

UNDERSTANDING THE ROLE OF REASONING DATA IN PRETRAINING

Anonymous authors

Paper under double-blind review

ABSTRACT

Reasoning data expose models to multi-step reasoning, but a reasoning trace may reveal only a subset of the underlying facts it composes. How should pretraining balance atomic facts with reasoning traces? We study this question by constructing a setting in which every fact provides an atomic training signal, while reasoning traces contain only a subset of the facts they compose, enabling us to measure generalization to unseen transitions. For a linear transformer, we derive the weight assigned to atomic training signals that minimizes held-transition squared error and show that its leading-order optimum increases with reasoning-chain length and decreases with the fraction of facts held out from reasoning traces. These trends persist in a full transformer layer with an MLP, softmax attention, and cross-entropy loss, as well as in a realistic mathematical reasoning setting. Together, our results provide a principled view of how reasoning-chain structure should shape the allocation of training data during pretraining.

1 INTRODUCTION

Reasoning traces teach models how to compose facts across multiple steps, while atomic examples provide broader coverage of those facts. Under a fixed training budget, these objectives compete: spending more data on reasoning traces leaves less budget for learning individual facts, while spending more on atomic examples provides broader factual coverage but less practice with composition. This raises a fundamental question for pretraining: **how should a model’s training budget be divided between atomic facts and reasoning traces?**

This question becomes especially important as reasoning data move into pretraining, rather than being reserved for the later post-training stage. Recent work has begun to examine how reasoning data should be incorporated into pretraining. Akter et al. (2026) study the effects of when reasoning data are introduced during training, as well as their diversity and quality, while Baek et al. (2026) investigate repeatedly sampled domain data during pretraining. Recent training recipes likewise incorporate instruction data and thinking traces into Olmo 3 midtraining (Olmo et al., 2025) and SFT-style examples into Nemotron 3 Nano pretraining (Blakeman et al., 2025). Together, these developments establish reasoning-oriented data as an increasingly important component of the pretraining mixture. Yet they leave open a basic question: **how should the optimal mixture depend on the structure and coverage of reasoning data?**

The tension is subtle because reasoning traces provide something atomic examples do not: they teach the model how to use knowledge across multiple steps. But reasoning traces can also be an expensive way to teach that knowledge. Longer traces consume more of a fixed training budget, while a limited collection of traces may cover only a subset of the facts needed at evaluation. Thus, longer reasoning can simultaneously make compositional training more valuable and make broad factual coverage more important. This suggests that the optimal allocation should depend on both **reasoning-chain length** and **reasoning-trace coverage**, namely, the fraction of facts that appear inside reasoning traces. Our goal is to characterize this allocation under a fixed training budget, where differences in

learning dynamics matter even when different data mixtures might eventually reach similar performance.

We study this question first in **state tracking**, a controlled setting for multi-step reasoning over a finite environment (Merrill et al., 2024; Merrill & Sabharwal, 2025; Forner et al., 2026; Huang et al., 2026c; Li et al., 2025). Starting from an initial state, the model receives a sequence of relations and must infer the state reached after executing the entire sequence. This formulation captures capabilities that arise in practical language-model tasks, including tracking entities through a narrative, reasoning over sequences of chess moves, and executing programs (Kim & Schuster, 2023; Merrill et al., 2024). It has consequently become a standard testbed for studying language-model reasoning, both theoretically (Liu et al., 2023; Merrill et al., 2024; Merrill & Sabharwal, 2025) and empirically (Zhang et al., 2022; Liu et al., 2023; Li et al., 2025).

Our construction makes the distinction between **learning facts** and **learning to compose them** explicit. Each state transition corresponds to an atomic fact. We include every transition as an atomic training example, but construct reasoning traces using only a subset of the transitions. The remaining transitions are *held transitions*: the model sees them individually during training, but never encounters them inside a reasoning trace. Thus, a held transition is not an unseen fact—the model has learned it atomically—but an unseen *use* of that fact within a reasoning chain. This construction cleanly separates factual coverage from reasoning composition and lets us ask whether learning to compose facts transfers to facts that have never been used in a reasoning trace.

We first explore this tradeoff empirically using a one-layer transformer with softmax attention and an MLP. We vary reasoning-chain length and reasoning-trace coverage. **Perhaps surprisingly, we find that longer reasoning chains favor allocating more training data to atomic facts, not less.** Intuitively, longer traces provide richer compositional training, but each trace also consumes more of the fixed budget, making broad factual coverage increasingly important. Conversely, when reasoning traces cover more of the available facts, reasoning data become relatively more valuable. We then test whether these effects extend beyond the controlled setting using FineMath (allal et al., 2025) and MathLake (Cai et al., 2025). These experiments examine analogous effects of reasoning length and data coverage, as well as the stage of training at which reasoning data are introduced.

To explain these trends, we turn to a linear transformer, a standard simplification for analyzing transformer learning dynamics (Zhang et al., 2024; Ahn et al., 2023; Huang et al., 2025). Atomic training examples directly reinforce transitions that reasoning traces may never contain, while reasoning traces train the model to compose transitions across multiple steps. We derive the atomic-data weight that minimizes squared prediction error on held transitions. To leading order, the optimal weight **increases with reasoning-chain length and decreases with the fraction of held transitions**. In other words, longer reasoning chains favor greater allocation to atomic facts, whereas broader reasoning-trace coverage makes reasoning data relatively more valuable. These theoretical predictions agree with the trends observed in both our controlled and realistic experiments.

Contributions. Our contributions are three-fold. First, we introduce a state-tracking setting that separates **factual coverage** from **reasoning composition**: all transitions are observed atomically, but only a subset appears inside reasoning traces. Using a one-layer softmax-MLP transformer, we show that an interior mixture of atomic and reasoning examples can outperform either extreme, and characterize how the preferred mixture changes with reasoning-chain length and reasoning-trace coverage. Second, for a linear-attention model, we derive an explicit mixture law whose optimum depends on reasoning-chain length and coverage, showing why longer reasoning chains can favor greater allocation to atomic facts. Third, we develop a fixed-token FineMath-MathLake protocol to study the same allocation question in a realistic setting, varying reasoning length, data coverage, and when reasoning data are introduced during training.

2 RELATED WORK

Reasoning data in pretraining. Recent work has increasingly explored the role of reasoning data during pretraining, rather than reserving them for post-training. Akter et al. (2026) systematically study how reasoning data should be allocated between pretraining and supervised fine-tuning, showing that reasoning examples introduced during pretraining provide benefits that cannot be fully recovered by shifting the same supervision to post-training. Wang et al. (2025) find that mathematical and long-chain-of-thought data during mid-training improve both base-model reasoning and subsequent reinforcement-learning performance, while Hatamizadeh et al. (2026) incorporate explicit reasoning directly into the pretraining objective. An earlier related observation comes from Allen-Zhu & Li (2024), who, although not studying reasoning data specifically, argue from controlled knowledge-learning experiments that instruction-finetuning data should be mixed into pretraining rather than reserved entirely for post-training. These works establish the value of introducing richer forms of supervision during pretraining, but leave open a more fundamental question: *how should reasoning data be balanced against atomic data when the training budget is fixed?* We study this tradeoff by separating factual coverage from reasoning composition: every fact is observed atomically, but only a subset is exposed within reasoning traces. This allows us to characterize when reasoning data are most valuable and how their benefit depends on reasoning-chain length and coverage.

Transformer theory. A growing body of work uses controlled synthetic tasks to understand how Transformers acquire, retrieve, and compose structured knowledge. Nichani et al. (2025) analyze factual recall through associative memories, showing how shallow Transformers can store and retrieve large collections of subject-relation-object associations. Huang et al. (2026b) study how attention learns to select relevant relations and retrieve corresponding facts from a pretrained associative memory. Most closely related to our setting, Huang et al. (2026c) prove that Transformers trained by gradient descent can learn chain-of-thought state tracking and characterize how the algebraic structure of the transition system governs length generalization through attention concentration. These works provide insight into how Transformers store and compose knowledge, but do not study how different forms of training data interact under a fixed budget. We instead consider a setting in which every transition is observed through atomic examples, while only a subset appears in reasoning traces. This cleanly separates **factual availability** from **exposure to multi-step composition**, allowing us to analyze how the allocation between atomic and reasoning data affects the ability to compose transitions that are known individually but absent from the reasoning data.

3 TOY SETTING: STATE TRACKING

3.1 PRELIMINARIES

Definition 1 (State-tracking world and held transitions). *Let $[K] = \{1, \dots, K\}$ and $N, M \geq 2$. Entities are indexed by $e \in [N]$, and each relation $r \in [M]$ acts through a fixed permutation $\pi_r : [N] \rightarrow [N]$. The full fact table $\mathcal{A} = [N] \times [M]$ contains $D = NM$ entity-relation pairs. A length- L chain $\chi = (r_1, \dots, r_L, e_0) \in [M]^L \times [N]$, with $L \geq 2$, induces the state sequence*

$$e_\ell = \pi_{r_\ell}(e_{\ell-1}), \quad \ell \in [L].$$

For an integer $1 \leq m < D$, draw the held set $\mathcal{H}$ and define its complement $\mathcal{S}$ by

$$\mathcal{H} \sim \text{Unif}\{H \subseteq \mathcal{A} : |H| = m\}, \quad q = \frac{m}{D}, \quad p = 1 - q, \quad \mathcal{S} = \mathcal{A} \setminus \mathcal{H}.$$

The held fraction q is the fraction of facts excluded from reasoning traces, and the reasoning coverage p is the fraction available to those traces. The mask is sampled once and fixed throughout each run.

