

# Sigmoid Self-Attention has Lower Sample Complexity than Softmax Self-Attention: A Mixture-of-Experts Perspective

Fanqi Yan<sup>◊,\*</sup> Huy Nguyen<sup>◊,\*</sup> Pedram Akbarian<sup>◊</sup> Nhat Ho<sup>◊,†</sup> Alessandro Rinaldo <sup>◊,†</sup>

The University of Texas at Austin<sup>◊</sup>  
May 27, 2025

## Abstract

At the core of the popular Transformer architecture is the self-attention mechanism, which dynamically assigns softmax weights to each input token so that the model can focus on the most salient information. However, the softmax structure slows down the attention computation due to its row-wise nature, and it inherently introduces competition among tokens: as the weight assigned to one token increases, the weights of others decrease. This competitive dynamic may narrow the focus of self-attention to a limited set of features, potentially overlooking other informative characteristics. Recent experimental studies have shown that using the element-wise sigmoid function helps eliminate token competition and reduce the computational overhead. Despite these promising empirical results, a rigorous comparison between sigmoid and softmax self-attention mechanisms remains absent in the literature. This paper closes this gap by theoretically demonstrating that sigmoid self-attention is more sample-efficient than its softmax counterpart. Toward that goal, we represent the self-attention matrix as a mixture of experts and show that “experts” in sigmoid self-attention require significantly less data to achieve the same approximation error as those in softmax self-attention.

## 1 Introduction

Transformer models [47] have been known as the state-of-the-art architecture for a wide range of machine learning and deep learning applications, including language modeling [13, 2, 42, 45], computer vision [14, 3, 41, 31], and reinforcement learning [4, 28, 22], etc. One of the central components that contribute to the success of the Transformer models is the self-attention mechanism, which enables sequence-to-sequence models to concentrate on relevant parts of the input data. In particular, for each token in an input sequence, the self-attention mechanism computes a context vector formulated as a weighted sum of the tokens, where more relevant tokens to the context are assigned larger weights than others (see Section 2 for a formal definition). Therefore, self-attention is able to capture long-range dependencies and complex relationships within the data.

However, since the weights in the context vector are normalized by the softmax function, there might be an undesirable competition among the tokens, that is, an increase in the weight of a token leads to a decrease in the weights of others. As a consequence, the traditional softmax self-attention mechanism might focus only on a few aspects of the data and possibly ignore other informative features [43]. Additionally, Gu et al. [19] also discovered that the tokens’ inner dependence on the attention scores owing to the softmax normalization partly causes the attention sink phenomenon occurring

---

\* Equal contribution. † Co-last authors.

Table 1: Summary of the sample complexity of sigmoid/softmax self-attention under the MoE perspective, that is, the expert sample complexity to attain the approximation error  $\epsilon$  (ignoring logarithmic factors of lower order). In this work, we consider two types of expert functions, including expert networks with ReLU, GELU activations; and polynomial experts. Below,  $\tau$  denotes some positive constant.

	ReLU, GELU Experts		Polynomial Experts	
	Sparse Regime	Dense Regime	Sparse Regime	Dense Regime
SigmoidAttn (ours)	$\mathcal{O}(\epsilon^{-4})$	$\mathcal{O}(\epsilon^{-2})$	$\mathcal{O}(\exp(\epsilon^{-1/\tau}))$	$\mathcal{O}(\epsilon^{-2})$
SoftmaxAttn [1]	$\mathcal{O}(\epsilon^{-4})$	$\mathcal{O}(\epsilon^{-2})$	$\mathcal{O}(\exp(\epsilon^{-1/\tau}))$	$\mathcal{O}(\exp(\epsilon^{-1/\tau}))$

when auto-regressive language models assign significant attention to the initial token regardless of their semantic importance. Furthermore, the softmax self-attention is not computationally efficient as it involves a row-wise softmax operation. In response to these issues, Ramapuram et al. [43] proposed replacing the row-wise softmax function with the element-wise sigmoid function, which not only alleviates unnecessary token competition but also speeds up the computation. Then, they extended the FlashAttention2 framework for the softmax self-attention [11, 10] to the setting of the sigmoid self-attention and showed that the latter shared the same performance as the former in several tasks but with faster training and inference. On the theoretical side, they demonstrated that the Transformer model equipped with the sigmoid self-attention was a universal function approximator on sequence-to-sequence tasks. Nevertheless, a theoretical comparison between the sigmoid self-attention and the softmax self-attention is lacking in the literature.

**Contributions.** In this paper, our objective is to compare the sample complexity of these two attention variants through their connection to the Mixture-of-Experts (MoE) model [23]. Our contributions are twofold and can be summarized as follows (see also Table 1):

1. *Connection between self-attention and MoE:* In Section 2, we provide a rigorous derivation for illustrating that each row of a sigmoid self-attention matrix (resp. a softmax self-attention) head can be represented as an MoE with the sigmoid gating (resp. the softmax gating) and quadratic affinity scores, where each row of the value matrix plays a role as an expert. Although this connection has been mentioned in previous works [1, 9, 48], there have not been any theoretical guarantee for this result, to the best of our knowledge.

2. *Sample complexity comparison between sigmoid self-attention and softmax self-attention:* From the MoE perspective, we then perform a comprehensive convergence analysis of sigmoid self-attention. The results of our analysis demonstrate that polynomially many data points – of order  $\mathcal{O}(\epsilon^{-4})$  – are needed to achieve an estimation accuracy of order  $\epsilon$  for the value matrix’s rows. On the other hand, a similar convergence analysis of softmax self-attention conducted in [1] reveals that, in order to reach the same approximation error, exponentially many data points – of order  $\mathcal{O}(\exp(\epsilon^{-1/\tau}))$ , for some constant  $\tau > 0$  – are needed. Thus, we claim that sigmoid self-attention is more sample-efficient than softmax self-attention. Furthermore, it is worth emphasizing that the convergence analysis performed in this work is more technically challenging than that in [1] due to an issue in the model convergence induced by the structure of the sigmoid function, which will be elaborated in Section 3.

## 2 Background: Self-Attention and MoEs

**Self-Attention Mechanism.** The self-attention mechanism plays a crucial role in the Transformer architecture [47]. Given an input sequence  $\mathbb{X} \in \mathbb{R}^{N \times d}$  of  $N$  feature vectors of dimension  $d$ , the self-attention mechanism first projects it into the query matrix  $Q \in \mathbb{R}^{N \times d_k}$ , the key matrix  $K \in \mathbb{R}^{N \times d_k}$ , and the value matrix  $V \in \mathbb{R}^{N \times d_v}$  through three following linear transformations:

$$Q = \mathbb{X}W_Q, \quad K = \mathbb{X}W_K, \quad V = \mathbb{X}W_V,$$

where  $W_Q \in \mathbb{R}^{d \times d_k}$ ,  $W_K \in \mathbb{R}^{d \times d_k}$ ,  $W_V \in \mathbb{R}^{d \times d_v}$  are learnable weight matrices. Then, the vanilla *softmax self-attention* can be compactly written as

$$\text{SoftmaxAttn}(\mathbb{X}) = \text{Softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V, \quad (1)$$

where the Softmax function acts row-wise on the matrix  $QK^T/\sqrt{d_k} \in \mathbb{R}^{N \times N}$  as follows:  $\text{Softmax}(u) = \frac{1}{\sum_{j=1}^N \exp(u_j)}(\exp(u_1), \dots, \exp(u_N))$ , where  $u = (u_1, \dots, u_N)$  is a row vector in  $\mathbb{R}^N$ . On the other hand, the *sigmoid self-attention* is computed by replacing the above softmax function with the sigmoid function  $\sigma : z \mapsto (1 + \exp(-z))^{-1}$  applied element-wise to the matrix  $QK^T/\sqrt{d_k}$ :

$$\text{SigmoidAttn}(\mathbb{X}) = \sigma\left(\frac{QK^T}{\sqrt{d_k}}\right)V. \quad (2)$$

**Mixture of Experts.** Mixture of experts (MoE) [23] is an extension of classical mixture models [30] that aggregates the power of  $N$  sub-models called experts, each of which can be formulated as a feed-forward network [44], a regression function [16], or a classifier [5] denoted by  $\mathcal{E}_i : \mathbb{R}^d \rightarrow \mathbb{R}^{d_v}$  for  $1 \leq i \leq N$ . For that purpose, the MoE employs an adaptive gating mechanism, denoted by  $\mathcal{G} : \mathbb{R}^d \rightarrow \mathbb{R}^N$ , to calculate input-dependent weights for these experts in a dynamic way, so that the more relevant experts to the input will be assigned larger weights. Then, given an input  $h \in \mathbb{R}^d$ , the MoE output  $y \in \mathbb{R}^{d_v}$  is expressed as a weighted sum of the expert outputs:

$$y = \sum_{j=1}^N \mathcal{G}(h)_j \cdot \mathcal{E}_j(h). \quad (3)$$

In practice, there are two main types of gating functions, namely *the sigmoid gating function* [12, 8, 7] and *the softmax gating function* [23, 27] defined as

- (i) *Sigmoid gating*:  $\mathcal{G}(h)_j := \sigma(s_j(h))$  for  $1 \leq j \leq N$ , where  $s_j(h) \in \mathbb{R}$  represents the affinity score between the input  $h$  and the  $j$ -th expert  $\mathcal{E}_j$ ;
- (ii) *Softmax gating*:  $\mathcal{G}(h)_j = \text{Softmax}(s(h))_j$  for  $1 \leq j \leq N$ , where  $s(h) := (s_1(h), s_2(h), \dots, s_N(h))$ .

Due to its adaptability and expressiveness, the MoE has been widely leveraged in several fields, including natural language processing [15, 17, 51, 25, 35, 40], multi-task learning [32, 20, 6], and speech recognition [24, 18, 50, 29], etc.

**Self-Attention meets Mixture of Experts.** We will now show that each row of the sigmoid/softmax self-attention matrix can be represented as an MoE with quadratic affinity scores. Since the subsequent arguments apply for both attention variants, we will present only the derivation using the sigmoid self-attention matrix, which we recall is defined as

$$\text{SigmoidAttn}(\mathbb{X}) = \sigma(\mathbb{X}B\mathbb{X}^\top)\mathbb{X}W_V,$$

where  $B := \frac{W_Q W_K^\top}{\sqrt{d_k}} \in \mathbb{R}^{d \times d}$ . Let  $x_i \in \mathbb{R}^{1 \times d}$  denote the  $i$ -th row vector of the input matrix  $\mathbb{X}$ . Since the sigmoid function is applied element-wise, the  $(i, j)$ -th element of the matrix  $\sigma(\mathbb{X} B \mathbb{X}^\top) \in \mathbb{R}^{N \times N}$  is

$$[\sigma(\mathbb{X} B \mathbb{X}^\top)]_{i,j} = \sigma(x_i B x_j^\top).$$

As a result, the  $i$ -th row vector of the matrix  $\text{SigmoidAttn}(\mathbb{X})$  takes the form

$$[\text{SigmoidAttn}(\mathbb{X})]_{i,:} = \sum_{j=1}^N \sigma(x_i B x_j^\top) \cdot x_j W_V.$$

Next, let  $X = [x_1, x_2, \dots, x_N] \in \mathbb{R}^{1 \times Nd}$  be the concatenation of  $N$  input tokens. For each  $1 \leq i \leq N$ , let  $E_i \in \mathbb{R}^{Nd \times d}$  be the matrix such that  $X E_i = x_i$ . Then, the  $i$ -th row vector of the sigmoid self-attention matrix can be rewritten as

$$[\text{SigmoidAttn}(\mathbb{X})]_{i,:} = \sum_{j=1}^N \sigma(X E_i B E_j^\top X^\top) \cdot X E_j W_V \sum_{j=1}^N \sigma(X M_{ij} X^\top) \cdot X P_j,$$

where  $M_{ij} := E_i B E_j^\top = \frac{E_i W_Q W_K^\top E_j^\top}{\sqrt{d_k}}$  and  $P_j := E_j W_V$ .

