

CRITICAL ATTENTION SCALING IN LONG-CONTEXT TRANSFORMERS

Anonymous authors

Paper under double-blind review

ABSTRACT

As large language models scale to longer contexts, attention layers suffer from a fundamental pathology: attention scores collapse toward uniformity as context length n increases, causing tokens to cluster excessively, a phenomenon known as rank-collapse. While *attention scaling* effectively addresses this deficiency by rescaling attention scores with a polylogarithmic factor β_n , theoretical justification for this approach remains lacking.

We analyze a simplified yet tractable model that magnifies the effect of attention scaling. In this model, attention exhibits a phase transition governed by the scaling factor β_n : insufficient scaling collapses all tokens to a single direction, while excessive scaling reduces attention to identity, thereby eliminating meaningful interactions between tokens. Our main result identifies the critical scaling $\beta_n \asymp \log n$ and provides a rigorous justification for attention scaling in YaRN and Qwen, clarifying why logarithmic scaling maintains sparse, content-adaptive attention at large context lengths.

1 INTRODUCTION

The attention mechanism is a cornerstone of modern transformer architectures on which Large Language Models (LLMs) rely. Mathematically, an attention layer is a nonlinear operator ATT that maps a collection of *tokens* $\{x_1, \dots, x_n\}$ from \mathbb{R}^d to \mathbb{R}^d . This operator is parametrized by three (possibly sparse) d by d matrices K, Q , and V and maps $\{x_1, \dots, x_n\}$ to $\{x'_1, \dots, x'_n\}$ using the following formula. Define the normalization operator $N(x) = x/\|x\|$ and for any $i = 1, \dots, n$ define $q_i = QN(x_i)$, $k_i = KN(x_i)$. Then $x'_i = \text{ATT}(x_1, \dots, x_n)_i$ is defined as

$$x'_i = V \sum_{j=1}^n N(x_j) A_{ij}, \quad A_{ij} = \frac{e^{a_{ij}}}{\sum_{k=1}^n e^{a_{ik}}}, \quad (1)$$

where the terms $a_{ij} = q_i^\top k_j$ are called *attention scores*.

A recent line of theoretical work has demonstrated that attention acts as a *contractive* operator that tends to cluster tokens together; see Dong et al. (2021); Geshkovski et al. (2024b; 2025); Karagodin et al. (2024); Geshkovski et al. (2024a); Bruno et al. (2025a); Polyanskiy et al. (2025); Chen et al. (2025a); Cowsik et al. (2024); Giorlandino & Goldt (2025); Arroyo et al. (2025). This clustering effect is also known as “rank-collapse” or “token uniformity” and arises because the distribution of attention scores tends to flatten as the sequence length n grows, causing each token to disperse its attention across too many other tokens rather than focusing selectively.

Various practical solutions have been proposed to curb this clustering behavior. In this work, we focus on simple context-length-aware modifications of the attention mechanism following ideas practically implemented as YaRN (Peng et al., 2023), Qwen (Bai et al., 2023), SSMaX (Nakanishi, 2025), and SWAN-GPT (Puvvada et al., 2025). These methods employ a straightforward strategy that rescales attention scores a_{ij} by a single poly-logarithmic factor β_n ; see Table 1. Our goal in this paper is to answer the following fundamental question:

What is the optimal order of magnitude of the β_n scaling?

To address this question, we propose a highly simplified yet completely tractable model for attention. This model exhibits a phase transition governed by the parameter β_n as $n \rightarrow \infty$: when β_n is below a critical threshold, attention becomes overly contractive and collapses all tokens to a single direction, while when β_n is too large, attention acts as an identity operator and fails to process information effectively. More precisely, we establish that the critical parameter β_n scales as $\log n$, which corroborates the empirical guidelines underlying YaRN, Qwen, SSMax, and SWAN-GPT.

Our work is intimately connected to the recent contributions of Giorlandino & Goldt (2025) and Cowsik et al. (2024), who investigate the contractive effects of attention mechanisms with random key and query matrices K and Q to establish proper initialization schemes for these parameters. A crucial insight from Cowsik et al. (2024) is that analyzing the evolution of symmetric token configurations provides a more mathematically tractable framework compared to the generic input distributions considered in Geshkovski et al. (2025). This symmetric setting, while simplified, captures essential dynamics of the attention mechanism and enables rigorous theoretical analysis; see also Karagodin et al. (2025).

The choice $\beta_n = \gamma \log n$ appears natural in retrospect. As noted in Nakanishi (2025), with such a scaling the attention weights A_{ij} in Equation (1) become

$$A_{ij} = \frac{n^{\gamma a_{ij}}}{\sum_{k=1}^n n^{\gamma a_{ik}}}.$$

To illustrate the resulting dynamics, consider a simplified regime where all attention scores a_{ij} are of order one: specifically, let $a_{ii} = 1$ and $a_{ij} = \rho > 0$ for $i \neq j$. In this setting, the off-diagonal weights satisfy

$$A_{ij} = \frac{n^{\gamma \rho}}{n^{\gamma} + (n-1)n^{\gamma \rho}} \sim \begin{cases} 1/n & \text{if } \gamma < \frac{1}{1-\rho} \\ 1/n^{\gamma(1-\rho)} & \text{if } \gamma > \frac{1}{1-\rho} \end{cases}$$

This analysis reveals two distinct regimes. When γ is small (subcritical regime), attention weights are asymptotically uniform, resulting in diffuse attention that, as we demonstrate below, leads to severe token contraction. Conversely, when γ is large (supercritical regime), off-diagonal weights become negligible with respect to the diagonal ones so that the attention mechanism is effectively suppressed.

The critical regime emerges at the phase boundary $\gamma = \frac{1}{1-\rho}$ where attention can concentrate on a sublinear yet nontrivial number of tokens so as to maintain sufficient connections to facilitate information flow from a small set of important tokens. This sparse attention is related to structured attention mechanisms employed in long-context architectures such as Longformer (Beltagy et al., 2020) and SWIN (Liu et al., 2021) which implement a sliding window over $k \ll n$ -nearest neighbors but where proximity is measured in terms of token position rather than embedding. Unlike these structurally constrained approaches that rely on fixed positional neighborhoods, the logarithmic scaling enables the attention pattern to be entirely *content-adaptive*, allowing each token to dynamically select its most relevant context based on semantic similarity rather than positional proximity.

Following similar motivations, Giorlandino & Goldt (2025) establish a compelling analogy between attention dynamics and the random energy model from statistical physics (Derrida, 1981). Using the replica method—an analytical heuristic from statistical physics—they identify a phase transition occurring at $\beta_n \sim \sqrt{\log n}$, which differs from the scalings presented in Table 1. This result represents a significant discrepancy from our findings and highlights fundamental differences in modeling assumptions. More specifically, their approach assumes that the attention scores a_{ij} are correlated Gaussian random variables. This assumption effectively induces a random geometry on the token space, where similarity between tokens is treated as fundamentally random. In this sense, their model bears closer resemblance to recent Kuramoto models on random graphs studied in Abdalla et al. (2022); Jain et al. (2025), where the authors investigate the synchronization of oscillators interacting across the edges of a (sparse) Erdős–Rényi random graph with unit edge weights. However, in the case of Giorlandino & Goldt (2025), the random graph is both directed and dense, with the

Method	β_n scaling
YaRN	$(\log n)^2$
Qwen	$\log n$
SSMax	$\log n$
SWAN-GPT	$\log n$

Table 1: Attention scaling factors for various methods. The standard attention score $\exp(k_i^\top q_j)$ is replaced with $\exp(C\beta_n k_i^\top q_j)$, $C > 0$.

edge pointing from token j to token i having weight given by

$$A_{ij} = \frac{e^{\beta_n a_{ij}}}{\sum_{k=1}^n e^{\beta_n a_{ik}}} \quad (2)$$

where a_{ij} are Gaussian random variables. While Giorlandino & Goldt (2025) assumes a specific correlation structure between the Gaussian random variables, the phase transition they uncover is expected to be universal within a large class of random matrices including Wigner ones. Crucially though, in such models, the interaction strength A_{ij} is independent of the positional relationship between tokens i and j , making this model qualitatively different from standard attention mechanisms where attention is focused on few (or all) of the preceding tokens. [For completeness, we refer readers to Appendix D for a derivation of the critical scaling \$\beta_n \asymp \sqrt{\log n}\$ in the i.i.d. Gaussian score model.](#)

Bruno et al. (2025b) adopt a different approach to studying the regime where $n \rightarrow \infty$ and $\beta_n \rightarrow \infty$, in a more general setting than ours. By considering various levels of generality for the matrices K, Q, V , this work identifies distinct regimes of token dynamics and relates them to the hardmax ($\beta = \infty$) limit. Importantly, the analysis is conducted in the subcritical regime and differs from the present work in focusing on a broader class of models, for which the critical regime has yet to be precisely characterized. We believe that combining the analytical tools developed in both papers could yield a deeper understanding of this critical regime and represents a promising direction for future research.

The remainder of the paper is organized as follows. Section 2 provides a precise mathematical formulation of the phase transition phenomena for the rescaled attention layer. We begin by analyzing token angles and the contractive behavior of tokens under two settings: an idealized but intuitive simplex model (Section 2.1) and a more realistic model with the simplex constraint relaxed (Section 2.2). In both cases, we identify three distinct regimes of the scaling parameter, each leading to qualitatively different contrastive behaviors of the self-attention layer. Section 2.3 turns to the gradient norm of the rescaled attention operator. Because rank collapse is often accompanied by vanishing gradients, we characterize the gradient dynamics across scaling regimes and show when gradients vanish, or stabilize to non-trivial limits. Section 3 presents our numerical experiments, which validate these theoretical predictions.

Throughout this paper, when we denote a quantity as $o_n(1)$, where n is the number of tokens, we mean there are positive constants C_1, C_2 independent of the dimension d , such that $|o_n(1)| \leq C_1 n^{-C_2}$. The constants C_1, C_2 depend on the assumptions in theorems.

2 A PHASE TRANSITION FOR ATTENTION

In this section, we establish the main theorem of this paper, namely a phase transition for the contractive properties of the attention layer when $\beta_n = \gamma \log n$ for some $\gamma > 0$.

Following Geshkovski et al. (2025), we study a simplified version of the attention layer with pre-layer norm that is described in the introduction by assuming that $K = Q = V = I_d$. More specifically, the model we study is given as follows.

For any two points $x, y \in \mathbb{R}^d$, let $\langle x, y \rangle = x^\top y$ denote the standard Euclidean inner product in \mathbb{R}^d , and $\|x\| = \sqrt{\langle x, x \rangle}$. Finally, recall that $N(x) := x/\|x\|$.

For any collection of tokens $\{x_1, \dots, x_n\}$ in \mathbb{R}^d , define $y_i = N(x_i) \in \mathbb{S}^{d-1}$ for $i = 1, \dots, n$ and

$$Z_i := \sum_{k=1}^n e^{a_{ik}}, \quad A_{ij} := \frac{e^{a_{ij}}}{Z_i}, \quad a_{ij} := \beta \langle y_i, y_j \rangle, \quad (3)$$

for $i, j = 1, \dots, n$. We then define

$$\text{ATT}(y_i) := \sum_{j=1}^n A_{ij} y_j. \quad (4)$$

Since the seminal work of He et al. (2016), residual connections are added to modern architectures and naturally act as a regularization scheme of the attention map towards the identity; see Chen et al.

(2025b). With said residual connections, each token x_i is mapped to x'_i using the following update rule

$$x'_i := \text{ATT}(y_i) + \alpha x_i, \quad \alpha \geq 0. \quad (5)$$

Our first goal is to understand where the angle $\angle(x'_i, x'_j)$ compares to $\angle(x_i, x_j)$. If $\angle(x'_i, x'_j) < \angle(x_i, x_j)$ —or equivalently $\langle y'_i, y'_j \rangle > \langle y_i, y_j \rangle$, with $y'_i = N(x'_i)$ —we say that attention is *contractive*.

The nonlinear update rule (5) can produce complex dynamics, in which some pairs of tokens move closer together while others drift apart. This diversity of motion is in fact the most desirable outcome in practice, and it emerges precisely at the phase transition identified in this study. Beyond this critical regime, the tokens exhibit an unexpectedly cohesive behavior. To delineate the boundaries of the critical regime, we assume that the size and relative positions of the initial tokens are governed by constants independent of the number n of tokens. As an analytically tractable extreme of this assumption, we first consider the case in which the tokens form a regular simplex in \mathbb{R}^d as in Cowsik et al. (2024). Despite its symmetry, this configuration is sufficient to capture and predict the onset of the phase transition. We subsequently relax this constraint in Section 2.2 to show that the same phase transition occurs in more realistic configurations.

2.1 THE SIMPLEX CASE

The following assumption was made in Cowsik et al. (2024) and subsequently in Giorlandino & Goldt (2025). While rather stringent—in particular, it requires $d \geq n$ —it turns out to provide a tractable yet predictive setup to study the contractive properties of attention.

Assumption 1 *There exists nonnegative constants $q \geq 0$ and $\rho \in (0, 1)$ such that $\|x_i\|^2 = q$ and $\langle y_i, y_j \rangle = \rho$, for any $i, j = 1, \dots, n$ and $i \neq j$.*

Under Assumption 1, it is easy to see that there are positive constants ρ' and q' such that $\langle y'_i, y'_j \rangle = \rho'$ for all $i \neq j$ and $\|x'_i\|^2 = q'$ for all i . This simplification gives rise to a tractable phase transition.

Theorem 2.1 *Under Assumption 1, there is a $\rho' \in (0, 1)$ such that $\langle y'_i, y'_j \rangle = \rho'$ for all $i \neq j$. Moreover, if $\beta = \gamma \log n$ where γ is a positive constant, then for any $i \neq j$, it holds*

$$\lim_{n \rightarrow +\infty} \langle y'_i, y'_j \rangle = \begin{cases} \frac{\rho(\alpha\sqrt{q}+1)^2}{\alpha^2 q + 2\alpha\sqrt{q}\rho + \rho} & \text{if } \gamma < \frac{1}{1-\rho}, \\ \frac{\rho(\alpha\sqrt{q}+1)^2}{\alpha^2 q + \alpha\sqrt{q}(1+\rho) + \frac{1+3\rho}{4}} & \text{if } \gamma = \frac{1}{1-\rho}, \\ \rho & \text{if } \gamma > \frac{1}{1-\rho}. \end{cases} \quad (6)$$

Note that when $\gamma \leq \frac{1}{1-\rho}$, the right hand sides of Equation (6) are strictly larger than ρ for any $\alpha \geq 0$. In other words, in the critical and subcritical regimes attention is contractive even in the presence of a residual connection. Of course, when $\alpha \rightarrow \infty$, the effects of attention dissipates and the limit tends to ρ for all phases. This is expected as the update from y_i to y'_i tends to the identity map, an effect known to mitigate oversmoothing” in residual neural networks; see Chen et al. (2025b).