Definition 2 (Atomic, reasoning, and evaluation populations). *An atomic example samples e_0, r_1 independently and uniformly and $r_2, \dots, r_L$ all belong to $\mathcal{S}$, but supervises only $e_1 = \pi_{r_1}(e_0)$. Its first entity-relation pair is therefore uniform over $\mathcal{A}$, including pairs in $\mathcal{H}$. A*

reasoning example is sampled uniformly from chains whose actual transitions all belong to $\mathcal{S}$ and supervises the complete state trace $e_1, \dots, e_L$. Formally,

$$\begin{aligned}\mathcal{D}_{\text{at}} &= \mathcal{D}_{\text{full}} = \text{Unif}([M]^L \times [N]), \\ \mathcal{C}_L(\mathcal{S}) &= \{\chi \in [M]^L \times [N] : (e_{\ell-1}, r_\ell) \in \mathcal{S} \text{ for every } \ell \in [L]\}, \\ \mathcal{D}_{\text{rea}}(\mathcal{S}, L) &= \text{Unif}(\mathcal{C}_L(\mathcal{S})).\end{aligned}$$

We write $\mathbb{E}_{\chi \sim \mathcal{D}_{\text{rea}}}$ for expectation over the reasoning population associated with the fixed set $\mathcal{S}$.

Learning objective. Entities and relations are represented by distinct typed identity tokens. Both data sources use the same relation-first input format:

$$\begin{aligned}\text{atomic: } & [r_1, \dots, r_L, e_0] \longrightarrow [e_1], \\ \text{reasoning: } & [r_1, \dots, r_L, e_0] \longrightarrow [e_1, \dots, e_L].\end{aligned}$$

The shared prefix places the first prediction at the same absolute position and supplies the same number of candidate relations in both sources. Atomic examples supervise only the first transition, whereas reasoning examples additionally supervise subsequent prediction roles. Every fact is available through atomic supervision, while reasoning traces cover only a subset, allowing us to study whether facts learned individually can be used within multi-step computations.

Let $\beta \in [0, 1]$ be the atomic *loss weight*. We average the reasoning loss over its L state predictions before combining the two sources:

$$\begin{aligned}\mathcal{L}_\beta^{\text{CE}}(\theta) &= \beta \mathbb{E}_{\chi \sim \mathcal{D}_{\text{at}}} \ell_{\text{CE},1}(\theta; \chi) \\ &\quad + (1 - \beta) \mathbb{E}_{\chi \sim \mathcal{D}_{\text{rea}}} \left[\frac{1}{L} \sum_{\ell=1}^L \ell_{\text{CE},\ell}(\theta; \chi) \right],\end{aligned}\tag{1}$$

where

$$\ell_{\text{CE},\ell}(\theta; \chi) = -\log p_\theta(e_\ell \mid [r_1, \dots, r_L, e_0, \dots, e_{\ell-1}]), \quad \ell \in [L].$$

Each cross-entropy term supervises the next entity using the true preceding states. A test chain is called *held* if at least one of its ground-truth transitions belongs to $\mathcal{H}$.

3.2 ONE FULL LAYER TRANSFORMER

Our learner consists of a single causal self-attention head followed by a position-wise multi-layer perceptron (MLP). We encode token identities and positions in separate one-hot blocks and directly parameterize the merged query-key matrix, following Zhang et al. (2024); Ahn et al. (2023); Huang et al. (2025; 2026c;b). We first define the input representation, followed by the attention and output maps.

Definition 3 (Token and positional encoding). *Let $u_r \in \mathbb{R}^M$ and $v_e \in \mathbb{R}^N$ be the standard basis vectors corresponding to relation r and entity e , respectively. Fix a maximum input-prefix length $T \geq 2L$, shared by atomic and reasoning examples. The token-content vectors are*

$$c(r) = \begin{pmatrix} u_r \\ 0_N \end{pmatrix}, \quad c(e) = \begin{pmatrix} 0_M \\ v_e \end{pmatrix}.$$

For a token x_t at position t , let $\mathbf{c}_t := c(x_t)$ and let $\mathbf{p}_t \in \mathbb{R}^T$ be the t -th standard basis vector. Concatenating content and position gives

$$\tilde{\mathbf{x}}_t = \begin{pmatrix} \mathbf{c}_t \\ \mathbf{p}_t \end{pmatrix} \in \mathbb{R}^d, \quad d = M + N + T, \quad \widetilde{\mathbf{X}}_t = [\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_t] \in \mathbb{R}^{d \times t}.$$

All content and positional vectors are fixed. Relation, entity, and position coordinates occupy disjoint blocks of sizes M , N , and T , respectively. We write $\mathbf{X}_t = [\mathbf{c}_1, \dots, \mathbf{c}_t] \in \mathbb{R}^{(M+N) \times t}$ for the corresponding content-only matrix.

Both data sources predict e_1 at position $L + 1$, and reasoning role ℓ predicts e_ℓ from the prefix ending at position $L + \ell$. The last supervised query is therefore at position $2L$, so $T \geq 2L$ covers all input prefixes. We reserve T for the maximum prefix length and use τ for the training horizon below.

We parameterize the merged query-key scores directly by a trainable matrix $\mathbf{A} \in \mathbb{R}^{d \times d}$. In particular, the score assigned to key position i by query position j is $\tilde{\mathbf{x}}_i^\top \mathbf{A} \tilde{\mathbf{x}}_j$.

Definition 4 (Causal self-attention head). For $\mathbf{A} \in \mathbb{R}^{d \times d}$ and a prefix $\tilde{\mathbf{X}}_t \in \mathbb{R}^{d \times t}$, define

$$\text{attn}(\tilde{\mathbf{X}}_t; \mathbf{A}) := \tilde{\mathbf{X}}_t \text{softmax} \left(\text{MASK}(\tilde{\mathbf{X}}_t^\top \mathbf{A} \tilde{\mathbf{X}}_t) \right) \in \mathbb{R}^{d \times t},$$

where, for $S \in \mathbb{R}^{t \times t}$,

$$[\text{MASK}(S)]_{ij} = \begin{cases} S_{ij}, & i \leq j, \\ -\infty, & i > j. \end{cases}$$

The softmax is applied column-wise.

The attention output is passed through a shared MLP to produce next-token logits over the entity and relation vocabulary.

Definition 5 (One-layer attention-MLP Transformer). Fix an MLP width $m \geq 1$ and a coordinatewise nonlinearity ϕ . The trainable matrices are

$$\mathbf{A} \in \mathbb{R}^{d \times d}, \quad \mathbf{W} \in \mathbb{R}^{m \times d}, \quad \mathbf{V} \in \mathbb{R}^{m \times (M+N)},$$

with parameters $\boldsymbol{\theta}_{\text{TF}} = (\mathbf{A}, \mathbf{V}, \mathbf{W})$. At query position t , the vocabulary-logit vector is

$$f_{\text{TF},t}(\tilde{\mathbf{X}}_t; \boldsymbol{\theta}_{\text{TF}}) = \mathbf{V}^\top \phi \left(\mathbf{W} [\text{attn}(\tilde{\mathbf{X}}_t; \mathbf{A})]_{:,t} \right) \in \mathbb{R}^{M+N}.$$

The next-token distribution is

$$p_{\boldsymbol{\theta}_{\text{TF}}}(\cdot | \tilde{\mathbf{X}}_t) = \text{softmax} \left(f_{\text{TF},t}(\tilde{\mathbf{X}}_t; \boldsymbol{\theta}_{\text{TF}}) \right).$$

The first M output coordinates index relations, and the last N index entities. Accordingly, the probability assigned to entity $e \in [N]$ is component $M + e$ of this distribution.

We train this model using the mixed cross-entropy objective in equation 1.

Effect of data mixture Figure 1 shows that mixing atomic and reasoning examples improves OOD accuracy over either pure-source baseline, suggesting that the two forms of supervision play complementary roles. The empirically preferred atomic example ratio generally increases with reasoning length L and decreases with the held-out fact fraction q . This pattern is not strictly monotone across all selected grid maximizers: the $L \in \{3, 4\}$ cases at $q = 30\%$ deviate from the held-fraction trend. However, these panels exhibit near-optimal plateaus, with several mixture ratios attaining very similar mean OOD accuracy.

3.3 LINEAR-ATTENTION AND THE OPTIMAL MIXTURE

We next examine how the allocation between atomic examples and reasoning traces affects optimization. Our goal is to characterize the mixture that best supports the use of held transitions after a common training horizon, and to understand how this allocation depends on reasoning length and coverage. To make this interaction tractable, we study a scalar linear-attention surrogate

Following prior block-structured analyses of Transformer learning (Nichani et al., 2024; Huang et al., 2025; 2026a;b; Javanmard et al., 2026b;a), we use a block-diagonal attention score matrix with token-token and position-position blocks, fixing the mixed token-position blocks to zero. These two components contribute additively to the softmax logits, but their contributions multiply after exponentiation. This motivates combining the bilinear content scores with a learned positional gate through the Hadamard product in the linear surrogate:

$$(\mathbf{X}_t^\top \mathbf{B} \mathbf{X}_t) \odot \tilde{\mathbf{C}}_{1:t, 1:t}.$$




Figure 1: **Effect of atomic-reasoning mixing on state-tracking accuracy.** Final-answer accuracy as a function of the atomic ratio in a synthetic world with 64 entities and 8 relations, each sampled as an independent random permutation. All runs use 4000 AdamW updates with learning rate 0.002, weight decay 0.01, and batch size 256. Curves and shaded regions indicate the mean and ± 1 standard deviation across three training seeds. The empirically optimal atomic example ratio generally increases with reasoning length L and decreases with the held-out fact fraction q . The selected grid maximizers for $L \in \{3, 4\}$ at $q = 30\%$ deviate from the latter trend; however, these panels exhibit near-optimal plateaus, with several tested mixture ratios achieving very similar mean OOD accuracy.

