\right)^2 + \mathcal{O}\left(\|u\|^3\right) \quad \text{as } u \to 0.$$

Proof of Proposition 4.1. Since $\pi_m \in \Delta^+(\mathcal{V})$ has full support, $\pi_m + u \in (0, \infty)^{\mathcal{V}}$ for every sufficiently small u. By Lemma A.2, ψ_α is infinitely differentiable on $(0, \infty)^{\mathcal{V}}$, so by Taylor's theorem around π_m,

$$\psi_\alpha(\pi_m + u) = \psi_\alpha(\pi_m) + \langle \nabla\psi_\alpha(\pi_m), u\rangle + \frac{1}{2} u^\top \nabla^2 \psi_\alpha(\pi_m)\, u + \mathcal{O}\left(\|u\|^3\right).$$

Thus, by the definition (10), and using the fact that $\nabla^2\psi_\alpha(\pi_m) = \operatorname{diag}\left(\pi_m(v)^{\alpha-2}\right)_{v\in\mathcal{V}}$ from Lemma A.2,

$$\begin{aligned}
D_{\psi_\alpha}(\pi_m + u, \pi_m) &= \psi_\alpha(\pi_m + u) - \psi_\alpha(\pi_m) - \langle \nabla\psi_\alpha(\pi_m), u\rangle \\
&= \frac{1}{2} u^\top \nabla^2 \psi_\alpha(\pi_m)\, u + \mathcal{O}\left(\|u\|^3\right) \\
&= \frac{1}{2} \sum_{v\in\mathcal{V}} \pi_m(v)^{\alpha-2}\, u(v)^2 + \mathcal{O}\left(\|u\|^3\right) \\
&= \frac{1}{2} \sum_{v\in\mathcal{V}} \pi_m(v)^\alpha \left(\frac{u(v)}{\pi_m(v)}\right)^2 + \mathcal{O}\left(\|u\|^3\right).
\end{aligned}$$

□

Proof of Proposition 4.2

Proposition. *Let $\alpha \in [0,1]$ and $\eta > 0$, and fix the score s_m. For every token $v \in \mathcal{V}$, the log-ratio of the mass assigned to v by the update $\hat{p}_m(v\,;\alpha)$ relative to $\pi_m(v)$ satisfies*

$$\ln \frac{\hat{p}_m(v\,;\alpha)}{\pi_m(v)} = \pi_m(v)^{1-\alpha}\left(\eta\, s_m(v) - \hat{\lambda}\right) + \mathcal{O}\left(\pi_m(v)^{2(1-\alpha)}\right) \quad \text{as } \pi_m(v) \to 0.$$

Proof of Proposition 4.2. Fix $v \in \mathcal{V}$, the step size $\eta > 0$ and the score s_m, and write $\hat{p}_m = \hat{p}_m(\,\cdot\,;\alpha)$ and $\psi = \psi_\alpha$. By Proposition 3.3, Equation (17) holds at $(p, \lambda) = (\hat{p}_m, \hat{\lambda})$, so for every $u \in \mathcal{V}$,

$$[\nabla\psi(\hat{p}_m)](u) - [\nabla\psi(\pi_m)](u) = \eta\, s_m(u) - \hat{\lambda}. \tag{38}$$

We first show that $\eta\, s_m(v) - \hat{\lambda}$ is bounded independently of $\pi_m(v)$. Since $\hat{p}_m$ and π_m both sum to one, there exist $u, u' \in \mathcal{V}$ with

$$\hat{p}_m(u) \geq \pi_m(u) \qquad \text{and} \qquad \hat{p}_m(u') \leq \pi_m(u').$$

By Lemma A.2, $[\nabla\psi(p)](u)$ depends only on $p(u)$ and is strictly increasing in it. Thus, the left-hand side of Equation (38) is non-negative at u and non-positive at u', i.e.,

$$\eta\, s_m(u') \leq \hat{\lambda} \leq \eta\, s_m(u).$$

Hence $\eta \min_w s_m(w) \leq \hat{\lambda} \leq \eta \max_w s_m(w)$, and so

$$\left|\eta\, s_m(v) - \hat{\lambda}\right| \leq C \coloneqq \eta\left(\max_{w\in\mathcal{V}} s_m(w) - \min_{w\in\mathcal{V}} s_m(w)\right), \tag{39}$$

where C does not depend on $\pi_m(v)$. We now consider two cases.

Case 1: $\alpha = 1$.

Dividing Equation (18) by $\pi_m(v)$ and taking logarithms,

$$\ln \frac{\hat{p}_m(v)}{\pi_m(v)} = \eta\, s_m(v) - \hat{\lambda} = \pi_m(v)^{1-\alpha}\left(\eta\, s_m(v) - \hat{\lambda}\right),$$

since $\pi_m(v)^0 = 1$. Thus Equation (21) holds with no remainder.

Case 2: $\alpha \in [0, 1)$.

Define

$$y \coloneqq (1-\alpha)\,\pi_m(v)^{1-\alpha}\left(\eta\, s_m(v) - \hat{\lambda}\right).$$

Then the base of the power in Equation (19) is $1 - y$, and dividing Equation (19) by $\pi_m(v)$ gives

$$\frac{\hat{p}_m(v)}{\pi_m(v)} = (1-y)^{-\frac{1}{1-\alpha}}.$$

Since $\hat{p}_m(v) > 0$, the base $1 - y$ is positive. Taking logarithms,

$$\ln \frac{\hat{p}_m(v)}{\pi_m(v)} = -\frac{\ln(1-y)}{1-\alpha}.$$

Since $\frac{y}{1-\alpha} = \pi_m(v)^{1-\alpha}\left(\eta\, s_m(v) - \hat{\lambda}\right)$, subtracting it from both sides gives

$$\ln \frac{\hat{p}_m(v)}{\pi_m(v)} - \pi_m(v)^{1-\alpha}\left(\eta\, s_m(v) - \hat{\lambda}\right) = \frac{-\ln(1-y) - y}{1-\alpha}. \tag{40}$$

It remains to bound the right-hand side. By Equation (39),

$$|y| \le (1-\alpha)\, C\, \pi_m(v)^{1-\alpha},$$

which tends to 0 as $\pi_m(v) \to 0$, since $1 - \alpha > 0$. Thus $|y| \le \frac{1}{2}$ for all sufficiently small $\pi_m(v)$. For such y, the Taylor series $-\ln(1-y) = \sum_{k\ge1} \frac{y^k}{k}$ converges, and

$$|-\ln(1-y) - y| = \left|\sum_{k\ge2} \frac{y^k}{k}\right| \le \frac{1}{2}\sum_{k\ge2} |y|^k = \frac{y^2}{2(1-|y|)} \le y^2.$$

Substituting this into Equation (40),

$$\left|\ln \frac{\hat{p}_m(v)}{\pi_m(v)} - \pi_m(v)^{1-\alpha}\left(\eta\, s_m(v) - \hat{\lambda}\right)\right| \le \frac{y^2}{1-\alpha} \le (1-\alpha)\, C^2\, \pi_m(v)^{2(1-\alpha)},$$

which is $\mathcal{O}\left(\pi_m(v)^{2(1-\alpha)}\right)$ as $\pi_m(v) \to 0$. This proves Equation (21). □

Proof of Proposition 4.3

Proposition. *Let $\alpha \in [0,1]$ and $\eta > 0$. If a token $v \in \mathcal{V}$ satisfies*

$$(1-\alpha)\,\eta\,\pi_m(v)^{1-\alpha}\left(s_m(v) - \bar{s}_m^{[\alpha]}\right) < 1,$$

then

$$\frac{\hat{p}_m(v\,;\alpha)}{\pi_m(v)} \le \begin{cases} \left(1 - (1-\alpha)\,\eta\,\pi_m(v)^{1-\alpha}\left(s_m(v) - \bar{s}_m^{[\alpha]}\right)\right)^{-\frac{1}{1-\alpha}} & \text{if } \alpha \in [0,1), \\ e^{\eta\left(s_m(v) - \bar{s}_m^{[1]}\right)} & \text{if } \alpha = 1. \end{cases}$$

Proof of Proposition 4.3. Fix $v \in \mathcal{V}$, write $\hat{p}_m = \hat{p}_m(\,\cdot\,;\alpha)$, and recall the weighted mean score $\bar{s}_m^{[\alpha]}$ of Equation (22). In both cases, we first show that

$$\hat{\lambda} \ge \eta\, \bar{s}_m^{[\alpha]}, \tag{41}$$

and then substitute this bound into the closed form of the update.

Case 1: $\alpha = 1$.

By Equation (18), $\ln \frac{\hat{p}_m(u)}{\pi_m(u)} = \eta\, s_m(u) - \hat{\lambda}$ for every $u \in \mathcal{V}$. Since the logarithm is concave, Jensen's inequality gives

$$\sum_{u\in\mathcal{V}} \pi_m(u) \left(\eta\, s_m(u) - \hat{\lambda}\right) = \sum_{u\in\mathcal{V}} \pi_m(u) \ln \frac{\hat{p}_m(u)}{\pi_m(u)} \le \ln \sum_{u\in\mathcal{V}} \pi_m(u) \frac{\hat{p}_m(u)}{\pi_m(u)} = \ln 1 = 0.$$

Further expanding the left-hand side above,

$$\sum_{u\in\mathcal{V}} \pi_m(u) \left(\eta\, s_m(u) - \hat{\lambda}\right) = \eta \left(\sum_{u\in\mathcal{V}} \pi_m(u)\, s_m(u)\right) - \hat{\lambda} \left(\sum_{u\in\mathcal{V}} \pi_m(u)\right) = \eta\, \bar{s}_m^{[1]} - \hat{\lambda}.$$

Thus, $\eta\, \bar{s}_m^{[1]} - \hat{\lambda} \le 0$, which proves Equation (41). Applying this to the ratio $\hat{p}_m/\pi_m$,

$$\frac{\hat{p}_m(v)}{\pi_m(v)} = e^{\eta\, s_m(v) - \hat{\lambda}} \le e^{\eta\left(s_m(v) - \bar{s}_m^{[1]}\right)}.$$

Case 2: $\alpha \in [0, 1)$.