Note also that for $\alpha = 0$, that is in absence of residual connections, the limit in Equation (6) reduces to

$$\lim_{n \rightarrow +\infty} \langle y'_i, y'_j \rangle = \begin{cases} 1 & \text{if } \gamma < \frac{1}{1-\rho}, \\ \frac{4\rho}{1+3\rho} & \text{if } \gamma = \frac{1}{1-\rho}, \\ \rho & \text{if } \gamma > \frac{1}{1-\rho}. \end{cases} \quad (7)$$

In the subcritical case, the tokens contract in one step towards a single cluster when $n \rightarrow \infty$ while in the supercritical case, their inner product does not change. In fact, a careful inspection of the proof reveals that in this supercritical regime the attention operator converges to the identity as $n \rightarrow \infty$. When $\alpha > 0$, the subcritical case is mitigated by the residual connection which prevents tokens to collapse to a single point in one step. Nevertheless, this singular behavior reveals a major limitation in the simplex case: since the tokens are equidistant the phase transition reveals an all-or-nothing

phenomenon where attention transitions from $A_{ij} \sim 1/n$ so that $\text{ATT}(y_i) = \bar{y} = \frac{1}{n} \sum_{j=1}^n y_j$ for all i to $A_{ij} = \delta_{ij}$ so that $\text{ATT}(y_i) = y_i$ for all i . In the next section, we present a similar result Theorem 2.3, where the simplex assumption is relaxed.

Before we end this section, we present the proof for Equation (7) as a special case of Theorem 2.1. The detailed proof for Theorem 2.1 and the later Theorem 2.3 in Section 2.2 is included in Appendix A.

Proof 1 (Proof of Equation (7)) In Equation (5), when $\alpha = 0$, we have that $x'_i = \text{ATT}(y_i)$ for each $i = 1, 2, \dots, n$. In Equation (3), under Assumption 1, we notice that the quantity $\sum_{k=1}^n e^{a_{ik}}$ in the denominator of A_{ij} is independent of the choice of i , and equals to $e^\beta + (n-1)e^{\rho\beta}$. Denote this as $Z := e^\beta + (n-1)e^{\rho\beta}$. Then Equation (4) and (5) become

$$x'_i = \text{ATT}(y_i) = \frac{1}{Z} \left(e^\beta y_i + \sum_{m \neq i} e^{\rho\beta} y_m \right).$$

Under Assumption 1, a direct computation shows that for any $i = 1, 2, \dots, n$,

$$\langle x'_i, x'_i \rangle = \frac{1}{Z^2} \left(e^{2\beta} + 2(n-1)\rho e^{(1+\rho)\beta} + (n-1)(1+(n-2)\rho)e^{2\rho\beta} \right),$$

and for any two different $i, j = 1, 2, \dots, n$,

$$\langle x'_i, x'_j \rangle = \frac{1}{Z^2} \left(\rho e^{2\beta} + 2(1+(n-2)\rho)e^{\beta(1+\rho)} + ((n-2) + (n^2 - 3n + 3)\rho)e^{2\rho\beta} \right).$$

See also Lemma A.3 and Lemma A.4 for more detailed computations for $\langle x'_i, x'_i \rangle$ and $\langle x'_i, x'_j \rangle$.

For $Z = e^\beta + (n-1)e^{\rho\beta}$, when we let $\beta = \gamma \log n$, we see that $e^\beta = n^\gamma$ and $ne^{\rho\beta} = n^{1+\rho\gamma}$ in Z . The largest term in Z then depends on the relation between γ and $1+\rho\gamma$: when $\gamma < \frac{1}{1-\rho}$, $n^{1+\rho\gamma}$ is the largest term; when $\gamma > \frac{1}{1-\rho}$, n^γ is the largest term. We then directly get the following three phases for Z from the above arguments:

$$Z = \begin{cases} (1 + o_n(1)) \cdot ne^{\rho\beta} & \text{if } \gamma < \frac{1}{1-\rho}, \\ (2 + o_n(1)) \cdot e^\beta & \text{if } \gamma = \frac{1}{1-\rho}, \\ (1 + o_n(1)) \cdot e^\beta & \text{if } \gamma > \frac{1}{1-\rho}, \end{cases} \quad (8)$$

where the terms $o_n(1)$ go to 0 as $n \rightarrow +\infty$. Similarly, we can get the the following three phases for $\langle x'_i, x'_i \rangle$:

$$\lim_{n \rightarrow +\infty} \langle x'_i, x'_i \rangle = \begin{cases} \rho & \text{if } \gamma < \frac{1}{1-\rho}, \\ \frac{1+3\rho}{4} & \text{if } \gamma = \frac{1}{1-\rho}, \\ 1 & \text{if } \gamma > \frac{1}{1-\rho}. \end{cases} \quad (9)$$

For $\langle x'_i, x'_j \rangle$, we always have that $\lim_{n \rightarrow +\infty} \langle x'_i, x'_j \rangle = \rho$ for γ in these three different regimes. Then Equation (7) follows from these two limits because $\langle y'_i, y'_j \rangle = \langle x'_i / \|x'_i\|, x'_j / \|x'_j\| \rangle$.

2.2 THE ALMOST-SIMPLEX CASE

In this section, we relax Assumption 1 to allow pairwise angles and lengths to vary slightly. This relaxation makes it possible for tokens to lie in a dimension $d \ll n$. Although the resulting bounds are not as sharp as those obtained under Assumption 1, they demonstrate that the critical scaling $\beta_n = \Theta(\log n)$ is intrinsic and not merely an artifact of a particular geometric construction.

Assumption 2 There exist constants $q_1, q_2 \in (0, \infty)$, $\rho_1, \rho_2 \in (0, 1)$ such that $q_1 \leq \|x_i\|^2 \leq q_2$ and $\rho_1 \leq \langle y_i, y_j \rangle \leq \rho_2$, for any $i, j = 1, \dots, n$ and $i \neq j$. Moreover, $\rho_1 = \langle y_i, y_j \rangle$ for some i, j .

Remark 2.2 Assumption 2 already allows near-ties when ρ_2 is close to 1, and it can be further generalized to allow multiple exact ties in the top scores. Specifically, one may assume that there exists a fixed $k \in \mathbb{Z}_+$ such that, for each $i \in \{1, \dots, n\}$, there are at most k indices j with $\langle y_i, y_j \rangle =$

1. Under this setting, all of our main results continue to hold with the same critical scaling order $\log n$ but with different constants depending on k . In fact, this setting is a special case of the more general setting discussed in Appendix C, where Assumption 3 partitions the inner products into three ranges, $[\rho_1, \rho_2]$, $[\rho_3, \rho_4]$, and $\{1\}$, with $0 \leq \rho_1 < \rho_2 < \rho_3 < \rho_4 < 1$. The phase transition behavior remains of order $\log n$ in that general setting as well. For clarity and readability of the main exposition, we keep Assumption 2 in the main text.

It is easy to see using standard probabilistic tools that Assumption 2 holds with high probability when the y_i 's are independent random vectors uniformly distributed on a half-sphere for example.

Theorem 2.3 Under Assumption 2, we have the following phase transition when $\beta = \gamma \log n$ for some fixed $\gamma > 0$.

If $\gamma < \frac{1}{1-\rho_1}$, then there is a constant $\varepsilon > 0$ depending on α, ρ_2, q_1, q_2 , such that

$$\lim_{n \rightarrow +\infty} \min_{i \neq j} \langle y'_i, y'_j \rangle \geq \rho_1 + \varepsilon > \rho_1, \quad (10)$$

which implies that the angle between tokens becomes strictly smaller after an attention layer Equation (5).

If $\gamma > \frac{1}{1-\rho_2}$, then for any $i \in \llbracket 1, n \rrbracket$,

$$\text{ATT}(y_i) = y_i + o_n(1), \text{ and hence } x'_i = y_i + \alpha x_i + o_n(1), \quad (11)$$

where the term $o_n(1)$ goes to 0 as $n \rightarrow +\infty$ with a speed uniform in i . Hence, when $\gamma > \frac{1}{1-\rho_2}$, for any two different $i, j \in \llbracket 1, n \rrbracket$,

$$\lim_{n \rightarrow +\infty} \langle y'_i, y'_j \rangle = \langle y_i, y_j \rangle. \quad (12)$$

which implies that the angle between tokens does not change after an attention layer Equation (5).

The proof for Theorem 2.3 is included in Appendix A, but the general intuition is similar to the proof for Equation (7) in Section 2.1. As we have seen in that proof, the first step to build up phase transition regimes for $\langle y'_i, y'_j \rangle$ is to study the phase transition regimes for Z_i in Equation (3). Adjusting the logarithmic scaling factor γ causes different phase transition regimes for Z_i first. When γ is small enough, the weights $e^{a_{ik}}$ consisting of Z_i are asymptotically uniform, and each token almost equally interacts with the other tokens. When γ is large enough, each token mostly focuses on itself.

Building on this observation, Theorem 2.1 and Theorem 2.3 together demonstrate that γ controls the effective interaction range of each token. In particular, we have seen in Theorem 2.1 the existence of the critical regime when $\gamma = \frac{1}{1-\rho}$. In this case, although the tokens continue to contract, their rate of shrinkage is evidently slower than in the subcritical regime, as shown in Equation (6) and Equation (7).

It is hence natural to ask whether further regimes emerge when γ is varied between the supercritical and subcritical threshold. Indeed, in Appendix C, we prove the existence of a nontrivial middle phase when γ is between the two extrema $\frac{1}{1-\rho_1}$ and $\frac{1}{1-\rho_2}$, under a refined assumption on the distribution of tokens, which allows for a sharper characterization of the transition. Under this refined assumption, Theorem C.2 show the existence of γ_1, γ_2 such that Equation (5) presents three different phases: $\gamma < \gamma_1$, $\gamma_1 < \gamma < \gamma_2$, and $\gamma > \gamma_2$. In the extreme regimes, when $\gamma < \gamma_1$, each token interacts with almost all the remaining tokens, while when $\gamma > \gamma_2$, each token only focuses on itself, consistent with Theorem 2.3. In the intermediate regime $\gamma_1 < \gamma < \gamma_2$, however, the weights $e^{a_{ik}}$ concentrate on only a small subset of tokens, so that each Z_i and hence the update in Equation (5) is dominated by a few highly relevant interactions. This shows that the logarithmic scaling enables each token to dynamically select its most relevant context.

We conclude by noting that those $o_n(1)$ terms in our theorems satisfy the bound $|o_n(1)| \leq C_1 n^{-C_2}$ for some positive constants C_1, C_2 that are independent of d (though varying across theorems). As a result, the simplex configuration (Assumption 1) and the almost simplex configuration (Assumption 2) remains valid under repeated application of the ATT operator up to $\text{poly}(n)$ iterations. In particular, the accumulated error remains negligible at this scale, so our theorems and arguments extend to transformers with many layers.

2.3 PROPAGATION OF GRADIENTS UNDER ATTENTION LAYER

In the previous section, we established how attention scaling influences the propagation of token representations, corresponding to running the Transformer in the forward (inference) direction. During training, however, the Transformer is also executed in the *backward* direction to compute gradients via backpropagation (Rumelhart et al., 1986). In this section, we show that a similar phase transition arises in the backward pass: in the subcritical regime—where token representations rapidly collapse in the forward pass—the gradients also collapse, whereas in the supercritical regime they retain their scale. The stability of gradients is a crucial computational consideration that strongly affects a model’s ability to be trained effectively. For this reason, several theoretical analyses of gradient dynamics in Transformers have been conducted, albeit without attention scaling; see, for example, Cowsik et al. (2024); Dong et al. (2021); Noci et al. (2022).

Let the input token configuration be denoted by $X(0)$, and let $X(t)$ represent the positions of all tokens at the output of Transformer layer t . To compute gradients, one needs to evaluate the end-to-end input–output Jacobian across L layers of the Transformer. By the chain rule, this Jacobian can be expressed as

$$\frac{\partial X(L)}{\partial X(0)} = \frac{\partial X(L)}{\partial X(L-1)} \frac{\partial X(L-1)}{\partial X(L-2)} \cdots \frac{\partial X(1)}{\partial X(0)}.$$

Thus, the end-to-end Jacobian can be obtained by recursively computing and multiplying the layer-wise Jacobians. This procedure is known as the *adjoint method* in dynamical systems theory (Lions, 1971), and as *backpropagation* in the machine learning community.

Our main result shows that when $\beta_n = \gamma \log n$ with subcritical γ , the typical singular values of $\frac{\partial X(t+1)}{\partial X(t)}$ are close to zero (apart from the contribution of the residual connection). In contrast, for supercritical values of γ , the contribution of the attention component to the Jacobian is non-trivial and behaves as a normalization map.

We now proceed with formal definitions. For $x \in \mathbb{R}^d$, let $(x)_u$ denote its u -th coordinate for $u = 1, 2, \dots, d$. The concatenation $X = (x_1, x_2, \dots, x_n) \in \mathbb{R}^{nd}$ represents the configuration of all tokens. The normalization map is defined by

$$\mathcal{N}(X) = \mathcal{N}(x_1, x_2, \dots, x_n) := (N(x_1), N(x_2), \dots, N(x_n)), \quad (13)$$

and the attention map by

$$\mathcal{ATT}(Y) = \mathcal{ATT}(y_1, y_2, \dots, y_n) := (\text{ATT}(y_1), \text{ATT}(y_2), \dots, \text{ATT}(y_n)), \quad (14)$$

where $\text{ATT}(y_i)$ is defined in equation 4 and $Y = (y_1, \dots, y_n)$. Under these definitions, the update equation 5 can be written compactly as

$$X' = \mathcal{ATT}(\mathcal{N}(X)) + \alpha X, \quad (15)$$

where $X' = (x'_1, x'_2, \dots, x'_n)$.

We define the $nd \times nd$ Jacobian matrix as

$$\nabla_X X' := \left(\frac{\partial (x'_j)_v}{\partial (x_i)_u} \right)_{(j,v),(i,u)}, \quad (16)$$

for $i, j = 1, \dots, n$ and $u, v = 1, \dots, d$. The matrix norm of $\nabla_X X'$ is given by

$$\|\nabla_X X'\|^2 := \text{tr}[(\nabla_X X')^\top \nabla_X X'] = \sum_{i,j=1}^n \sum_{u,v=1}^d \left(\frac{\partial (x'_j)_v}{\partial (x_i)_u} \right)^2. \quad (17)$$

Let $\sigma_1, \sigma_2, \dots, \sigma_{nd}$ denote the singular values of $\nabla_X X'$. Then the normalized Jacobian norm satisfies

$$\frac{1}{nd} \|\nabla_X X'\|^2 = \frac{1}{nd} \sum_{i=1}^{nd} \sigma_i^2, \quad (18)$$

which represents the mean squared singular value of the Jacobian.

Before stating our results on $\frac{1}{nd} \|\nabla_X X'\|^2$, we note that the Jacobian $\nabla_X X'$ can be decomposed into the residual part αI_{nd} and the attention part $\nabla_X(\mathcal{ATT}(\mathcal{N}(X)))$. As shown in Theorems 2.1 and 2.3, the residual component αI_{nd} does not affect the phase transition behavior. Therefore, to streamline the analysis, we focus exclusively on the attention term $\nabla_X(\mathcal{ATT}(\mathcal{N}(X)))$ by setting $\alpha = 0$ in equation 15. The following theorems characterize $\frac{1}{nd} \|\nabla_X X'\|^2$ under this setting.