Here $\mathbf{B}$ controls token interactions, while $\tilde{\mathbf{C}}$ controls positional routing. In particular, a zero gate entry excludes a connection, whereas an additive zero would leave its content score unchanged. Removing attention normalization and replacing the nonlinear output map with a scalar linear readout gives the following model.

Definition 6 (Scalar linear-attention model). Let $\mathbf{X}_t \in \mathbb{R}^{(M+N) \times t}$ contain the fixed typed one-hot content vectors, with $t \leq T$ and $T \geq 2L$. For $\mathbf{W} \in \mathbb{R}^{1 \times (M+N)}$, $\mathbf{B} \in \mathbb{R}^{(M+N) \times (M+N)}$, and $\tilde{\mathbf{C}} \in \mathbb{R}^{T \times T}$, define

$$f_{\text{lin}}(\mathbf{X}_t; \mathbf{W}, \mathbf{B}, \tilde{\mathbf{C}}) = \mathbf{W} \mathbf{X}_t \left[(\mathbf{X}_t^\top \mathbf{B} \mathbf{X}_t) \odot \tilde{\mathbf{C}}_{1:t, 1:t} \right] \in \mathbb{R}^{1 \times t}. \quad (2)$$

The trainable coordinates are the relation values $w_r = \mathbf{W}_{1,r}$, the relation-entity scores $b_{r,e} = \mathbf{B}_{r,M+e}$, and the routing coefficients $C_{j\ell} = \tilde{\mathbf{C}}_{j,L+\ell}$ for $j, \ell \in [L]$; all remaining entries are fixed at zero. Write $\theta = (w, b, C)$.

At role ℓ , the teacher-forced prefix ends at $e_{\ell-1}$ and has length $L + \ell$. Let $f_\ell(\chi; \theta)$ be entry $(1, L + \ell)$ of $f_{\text{lin}}(\mathbf{X}_{L+\ell}; \theta)$. For entity $e \in [N]$, define

$$\zeta_e = \frac{2e - N - 1}{N}, \quad \hat{e}_\ell \in \arg \min_{e \in [N]} |f_\ell(\chi; \theta) - \zeta_e|. \quad (3)$$

Squared loss and joint training. Atomic prefixes sample the initial entity and all relation labels independently and uniformly; reasoning prefixes are uniform over chains whose actual transitions belong to $\mathcal{S}$. For an atomic loss weight $\beta \in [0, 1]$, the mixed population

loss is

$$L_\beta(\theta; \mathcal{S}) = \frac{\beta}{2} \mathbb{E}_{\chi \sim \mathcal{D}_{\text{at}}} (f_1(\chi; \theta) - \zeta_{e_1})^2 + \frac{1-\beta}{2L} \sum_{\ell=1}^L \mathbb{E}_{\chi \sim \mathcal{D}_{\text{rea}}} (f_\ell(\chi; \theta) - \zeta_{e_\ell})^2. \quad (4)$$

Atomic examples supervise only the first state, whereas reasoning examples supervise the complete trace. Averaging the reasoning loss over its L predictions makes β the relative weight of the two source losses.

All active entries of $\mathbf{W}, \mathbf{B}, \tilde{\mathbf{C}}$ are jointly trained by gradient flow:

$$\dot{\theta}_\beta(t) = -\nabla_\theta L_\beta(\theta_\beta(t); \mathcal{S}), \quad \theta_\beta(0) = \theta_0. \quad (5)$$

We initialize $\mathbf{B}$ to zero and use the same initialization for every mixture:

$$\mathbf{B}(0) = 0, \quad w_r(0) = 1 \quad (r \in [M]), \quad C(0) = E_{11}. \quad (6)$$

Here E_{11} has a single unit entry at $(1, 1)$ and zeros elsewhere.

Held-transition error and optimal allocation. The training objective alone does not measure whether facts learned from atomic supervision can be used at later prediction roles. We therefore evaluate later transitions that were excluded from reasoning traces. For each $\ell \geq 2$, let

$$\mathcal{D}_{H,\ell} := \mathcal{D}_{\text{full}}(\cdot \mid (e_{\ell-1}, r_\ell) \in \mathcal{H}).$$

The conditioning event has probability $q > 0$, so this distribution is well-defined. Define the held-transition error by

$$\mathcal{E}_H(\theta) := \frac{1}{2(L-1)} \sum_{\ell=2}^L \mathbb{E}_{\chi \sim \mathcal{D}_{H,\ell}} (f_\ell(\chi; \theta) - \zeta_{e_\ell})^2.$$

At a training horizon $\tau > 0$, define the optimal atomic weight as

$$\beta_{\tau,S}^* \in \arg \min_{\beta \in [0,1]} \mathcal{E}_H(\theta_\beta(\tau)). \quad (7)$$

We now characterize how reasoning length and coverage affect the optimal atomic loss weight. The following theorem shows that, at a fixed training horizon, $\beta_{\tau,S}^*$ concentrates around a deterministic center as the numbers of entities and relations grow. This center increases with reasoning length and decreases with the held fraction, describing how the preferred allocation changes with the task and the available reasoning supervision.

Theorem 1 (Length–coverage trend). *Fix $L \geq 2$, $p = 1 - q \in (0, 1)$, $\tau > 0$, and $\delta \in (0, 1)$. W.L.O.G, assume $\pi_r(e) = 1 + ((e + r - 2) \bmod N)$, $M = N$. Consider gradient flow on Equation (4), and $B(0) = 0$, $w_r(0) = 1$, $C(0) = E_{11}$. Let $\mathcal{H}$ be uniform among sets of qN^2 facts, where qN^2 is an integer, and put $\mathcal{S} = \mathcal{A} \setminus \mathcal{H}$. For all sufficiently large N , with probability at least $1 - \delta$, the finite-time MSE-optimal atomic weight $\beta_{\tau,S}^*$ is unique in $(1/3, 2/3)$, and*

$$|\beta_{\tau,S}^* - \beta(L(1-q))| \leq \mathfrak{R}_N, \quad \mathfrak{R}_N = O_{L,p,\tau} \left(\frac{\sqrt{\log(2/\delta)}}{N} \right), \quad (8)$$

where

$$\beta(z) = \frac{1}{2 - z + \sqrt{z^2 - z + 1}}, \quad z > 0. \quad (9)$$

The center increases with L and decreases with q .

We refer the complete proof to Section A. The theorem predicts a larger leading optimal atomic weight for longer reasoning chains and a smaller weight when a larger fraction of facts is excluded from reasoning traces, which is consistent with Figure 1.

4 EXPERIMENTS ON NATURAL LANGUAGE

Balancing general math text and reasoning data We pretrain Olmo-3 150M on mixtures of FineMath, a general math web-text corpus, and the OpenMathInstruct-2 subset of MathLake, a mathematical reasoning dataset. We sweep the MathLake fraction from 0% to 100% in 10% steps. The total budget is fixed at 10B tokens, so higher MathLake fractions also mean more repetition of its 277M tokens (up to ~ 36 epochs). Prior work either varies composition and token count together (Zhao et al., 2025), uses a single reasoning fraction (Akter et al., 2026), or tests only small domain fractions (Baek et al., 2026); none includes both pure-source endpoints, so none shows that a mixture beats either source alone.

Table 1 shows that both extremes are suboptimal. Pure FineMath reaches a GSM8K CE of 1.785, while pure MathLake degrades to 7.513. Although its training loss falls to 0.2, GSM8K CE rises throughout training, consistent with overfitting to a heavily repeated corpus. A small share of reasoning data helps significantly, with CE dropping to 1.613 at 10% MathLake and reaching its minimum of 1.608 at 20%. Beyond this point, CE rises monotonically, returning to the pure FineMath level at 50%. We therefore run subsequent experiments at 10%, 20%, and 30% MathLake, covering the optimum and the start of the degradation regime.

Table 1: Mixture-ratio sweep between FineMath and MathLake (Olmo-3 150M, 10B tokens). Zero-shot GSM8K test CE and PPL over the solution sequence. Best in bold.

FM (%)	ML (%)	CE $\downarrow$	PPL $\downarrow$
100	0	1.785	5.96
90	10	1.613	5.02
80	20	1.608	4.99
70	30	1.674	5.33
60	40	1.689	5.41
50	50	1.768	5.86
40	60	1.817	6.16
30	70	2.035	7.66
20	80	2.148	8.57
10	90	2.402	11.04
0	100	7.513	1831.40




Figure 2: GSM8K CE at step 5000. (a) Varying the fraction p of the MathLake pool available for training, at three FineMath/MathLake ratios. (b) Training on the easy or hard half of MathLake versus a random 50% subset.

Quantity and difficulty of reasoning data At each of the mixture ratios above, we vary two properties of the reasoning data.

Held-out ratio. We control the fraction $p \in \{25\%, 50\%, 75\%, 100\%\}$ of the MathLake pool available for training and withhold the rest. Since the mixture ratio and total budget are fixed, every run sees the same number of reasoning tokens. To make the held-out ratio a monotone intervention, the pools are nested. We shuffle MathLake once at the document level and take each pool as a prefix of the next, where the 25% pool is a subset of the 50% pool, which is a subset of the 75% pool. Thus, moving from a smaller to a larger p only adds unseen documents and never swaps one sample for another.

Difficulty. We use MathLake’s per-problem difficulty scores (1–10) as a proxy for reasoning length and complexity. The scores closely track solution length (258 tokens at level 1 vs. 680 at level 9), so splitting by level alone would also split by token volume. We instead cut the corpus at its token-weighted median, which falls between levels 4 and 5, and cap

Table 2: Timing of reasoning data. Each run trains Olmo-3 150M for 10k steps with MathLake making up 20% of all tokens; the first K steps use pure FineMath, so the MathLake share while mixing rises with K . Zero-shot GSM8K test CE and PPL over the solution sequence at the final checkpoint. Best in bold.