Hence, each row of the sigmoid self-attention matrix can be represented as an MoE with quadratic affinity scores.

### 3 Problem Setup

Suppose that the data  $(X_1, Y_1), (X_2, Y_2), \dots, (X_n, Y_n) \in \mathbb{R}^d \times \mathbb{R}$  are generated according to the regression model

$$Y_i = f_{G_*}(X_i) + \varepsilon_i, \quad i = 1, 2, \dots, n, \quad (4)$$

where  $X_1, X_2, \dots, X_n$  are i.i.d. samples from a probability distribution  $\mu$  on  $\mathbb{R}^d$ , and  $\varepsilon_1, \varepsilon_2, \dots, \varepsilon_n$  are i.i.d. Gaussian noise variables with  $\mathbb{E}[\varepsilon_i | X_i] = 0$  and  $\text{Var}[\varepsilon_i | X_i] = \nu$ , for  $1 \leq i \leq n$ . Additionally, the regression function  $f_{G_*} : \mathbb{R}^d \rightarrow \mathbb{R}$  is unknown and formulated as a sigmoid gating mixture of  $N^*$  experts, i.e.

$$f_{G_*}(x) := \sum_{i=1}^{N^*} \frac{1}{1 + \exp(-s(x, \theta_i^*))} \cdot \mathcal{E}(x, \eta_i^*), \quad (5)$$

where  $\mathcal{E}(x, \eta_i^*)$  denotes the parametric expert function, while the affinity score function  $s(x, \theta_i^*)$  takes a quadratic form. In this work, we consider two types of affinity score functions:

- (i) the *fully quadratic score*:  $s(x, \theta_i^*) = x^\top A_i^* x + (b_i^*)^\top x + c_i^*$ ;
- (ii) and the *partially quadratic score*:  $s(x, \theta_i^*) = x^\top A_i^* x + c_i^*$ ;

where  $\Theta = \{(A, b, c, \eta) \in \mathbb{R}^{d \times d} \times \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^q\}$  denotes the parameter space. Given this setup, our goal is to determine the sample size necessary for the estimators of the parameters and experts in model (5) to reach an approximation error  $\epsilon > 0$ . Due to space limitation, we will present only the results for the case of a fully quadratic affinity score function in Section 4 and defer those for its

partial version to Appendix C. We also refer to Akbarian et al. [1] for a similar convergence analysis for softmax gating MoE with quadratic affinity scores.

**Notation.** We let  $[n]$  be the set  $\{1, 2, \dots, n\}$  for any  $n \in \mathbb{N}$ . Next, for any set  $S$ , we denote  $|S|$  as its cardinality. For any vectors  $v := (v_1, v_2, \dots, v_d) \in \mathbb{R}^d$  and  $\alpha := (\alpha_1, \alpha_2, \dots, \alpha_d) \in \mathbb{N}^d$ , we let  $v^\alpha = v_1^{\alpha_1} v_2^{\alpha_2} \dots v_d^{\alpha_d}$ ,  $|v| := \sum_{i=1}^d v_i$  and  $\alpha! := \alpha_1! \alpha_2! \dots \alpha_d!$ , while  $\|v\|$  denotes its 2-norm value. Lastly, for any positive sequences  $(a_n)_{n \geq 1}$  and  $(b_n)_{n \geq 1}$ , we write  $a_n = \mathcal{O}(b_n)$  or  $a_n \lesssim b_n$  if  $a_n \leq Cb_n$  for all  $n \in \mathbb{N}$ , where  $C > 0$  is some universal constant. For a sequence  $(A_n)_{n \geq 1}$  of positive random variables, the notation  $A_n = \mathcal{O}_P(b_n)$  signifies  $A_n/b_n$  is stochastically bounded, i.e., for any  $\epsilon > 0$ , there exists an  $M > 0$  such that  $\mathbb{P}(A_n/b_n > M) < \epsilon$  for all sufficiently large  $n$ .

**Least squares estimation.** To estimate the unknown ground-truth parameters  $\{A_i^*, b_i^*, c_i^*, \eta_i^*\}_{i=1}^{N^*}$ , we use the least squares method [46] to compute the estimator

$$\hat{G}_n := \arg \min_{G \in \mathcal{M}_N(\Theta)} \sum_{i=1}^n (Y_i - f_G(X_i))^2, \quad (6)$$

where  $\mathcal{M}_N(\Theta) := \{G = \sum_{i=1}^{N'} \frac{1}{1+\exp(-c_i)} \delta_{(A_i, b_i, \eta_i)} : 1 \leq N' \leq N, (A_i, b_i, \eta_i) \in \Theta\}$  is the set of all mixing measures with at most  $N$  atoms. As the number of ground-truth experts  $N^*$  is unknown in practice, we assume that the number of fitted experts  $N$  is larger than  $N^*$ , i.e.  $N > N^*$ .

**A challenge in the convergence of the regression function estimator.** Since the number of fitted experts is larger than the number of ground-truth experts, there must exist some atom  $(A_i^*, b_i^*, \eta_i^*)$  of the mixing measure  $G_*$  that is fitted by at least two atoms of  $\hat{G}_n$ ; we will refer to  $(A_i^*, b_i^*, \eta_i^*)$  as an *over-specified* atom of  $G_*$ . For example, if, say,  $(\hat{A}_i^n, \hat{b}_i^n, \hat{\eta}_i^n) \rightarrow (A_i^*, b_i^*, \eta_i^*)$  in probability as  $n \rightarrow \infty$ , for  $i \in \{1, 2\}$ , then the corresponding estimators of the expert functions converge in probability as well, i.e.  $\mathcal{E}(x, \hat{\eta}_i^n) \rightarrow \mathcal{E}(x, \eta_i^*)$ , in probability, for  $i = 1, 2$ . This will, in turn, ensure the convergence of the regression function, i.e.  $\|f_{\hat{G}_n} - f_{G_*}\|_{L^2(\mu)} \rightarrow 0$  in probability, provided that, for  $\mu$ -almost every  $x$ ,

$$\sum_{i=1}^2 \frac{1}{1 + \exp(-x^\top \hat{A}_i^n x - (\hat{b}_i^n)^\top x - \hat{c}_i^n)} \rightarrow \frac{1}{1 + \exp(-x^\top A_1^* x - (b_1^*)^\top x - c_1^*)},$$

in probability as  $n \rightarrow \infty$ . Notably, the above limit holds only if  $A_1^* = 0_{d \times d}$  and  $b_1^* = 0_d$  (see Appendix B.5 for further details). Thus, we will divide our analysis into two complementary regimes of the gating parameters:

- (i) *Sparse regime*: all the over-specified gating parameters are zero,  $(A_i^*, b_i^*) = (0_{d \times d}, 0_d)$ ;
- (ii) *Dense regime*: at least one among the over-specified gating parameters is non-zero,  $(A_i^*, b_i^*) \neq (0_{d \times d}, 0_d)$ .

The next results derive separate convergence rates for the regression function estimator  $f_{\hat{G}_n}$  under the two regimes.

**Proposition 1.** *Under the sparse regime of the gating parameters,*

$$\|f_{\hat{G}_n} - f_{G_*}\|_{L^2(\mu)} = \mathcal{O}_P(\sqrt{\log(n)/n}). \quad (7)$$

The proof of Proposition 1 can be found in Appendix B.1. The bound (7) reveals that the convergence rate of the regression function estimator  $f_{\hat{G}_n}$  to the ground-truth regression function is of parametric order  $\mathcal{O}_P(\sqrt{\log(n)/n})$  under the sparse regime. In contrast, under the dense

regime, the smallest  $L_2(\mu)$  distance between the regression function estimator  $f_{\widehat{G}_n}$  and the set of MoE regression functions  $f_{\overline{G}}$  with more than  $N^*$  experts functions where  $\overline{G} \in \overline{\mathcal{M}}_N(\Theta) := \arg \min_{G \in \mathcal{M}_N(\Theta) \setminus \mathcal{M}_{N^*}(\Theta)} \|f_G - f_{G_*}\|_{L_2(\mu)}$  vanishes to 0 as  $n \rightarrow \infty$ . The rate for that convergence is given by the following result.

**Corollary 1.** *Under the dense regime of the gating parameters,*

$$\inf_{\overline{G} \in \overline{\mathcal{M}}_N(\Theta)} \|f_{\widehat{G}_n} - f_{\overline{G}}\|_{L^2(\mu)} = \mathcal{O}_P(\sqrt{\log(n)/n}).$$

**From regression convergence rate to expert rate.** In light of Proposition 1 and Corollary 1, we conclude that, if there exists a loss function, say,  $\mathcal{L}(\cdot, \cdot)$  between the parameters  $G$  and  $G_*$  such that the lower bound  $\|f_{\widehat{G}_n} - f_{G_*}\|_{L^2(\mu)} \gtrsim \mathcal{L}(\widehat{G}_n, G_*)$  holds true, then we immediately deduce  $\mathcal{L}(\widehat{G}_n, G_*) = \mathcal{O}_P(\sqrt{\log(n)/n})$ , thereby obtaining a convergence rate for the expert parameters. We carry out this program in the next section.

## 4 Sample Complexity of Sigmoid Self-Attention

In this section, we investigate the sample complexity of the sigmoid self-attention through the perspective of the sigmoid gating MoE with fully quadratic affinity score function. In particular, we determine how much data the experts need to reach an approximation error  $\epsilon$ , which can be deduced from the expert convergence rate. We start with the sparse regime of the gating parameters in Section 4.1, and then proceed with the dense regime in Section 4.2.

### 4.1 Sparse Regime of Gating Parameters

Recall that under the sparse regime, all the over-specified gating parameters are zero, that is,  $(A_i^*, b_i^*) = (0_{d \times d}, 0_d)$ . Suppose that  $\{A_i^*, b_i^*\}_{i=1}^{\overline{N}}$  are *over-specified* parameters, i.e., those fitted by at least two estimators, where  $1 \leq \overline{N} \leq N^*$ . The remaining gating parameters are *exactly-specified* parameters  $\{A_i^*, b_i^*\}_{\overline{N}+1}^{N^*}$ , i.e., those fitted by exactly one estimator.

$$\underbrace{(A_1^*, b_1^*), \dots, (A_{\overline{N}}^*, b_{\overline{N}}^*)}_{\text{over-specified}}, \underbrace{(A_{\overline{N}+1}^*, b_{\overline{N}+1}^*), \dots, (A_{N^*}^*, b_{N^*}^*)}_{\text{exactly-specified}}$$

As mentioned in Section 3, in order to obtain the expert convergence rate, it is sufficient to establish the lower bound  $\|f_{\widehat{G}_n} - f_{G_*}\|_{L^2(\mu)} \gtrsim \mathcal{L}(\widehat{G}_n, G_*)$ . A popular approach adopted in previous works [33, 38] for this problem is to decompose the difference  $f_{\widehat{G}_n}(x) - f_{G_*}(x)$  by applying a Taylor expansion to the product of the sigmoid gating function and the expert function given by

$$F(x; A, b, c, \eta) := \sigma(x^\top A x + b^\top x + c) \cdot \mathcal{E}(x, \eta).$$

This decomposition of the regression function is expected to consist of linearly independent terms so that when  $\|f_{\widehat{G}_n} - f_{G_*}\|_{L^2(\mu)} \rightarrow 0$  as  $n \rightarrow \infty$ , the parameter discrepancies in the decomposition will also converge to zero, leading to both parameter and expert convergence. To secure such linear independence, we need to impose a *strong identifiability* condition on the function  $x \mapsto F(x; A, b, c, \eta)$ .