Dividing Equation (19) by $\pi_m(u)$ and raising both sides to the power $-(1-\alpha)$, for every $u \in \mathcal{V}$,

$$\left(\frac{\hat{p}_m(u)}{\pi_m(u)}\right)^{-(1-\alpha)} = 1 + (1-\alpha)\, \pi_m(u)^{1-\alpha} \left(\hat{\lambda} - \eta\, s_m(u)\right). \tag{42}$$

For $x > 0$, define

$$\phi(x) \coloneqq \frac{x^{-(1-\alpha)} - 1}{1-\alpha}.$$

Then ϕ is convex since $\phi''(x) = (2-\alpha)\, x^{-(3-\alpha)} > 0$, and it satisfies $\phi(1) = 0$. Rearranging Equation (42) and multiplying by $\pi_m(u)$,

$$\pi_m(u)^{2-\alpha} \left(\hat{\lambda} - \eta\, s_m(u)\right) = \pi_m(u)\, \phi\left(\frac{\hat{p}_m(u)}{\pi_m(u)}\right).$$

Summing over $u \in \mathcal{V}$ and applying Jensen's inequality,

$$\begin{aligned}\hat{\lambda} \left(\sum_{u\in\mathcal{V}} \pi_m(u)^{2-\alpha}\right) - \eta \left(\sum_{u\in\mathcal{V}} \pi_m(u)^{2-\alpha}\, s_m(u)\right) &= \sum_{u\in\mathcal{V}} \pi_m(u)\, \phi\left(\frac{\hat{p}_m(u)}{\pi_m(u)}\right) \\ &\ge \phi\left(\sum_{u\in\mathcal{V}} \pi_m(u) \frac{\hat{p}_m(u)}{\pi_m(u)}\right) = \phi(1) = 0.\end{aligned}$$

Thus,

$$\hat{\lambda} \ge \eta\, \frac{\sum_{u\in\mathcal{V}} \pi_m(u)^{2-\alpha} s_m(u)}{\sum_{u\in\mathcal{V}} \pi_m(u)^{2-\alpha}} = \eta\, \bar{s}_m^{[\alpha]},$$

which proves Equation (41). Substituting it into Equation (42) at $u = v$,

$$\left(\frac{\hat{p}_m(v)}{\pi_m(v)}\right)^{-(1-\alpha)} \ge 1 - (1-\alpha)\, \eta\, \pi_m(v)^{1-\alpha} \left(s_m(v) - \bar{s}_m^{[\alpha]}\right).$$

The right-hand side is positive by assumption, and $x \mapsto x^{-\frac{1}{1-\alpha}}$ is decreasing on $(0, \infty)$, so raising both sides to the power $-\frac{1}{1-\alpha}$ gives

$$\frac{\hat{p}_m(v)}{\pi_m(v)} \le \left(1 - (1-\alpha)\, \eta\, \pi_m(v)^{1-\alpha} \left(s_m(v) - \bar{s}_m^{[\alpha]}\right)\right)^{-\frac{1}{1-\alpha}}. \qquad \square$$

A.2.4 Section 4.2

Proof of Proposition 4.4

Proposition. *Let $\alpha \in [0,1]$ and $\eta > 0$, and write $\hat{p}_m = \hat{p}_m(\,\cdot\,;\alpha,\eta)$. Then*

$$\tilde{A}_m(\hat{p}_m) - \tilde{A}_m(\pi_m) = \frac{1}{\eta}\Big(D_{\psi_\alpha}(\hat{p}_m,\pi_m) + D_{\psi_\alpha}(\pi_m,\hat{p}_m)\Big) \geq 0.$$

Further, if the score s_m is non-constant over $\mathcal{V}$, the gain is strictly positive, $\tilde{A}_m(\hat{p}_m) - \tilde{A}_m(\pi_m) > 0$.

Proof of Proposition 4.4. Write $\psi = \psi_\alpha$. By the definition (10),

$$\begin{aligned} D_\psi(\hat{p}_m,\pi_m)&+D_\psi(\pi_m,\hat{p}_m) \\ &= \psi(\hat{p}_m) - \psi(\pi_m) - \langle \nabla\psi(\pi_m), \hat{p}_m - \pi_m\rangle + \psi(\pi_m) - \psi(\hat{p}_m) - \langle \nabla\psi(\hat{p}_m), \pi_m - \hat{p}_m\rangle \\ &= \langle \nabla\psi(\hat{p}_m) - \nabla\psi(\pi_m), \hat{p}_m - \pi_m\rangle. \end{aligned}$$

By Proposition 3.3, Equation (17) holds at $(p,\lambda) = (\hat{p}_m, \hat{\lambda})$, i.e.,

$$\nabla\psi(\hat{p}_m) - \nabla\psi(\pi_m) = \eta\, s_m - \hat{\lambda}\mathbf{1}.$$

Substituting this above,

$$\begin{aligned} D_\psi(\hat{p}_m,\pi_m) + D_\psi(\pi_m,\hat{p}_m) &= \langle \eta\, s_m - \hat{\lambda}\mathbf{1}, \hat{p}_m - \pi_m\rangle \\ &= \eta\langle s_m, \hat{p}_m - \pi_m\rangle - \hat{\lambda}\langle \mathbf{1}, \hat{p}_m - \pi_m\rangle \\ &= \eta\langle s_m, \hat{p}_m - \pi_m\rangle \\ &= \eta\left(\tilde{A}_m(\hat{p}_m) - \tilde{A}_m(\pi_m)\right), \end{aligned}$$

where the third step uses $\langle \mathbf{1}, \hat{p}_m\rangle = \langle \mathbf{1}, \pi_m\rangle = 1$, and the last step uses the definition of $\tilde{A}_m$ in Equation (9). Thus,

$$\tilde{A}_m(\hat{p}_m) - \tilde{A}_m(\pi_m) = \frac{1}{\eta}\left(D_\psi(\hat{p}_m,\pi_m) + D_\psi(\pi_m,\hat{p}_m)\right).$$

By Lemma A.1(i),

$$D_\psi(\hat{p}_m,\pi_m) + D_\psi(\pi_m,\hat{p}_m) \geq 0,$$

and further

$$D_\psi(\hat{p}_m,\pi_m) + D_\psi(\pi_m,\hat{p}_m) = 0 \iff \hat{p}_m = \pi_m.$$

To complete the proof, we show that if s_m is non-constant, then $\hat{p}_m \neq \pi_m$ and thus $\tilde{A}_m(\hat{p}_m) - \tilde{A}_m(\pi_m) = \frac{1}{\eta}\left(D_\psi(\hat{p}_m,\pi_m) + D_\psi(\pi_m,\hat{p}_m)\right) > 0$.

Assume for contradiction that $\hat{p}_m = \pi_m$. Then Equation (17) at $(p,\lambda) = (\hat{p}_m, \hat{\lambda})$ reads

$$\nabla\psi(\pi_m) = \nabla\psi(\pi_m) + \eta\, s_m - \hat{\lambda}\mathbf{1},$$

which implies that

$$s_m = \frac{\hat{\lambda}}{\eta}\mathbf{1},$$

which contradicts that s_m is non-constant. Thus $\hat{p}_m \neq \pi_m$, so by Lemma A.1(i) both divergences are strictly positive, and thus, $\tilde{A}_m(\hat{p}_m) - \tilde{A}_m(\pi_m) > 0$. □

Proof of Theorem 4.5

Theorem. *Let $\alpha \in [0,1]$, and let the score s_m be non-constant over the vocabulary $\mathcal{V}$. Then the derivatives in Equation (26) exist, $\tilde{\Gamma}_m(\alpha) > 0$, and*

$$\frac{\Gamma_m(\alpha)}{\tilde{\Gamma}_m(\alpha)} = \frac{\mathrm{Cov}_{\mathsf{V}\sim\pi_m^{[\alpha]}}\left(s_m(\mathsf{V}),\, a_m(\mathsf{V})\right)}{\mathrm{Var}_{\mathsf{V}\sim\pi_m^{[\alpha]}}\left(s_m(\mathsf{V})\right)}.$$

Proof of Theorem 4.5. Write $\psi = \psi_\alpha$. For $\eta > 0$, write $p(\eta) \coloneqq \hat{p}_m(\,\cdot\,;\alpha,\eta)$ and $\lambda(\eta)$ for its multiplier $\hat{\lambda}$ in Proposition 3.3, and set $p(0) \coloneqq \pi_m$ and $\lambda(0) \coloneqq 0$. The derivatives in Equation (26) are taken from the right at $\eta = 0$. We first show that $\eta \mapsto p(\eta)$ is differentiable from the right at $\eta = 0$ and compute its derivative, and then compute $\Gamma_m(\alpha)$ and $\tilde{\Gamma}_m(\alpha)$.

Define $F : (0,\infty)^{\mathcal{V}} \times \mathbb{R} \times \mathbb{R} \to \mathbb{R}^{|\mathcal{V}|+1}$ by

$$F(p,\lambda,\eta) \coloneqq \begin{pmatrix} \nabla\psi(p) - \nabla\psi(\pi_m) - \eta\, s_m + \lambda\mathbf{1} \\ \langle \mathbf{1}, p\rangle - 1 \end{pmatrix},$$

where the first block has $|\mathcal{V}|$ entries and the second block has one entry. Thus $F(p,\lambda,\eta) = 0$ if and only if p sums to one and (p,λ) satisfies Equation (17) with step size η. By Proposition 3.3, for every $\eta > 0$,

$$F(p(\eta),\lambda(\eta),\eta) = 0.$$

This also holds at $\eta = 0$, since $F(\pi_m, 0, 0) = 0$. By Lemma A.2, ψ is infinitely differentiable on $(0,\infty)^{\mathcal{V}}$, so F is continuously differentiable.

Next, we compute the Jacobian matrix $J \in \mathbb{R}^{(|\mathcal{V}|+1)\times(|\mathcal{V}|+1)}$ of F with respect to (p,λ) at $(\pi_m, 0, 0)$. In block form, this gives that

$$J = \begin{pmatrix} H & \mathbf{1} \\ \mathbf{1}^\top & 0 \end{pmatrix}, \qquad \text{for } H \coloneqq \mathrm{diag}\left(\pi_m(v)^{\alpha-2}\right)_{v\in\mathcal{V}} \in \mathbb{R}^{\mathcal{V}\times\mathcal{V}}.$$

We show that J is invertible. Suppose that $x \in \mathbb{R}^{\mathcal{V}}$ and $\mu \in \mathbb{R}$ satisfy $J\begin{pmatrix} x \\ \mu \end{pmatrix} = 0$. This gives the pair of equations

$$Hx + \mu\mathbf{1} = 0, \qquad \mathbf{1}^\top x = 0.$$

Since H is diagonal with positive entries, it is invertible, and the first equation gives $x = -\mu H^{-1}\mathbf{1}$. Substituting this into the second equation,

$$0 = \mathbf{1}^\top x = -\mu\, \mathbf{1}^\top H^{-1}\mathbf{1} = -\mu \sum_{v\in\mathcal{V}} \pi_m(v)^{2-\alpha}.$$

Since $\sum_v \pi_m(v)^{2-\alpha} > 0$, we have $\mu = 0$, and thus $x = 0$. Thus $J\begin{pmatrix} x \\ \mu \end{pmatrix} = 0$ implies $x = 0$ and $\mu = 0$, so J is invertible.