Theorem 2.4 *Adopt Assumption 1 and Equation (15) with $\alpha = 0$. Then, we have the following phase transition phenomenon: let $\beta = \gamma \log n$ where γ is a positive constant.*

If $\gamma < \frac{1}{1-\rho}$,

$$\frac{1}{nd} \|\nabla_X X'\|^2 = 0 + o_n(1). \quad (19)$$

If $\gamma = \frac{1}{1-\rho}$

$$\frac{1}{nd} \|\nabla_X X'\|^2 = \frac{1}{4q} \left(1 - \frac{1}{d}\right) + o_n(1). \quad (20)$$

If $\gamma > \frac{1}{1-\rho}$

$$\frac{1}{nd} \|\nabla_X X'\|^2 = \frac{1}{q} \left(1 - \frac{1}{d}\right) + o_n(1). \quad (21)$$

In both cases, the terms $o_n(1)$ go to 0 as $n \rightarrow +\infty$, with speeds depending on γ, ρ, q .

The results of the previous theorem show that under the simplex assumption, the phase transition in the backward dynamics (for gradients) is as sharp as for the forward pass: for small γ , gradients do not flow through the attention block.

We can also extend the analysis for Theorem 2.4 to the relaxed Assumption 2.

Theorem 2.5 *Adopt Assumption 2 and Equation (15) with $\alpha = 0$. Then, we have the following phase transition phenomenon: let $\beta = \gamma \log n$ where γ is a positive constant.*

If $\gamma < \frac{1}{1-\rho_1}$,

$$\frac{1}{nd} \|\nabla_X X'\|^2 \leq 4 \frac{\gamma^2 (\log(n))^2}{q_1 d} + o_n(1), \quad (22)$$

If $\gamma > \frac{1}{1-\rho_2}$,

$$\frac{1}{nd} \|\nabla_X X'\|^2 \geq \frac{1}{q_2} \left(1 - \frac{1}{d}\right) + o_n(1), \quad (23)$$

which is away from 0 even when d, n is very large. Indeed, when $\gamma > \frac{1}{1-\rho_2}$, for any fixed $i, j \in \llbracket 1, n \rrbracket$,

$$\left(\frac{\partial(\mathcal{ATT}(\mathcal{N}(x_j)))_v}{\partial(x_i)_u} \right)_{d \times d} = \frac{\delta_{ij}}{\|x_i\|} (I_d - y_i y_i^T) + \mathbf{o}_n(1) + o_n(1) \cdot I_d, \quad (24)$$

where the leading order term is exactly $\frac{\partial(\mathcal{N}(x_j))_v}{\partial(x_i)_u}$ as shown in Proposition B.1. Here, I_d is the $d \times d$ identity matrix, the term $\mathbf{o}_n(1)$ ($o_n(1)$, respectively) is a $d \times d$ matrix (constant, respectively) with matrix norm as defined in Equation (17) (value, respectively) going to 0 as $n \rightarrow +\infty$, with a speed independent of i, j but only depending on γ, ρ_2, q_1 .

We present the proofs for Theorem 2.4 and Theorem 2.5 in Appendix B. Note that the $\frac{\log^2 n}{d}$ term in equation 22 is small for typical values of n and d used in Transformers. Theorem 2.4 and Theorem 2.5 also corroborate the fact that tokens collapse fast when γ is in the subcritical regime, while each token only focuses on itself when γ is in the supercritical regime.

3 NUMERICAL EXPERIMENTS

This section reports numerical experiments designed to support our theoretical predictions. In the following numerical experiments, we test the phase transition in the almost-simplex case as Section 2.2. We generate samples $\{x_1, \dots, x_n\} \subset \mathbb{R}^d$ such that the expectations $\mathbb{E}\|x_i\|^2 = 1$ and $\mathbb{E}\langle x_i, x_j \rangle = \rho \in [0, 1]$ for $i \neq j$. More precisely, we generate x_i according to

$$x_i = \sqrt{\rho} z_0 + \sqrt{1 - \rho} z_i, \quad (25)$$

where z_0, z_1, \dots, z_n are i.i.d. standard Gaussian vectors in \mathbb{R}^d . The generated samples satisfy the Assumption 2 with high probability.

In Figure 1, we plot the input-to-output angle ratio λ , defined as

$$\lambda = \frac{2}{n(n-1)} \sum_{1 \leq i < j \leq n} \frac{1 - \langle y'_i, y'_j \rangle}{1 - \langle y_i, y_j \rangle}, \quad (26)$$

for samples processed through a single self-attention layer with different γ and of different dimensions d . Consistent with our theoretical predictions, the layer acts as a contraction mapping when γ is small, reducing pairwise output angles, whereas for large γ the output angles remain nearly unchanged from the input. Moreover, in the large d regime the angle between input tokens $\langle y_i, y_j \rangle$ ($i \neq j$) concentrate near ρ , so that the simplex Assumption 1 is effectively satisfied. In this setting, we observe a sharp phase transition in agreement with Theorem 2.1. In the small d regime, however, the input tokens $\langle y_i, y_j \rangle$ randomly distributed in an interval (ρ_1, ρ_2) , and an intermediate phase emerges in which the contraction is only partial: some angles shrink significantly while others remain close to their original values, which smooths out the transition.




Figure 1: Plots of the input-to-output angle ratio λ , defined in Equation (26), as a function of ρ and γ . The tokens are first normalized by a pre-layer normalization and then passed through a single self-attention layer (4), with residual connections and MLP layers omitted. The dashed curve corresponds to $\gamma = \frac{1}{1-\rho}$, which approximates the actual phase transition with increasing accuracy as d grows, as implied by Theorem 2.1.

In Figure 2, we plot the normalized matrix norm for the $nd \times nd$ matrix $\nabla_X X'$, defined as

$$\eta = \frac{1}{nd} \|\nabla_X X'\|^2, \quad (27)$$

for samples passed through a single self-attention layer with varying γ and dimension d . Across all three plots, the normalized gradient norm remains close to 0 when γ is small, while for large γ it approaches $1 - 1/d$, consistent with Theorem 2.5. Similar to the token-angle behavior, a sharp phase transition emerges near $\gamma = \frac{1}{1-\rho}$ in the large- d regime, in agreement with the predictions under the simplex assumption. In lower dimensions, fluctuations in the pairwise angle prevent perfect concentration, and the transition is smoothed into an intermediate regime where the gradient norm only partially stabilizes.




Figure 2: Plots of the normalized norm η of the gradient, defined by Equation (27), as a function of ρ and γ . The tokens are first normalized by a pre-layer normalization and then passed through a single self-attention layer (4), with residual connections and MLP layers omitted. The dash curve shows $\frac{1}{1-\rho}$, which approximate the actual phase transition with increasing accuracy as d grows, as implied by Theorem 2.4. The matrix norm η is computed by the Hutchinson trace estimator (Hutchinson, 1989), based on the definition in Equation (17).

4 CONCLUSION

This paper develops a framework for understanding phase transitions in self-attention as the context length n grows, identifying a critical scaling $\beta_n \asymp \log n$ that separates a subcritical contractive regime from a supercritical unchanged regime. A central message is that this transition is rooted in the geometry of the score landscape: in our model, the gaps between the top ordered scores remain $O(1)$, which leads to the $\log n$ scaling. We show that this scaling is robust under various perturbations and in settings permitting multiple near-ties, demonstrating that the $\log n$ law is a structural consequence of content-adaptive interactions.

These results also connect to two lines of work. First, the subcritical regime parallels the over-smoothing phenomenon in graph neural networks (GNNs), where indiscriminate averaging leads to rank collapse (Wu et al., 2024; Arroyo et al., 2025; DeZoort & Hanin, 2023; Chen et al., 2025b). The Jacobian-norm phase transition in our analysis mirrors vanishing-gradient thresholds in GNN. Despite this conceptual similarity, the underlying mechanisms are fundamentally different. In GNNs, the over-smoothing is *depth-driven*: the network dynamics can be viewed as a multiplicative composition of fixed matrices, largely insensitive to the geometric arrangement of the input features. In contrast, self-attention dynamics exhibit *length-driven* collapse, where the rank collapse can occur within a single layer and depends strongly on token positions and the context length. This is precisely the form of rank collapse identified in our analysis. Second, recent analysis of attention sinks (Xiao et al., 2023; Gu et al., 2024; Barbero et al., 2025) introduce additive stabilizing tokens to prevent excessive mixing, whereas our results focus on a multiplicative mechanism with temperature rescaling. Both approaches aim to maintain meaningful attention patterns at long but finite context lengths, and they address rank collapse from distinct and potentially complementary angles.

Taken together, these results provide a mathematically grounded view of long-context attention and suggest several directions for future work, including alternative token geometries, causal and multi-head extensions, and a deeper study of how depth, context length, and aggregation dynamics interact in self-attention dynamics.

LLM USAGE

Large Language Models (LLMs) were used during peer review for grammar and syntax refinement only. All ideas, technical content, analyses and conclusions remain the authors’ work.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their careful reading of our manuscript and for several insightful suggestions that have improved the presentation of the article.

REFERENCES

- Pedro Abdalla, Afonso S Bandeira, Martin Kassabov, Victor Souza, Steven H Strogatz, and Alex Townsend. Expander graphs are globally synchronising, 2022. *arXiv:2210.12788*.
- Álvaro Arroyo, Alessio Gravina, Benjamin Gutteridge, Federico Barbero, Claudio Gallicchio, Xiaowen Dong, Michael Bronstein, and Pierre Vandergheynst. On vanishing gradients, over-smoothing, and over-squashing in gnns: Bridging recurrent and graph learning. *arXiv preprint arXiv:2502.10818*, 2025.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. Qwen technical report, 2023. *arXiv:2309.16609*.
- Federico Barbero, Alvaro Arroyo, Xiangming Gu, Christos Perivolaropoulos, Michael Bronstein, Petar Veličković, and Razvan Pascanu. Why do llms attend to the first token? *arXiv preprint arXiv:2504.02732*, 2025.
- Iz Beltagy, Matthew E Peters, and Arman Cohan. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*, 2020.
- Giuseppe Bruno, Federico Pasqualotto, and Andrea Agazzi. Emergence of meta-stable clustering in mean-field transformer models. In *International Conference on Learning Representations*, 2025a.
- Giuseppe Bruno, Federico Pasqualotto, and Andrea Agazzi. A multiscale analysis of mean-field transformers in the moderate interaction regime. *NeurIPS*, 2025b.
- Shi Chen, Zhengjiang Lin, Yury Polyanskiy, and Philippe Rigollet. Quantitative clustering in mean-field transformer models, 2025a. *arXiv:2504.14697*.
- Ziang Chen, Zhengjiang Lin, Shi Chen, Yury Polyanskiy, and Philippe Rigollet. Residual connections provably mitigate oversmoothing in graph neural networks, 2025b. *arXiv:2501.00762*.
- Aditya Cowsik, Tamra Nebabu, Xiao-Liang Qi, and Surya Ganguli. Geometric dynamics of signal propagation predict trainability of transformers, 2024. *arXiv:2403.02579*.
- Laurens De Haan and Ana Ferreira. *Extreme value theory: an introduction*. Springer, 2006.
- B. Derrida. Random-energy model: An exactly solvable model of disordered systems. *Phys. Rev. B*, 24:2613–2626, 1981. doi: 10.1103/PhysRevB.24.2613.
- Gage DeZoort and Boris Hanin. Principles for initialization and architecture selection in graph neural networks with relu activations. *arXiv preprint arXiv:2306.11668*, 2023.
- Yihe Dong, Jean-Baptiste Cordonnier, and Andreas Loukas. Attention is not all you need: Pure attention loses rank doubly exponentially with depth. In *International Conference on Machine Learning*, pp. 2793–2803. PMLR, 2021.
- Borjan Geshkovski, Hugo Koubbi, Yury Polyanskiy, and Philippe Rigollet. Dynamic metastability in the self-attention model, 2024a. *arXiv:2410.06833*.
- Borjan Geshkovski, Cyril Letrouit, Yury Polyanskiy, and Philippe Rigollet. The emergence of clusters in self-attention dynamics. *Advances in Neural Information Processing Systems*, 36, 2024b.
- Borjan Geshkovski, Cyril Letrouit, Yury Polyanskiy, and Philippe Rigollet. A mathematical perspective on transformers. *Bull. Amer. Math. Soc.*, 2025.
- Alessio Giorlandino and Sebastian Goldt. Two failure modes of deep transformers and how to avoid them: a unified theory of signal propagation at initialisation, 2025. *arXiv:2505.24333*.

- Xiangming Gu, Tianyu Pang, Chao Du, Qian Liu, Fengzhuo Zhang, Cunxiao Du, Ye Wang, and Min Lin. When attention sink emerges in language models: An empirical view. *arXiv preprint arXiv:2410.10781*, 2024.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.
- Michael F Hutchinson. A stochastic estimator of the trace of the influence matrix for laplacian smoothing splines. *Communications in Statistics-Simulation and Computation*, 18(3):1059–1076, 1989.
- Vishesh Jain, Clayton Mizgerd, and Mehtaab Sawhney. The random graph process is globally synchronizing, 2025. *arXiv:2501.12205*.
- Nikita Karagodin, Yury Polyanskiy, and Philippe Rigollet. Clustering in causal attention masking, 2024. *arXiv:2411.04990*.
- Nikita Karagodin, Shu Ge, Yury Polyanskiy, and Philippe Rigollet. Normalization in attention dynamics, 2025. *arXiv:2510.22026*.
- Jacques Louis Lions. *Optimal control of systems governed by partial differential equations*, volume 170. Springer, 1971.
- Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 10012–10022, 2021.
- Ken M Nakanishi. Scalable-softmax is superior for attention, 2025. *arXiv:2501.19399*.
- Lorenzo Noci, Sotiris Anagnostidis, Luca Biggio, Antonio Orvieto, Sidak Pal Singh, and Aurelien Lucchi. Signal propagation in transformers: Theoretical perspectives and the role of rank collapse. *Advances in Neural Information Processing Systems*, 35:27198–27211, 2022.
- Bowen Peng, Jeffrey Quesnelle, Honglu Fan, and Enrico Shippole. Yarn: Efficient context window extension of large language models, 2023. *arXiv:2309.00071*.
- Yury Polyanskiy, Philippe Rigollet, and Andrew Yao. Synchronization of mean-field models on the circle, 2025. *arXiv:2507.22857*.
- Krishna C Puvvada, Faisal Ladhak, Santiago Akle Serrano, Cheng-Ping Hsieh, Shantanu Acharya, Somshubra Majumdar, Fei Jia, Samuel Kriman, Simeng Sun, Dima Rekesh, et al. Swan-gpt: An efficient and scalable approach for long-context language modeling, 2025. *arXiv:2504.08719*.
- David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-propagating errors. *nature*, 323(6088):533–536, 1986.
- Xinyi Wu, Amir Ajorlou, Zihui Wu, and Ali Jadbabaie. Demystifying oversmoothing in attention-based graph neural networks. *Advances in Neural Information Processing Systems*, 36, 2024.
- Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. Efficient streaming language models with attention sinks. *arXiv preprint arXiv:2309.17453*, 2023.

A PROOF OF THEOREM 2.1 AND THEOREM 2.3

In this section, we adopt Assumption 2 and prove Theorem 2.3 first. Then, we prove Theorem 2.1. To simplify notations, we define $\llbracket 1, n \rrbracket := \{1, 2, \dots, n\}$ for any $n \in \mathbb{Z}_+$.