**Definition 1** (Strong identifiability). *We call an expert function  $x \mapsto \mathcal{E}(x, \eta)$  strongly identifiable if it is twice differentiable w.r.t its parameter  $\eta$  for  $\mu$ -almost all  $x$  and, for any natural number  $\ell$  and any distinct parameters  $\{(A_i, b_i, c_i, \eta_i)\}_{i=1}^\ell$ , the sets of functions*

$$\begin{aligned} & \left\{ \frac{\partial^{|\gamma_1|+|\gamma_2|+|\gamma_3|} F}{\partial A^{\gamma_1} \partial b^{\gamma_2} \partial \eta^{\gamma_3}}(x, 0_{d \times d}, 0_d, c_i, \eta_i) : i \in [\ell], \sum_{j=1}^3 |\gamma_j| \in [2] \right\} \quad \text{and} \\ & \left\{ \frac{\partial F}{\partial A^{\tau_1} \partial b^{\tau_2} \partial c^{\tau_3} \partial \eta^{\tau_4}}(x; A_i, b_i, c_i, \eta_i) : i \in [\ell], \sum_{j=1}^4 |\tau_j| = 1 \right\} \end{aligned}$$

each contains linearly independent, functions for  $\mu$ -almost all  $x$ , where  $(\gamma_1, \gamma_2, \gamma_3) \in \mathbb{N}^{d \times d} \times \mathbb{N}^d \times \mathbb{N}^q$  and  $(\tau_1, \tau_2, \tau_3, \tau_4) \in \mathbb{N}^{d \times d} \times \mathbb{N}^d \times \mathbb{N} \times \mathbb{N}^q$ .

**Examples.** Let us consider two-layer neural networks of the form  $\mathcal{E}(x, (\alpha, \beta, \lambda)) = \lambda \phi(\alpha^\top x + \beta)$ , where  $\phi$  is some activation function and  $(\alpha, \beta, \lambda) \in \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}$ . It can be verified that if  $\phi$  is the ReLU or GELU function,  $\alpha \neq 0_d$ , and  $\lambda \neq 0$ , then the function  $x \mapsto \mathcal{E}(x, (\alpha, \beta, \lambda))$  is strongly identifiable. In contrast, if the expert function is of polynomial form  $\mathcal{E}(x, (\alpha, \beta)) = (\alpha^\top x + \beta)^p$  for some  $p \in \mathbb{N}$ , then it fails to satisfy the strong identifiability condition. For the case of  $p = 1$ , this is due to the linear dependence expressed by the partial differential equations (PDEs)

$$\frac{\partial^2 F}{\partial A \partial c} = \frac{\partial^2 F}{\partial b \partial b^\top}, \quad \frac{\partial^2 F}{\partial A \partial \beta} = \frac{\partial^2 F}{\partial b \partial \alpha}, \quad \frac{\partial^2 F}{\partial b \partial \beta} = \frac{\partial^2 F}{\partial c \partial \alpha}. \quad (8)$$

Intuitively, the strong identifiability condition helps eliminate potential interactions among parameters expressed in the language of PDEs, namely those in equation (8), where gating parameters interact with themselves and with expert parameters. We will show later in Theorem 2 that those interactions lead to strikingly slow expert convergence rates, thereby reducing the model sample complexity.

In the next sections, we will study the convergence behavior of strongly identifiable experts and polynomial experts.

#### 4.1.1 Strongly Identifiable Experts

**Voronoi loss.** For a mixing measure  $G$  with  $1 \leq N' \leq N$  atoms, we allocate its atoms across the Voronoi cells  $\{\mathcal{A}_j \equiv \mathcal{A}_j(G), j \in [N']\}$  generated by the atoms of  $G_*$ , where

$$\mathcal{A}_j := \{i \in [N'] : \|\theta_i - \theta_j^*\| \leq \|\theta_i - \theta_\ell^*\|, \forall \ell \neq j\}, \quad (9)$$

with  $\theta_i := (A_i, b_i, \eta_i)$  and  $\theta_j^* := (A_j^*, b_j^*, \eta_j^*)$  for all  $j \in [N']$ . Then, the Voronoi loss function is defined as

$$\begin{aligned} \mathcal{L}_1(G, G_*) := & \sum_{j=1}^{\bar{N}} \left| \sum_{i \in \mathcal{A}_j} \frac{1}{1 + \exp(-c_i)} - \frac{1}{1 + \exp(-c_j^*)} \right| + \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} \left[ \|\Delta A_{ij}\|^2 + \|\Delta b_{ij}\|^2 \right. \\ & \left. + \|\Delta \eta_{ij}\|^2 \right] + \sum_{j=\bar{N}+1}^{N^*} \sum_{i \in \mathcal{A}_j} \left[ \|\Delta A_{ij}\| + \|\Delta b_{ij}\| + |\Delta c_{ij}| + \|\Delta \eta_{ij}\| \right], \end{aligned}$$

where we denote  $\Delta A_{ij} := A_i - A_j^*$ ,  $\Delta b_{ij} := b_i - b_j^*$ ,  $\Delta c_{ij} := c_i - c_j^*$  and  $\Delta \eta_{ij} := \eta_i - \eta_j^*$ . In the statement above, if the Voronoi cell  $\mathcal{A}_j$  is empty, the corresponding summation term is conventionally

defined to be zero. Additionally, the Voronoi loss function  $\mathcal{L}_1$  can be computed efficiently, with a computational complexity of  $\mathcal{O}(N \times N^*)$ .

With the above Voronoi loss at hand, we finally obtain the following parameter convergence rate.

**Theorem 1.** *If the expert function  $x \mapsto \mathcal{E}(x, \eta)$  is strongly identifiable, then the lower bound  $\|f_G - f_{G_*}\|_{L^2(\mu)} \gtrsim \mathcal{L}_1(G, G_*)$  holds true for any  $G \in \mathcal{M}_N(\Theta)$ , then*

$$\mathcal{L}_1(\widehat{G}_n, G_*) = \mathcal{O}_P(\sqrt{\log(n)/n}).$$

The proof of Theorem 1 is in Appendix B.2. A few comments regarding the above result are in order.

(i) *Parameter convergence rates:* From the construction of the Voronoi loss  $\mathcal{L}_1$ , the convergence rates for estimating the over-specified parameters  $A_i^*, b_i^*, \eta_i^*, i \in [\bar{N}]$ , are all of the same order  $\mathcal{O}_P([\log(n)/n]^{\frac{1}{4}})$ . On the other hand, those for the exactly-specified parameters  $A_i^*, b_i^*, \eta_i^*, \bar{N} + 1 \leq i \leq N^*$  are faster, of the order  $\mathcal{O}_P([\log(n)/n]^{\frac{1}{2}})$ .

(ii) *Expert convergence rates:* Since the expert function  $\mathcal{E}(\cdot, \eta)$  is twice differentiable w.r.t  $\eta$  over a bounded domain, it is also a Lipschitz function w.r.t  $\eta$ . Therefore, by denoting  $\widehat{G}_n = \sum_{i=1}^{\bar{N}_n} \frac{1}{1 + \exp(-\widehat{c}_i^n)} \delta_{(\widehat{A}_i^n, \widehat{b}_i^n, \widehat{\eta}_i^n)}$ , we conclude that

$$\sup_x |\mathcal{E}(x, \widehat{\eta}_i^n) - \mathcal{E}(x, \eta_j^*)| \leq L_1 \|\widehat{\eta}_i^n - \eta_j^*\|, \quad (10)$$

for any  $i \in \mathcal{A}_j(\widehat{G}_n)$ , where  $L_1 \geq 0$  is a Lipschitz constant. The above bound implies that the convergence rates for estimating the exactly-specified experts and over-specified experts are of orders  $\mathcal{O}_P([\log(n)/n]^{\frac{1}{2}})$  and  $\mathcal{O}_P([\log(n)/n]^{\frac{1}{4}})$ , respectively. Thus, it takes the exactly-specified experts a polynomial number  $\mathcal{O}(\epsilon^{-2})$  of data points to achieve an approximation error of  $\epsilon$ , while the over-specified experts need a polynomial number  $\mathcal{O}(\epsilon^{-4})$  of data to achieve the same error.

#### 4.1.2 Polynomial Experts

We now investigate polynomial experts of the form  $\mathcal{E}(x, (\alpha, \beta)) = (\alpha^\top x + \beta)^p$ , where  $p \in \mathbb{N}$ . As mentioned in the example paragraph following the Definition 1, the polynomial experts do not meet the strong identifiability condition due to the linear dependence among the derivative of the function  $F(x; A, b, c, \eta)$ . For example, when  $p = 1$ , such linear dependence is exhibited via the interaction among the gating parameters (see the first PDE in equation (8)) and the interaction between the gating parameters and the expert parameters (see the last two PDEs in equation (8)). Those PDEs account for the non-strong identifiability of the polynomial experts. Notably, we will demonstrate in Theorem 2 that such parameter interactions lead to slow convergence rates for parameter estimation and expert estimation. Toward that goal, let us introduce the Voronoi loss function

$$\begin{aligned} \mathcal{L}_{2,r}(G, G_*) := & \sum_{j=1}^{\bar{N}} \left| \sum_{i \in \mathcal{A}_j} \frac{1}{1 + \exp(-c_i)} - \frac{1}{1 + \exp(-c_j^*)} \right| + \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} \left[ \|\Delta A_{ij}\|^r + \|\Delta b_{ij}\|^r \right. \\ & \left. + \|\Delta \alpha_{ij}\|^r + \|\Delta \beta_{ij}\|^r \right] + \sum_{j=\bar{N}+1}^{N^*} \sum_{i \in \mathcal{A}_j} \left[ \|\Delta A_{ij}\|^r + \|\Delta b_{ij}\|^r + \|\Delta c_{ij}\|^r + \|\Delta \alpha_{ij}\|^r + \|\Delta \beta_{ij}\|^r \right], \end{aligned} \quad (11)$$

where we denote  $\Delta \alpha_{ij} := \alpha_i - \alpha_j^*$  and  $\Delta \beta_{ij} := \beta_i - \beta_j^*$ .

**Theorem 2.** Suppose that the expert function takes a polynomial form  $\mathcal{E}(x, \alpha, \beta) = (\alpha^\top x + \beta)^p$ , for some  $p \in \mathbb{N}$ . Then, for any  $r \geq 1$ ,

$$\inf_{\tilde{G}_n \in \mathcal{M}_N(\Theta)} \sup_{G \in \mathcal{M}_N(\Theta) \setminus \mathcal{M}_{N^*-1}(\Theta)} \mathbb{E}_{f_G} [\mathcal{L}_{2,r}(\tilde{G}_n, G)] \gtrsim \frac{1}{\sqrt{n}},$$

where  $\mathbb{E}_{f_G}$  indicates the expectation taken w.r.t. the product measure with  $f_G^n$ .

The proof of Theorem 2 is in Appendix B.3. Below we highlight some important implications of the above result.

(i) *Parameter convergence rates:* The minimax lower bound in Theorem 2 implies that the convergence rates for estimating parameters  $A_j^*, b_j^*, \alpha_j^*, \beta_j^*$  are slower than any polynomial rates  $\mathcal{O}_P(n^{-1/2r})$  for any  $r \geq 1$ , potentially as slow as  $\mathcal{O}_P(1/\log^\tau(n))$ , for some constant  $\tau > 0$ .