Since F is continuously differentiable, $F(\pi_m, 0, 0) = 0$, and J is invertible, the implicit function theorem gives $\epsilon > 0$ and a continuously differentiable map $\eta \mapsto (\tilde{p}(\eta), \tilde{\lambda}(\eta))$ from $(-\epsilon,\epsilon)$ to $(0,\infty)^{\mathcal{V}} \times \mathbb{R}$ with $(\tilde{p}(0), \tilde{\lambda}(0)) = (\pi_m, 0)$ and

$$F(\tilde{p}(\eta), \tilde{\lambda}(\eta), \eta) = 0 \qquad \text{for every } \eta \in (-\epsilon,\epsilon).$$

For $\eta \in (0, \epsilon)$, the last entry of F gives $\langle \mathbf{1}, \tilde{p}(\eta) \rangle = 1$, so $\tilde{p}(\eta) \in \Delta^+(\mathcal{V})$, and the first block shows that $(\tilde{p}(\eta), \tilde{\lambda}(\eta))$ satisfies Equation (17). By the uniqueness of this pair in Proposition 3.3,

$$(\tilde{p}(\eta), \tilde{\lambda}(\eta)) = (p(\eta), \lambda(\eta)) \qquad \text{for every } \eta \in [0, \epsilon),$$

where the case $\eta = 0$ holds by construction. Thus $\eta \mapsto (p(\eta), \lambda(\eta))$ is continuously differentiable on $[0, \epsilon)$, with the derivative at $\eta = 0$ taken from the right.

We now compute the derivatives at $\eta = 0$. Differentiating $F(p(\eta), \lambda(\eta), \eta) = 0$ with respect to η at $\eta = 0$ by the chain rule,

$$J \begin{pmatrix} \frac{dp}{d\eta}\Big|_{\eta=0} \\ \frac{d\lambda}{d\eta}\Big|_{\eta=0} \end{pmatrix} + \frac{\partial F}{\partial \eta}(\pi_m, 0, 0) = 0.$$

By the definition of F,

$$\frac{\partial F}{\partial \eta}(\pi_m, 0, 0) = \begin{pmatrix} -s_m \\ 0 \end{pmatrix}.$$

Further, multiplying out the blocks of J gives the pair of equations

$$H \frac{dp}{d\eta}\Big|_{\eta=0} + \frac{d\lambda}{d\eta}\Big|_{\eta=0} \mathbf{1} = s_m, \qquad \mathbf{1}^\top \frac{dp}{d\eta}\Big|_{\eta=0} = 0.$$

Multiplying the first equation by $H^{-1} = \operatorname{diag}\left(\pi_m(v)^{2-\alpha}\right)_{v \in \mathcal{V}}$, for every $v \in \mathcal{V}$,

$$\frac{d}{d\eta} p(v)\Big|_{\eta=0} = \pi_m(v)^{2-\alpha} \left(s_m(v) - \frac{d\lambda}{d\eta}\Big|_{\eta=0} \right). \tag{43}$$

Substituting this into the second equation,

$$0 = \sum_{v \in \mathcal{V}} \pi_m(v)^{2-\alpha} \, s_m(v) - \frac{d\lambda}{d\eta}\Big|_{\eta=0} \sum_{v \in \mathcal{V}} \pi_m(v)^{2-\alpha}.$$

Rearranging, and using the definition (22) of $\pi_m^{[\alpha]}$,

$$\frac{d\lambda}{d\eta}\Big|_{\eta=0} = \frac{\sum_{v \in \mathcal{V}} \pi_m(v)^{2-\alpha} \, s_m(v)}{\sum_{u \in \mathcal{V}} \pi_m(u)^{2-\alpha}} = \mathbb{E}_{\mathsf{V} \sim \pi_m^{[\alpha]}} \left[s_m(\mathsf{V}) \right].$$

Substituting this into Equation (43),

$$\frac{d}{d\eta} p(v)\Big|_{\eta=0} = \pi_m(v)^{2-\alpha} \left(s_m(v) - \mathbb{E}_{\mathsf{V} \sim \pi_m^{[\alpha]}} \left[s_m(\mathsf{V}) \right] \right). \tag{44}$$

Since $A_m(p) = \langle a_m, p \rangle$ and $\tilde{A}_m(p) = \langle s_m, p \rangle$ are linear in p, by the chain rule both derivatives in Equation (26) exist, with

$$\Gamma_m(\alpha) = \left\langle a_m, \, \frac{d}{d\eta} p\Big|_{\eta=0} \right\rangle, \qquad \tilde{\Gamma}_m(\alpha) = \left\langle s_m, \, \frac{d}{d\eta} p\Big|_{\eta=0} \right\rangle.$$

Let $c \coloneqq \sum_{u \in \mathcal{V}} \pi_m(u)^{2-\alpha} > 0$, so that $\pi_m^{[\alpha]}(v) = \pi_m(v)^{2-\alpha}/c$ by Equation (22). Substituting Equation (44),

$$\begin{aligned}
\Gamma_m(\alpha) &= \sum_{v \in \mathcal{V}} a_m(v)\, \pi_m(v)^{2-\alpha} \left(s_m(v) - \mathbb{E}_{\mathsf{V} \sim \pi_m^{[\alpha]}} \left[s_m(\mathsf{V}) \right] \right) \\
&= c \sum_{v \in \mathcal{V}} \pi_m^{[\alpha]}(v)\, a_m(v) \left(s_m(v) - \mathbb{E}_{\mathsf{V} \sim \pi_m^{[\alpha]}} \left[s_m(\mathsf{V}) \right] \right) \\
&= c\, \mathbb{E}_{\mathsf{V} \sim \pi_m^{[\alpha]}} \left[a_m(\mathsf{V}) \left(s_m(\mathsf{V}) - \mathbb{E}_{\mathsf{V}' \sim \pi_m^{[\alpha]}} \left[s_m(\mathsf{V}') \right] \right) \right] \\
&= c\, \mathrm{Cov}_{\mathsf{V} \sim \pi_m^{[\alpha]}} \left(s_m(\mathsf{V}),\, a_m(\mathsf{V}) \right),
\end{aligned}$$

where the last step uses $\mathrm{Cov}(X, Y) = \mathbb{E}\big[Y\,(X - \mathbb{E}[X])\big]$. By the same steps with s_m in place of a_m,

$$\tilde{\Gamma}_m(\alpha) = c\, \mathbb{E}_{\mathsf{V} \sim \pi_m^{[\alpha]}} \left[s_m(\mathsf{V}) \left(s_m(\mathsf{V}) - \mathbb{E}_{\mathsf{V}' \sim \pi_m^{[\alpha]}} \left[s_m(\mathsf{V}') \right] \right) \right] = c\, \mathrm{Var}_{\mathsf{V} \sim \pi_m^{[\alpha]}} \left(s_m(\mathsf{V}) \right).$$

Since $\pi_m^{[\alpha]} \in \Delta^+(\mathcal{V})$ has full support and s_m is non-constant, the variance is positive, so $\tilde{\Gamma}_m(\alpha) > 0$. Dividing,

$$\frac{\Gamma_m(\alpha)}{\tilde{\Gamma}_m(\alpha)} = \frac{c\, \mathrm{Cov}_{\mathsf{V} \sim \pi_m^{[\alpha]}} \left(s_m(\mathsf{V}),\, a_m(\mathsf{V}) \right)}{c\, \mathrm{Var}_{\mathsf{V} \sim \pi_m^{[\alpha]}} \left(s_m(\mathsf{V}) \right)} = \frac{\mathrm{Cov}_{\mathsf{V} \sim \pi_m^{[\alpha]}} \left(s_m(\mathsf{V}),\, a_m(\mathsf{V}) \right)}{\mathrm{Var}_{\mathsf{V} \sim \pi_m^{[\alpha]}} \left(s_m(\mathsf{V}) \right)}.$$

□

Proof of Corollary 4.5.1

Corollary. *Let $\alpha \in [0, 1]$, and suppose $\mathrm{Cov}_{\mathsf{V} \sim \pi_m^{[\alpha]}} \left(s_m(\mathsf{V}),\, a_m(\mathsf{V}) \right) > 0$. Then there exists $\eta_0 > 0$ such that $A_m\left(\hat{p}_m(\,\cdot\,; \alpha, \eta)\right) > A_m(\pi_m)$ for every $\eta \in (0, \eta_0)$.*

Proof of Corollary 4.5.1. A constant score has zero covariance with a_m, so the assumption implies that s_m is non-constant, and Theorem 4.5 applies. By its proof,

$$\Gamma_m(\alpha) = c\, \mathrm{Cov}_{\mathsf{V} \sim \pi_m^{[\alpha]}} \left(s_m(\mathsf{V}),\, a_m(\mathsf{V}) \right) > 0, \qquad c = \sum_{u \in \mathcal{V}} \pi_m(u)^{2-\alpha} > 0.$$

Define $h(\eta) \coloneqq A_m\left(\hat{p}_m(\,\cdot\,; \alpha, \eta)\right)$ for $\eta \geq 0$, so that $h(0) = A_m(\pi_m)$. By the definition (26) of $\Gamma_m(\alpha)$,

$$\lim_{\eta \to 0^+} \frac{h(\eta) - h(0)}{\eta} = \Gamma_m(\alpha) > 0.$$

Thus there exists $\eta_0 > 0$ such that for every $\eta \in (0, \eta_0)$,

$$\frac{h(\eta) - h(0)}{\eta} > \frac{\Gamma_m(\alpha)}{2} > 0.$$

Multiplying by $\eta > 0$ gives $h(\eta) > h(0)$, i.e., $A_m\left(\hat{p}_m(\,\cdot\,; \alpha, \eta)\right) > A_m(\pi_m)$. □

A.3 Solving for $\hat{\lambda}$

By Proposition 3.3, the Bregman tilting update is determined by the multiplier $\hat{\lambda}$, the unique scalar for which the right-hand side of Equation (18) or (19) sums to one. We describe the properties of $\hat{\lambda}$ and the algorithm we use to compute it at each decoding step.