We study the asymptotics of the quantity $\langle x'_i, x'_j \rangle$ as $n \rightarrow +\infty$. We use the notation

$$Z_i := \sum_{k=1}^n e^{a_{ik}} = e^\beta + \sum_{k \neq i} e^{a_{ik}}. \quad (28)$$

Lemma A.1 Let $\beta = \gamma \log n$ where γ is a positive constant. Under Assumption 2 and Equation (5), for any $i \in \llbracket 1, n \rrbracket$,

$$Z_i = \begin{cases} (1 + o_n(1)) \cdot \left(\sum_{k \neq i} e^{a_{ik}} \right) & \text{if } \gamma < \frac{1}{1-\rho_1}, \\ (1 + o_n(1)) \cdot e^\beta & \text{if } \gamma > \frac{1}{1-\rho_2}, \end{cases} \quad (29)$$

where the terms $o_n(1)$ go to 0 as $n \rightarrow +\infty$ with speeds independent of i but only depending on γ, ρ_1, ρ_2 .

Proof 2 (Proof of Lemma A.1) We notice that

$$Z_i = e^\beta + \sum_{k \neq i} e^{a_{ik}}. \quad (30)$$

We also notice that $e^{\beta t} = n^{\gamma t}$ for any t . It then holds that $e^\beta = n^\gamma$ and

$$n^{\gamma \rho_1} (n-1) \leq \sum_{k \neq i} e^{a_{ik}} \leq n^{\gamma \rho_2} (n-1). \quad (31)$$

Hence, when $\gamma < \frac{1}{1-\rho_1}$, $n^\gamma < n^{1+\gamma \rho_1}$, the leading order term in Z_i is $\sum_{k \neq i} e^{a_{ik}}$. We also see that

$$Z_i = \left(\frac{e^\beta}{\left(\sum_{k \neq i} e^{a_{ik}} \right)} + 1 \right) \cdot \left(\sum_{k \neq i} e^{a_{ik}} \right), \quad (32)$$

with

$$\frac{e^\beta}{\left(\sum_{k \neq i} e^{a_{ik}} \right)} \leq \frac{n^\gamma}{n^{\gamma \rho_1} (n-1)}, \quad (33)$$

which goes to 0 as $n \rightarrow +\infty$, and is independent of i but only depending on γ, ρ_1 . Similarly, when $\gamma > \frac{1}{1-\rho_2}$, $n^\gamma > n^{1+\gamma \rho_2}$, the leading order term in Z_i is e^β , and similar arguments hold true.

Lemma A.2 Let $\beta = \gamma \log n$ where γ is a positive constant. Under Assumption 2 and Equation (5), if $\gamma > \frac{1}{1-\rho_2}$, then for any $i \in \llbracket 1, n \rrbracket$,

$$\text{ATT}(y_i) = y_i + o_n(1), \text{ and hence } x'_i = y_i + \alpha x_i + o_n(1), \quad (34)$$

where the term $o_n(1)$ goes to 0 as $n \rightarrow +\infty$ with a speed independent of i but only depending on γ, ρ_2 .

Proof 3 (Proof of Lemma A.2) According to Lemma A.1, we see that when $\gamma > \frac{1}{1-\rho_2}$, $n^\gamma > n^{1+\gamma \rho_2}$, and hence,

$$\text{ATT}(y_i) = Z_i^{-1} \left(e^\beta y_i + \sum_{j \neq i} e^{a_{ij}} y_j \right) = (1 + o_n(1)) \left(y_i + e^{-\beta} \sum_{j \neq i} e^{a_{ij}} y_j \right). \quad (35)$$

Because $\|y_j\| = 1$,

$$\left\| e^{-\beta} \sum_{j \neq i} e^{a_{ij}} y_j \right\| \leq e^{-\beta} \sum_{j \neq i} e^{a_{ij}} \leq n^{-\gamma} \cdot n^{\gamma \rho_2} (n-1), \quad (36)$$

which goes to 0 as $n \rightarrow +\infty$, and is independent of i but only depending on γ, ρ_2 . This shows that when $\gamma > \frac{1}{1-\rho_2}$,

$$\text{ATT}(y_i) = (1 + o_n(1))(y_i + o_n(1)) = y_i + o_n(1). \quad (37)$$

Lemma A.3 Under Assumption 2 and Equation (5), for any $i \in \llbracket 1, n \rrbracket$,

$$\begin{aligned} \langle x'_i, x'_i \rangle &= \alpha^2 \|x_i\|^2 + \frac{2\alpha \|x_i\|}{Z_i} \left(e^\beta + \sum_{j \neq i} e^{a_{ij}} \langle y_i, y_j \rangle \right) \\ &\quad + \frac{1}{Z_i^2} \left(e^{2\beta} + 2e^\beta \sum_{j \neq i} e^{a_{ij}} \langle y_i, y_j \rangle + \sum_{j \neq i} \sum_{k \neq i} e^{a_{ij} + a_{ik}} \langle y_k, y_j \rangle \right). \end{aligned} \quad (38)$$

Let $\beta = \gamma \log n$ where γ is a positive constant. When $\gamma < \frac{1}{1-\rho_1}$,

$$\langle x'_i, x'_i \rangle = \alpha^2 \|x_i\|^2 + 2\alpha \|x_i\| \frac{\sum_{k \neq i} e^{a_{ik}} \langle y_i, y_k \rangle}{\sum_{k \neq i} e^{a_{ik}}} + \frac{\sum_{k \neq i} \sum_{l \neq i} e^{a_{ik} + a_{il}} \langle y_k, y_l \rangle}{\left(\sum_{k \neq i} e^{a_{ik}} \right)^2} + o_n(1). \quad (39)$$

When $\gamma > \frac{1}{1-\rho_2}$,

$$\langle x'_i, x'_i \rangle = (\alpha \|x_i\| + 1)^2 + o_n(1). \quad (40)$$

In both cases, the terms $o_n(1)$ go to 0 as $n \rightarrow +\infty$ with speeds independent of i but only depending on $\gamma, \rho_1, \rho_2, \alpha$.

Proof 4 (Proof of Lemma A.3) According to Equation (5), we see that

$$\langle x'_i, x'_i \rangle = \alpha^2 \|x_i\|^2 + 2\alpha \langle x_i, \text{ATT}(y_i) \rangle + \langle \text{ATT}(y_i), \text{ATT}(y_i) \rangle. \quad (41)$$

Equation (38) follows from direct computations. Two phase transitions Equation (39) and Equation (40) follow from similar arguments as in Lemma A.1.

Lemma A.4 Under Assumption 2 and Equation (5), for any two different $i, j \in \llbracket 1, n \rrbracket$,

$$\begin{aligned} \langle x'_i, x'_j \rangle &= \alpha^2 \langle x_i, x_j \rangle + \frac{\alpha \|x_j\|}{Z_i} \left(e^\beta \langle y_j, y_i \rangle + \sum_{k \neq i} e^{a_{ik}} \langle y_j, y_k \rangle \right) + \frac{\alpha \|x_i\|}{Z_j} \left(e^\beta \langle y_i, y_j \rangle + \sum_{l \neq j} e^{a_{jl}} \langle y_i, y_l \rangle \right) \\ &\quad + \frac{1}{Z_i Z_j} \left(e^{2\beta} \langle y_i, y_j \rangle + e^\beta \sum_{k \neq i} e^{a_{ik}} \langle y_j, y_k \rangle + e^\beta \sum_{l \neq j} e^{a_{jl}} \langle y_i, y_l \rangle + \sum_{k \neq i} \sum_{l \neq j} e^{a_{ik} + a_{jl}} \langle y_k, y_l \rangle \right). \end{aligned} \quad (42)$$

Let $\beta = \gamma \log n$ where γ is a positive constant. When $\gamma < \frac{1}{1-\rho_1}$,

$$\begin{aligned} \langle x'_i, x'_j \rangle &= \alpha^2 \langle x_i, x_j \rangle + \alpha \|x_j\| \frac{\sum_{k \neq i} e^{a_{ik}} \langle y_j, y_k \rangle}{\sum_{k \neq i} e^{a_{ik}}} + \alpha \|x_i\| \frac{\sum_{l \neq j} e^{a_{jl}} \langle y_i, y_l \rangle}{\sum_{l \neq j} e^{a_{jl}}} \\ &\quad + \frac{\sum_{k \neq i} \sum_{l \neq j} e^{a_{ik} + a_{jl}} \langle y_k, y_l \rangle}{\left(\sum_{k \neq i} e^{a_{ik}} \right) \cdot \left(\sum_{l \neq j} e^{a_{jl}} \right)} + o_n(1). \end{aligned} \quad (43)$$

When $\gamma > \frac{1}{1-\rho_2}$,

$$\langle x'_i, x'_j \rangle = (\alpha \|x_i\| + 1)(\alpha \|x_j\| + 1) \langle y_i, y_j \rangle + o_n(1). \quad (44)$$

Proof 5 (Proof of Lemma A.4) According to Equation (5), we see that for two different $i, j \in \llbracket 1, n \rrbracket$,

$$\langle x'_i, x'_j \rangle = \alpha^2 p + \alpha \langle x_i, \text{ATT}(y_j) \rangle + \alpha \langle x_j, \text{ATT}(y_i) \rangle + \langle \text{ATT}(y_i), \text{ATT}(y_j) \rangle. \quad (45)$$

Equation (42) follows from direct computations. Two phase transitions Equation (43) and Equation (44) follow from similar arguments as in Lemma A.1.

Next, we prove Theorem 2.3.

Proof 6 (Proof of Theorem 2.3) We first discuss the case when $\gamma < \frac{1}{1-\rho_1}$. According to Equation (43) and Assumption 2, we see that

$$\begin{aligned}\langle x'_i, x'_j \rangle &\geq \alpha^2 \|x_i\| \|x_j\| \rho_1 + \alpha \|x_j\| \rho_1 + \alpha \|x_i\| \rho_1 + \rho_1 + o_n(1) \\ &= \rho_1 (\alpha \|x_i\| + 1) (\alpha \|x_j\| + 1) + o_n(1).\end{aligned}\quad (46)$$

By Equation (39), we see that

$$\begin{aligned}\langle x'_i, x'_i \rangle &\leq \alpha^2 \|x_i\|^2 + 2\alpha \|x_i\| \rho_2 + \rho_2 + o_n(1) \\ &= \alpha^2 \|x_i\|^2 + 2\alpha \|x_i\| + 1 - (1 - \rho_2)(1 + 2\alpha \|x_i\|) + o_n(1) \\ &\leq (\alpha \|x_i\| + 1)^2 - (1 - \rho_2)(1 + 2\alpha q_1) + o_n(1).\end{aligned}\quad (47)$$

We have a similar inequality for $\langle x'_j, x'_j \rangle$. So, there is a constant $\delta > 0$ depending on ρ_2, α, q_1, q_2 and independent of n , such that

$$\frac{1}{\|x'_i\|} \geq \frac{1 + \delta}{\alpha \|x_i\| + 1} + o_n(1), \text{ and } \frac{1}{\|x'_j\|} \geq \frac{1 + \delta}{\alpha \|x_j\| + 1} + o_n(1).\quad (48)$$

Hence,

$$\langle y'_i, y'_j \rangle \geq \rho_1 (1 + \delta)^2 + o_n(1) \geq \rho_1 + \varepsilon + o_n(1),\quad (49)$$

for $\varepsilon = \rho_1 (1 + 2\delta)\delta > 0$ independent of n .

For the case when $\gamma < \frac{1}{1-\rho_2}$, Equation (11) and Equation (12) follow directly from Lemma A.2, Lemma A.3, and Lemma A.4.

Proof 7 (Proof of Theorem 2.1) We notice that Assumption 1 corresponds to the special case when $q_1 = q_2 = q$ and $\rho_1 = \rho_2 = \rho$ in Assumption 2. Clearly, Z_i is independent of the choice of $i \in \llbracket 1, n \rrbracket$ by its definition Equation (28). According to the explicit forms Equation (38) in Lemma A.3 and Equation (42) in Lemma A.4, one directly sees that both $\langle x_i, x_i \rangle$ and $\langle x_i, x_j \rangle$ are independent of the choices of $i, j \in \llbracket 1, n \rrbracket$. We can further compute that for any $i \in \llbracket 1, n \rrbracket$,

$$\lim_{n \rightarrow +\infty} \langle x'_i, x'_i \rangle = \begin{cases} \alpha^2 q + 2\alpha \sqrt{q} \rho + \rho & \text{if } \gamma < \frac{1}{1-\rho}, \\ \alpha^2 q + \alpha \sqrt{q} (1 + \rho) + \frac{1+3\rho}{4} & \text{if } \gamma = \frac{1}{1-\rho}, \\ (\alpha \sqrt{q} + 1)^2 & \text{if } \gamma > \frac{1}{1-\rho}, \end{cases}\quad (50)$$

and for any two different $i, j \in \llbracket 1, n \rrbracket$,

$$\lim_{n \rightarrow +\infty} \langle x'_i, x'_j \rangle = \rho (\alpha \sqrt{q} + 1)^2.\quad (51)$$

Equation (6) follows from Equation (50) and Equation (51).

When $\gamma < \frac{1}{1-\rho}$, we see that

$$\lim_{n \rightarrow +\infty} \langle y'_i, y'_j \rangle = \frac{\rho (\alpha \sqrt{q} + 1)^2}{\alpha^2 q + 2\alpha \sqrt{q} \rho + \rho} > \frac{\rho (\alpha \sqrt{q} + 1)^2}{\alpha^2 q + 2\alpha \sqrt{q} + 1} = \rho,\quad (52)$$

where the strict inequality is because $\rho < 1$. When $\gamma = \frac{1}{1-\rho}$, we can similarly show that $\lim_{n \rightarrow +\infty} \langle y'_i, y'_j \rangle > \rho$. This completes the proof for Theorem 2.1.

B PROOF OF THEOREM 2.4 AND THEOREM 2.5

We prove Theorem 2.5 first. We need to explicitly compute terms in $\frac{\partial(\text{ATT}(N(x_j)))_v}{\partial(x_i)_u}$, for which we need the following lemmas.