(ii) *Expert convergence rates:* Following the parameter convergence rates and using the same arguments as in equation (10), we deduce that the convergence rates for estimating experts could also be as slow as  $\mathcal{O}_P(1/\log^\tau(n))$ , for some constant  $\tau > 0$ . Therefore, it may require the experts an exponential number of data  $\mathcal{O}(\exp(\epsilon^{-1/\tau}))$  to obtain the approximation error  $\epsilon$ .

**Sample complexity comparison under the sparse regime:** According to the results in [1], estimating the experts in the softmax self-attention share the same sample complexity as estimating those in the sigmoid version. In particular, strongly identifiable experts and polynomial experts need  $\mathcal{O}(\epsilon^{-4})$  and  $\mathcal{O}(\exp(\epsilon^{-1/\tau}))$  data points to achieve the approximation error  $\epsilon$ . Thus, we claim that the sigmoid self-attention is as sample-efficient as the softmax self-attention under the sparse regime of gating parameters. Note that it is unlikely for all gating parameters to vanish in practice, as typical initialization schemes, training dynamics, or architectural designs (e.g., ReLU activations, residual connections) generally ensure that all gating values (mixture weights) are input-dependent. As a result, the model rarely enters an extremely sparse regime. That is, the sparse regime is less common than the dense regime, which will be studied in the next section.

## 4.2 Dense Regime of Gating Parameters

We now turn our attention to the dense regime where we assume that there exists some over-specified gating parameter different from zero, that is,  $(A_i^*, b_i^*) \neq (0_{d \times d}, 0_d)$ , for some  $i \in [N^*]$ . As demonstrated in Section 3, the smallest distance between the regression function estimator  $f_{\tilde{G}_n}$  and the set of the regression functions  $f_{\bar{G}}$  goes to 0 where  $\bar{G} \in \overline{\mathcal{M}}_N(\Theta) := \arg \min_{G \in \mathcal{M}_N(\Theta) \setminus \mathcal{M}_{N^*}(\Theta)} \|f_G - f_{G_*}\|_{L^2(\mu)}$ . Without loss of generality (WLOG), we assume that

$$\bar{G} := \sum_{i=1}^N \frac{1}{1 + \exp(-\bar{c}_i)} \delta_{(\bar{A}_i, \bar{b}_i, \bar{\eta}_i)}.$$

Similar to the sparse regime of gating parameters, we also establish an identifiability condition on the expert function to avoid any interaction among the parameters expressed via PDEs under the dense regime. Since the expert function is required to satisfy only a subset of the strong identifiability condition in Definition 1 in this case, we refer to the condition as *weak identifiability*.

**Definition 2** (Weak identifiability). *We call an expert function  $x \mapsto \mathcal{E}(x, \eta)$  weakly identifiable if it is differentiable w.r.t its parameter  $\eta$  for  $\mu$ -almost all  $x$  and, for any positive integer  $\ell$  and any*

distinct parameters  $\{(A_i, b_i, c_i, \eta_i)\}_{i=1}^\ell$ , the functions in the family

$$\left\{ \frac{\partial F}{\partial A^{\tau_1} \partial b^{\tau_2} \partial c^{\tau_3} \partial \eta^{\tau_4}}(x, A_i, b_i, c_i, \eta_i) : i \in [\ell], \sum_{j=1}^4 |\tau_j| = 1 \right\}$$

are linearly independent, for  $\mu$ -almost all  $x$ , where  $(\tau_1, \tau_2, \tau_3, \tau_4) \in \mathbb{N}^{d \times d} \times \mathbb{N}^d \times \mathbb{N} \times \mathbb{N}^q$ .

**Examples.** Strongly identifiable experts meet the weak identifiability condition. For instance, it can be verified that the previously mentioned two-layer neural networks of the form  $\mathcal{E}(x, (\alpha, \beta, \lambda)) = \lambda \phi(\alpha^\top x + \beta)$ , where  $\phi$  is ReLU or GELU function,  $\alpha \neq 0_d$ , and  $\lambda \neq 0$  are weakly identifiable. Moreover, polynomial experts  $\mathcal{E}(x, (\alpha, \beta)) = (\alpha^\top x + \beta)^p$ , for  $p \in \mathbb{N}$ , also satisfy the weak identifiability condition although they are not strongly identifiable.

In Theorem 3 below, we provide convergence rates for weakly identifiable experts based on the Voronoi loss

$$\mathcal{L}_3(G, \bar{G}) := \sum_{j=1}^N \sum_{i \in \mathcal{A}_j} [\|A_i - \bar{A}_i\| + \|b_i - \bar{b}_i\| + |c_i - \bar{c}_i| + \|\eta_i - \bar{\eta}_i\|].$$

**Theorem 3.** *If the function  $x \mapsto \mathcal{E}(x, \eta)$  is weakly identifiable, then the lower bound  $\inf_{\bar{G} \in \bar{\mathcal{M}}_N(\Theta)} \|f_G - f_{\bar{G}}\|_{L^2(\mu)} \gtrsim \mathcal{L}_3(G, \bar{G})$  holds true for any mixing measure  $G \in \mathcal{M}_N(\Theta)$ . As a consequence,*

$$\inf_{\bar{G} \in \bar{\mathcal{M}}_N(\Theta)} \mathcal{L}_3(\hat{G}_n, \bar{G}) = \mathcal{O}_P(\sqrt{\log(n)/n}).$$

The proof of Theorem 3 is in Appendix B.4. Since the convergence rates for the parameter estimators  $\hat{\eta}_i^n$  are of order  $\mathcal{O}_P([\log(n)/n]^{1/2})$ , the weakly identifiable expert estimators  $\mathcal{E}(x, \hat{\eta}_i^n)$  also admit the same convergence rates, as indicated by the inequality (10). Therefore, it takes those experts only a polynomial number of samples  $\mathcal{O}(\epsilon^{-2})$  to achieve an approximation error of  $\epsilon$ .

**Sample complexity comparison under the dense regime:** Recall that it costs strongly identifiable experts and polynomial experts in the softmax self-attention  $\mathcal{O}(\epsilon^{-4})$  and  $\mathcal{O}(\exp(\epsilon^{-1/\tau}))$  data points to reach the approximation error  $\epsilon$ , respectively [1]. On the other hand, as those experts satisfy the weak identifiability condition, they need only a sample size of order  $\mathcal{O}(\epsilon^{-2})$  to achieve the same error under the dense regime of sigmoid self-attention. For that reason, we claim that the sigmoid self-attention is more sample efficient than its softmax counterpart under the dense regime, which is more likely to occur in practice than the sparse regime.

## 5 Numerical Experiments

Below, we present some numerical experiments confirming our theoretical findings that on the superior sample complexity of sigmoid attention versus softmax attention. In particular, we capture the empirical convergence rates of parameter estimation in sigmoid quadratic gating and softmax quadratic gating MoE models with two expert configurations: expert networks with ReLU activation (ReLU experts) and those with linear activation (linear experts).

**Setup.** We generated synthetic data from a softmax quadratic gating MoE model. Both sigmoid quadratic gating MoE and softmax quadratic gating MoE models were fitted, with varying sample sizes. The empirical convergence rates are then evaluated for each model across the two expert

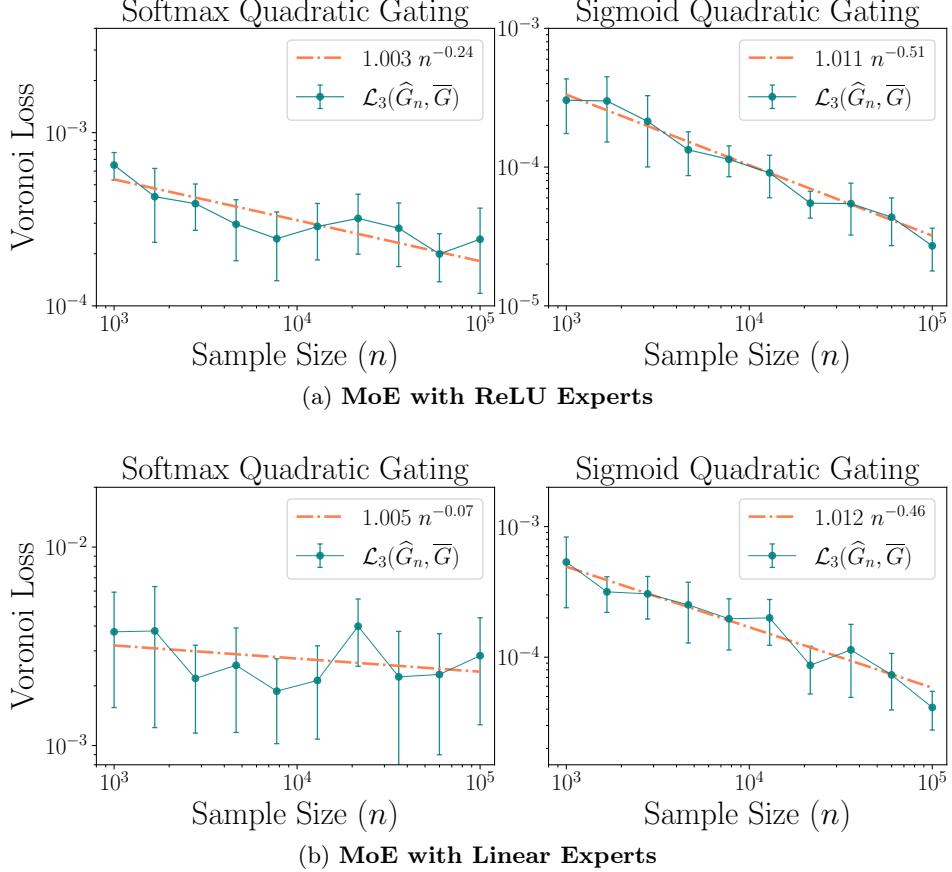


Figure 1: Log-log plots of the convergence rates of Voronoi losses for softmax and sigmoid quadratic gating MoE models. 1a Comparison between softmax quadratic gating and sigmoid quadratic gating MoE with ReLU experts. 1b Comparison between softmax quadratic gating and sigmoid quadratic gating with linear experts. Each plot illustrates the empirical Voronoi loss convergence rates, with solid lines representing the Voronoi losses and dash-dotted lines showing fitted trends.

configurations (ReLU experts and linear experts), providing insights into sigmoid quadratic gating performance. The details about the values of the ground-truth parameters and the training procedure are presented in Appendix E.

**Results.** Figure 1 shows the empirical convergence rates of Voronoi loss for sigmoid and softmax quadratic gating MoE models, with error bars representing three standard deviations to account for variability across runs. In Figure 1a for MoE models with ReLU experts, we observe that the softmax quadratic gating MoE converges at the rate of order  $\mathcal{O}(n^{-0.24})$ , whereas the sigmoid version achieves a significantly faster rate of order  $\mathcal{O}(n^{-0.51})$ . Similarly, in Figure 1b for MoE models with linear experts, the sigmoid quadratic gating MoE attains a convergence rate of order  $\mathcal{O}(n^{-0.46})$ , while the softmax counterpart exhibits a significantly slower convergence rate of order  $\mathcal{O}(n^{-0.07})$ . These results empirically validate our theoretical findings that sigmoid quadratic gating MoE admits faster parameter estimation rates across both the configurations of ReLU experts and linear experts.