Properties of $\hat{\lambda}$. From the proof of Proposition 3.3 (Appendix A.2.2), we have that $\hat{p}_m(\cdot\,;\alpha) = \nu_\alpha(\cdot\,;\hat{\lambda})$, where $\nu_\alpha(\cdot\,;\lambda)$ and its total mass $Z_\alpha(\lambda)$ are defined in Equation (35), and $\hat{\lambda}$ solves $Z_\alpha(\lambda) = 1$. For $\alpha = 1$, this gives the closed form

$$\hat{\lambda} = \ln \sum_{v \in \mathcal{V}} \pi_m(v)\, e^{\eta\, s_m(v)}.$$

For $\alpha \in [0, 1)$, $\hat{\lambda}$ has no closed form. However, the same proof shows that $\lambda_{\text{lo}} \le \hat{\lambda} \le \lambda_{\text{hi}}$ for the endpoints λ_{lo} and λ_{hi} of Equation (36). This gives a computable bound for $\hat{\lambda}$ at inference time:

$$\lambda_{\text{lo}} = \left(\max_{v \in \mathcal{V}} \eta s_m(v) - \frac{\pi_m(v)^{\alpha-1}}{1-\alpha}\right) + \frac{1}{1-\alpha} \le \hat{\lambda} \le \eta \left(\max_{v \in \mathcal{V}} s_m(v)\right) = \lambda_{\text{hi}}. \tag{45}$$

The same proof also shows that Z_α is well defined and strictly decreasing on the interval $[\lambda_{\text{lo}}, \lambda_{\text{hi}}]$ (Equation 37), so $\hat{\lambda}$ is the unique solution of $Z_\alpha(\lambda) = 1$ in this interval. Further, Z_α is strictly convex on $[\lambda_{\text{lo}}, \lambda_{\text{hi}}]$, since differentiating Equation (37) again gives

$$\frac{d^2}{d\lambda^2} Z_\alpha(\lambda) = (2-\alpha) \sum_{v \in \mathcal{V}} \nu_\alpha(v\,;\lambda)^{3-2\alpha} > 0.$$

Root Finding Algorithm. For $\alpha \in [0, 1)$, we use the properties above to approximate $\hat{\lambda}$ at each decoding step in three stages: bisection, Newton's method, and normalization.

First, we run bisection on an interval $[\ell, u]$, initialized at the endpoints of Equation (45), $\ell = \lambda_{\text{lo}}$ and $u = \lambda_{\text{hi}}$. Since Z_α is strictly decreasing with $Z_\alpha(\hat{\lambda}) = 1$, every step maintains $\ell \le \hat{\lambda} \le u$ while halving the width of the interval. We use 6 bisection steps in our experiments.

Second, we run Newton's method to find the root of $Z_\alpha(\lambda) - 1$, starting from the lower end $\lambda_0 = \ell$ of the final bisection interval. By Equation (37), the Newton iterates are

$$\lambda_{k+1} := \lambda_k - \frac{Z_\alpha(\lambda_k) - 1}{Z'_\alpha(\lambda_k)} = \lambda_k + \frac{Z_\alpha(\lambda_k) - 1}{\sum_{v \in \mathcal{V}} \nu_\alpha(v\,;\lambda_k)^{2-\alpha}}.$$

Since Z_α is strictly decreasing and strictly convex on $[\lambda_{\text{lo}}, \lambda_{\text{hi}}]$ and $\lambda_0 \le \hat{\lambda}$, the Newton iterates increase monotonically toward $\hat{\lambda}$ without overshoot, i.e., $\lambda_k \le \lambda_{k+1} \le \hat{\lambda}$ for every $k \ge 0$. Further, Newton's method achieves a quadratic convergence rate when the iterates are sufficiently close to $\hat{\lambda}$ [SM03]. We use 6 Newton steps in our experiments.

Lastly, we normalize the final iterate to ensure a valid PMF. Writing λ_K for the final Newton iterate, we sample the next token from the PMF $\nu_\alpha(v\,;\lambda_K)/Z_\alpha(\lambda_K)$. This lies in $\Delta^+(\mathcal{V})$ for any number of steps of our procedure, and it is equal to $\hat{p}_m(v\,;\alpha)$ when $\lambda_K = \hat{\lambda}$ recovers the exact root.

This procedure is cheap relative to the LLM decoding step. Each bisection step evaluates Z_α once, and each Newton step evaluates Z_α and its derivative (37) once. Each evaluation is a single elementwise pass over the vectors π_m and s_m, which the update at $\alpha = 1$ computes as well. Thus, finding $\hat{\lambda}$ requires no additional forward pass of the LLM, and its cost of $\mathcal{O}(|\mathcal{V}|)$ operations per step, for a fixed number of steps, is negligible next to the two forward passes which produce π_m and s_m.

B Experimental Details

We run all experiments on NVIDIA H100 GPUs. All of the models and benchmarks we use are open source. The rest of this section gives the full details of our experiments.

B.1 Recursive Self-Aggregation

RSA [Ven+26] scales the iterative refinement procedure along an additional parallel axis. It solves $Q > 1$ chains in parallel that interact through their prompts, so we write $h^{(t,j)}$, $r^{(t,j)}$, and $f^{(t,j)} = \text{Eval}(r^{(t,j)})$ for the prompt, rollout, and feedback of chain $j \in [Q]$ at round t. At round 1, every chain is prompted with $h^{(1,j)} = h^{(1)}$. After round t, RSA draws for each chain j a set $\mathcal{S}^{(t,j)} \subseteq [Q]$ of $K > 1$ distinct chains, uniformly at random from all subsets of $[Q]$ of size K and independently across chains, so the same rollout can appear in the prompts of several chains. The prompt $h^{(t+1,j)}$ then contains the question, the K rollouts $r^{(t,i)}$ for $i \in \mathcal{S}^{(t,j)}$, and instructions to aggregate them into a single improved solution. While Self-Refine independently refines one chain at a time, RSA mixes knowledge across chains, so a correct step found by one chain can spread to the others. We modify RSA to incorporate feedback by showing each rollout $r^{(t,i)}$ followed by its feedback $f^{(t,i)} = \text{Eval}(r^{(t,i)})$, which reviews that rollout on its own. We use $K = 4$, following the prior work [Ven+26].

B.2 Benchmarks and Grading

LiveCodeBench. We use LiveCodeBench [JHGLYZWSSS25], which collects new problems from programming contests over time so that models can be tested on problems released after their training data. We use the 131 problems of its sixth release that were published from February to April 2025 following the evaluations reported in the model card of Qwen3-4B-Instruct-2507[2] as well as prior work [Hüb+26]. Of these problems, 82 are from AtCoder, where the program reads standard input and writes standard output, and 49 are from LeetCode, where the program completes a given function. By the labels of the benchmark, 31 problems are easy, 39 are medium, and 61 are hard. Each problem has a few public tests that appear as examples in the problem statement (a median of 3) and a larger set of private tests (a median of 40). We take the last Python code block of each rollout and run it on every one of the exposed and hidden tests in a sandbox, with a limit of 6 seconds of CPU time per test. A rollout is correct if its program passes every test, and a rollout without a code block is incorrect. We reimplement the grading of the official LiveCodeBench evaluation code so that programs run in a sandbox and the feedback can use the result of every public test. On the zero-shot rollouts, our grader agrees with the official code on all but one of the 2069 programs; the exception has a test that runs at the 6 s limit.

MBPP+. We use MBPP+ [LXWZ23], the 378 tasks of the EvalPlus release (v0.2.0) of MBPP [Aus+21]. Each task is a one-sentence description of a Python function to write, with its name, and 3 to 7 hand-written test assertions (a median of 3), which the prompt shows as examples; EvalPlus

[2] https://huggingface.co/Qwen/Qwen3-4B-Instruct-2507

adds about 105 generated tests per task that probe edge cases, which the prompt never shows. We grade as the EvalPlus leaderboard does: the program must define the named function, and a rollout is correct if its function returns the reference output on every hand-written and generated test, compared with EvalPlus's rules (exact equality, with a tolerance for floating-point outputs), each test with EvalPlus's time allowance of max(1 s, 4 times the reference solution's time) of CPU time. Our grader agrees with the EvalPlus checker on 6031 of the 6048 zero-shot programs; the other 17 are cases where the checker's 4 GB memory limit fails its own reference solution. The solver already solves most of MBPP+ zero-shot: its mean pass@1 is 75.5% (averaged over the four seeds), and at seed 0 all 16 of its zero-shot rollouts are correct on 223 of the 378 tasks (59%), so on these tasks the mean pass@1 of every method is pinned near 100% and only losses are possible. We therefore report MBPP+ on its *hard set*: the 155 tasks that were not solved by all 16 rollouts of the seed-0 zero-shot run (its mean pass@1 on them is 39.6%). The set is fixed before the runs and is the same for every method and seed; since it is chosen from the seed-0 zero-shot rollouts, its round-1 score at seed 0 is selected low, while seeds 1 to 3 draw their own round 1 on the same tasks. The results on the full set of 378 tasks are in Appendix B.8.

AIME 2025. We use the 30 problems of AIME 2025 (I and II), as distributed by MathArena [BDPJV25]. Each answer is an integer between 0 and 999. We grade the last `\boxed{}` expression of each rollout against the reference answer with Math-Verify,[3] and a rollout without one is incorrect.

HMMT February 2025. We use the 30 problems of HMMT February 2025, also as distributed by MathArena. These are the problems of the three individual rounds of the February 2025 Harvard-MIT Mathematics Tournament, 10 each in algebra and number theory, combinatorics, and geometry. Only 14 of the 30 answers are integers; the others are fractions, radicals, other closed-form expressions, and for one problem a list of two numbers. We grade the last `\boxed{}` expression of each rollout against the reference answer with the answer parser of MathArena, which is how its leaderboard grades this competition, and a rollout without one is incorrect.

B.3 Solver and Feedback

Solver. We use Qwen3-4B-Instruct-2507 [Yan+25] as the solver on every benchmark, sampled at temperature 1 and at most 8192 new tokens per rollout on AIME, LiveCodeBench, and MBPP+ and 12288 on HMMT. On HMMT, we observed that an 8192 token cap led to a high truncation rate, so we increase the limit to 12288 in that setting.

Execution feedback. On LiveCodeBench and MBPP+, the feedback comes from running the program rather than from an LLM. Eval($r^{(t)}$) runs the program of $r^{(t)}$ on the exposed tests and reports the first test it fails. For a wrong answer, it shows the test input, the output of the program, and the expected output. For a runtime error, it shows the error and the line of the program that raised it, and for a time limit, it shows the test input. If the program passes every exposed test, the feedback says so and reminds the model that the hidden tests may still cover edge cases or larger inputs. The feedback never shows the hidden tests. The exact messages are given in Appendix B.5.1.