B.1 PROOF OF THEOREM 2.5

Lemma B.1 For any $i, k \in \llbracket 1, n \rrbracket$ and $u, w \in \llbracket 1, d \rrbracket$,

$$\frac{\partial(N(x_k))_w}{\partial(x_i)_u} = \delta_{ik} \frac{\delta_{uw} \|x_k\|^2 - (x_k)_w (x_k)_u}{\|x_k\|^3}.\quad (53)$$

Proof 8 (Proof of Lemma B.1)

$$\frac{\partial(N(x_k))_w}{\partial(x_i)_u} = \frac{\partial((x_k)_w \cdot \|x_k\|^{-1})}{\partial(x_i)_u} = \delta_{ik} \frac{\delta_{wu} \|x_k\| - (x_k)_w \cdot \frac{(x_k)_u}{\|x_k\|}}{\|x_k\|^2}. \quad (54)$$

Lemma B.2 For any $k, j \in \llbracket 1, n \rrbracket$ and $w, v \in \llbracket 1, d \rrbracket$,

$$\begin{aligned} & \frac{\partial(\text{ATT}(y_j))_v}{\partial(y_k)_w} \\ &= \left[\left(\delta_{kj} \beta \left(\sum_{m=1}^n e^{\beta \langle y_j, y_m \rangle} (y_m)_w (y_m)_v \right) + e^{\beta \langle y_j, y_k \rangle} (\beta(y_j)_w (y_k)_v + \delta_{wv}) \right) \cdot \left(\sum_{l=1}^n e^{\beta \langle y_j, y_l \rangle} \right) \right. \\ & \quad \left. - \left(\delta_{kj} \beta \left(\sum_{l=1}^n e^{\beta \langle y_j, y_l \rangle} (y_l)_w \right) + \beta e^{\beta \langle y_j, y_k \rangle} (y_j)_w \right) \cdot \left(\sum_{m=1}^n e^{\beta \langle y_j, y_m \rangle} (y_m)_v \right) \right] \\ & \quad \cdot \left(\sum_{l=1}^n e^{\beta \langle y_j, y_l \rangle} \right)^{-2}. \end{aligned} \quad (55)$$

Proof 9 (Proof of Lemma B.2) By Equation (4),

$$(\text{ATT}(y_j))_v = \frac{\sum_{m=1}^n e^{\beta \langle y_j, y_m \rangle} (y_m)_v}{\sum_{l=1}^n e^{\beta \langle y_j, y_l \rangle}}. \quad (56)$$

A direct computation shows that

$$\begin{aligned} & \frac{\partial(\text{ATT}(y_j))_v}{\partial(y_k)_w} \\ &= \left[\left(\sum_{m=1}^n (\delta_{kj} \beta (y_m)_w (y_m)_v + \delta_{km} \beta (y_j)_w (y_m)_v + \delta_{km} \delta_{wv}) e^{\beta \langle y_j, y_m \rangle} \right) \cdot \left(\sum_{l=1}^n e^{\beta \langle y_j, y_l \rangle} \right) \right. \\ & \quad \left. - \left(\sum_{l=1}^n (\delta_{kj} (y_l)_w + \delta_{kl} (y_j)_w) \beta e^{\beta \langle y_j, y_l \rangle} \right) \cdot \left(\sum_{m=1}^n e^{\beta \langle y_j, y_m \rangle} (y_m)_v \right) \right] \\ & \quad \cdot \left(\sum_{l=1}^n e^{\beta \langle y_j, y_l \rangle} \right)^{-2} \\ &= \left[\left(\delta_{kj} \beta \sum_{m=1}^n e^{\beta \langle y_j, y_m \rangle} (y_m)_w (y_m)_v + e^{\beta \langle y_j, y_k \rangle} (\beta(y_j)_w (y_k)_v + \delta_{wv}) \right) \cdot \left(\sum_{l=1}^n e^{\beta \langle y_j, y_l \rangle} \right) \right. \\ & \quad \left. - \left(\delta_{kj} \beta \sum_{l=1}^n e^{\beta \langle y_j, y_l \rangle} (y_l)_w + \beta e^{\beta \langle y_j, y_k \rangle} (y_j)_w \right) \cdot \left(\sum_{m=1}^n e^{\beta \langle y_j, y_m \rangle} (y_m)_v \right) \right] \\ & \quad \cdot \left(\sum_{l=1}^n e^{\beta \langle y_j, y_l \rangle} \right)^{-2}. \end{aligned} \quad (57)$$

For $x, y \in \mathbb{R}^d$, we use $x \otimes y$ to denote the $d \times d$ matrix with (u, v) -th element $(x \otimes y)_{uv} = (x)_u (y)_v$, i.e., $x \otimes y := xy^T$. We then have the following proposition.

Lemma B.3 Adopt Assumption 2 and Equation (5). For any $i, j \in \llbracket 1, n \rrbracket$, consider the $d \times d$ matrix formed by $\frac{\partial(\text{ATT}(N(x_j)))_v}{\partial(x_i)_u}$, for $u, v \in \llbracket 1, d \rrbracket$. Denote $y_k = N(x_k)$ for each $k \in \llbracket 1, n \rrbracket$. Then, this matrix has the following form:

$$\left(\frac{\partial(\text{ATT}(N(x_j)))_v}{\partial(x_i)_u} \right)_{d \times d} = \|x_i\|^{-\frac{1}{2}} [(\mathbf{R}_1 + \mathbf{R}_2)Z_j - (\mathbf{U}_1 + \mathbf{U}_2) \otimes \mathbf{V}_j] \cdot Z_j^{-2}, \quad (58)$$

where $Z_j = \sum_{l=1}^n e^{\beta \langle y_j, y_l \rangle}$ as in Equation (28),

$$\mathbf{R}_1 := \delta_{ij} \beta (\mathbf{W}_j - y_i \otimes (\mathbf{W}_j y_i)), \quad \mathbf{R}_2 := e^{\beta \langle y_j, y_i \rangle} ((-y_i + \beta \mathbf{P}_{y_i} y_j) \otimes y_i + I_d), \quad (59)$$

and

$$\mathbf{U}_1 := \delta_{ij} \beta (\mathbf{P}_{y_i} \mathbf{V}_j), \quad \mathbf{U}_2 := \beta e^{\beta \langle y_j, y_i \rangle} (\mathbf{P}_{y_i} y_j). \quad (60)$$

In Equation (59) and Equation (60),

$$\mathbf{V}_j := \sum_{m=1}^n e^{\beta \langle y_j, y_m \rangle} y_m, \quad \mathbf{W}_j := \sum_{m=1}^n e^{\beta \langle y_j, y_m \rangle} y_m \otimes y_m, \quad \mathbf{P}_x y := y - \langle y, x \rangle x. \quad (61)$$

Proof 10 (Proof of Lemma B.3) By chain rule and Proposition B.1, we have that

$$\begin{aligned} \frac{\partial(\text{ATT}(N(x_j)))_v}{\partial(x_i)_u} &= \sum_{k=1}^n \sum_{w=1}^d \frac{\partial(\text{ATT}(y_j))_v}{\partial(y_k)_w} \Big|_{Y=\mathcal{N}(X)} \cdot \frac{\partial(N(x_k))_w}{\partial(x_i)_u} \\ &= \|x_i\|^{-\frac{3}{2}} \left(\|x_i\| \cdot \frac{\partial(\text{ATT}(y_j))_v}{\partial(y_i)_u} - \sum_{w=1}^d (x_i)_u (x_i)_w \frac{\partial(\text{ATT}(y_j))_v}{\partial(y_i)_w} \right) \Big|_{Y=\mathcal{N}(X)} \\ &= \|x_i\|^{-\frac{1}{2}} \left(\frac{\partial(\text{ATT}(y_j))_v}{\partial(y_i)_u} - (y_i)_u \sum_{w=1}^d (y_i)_w \frac{\partial(\text{ATT}(y_j))_v}{\partial(y_i)_w} \right) \Big|_{Y=\mathcal{N}(X)}. \end{aligned} \quad (62)$$

According to Proposition B.2 and the notation $Z_j = \sum_{l=1}^n e^{a_{jl}}$, we see that

$$\begin{aligned} &\sum_{w=1}^d (y_i)_w \frac{\partial(\text{ATT}(y_j))_v}{\partial(y_i)_w} \\ &= \left[\left(\delta_{ij} \beta \left(\sum_{m=1}^n e^{\beta \langle y_j, y_m \rangle} \langle y_m, y_i \rangle (y_m)_v \right) + e^{\beta \langle y_j, y_i \rangle} (\beta \langle y_j, y_i \rangle + 1) (y_i)_v \right) \cdot Z_j \right. \\ &\quad \left. - \left(\delta_{ij} \beta \left(\sum_{l=1}^n e^{\beta \langle y_j, y_l \rangle} \langle y_l, y_i \rangle \right) + \beta e^{\beta \langle y_j, y_i \rangle} \langle y_j, y_i \rangle \right) \cdot \left(\sum_{m=1}^n e^{\beta \langle y_j, y_m \rangle} (y_m)_v \right) \right] \cdot Z_j^{-2}. \end{aligned} \quad (63)$$

Hence,

$$\begin{aligned} &\|x_i\|^{\frac{1}{2}} \cdot \frac{\partial(\text{ATT}(N(x_j)))_v}{\partial(x_i)_u} \\ &= \left(\frac{\partial(\text{ATT}(y_j))_v}{\partial(y_i)_u} - (y_i)_u \sum_{w=1}^d (y_i)_w \frac{\partial(\text{ATT}(y_j))_v}{\partial(y_i)_w} \right) \Big|_{Y=\mathcal{N}(X)} \\ &= \left[\left[\delta_{ij} \beta \left(\sum_{m=1}^n e^{\beta \langle y_j, y_m \rangle} ((y_m)_u (y_m)_v - \langle y_m, y_i \rangle (y_m)_v (y_i)_u) \right) \right. \right. \\ &\quad \left. \left. + e^{\beta \langle y_j, y_i \rangle} (\beta (y_j)_u (y_i)_v + \delta_{uv} - (\beta \langle y_j, y_i \rangle + 1) (y_i)_v (y_i)_u) \right] \cdot Z_j \right. \\ &\quad \left. - \left[\delta_{ij} \beta \left(\sum_{l=1}^n e^{\beta \langle y_j, y_l \rangle} ((y_l)_u - \langle y_l, y_i \rangle (y_i)_u) \right) \right. \right. \\ &\quad \left. \left. + \beta e^{\beta \langle y_j, y_i \rangle} ((y_j)_u - \langle y_j, y_i \rangle (y_i)_u) \right] \cdot \left(\sum_{m=1}^n e^{\beta \langle y_j, y_m \rangle} (y_m)_v \right) \right] \cdot Z_j^{-2}. \end{aligned} \quad (64)$$

We then adopt the notation Equation (61), i.e.,

$$\mathbf{V}_j = \sum_{m=1}^n e^{\beta \langle y_j, y_m \rangle} y_m, \quad \mathbf{W}_j = \sum_{m=1}^n e^{\beta \langle y_j, y_m \rangle} y_m \otimes y_m, \quad \mathbf{P}_x y := y - \langle y, x \rangle x. \quad (65)$$

So, the matrix form of Equation (64) becomes

$$\begin{aligned}
& \left[\left[\delta_{ij} \beta (\mathbf{W}_j - y_i \otimes (\mathbf{W}_j y_i)) + e^{\beta \langle y_j, y_i \rangle} (\beta y_j \otimes y_i + I_d - (\beta \langle y_j, y_i \rangle + 1) y_i \otimes y_i) \right] \cdot Z_j \right. \\
& \quad \left. - \left[\delta_{ij} \beta (\mathbf{V}_j - \langle \mathbf{V}_j, y_i \rangle y_i) + \beta e^{\beta \langle y_j, y_i \rangle} (y_j - \langle y_j, y_i \rangle y_i) \right] \otimes \mathbf{V}_j \right] \cdot Z_j^{-2} \\
& = \left[\left[\delta_{ij} \beta (\mathbf{W}_j - y_i \otimes (\mathbf{W}_j y_i)) + e^{\beta \langle y_j, y_i \rangle} ((-y_i + \beta \mathbf{P}_{y_i} y_j) \otimes y_i + I_d) \right] \cdot Z_j \right. \\
& \quad \left. - \left[\delta_{ij} \beta (\mathbf{P}_{y_i} \mathbf{V}_j) + \beta e^{\beta \langle y_j, y_i \rangle} (\mathbf{P}_{y_i} y_j) \right] \otimes \mathbf{V}_j \right] \cdot Z_j^{-2}.
\end{aligned} \tag{66}$$

We further use the notations in Equation (59) and Equation (60), i.e.,

$$\mathbf{R}_1 = \delta_{ij} \beta e^{\beta \rho} (\mathbf{W}_j - y_i \otimes (\mathbf{W}_j y_i)), \quad \mathbf{R}_2 = e^{\beta \langle y_j, y_i \rangle} ((-y_i + \beta \mathbf{P}_{y_i} y_j) \otimes y_i + I_d), \tag{67}$$

and

$$\mathbf{U}_1 = \delta_{ij} \beta (\mathbf{P}_{y_i} \mathbf{V}_j), \quad \mathbf{U}_2 = \beta e^{\beta \langle y_j, y_i \rangle} (\mathbf{P}_{y_i} y_j). \tag{68}$$

Finally, the matrix form of Equation (64) becomes

$$[(\mathbf{R}_1 + \mathbf{R}_2)Z - (\mathbf{U}_1 + \mathbf{U}_2) \otimes \mathbf{V}_j] \cdot Z_j^{-2}. \tag{69}$$

Lemma B.4 Let $\beta = \gamma \log n$ where γ is a positive constant. Under Assumption 2 and Equation (5), if $\gamma > \frac{1}{1-\rho_2}$, then for any fixed $i, j \in \llbracket 1, n \rrbracket$, the $d \times d$ matrix satisfies

$$\left(\frac{\partial(\text{ATT}(N(x_j)))_v}{\partial(x_i)_u} \right)_{d \times d} = \frac{\delta_{ij}}{\|x_i\|} (I_d - y_i \otimes y_i) + \mathbf{o}_n(1) + o_n(1) \cdot I_d, \tag{70}$$

where the leading order term is exactly $\frac{\partial(N(x_j))_v}{\partial(x_i)_u}$. The term $\mathbf{o}_n(1)$ ($o_n(1)$, respectively) is a $d \times d$ matrix (constant, respectively) with matrix norm as defined in Equation (17) (value, respectively) going to 0 as $n \rightarrow +\infty$, with a speed independent of i, j but only depending on γ, ρ_2, q_1 .

Proof 11 (Proof of Lemma B.4) We frequently use this formula: for two vectors V_1, V_2 , the matrix norm of $V_1 \otimes V_2$ as defined in Equation (17) is $\|V_1\| \|V_2\|$. When $\gamma > \frac{1}{1-\rho_2}$, $n^\gamma > n^{1+\gamma\rho_2}$, and we know from Lemma A.1 that $Z_j = (1 + o_n(1)) \cdot e^\beta$ for any $j \in \llbracket 1, n \rrbracket$. Adopt the notations in Proposition B.3, we then show the following facts when $\gamma > \frac{1}{1-\rho_2}$:

$$\mathbf{R}_1 Z_j^{-1} = \mathbf{o}_n(1), \quad \mathbf{R}_2 Z_j^{-1} = \delta_{ij} (-y_i \otimes y_i + I_d) + \mathbf{o}_n(1) + o_n(1) \cdot I_d, \tag{71}$$

and

$$[(\mathbf{U}_1 + \mathbf{U}_2) \otimes \mathbf{V}_j] \cdot Z_j^{-2} = \mathbf{o}_n(1). \tag{72}$$

First, for $\mathbf{R}_1 Z_j^{-1}$, when $i \neq j$, we have that $\mathbf{R}_1 = 0$ by its definition. When $i = j$, $\mathbf{R}_1 = \beta \sum_{m=1}^n e^{\beta \langle y_i, y_m \rangle} (y_m \otimes y_m - \langle y_m, y_i \rangle y_i \otimes y_m)$ and we notice that the term when $m = i$ is 0. So, because $\|y_m \otimes y_m - \langle y_m, y_i \rangle y_i \otimes y_m\| \leq \|y_m\|^2 + \|y_m\|^2 \|y_i\|^2 = 2$, $e^\beta = n^\gamma$,

$$\|\mathbf{R}_1\| Z_j^{-1} \leq \beta(n-1) e^{\beta \rho_2} \cdot 2 Z_j^{-1} \leq 2\gamma \log(n) \cdot n^{\gamma \rho_2 + 1 - \gamma} (1 + o_n(1)), \tag{73}$$

which goes to 0 with a speed independent of i, j , because $\gamma \rho_2 + 1 - \gamma < 0$.