## 6 Conclusion

In this paper, we establish a mathematical connection between the self-attention mechanism in the Transformer architecture and the Mixture-of-Experts (MoE) model. Within the MoE framework, we investigate the sample complexity of the sigmoid self-attention by conducting a convergence analysis of parameter and expert estimation under the sigmoid gating MoE models with fully and partially quadratic affinity score functions in both the sparse and dense regimes for the gating parameters. Our results show that sigmoid self-attention has a higher sample complexity than the softmax version in the more common dense regime, a finding further confirmed by our simulations. A limitation of our work is that we consider only a single attention head rather than a more popular multi-head attention mechanism [47]. However, we believe this problem can be overcome by formulating the multi-head attention as a hierarchical MoE [27, 37], which we leave to future work.

## Supplement to “Sigmoid Self-Attention has Lower Sample Complexity than Softmax Self-Attention: A Mixture-of-Experts Perspective”

The supplementary material is structured as follows: Appendix A provides a discussion of related works on the self-attention mechanism and Mixture of Experts (MoE) models. Appendix B contains detailed proofs of the results presented in Sections 3 and 4. In Appendix C, we investigate the sample complexity of sigmoid self-attention under the partially quadratic score for the gating mechanism of the MoE model, accompanied by proofs of the relevant results. Finally, Appendix D presents additional experimental results, while Appendix E provides the experimental details.

## A Related Works

Attention mechanisms have become a cornerstone of the Transformer architecture. In its standard formulation, attention weights are calculated as the softmax of the dot products between keys and queries [47]. More recently, Ramapuram et al. [43] established that Transformers with sigmoid-based attention are universal function approximators and exhibit enhanced regularity compared to those using softmax attention. Beyond their fundamental role in Transformers, attention mechanisms have been closely linked to Mixture of Experts (MoE) architectures. Csordas et al. [9] introduces SwitchHead, an MoE-based attention method designed to reduce both the computation and storage requirements for attention matrices. Wu et al. [48] proposes Multi-Head MoE (MH-MoE), where each parallel layer contains a set of  $N$  experts decoupled from the head in multi-head self-attention, focusing on scalability and architectural improvements. Jin et al. [26] demonstrates that multi-head attention can be written in summation form, motivating the Mixture-of-Head (MoH) architecture, which treats each attention head as an MoE expert. This design boosts inference efficiency without compromising accuracy or increasing parameter counts. Le et al. Furthermore, Akbarian et al. [1] examined the connection between softmax self-attention and softmax quadratic gating MoE, showing that MoE models with softmax quadratic gating outperform their counterparts using traditional softmax linear gating [38].

Another line of research has also explored convergence rates for expert estimation in Gaussian Mixture of Experts models. First, Mendes et al. [34] examined maximum likelihood estimation for Mixture of Experts models with polynomial regression experts. They analyzed how quickly the estimated density converges to the true density under the Kullback–Leibler (KL) divergence and offered insights on selecting the appropriate number of experts. Ho et al. [21] derived convergence rates for parameter estimation by utilizing a connection between the algebraic independence of expert functions and model parameters, which they formulated through a class of partial differential equations (PDEs). Building on this foundation, Nguyen et al. [39] investigated Gaussian MoE models with softmax gating, uncovering that expert estimation rates are influenced by the solvability of a system of polynomial equations resulting from the interplay between gating and expert parameters. Yan et al. [49] analyzed the convergence of parameter estimation in a contaminated Gaussian mixture of experts, where a pre-trained Gaussian expert is mixed with a prompt Gaussian expert by an unknown proportion. Their approach leverages the algebraic interaction between the pre-trained model’s parameters and the prompt. For Gaussian MoE models incorporating Top-K sparse softmax gates, Nguyen et al. [36] analyzed convergence rates using novel loss functions specifically designed for sparse gating mechanisms.

## B Proof of the Results in Sections 3 and 4

### B.1 Proof of Proposition 1

First we will introduce some necessary notations used throughout this appendix. We let  $\mathcal{F}_N(\Theta) := \{f_G(x) : G \in \mathcal{M}_N(\Theta)\}$  to be the set of all regression functions in  $\mathcal{M}_N(\Theta)$ . Then we consider the intersection between  $L^2(\mu)$ -ball centered around the regression function  $f_{G_*}(x)$  with radius  $\delta > 0$  and the set  $\mathcal{F}_N(\Theta)$  and define as

$$\mathcal{F}_N(\Theta, \delta) := \{f \in \mathcal{F}_N(\Theta) : \|f - f_{G_*}\|_{L^2(\mu)} \leq \delta\}.$$

We assess the complexity of the above class using the bracketing entropy integral in [46],

$$\mathcal{J}_B(\delta, \mathcal{F}_N(\Theta, \delta)) := \int_{\delta^2/2^{13}}^{\delta} H_B^{1/2}(t, \mathcal{F}_N(\Theta, t), \|\cdot\|_{L^2(\mu)}) dt \vee \delta, \quad (12)$$

where  $H_B(t, \mathcal{F}_N(\Theta, t), \|\cdot\|_{L^2(\mu)})$  represents the bracketing entropy [46] of  $\mathcal{F}_N(\Theta, t)$  under the  $L^2(\mu)$ -norm, and  $t \vee \delta := \max\{t, \delta\}$ . Using similar arguments as in Theorems 7.4 and 9.2 of [46], with the notation adapted to our context, we derive the following lemma:

**Lemma 1.** *Let  $\Psi(\delta) \geq \mathcal{J}_B(\delta, \mathcal{F}_N(\Theta, \delta))$  such that  $\Psi(\delta)/\delta^2$  is a non-increasing function of  $\delta$ . Then, for some universal constant  $c$  and for some sequence  $(\delta_n)$  that satisfied  $\sqrt{n}\delta_n^2 \geq c\Psi(\delta_n)$ , the following holds for any  $\delta \geq \delta_n$ :*

$$\mathbb{P}\left(\|f_{\widehat{G}_n} - f_{G_*}\|_{L^2(\mu)} > \delta\right) \leq c \exp\left(-\frac{n\delta^2}{c^2}\right).$$

**Outline of Proof.** To prove the result, it suffices to establish the following bound for the bracketing entropy  $H_B(\cdot)$  of the function class  $\mathcal{F}_N(\Theta)$ :

$$H_B(\epsilon, \mathcal{F}_N(\Theta), \|\cdot\|_{L^2(\mu)}) \lesssim \log(1/\epsilon), \quad \forall \epsilon \in (0, 1/2]. \quad (13)$$

Using equation (13), the integral term  $\mathcal{J}_B(\delta, \mathcal{F}_N(\Theta, \delta))$  can be bounded as:

$$\mathcal{J}_B(\delta, \mathcal{F}_N(\Theta, \delta)) = \int_{\delta^2/2^{13}}^{\delta} H_B^{1/2}(t, \mathcal{F}_N(\Theta, t), \|\cdot\|_{L^2(\mu)}) dt \vee \delta \lesssim \int_{\delta^2/2^{13}}^{\delta} \sqrt{\log(1/t)} dt \vee \delta. \quad (14)$$

This integral evaluates to:

$$\mathcal{J}_B(\delta, \mathcal{F}_N(\Theta, \delta)) \lesssim \delta \cdot \log(1/\delta)^{1/2}.$$

To ensure  $\Psi(\delta) \geq \mathcal{J}_B(\delta, \mathcal{F}_N(\Theta, \delta))$ , we define  $\Psi(\delta) := \delta \cdot \log(1/\delta)^{1/2}$ , which satisfies the condition that  $\Psi(\delta)/\delta^2$  is a non-increasing function.

**Choice of Sequence** ( $\delta_n$ ). Set  $\delta_n := \sqrt{\log(n)/n}$ . By construction, this sequence satisfies  $\sqrt{n}\delta_n^2 \geq c\Psi(\delta_n)$  for some universal constant  $c$ . Substituting  $\delta = \delta_n$  into Lemma 1 yields the desired probability bound.

**Conclusion.** By applying Lemma 1, we reduce the problem to verifying the entropy bound (13), which will be established now:

### Proof of Inequality (13):

*Proof.* To prove inequality (13), we begin by noting that the expert functions are bounded, implying

$$|f_G(x)| \leq M \quad \text{for almost every } x,$$

where  $M > 0$  is a constant.

**Step 1: Covering the Function Class.** Let  $\tau \leq \epsilon$ , and consider a  $\tau$ -cover  $\{\zeta_1, \dots, \zeta_V\}$  of the function class  $\mathcal{F}_N(\Theta)$  under the  $L^2(\mu)$ -norm. The covering number  $V := V(\tau, \mathcal{F}_N(\Theta), \|\cdot\|_{L^\infty})$  represents the  $\tau$ -covering number of the metric space  $\mathcal{F}_N(\Theta)$  with the  $L^\infty$ -norm.

We construct brackets of the form  $[L_i(x), U_i(x)]$  for  $i \in [V]$ , where

$$L_i(x) := \max\{\zeta_i(x) - \tau, 0\}, \quad U_i(x) := \max\{\zeta_i(x) + \tau, M\}.$$

**Step 2: Properties of the Brackets.** It follows that for all  $i \in [V]$ :

- $\mathcal{F}_N(\Theta) \subseteq \bigcup_{i=1}^V [L_i(x), U_i(x)]$ ,
- $U_i(x) - L_i(x) \leq \min\{2\tau, M\}$ .

Thus, the bracket size satisfies

$$\|U_i - L_i\|_{L^2(\mu)}^2 = \int (U_i - L_i)^2 d\mu(x) \leq \int 4\tau^2 d\mu(x) = 4\tau^2.$$

This implies

$$\|U_i - L_i\|_{L^2(\mu)} \leq 2\tau.$$

**Step 3: Bracketing Entropy.** From the definition of bracketing entropy, we have

$$H_B(2\tau, \mathcal{F}_N(\Theta), \|\cdot\|_{L^2(\mu)}) \leq \log V = \log V(\tau, \mathcal{F}_N(\Theta), \|\cdot\|_{L^\infty}),$$

leading to the bound:

$$H_B(2\tau, \mathcal{F}_N(\Theta), \|\cdot\|_{L^2(\mu)}) \lesssim \log V. \quad (15)$$

**Step 4: Covering Compact Sets.** Next, consider the parameter space  $\Theta$ , which we decompose into:

$$\Delta := \{(A, b, c) \in \mathbb{R}^{d \times d} \times \mathbb{R}^d \times \mathbb{R} : (A, b, c, \eta) \in \Theta\}, \quad \Omega := \{\eta \in \mathbb{R}^q : (A, b, c, \eta) \in \Theta\}.$$

Both  $\Delta$  and  $\Omega$  are compact sets. For any  $\tau > 0$ , there exist  $\tau$ -covers  $\Delta_\tau$  and  $\Omega_\tau$  such that:

$$|\Delta_\tau| \lesssim \mathcal{O}_P(\tau^{-(d^2+d+1)N}), \quad |\Omega_\tau| \lesssim \mathcal{O}_P(\tau^{-qN}).$$

**Step 5: Difference of the Functions.** Define  $\check{\eta}_i \in \Omega_\tau$  such that

$$\check{\eta}_i = \arg \min_{\eta \in \Omega_\tau} \|\eta - \eta_i\|,$$

and  $(\check{A}_i, \check{b}_i, \check{c}_i) \in \Delta_\tau$  such that

$$(\check{A}_i, \check{b}_i, \check{c}_i) = \arg \min_{(A, b, c) \in \Delta_\tau} \|(A, b, c) - (A_i, b_i, c_i)\|.$$

Now we consider mixing measures as

$$\tilde{G} = \sum_{i=1}^N \frac{1}{1 + \exp(-c_i)} \delta_{(A_i, b_i, \check{\eta}_i)}, \quad \check{G} = \sum_{i=1}^N \frac{1}{1 + \exp(-\check{c}_i)} \delta_{(\check{A}_i, \check{b}_i, \check{\eta}_i)}$$

given  $G = \sum_{i=1}^N \frac{1}{1 + \exp(-c_i)} \delta_{(A_i, b_i, \eta_i)} \in \mathcal{M}_N(\Theta)$  as a mixing measure. The above formulations imply that

$$\begin{aligned} \|f_G - f_{\tilde{G}}\|_{L^\infty} &= \sup_{x \in \mathcal{X}} \left| \sum_{i=1}^N \frac{1}{1 + \exp(-x^\top A_i x - (b_i)^\top x - c_i)} \cdot [\mathcal{E}(x, \eta_i) - \mathcal{E}(x, \check{\eta}_i)] \right| \\ &\leq \sum_{i=1}^N \sup_{x \in \mathcal{X}} \frac{1}{1 + \exp(-x^\top A_i x - (b_i)^\top x - c_i)} \cdot |\mathcal{E}(x, \eta_i) - \mathcal{E}(x, \check{\eta}_i)| \\ &\leq \sum_{i=1}^N \sup_{x \in \mathcal{X}} |\mathcal{E}(x, \eta_i) - \mathcal{E}(x, \check{\eta}_i)| \\ &\leq \sum_{i=1}^N \sup_{x \in \mathcal{X}} L_1 \cdot \|\eta_i - \check{\eta}_i\| \\ &\leq N L_1 \tau \lesssim \tau. \end{aligned}$$

The second inequality holds because the sigmoid weight is bounded by one, while the third inequality stems from the expert function  $\mathcal{E}(x, \cdot)$  being Lipschitz continuous with a Lipschitz constant  $L_1 > 0$ .