[3] https://github.com/huggingface/Math-Verify

The benchmark only labels about 3 tests per problem as public, which can miss many errors. Following prior work [Hüb+26], we construct the exposed test set as the benchmark's public tests along with a random half of the benchmark's private tests, drawn once and fixed for every run. Thus, we treat half of the benchmark's private tests as hidden, with the other half exposed. We still grade each program on all of its tests. Since the added tests are part of the grade, this setting measures how well each method uses informative feedback, not how well its programs generalize to unseen tests. On MBPP+ we expose a random half of the generated tests in the same way (a median of 55 exposed and 53 hidden tests per task).

Feedback model. On AIME and HMMT, the feedback comes from an LLM that reviews each rollout. We use Qwen3.5-27B to implement the feedback mechanism Eval, with thinking disabled and the sampling parameters recommended for it (temperature 1, top-p 0.95, top-k 20, presence penalty 1.5). The feedback model sees the question and a single rollout, but never the reference answer, and replies in two parts. The first part is a private verification of at most 8192 tokens, in which it may redo the rollout's computations itself. The second part, of at most 1024 tokens, is feedback to the author: it points to the first step whose conclusion is false or unjustified, explains why, and says what to reconsider, without stating an answer of its own. The solver sees only the second part as $f^{(t)}$. When we evaluate the feedback model zero-shot, we give it at most 32768 new tokens per rollout rather than 8192 as it is biased toward longer answers. HMMT uses the same feedback model, prompts, and token limits for the feedback as AIME. See Appendix B.5.2 for the exact prompts we use.

B.4 Protocol, Metrics, and Statistics

We evaluate the solver zero-shot on every benchmark, using the prompt $h^{(1)}$ with $Q = 16$ rollouts per question, and the feedback model zero-shot on AIME and HMMT with $Q = 8$ rollouts per question. We then evaluate the prompting-only versions of Self-Refine and RSA with feedback ($\eta = 0$), and each of them with Bregman tilting at $\alpha = 0$ (Equation (19)) with $\eta = 48$ and at $\alpha = 1$ with $\eta = 1$. We chose η by calibrating its value on a single seed of AIME 2025, balancing divergence to the input-prompting LLM, truncation rates, and performance. We then keep it unchanged for every benchmark and method. At $\alpha = 0$, we find $\hat{\lambda}$ at each position with 6 bisection steps followed by 6 Newton steps (Appendix A.3). On LiveCodeBench and MBPP+, the prompts are the code generation prompts of RSA [Ven+26], and Self-Refine uses their aggregation prompt with a single candidate (Appendix B.5.1). The prompts with feedback show each candidate followed by the result of running its program on the exposed tests, and the prompt without feedback $h_{\setminus f}^{(t+1)}$ is the same prompt without these results. We run every method with four seeds on every benchmark. Within a seed, all methods share the same round-1 rollouts and the same feedback on them, so round 1 is the zero-shot solver for every method and the comparisons between methods are paired by question; each seed draws its own round 1. We report means over the four seeds. All standard errors treat the questions of a benchmark as fixed and measure the variation across independent runs on them. For mean pass@1 under Self-Refine, whose chains are independent, we compute them over the 64 chains of the four seeds, and for every other quantity over the four seeds.

On every benchmark, we perform each method for $T = 5$ rounds with $Q = 16$ chains per question, and then compute metrics on each round's Q rollouts, averaged over the questions of the benchmark:

- **mean pass@1**: the fraction of the Q rollouts that are correct;
- **pass@Q**: whether at least one of the Q rollouts is correct.

Both metrics are computed directly from the Q rollouts of a round. Under Self-Refine, the Q chains of a question are independent, so pass@Q is the usual pass@k with $k = Q$. Under RSA, the chains share rollouts through their prompts after round 1, so pass@Q instead measures how many questions the population that RSA produces covers.

B.5 Prompts

See below for the various prompts used in our experiments. LiveCodeBench and MBPP+ use the prompts of Appendix B.5.1, and AIME and HMMT use the same prompts (Appendix B.5.2). None of the prompts uses a system prompt, and the question, rollouts, and feedback are inserted as plain text.

B.5.1 Code

The LiveCodeBench and MBPP+ prompts are the code generation prompts used in prior work [Ven+26]. Each ends with a format instruction, written *fmt* below. On LiveCodeBench it is the format instruction of the LiveCodeBench evaluation code; for a problem whose program reads standard input, it is

> Read the inputs from stdin solve the problem and write the answer to stdout (do not directly test on the sample inputs). Enclose your code within delimiters as follows. Ensure that when the python program runs, it reads the inputs, runs the algorithm and writes output to STDOUT.
>
> ```python
> # YOUR CODE HERE
> ```

and for a problem whose program completes a given function, it is

> You will use the following starter code to write the solution to the problem and enclose your code within delimiters.
> ```python
> *starter code*
> ```

On MBPP+, *fmt* shows the task's hand-written test assertions and names the function to define:

Your code should pass these tests:

```python
the task's test assertions, one per line
```

Enclose your code within delimiters as follows. Ensure that the code block defines the function `function name` the tests call, with every import and helper it needs.

```python
# YOUR CODE HERE
```

The question of an MBPP+ task is its one-sentence description.

Zero-shot. Round 1 of every method, and the zero-shot baseline, use the prompt $h^{(1)}$:

User: You will be given a question (problem specification) and will generate a correct Python program that matches the specification and passes all tests.

Question:
question

fmt

Self-Refine and RSA. The prompt without feedback, $h_{\setminus f}^{(t+1)}$, is the aggregation prompt of RSA [Ven+26]. RSA shows $K = 4$ rollouts drawn from the previous round. Self-Refine uses the same prompt with $K = 1$, where the one rollout is the chain's own rollout $r^{(t)}$.

User: You are given a python code implementation problem and several candidate code blocks with their reasoning. Some candidates may be incorrect or contain errors. Aggregate the useful ideas and produce a single, high-quality solution. Reason carefully; if candidates disagree, choose the correct path.
You will be given a question (problem specification) and will generate a correct Python program that matches the specification and passes all tests.

Question:
question

Candidate solutions (may contain mistakes):

---- Solution 1 ----
rollout 1

⋮

---- Solution K ----
rollout K

Now provide an improved and correct solution along with its reasoning.

fmt

The prompt with feedback, $h^{(t+1)}$, shows each of the same rollouts followed by its feedback. The

text that appears only in this prompt is in bold:

> **User:** You are given a python code implementation problem and several candidate code blocks with their reasoning**, each followed by the result of running its program on the problem's example tests**. Some candidates may be incorrect or contain errors. **Use the test results to tell which ideas are sound, aggregate the useful ones** and produce a single, high-quality solution. Reason carefully; if candidates disagree, choose the correct path.
> You will be given a question (problem specification) and will generate a correct Python program that matches the specification and passes all tests.
>
> Question:
> *question*
>
> Candidate solutions (may contain mistakes):
>
> ---- Solution 1 ----
> *rollout 1*
>
> **---- Result of running solution 1 on the example tests ----**
> ***feedback on rollout 1***
>
> ⋮
>
> ---- Solution K ----
> *rollout K*
>
> **---- Result of running solution K on the example tests ----**
> ***feedback on rollout K***
>
> Now provide an improved and correct solution along with its reasoning.
>
> *fmt*

Execution feedback. The feedback $f^{(t)} = \text{Eval}(r^{(t)})$ is one of the following messages, where the example tests are the public tests of the problem together with the added half of its private tests (Appendix B.3), and n is their number. If the program passes every example test, the feedback is

> All n example tests passed. The hidden tests remain, and the examples are few and small: a program that passes them can still be wrong on edge cases or too slow for the largest inputs allowed by the constraints. Re-check the constraints, boundary cases and time complexity before keeping this solution.

Otherwise, it reports the first example test i that the program fails, with the number m of example tests it passes. For a wrong answer, the feedback is

> m of n example tests passed. Example test i: Wrong Answer
>
> Input
> *test input*
>
> Output
> *output of the program*
>
> Expected
> *expected output*

For a runtime error, the first line is the same with "Runtime Error" in place of "Wrong Answer", followed by the error message, "Line ℓ of your program" for the line ℓ that raised it, and the test input. For a time limit, it is the same with "Time Limit Exceeded (6 s)", followed by the test input. If the program does not compile, the feedback is "Compilation Error (line ℓ)" followed by the error message. If the rollout has no code block, the feedback is

> Incorrect Format: no code block was found in your previous response, so no program could be tested. Your complete program must appear in a single ```` ```python ... ``` ```` code block, in the format the question specifies (reading standard input and writing standard output, or completing the given starter code).

and if the rollout was truncated before a complete code block, the feedback is

> Truncated Attempt: your previous response ran past the length limit before a complete code block appeared, so no program could be tested. Keep the reasoning brief and finish with the complete program. Your complete program must appear in a single ```` ```python ... ``` ```` code block, in the format the question specifies (reading standard input and writing standard output, or completing the given starter code).

B.5.2 Math

Zero-shot. Round 1 of every method, and the zero-shot baselines, use the prompt $h^{(1)}$:

> **User:** *question*
>
> Put your final answer within `\boxed{}`.

Self-Refine. The Self-Refine prompt without feedback, $h_f^{(t+1)}$, is the self-refinement prompt of RSA [Ven+26]:

> **User:** You are given a math problem and a candidate solution. The candidate may be incomplete or contain errors. Refine this trajectory and produce an improved, higher-quality solution. If it is entirely wrong, attempt a new strategy. End with the final result in `\boxed{}`.
>
> Problem:
> *question*
>
> Candidate solution (may contain mistakes):
> ---- Candidate ----
> $r^{(t)}$
>
> Now refine the candidate into an improved solution. Provide clear reasoning and end with the final answer in `\boxed{}`.

The Self-Refine prompt with feedback, $h^{(t+1)}$, shows the same rollout followed by its feedback. The text that appears only in this prompt is in bold:

User: You are given a math problem, a candidate solution**, and a reviewer's feedback on that candidate**. The candidate may be incomplete or contain errors**, and the reviewer may or may not have found a problem with it**. **Use the feedback to** refine this trajectory and produce an improved, higher-quality solution. If **the candidate** is entirely wrong, attempt a new strategy. End with the final result in \boxed{}.

Problem:
question

Candidate solution (may contain mistakes):
---- Candidate ----
$r^{(t)}$

---- Reviewer's feedback on the candidate ----
$\boldsymbol{f^{(t)}}$

Now refine the candidate into an improved solution **that addresses the feedback**. Provide clear reasoning and end with the final answer in \boxed{}.