Pure FM steps (K)	ML share while mixing	CE ↓	PPL ↓
0	20.0%	1.744	5.72
1k	22.2%	1.749	5.75
2k	25.0%	1.754	5.78
3k	28.6%	1.767	5.85
4k	33.3%	1.771	5.88
6k	50.0%	1.853	6.38
8k	100.0%	2.411	11.14

both halves at 133.9M tokens, so easy and hard runs see identical reasoning-token budgets at identical repetition rates. A random 50% pool (138.7M tokens) serves as the control.

As shown by Figure 2a, the optimal mixing ratio increases with p , therefore decreases with held-out ratio. In Figure 2b, the optimal mixing ratio increases with difficulty. The results are consistent with Theorem 1.

When to introduce reasoning data Finally, we fix the total reasoning budget and vary when it enters training. Each run lasts 10k steps with MathLake making up 20% of all training tokens, but the first K steps use FineMath only. The full reasoning budget must then fit into the remaining $10k - K$ steps, so MathLake’s share of the post-switch mixture is 20%, 25%, 33%, 50%, and 100% for $K = 0, 2k, 4k, 6k, 8k$. We implement this by resuming each arm from the step- K checkpoint of a single pure-FineMath run, carrying over the step count and learning-rate schedule.

Delaying reasoning data brings no benefit (Table 2): GSM8K CE increases monotonically with K . Since starting early is never worse, we recommend mixing in reasoning data from the beginning of training, consistent with Baek et al. (2026), who find that domain data is best introduced early.

5 CONCLUSION

We studied a fundamental question in reasoning-oriented pretraining: how should a fixed training budget be divided between atomic facts and reasoning traces? By separating factual coverage from exposure to multi-step composition, we showed that the optimal allocation depends systematically on reasoning-chain length and reasoning-trace coverage. Perhaps surprisingly, longer reasoning chains can favor greater allocation to atomic facts, as richer compositional training comes at the cost of broader factual coverage. Our theoretical analysis explains this effect in a simplified linear setting, and experiments with full Transformers and realistic mathematical data demonstrate that the same qualitative trends persist beyond the controlled regime. Together, our results provide a principled perspective on how the structure and coverage of reasoning data should shape pretraining mixtures.

Limitations. Our controlled setting abstracts away heterogeneous data sources and varying data quality in large-scale pretraining. We also focus on pretraining and do not study how different data mixtures affect subsequent SFT and RL, which we leave for future work.

AI USE STATEMENT

We have not used generative AI tools for tasks with required disclosure. Additionally, we used generative AI tools for polishing the paper. We have reviewed all AI-assisted work. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REFERENCES

- Kwangjun Ahn, Xiang Cheng, Hadi Daneshmand, and Suvrit Sra. Transformers learn to implement preconditioned gradient descent for in-context learning. *Advances in Neural Information Processing Systems*, 36:45614–45650, 2023.
- Syeda Nahida Akter, Shrimai Prabhumoye, Eric Nyberg, Mostofa Patwary, Mohammad Shoeybi, Yejin Choi, and Bryan Catanzaro. Front-loading reasoning: The synergy between pretraining and post-training data. In *International Conference on Learning Representations*, volume 2026, pp. 80909–80928, 2026.
- Loubna Ben allal, Anton Lozhkov, Elie Bakouch, Gabriel Martin Blazquez, Guilherme Penedo, Lewis Tunstall, Andrés Marafioti, Agustín Piqueres Lajarín, Hynek Kydlíček, Vaibhav Srivastav, Joshua Lochner, Caleb Fahlgren, Xuan Son NGUYEN, Ben Burtenshaw, Clémentine Fourrier, Haojun Zhao, Hugo Larcher, Mathieu Morlon, Cyril Zakka, Colin Raffel, Leandro Von Werra, and Thomas Wolf. SmoLLM2: When smol goes big — data-centric training of a fully open small language model. In *Second Conference on Language Modeling*, 2025. URL <https://openreview.net/forum?id=3JiCl2A14H>.
- Zeyuan Allen-Zhu and Yuanzhi Li. Physics of language models: Part 3.1, knowledge storage and extraction. In *Forty-first International Conference on Machine Learning*, 2024. URL <https://openreview.net/forum?id=5x788rqbcbj>.
- Christina Baek, Ricardo Pio Monti, David J. Schwab, Amro Abbas, Rishabh Adiga, Cody Blakeney, Maximilian Böther, Paul Burststein, Aldo Gael Carranza, Alvin Deng, Parth Doshi, Vineeth Dorna, Alex Fang, Tony Jiang, Siddharth Joshi, Brett W. Larsen, Jason Chan Lee, Katherine L. Mentzer, Luke Merrick, Haakon Mongstad, Fan Pan, Anshuman Suri, Darren Teh, Jason Telanoff, Jack Urbanek, Zhengping Wang, Josh Wills, Haoli Yin, Aditi Raghunathan, J Zico Kolter, Bogdan Giza, Ari S. Morcos, Matthew L Leavitt, and Pratyush Maini. The finetuner’s fallacy: When to pretrain with your finetuning data. In *Third Conference on Language Modeling*, 2026. URL <https://openreview.net/forum?id=HgDsWB3KFB>.
- Aaron Blakeman, Aaron Grattafiori, Aarti Basant, Abhibha Gupta, Abhinav Khattar, Adi Renduchintala, Aditya Vavre, Akanksha Shukla, Akhiad Bercovich, Aleksander Ficek, et al. Nemotron 3 nano: Open, efficient mixture-of-experts hybrid mamba-transformer model for agentic reasoning. *arXiv preprint arXiv:2512.20848*, 2025.
- Mengzhang Cai, Xin Gao, Yu Li, Honglin Lin, Zheng Liu, Zhuoshi Pan, Qizhi Pei, Xiaoran Shang, Mengyuan Sun, Zinan Tang, et al. Opendataarena: A fair and open arena for benchmarking post-training dataset value. *arXiv preprint arXiv:2512.14051*, 2025.
- Niklas Forner, Marcel Kühn, Matthias Thamm, and Bernd Rosenow. Learning dynamics of chain-of-thought state tracking in a solvable transformer model. *arXiv preprint arXiv:2606.18164*, 2026.
- Ali Hatamizadeh, Syeda Nahida Akter, Shrimai Prabhumoye, Jan Kautz, Mostofa Patwary, Mohammad Shoeybi, Bryan Catanzaro, and Yejin Choi. Rlp: Reinforcement as a pretraining objective. In *International Conference on Learning Representations*, volume 2026, pp. 41375–41402, 2026.
- Jianhao Huang, Zixuan Wang, and Jason D. Lee. Transformers learn to implement multi-step gradient descent with chain of thought. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=r3DF5s0o5B>.
- Jianhao Huang, Zhanpeng Zhou, Renqiu Xia, Baharan Mirzasoleiman, Weijie J Su, and Wei Huang. How transformers learn to plan via multi-token prediction. In *Third Conference on Language Modeling*, 2026a. URL <https://openreview.net/forum?id=1uj8eRv3aJ>.
- Ruomin Huang, Eshaan Nichani, Jason D Lee, and Rong Ge. Fine-tuning dynamics of in-context factual recall in transformers. *arXiv preprint arXiv:2605.27774*, 2026b.

- Yu Huang, Zixin Wen, Aarti Singh, Yuejie Chi, and Yuxin Chen. Transformers provably learn chain-of-thought reasoning with length generalization. *Advances in Neural Information Processing Systems*, 38:126815–126965, 2026c.
- Adel Javanmard, Baharan Mirzasoleiman, and Vahab Mirrokni. Theoretical perspectives on data quality and synergistic effects in pre- and post-training reasoning models. In *Forty-third International Conference on Machine Learning*, 2026a. URL <https://openreview.net/forum?id=URIdUFsz9E>.
- Adel Javanmard, Baharan Mirzasoleiman, and Vahab Mirrokni. Understanding the role of training data in test-time scaling. In *The Fourteenth International Conference on Learning Representations*, 2026b. URL <https://openreview.net/forum?id=Y9FfDNa2nJ>.
- Najoung Kim and Sebastian Schuster. Entity tracking in language models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 3835–3855, 2023.
- Belinda Z Li, Zifan Carl Guo, and Jacob Andreas. (how) do language models track state? *arXiv preprint arXiv:2503.02854*, 2025.
- Bingbin Liu, Jordan T. Ash, Surbhi Goel, Akshay Krishnamurthy, and Cyril Zhang. Transformers learn shortcuts to automata. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=De4FYqjFueZ>.
- William Merrill and Ashish Sabharwal. A little depth goes a long way: The expressive power of log-depth transformers. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL <https://openreview.net/forum?id=5pHfYe10iX>.
- William Merrill, Jackson Petty, and Ashish Sabharwal. The illusion of state in state-space models. In *Forty-first International Conference on Machine Learning*, 2024. URL <https://openreview.net/forum?id=QZgo9JZpLq>.
- Eshaan Nichani, Alex Damian, and Jason D. Lee. How transformers learn causal structure with gradient descent. In *Forty-first International Conference on Machine Learning*, 2024. URL <https://openreview.net/forum?id=jNM4imlHZv>.
- Eshaan Nichani, Jason Lee, and Alberto Bietti. Understanding factual recall in transformers via associative memories. In *International Conference on Learning Representations*, volume 2025, pp. 4166–4207, 2025.
- Team Olmo, Allyson Ettinger, Amanda Bertsch, Bailey Kuehl, David Graham, David Heine-man, Dirk Groeneveld, Faeze Brahman, Finbarr Timbers, Hamish Ivison, et al. Olmo 3. *arXiv preprint arXiv:2512.13961*, 2025.
- Zengzhi Wang, Fan Zhou, Xuefeng Li, and Pengfei Liu. Octothinker: Mid-training incentivizes reinforcement learning scaling. In *2nd AI for Math Workshop @ ICML 2025*, 2025. URL <https://openreview.net/forum?id=chCeUHjLzs>.
- Ruiqi Zhang, Spencer Frei, and Peter L Bartlett. Trained transformers learn linear models in-context. *Journal of Machine Learning Research*, 25(49):1–55, 2024.
- Yi Zhang, Arturs Backurs, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, and Tal Wagner. Unveiling transformers with lego: a synthetic reasoning task. *arXiv preprint arXiv:2206.04301*, 2022.
- Rosie Zhao, Alexandru Meterez, Sham Kakade, Cengiz Pehlevan, Samy Jelassi, and Eran Malach. Echo chamber: Rl post-training amplifies behaviors learned in pretraining. *arXiv preprint arXiv:2504.07912*, 2025.