Then we will have

$$\begin{aligned}
& \|f_{\tilde{G}} - f_{\check{G}}\|_{L^\infty} \\
&= \sup_{x \in \mathcal{X}} \left| \sum_{i=1}^N \left[ \frac{1}{1 + \exp(-x^\top A_i x - (b_i)^\top x - c_i)} - \frac{1}{1 + \exp(-x^\top \check{A}_i x - (\check{b}_i)^\top x - \check{c}_i)} \right] \cdot \mathcal{E}(x, \check{\eta}_i) \right| \\
&\leq \sum_{i=1}^N \sup_{x \in \mathcal{X}} \left[ \frac{1}{1 + \exp(-x^\top A_i x - (b_i)^\top x - c_i)} - \frac{1}{1 + \exp(-x^\top \check{A}_i x - (\check{b}_i)^\top x - \check{c}_i)} \right] \cdot |\mathcal{E}(x, \check{\eta}_i)| \\
&\leq \sum_{i=1}^N \sup_{x \in \mathcal{X}} M' \left[ \frac{1}{1 + \exp(-x^\top A_i x - (b_i)^\top x - c_i)} - \frac{1}{1 + \exp(-x^\top \check{A}_i x - (\check{b}_i)^\top x - \check{c}_i)} \right] \\
&\leq \sum_{i=1}^N \sup_{x \in \mathcal{X}} M' L_2 (\|A_i - \check{A}_i\| \cdot \|x\|^2 + \|b_i - \check{b}_i\| \cdot \|x\| + |c_i - \check{c}_i|) \\
&\leq NM'(\tau B^2 + \tau B + \tau) \lesssim \tau
\end{aligned}$$

The second inequality holds because the expert function is bounded, i.e.,  $|\mathcal{E}(x, \check{\eta}_i)| \leq M'$ . The third inequality follows from the sigmoid function being Lipschitz continuous with a Lipschitz constant  $L_2 > 0$ , and the fourth inequality arises due to the boundedness of the input space. Then we know that from the triangle inequality

$$\|f_G - f_{\check{G}}\|_{L^\infty} \leq \|f_G - f_{\tilde{G}}\|_{L^\infty} + \|f_{\tilde{G}} - f_{\check{G}}\|_{L^\infty} \leq \tau.$$

**Step 6: Conclusion.** Now we recall the definition of the covering number, we will have

$$N(\tau, \mathcal{F}_N(\Theta), \|\cdot\|_{L^\infty}) \leq |\Delta_\tau| \times |\Omega_\tau| \leq \mathcal{O}_P(n^{-(d^2+d+1)N}) \times \mathcal{O}_P(n^{-qN}) \leq \mathcal{O}_P(n^{-(d^2+d+1+q)N}).$$

Recall the equation (15), the bracketing entropy of  $\mathcal{F}_N(\Theta)$  under  $L^2(\mu)$  is bounded by  $\log(1/\epsilon)$  by setting  $\tau = \epsilon/2$ , completing the proof.  $\square$

## B.2 Proof of Theorem 1

*Proof.* In order to prove  $\|f_G - f_{G_*}\|_{L^2(\mu)} \gtrsim \mathcal{L}_1(G, G_*)$  for any  $G \in \mathcal{M}_N(\Theta)$ , it is sufficient to show that

$$\inf_{G \in \mathcal{M}_N(\Theta)} \frac{\|f_G - f_{G_*}\|_{L^2(\mu)}}{\mathcal{L}_1(G, G_*)} > 0. \quad (16)$$

To prove the above inequality, we consider two cases for the denominator  $\mathcal{L}_1(G, G_*)$ : either it lies within a ball  $B(0, \varepsilon)$  where the loss is sufficiently small, or it falls outside this region, where  $\mathcal{L}_1(G, G_*)$  will not vanish.

**Local part:** At first, we focus on that

$$\lim_{\varepsilon \rightarrow 0} \inf_{G \in \mathcal{M}_N(\Theta): \mathcal{L}_1(G, G_*) \leq \varepsilon} \frac{\|f_G - f_{G_*}\|_{L^2(\mu)}}{\mathcal{L}_1(G, G_*)} > 0. \quad (17)$$

Assume, for contradiction, that the above claim does not hold. Then there exists a sequence of mixing measures  $G_n = \sum_{i=1}^{N^*} \frac{1}{1+\exp(-c_i^n)} \delta_{(A_i^n, b_i^n, \eta_i^n)}$  in  $\mathcal{M}_N(\Theta)$  such that as  $n \rightarrow \infty$ , we get

$$\begin{cases} \mathcal{L}_{1n} := \mathcal{L}_1(G_n, G_*) \rightarrow 0, \\ \|f_{G_n} - f_{G_*}\|_{L^2(\mu)} / \mathcal{L}_{1n} \rightarrow 0. \end{cases} \quad (18)$$

Let us recall that

$$\begin{aligned} \mathcal{L}_{1n} &:= \sum_{j=1}^{\bar{N}} \left| \sum_{i \in \mathcal{A}_j} \frac{1}{1 + \exp(-c_i^n)} - \frac{1}{1 + \exp(-c_j^*)} \right| + \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} [\|\Delta A_{ij}^n\|^2 + \|\Delta b_{ij}^n\|^2 + \|\Delta \eta_{ij}^n\|^2] \\ &+ \sum_{j=\bar{N}+1}^{N^*} \sum_{i \in \mathcal{A}_j} [\|\Delta A_{ij}^n\| + \|\Delta b_{ij}^n\| + |\Delta c_{ij}^n| + \|\Delta \eta_{ij}^n\|], \end{aligned}$$

where  $\Delta A_{ij}^n := A_i^n - A_j^*$ ,  $\Delta b_{ij}^n := b_i^n - b_j^*$ ,  $\Delta c_{ij}^n := c_i^n - c_j^*$ ,  $\Delta \eta_{ij}^n := \eta_i^n - \eta_j^*$ . Since  $\mathcal{L}_{1n} \rightarrow 0$  as  $n \rightarrow 0$ , there are two different situation for parameter convergence:

- For  $j = 1, \dots, \bar{N}$ , parameters are over-fitted:  $\sum_{i \in \mathcal{A}_j} \frac{1}{1 + \exp(-c_i^n)} \rightarrow \frac{1}{1 + \exp(-c_j^*)}$  and  $(A_i^n, b_i^n, \eta_i^n) \rightarrow (A_j^*, b_j^*, \eta_j^*)$ ,  $\forall i \in \mathcal{A}_j$ ;
- For  $j = \bar{N} + 1, \dots, N^*$ , parameters are exact-fitted:  $(A_i^n, b_i^n, c_i^n, \eta_i^n) \rightarrow (A_j^*, b_j^*, c_j^*, \eta_j^*)$ ,  $\forall i \in \mathcal{A}_j$ .

**Step 1 - Taylor expansion:** In this step, we decompose the term  $f_{G_n}(x) - f_{G_*}(x)$  using a Taylor expansion. First, let us denote

$$\begin{aligned} f_{G_n}(x) - f_{G_*}(x) &= \sum_{j=1}^{\bar{N}} \left[ \sum_{i \in \mathcal{A}_j} \frac{1}{1 + \exp(-x^\top A_i^n x - (b_i^n)^\top x - c_i^n)} \cdot \mathcal{E}(x, \eta_i^n) - \frac{1}{1 + \exp(-x^\top A_j^* x - (b_j^*)^\top x - c_j^*)} \cdot \mathcal{E}(x, \eta_j^*) \right] \\ &+ \sum_{j=\bar{N}+1}^{N^*} \left[ \sum_{i \in \mathcal{A}_j} \frac{1}{1 + \exp(-x^\top A_i^n x - (b_i^n)^\top x - c_i^n)} \cdot \mathcal{E}(x, \eta_i^n) - \frac{1}{1 + \exp(-x^\top A_j^* x - (b_j^*)^\top x - c_j^*)} \cdot \mathcal{E}(x, \eta_j^*) \right] \\ &= \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} \left[ \frac{1}{1 + \exp(-x^\top A_i^n x - (b_i^n)^\top x - c_i^n)} \cdot \mathcal{E}(x, \eta_i^n) - \frac{1}{1 + \exp(-c_i^n)} \cdot \mathcal{E}(x, \eta_j^*) \right] := \text{I}_n \\ &+ \sum_{j=1}^{\bar{N}} \left[ \sum_{i \in \mathcal{A}_j} \frac{1}{1 + \exp(-c_i^n)} - \frac{1}{1 + \exp(-c_j^*)} \right] \cdot \mathcal{E}(x, \eta_j^*) := \text{II}_n \\ &+ \sum_{j=\bar{N}+1}^{N^*} \left[ \sum_{i \in \mathcal{A}_j} \frac{1}{1 + \exp(-x^\top A_i^n x - (b_i^n)^\top x - c_i^n)} \cdot \mathcal{E}(x, \eta_i^n) - \frac{1}{1 + \exp(-x^\top A_j^* x - (b_j^*)^\top x - c_j^*)} \cdot \mathcal{E}(x, \eta_j^*) \right] := \text{III}_n \end{aligned} \quad (19)$$