RSA. Without feedback, we use the aggregation prompt of RSA [Ven+26]. At each round after the first, it shows $K = 4$ rollouts drawn from the previous round:

User: You are given a math problem and several candidate solutions. Some candidates may be incorrect or contain errors. Aggregate the useful ideas and produce a single, high-quality solution. Reason carefully; if candidates disagree, choose the correct path. If all are incorrect, then attempt a different strategy. End with the final result in \boxed{}.

Problem:
question

Candidate solutions (may contain mistakes):
---- Solution 1 ----
rollout 1

⋮

---- Solution K ----
rollout K

Now write a single improved solution. Provide clear reasoning and end with the final answer in \boxed{}.

The RSA prompt with feedback shows each of the same rollouts followed by the feedback model's reply on it. The text that appears only in this prompt is in bold:

User: You are given a math problem, several candidate solutions**, and a reviewer's feedback on each candidate**. Some candidates may be incorrect or contain errors**, and the reviewer may or may not have found a problem with each one**. **Use the feedback to judge the candidates,** aggregate the useful ideas**,** and produce a single, high-quality solution. Reason carefully; if candidates disagree, choose the correct path. If all are incorrect, then attempt a different strategy. End with the final result in \boxed{}.

Problem:
question

Candidate solutions (may contain mistakes)**, each followed by the reviewer's feedback on it**:
---- Solution 1 ----
rollout 1

---- Reviewer's feedback on Solution 1 ----
feedback on rollout 1

⋮

---- Solution K ----
rollout K

---- Reviewer's feedback on Solution K ----
feedback on rollout K

Now write a single improved solution **that addresses the feedback**. Provide clear reasoning and end with the final answer in \boxed{}.

Feedback model. The feedback model $\text{Eval}(r^{(t)})$ reviews one rollout at a time with the following prompt:

User: You are reviewing a candidate solution to a math problem. Your job is to decide whether its final answer is correct and, if not, to tell the author where it went wrong. You are not told whether the candidate is correct. The candidate may also have been cut off before finishing.

Work in two parts.

PART 1: VERIFICATION. This part is for you alone; the author will not see it. Check the candidate as carefully as you need to, including redoing computations and derivations yourself.
- Restate what the problem asks and check that the candidate solves that problem, not a misreading of it.
- Go through the candidate's steps in order. For each important step, verify that its conclusion is true and follows from what came before. Recompute the key quantities rather than taking them on trust.
- Check for missed cases, unjustified assumptions and sign or counting errors.
- Check the final answer against the problem's constraints and the form it requires, by substitution or an independent route where possible.

A step that is terse but true is not an error. An error is a conclusion that is false, or a claim the solution depends on that does not follow.
Stop rule: as soon as you have confirmed one error that the final answer depends on, and that the candidate does not correct later in the solution, stop verifying and go to Part 2. Everything after a broken step is moot, and you do not need to work out the right answer yourself.

PART 2: FEEDBACK TO THE AUTHOR. The author will see only this part. Do not state, hint at, or box a final answer of your own, and do not write out a full alternative solution; the author must redo the work. Point to the first step whose conclusion is false or unjustified (by the claim it makes, not by a step number), explain why it fails, and say what the author should check or reconsider there. If you found no error, say so and name the step most worth double-checking.

Finish with exactly one line in this form:
Verdict: correct | wrong | missing
"correct" means you verified the boxed answer; "wrong" means you believe it is not the right answer; "missing" means the candidate never reaches a final answer. Give "correct" only if Part 1 actually reproduced the key steps.

Problem:
question

Candidate solution:
$r^{(t)}$

Begin with "PART 1: VERIFICATION".

We generate the reply in two stages. The first stage generates Part 1 until the model begins Part 2 or reaches 8192 tokens. We then append the header "PART 2: FEEDBACK TO THE AUTHOR" and generate Part 2 for at most 1024 tokens. The feedback $f^{(t)}$ that the solver sees is Part 2 with its verdict line removed, so the solver reads the critique but not the feedback model's final judgment.

B.6 Results by Round

Tables 4, 5, 6, and 7 give mean pass@1 and pass@16 at every round for the runs of Table 1. On every benchmark and for both input-prompting methods, Bregman tilting at either value of α improves over input prompting at every round in both metrics.

Table 4: Mean pass@1 and pass@16 (%) by round on LiveCodeBench, averaged over four seeds, with $Q = 16$ rollouts per question, along with the standard error. Zero-shot is the round-1 solver which all methods share within a seed. The shaded columns correspond to our Bregman tilting. The best score in each row and metric is highlighted in **green**.

		Pass@1			Pass@16		
			Bregman tilting			Bregman tilting	
Method	Round	Prompting	$\alpha = 1$	$\alpha = 0$	Prompting	$\alpha = 1$	$\alpha = 0$
Zero-shot (round 1)		34.1 ±0.2			47.5 ±0.7		
Self-Refine	2	36.2 ±0.2	37.2 ±0.2	**37.9** ±0.2	48.9 ±0.7	51.3 ±0.7	**51.7** ±0.6
	3	37.3 ±0.2	38.1 ±0.2	**39.0** ±0.2	50.8 ±0.8	52.9 ±0.7	**54.2** ±0.7
	4	38.0 ±0.2	39.4 ±0.2	**40.3** ±0.2	51.5 ±0.7	53.6 ±0.8	**56.5** ±0.7
	5	38.6 ±0.2	39.9 ±0.2	**41.0** ±0.3	52.1 ±0.7	55.5 ±0.8	**59.0** ±0.7
RSA	2	38.3 ±0.3	40.0 ±0.3	**40.4** ±0.3	48.3 ±0.8	49.2 ±0.8	**50.4** ±0.8
	3	40.5 ±0.4	42.9 ±0.4	**43.3** ±0.4	48.1 ±0.7	49.4 ±0.8	**50.4** ±0.8
	4	42.4 ±0.4	44.8 ±0.5	**45.2** ±0.5	48.1 ±0.7	49.2 ±0.8	**51.0** ±0.8
	5	43.6 ±0.5	45.9 ±0.6	**47.0** ±0.6	47.9 ±0.8	49.6 ±0.8	**51.5** ±0.8

Table 5: Mean pass@1 and pass@16 (%) by round on the MBPP+ hard set, averaged over four seeds, with $Q = 16$ rollouts per question, along with the standard error. Zero-shot is the round-1 solver which all methods share within a seed. The shaded columns correspond to our Bregman tilting. The best score in each row and metric is highlighted in **green**.

		Pass@1			Pass@16		
			Bregman tilting			Bregman tilting	
Method	Round	Prompting	$\alpha = 1$	$\alpha = 0$	Prompting	$\alpha = 1$	$\alpha = 0$
Zero-shot (round 1)		40.8 ±0.3			62.3 ±0.7		
Self-Refine	2	55.4 ±0.3	59.8 ±0.3	**62.0** ±0.3	76.1 ±0.7	79.4 ±0.5	**80.0** ±0.5
	3	59.4 ±0.3	63.4 ±0.3	**65.5** ±0.3	79.5 ±0.6	80.8 ±0.4	**83.2** ±0.5
	4	62.4 ±0.3	66.0 ±0.3	**67.5** ±0.3	81.6 ±0.5	82.6 ±0.5	**84.5** ±0.5
	5	64.3 ±0.3	67.1 ±0.3	**67.8** ±0.3	83.4 ±0.5	83.5 ±0.5	**85.0** ±0.5
RSA	2	59.4 ±0.4	63.1 ±0.3	**65.0** ±0.3	74.8 ±0.7	77.4 ±0.6	**77.9** ±0.5
	3	65.9 ±0.5	71.0 ±0.4	**73.0** ±0.3	77.4 ±0.7	80.2 ±0.5	**80.8** ±0.5
	4	70.2 ±0.5	75.4 ±0.5	**76.5** ±0.4	77.9 ±0.6	80.8 ±0.6	**81.9** ±0.4
	5	72.9 ±0.6	77.5 ±0.5	**78.9** ±0.4	77.7 ±0.5	81.1 ±0.6	**82.6** ±0.4

Table 6: Mean pass@1 and pass@16 (%) by round on AIME 2025, averaged over four seeds, with $Q = 16$ rollouts per question, along with the standard error. Zero-shot is the round-1 solver which all methods share within a seed. The shaded columns correspond to our Bregman tilting. The best score in each row and metric is highlighted in **green**. For reference, the zero-shot mean pass@1 of the Qwen3.5-27B feedback model is 75.8%.

		Pass@1			Pass@16		
			Bregman tilting			Bregman tilting	
Method	Round	Prompting	$\alpha = 1$	$\alpha = 0$	Prompting	$\alpha = 1$	$\alpha = 0$
Zero-shot (round 1)		41.4 ±0.6			67.5 ±2.0		
Self-Refine	2	56.2 ±0.7	56.7 ±0.7	**57.3** ±0.7	81.7 ±1.9	83.3 ±1.9	**84.2** ±2.0
	3	62.6 ±0.7	65.1 ±0.7	**65.7** ±0.7	84.2 ±1.7	87.5 ±1.4	**88.3** ±1.5
	4	67.2 ±0.6	69.9 ±0.7	**70.6** ±0.7	86.7 ±1.5	90.0 ±1.2	**90.8** ±1.3
	5	70.3 ±0.6	71.8 ±0.6	**73.2** ±0.6	87.5 ±1.6	90.0 ±1.4	**92.5** ±0.8
RSA	2	59.4 ±0.9	62.7 ±1.0	**64.0** ±0.9	75.8 ±1.9	**83.3** ±1.7	82.5 ±1.4
	3	71.5 ±1.3	**76.7** ±0.9	76.5 ±1.2	83.3 ±1.8	85.8 ±1.3	**86.7** ±1.5
	4	76.7 ±1.3	81.5 ±1.0	**82.8** ±1.2	84.2 ±1.4	88.3 ±1.5	**90.0** ±1.4
	5	80.7 ±0.9	84.5 ±1.0	**85.9** ±1.2	85.8 ±1.6	89.2 ±1.4	**90.0** ±1.5

Table 7: Mean pass@1 and pass@16 (%) by round on HMMT February 2025, averaged over four seeds, with $Q = 16$ rollouts per question, along with the standard error. Zero-shot is the round-1 solver which all methods share within a seed. The shaded columns correspond to our Bregman tilting. The best score in each row and metric is highlighted in **green**. For reference, the zero-shot mean pass@1 of the Qwen3.5-27B feedback model is 68.3%.