A PROOF OF THE LENGTH-COVERAGE THEOREM

Throughout, $y_{r,e} = \zeta_{\pi_r(e)}$, $a_{\ell j} = (e_{\ell-1}, r_j)$, $a_\ell = a_{\ell\ell}$, $D = N^2$, $\lambda = 1 - \beta$, and $B(0) = 0$. For a family $h(\beta)$, define

$$\|h\|_{[2]} := \sum_{k=0}^2 \frac{1}{k!} \sup_{\beta \in [0,1]} \|\partial_\beta^k h(\beta)\|. \quad (10)$$

The base norm is entrywise maximum for tables/vectors and $\|C\|_1 = \sum_{j,\ell} |C_{j\ell}|$. A multilinear map of base operator norm c satisfies

$$\begin{aligned} \|\mathcal{B}(h_1, h_2)\|_{[2]} &\leq c \sum_{i+j \leq 2} \frac{\sup_\beta \|\partial_\beta^i h_1\|}{i!} \frac{\sup_\beta \|\partial_\beta^j h_2\|}{j!} \leq c \|h_1\|_{[2]} \|h_2\|_{[2]}, \\ \|\beta\|_{[2]} &= \|1 - \beta\|_{[2]} = 2. \end{aligned} \quad (11)$$

The same calculation applies to higher multilinear products.

Lemma 1 (Balance and the fixed-mask cubic). *For a fixed mask with a nonempty reasoning population, put*

$$\begin{aligned} \nu &= \frac{N^2 - 1}{3N^2}, \quad \nu_H = \frac{1}{qD} \sum_{a \in H} y_a^2, \quad \omega_a = \frac{D}{L} \Pr(a_1 = a), \\ K_{j\ell} &= \mathbb{E}_S[y_{a_{\ell j}} y_{a_\ell}], \quad U_{j\ell} = \mathbb{E}_S[\omega_{a_{\ell j}} y_{a_{\ell j}} y_{a_\ell}], \\ Z_S &= \frac{1}{qDN} \sum_e \left(\sum_{r: (e,r) \in H} y_{r,e} \right) \left(\sum_r \omega_{r,e} y_{r,e} \right). \end{aligned} \quad (12)$$

For $\ell \geq 2$,

$$\begin{aligned} \mathbb{E}_{H,\ell}[y_{a_{\ell j}} y_{a_\ell}] &= \nu_H \mathbf{1}\{j = \ell\}, \\ \mathbb{E}_{H,\ell}[\omega_{a_{\ell j}} y_{a_{\ell j}} y_{a_\ell}] &= Z_S \mathbf{1}\{j \neq \ell\}. \end{aligned} \quad (13)$$

Define

$$\begin{aligned} P_S(\beta) &= \frac{\nu_H \beta \lambda}{L-1} \sum_{\ell=2}^L [\beta K_{\ell\ell} + \lambda U_{\ell\ell}] \\ &\quad + \frac{Z_S \lambda^2}{L-1} \sum_{\ell=2}^L \sum_{j \neq \ell} [\beta K_{j\ell} + \lambda U_{j\ell}]. \end{aligned} \quad (14)$$

This is the leading gain polynomial identified in Lemma 3.

Proof. Complete shifts and the scalar code give

$$\frac{1}{N} \sum_r y_{r,e} = \frac{1}{N} \sum_{x=1}^N \zeta_x = 0, \quad \frac{1}{N} \sum_r y_{r,e}^2 = \frac{1}{N} \sum_{x=1}^N \zeta_x^2 = \nu, \quad \frac{1}{4} \leq \nu < \frac{1}{3}. \quad (15)$$

For every relation tuple $\mathbf{r}$,

$$\Pr_{\text{full}}(e_{\ell-1} = e \mid \mathbf{r}) = \frac{1}{N}, \quad \Pr_{\text{full}}(a_\ell = a) = \frac{1}{D}, \quad \Pr_{\text{full}}(r_j = t \mid a_\ell = (e, r)) = \frac{1}{N} \quad (j \neq \ell). \quad (16)$$

Consequently, for $j \neq \ell$,

$$\begin{aligned} \mathbb{E}_{H,\ell}[y_{a_{\ell j}} y_{a_\ell}] &= \frac{1}{qD} \sum_{(e,r) \in H} y_{r,e} \underbrace{\frac{1}{N} \sum_t y_{t,e}}_0 = 0, \\ \mathbb{E}_{H,\ell}[\omega_{a_{\ell j}} y_{a_{\ell j}} y_{a_\ell}] &= \frac{1}{qDN} \sum_{(e,r) \in H} y_{r,e} \sum_t \omega_{t,e} y_{t,e} = Z_S. \end{aligned} \quad (17)$$

For $j = \ell$, the two averages are ν_H and 0, since $a_\ell \in H \Rightarrow \omega_{a_\ell} = 0$. Substitution in

$$\frac{\lambda}{L-1} \sum_{\ell=2}^L \sum_j (\beta K_{j\ell} + \lambda U_{j\ell}) \mathbb{E}_{H,\ell}[(\beta + \lambda \omega_{a_{\ell j}}) y_{a_{\ell j}} y_{a_\ell}]$$

gives equation 14. In general, $Z_S \neq 0$. $\square$

Lemma 2 (Concentration of the balanced cubic). *Fix $L \geq 2$, $p \in (0, 1)$, and $\delta \in (0, 1)$. For sufficiently large N , with probability at least $1 - \delta$,*

$$F_0(S) := \Pr_{\text{full}}(\chi \text{ is valid}) \geq p^L/2, \quad D \Pr_S(a_{\ell j} = a) \leq J_0 := 2p^{-L},$$

$$\|P_S/\nu^2 - g\|_{[2]} \leq \eta_{\text{mask}} = O_{L,p} \left(\frac{\sqrt{\log(2/\delta)}}{N} \right), \quad g(\beta) = \beta \lambda \left(\beta + \frac{\lambda}{pL} \right). \quad (18)$$

An explicit η_{mask} is given in equation 37.

Proof. 1. Raw moments and exact normalizations. Let $I_S(\chi)$ indicate validity and set

$$F_\ell = \mathbb{E}_{\text{full}}[I_S(\chi) y_{a_\ell}^2], \quad h_S(a) = \Pr_{\text{full}}(I_S(\chi) = 1 \mid a_1 = a).$$

Given χ , let $\chi^{(\ell)}$ start at $e_{\ell-1}$ with its first edge forced to be a_ℓ and its remaining $L-1$ relations fresh. Let r_0 be independent and uniform. Define

$$T_\ell = \mathbb{E}[I_S(\chi) I_S(\chi^{(\ell)}) y_{a_\ell}^2],$$

$$M_S = \mathbb{E}[I_S(\chi) \mathbf{1}\{(e_0, r_0) \in H\} y_{r_0, e_0} y_{r_1, e_0}]. \quad (19)$$

Then

$$\Pr_S(a_1 = a) = \frac{h_S(a)}{DF_0}, \quad \omega_a = \frac{h_S(a)}{LF_0},$$

$$K_{\ell\ell} = \frac{F_\ell}{F_0}, \quad U_{\ell\ell} = \frac{\mathbb{E}_{\text{full}}[I_S h_S(a_\ell) y_{a_\ell}^2]}{LF_0^2} = \frac{T_\ell}{LF_0^2},$$

$$M_S = \frac{1}{DN} \sum_e \sum_{r: (e, r) \in H} y_{r, e} \sum_t h_S(e, t) y_{t, e}, \quad Z_S = \frac{M_S}{qLF_0}. \quad (20)$$

2. Explicit expectation-bias bounds. For complete shifts,

$$(e_0, r_1, \dots, r_L) \longleftrightarrow (e_0, e_1, \dots, e_L), \quad r_i = 1 + ((e_i - e_{i-1}) \bmod N),$$

$$\Pr(e_0 = x_0, \dots, e_L = x_L) = N^{-(L+1)}. \quad (21)$$

The independent vertex samples underlying each configuration are

statistic	independent samples	distinct requests if no collisions
F_0, F_ℓ	$e_0, \dots, e_L$	L retained
T_ℓ	$e_0, \dots, e_L, z_2, \dots, z_L$	$(2L-1)$ retained
M_S	$e_0, \dots, e_L, z$	L retained, 1 held