For $\mathbf{R}_2 Z_j^{-1}$, we notice that when $i \neq j$, $e^{\beta \langle y_i, y_j \rangle} Z_j^{-1} \leq e^{\beta(\rho_2-1)} (1 + o_n(1)) = n^{\gamma(\rho_2-1)} (1 + o_n(1))$, which goes to 0 with a speed independent of i, j . So, $\mathbf{R}_2 Z_j^{-1} = \mathbf{o}_n(1) + o_n(1) \cdot I_d$ when $i \neq j$. When $i = j$, $\mathbf{R}_2 = e^\beta (-y_i \otimes y_i + I_d)$, and so $\mathbf{R}_2 Z_j^{-1} = (-y_i \otimes y_i + I_d) + \mathbf{o}_n(1) + o_n(1) \cdot I_d$.

For $[(\mathbf{U}_1 + \mathbf{U}_2) \otimes \mathbf{V}_j] \cdot Z_j^{-2}$, we see that when $i \neq j$, $\mathbf{U}_1 = 0$, and so

$$\begin{aligned}
& \|(\mathbf{U}_1 + \mathbf{U}_2) \otimes \mathbf{V}_j\| \cdot Z_j^{-2} \leq Z_j^{-2} \beta \sum_{m=1}^n e^{\beta \langle y_j, y_m + y_i \rangle} \|y_m\| \|\mathbf{P}_{y_i} y_j\| \\
& \leq Z_j^{-2} \beta e^{\beta(1+\rho_2)} n = \gamma \log(n) n^{\gamma(\rho_2-1)+1} (1 + o_n(1)),
\end{aligned} \tag{74}$$

which goes to 0 with a speed independent of i, j because $\gamma > \frac{1}{1-\rho_2}$. When $i = j$, $\mathbf{U}_2 = 0$, and so

$$\begin{aligned} \|(\mathbf{U}_1 + \mathbf{U}_2) \otimes \mathbf{V}_j\| \cdot Z_j^{-2} &\leq Z_j^{-2} \beta \|\mathbf{P}_{y_i} \mathbf{V}_i\| \|\mathbf{V}_i\| \\ &\leq Z_j^{-2} \beta \left(\sum_{m \neq i} e^{\beta \langle y_i, y_m \rangle} \|\mathbf{P}_{y_i} y_m\| \right) \cdot \left(e^\beta + \sum_{m \neq i} e^{\beta \langle y_i, y_m \rangle} \|\mathbf{P}_{y_i} y_m\| \right) \\ &\leq Z_j^{-2} \beta (e^{\beta \rho_2 n}) \cdot (e^\beta + e^{\beta \rho_2 n}) \\ &= \gamma \log(n) n^{\gamma(\rho_2-1)+1} (1 + n^{\gamma(\rho_2-1)+1}) (1 + o_n(1)), \end{aligned} \quad (75)$$

which goes to 0 with a speed independent of i, j because $\gamma > \frac{1}{1-\rho_2}$. Hence, $[(\mathbf{U}_1 + \mathbf{U}_2) \otimes \mathbf{V}_j] \cdot Z_j^{-2} = o_n(1)$.

Lemma B.5 Let $\beta = \gamma \log n$ where γ is a positive constant. Under Assumption 2 and Equation (5), if $\gamma < \frac{1}{1-\rho_1}$, then for fixed $i, j \in \llbracket 1, n \rrbracket$, the $d \times d$ matrix satisfies

$$\left\| \left(\frac{\partial(\text{ATT}(N(x_j)))_v}{\partial(x_i)_u} \right)_{d \times d} \right\| \leq \|x_i\|^{-\frac{1}{2}} \cdot (2\beta \delta_{ij} + (2\beta + \sqrt{d}) e^{a_{ij}} Z_j^{-1}). \quad (76)$$

Proof 12 (Proof of Lemma B.5) According to Lemma A.1, when $\gamma < \frac{1}{1-\rho_1}$, $Z_j = (1 + o_n(1)) \cdot (\sum_{k \neq j} e^{a_{jk}})$ for any $j \in \llbracket 1, n \rrbracket$, and $Z_j \geq n^{\gamma \rho_1 + 1} (1 + o_n(1)) > n^\gamma (1 + o_n(1))$, because $\gamma \rho_1 + 1 > \gamma$. Adopt the notations in Proposition B.3, we then show the following facts when $\gamma < \frac{1}{1-\rho_1}$:

$$\|\mathbf{R}_1\| Z_j^{-1} \leq \delta_{ij} \beta, \quad \|\mathbf{R}_2\| Z_j^{-1} \leq Z_j^{-1} e^{a_{ij}} (\beta + \sqrt{d-1}), \quad (77)$$

and

$$\|(\mathbf{U}_1 + \mathbf{U}_2) \otimes \mathbf{V}_j\| \cdot Z_j^{-2} \leq \beta (\delta_{ij} + e^{a_{ij}} Z_j^{-1}). \quad (78)$$

First, for $\mathbf{R}_1 Z_j^{-1}$, when $i \neq j$, we have that $\mathbf{R}_1 = 0$ by its definition. When $i = j$, $\mathbf{R}_1 = \beta \sum_{m=1}^n e^{\beta \langle y_i, y_m \rangle} (y_m \otimes y_m - \langle y_m, y_i \rangle y_i \otimes y_m)$. So, because we have that $\|y_m \otimes y_m - \langle y_m, y_i \rangle y_i \otimes y_m\| = \|\mathbf{P}_{y_i} y_m \otimes y_m\| = \|\mathbf{P}_{y_i} y_m\| \|\mathbf{V}_i\| \leq 1$,

$$\|\mathbf{R}_1\| Z_j^{-1} \leq \beta Z_j \cdot Z_j^{-1} = \beta. \quad (79)$$

For $\mathbf{R}_2 Z_j^{-1}$, because $\| -y_i \otimes y_i + I_d \| = \sqrt{d-1}$, we have that

$$\|\mathbf{R}_2\| Z_j^{-1} \leq Z_j^{-1} e^{a_{ij}} (\beta + \sqrt{d-1}). \quad (80)$$

For $[(\mathbf{U}_1 + \mathbf{U}_2) \otimes \mathbf{V}_j] \cdot Z_j^{-2}$, we see that $\|\mathbf{V}_j\| \leq \sum_{m=1}^n e^{\beta \langle y_j, y_m \rangle} = Z_j$. Also, $\|\mathbf{U}_1\| Z_j^{-1} \leq \delta_{ij} \beta \|\mathbf{V}_j\| Z_j^{-1} \leq \delta_{ij} \beta$, $\|\mathbf{U}_2\| Z_j^{-1} \leq \beta e^{a_{ij}} Z_j^{-1}$. Hence, we have that

$$\|(\mathbf{U}_1 + \mathbf{U}_2) \otimes \mathbf{V}_j\| \cdot Z_j^{-2} \leq \beta (\delta_{ij} + e^{a_{ij}} Z_j^{-1}). \quad (81)$$

Proof 13 (Proof of Theorem 2.5) Theorem 2.5 follows directly from Lemma B.4 and Lemma B.5.

B.2 PROOF OF THEOREM 2.4

The proof for Theorem 2.4 requires more delicate arguments. The part when $\gamma > \frac{1}{1-\rho}$ in Theorem 2.4 directly follows from Lemma B.4, so we only focus on the part when $\gamma \leq \frac{1}{1-\rho}$. We remark that when $\gamma < \frac{1}{1-\rho}$, our result is that $\frac{1}{nd} \|\nabla_X X'\|^2 = 0 + o_n(1)$, which is a better estimate than Equation (22) in Theorem 2.5.

We first have the following lemma which replaces Lemma B.3 when we adopt Assumption 1.

Lemma B.6 *Adopt Assumption 1 and Equation (5). For any $i, j \in \llbracket 1, n \rrbracket$, consider the $d \times d$ matrix formed by $\frac{\partial(\text{ATT}(N(x_j)))_v}{\partial(x_i)_u}$, for $u, v \in \llbracket 1, d \rrbracket$. Denote $y_k = N(x_k)$ for each $k \in \llbracket 1, n \rrbracket$. Then, this matrix has the following form:*

$$\left(\frac{\partial(\text{ATT}(N(x_j)))_v}{\partial(x_i)_u} \right)_{d \times d} = q^{-\frac{1}{2}} [(\mathbf{R}_1 + \mathbf{R}_2)Z - (\mathbf{U}_1 + \mathbf{U}_2) \otimes (\mathbf{U}_3 + \mathbf{U}_4)] \cdot Z^{-2}, \quad (82)$$

where $Z = e^\beta + (n-1)e^{\beta\rho}$,

$$\mathbf{R}_1 := \delta_{ij}\beta e^{\beta\rho}(\mathbf{W} - y_i \otimes (\mathbf{W}y_i)), \quad \mathbf{R}_2 := e^{\beta\langle y_j, y_i \rangle}((-y_i + \beta\mathbf{P}_{y_i}y_j) \otimes y_i + I_d), \quad (83)$$

and

$$\begin{aligned} \mathbf{U}_1 &:= \delta_{ij}\beta e^{\beta\rho}(\mathbf{P}_{y_i}\mathbf{V}), \quad \mathbf{U}_2 := \beta e^{\beta\langle y_j, y_i \rangle}(\mathbf{P}_{y_i}y_j), \\ \mathbf{U}_3 &:= (e^\beta - e^{\beta\rho})y_j, \quad \mathbf{U}_4 := e^{\beta\rho}\mathbf{V}. \end{aligned} \quad (84)$$

In Equation (83) and Equation (84),

$$\mathbf{V} := \sum_{m=1}^n y_m, \quad \mathbf{W} := \sum_{m=1}^n y_m \otimes y_m, \quad \mathbf{P}_x y := y - \langle y, x \rangle x. \quad (85)$$

Proof 14 (Proof of Lemma B.6) We first apply Lemma B.3 to get Equation (58). After replacing $\langle y_j, y_m \rangle = \rho$ for $m \neq j$, we can obtain Equation (82). The only remark is that the term $\delta_{ij}(\mathbf{W}_j - y_i \otimes (\mathbf{W}_j y_i))$ in \mathbf{R}_1 of Equation (59) is nonzero when $i = j$. Then, when $i = j$, $\mathbf{W}_i - y_i \otimes (\mathbf{W}_i y_i) = \sum_{m=1}^n e^{\beta\langle y_i, y_m \rangle}(y_m \otimes y_m - y_i \otimes y_m \langle y_m, y_i \rangle)$. If $m = i$, the summand $(y_m \otimes y_m - y_i \otimes y_m \langle y_m, y_i \rangle)$ becomes 0. Hence, $\mathbf{W}_i - y_i \otimes (\mathbf{W}_i y_i) = \sum_{m \neq i} e^{\beta\langle y_i, y_m \rangle}(y_m \otimes y_m - y_i \otimes y_m \langle y_m, y_i \rangle) = e^{\beta\rho} \sum_{m \neq i} (y_m \otimes y_m - y_i \otimes y_m \langle y_m, y_i \rangle) = e^{\beta\rho}(\mathbf{W} - y_i \otimes (\mathbf{W}y_i))$.

Next, to compute the matrix norm of Equation (82), we see that for any matrix K , its matrix norm square equals to $\text{Tr}(K^T K)$. Hence, the matrix norm square of Equation (82) equals to

$$\begin{aligned} & q^{-1}Z^{-4} \cdot \left(\text{Tr}[Z^2(\mathbf{R}_1 + \mathbf{R}_2)^T(\mathbf{R}_1 + \mathbf{R}_2)] - 2Z(\mathbf{U}_1 + \mathbf{U}_2)^T(\mathbf{R}_1 + \mathbf{R}_2)(\mathbf{U}_3 + \mathbf{U}_4) \right. \\ & \left. + \|\mathbf{U}_1 + \mathbf{U}_2\|^2 \|\mathbf{U}_3 + \mathbf{U}_4\|^2 \right). \end{aligned} \quad (86)$$

We then compute these terms separately, and sum them in i, j . We first have the following basic equalities for the notations \mathbf{V}, \mathbf{W} in Equation (85).

Lemma B.7 *For the notations in Equation (85), i.e.,*

$$\mathbf{V} := \sum_{m=1}^n y_m, \quad \mathbf{W} := \sum_{m=1}^n y_m \otimes y_m, \quad \mathbf{P}_x y := y - \langle y, x \rangle x, \quad (87)$$

we have that

$$\begin{aligned} \text{Tr}(\mathbf{W}^2) &= \sum_{m,l} \langle y_m, y_l \rangle^2 = n(n\rho^2 + (1-\rho^2)), \\ \text{Tr}(\mathbf{W}) &= n, \quad \text{Tr}(\mathbf{W}y_i y_i^T) = n\rho^2 + (1-\rho^2), \quad \|\mathbf{P}_{y_i}y_j\|^2 = 1 - \rho^2. \end{aligned} \quad (88)$$

Also,

$$\begin{aligned} \mathbf{W}y_i &= \sum_{m=1}^n \langle y_m, y_i \rangle y_m = (1-\rho)y_i + \rho\mathbf{V}, \\ \langle \mathbf{V}, y_i \rangle &= \sum_{m=1}^n \langle y_m, y_i \rangle = n\rho + (1-\rho), \\ \|\mathbf{V}\|^2 &= \sum_{m,l} \langle y_m, y_l \rangle = n + \rho n(n-1) = n(n\rho + (1-\rho)), \\ \|\mathbf{P}_{y_i}\mathbf{V}\|^2 &= \|\mathbf{V}\|^2 - \langle \mathbf{V}, y_i \rangle^2 = (n-1)(n\rho + (1-\rho))(1-\rho), \\ \|\mathbf{W}y_i\|^2 &= n^2\rho^3 + 3n\rho^2(1-\rho) + (1+2\rho)(1-\rho)^2. \end{aligned} \quad (89)$$

Proof 15 (Proof of Lemma B.7) *Direct Computations.*

Lemma B.8 *For terms $\mathbf{R}_1, \mathbf{R}_2$ in Lemma B.6, we have that*

$$\begin{aligned} \sum_{i,j} \text{Tr} [(\mathbf{R}_1 + \mathbf{R}_2)^T (\mathbf{R}_1 + \mathbf{R}_2)] \\ = \beta^2 e^{2\beta\rho} n [n^2 \rho^2 (1 - \rho) + n(1 - \rho)(1 + \rho - 3\rho^2) - (1 + 2\rho)(1 - \rho)^2] \\ + \beta e^{\beta(\rho+1)} n(n-1)(1 - \rho^2) \\ + e^{2\beta} (d-1)n + e^{2\beta\rho} [\beta^2(1 - \rho^2) + d - 1] n(n-1). \end{aligned} \quad (90)$$

As a corollary, when we pick $\beta = \gamma \log n$, we have the following phase transition limits as $n \rightarrow +\infty$:

$$\frac{1}{nZ^2} \sum_{i,j} \text{Tr} [(\mathbf{R}_1 + \mathbf{R}_2)^T (\mathbf{R}_1 + \mathbf{R}_2)] = \begin{cases} \beta^2 \rho^2 (1 - \rho) + o_n(1) & \text{if } \gamma < \frac{1}{1-\rho}, \\ \frac{d-1+\beta^2 \rho^2 (1-\rho)}{4} + o_n(1) & \text{if } \gamma = \frac{1}{1-\rho}, \\ d-1 + o_n(1) & \text{if } \gamma > \frac{1}{1-\rho}. \end{cases} \quad (91)$$