		Pass@1			Pass@16		
			Bregman tilting			Bregman tilting	
Method	Round	Prompting	$\alpha = 1$	$\alpha = 0$	Prompting	$\alpha = 1$	$\alpha = 0$
Zero-shot (round 1)		29.8 ±0.5			52.5 ±2.3		
Self-Refine	2	38.3 ±0.6	40.4 ±0.7	**40.9** ±0.7	65.0 ±1.8	66.7 ±1.7	**70.0** ±1.7
	3	45.2 ±0.7	**48.7** ±0.7	**48.7** ±0.7	70.8 ±2.0	75.8 ±1.9	**80.0** ±1.9
	4	52.1 ±0.7	**54.5** ±0.7	54.3 ±0.7	77.5 ±1.7	80.8 ±1.3	**82.5** ±1.7
	5	55.3 ±0.7	57.8 ±0.7	**58.2** ±0.7	78.3 ±1.5	**83.3** ±1.5	**83.3** ±1.5
RSA	2	42.2 ±0.9	45.6 ±1.0	**47.0** ±0.9	62.5 ±1.7	64.2 ±1.7	**65.0** ±1.7
	3	55.0 ±1.2	59.8 ±1.1	**60.9** ±1.1	67.5 ±2.0	73.3 ±1.8	**76.7** ±1.5
	4	63.0 ±1.3	69.0 ±1.2	**69.8** ±1.0	74.2 ±1.7	**79.2** ±1.7	78.3 ±1.2
	5	68.7 ±1.6	74.5 ±1.4	**74.6** ±1.0	75.0 ±1.9	80.8 ±1.6	**83.3** ±1.5

Table 8 gives the majority vote accuracy at every round on AIME 2025 and HMMT February 2025, where a question counts as correct if the most common final answer among its $Q = 16$ rollouts is correct. As with the other two metrics, Bregman tilting at either value of α improves over input prompting at every round for both input-prompting methods.

Table 8: Majority vote accuracy (%) by round on AIME 2025 and HMMT February 2025, averaged over four seeds, with $Q = 16$ rollouts per question, along with the standard error. Ties go to the answer that appears first. Zero-shot is the round-1 solver which all methods share within a seed. The shaded columns correspond to our Bregman tilting. The best score in each row and benchmark is highlighted in **green**.

		AIME 2025			HMMT Feb 2025		
			Bregman tilting			Bregman tilting	
Method	Round	Prompting	$\alpha = 1$	$\alpha = 0$	Prompting	$\alpha = 1$	$\alpha = 0$
Zero-shot (round 1)		57.5 ±2.5			34.2 ±1.4		
Self-Refine	2	70.0 ±1.8	70.8 ±2.0	**72.5** ±1.7	43.3 ±2.5	52.5 ±2.3	**53.3** ±2.2
	3	77.5 ±1.6	**80.8** ±1.7	**80.8** ±0.8	54.2 ±1.3	62.5 ±1.7	**64.2** ±1.9
	4	80.0 ±1.4	84.2 ±1.4	**85.0** ±1.5	65.8 ±1.4	66.7 ±1.4	**68.3** ±1.9
	5	81.7 ±1.2	84.2 ±0.8	**85.8** ±1.3	67.5 ±1.4	**70.8** ±1.7	**70.8** ±1.4
RSA	2	64.2 ±2.2	69.2 ±2.3	**70.8** ±2.0	47.5 ±2.2	52.5 ±2.5	**55.8** ±2.2
	3	74.2 ±1.9	**83.3** ±1.7	82.5 ±1.7	58.3 ±2.0	65.8 ±1.7	**66.7** ±1.5
	4	78.3 ±1.9	85.0 ±1.5	**87.5** ±1.6	66.7 ±2.0	72.5 ±1.7	**73.3** ±1.8
	5	82.5 ±0.8	**87.5** ±1.3	**87.5** ±1.6	70.0 ±2.3	76.7 ±1.9	**77.5** ±1.3

B.7 Reference Prompt by Round

Figure 2 gives mean pass@1 and pass@16 at every round for the runs of Table 2.

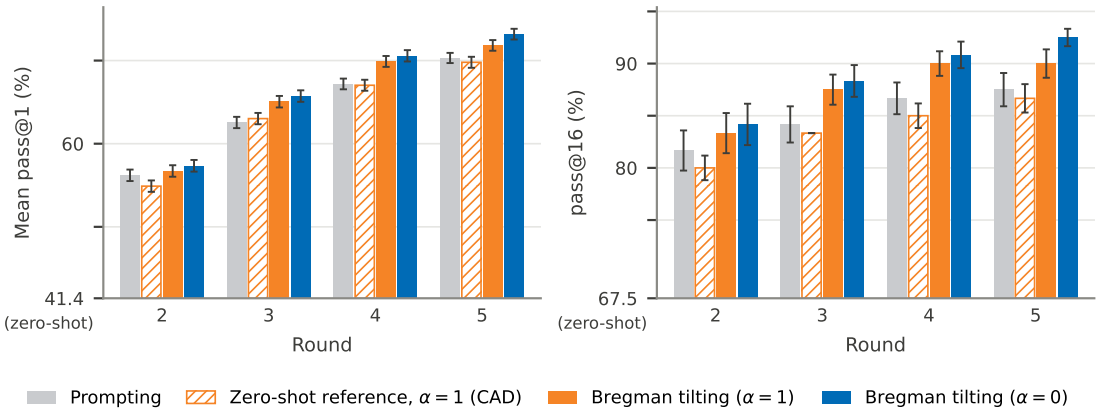

Figure 2: Mean pass@1 (left) and pass@16 (right) by round on AIME 2025 with Self-Refine and feedback from Qwen3.5-27B, for input prompting, Bregman tilting at $\alpha = 1$ with the zero-shot reference prompt $h^{(1)}$ (CAD), and Bregman tilting with our no-feedback reference prompt $h_{\setminus f}^{(t+1)}$ at $\alpha = 1$ and $\alpha = 0$. Bars are means over four seeds, with standard error bars. Within each seed, all methods use the same round-1 rollouts and feedback, so each y-axis starts at the round-1 score, marked "(zero-shot)". At round 3, the zero-shot reference reaches the same questions at every seed, so its pass@16 has zero standard error.

B.8 MBPP+ Full Set

Table 9 gives mean pass@1 and pass@16 at every round on the full set of 378 MBPP+ tasks, for the runs of Table 5. The full set adds the 223 tasks which the solver already solves in all 16 of its seed-0 zero-shot rollouts (Appendix B.2).

Table 9: Mean pass@1 and pass@16 (%) by round on the full set of 378 MBPP+ tasks, averaged over four seeds, with $Q = 16$ rollouts per question, along with the standard error. Zero-shot is the round-1 solver which all methods share within a seed. The shaded columns correspond to our Bregman tilting. The best score in each row and metric is highlighted in **green**.

		Pass@1			Pass@16		
			Bregman tilting			Bregman tilting	
Method	Round	Prompting	$\alpha = 1$	$\alpha = 0$	Prompting	$\alpha = 1$	$\alpha = 0$
Zero-shot (round 1)		75.5 ±0.1			84.5 ±0.3		
Self-Refine	2	81.3 ±0.1	83.1 ±0.1	**84.0** ±0.1	90.2 ±0.3	91.5 ±0.2	**91.8** ±0.2
	3	82.9 ±0.1	84.1 ±0.1	**84.8** ±0.1	91.6 ±0.2	92.1 ±0.2	**93.1** ±0.2
	4	84.2 ±0.1	85.1 ±0.1	**85.2** ±0.1	92.5 ±0.2	92.9 ±0.2	**93.6** ±0.2
	5	85.0 ±0.1	**85.6** ±0.1	85.3 ±0.1	93.2 ±0.2	93.3 ±0.2	**93.8** ±0.2
RSA	2	83.0 ±0.2	84.6 ±0.1	**85.4** ±0.1	89.7 ±0.3	90.7 ±0.3	**90.9** ±0.2
	3	85.7 ±0.2	87.9 ±0.2	**88.8** ±0.1	90.7 ±0.3	91.9 ±0.2	**92.1** ±0.2
	4	87.5 ±0.2	89.8 ±0.2	**90.3** ±0.2	90.9 ±0.3	92.1 ±0.2	**92.6** ±0.2
	5	88.7 ±0.2	90.6 ±0.2	**91.2** ±0.2	90.9 ±0.2	92.3 ±0.2	**92.9** ±0.2

B.9 Runtime Experiment Details

We give the details of the runtime comparison of Table 3. The rollouts are the round-2 rollouts of the seed-0 input-prompting Self-Refine run on AIME 2025. For each of them, the prompt $h^{(2)}$ contains the question, the round-1 rollout, and its feedback, and the reference prompt $h^{(2)}_{\setminus f}$ is the same prompt without the feedback. With one chain, we decode the first chain of each of the 30 questions, one at a time (97,138 tokens in total). With 8 chains, we decode the first 8 chains of 10 questions drawn at random, with the 8 chains of a question together in one batch (234,424 tokens in total), where a chain leaves the batch when its recorded rollout ends. All methods use the same decoding implementation with key-value caching, each on an NVIDIA H100 NVL GPU. Input prompting decodes $h^{(2)}$ alone, while Bregman tilting decodes $h^{(2)}$ and $h^{(2)}_{\setminus f}$ together, so each chain takes two sequences of the batch. CAD does the same with the zero-shot prompt $h^{(1)}$, which contains the question alone, in place of $h^{(2)}_{\setminus f}$. As in our experiments, we use $\eta = 1$ at $\alpha = 1$, and $\eta = 48$ with 6 bisection steps followed by 6 Newton steps for $\hat{\lambda}$ at $\alpha = 0$ (Appendix B.4). At every position, each method computes its next-token PMF and samples a token from it, and we then replace the sampled token with the recorded one, so the timed computation is that of Algorithm 1 while the prefixes are identical across methods.

For each method and number of chains, we make one untimed pass over the rollouts followed by two timed passes, and report the mean of the two, which differ by less than 0.3%. The time per token is

the wall-clock decoding time divided by the number of generated tokens, summed over the chains of a batch, and it excludes the initial processing of the prompts, which Table 3 reports separately. The fixed cost of a decoding step, which is shared by all sequences in the batch, includes loading the model weights from GPU memory and launching the GPU kernels, while each sequence adds its own attention over its key-value cache and its own next-token PMF. The prompt processing time doubles under Bregman tilting, since the reference prompt is about as long as the prompt, but barely changes under CAD, since its reference prompt is more than 20 times shorter. The key-value cache, which is allocated in advance for a context of 20,480 tokens per sequence, doubles under both CAD and Bregman tilting, while the 7.5 GiB of model weights are stored once.