The auxiliary path starts with $(e_{\ell-1}, e_\ell)$; the mixed test edge is (e_0, z) . For the event $\mathcal{C}$ that two independent samples coincide,

$$\Pr(\mathcal{C}) \leq \gamma_N := \frac{\binom{2L}{2}}{N}. \quad (22)$$

Assume $pD \geq 2L$. If $k \leq 2L$ distinct retained facts are requested,

$$p_k := \frac{(pD)_k}{(D)_k}, \quad \frac{p_k}{p^k} = \prod_{j=0}^{k-1} \frac{1 - j/(pD)}{1 - j/D} \leq 1,$$

$$\frac{p_k}{p^k} \geq \prod_{j=0}^{k-1} \left(1 - \frac{j}{pD}\right) \geq 1 - \sum_{j=0}^{k-1} \frac{j}{pD} = 1 - d_k, \quad d_k := \frac{k(k-1)}{2pD}. \quad (23)$$

For the collision-free mixed configuration,

$$\frac{\Pr(\text{retained path, held test edge})}{qp^L} = \frac{p_L/p^L}{1-L/D}, \quad \left| \frac{p_L/p^L}{1-L/D} - 1 \right| \leq \frac{d_L + L/D}{1-L/D} \leq 2d_L + 2L/D. \quad (24)$$

If R is a normalized inclusion probability and $|w| \leq 1$ a scalar weight, then

$$|\mathbb{E}[wR] - \mathbb{E}w| \leq \mathbb{E}[|R - 1| \mathbf{1}_{C^c}] + \mathbb{E}[|R - 1| \mathbf{1}_C]. \quad (25)$$

On $\mathcal{C}$, use respectively $0 \leq R \leq p^{-L}$, $p^{-(2L-1)}$, or $(qp^L)^{-1}$. Therefore, with expectation over masks,

statistic X	reference x_o	$ \mathbb{E}_{\text{mask}} X - x_o \leq$
F_0/p^L	1	$b_0 := (1 + p^{-L})\gamma_N + d_L$
F_ℓ/p^L	ν	b_0
T_ℓ/p^{2L-1}	ν	$b_T := (1 + p^{-(2L-1)})\gamma_N + d_{2L-1}$
$M_S/(qp^L)$	0	$b_M := (1 + (qp^L)^{-1})\gamma_N + 2d_L + 2L/D$
ν_H	ν	0

(26)

The reference means use $\mathbb{E}\zeta_{e_\ell}^2 = \nu$ and $\mathbb{E}[\zeta_z \zeta_{e_1}] = \mathbb{E}\zeta_z \mathbb{E}\zeta_{e_1} = 0$; z, e_1 are independent before masking.

3. Global sensitivity and concentration. If masks S, S' differ by one exchange, $|S \Delta S'| = 2$. A configuration with k edge requests can change only if one request hits a changed fact. Each request is marginally uniform, so

$$|\mathbb{E}[wI_S] - \mathbb{E}[wI_{S'}]| \leq \Pr(\exists i \leq k : a_i \in S \Delta S') \leq \sum_{i=1}^k \frac{2}{D} = \frac{2k}{D}. \quad (27)$$

The two forced copies of a diagonal-pair edge count as one request. Thus the global exchange constants are

$F_0/p^L, F_\ell/p^L$	$c_0 = 2L/(Dp^L)$
T_ℓ/p^{2L-1}	$c_T = 2(2L-1)/(Dp^{2L-1})$
$M_S/(qp^L)$	$c_M = 2(L+1)/(Dqp^L)$
ν_H	$c_H = 2/(qD)$

(28)

For a function X with exchange constant c , reveal the smaller subset in $k_* = \min(p, q)D$ steps. Its Doob increments Δ_i obey

$$\mathbb{E}[\Delta_i | \mathcal{F}_{i-1}] = 0, \quad \text{range}(\Delta_i | \mathcal{F}_{i-1}) \leq c. \quad (29)$$

The range bound follows by coupling two possible next reveals with one exchange in their uniform completions. For the conditional log moment function $\psi_i(s)$,

$$\psi_i(0) = \psi'_i(0) = 0, \quad \psi''_i(s) = \text{Var}_s(\Delta_i | \mathcal{F}_{i-1}) \leq c^2/4, \quad \psi_i(s) \leq s^2 c^2/8. \quad (30)$$

Iteration and exponential Markov yield

$$\begin{aligned} \mathbb{E}e^{s(X - \mathbb{E}X)} &\leq e^{k_* s^2 c^2/8}, \\ \Pr(|X - \mathbb{E}X| > x) &\leq 2 \inf_{s>0} e^{-sx + k_* s^2 c^2/8} = 2e^{-2x^2/(k_* c^2)}. \end{aligned} \quad (31)$$

There are $J_* = 2L + 1$ statistics in equation 26. Set

$$\begin{aligned} v_\delta &= \sqrt{\frac{k_*}{2} \log \frac{2J_*}{\delta}}, \\ \eta_0 &= b_0 + c_0 v_\delta, \quad \eta_T = b_T + c_T v_\delta, \quad \eta_M = b_M + c_M v_\delta, \quad \eta_H = c_H v_\delta. \end{aligned} \quad (32)$$

A union bound gives an event $\mathcal{E}$, $\Pr(\mathcal{E}) \geq 1 - \delta$, on which

$$\begin{aligned} |F_0/p^L - 1|, |F_\ell/p^L - \nu| &\leq \eta_0, \quad |T_\ell/p^{2L-1} - \nu| \leq \eta_T, \\ |M_S/(qp^L)| &\leq \eta_M, \quad |\nu_H - \nu| \leq \eta_H. \end{aligned} \quad (33)$$

4. Ratios and cubic coefficients. Assume $\eta_0 \leq 1/2$ and write $z = F_0/p^L$. On $\mathcal{E}$,

$$\begin{aligned} F_0 &\geq p^L(1 - \eta_0) \geq p^L/2, \\ |K_{\ell\ell} - \nu| &= \left| \frac{F_\ell/p^L - \nu z}{z} \right| \leq k_\Delta := \frac{(1 + \nu)\eta_0}{1 - \eta_0}, \\ \left| U_{\ell\ell} - \frac{\nu}{pL} \right| &= \frac{1}{pL} \left| \frac{T_\ell/p^{2L-1} - \nu z^2}{z^2} \right| \leq u_\Delta := \frac{\eta_T + \nu(2\eta_0 + \eta_0^2)}{pL(1 - \eta_0)^2}, \\ |Z_S| &\leq z_\Delta := \frac{\eta_M}{L(1 - \eta_0)}, \quad 0 \leq \omega_a \leq \omega_+ := \frac{1}{Lp^L(1 - \eta_0)}, \\ |K_{j\ell}| &\leq 1, \quad |U_{j\ell}| \leq \omega_+. \end{aligned} \tag{34}$$

Introduce scalar coefficients, not parameter matrices:

$$A_S = \frac{\nu_H}{L-1} \sum_{\ell=2}^L K_{\ell\ell}, \quad B_S = \frac{\nu_H}{L-1} \sum_{\ell=2}^L U_{\ell\ell}, \quad K_O = \frac{1}{L-1} \sum_{\ell=2}^L \sum_{j \neq \ell} K_{j\ell}, \quad U_O = \frac{1}{L-1} \sum_{\ell=2}^L \sum_{j \neq \ell} U_{j\ell}.$$

The product inequality $|xy - x_0y_0| \leq |x_0||y - y_0| + |y_0||x - x_0| + |x - x_0||y - y_0|$ gives

$$\begin{aligned} |A_S - \nu^2| &\leq a_\Delta := \nu k_\Delta + \nu \eta_H + k_\Delta \eta_H, \\ |B_S - \nu^2/(pL)| &\leq b_\Delta := \nu u_\Delta + \frac{\nu}{pL} \eta_H + u_\Delta \eta_H, \\ |K_O| &\leq L-1, \quad |U_O| \leq (L-1)\omega_+. \end{aligned} \tag{35}$$

Now

$$P_S - \nu^2 g = (A_S - \nu^2)\beta^2 \lambda + (B_S - \nu^2/(pL))\beta \lambda^2 + Z_S(K_O \beta \lambda^2 + U_O \lambda^3). \tag{36}$$

Direct differentiation supplies all three norm constants:

$\frac{h}{\beta^2 \lambda, \beta \lambda^2, \lambda^3}$	$\ h\ _\infty$	$\ h'\ _\infty$	$\frac{1}{2}\ h''\ _\infty$	$\ h\ _{[2]}$
	4/27	1	2	85/27
	1	3	3	7

Hence an explicit mask error is

$$\eta_{\text{mask}} := \frac{1}{\nu^2} \left[\frac{85}{27}(a_\Delta + b_\Delta) + (L-1)z_\Delta \left(\frac{85}{27} + 7\omega_+ \right) \right]. \tag{37}$$

For fixed L, p , all four errors in equation 32 are $O_{L,p}(\sqrt{\log(2/\delta)}/N)$; equations equation 34–equation 37 give the same rate. Finally,

$$D \Pr_S(a_{\ell j} = a) = D \frac{\Pr_{\text{full}}(a_{\ell j} = a, I_S = 1)}{F_0} \leq \frac{D \Pr_{\text{full}}(a_{\ell j} = a)}{F_0} = F_0^{-1} \leq 2p^{-L}.$$

□

Lemma 3 (Uniform finite-time cubic approximation). *Fix a mask with $F_0 \geq p^L/2$ and a horizon $\tau > 0$. For sufficiently large N , the joint raw flow satisfies*

$$-\frac{2LD^2}{\nu^2 \tau^3} \left[\mathcal{E}_H(\theta_\beta(\tau)) - \frac{\nu_H}{2} \right] = \frac{P_S(\beta)}{\nu^2} + r_{\text{time}}(\beta), \quad \|r_{\text{time}}\|_{[2]} \leq \eta_{\text{time}} = O_{L,p,\tau}(D^{-1}). \tag{38}$$

An explicit bound is equation 59.