Proof 16 (Proof of Lemma B.8) *We first notice that \mathbf{W} is a symmetric matrix and $\|y_i\| = 1$. We then expand each term in Lemma B.8 and use Lemma B.7.*

$$\begin{aligned} \sum_{i,j} \text{Tr} [(\mathbf{R}_1)^T \mathbf{R}_1] &= \beta^2 e^{2\beta\rho} \sum_i (\text{Tr} (\mathbf{W}^2 - 2\mathbf{W}y_i(\mathbf{W}y_i)^T) + \|y_i\|^2 \|\mathbf{W}y_i\|^2) \\ &= \beta^2 e^{2\beta\rho} \sum_i (\text{Tr} \mathbf{W}^2 - 2\|\mathbf{W}y_i\|^2 + \|\mathbf{W}y_i\|^2) \\ &= \beta^2 e^{2\beta\rho} n [n^2 \rho^2 (1 - \rho) + n(1 - \rho)(1 + \rho - 3\rho^2) - (1 + 2\rho)(1 - \rho)^2]. \end{aligned} \quad (92)$$

Then,

$$\begin{aligned} \sum_{i,j} \text{Tr} [(\mathbf{R}_1)^T \mathbf{R}_2] &= \text{Tr} \sum_{i,j} \delta_{ij} \beta e^{\beta\rho} (\mathbf{W} - \mathbf{W}y_i y_i^T) e^{\beta\langle y_j, y_i \rangle} ((-y_i + \beta \mathbf{P}_{y_i} y_j) \otimes y_i + I_d) \\ &= \beta e^{\beta(\rho+1)} \text{Tr} \sum_i (\mathbf{W} - \mathbf{W}y_i y_i^T) (-y_i y_i^T + I_d) = \beta e^{\beta(\rho+1)} \text{Tr} \sum_i (\mathbf{W} - \mathbf{W}y_i y_i^T) \\ &= \beta e^{\beta(\rho+1)} n(n-1)(1 - \rho^2), \end{aligned} \quad (93)$$

where the second equality is because $\mathbf{P}_{y_i} y_i = 0$.

$$\begin{aligned} \sum_{i,j} \text{Tr} [(\mathbf{R}_2)^T \mathbf{R}_2] &= \sum_{i,j} e^{2\beta\langle y_j, y_i \rangle} \text{Tr} [(-y_i + \beta \mathbf{P}_{y_i} y_j) y_i^T + I_d] (y_i (-y_i + \beta \mathbf{P}_{y_i} y_j)^T + I_d) \\ &= \sum_{i \neq j} e^{2\beta\rho} [(1 + \beta^2(1 - \rho^2)) - 2 + d] + \sum_i e^{2\beta} (d-1) \\ &= e^{2\beta} (d-1)n + e^{2\beta\rho} [\beta^2(1 - \rho^2) + d - 1] n(n-1). \end{aligned} \quad (94)$$

Next, we show the asymptotics Equation (91) as $n \rightarrow +\infty$. According to Lemma A.1, we have that

$$Z = \begin{cases} (1 + o_n(1)) \cdot n e^{\beta\rho} & \text{if } \gamma < \frac{1}{1-\rho}, \\ (1 + o_n(1)) \cdot e^\beta & \text{if } \gamma > \frac{1}{1-\rho}. \end{cases} \quad (95)$$

That is, when $\gamma < \frac{1}{1-\rho}$, the leading order terms are those terms involving $n e^{\beta\rho}$, and all the remaining terms go to 0 after dividing $n e^{\beta\rho}$; when $\gamma > \frac{1}{1-\rho}$, the leading order terms are those terms involving e^β , and all the remaining terms go to 0 after dividing e^β . Hence, when $\gamma < \frac{1}{1-\rho}$, the leading order term in Equation (90) is the term $\beta^2 e^{2\beta\rho} n^3 \rho^2 (1 - \rho)$; when $\gamma > \frac{1}{1-\rho}$, the leading order term is $e^{2\beta} (d-1)n$. This proves Equation (91).

Lemma B.9 For terms $\mathbf{R}_1, \mathbf{R}_2, \mathbf{U}_1, \mathbf{U}_2, \mathbf{U}_3, \mathbf{U}_4$ in Lemma B.6, we have that

$$\begin{aligned} & \sum_{i,j} (\mathbf{U}_1 + \mathbf{U}_2)^T (\mathbf{R}_1 + \mathbf{R}_2) (\mathbf{U}_3 + \mathbf{U}_4) \\ &= \rho \beta^2 e^{2\beta\rho} (e^\beta - e^{\beta\rho}) n(n-1)(n\rho + (1-\rho))(1-\rho) \\ & \quad + \beta^2 e^{3\beta\rho} n(n-1)(n\rho + (1-\rho))^2 (1-\rho) \\ & \quad + \beta e^{\beta(2\rho+1)} n(n-1)(n\rho + (1-\rho))(1-\rho) \\ & \quad + \beta e^{2\beta\rho} (e^\beta - e^{\beta\rho}) (\beta\rho + 1) n(n-1)(1-\rho^2) \\ & \quad + \beta e^{3\beta\rho} n(n-1)(n\rho + (1-\rho))(\beta(1-\rho^2) + (1-\rho)). \end{aligned} \quad (96)$$

As a corollary, when we pick $\beta = \gamma \log n$, we have the following phase transition limits as $n \rightarrow +\infty$:

$$\frac{1}{nZ^3} \sum_{i,j} (\mathbf{U}_1 + \mathbf{U}_2)^T (\mathbf{R}_1 + \mathbf{R}_2) (\mathbf{U}_3 + \mathbf{U}_4) = \begin{cases} \beta^2 \rho^2 (1-\rho) + o_n(1) & \text{if } \gamma < \frac{1}{1-\rho}, \\ \frac{\beta^2 \rho^2 (1-\rho)}{4} + o_n(1) & \text{if } \gamma = \frac{1}{1-\rho}, \\ 0 + o_n(1) & \text{if } \gamma > \frac{1}{1-\rho}. \end{cases} \quad (97)$$

Proof 17 (Proof of Lemma B.9) We expand each term in Lemma B.9 and also apply Lemma B.7 to each term. We first estimate terms involving \mathbf{U}_1 .

$$\begin{aligned} \sum_{i,j} \mathbf{U}_1^T \mathbf{R}_1 \mathbf{U}_3 &= \sum_i \beta^2 e^{2\beta\rho} (e^\beta - e^{\beta\rho}) (\mathbf{P}_{y_i} \mathbf{V})^T (\mathbf{W} - y_i \otimes (\mathbf{W} y_i)) y_i \\ &= \beta^2 e^{2\beta\rho} (e^\beta - e^{\beta\rho}) \sum_i (\mathbf{P}_{y_i} \mathbf{V})^T \mathbf{W} y_i = \rho \beta^2 e^{2\beta\rho} (e^\beta - e^{\beta\rho}) \sum_i (\mathbf{P}_{y_i} \mathbf{V})^T \mathbf{V} \\ &= \rho \beta^2 e^{2\beta\rho} (e^\beta - e^{\beta\rho}) n(n-1)(n\rho + (1-\rho))(1-\rho), \end{aligned} \quad (98)$$

where the second and the third equality is because $\langle \mathbf{P}_{y_i} \mathbf{V}, y_i \rangle = 0$.

$$\begin{aligned} \sum_{i,j} \mathbf{U}_1^T \mathbf{R}_1 \mathbf{U}_4 &= \sum_i \beta^2 e^{3\beta\rho} (\mathbf{P}_{y_i} \mathbf{V})^T (\mathbf{W} - y_i \otimes (\mathbf{W} y_i)) \mathbf{V} \\ &= \beta^2 e^{3\beta\rho} \sum_i (\mathbf{P}_{y_i} \mathbf{V})^T \mathbf{W} \mathbf{V} = \beta^2 e^{3\beta\rho} (n\rho + (1-\rho)) \sum_i \|\mathbf{P}_{y_i} \mathbf{V}\|^2 \\ &= \beta^2 e^{3\beta\rho} n(n-1)(n\rho + (1-\rho))^2 (1-\rho), \end{aligned} \quad (99)$$

where the second equality is because $\langle \mathbf{P}_{y_i} \mathbf{V}, y_i \rangle = 0$.

$$\sum_{i,j} \mathbf{U}_1^T \mathbf{R}_2 \mathbf{U}_3 = \sum_i \beta e^{\beta(\rho+1)} (e^\beta - e^{\beta\rho}) (\mathbf{P}_{y_i} \mathbf{V})^T ((-y_i + \beta \mathbf{P}_{y_i} y_i) \otimes y_i + I_d) y_i = 0, \quad (100)$$

where the second equality is because $\langle \mathbf{P}_{y_i} \mathbf{V}, y_i \rangle = 0$ and $\mathbf{P}_{y_i} y_i = 0$.

$$\begin{aligned} \sum_{i,j} \mathbf{U}_1^T \mathbf{R}_2 \mathbf{U}_4 &= \sum_i \beta e^{\beta(2\rho+1)} (\mathbf{P}_{y_i} \mathbf{V})^T ((y_i + \beta \mathbf{P}_{y_i} y_i) \otimes y_i + I_d) \mathbf{V} \\ &= \beta e^{\beta(2\rho+1)} \sum_i \|\mathbf{P}_{y_i} \mathbf{V}\|^2 = \beta e^{\beta(2\rho+1)} n(n-1)(n\rho + (1-\rho))(1-\rho), \end{aligned} \quad (101)$$

where the second equality is because $\langle \mathbf{P}_{y_i} \mathbf{V}, y_i \rangle = 0$ and $\mathbf{P}_{y_i} y_i = 0$.

Next, we estimate the terms involving \mathbf{U}_2 . We first recall that $\mathbf{U}_2 = \beta e^{\beta \langle y_j, y_i \rangle} (\mathbf{P}_{y_i} y_j)$. Because $\mathbf{P}_{y_i} y_j = 0$ when $i = j$, we can just replace $e^{\beta \langle y_j, y_i \rangle}$ with $e^{\beta\rho}$ in \mathbf{U}_2 , i.e., $\mathbf{U}_2 = \beta e^{\beta\rho} (\mathbf{P}_{y_i} y_j)$. Hence,

$$\sum_{i,j} \mathbf{U}_2^T \mathbf{R}_1 \mathbf{U}_3 = 0, \quad \sum_{i,j} \mathbf{U}_2^T \mathbf{R}_1 \mathbf{U}_4 = 0, \quad (102)$$

because $\delta_{ij}(\mathbf{P}_{y_i} y_j) = 0$ for any i, j in $\mathbf{U}_2^T \mathbf{R}_1$.

$$\begin{aligned}
\sum_{i,j} \mathbf{U}_2^T \mathbf{R}_2 \mathbf{U}_3 &= \sum_{i,j} \beta e^{\beta(\rho + \langle y_j, y_i \rangle)} (e^\beta - e^{\beta\rho}) (\mathbf{P}_{y_i} y_j)^T ((-y_i + \beta \mathbf{P}_{y_i} y_j) \otimes y_i + I_d) y_j \\
&= \beta e^{2\beta\rho} (e^\beta - e^{\beta\rho}) \sum_{i \neq j} (\mathbf{P}_{y_i} y_j)^T ((-y_i + \beta \mathbf{P}_{y_i} y_j) \rho + y_j) \\
&= \beta e^{2\beta\rho} (e^\beta - e^{\beta\rho}) (\beta\rho + 1) \sum_{i \neq j} \|\mathbf{P}_{y_i} y_j\|^2 \\
&= \beta e^{2\beta\rho} (e^\beta - e^{\beta\rho}) (\beta\rho + 1) n(n-1)(1 - \rho^2).
\end{aligned} \tag{103}$$

where the second equality is because $\mathbf{P}_{y_i} y_j \neq 0$ only when $i \neq j$, on which $\langle y_j, y_i \rangle = \rho$, and the third equality is because $\langle \mathbf{P}_{y_i} y_j, y_i \rangle = 0$.

$$\begin{aligned}
\sum_{i,j} \mathbf{U}_2^T \mathbf{R}_2 \mathbf{U}_4 &= \sum_{i,j} \beta e^{\beta(2\rho + \langle y_j, y_i \rangle)} (\mathbf{P}_{y_i} y_j)^T ((-y_i + \beta \mathbf{P}_{y_i} y_j) \otimes y_i + I_d) \mathbf{V} \\
&= \beta e^{3\beta\rho} \sum_{i \neq j} \left(\beta \|\mathbf{P}_{y_i} y_j\|^2 (n\rho + (1 - \rho)) + (\mathbf{P}_{y_i} y_j)^T \mathbf{V} \right) \\
&= \beta e^{3\beta\rho} \sum_{i \neq j} (\beta(1 - \rho^2)(n\rho + (1 - \rho)) + (1 - \rho)(n\rho + (1 - \rho))) \\
&= \beta e^{3\beta\rho} n(n-1)(n\rho + (1 - \rho))(\beta(1 - \rho^2) + (1 - \rho)).
\end{aligned} \tag{104}$$

where the second equality is because $\mathbf{P}_{y_i} y_j \neq 0$ only when $i \neq j$, on which $\langle y_j, y_i \rangle = \rho$, and the third equality is because $\langle \mathbf{P}_{y_i} y_j, y_i \rangle = 0$.

The proof for Equation (97) is similar to the proof for Equation (91) in Lemma B.8. Notice that when $\gamma < \frac{1}{1-\rho}$, we need to pick up terms involving $ne^{\beta\rho}$, and the leading order term in Equation (96) is the one in the second line of Equation (96), which is $\beta^2 n^4 e^{3\beta\rho} \rho^2 (1 - \rho)$; when $\gamma > \frac{1}{1-\rho}$, after diving nZ^3 , all terms in Equation (96) are $o_n(1)$ terms.