Let us denote  $\sigma(x, A, b, c) = 1/[1 + \exp(-x^\top A x - b^\top x - c)]$ , then  $\text{I}_n$ ,  $\text{II}_n$  and  $\text{III}_n$  could be denoted

as

$$\begin{aligned}
\text{I}_n &= \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} \left[ \frac{1}{1 + \exp(-x^\top A_i^n x - (b_i^n)^\top x - c_i^n)} \cdot \mathcal{E}(x, \eta_i^n) - \frac{1}{1 + \exp(-c_i^n)} \cdot \mathcal{E}(x, \eta_j^*) \right] \\
&= \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} [\sigma(x, A_i^n, b_i^n, c_i^n) \mathcal{E}(x, \eta_i^n) - \sigma(x, 0_{d \times d}, 0_d, c_i^n) \mathcal{E}(x, \eta_j^*)] \\
&= \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} \sum_{|\gamma|=1}^2 \frac{1}{\gamma!} (\Delta A_{ij}^n)^{\gamma_1} (\Delta b_{ij}^n)^{\gamma_2} (\Delta \eta_{ij}^n)^{\gamma_3} \frac{\partial^{|\gamma_1|+|\gamma_2|} \sigma}{\partial A^{\gamma_1} \partial b^{\gamma_2}}(x, 0_{d \times d}, 0_d, c_i^n) \frac{\partial^{|\gamma_3|} \mathcal{E}}{\partial \eta^{\gamma_3}}(x, \eta_j^*) + R_1(x), \quad (20)
\end{aligned}$$

$$\text{II}_n = \sum_{j=1}^{\bar{N}} \left[ \sum_{i \in \mathcal{A}_j} \frac{1}{1 + \exp(-c_i^n)} - \frac{1}{1 + \exp(-c_j^*)} \right] \cdot \mathcal{E}(x, \eta_j^*), \quad (21)$$

$$\begin{aligned}
\text{III}_n &= \sum_{j=\bar{N}+1}^{N^*} \left[ \sum_{i \in \mathcal{A}_j} \frac{1}{1 + \exp(-x^\top A_i^n x - (b_i^n)^\top x - c_i^n)} \cdot \mathcal{E}(x, \eta_i^n) - \frac{1}{1 + \exp(-x^\top A_j^* x - (b_j^*)^\top x - c_j^*)} \cdot \mathcal{E}(x, \eta_j^*) \right] \\
&= \sum_{j=\bar{N}+1}^{N^*} \left[ \sum_{i \in \mathcal{A}_j} \sigma(x, A_i^n, b_i^n, c_i^n) \mathcal{E}(x, \eta_i^n) - \sigma(x, A_j^*, b_j^*, c_j^*) \mathcal{E}(x, \eta_j^*) \right] \\
&= \sum_{j=\bar{N}+1}^{N^*} \sum_{|\tau|=1} \left[ \sum_{i \in \mathcal{A}_j} \frac{1}{\tau!} (\Delta A_{ij}^n)^{\tau_1} (\Delta b_{ij}^n)^{\tau_2} (\Delta c_{ij}^n)^{\tau_3} (\Delta \eta_{ij}^n)^{\tau_4} \right] \frac{\partial^{|\tau_1|+|\tau_2|+|\tau_3|} \sigma}{\partial A^{\tau_1} \partial b^{\tau_2} \partial c^{\tau_3}}(x, A_j^*, b_j^*, c_j^*) \frac{\partial^{|\tau_4|} \mathcal{E}}{\partial \eta^{\tau_4}}(x, \eta_j^*) + R_2(x), \quad (22)
\end{aligned}$$

where  $R_i(x), i = 1, 2$  are Taylor remainder such that  $R_i(x)/\mathcal{L}_{1n} \rightarrow 0$  as  $n \rightarrow \infty$  for  $i = [2]$ .

Now we could denote

$$K_{j,i,\gamma_1:3}^n = \frac{1}{\gamma!} (\Delta A_{ij}^n)^{\gamma_1} (\Delta b_{ij}^n)^{\gamma_2} (\Delta \eta_{ij}^n)^{\gamma_3}, \quad j \in [\bar{N}], i \in \mathcal{A}_j, \gamma \in \mathbb{R}^{d \times d} \times \mathbb{R}^d \times \mathbb{R}^q \quad (23)$$

$$L_j^n = \sum_{i \in \mathcal{A}_j} \frac{1}{1 + \exp(-c_i^n)} - \frac{1}{1 + \exp(-c_j^*)}, \quad j \in [\bar{N}] \quad (24)$$

$$T_{j,\tau_1:4}^n = \sum_{i \in \mathcal{A}_j} \frac{1}{\tau!} (\Delta A_{ij}^n)^{\tau_1} (\Delta b_{ij}^n)^{\tau_2} (\Delta c_{ij}^n)^{\tau_3} (\Delta \eta_{ij}^n)^{\tau_4}, \quad j \in \{\bar{N}+1, \dots, N^*\}, \tau \in \mathbb{R}^{d \times d} \times \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^q \quad (25)$$

where  $1 \leq \sum_{i=1}^3 |\gamma_i| \leq 2$  and  $\sum_{i=1}^4 |\tau_i| = 1$ . Also recall that for the last term,  $i$  was settled directly by  $j$  since all the parameters are exact-fitted when  $\bar{N}+1 \leq j \leq N^*$ , i.e.  $|\mathcal{A}_j| = 1$ .

Using these notations (23) - (25) we can now rewrite the difference  $f_{G_n}(x) - f_{G_*}(x)$  as

$$\begin{aligned}
f_{G_n}(x) - f_{G_*}(x) &= \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} \sum_{|\gamma|=1}^2 K_{j,i,\gamma_{1:3}}^n \frac{\partial^{|\gamma_1|+|\gamma_2|} \sigma}{\partial A^{\gamma_1} \partial b^{\gamma_2}}(x, 0_{d \times d}, 0_d, c_i^n) \frac{\partial^{|\gamma_3|} \mathcal{E}}{\partial \eta^{\gamma_3}}(x, \eta_j^*) + \sum_{j=1}^{\bar{N}} L_j^n \cdot \mathcal{E}(x, \eta_j^*) \\
&+ \sum_{j=\bar{N}+1}^{N^*} \sum_{|\tau|=1} T_{j,\tau_{1:4}}^n \frac{\partial^{|\tau_1|+|\tau_2|+|\tau_3|} \sigma}{\partial A^{\tau_1} \partial b^{\tau_2} \partial c^{\tau_3}}(x, A_j^*, b_j^*, c_j^*) \frac{\partial^{|\tau_4|} \mathcal{E}}{\partial \eta^{\tau_4}}(x, \eta_j^*) + R_1(x) + R_2(x).
\end{aligned} \tag{26}$$

**Step 2 - Non-vanishing coefficients:** Now we claim that at least one in the set

$$\mathcal{S} = \left\{ \frac{K_{j,i,\gamma_{1:3}}^n}{\mathcal{L}_{1n}}, \frac{L_j^n}{\mathcal{L}_{1n}}, \frac{T_{j,\tau_{1:4}}^n}{\mathcal{L}_{1n}} \right\}$$

will not vanish as  $n$  goes to infinity. We prove by contradiction that all of them converge to zero when  $n \rightarrow 0$ :

$$\frac{K_{j,i,\gamma_{1:3}}^n}{\mathcal{L}_{1n}} \rightarrow 0, \quad \frac{L_j^n}{\mathcal{L}_{1n}} \rightarrow 0, \quad \frac{T_{j,\tau_{1:4}}^n}{\mathcal{L}_{1n}} \rightarrow 0.$$

Then follows directly from  $L_j^n / \mathcal{L}_{1n} \rightarrow 0$  we could conclude that

$$\frac{1}{\mathcal{L}_{1n}} \sum_{j=1}^{\bar{N}} \left| \sum_{i \in \mathcal{A}_j} \frac{1}{1 + \exp(-c_i^n)} - \frac{1}{1 + \exp(-c_j^*)} \right| = \frac{1}{\mathcal{L}_{1n}} \sum_{j=1}^{\bar{N}} |L_j^n| \rightarrow 0. \tag{27}$$

Before consider other coefficients, for simplicity, we denote  $e_{d,u} = \overbrace{(0, \dots, 0, \underbrace{1}_{u\text{-th}}, 0, \dots, 0)}^{d\text{-tuple}} \in \mathbb{R}^d$  as a  $d$ -tuple with all components equal to 0, except the  $u$ -th, which is 1; and  $e_{d \times d, uv}$  as a  $d \times d$  matrix with all components equal to 0, except the element in the  $u$ -th row and  $v$ -th column, which is 1, i.e.  $e_{d \times d, uv} = \overbrace{(0_d^\top, \dots, 0_d^\top, \underbrace{e_{d,u}^\top, 0_d^\top, \dots, 0_d^\top}_{v\text{-th}})}^{d\text{-column}} \in \mathbb{R}^{d \times d}$ .

Now consider for arbitrary  $u, v \in [d]$ , let  $\gamma_1 = 2e_{d \times d, uv}$ ,  $\gamma_2 = 0_d$  and  $\gamma_3 = 0_q$ , we will have

$$\frac{1}{\mathcal{L}_{1n}} \left| \Delta(A_{ij}^n)^{(uv)} \right|^2 = \frac{1}{2\mathcal{L}_{1n}} \left| K_{j,i,2e_{d \times d, uv}, 0_d, 0_q}^n \right| \rightarrow 0, \quad n \rightarrow \infty.$$

Then by taking the summation of the term with  $u, v \in [d]$ , we will have

$$\frac{1}{\mathcal{L}_{1n}} \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} \|\Delta A_{ij}^n\|^2 = \frac{1}{2\mathcal{L}_{1n}} \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} \sum_{u=1}^d \sum_{v=1}^d |K_{j,i,2e_{d \times d, uv}, 0_d, 0_q}^n| \rightarrow 0. \tag{28}$$

For arbitrary  $u \in [d]$ , let  $\gamma_1 = 0_{d \times d}$ ,  $\gamma_2 = 2e_{d,u}$  and  $\gamma_3 = 0_q$ , we will have

$$\frac{1}{\mathcal{L}_{1n}} \left| \Delta(b_{ij}^n)^{(u)} \right|^2 = \frac{1}{2\mathcal{L}_{1n}} \left| K_{j,i,0_{d \times d}, 2e_{d,u}, 0_q}^n \right| \rightarrow 0, \quad n \rightarrow \infty,$$

Then by taking the summation of the previous term with  $u \in [d]$ , we will have

$$\frac{1}{\mathcal{L}_{1n}} \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} \|\Delta b_{ij}^n\|^2 = \frac{1}{2\mathcal{L}_{1n}} \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} \sum_{u=1}^d |K_{j,i,0_{d \times d},2e_{d,u},0_q}^n| \rightarrow 0. \quad (29)$$

Similarly we will have

$$\frac{1}{\mathcal{L}_{1n}} \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} \|\Delta \eta_{ij}^n\|^2 = \frac{1}{2\mathcal{L}_{1n}} \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} \sum_{w=1}^q |K_{j,i,0_{d \times d},0_d,2e_{q,w}}^n| \rightarrow 0. \quad (30)$$

Now, consider  $\bar{N} + 1 \leq j \leq N^*$ , such that its corresponding Voronoi cell has only one element, i.e.  $|\mathcal{A}_j| = 1$ . For arbitrary  $u, v \in [d]$ , let  $\tau_1 = e_{d \times d,uv}, \tau_2 = 0_d, \tau_3 = 0$  and  $\tau_4 = 0_q$ , we will have

$$\frac{1}{\mathcal{L}_{1n}} \sum_{i \in \mathcal{A}_j} \left| \Delta (A_{ij}^n)^{(uv)} \right| = \frac{1}{\mathcal{L}_{1n}} \left| T_{j,e_{d \times d,uv},0_d,0,0_q}^n \right| \rightarrow 0, \quad n \rightarrow \infty.$$

Then by taking the summation of the term with  $u, v \in [d]$ , we will have

$$\frac{1}{\mathcal{L}_{1n}} \sum_{j=\bar{N}+1}^{N^*} \sum_{i \in \mathcal{A}_j} \|\Delta A_{ij}^n\|_1 = \frac{1}{\mathcal{L}_{1n}} \sum_{j=\bar{N}+1}^{N^*} \sum_{u=1}^d \sum_{v=1}^d |T_{j,e_{d \times d,uv},0_d,0,0_q}^n| \rightarrow 0.$$