Proof. 1. Exact raw gradients. For a chain function h , define the weighted linear functional

$$\mathcal{T}_{\beta,\ell}[h] := \beta \mathbf{1}\{\ell = 1\} \mathbb{E}_{\text{at}} h + \frac{\lambda}{L} \mathbb{E}_S h.$$

Write $G_{r,e} = w_r b_{r,e}$ and $R_\ell = f_\ell - y_{a_\ell}$. Then

$$\begin{aligned} Q_a &= \sum_{\ell,j} C_{j\ell} \mathcal{T}_{\beta,\ell}[\mathbf{1}\{a_{\ell j} = a\} R_\ell], \\ \dot{b}_{r,e} &= -w_r Q_{r,e}, \quad \dot{w}_r = -\sum_e b_{r,e} Q_{r,e}, \\ \dot{C}_{j\ell} &= -\mathcal{T}_{\beta,\ell}[G_{a_{\ell j}} R_\ell]. \end{aligned} \tag{39}$$

For $\langle h, k \rangle_\beta := \sum_\ell \mathcal{T}_{\beta, \ell}[h_\ell k_\ell]$,

$$\|y\|_\beta^2 = \beta \mathbb{E}_{\text{at}} y_{a_1}^2 + \frac{\lambda}{L} \sum_\ell \mathbb{E}_S y_{a_\ell}^2 \leq 1,$$

$$\frac{d}{dt}(\|w\|_2^2 + \|b\|_F^2 + \|C\|_F^2) = -6\langle f, f - y \rangle_\beta = -6\|f - y/2\|_\beta^2 + \frac{3}{2}\|y\|_\beta^2 \leq \frac{3}{2}. \quad (40)$$

The polynomial ODE thus has a global real solution, with smooth β -dependence on finite intervals. Here $y_\ell = y_{a_\ell}$.

2. Rescaled equations and their operator bounds. Use $Y = Db$ only as an auxiliary coordinate. Define

$$\begin{aligned} [\mathcal{A}_\beta(C)]_a &:= D \sum_{\ell, j} C_{j\ell} \mathcal{T}_{\beta, \ell}[\mathbf{1}\{a_{\ell j} = a\} y_{a_\ell}], \\ [\mathcal{B}_\beta(H)]_{j\ell} &:= \mathcal{T}_{\beta, \ell}[H_{a_{\ell j}} y_{a_\ell}], \\ [\mathcal{U}_\beta(C_1, C_2, H)]_a &:= D \sum_{\ell, j} (C_1)_{j\ell} \mathcal{T}_{\beta, \ell} \left[\mathbf{1}\{a_{\ell j} = a\} \sum_k (C_2)_{k\ell} H_{a_{\ell k}} \right], \\ [\mathcal{V}_\beta(H_1, C, H_2)]_{j\ell} &:= \mathcal{T}_{\beta, \ell} \left[(H_1)_{a_{\ell j}} \sum_k C_{k\ell} (H_2)_{a_{\ell k}} \right]. \end{aligned} \quad (41)$$

The arguments H, H_1, H_2 in these maps are scalar tables, not held sets. Let $J_0 = 2p^{-L}$; subscripts 1, 0 below denote $\beta = 1, 0$. The candidate probabilities give the base operator bounds

$$\begin{aligned} \|\mathcal{A}_1(C)\|_{\max} &\leq D \frac{1}{D} \sum_j |C_{j1}| \leq \|C\|_1, & \|\mathcal{A}_0(C)\|_{\max} &\leq \frac{J_0}{L} \|C\|_1, \\ \|\mathcal{B}_1(H)\|_1 &\leq L \|H\|_{\max}, & \|\mathcal{B}_0(H)\|_1 &\leq \frac{L^2}{L} \|H\|_{\max}, \\ \|\mathcal{U}_1(C_1, C_2, H)\|_{\max} &\leq \|C_1\|_1 \|C_2\|_1 \|H\|_{\max}, & \|\mathcal{U}_0(C_1, C_2, H)\|_{\max} &\leq \frac{J_0}{L} \|C_1\|_1 \|C_2\|_1 \|H\|_{\max}, \\ \|\mathcal{V}_1(H_1, C, H_2)\|_1 &\leq L \|H_1\|_{\max} \|C\|_1 \|H_2\|_{\max}, & \|\mathcal{V}_0(H_1, C, H_2)\|_1 &\leq \|H_1\|_{\max} \|C\|_1 \|H_2\|_{\max}. \end{aligned} \quad (42)$$

For example, the final coefficient is $L^{-1} \sum_{j, \ell, k} |C_{k\ell}| = \|C\|_1$. Apply equation 11 to the affine mixture of these operators:

$$\begin{aligned} \|\mathcal{A}_\beta(C)\|_{[2]} &\leq 2(1 + J_0/L) \|C\|_{[2]}, \\ \|\mathcal{U}_\beta(C_1, C_2, H)\|_{[2]} &\leq 2(1 + J_0/L) \|C_1\|_{[2]} \|C_2\|_{[2]} \|H\|_{[2]}, \\ \|\mathcal{B}_\beta(H)\|_{[2]} &\leq (2L + 2L) \|H\|_{[2]}, \\ \|\mathcal{V}_\beta(H_1, C, H_2)\|_{[2]} &\leq (2L + 2) \|H_1\|_{[2]} \|C\|_{[2]} \|H_2\|_{[2]}. \end{aligned} \quad (43)$$

Set

$$K := 4(L + 1 + J_0) \geq \max\{2(1 + J_0/L), 4L, 2L + 2\}. \quad (44)$$

Row scaling is $(wY)_{r,e} = w_r Y_{r,e}$. Substitution into the raw flow gives

$$\begin{aligned} \dot{Y} &= w \mathcal{A}_\beta(C) - D^{-1} w \mathcal{U}_\beta(C, C, wY), \\ \dot{w}_r &= D^{-2} \sum_e Y_{r,e} [\mathcal{A}_\beta(C)]_{r,e} - D^{-3} \sum_e Y_{r,e} [\mathcal{U}_\beta(C, C, wY)]_{r,e}, \\ \dot{C} &= D^{-1} \mathcal{B}_\beta(wY) - D^{-2} \mathcal{V}_\beta(wY, C, wY). \end{aligned} \quad (45)$$

Thus, for $s_Y = \|Y\|_{[2]}$, $s_w = \|w\|_{[2]}$, $s_C = \|C\|_{[2]}$,

$$\begin{aligned} \|\dot{Y}\|_{[2]} &\leq K s_w s_C (1 + s_w s_C s_Y / D), \\ \|\dot{C}\|_{[2]} &\leq (K/D) s_w s_Y (1 + s_w s_C s_Y / D), \\ \|\dot{w}\|_{[2]} &\leq (KN/D^2) s_Y s_C (1 + s_w s_C s_Y / D). \end{aligned} \quad (46)$$

3. Close the bootstrap with a stated dimension threshold. Let $R_Y = 1 + 8K\tau$ and assume

$$D \geq \max\{4R_Y, 8KR_Y\tau\}. \quad (47)$$

Up to the first exit from $s_w, s_C \leq 2, s_Y \leq R_Y$,

$$\begin{aligned} s_w s_C s_Y / D &\leq 4R_Y / D \leq 1, & \|\dot{Y}\|_{[2]} &\leq 8K, \\ \|\dot{C}\|_{[2]} &\leq 4KR_Y / D, & \|\dot{w}\|_{[2]} &\leq 4KNR_Y / D^2 \leq 4KR_Y / D, \\ s_Y(t) &\leq 8Kt < R_Y, & s_C(t), s_w(t) &\leq 1 + 4KR_Y\tau / D \leq \frac{3}{2} < 2. \end{aligned} \quad (48)$$

No exit occurs on $[0, \tau]$. Substituting $s_Y(s) \leq 8Ks$ back into equation 46 sharpens the other two estimates:

$$\begin{aligned} \|Y(t)\|_{[2]} &\leq 8Kt, \\ \|C(t) - E_{11}\|_{[2]} &\leq \frac{4K}{D} \int_0^t 8Ks \, ds = \frac{16K^2 t^2}{D}, \\ \|w(t) - \mathbf{1}\|_{[2]} &\leq \frac{4KN}{D^2} \int_0^t 8Ks \, ds = \frac{16K^2 N t^2}{D^2}. \end{aligned} \quad (49)$$

4. Derive the first two expansions by integration. Let

$$v_a := [\beta + \lambda\omega_a]y_a = [\mathcal{A}_\beta(E_{11})]_a, \quad \|v\|_{[2]} \leq K. \quad (50)$$

Using $N \leq D$, equations equation 45 and equation 49 imply

$$\begin{aligned} \|\dot{Y} - v\|_{[2]} &\leq \|w - \mathbf{1}\|_{[2]} K s_C + K \|C - E_{11}\|_{[2]} + \frac{K}{D} s_w^2 s_C^2 s_Y \\ &\leq \frac{32K^3 N t^2}{D^2} + \frac{16K^3 t^2}{D} + \frac{128K^2 t}{D} \leq \frac{128K^2 t + 48K^3 t^2}{D}, \\ \|Y(t) - tv\|_{[2]} &\leq \frac{64K^2 t^2 + 16K^3 t^3}{D}. \end{aligned} \quad (51)$$

Consequently,

$$\begin{aligned} \|wY - tv\|_{[2]} &\leq \|Y - tv\|_{[2]} + \|w - \mathbf{1}\|_{[2]} \|Y\|_{[2]} \\ &\leq \frac{64K^2 t^2 + 16K^3 t^3}{D} + \frac{128K^3 N t^3}{D^2} \leq \frac{E_G(t)}{D}, \quad E_G(t) := 64K^2 t^2 + 144K^3 t^3. \end{aligned} \quad (52)$$