Lemma B.10 For terms $\mathbf{U}_1, \mathbf{U}_2, \mathbf{U}_3, \mathbf{U}_4$ in Lemma B.6, we have that

$$\begin{aligned}
&\sum_{i,j} \|\mathbf{U}_1 + \mathbf{U}_2\|^2 \|\mathbf{U}_3 + \mathbf{U}_4\|^2 \\
&= \beta^2 e^{2\beta\rho} n(n-1)(n\rho + 2)(1 - \rho) \\
&\quad \cdot [(e^\beta - e^{\beta\rho})^2 + 2e^{\beta\rho}(e^\beta - e^{\beta\rho})(n\rho + (1 - \rho)) + e^{2\beta\rho} n(n\rho + (1 - \rho))].
\end{aligned} \tag{105}$$

As a corollary, when we pick $\beta = \gamma \log n$, we have the following phase transition limits as $n \rightarrow +\infty$:

$$\frac{1}{nZ^4} \sum_{i,j} \|\mathbf{U}_1 + \mathbf{U}_2\|^2 \|\mathbf{U}_3 + \mathbf{U}_4\|^2 = \begin{cases} \beta^2 \rho^2 (1 - \rho) + o_n(1) & \text{if } \gamma < \frac{1}{1-\rho}, \\ \frac{\beta^2 \rho (1-\rho)(1+3\rho)}{16} + o_n(1) & \text{if } \gamma = \frac{1}{1-\rho}, \\ 0 + o_n(1) & \text{if } \gamma > \frac{1}{1-\rho}. \end{cases} \tag{106}$$

Proof 18 (Proof of Lemma B.10) We notice that $\langle \mathbf{U}_1, \mathbf{U}_2 \rangle = 0$ because $\delta_{ij} \mathbf{P}_{y_i} y_j = 0$ for any i, j . So,

$$\begin{aligned}
\|\mathbf{U}_1 + \mathbf{U}_2\|^2 &= \delta_{ij} \beta^2 e^{2\beta\rho} \|\mathbf{P}_{y_i} \mathbf{V}\|^2 + \beta^2 e^{2\beta\langle y_j, y_i \rangle} \|\mathbf{P}_{y_i} y_j\|^2 \\
&= \delta_{ij} \beta^2 e^{2\beta\rho} (n-1)(n\rho + (1 - \rho))(1 - \rho) + (1 - \delta_{ij}) \beta^2 e^{2\beta\rho} (1 - \rho^2),
\end{aligned} \tag{107}$$

where the second equality is because $e^{2\beta\langle y_j, y_i \rangle} \|\mathbf{P}_{y_i} y_j\|^2 \neq 0$ only if $i \neq j$, on which $e^{2\beta\langle y_j, y_i \rangle} \|\mathbf{P}_{y_i} y_j\|^2 = e^{2\beta\rho} (1 - \rho^2)$.

$$\begin{aligned}
\|\mathbf{U}_3 + \mathbf{U}_4\|^2 &= (e^\beta - e^{\beta\rho})^2 + 2e^{\beta\rho}(e^\beta - e^{\beta\rho}) \langle \mathbf{V}, y_j \rangle + e^{2\beta\rho} \|\mathbf{V}\|^2 \\
&= (e^\beta - e^{\beta\rho})^2 + 2e^{\beta\rho}(e^\beta - e^{\beta\rho})(n\rho + (1 - \rho)) + e^{2\beta\rho} n(n\rho + (1 - \rho)),
\end{aligned} \tag{108}$$

which is independent of i, j . Hence,

$$\begin{aligned}
& \sum_{i,j} \|\mathbf{U}_1 + \mathbf{U}_2\|^2 \|\mathbf{U}_3 + \mathbf{U}_4\|^2 \\
&= [\beta^2 e^{2\beta\rho} n(n-1)(n\rho + (1-\rho))(1-\rho) + n(n-1)\beta^2 e^{2\beta\rho}(1-\rho^2)] \|\mathbf{U}_3 + \mathbf{U}_4\|^2 \\
&= \beta^2 e^{2\beta\rho} n(n-1)(n\rho + 2)(1-\rho) \|\mathbf{U}_3 + \mathbf{U}_4\|^2 \\
&= \beta^2 e^{2\beta\rho} n(n-1)(n\rho + 2)(1-\rho) \\
&\quad \cdot [(e^\beta - e^{\beta\rho})^2 + 2e^{\beta\rho}(e^\beta - e^{\beta\rho})(n\rho + (1-\rho)) + e^{2\beta\rho}n(n\rho + (1-\rho))] .
\end{aligned} \tag{109}$$

The proof for Equation (106) is similar to the proof for Equation (91) in Lemma B.8. Notice that when $\gamma < \frac{1}{1-\rho}$, we need to pick up terms involving $ne^{\beta\rho}$, and the leading order term in Equation (105) is $\beta^2 n^5 e^{4\beta\rho} \rho^2 (1-\rho)$; when $\gamma > \frac{1}{1-\rho}$, after diving nZ^4 , all terms in Equation (96) are $o_n(1)$ terms.

Proof 19 (Proof of Theorem 2.4) As we have mentioned at the beginning of Appendix B.2, we only need to focus the case when $\gamma \leq \frac{1}{1-\rho}$, which follows directly from Lemma B.8, Lemma B.9, and Lemma B.10. We notice that, in these three lemmas, the leading order terms are the same, $\beta^2 \rho^2 (1-\rho)$, which cancels in Equation (86). Hence, when $\gamma < \frac{1}{1-\rho}$, $\frac{1}{nd} \|\nabla_X X'\|^2 = 0 + o_n(1)$. When $\gamma = \frac{1}{1-\rho}$, we also only need to use the corresponding cases in these three lemmas and combine them in Equation (86) to get the conclusion in Theorem 2.4. One remark is that under Assumption 1, we have that $n \leq d$ implicitly. So, when $\gamma = \frac{1}{1-\rho}$, terms in Equation (86) involving $\frac{\beta^2}{d} = \frac{\gamma^2 (\log n)^2}{d}$ also become $o_n(1)$.

C MODIFIED ASSUMPTIONS WITH MORE MEDIAN PHASES

In this section, we modify Assumption 2, so that we can prove the existence of three different phases like Lemma A.1, Theorem 2.3, Theorem 2.5. We remark that we only showed the existence of two phases (two extrema) in Lemma A.1, Theorem 2.3, Theorem 2.5, but it doesn't mean under Assumption 2, there is no other transition phase between these two phases (two extrema). Under the following Assumption 3, we can show there are indeed at least three phases. Recall that for any $i \in \llbracket 1, n \rrbracket$, we defined $y_i = N(x_i)$.

Assumption 3

- For any $i \in \llbracket 1, n \rrbracket$, $\|x_i\|^2 \in [q_1, q_2]$ for some positive constants $q_1 \leq q_2$.
- There is a $\tau \in (0, 1]$, four positive constants $\rho_3, \rho_4, \kappa_3, \kappa_4$ with $\rho_3 \leq \rho_4$, $\kappa_3 \leq \kappa_4$, and $\rho_4 < 1$, such that for any $i \in \llbracket 1, n \rrbracket$, if we define

$$\mathcal{K}_i = \{m \neq i \mid \langle y_m, y_i \rangle \in [\rho_3, \rho_4]\}, \tag{110}$$

then we have that

$$\kappa_3 \leq \frac{|\mathcal{K}_i|}{n^\tau} \leq \kappa_4. \tag{111}$$

- For any $i \in \llbracket 1, n \rrbracket$ and any $j \notin \mathcal{K}_i \cup \{i\}$, $\langle y_i, y_j \rangle \in [\rho_1, \rho_2]$ for some nonnegative constants ρ_1, ρ_2 satisfying $\rho_1 \leq \rho_2 < \rho_3 \leq \rho_4$.
- For technical reason, we further assume that $(1-\tau)(1-\rho_2) + \rho_2 < \rho_3$.

Lemma C.1 Let $\beta = \gamma \log n$ where γ is a positive constant. Under Assumption 2 and Equation (5), for any $i \in \llbracket 1, n \rrbracket$,

$$Z_i = \begin{cases} (1 + o_n(1)) \cdot \left(\sum_{m \notin \mathcal{K}_i \cup \{i\}} e^{a_{im}} \right) & \text{if } \gamma < \min \left\{ \frac{1}{1-\rho_1}, \frac{1-\tau}{\rho_4-\rho_1} \right\}, \\ (1 + o_n(1)) \cdot \left(\sum_{m \in \mathcal{K}_i} e^{a_{im}} \right) & \text{if } \frac{1-\tau}{\rho_3-\rho_2} < \gamma < \frac{\tau}{1-\rho_3}, \\ (1 + o_n(1)) \cdot e^\beta & \text{if } \gamma > \max \left\{ \frac{1}{1-\rho_2}, \frac{\tau}{1-\rho_4} \right\}, \end{cases} \tag{112}$$

where the terms $o_n(1)$ go to 0 as $n \rightarrow +\infty$ with speeds independent of i but only depending on $\gamma, \rho_1, \rho_2, \rho_3, \rho_4, \tau, \kappa_3, \kappa_4$.

Proof 20 The proof is similar to Lemma A.1. We notice that

$$\begin{aligned} Z_i &= e^\beta + \sum_{m \in \mathcal{K}_i} e^{a_{im}} + \sum_{m \notin \mathcal{K}_i \cup \{i\}} e^{a_{im}} \\ &= n^\gamma + \sum_{m \in \mathcal{K}_i} n^{\gamma \langle y_i, y_m \rangle} + \sum_{m \notin \mathcal{K}_i \cup \{i\}} n^{\gamma \langle y_i, y_m \rangle}. \end{aligned} \quad (113)$$

We also notice that $\kappa_3 n^\tau \leq |\mathcal{K}_i| \leq \kappa_4 n^\tau$ according to Assumption 3. Hence,

$$\kappa_3 n^{\tau + \gamma \rho_3} \leq |\mathcal{K}_i| \cdot n^{\gamma \rho_3} \leq \sum_{m \in \mathcal{K}_i} n^{\gamma \langle y_i, y_m \rangle} \leq |\mathcal{K}_i| \cdot n^{\gamma \rho_4} \leq \kappa_4 n^{\tau + \gamma \rho_4}, \quad (114)$$

and

$$(n - \kappa_4 n^\tau - 1) \cdot n^{\gamma \rho_1} \leq \sum_{m \notin \mathcal{K}_i \cup \{i\}} n^{\gamma \langle y_i, y_m \rangle} \leq (n - |\mathcal{K}_i| - 1) \cdot n^{\gamma \rho_2} \leq n^{1 + \gamma \rho_2}. \quad (115)$$

When $\gamma < \min \left\{ \frac{1}{1-\rho_1}, \frac{1-\tau}{\rho_4-\rho_1} \right\}$, the leading order term in Z_i is $\sum_{m \notin \mathcal{K}_i \cup \{i\}} n^{\gamma \langle y_i, y_m \rangle}$; when $\frac{1-\tau}{\rho_3-\rho_2} < \gamma < \frac{\tau}{1-\rho_3}$, the leading order term in Z_i is $\sum_{m \in \mathcal{K}_i} n^{\gamma \langle y_i, y_m \rangle}$; when $\gamma > \max \left\{ \frac{1}{1-\rho_2}, \frac{\tau}{1-\rho_4} \right\}$, the leading order term in Z_i is n^γ . We also remark that the last assumption in Assumption 3 is to ensure the existence of the middle phase, i.e., $\frac{1-\tau}{\rho_3-\rho_2} < \gamma < \frac{\tau}{1-\rho_3}$. This finishes the proof for Lemma C.1 by similar arguments as in Lemma A.1.

A direct corollary of Lemma C.1 is the following theorem.

Theorem C.2 Under Assumption 2 and Equation (5) we have the following phase transition phenomena: let $\beta = \gamma \log n$ where γ is a positive constant. For any $i \in \llbracket 1, n \rrbracket$ the updating dynamics Equation (5) can be written as

$$x'_i = \alpha x_i + \begin{cases} \frac{\sum_{m \notin \mathcal{K}_i \cup \{i\}} e^{a_{im}} y_m}{\sum_{m \notin \mathcal{K}_i \cup \{i\}} e^{a_{im}}} + \mathbf{o}_n(1) & \text{if } \gamma < \min \left\{ \frac{1}{1-\rho_1}, \frac{1-\tau}{\rho_4-\rho_1} \right\}, \\ \frac{\sum_{m \in \mathcal{K}_i} e^{a_{im}} y_m}{\sum_{m \in \mathcal{K}_i} e^{a_{im}}} + \mathbf{o}_n(1) & \text{if } \frac{1-\tau}{\rho_3-\rho_2} < \gamma < \frac{\tau}{1-\rho_3}, \\ y_i + \mathbf{o}_n(1) & \text{if } \gamma > \max \left\{ \frac{1}{1-\rho_2}, \frac{\tau}{1-\rho_4} \right\}, \end{cases} \quad (116)$$

The terms $\mathbf{o}_n(1)$ represent vectors in \mathbb{R}^d with norms going to 0 as $n \rightarrow +\infty$, with a speed independent of i but only depending on $\gamma, \rho_1, \rho_2, \rho_3, \rho_4, \tau, \kappa_3, \kappa_4$.

The proof of Theorem C.2 is similar to Lemma C.1 so we omit its proof.

D ANALYSIS OF THE $\beta_n \asymp \sqrt{\log n}$ SCALING FOR I.I.D. GAUSSIAN SCORES

This appendix provides a short heuristic derivation of the $\beta_n \asymp \sqrt{\log n}$ scaling for softmax attention when the raw attention scores a_1, \dots, a_n are modeled as independent $\mathcal{N}(0, 1)$ random variables. The purpose is to contrast the behavior of this Gaussian setting with the geometric setting analyzed in the main text, where pairwise score gaps remain $O(1)$ and the critical scale becomes $\beta_n \asymp \log n$.

Let a_1, \dots, a_n be i.i.d. $\mathcal{N}(0, 1)$ and denote by

$$a_1^\downarrow \geq a_2^\downarrow \geq \dots \geq a_n^\downarrow$$

their order statistics. Set

$$t_n := \Phi^{-1} \left(1 - \frac{1}{n} \right) \sim \sqrt{2 \log n},$$

where Φ is the standard normal CDF. It is classical (see Theorem 2.1.1 in De Haan & Ferreira (2006)) that for any fixed $k \geq 2$,

$$(t_n(a_1^\downarrow - a_2^\downarrow), \dots, t_n(a_1^\downarrow - a_k^\downarrow)) \Rightarrow (E_1, E_1 + E_2, \dots, E_1 + \dots + E_{k-1}), \quad (117)$$

where $(E_i)_{i \geq 1}$ are i.i.d. $\text{Exp}(1)$ random variables. Thus the gap scale between neighboring scores in the Gaussian model is $1/\sqrt{\log n}$.

Define the softmax weights and the top weight

$$A_j(n) := \frac{\exp(\beta_n a_j)}{\sum_{k=1}^n \exp(\beta_n a_k)}, \quad A_1^\downarrow(n) := \max_{1 \leq j \leq n} A_j(n).$$

Using Equation (117), one can write heuristically

$$A_1^\downarrow(n) \approx \frac{1}{1 + \sum_{k=2}^{\infty} \exp\{-(\beta_n/\sqrt{2 \log n}) S_{k-1}\}}, \quad S_{k-1} := E_1 + \dots + E_{k-1}. \quad (118)$$

where we have replaced the finite sum by an infinite series, which captures the leading asymptotics. The critical scaling $\sqrt{\log n}$ shows up in Equation (118). For example, in the supercritical regime $\beta_n \gg \sqrt{\log n}$, we have the following:

Proposition D.1 (Supercritical regime) *If*

$$\frac{\beta_n}{\sqrt{\log n}} \rightarrow \infty,$$

then $A_1^\downarrow(n) \rightarrow 1$ as $n \rightarrow \infty$. In other words, the attention weights concentrate on the top-scoring token.

Proof 21 (Sketch of proof) *In Equation (118), we know the sum*

$$Z = \sum_{k=2}^{\infty} \exp\{-(\beta_n/\sqrt{2 \log n}) S_{k-1}\}$$

is almost surely finite. Since $\beta_n/\sqrt{\log n} \rightarrow \infty$, every such term in Z vanishes in probability. Hence the denominator in Equation (118) tends to 1, and $A_1^\downarrow(n) \rightarrow 1$ as $n \rightarrow \infty$.

Proposition D.1 shows that the softmax enters a regime in which a single index captures asymptotically all attention mass as soon as $\beta_n \gg \sqrt{\log n}$. The critical scale is thus determined by the requirement

$$\beta_n(a_1^\downarrow - a_k^\downarrow) = O(k) \iff \beta_n \asymp \sqrt{\log n}.$$

In the geometric models analyzed in the main text (simplex and almost-simplex assumptions), the gaps between the top pairwise inner products remain of order $O(1)$ as $n \rightarrow \infty$. Consequently, the balancing argument in the Introduction section yields the critical scaling $\beta_n \asymp \log n$. This stands in clear contrast to REM-type Gaussian models, where the top gaps shrink at the $1/\sqrt{\log n}$ scale and thus produce the critical regime $\beta_n \asymp \sqrt{\log n}$. This explains why REM-type Gaussian models lead to the $\sqrt{\log n}$ scale, whereas geometric models naturally produce the $\log n$ scale that aligns with many practical strategies to avoid mixing in long-context attention.