References

[Agr+26] L. A. Agrawal et al. "GEPA: Reflective Prompt Evolution Can Outperform Reinforcement Learning". In: *The Fourteenth International Conference on Learning Representations*. 2026.

[Aus+21] J. Austin, A. Odena, M. Nye, M. Bosma, H. Michalewski, D. Dohan, E. Jiang, C. Cai, M. Terry, Q. Le, and C. Sutton. *Program Synthesis with Large Language Models*. 2021. arXiv: 2108.07732 [cs.PL].

[BDPJV25] M. Balunović, J. Dekoninck, I. Petrov, N. Jovanović, and M. Vechev. *MathArena: Evaluating LLMs on Uncontaminated Math Competitions*. 2025. arXiv: 2505.23281 [cs.AI].

[BT03] A. Beck and M. Teboulle. "Mirror descent and nonlinear projected subgradient methods for convex optimization". *Operations Research Letters* 31.3 (2003), pp. 167–175.

[BMMS26] J. L. Block, M. Mohri, A. Mokhtari, and S. Shakkottai. "Temper-Then-Tilt: Principled Unlearning for Generative Models through Tempering and Classifier Guidance". In: *Proceedings of the 43rd International Conference on Machine Learning*. Vol. 306. Proceedings of Machine Learning Research. PMLR, July 2026, pp. 8558–8593.

[BV04] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[BHPSRK26] T. K. Buening, J. Hübotter, B. Pásztor, I. Shenfeld, G. Ramponi, and A. Krause. "Aligning Language Models from User Interactions". In: *Third Conference on Language Modeling*. 2026.

[CLSZ24] X. Chen, M. Lin, N. Schärli, and D. Zhou. "Teaching large language models to self-debug". In: *The Twelfth International Conference on Learning Representations*. 2024.

[DN21] P. Dhariwal and A. Nichol. "Diffusion models beat GANs on image synthesis". *Advances in Neural Information Processing Systems* 34 (2021), pp. 8780–8794.

[FBMOR24] C. Fernando, D. S. Banarse, H. Michalewski, S. Osindero, and T. Rocktäschel. "Promptbreeder: Self-Referential Self-Improvement via Prompt Evolution". In: *Forty-first International Conference on Machine Learning*. 2024.

[GSGSYDC24] Z. Gou, Z. Shao, Y. Gong, Y. Shen, Y. Yang, N. Duan, and W. Chen. "CRITIC: Large Language Models Can Self-Correct with Tool-Interactive Critiquing". In: *The Twelfth International Conference on Learning Representations*. 2024.

[HS22] J. Ho and T. Salimans. *Classifier-free diffusion guidance*. 2022. arXiv: 2207.12598 [cs.LG].

[Hüb+26] J. Hübotter, F. Lübeck, L. D. Behric, A. Baumann, M. Bagatella, D. Marta, I. Hakimi, I. Shenfeld, T. K. Buening, C. Guestrin, and A. Krause. "Reinforcement Learning via Self-Distillation". In: *Forty-third International Conference on Machine Learning*. 2026.

[JHGLYZWSSS25] N. Jain, K. Han, A. Gu, W.-D. Li, F. Yan, T. Zhang, S. Wang, A. Solar-Lezama, K. Sen, and I. Stoica. "LiveCodeBench: Holistic and Contamination Free Evaluation of Large Language Models for Code". In: *The Thirteenth International Conference on Learning Representations*. 2025.

[KBL24] M. Khanov, J. Burapacheep, and Y. Li. "ARGS: Alignment as Reward-Guided Search". In: *The Twelfth International Conference on Learning Representations*. 2024.

[KKLY24] T. Kim, J. Kim, G. Lee, and S.-Y. Yun. "Instructive decoding: Instruction-tuned large language models are self-refiner from noisy instructions". In: *The Twelfth International Conference on Learning Representations*. 2024.

[LHFLEHZL23] X. L. Li, A. Holtzman, D. Fried, P. Liang, J. Eisner, T. Hashimoto, L. Zettlemoyer, and M. Lewis. "Contrastive Decoding: Open-ended Text Generation as Optimization". In: *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Toronto, Canada: Association for Computational Linguistics, July 2023, pp. 12286–12312.

[Lin+26] Y. Lin et al. "Goedel-Prover-V2: Scaling Formal Theorem Proving with Scaffolded Data Synthesis and Self-Correction". In: *The Fourteenth International Conference on Learning Representations*. 2026.

[LXWZ23] J. Liu, C. S. Xia, Y. Wang, and L. Zhang. "Is Your Code Generated by ChatGPT Really Correct? Rigorous Evaluation of Large Language Models for Code Generation". In: *Thirty-seventh Conference on Neural Information Processing Systems*. 2023.

[Mad+23] A. Madaan et al. "Self-Refine: Iterative Refinement with Self-Feedback". In: *Thirty-seventh Conference on Neural Information Processing Systems*. 2023.

[Mud+24] S. Mudgal, J. Lee, H. Ganapathy, Y. Li, T. Wang, Y. Huang, Z. Chen, H.-T. Cheng, M. Collins, T. Strohman, J. Chen, A. Beutel, and A. Beirami. "Controlled Decoding from Language Models". In: *Forty-first International Conference on Machine Learning*. 2024.

[NY83] A. S. Nemirovsky and D. B. Yudin. *Problem Complexity and Method Efficiency in Optimization*. Wiley Series in Discrete Mathematics. Chichester; New York: Wiley-Interscience, 1983. ISBN: 978-0471103455.

[OIWGS24] T. X. Olausson, J. P. Inala, C. Wang, J. Gao, and A. Solar-Lezama. "Is self-repair a silver bullet for code generation?" In: *The Twelfth International Conference on Learning Representations*. 2024.

[PWSNHW25] D. Pandita, T. C. Weerasooriya, A. Shah, I. D. M.-X. Ng, C. M. Homan, and W. Wei. "ProRefine: Inference-time Prompt Refinement with Textual Feedback". In: *NeurIPS 2025 Workshop on Efficient Reasoning*. 2025.

[PVGLCC26] E. Penaloza, D. Vattikonda, N. Gontier, A. Lacoste, L. Charlin, and M. Caccia. *Privileged Information Distillation for Language Models*. 2026. arXiv: 2602.04942 [cs.LG].

[Qwe26] Qwen Team. *Qwen3.5-27B*. https://huggingface.co/Qwen/Qwen3.5-27B. 2026.

[RSMMEF23] R. Rafailov, A. Sharma, E. Mitchell, C. D. Manning, S. Ermon, and C. Finn. "Direct Preference Optimization: Your Language Model is Secretly a Reward Model". In: *Advances in Neural Information Processing Systems*. Vol. 36. 2023, pp. 53728–53741.

[RWGKP25] A. Rashid, R. Wu, J. Grosse, A. Kristiadi, and P. Poupart. "A Critical Look At Tokenwise Reward-Guided Text Generation". In: *Second Conference on Language Modeling*. 2025.

[SSFLB24] G. Sanchez, A. Spangher, H. Fan, E. Levi, and S. Biderman. "Stay on Topic with Classifier-Free Guidance". In: *Proceedings of the 41st International Conference on Machine Learning*. Vol. 235. Proceedings of Machine Learning Research. PMLR, July 2024, pp. 43197–43234.

[SWDRK17] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. *Proximal Policy Optimization Algorithms*. 2017. arXiv: 1707.06347 [cs.LG].

[Sha+24] Z. Shao, P. Wang, Q. Zhu, R. Xu, J. Song, X. Bi, H. Zhang, M. Zhang, Y. K. Li, Y. Wu, and D. Guo. *DeepSeekMath: Pushing the Limits of Mathematical Reasoning in Open Language Models*. 2024. arXiv: 2402.03300 [cs.CL].

[SHLTZY24] W. Shi, X. Han, M. Lewis, Y. Tsvetkov, L. Zettlemoyer, and W.-t. Yih. "Trusting Your Evidence: Hallucinate Less with Context-aware Decoding". In: *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*. 2024, pp. 783–791.

[SCGNY23] N. Shinn, F. Cassano, A. Gopinath, K. R. Narasimhan, and S. Yao. "Reflexion: language agents with verbal reinforcement learning". In: *Thirty-seventh Conference on Neural Information Processing Systems*. 2023.

[SKZ22] C. Snell, D. Klein, and R. Zhong. *Learning by Distilling Context*. 2022. arXiv: 2209.15189 [cs.CL].

[SM03] E. Süli and D. F. Mayers. *An Introduction to Numerical Analysis*. Cambridge University Press, 2003.

[Tsa88] C. Tsallis. "Possible generalization of Boltzmann-Gibbs statistics". *Journal of Statistical Physics* 52.1 (1988), pp. 479–487.

[Ven+26] S. Venkatraman, V. Jain, S. Mittal, V. Shah, J. Obando-Ceron, Y. Bengio, B. R. Bartoldson, B. Kailkhura, G. Lajoie, G. Berseth, N. Malkin, and M. Jain. *Recursive Self-Aggregation Unlocks Deep Thinking in Large Language Models*. 2026. arXiv: 2509.26626 [cs.LG].

[WL18] C.-Y. Wei and H. Luo. "More adaptive algorithms for adversarial bandits". In: *Conference On Learning Theory*. PMLR. 2018, pp. 1263–1291.

[XSKZAHG25] Y. Xu, U. M. Sehwag, A. Koppel, S. Zhu, B. An, F. Huang, and S. Ganesh. "GenARM: Reward Guided Generation with Autoregressive Reward Model for Test-Time Alignment". In: *The Thirteenth International Conference on Learning Representations*. 2025.

[Yan+25] A. Yang et al. *Qwen3 Technical Report*. 2025. arXiv: 2505.09388 [cs.CL].

[YK21] K. Yang and D. Klein. "FUDGE: Controlled Text Generation With Future Discriminators". In: *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 2021, pp. 3511–3535.

[YDWHW26] T. Ye, L. Dong, X. Wu, S. Huang, and F. Wei. *On-Policy Context Distillation for Language Models*. 2026. arXiv: 2602.12275 [cs.CL].

[ZXLHPCG26] S. Zhao, Z. Xie, M. Liu, J. Huang, G. Pang, F. Chen, and A. Grover. "Self-Distilled Reasoner: On-Policy Self-Distillation for Large Language Models". In: *Forty-third International Conference on Machine Learning*. 2026.

[ZS21] J. Zimmert and Y. Seldin. "Tsallis-inf: An optimal algorithm for stochastic and adversarial bandits". *Journal of Machine Learning Research* 22.28 (2021), pp. 1–49.