For routing, the exact equation gives

$$\begin{aligned} \left\| \dot{C} - \frac{t}{D} \mathcal{B}_\beta(v) \right\|_{[2]} &\leq \frac{K}{D} \|wY - tv\|_{[2]} + \frac{K}{D^2} s_w^2 s_Y^2 s_C \\ &\leq \frac{KE_G(t) + 512K^3 t^2}{D^2}, \\ \left\| C - E_{11} - \frac{t^2}{2D} \mathcal{B}_\beta(v) \right\|_{[2]} &\leq \frac{1}{D^2} \int_0^t [KE_G(s) + 512K^3 s^2] \, ds = \frac{E_C(t)}{D^2}, \\ E_C(t) &:= 192K^3 t^3 + 36K^4 t^4. \end{aligned} \quad (53)$$

In particular, the last two bounds are

$$\left\| G - \frac{t}{D} v \right\|_{[2]} \leq \frac{E_G(t)}{D^2}, \quad \left\| C - E_{11} - \frac{t^2}{2D} \mathcal{B}_\beta(v) \right\|_{[2]} \leq \frac{192K^3 t^3 + 36K^4 t^4}{D^2}. \quad (54)$$

5. Multiply the expansions and project onto the target. Let $C^{(2)} = t^2 \mathcal{B}_\beta(v) / (2D)$ and $G^{(1)} = tv / D$. For $\ell \geq 2$, $C_{:\ell}(0) = 0$ and

$$[\mathcal{B}_\beta(v)]_{j\ell} = \frac{\lambda}{L} [\beta K_{j\ell} + \lambda U_{j\ell}].$$

The cubic prediction is therefore

$$f_\ell^{(3)}(t) := \sum_j C_{j\ell}^{(2)} G_{a_{\ell j}}^{(1)} = \frac{t^3 \lambda}{2LD^2} \sum_j (\beta K_{j\ell} + \lambda U_{j\ell}) v_{a_{\ell j}}. \quad (55)$$

The product remainder satisfies

$$\begin{aligned} \|f_\ell - f_\ell^{(3)}\|_{[2]} &\leq \|C_{:\ell} - C_{:\ell}^{(2)}\|_{[2]} \|G\|_{[2]} + \|C_{:\ell}^{(2)}\|_{[2]} \|G - G^{(1)}\|_{[2]} \\ &\leq \frac{E_C(t)}{D^2} \frac{16Kt}{D} + \frac{K^2 t^2}{2D} \frac{E_G(t)}{D^2} = \frac{E_f(t)}{D^3}, \\ E_f(t) &:= 16KtE_C(t) + \frac{1}{2}K^2 t^2 E_G(t) = 3104K^4 t^4 + 648K^5 t^5, \\ \|f_\ell(t)\|_{[2]} &\leq \frac{16K^2 t^2}{D} \frac{16Kt}{D} = \frac{F(t)}{D^2}, \quad F(t) := 256K^3 t^3. \end{aligned} \quad (56)$$

Lemma 1 supplies the missing target-projection step:

$$\begin{aligned} \mathbb{E}_{H,\ell}[v_{a_{\ell j}} y_{a_\ell}] &= \beta \nu_H \mathbf{1}\{j = \ell\} + \lambda Z_S \mathbf{1}\{j \neq \ell\}, \\ \frac{1}{L-1} \sum_{\ell=2}^L \mathbb{E}_{H,\ell}[y_{a_\ell} f_\ell^{(3)}] &= \frac{t^3 \lambda}{2LD^2(L-1)} \sum_{\ell=2}^L \sum_j (\beta K_{j\ell} + \lambda U_{j\ell}) (\beta \nu_H \mathbf{1}\{j = \ell\} + \lambda Z_S \mathbf{1}\{j \neq \ell\}) \\ &= \frac{t^3}{2LD^2} P_S(\beta). \end{aligned} \quad (57)$$

Finally,

$$\begin{aligned} \mathcal{E}_H - \frac{\nu_H}{2} + \frac{t^3 P_S}{2LD^2} &= \frac{1}{L-1} \sum_{\ell=2}^L \mathbb{E}_{H,\ell}[-y_{a_\ell} (f_\ell - f_\ell^{(3)}) + f_\ell^2/2], \\ \left\| \mathcal{E}_H - \frac{\nu_H}{2} + \frac{t^3 P_S}{2LD^2} \right\|_{[2]} &\leq \frac{E_f(t)}{D^3} + \frac{F(t)^2}{2D^4}. \end{aligned} \quad (58)$$

Thus an explicit time error is

$$\begin{aligned} \eta_{\text{time}} &:= \frac{2L}{\nu^2 \tau^3} \left[\frac{E_f(\tau)}{D} + \frac{F(\tau)^2}{2D^2} \right] \\ &= \frac{2L}{\nu^2} \left[\frac{3104K^4 \tau + 648K^5 \tau^2}{D} + \frac{32768K^6 \tau^3}{D^2} \right] = O_{L,p,\tau}(D^{-1}). \end{aligned} \quad (59)$$

All bounds include two derivatives in β by equation 11. $\square$

Lemma 4 (Explicit center and stable ordering). *Set $a = (pL)^{-1}$, $g(\beta) = \beta(1-\beta)[\beta + a(1-\beta)]$, $\mu = 2 \min\{a, 1\}$, and $m = \min\{a, 1\}/3$. If $J = g + r$ and $\|r\|_{[2]} \leq \eta < m$, then J has a unique maximizer $\beta^* \in (1/3, 2/3)$ and*

$$|\beta^* - \beta(pL)| \leq \frac{\eta}{\mu - 2\eta}, \quad \beta'(z) > 0 \quad (z > 0). \quad (60)$$

Proof. Direct differentiation gives

$$\begin{aligned} g' &= a + 2(1-2a)\beta + 3(a-1)\beta^2, & g'' &= 2(1-2a) + 6(a-1)\beta, \\ g'(0) &= a, & g'(1/3) &= 1/3, & g'(2/3) &= -a/3, & g'(1) &= -1, \\ -g''(1/3) &= 2a, & -g''(2/3) &= 2, & \inf_{[1/3, 2/3]} (-g'') &= \mu. \end{aligned} \quad (61)$$

A quadratic with opposite endpoint signs has exactly one root in $(0, 1)$. Writing $b \in [0, 1]$ for a candidate mixture and $x = b/(1-b)$,

$$g'(b) = 0 \iff -x^2 + 2(1-a)x + a = 0 \iff x = 1-a + \sqrt{1-a+a^2} \iff b = \beta(pL) \in (1/3, 2/3). \quad (62)$$

For the exterior slopes,

$$\begin{aligned} \inf_{[0,1/3]} g' &\geq \begin{cases} \min\{a, 1/3\}, & a \leq 1, \\ 1/3, & a \geq 1, \end{cases} \geq m, \\ \inf_{[2/3,1]} (-g') &\geq \begin{cases} a/3, & a \leq 1, \\ \min\{a/3, 1\}, & a \geq 1, \end{cases} \geq m. \end{aligned} \quad (63)$$

These follow from the sign of $g''' = 6(a-1)$ and the endpoint values in equation 61. Hence

$$\begin{aligned} J' &\geq m - \eta > 0 \quad \text{on } [0, 1/3], \quad J' \leq -m + \eta < 0 \quad \text{on } [2/3, 1], \\ -J'' &\geq \mu - 2\eta > 0 \quad \text{on } [1/3, 2/3]. \end{aligned} \quad (64)$$

Existence follows by compactness; these inequalities give interiority and uniqueness. Writing $\theta = \beta(pL)$, we have

$$(\mu - 2\eta)|\beta^* - \theta| \leq |J'(\beta^*) - J'(\theta)| = |r'(\theta)| \leq \eta. \quad (65)$$

For $\Delta = z^2 - z + 1$,

$$\beta'(z) = \beta(z)^2 \left(1 - \frac{2z-1}{2\sqrt{\Delta}}\right) > 0, \quad 4\Delta - (2z-1)^2 = 3. \quad (66)$$

□

Proof of Theorem 1. On the event of Lemma 2, let

$$\begin{aligned} J_\tau(\beta) &:= -\frac{2LD^2}{\nu^2\tau^3} [\mathcal{E}_H(\theta_\beta(\tau)) - \nu_H/2] = g(\beta) + r(\beta), \\ \|r\|_{[2]} &\leq \eta_N := \eta_{\text{mask}} + \eta_{\text{time}} = O_{L,p,\tau} \left(\frac{\sqrt{\log(2/\delta)}}{N} + \frac{1}{N^2} \right). \end{aligned} \quad (67)$$

The explicit sufficient conditions are

$$pN^2 \geq 2L, \quad \eta_0 \leq \frac{1}{2}, \quad N^2 \geq \max\{4R_Y, 8KR_Y\tau\}, \quad \eta_N < m. \quad (68)$$

They hold for sufficiently large N with L, p, τ, δ fixed. Lemma 4 gives

$$\mathfrak{R}_N := \frac{\eta_N}{\mu - 2\eta_N} = O_{L,p,\tau} \left(\frac{\sqrt{\log(2/\delta)}}{N} \right), \quad |\beta_{\tau,S}^* - \beta(pL)| \leq \mathfrak{R}_N. \quad (69)$$

Maximizing J_τ is equivalent to minimizing $\mathcal{E}_H$. Moreover,

$$\partial_L \beta(pL) = p\beta'(pL) > 0, \quad \partial_q \beta(L(1-q)) = -L\beta'(pL) < 0. \quad (70)$$

□