Recall the topological equivalence between  $L_1$ -norm and  $L_2$ -norm on finite-dimensional vector space over  $\mathbb{R}$ , we will have

$$\frac{1}{\mathcal{L}_{1n}} \sum_{j=\bar{N}+1}^{N^*} \sum_{i \in \mathcal{A}_j} \|\Delta A_{ij}^n\| \rightarrow 0. \quad (31)$$

Following a similar argument, since

$$\begin{aligned} \frac{1}{\mathcal{L}_{1n}} \sum_{j=\bar{N}+1}^{N^*} \sum_{i \in \mathcal{A}_j} \|\Delta b_{ij}^n\|_1 &= \frac{1}{\mathcal{L}_{1n}} \sum_{j=\bar{N}+1}^{N^*} \sum_{u=1}^d |T_{j,0_{d \times d},e_{d,u},0,0_q}^n| \rightarrow 0, \\ \frac{1}{\mathcal{L}_{1n}} \sum_{j=\bar{N}+1}^{N^*} \sum_{i \in \mathcal{A}_j} |\Delta c_{ij}^n| &= \frac{1}{\mathcal{L}_{1n}} \sum_{j=\bar{N}+1}^{N^*} |T_{j,0_{d \times d},0_d,1,0_q}^n| \rightarrow 0, \\ \frac{1}{\mathcal{L}_{1n}} \sum_{j=\bar{N}+1}^{N^*} \sum_{i \in \mathcal{A}_j} \|\Delta \eta_{ij}^n\|_1 &= \frac{1}{\mathcal{L}_{1n}} \sum_{j=\bar{N}+1}^{N^*} \sum_{w=1}^q |T_{j,0_{d \times d},0_d,0,e_{q,w}}^n| \rightarrow 0, \end{aligned}$$

we obtain that

$$\frac{1}{\mathcal{L}_{1n}} \sum_{j=\bar{N}+1}^{N^*} \sum_{i \in \mathcal{A}_j} \|\Delta b_{ij}^n\| \rightarrow 0, \quad \frac{1}{\mathcal{L}_{1n}} \sum_{j=\bar{N}+1}^{N^*} \sum_{i \in \mathcal{A}_j} |\Delta c_{ij}^n| \rightarrow 0, \quad \frac{1}{\mathcal{L}_{1n}} \sum_{j=\bar{N}+1}^{N^*} \sum_{i \in \mathcal{A}_j} \|\Delta \eta_{ij}^n\| \rightarrow 0. \quad (32)$$

Now taking the summation of limits in equations (27) - (32), we could deduce that

$$1 = \frac{\mathcal{L}_{1n}}{\mathcal{L}_{1n}} = \frac{1}{\mathcal{L}_{1n}} \sum_{j=1}^{\bar{N}} \left| \sum_{i \in \mathcal{A}_j} \frac{1}{1 + \exp(-c_i^n)} - \frac{1}{1 + \exp(-c_j^*)} \right| + \frac{1}{\mathcal{L}_{1n}} \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} [\|\Delta A_{ij}^n\|^2 + \|\Delta b_{ij}^n\|^2 + \|\Delta \eta_{ij}^n\|^2] \\ + \frac{1}{\mathcal{L}_{1n}} \sum_{j=\bar{N}+1}^{N^*} \sum_{i \in \mathcal{A}_j} [\|\Delta A_{ij}^n\| + \|\Delta b_{ij}^n\| + |\Delta c_{ij}^n| + \|\Delta \eta_{ij}^n\|] \rightarrow 0,$$

as  $n \rightarrow \infty$ , which is a contradiction. Thus, not all the coefficients of elements in the set

$$\mathcal{S} = \left\{ \frac{K_{j,i,\gamma_{1:3}}^n}{\mathcal{L}_{1n}}, \frac{L_j^n}{\mathcal{L}_{1n}}, \frac{T_{j,\tau_{1:4}}^n}{\mathcal{L}_{1n}} \right\}$$

tend to 0 as  $n \rightarrow \infty$ . Let us denote by  $m_n$  the maximum of the absolute values of those elements. It follows from the previous result that  $1/m_n \not\rightarrow \infty$  as  $n \rightarrow \infty$ .

**Step 3 - Application of Fatou's lemma:** In this step, we apply Fatou's lemma to obtain the desired inequality in equation (17). Recall in the beginning we have assumed that  $\|f_{G_n} - f_{G_*}\|_{L^2(\mu)/\mathcal{L}_{1n}} \rightarrow 0$  and the topological equivalence between  $L_1$ -norm and  $L_2$ -norm on finite-dimensional vector space over  $\mathbb{R}$ , we obtain  $\|f_{G_n} - f_{G_*}\|_{L^1(\mu)/\mathcal{L}_{1n}} \rightarrow 0$ . By applying the Fatou's lemma, we get

$$0 = \lim_{n \rightarrow \infty} \frac{\|f_{G_n} - f_{G_*}\|_{L^1(\mu)}}{m_n \mathcal{L}_{1n}} \geq \int \liminf_{n \rightarrow \infty} \frac{|f_{G_n}(x) - f_{G_*}(x)|}{m_n \mathcal{L}_{1n}} d\mu(x) \geq 0.$$

This result suggests that for almost every  $x$ ,

$$\frac{f_{G_n}(x) - f_{G_*}(x)}{m_n \mathcal{L}_{1n}} \rightarrow 0, \quad (33)$$

recall equation (26), we deduce that

$$\frac{f_{G_n}(x) - f_{G_*}(x)}{m_n \mathcal{L}_{1n}} = \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} \sum_{|\gamma|=1}^2 \frac{K_{j,i,\gamma_{1:3}}^n}{m_n \mathcal{L}_{1n}} \frac{\partial^{|\gamma_1|+|\gamma_2|} \sigma}{\partial A^{\gamma_1} \partial b^{\gamma_2}}(x, 0_{d \times d}, 0_d, c_i^n) \frac{\partial^{|\gamma_3|} \mathcal{E}}{\partial \eta^{\gamma_3}}(x, \eta_j^*) + \sum_{j=1}^{\bar{N}} \frac{L_j^n}{m_n \mathcal{L}_{1n}} \cdot \mathcal{E}(x, \eta_j^*) \\ + \sum_{j=\bar{N}+1}^{N^*} \sum_{|\tau|=1} \frac{T_{j,\tau_{1:4}}^n}{m_n \mathcal{L}_{1n}} \frac{\partial^{|\tau_1|+|\tau_2|+|\tau_3|} \sigma}{\partial A^{\tau_1} \partial b^{\tau_2} \partial c^{\tau_3}}(x, A_j^*, b_j^*, c_j^*) \frac{\partial^{|\tau_4|} \mathcal{E}}{\partial \eta^{\tau_4}}(x, \eta_j^*) + \frac{R_1(x)}{m_n \mathcal{L}_{1n}} + \frac{R_2(x)}{m_n \mathcal{L}_{1n}}.$$

Let us denote

$$\frac{K_{j,i,\gamma_{1:3}}^n}{m_n \mathcal{L}_{1n}} \rightarrow k_{j,i,\gamma_{1:3}}, \quad \frac{L_j^n}{m_n \mathcal{L}_{1n}} \rightarrow l_j, \quad \frac{T_{j,\tau_{1:4}}^n}{m_n \mathcal{L}_{1n}} \rightarrow t_{j,\tau_{1:4}},$$

then from equation (33), we will have

$$\sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} \sum_{|\gamma|=1}^2 k_{j,i,\gamma_{1:3}} \frac{\partial^{|\gamma_1|+|\gamma_2|} \sigma}{\partial A^{\gamma_1} \partial b^{\gamma_2}}(x, 0_{d \times d}, 0_d, c_i^n) \frac{\partial^{|\gamma_3|} \mathcal{E}}{\partial \eta^{\gamma_3}}(x, \eta_j^*) + \sum_{j=1}^{\bar{N}} l_j^n \cdot \mathcal{E}(x, \eta_j^*) \\ + \sum_{j=\bar{N}+1}^{N^*} \sum_{|\tau|=1} t_{j,\tau_{1:4}}^n \frac{\partial^{|\tau_1|+|\tau_2|+|\tau_3|} \sigma}{\partial A^{\tau_1} \partial b^{\tau_2} \partial c^{\tau_3}}(x, A_j^*, b_j^*, c_j^*) \frac{\partial^{|\tau_4|} \mathcal{E}}{\partial \eta^{\tau_4}}(x, \eta_j^*) = 0,$$

for almost every  $x$ . Note that the expert function  $\mathcal{E}(\cdot, \eta)$  is strongly identifiable, then the above equation implies that

$$k_{j,i,\gamma_{1:3}} = l_j = t_{j,\tau_{1:4}} = 0,$$

for any  $j \in [N^*]$ ,  $(\gamma_1, \gamma_2, \gamma_3) \in \mathbb{N}^{d \times d} \times \mathbb{N}^d \times \mathbb{N}^q$  and  $(\tau_1, \tau_2, \tau_3, \tau_4) \in \mathbb{N}^{d \times d} \times \mathbb{N}^d \times \mathbb{N} \times \mathbb{N}^q$  such that  $1 \leq \sum_{i=1}^3 |\gamma_i| \leq 2$  and  $\sum_{i=1}^4 |\tau_i| = 1$ . This violates that at least one among the limits in the set  $\{k_{j,i,\gamma_{1:3}}, l_j, t_{j,\tau_{1:4}}\}$  is different from zero.

Thus, we obtain the local inequality in equation (17). Consequently, there exists some  $\varepsilon' > 0$  such that

$$\inf_{G \in \mathcal{M}_N(\Theta): \mathcal{L}_1(G, G_*) \leq \varepsilon'} \frac{\|f_G - f_{G_*}\|_{L^2(\mu)}}{\mathcal{L}_1(G, G_*)} > 0.$$

**Global part:** We now proceed to demonstrate equation (16) for the case where the denominator does not vanish, i.e.

$$\inf_{G \in \mathcal{M}_N(\Theta): \mathcal{L}_1(G, G_*) > \varepsilon'} \frac{\|f_G - f_{G_*}\|_{L^2(\mu)}}{\mathcal{L}_1(G, G_*)} > 0. \quad (34)$$

Suppose, for contradiction, that inequality (34) does not hold. Then there exists a sequence of mixing measures  $G'_n \in \mathcal{M}_N(\Theta)$  such that  $\mathcal{L}_1(G'_n, G_*) > \varepsilon'$  and

$$\lim_{n \rightarrow \infty} \frac{\|f_{G'_n} - f_{G_*}\|_{L^2(\mu)}}{\mathcal{L}_1(G'_n, G_*)} = 0.$$

Under this situation, we could deduce that  $\|f_{G'_n} - f_{G_*}\|_{L^2(\mu)} \rightarrow 0$  as  $n \rightarrow \infty$ . Since  $\Theta$  is a compact set, we can replace the sequence  $G'_n$  with a convergent subsequence, which approaches a mixing measure  $G' \in \mathcal{M}_N(\Theta)$ . Given that  $\mathcal{L}_1(G'_n, G_*) > \varepsilon'$ , we conclude that  $\mathcal{L}_1(G', G_*) > \varepsilon'$ . Then, using Fatou's lemma, we deduce:

$$0 = \lim_{n \rightarrow \infty} \|f_{G'_n} - f_{G_*}\|_{L^2(\mu)}^2 \geq \int \liminf_{n \rightarrow \infty} |f_{G'_n}(x) - f_{G_*}(x)|^2 d\mu(x),$$

which indicates that  $f_{G'}(x) = f_{G_*}(x)$  for almost every  $x$ . Based on the Proposition 2 followed, we will have that  $G' \equiv G_*$ . □

**Proposition 2.** *If the equation  $f_G(x) = f_{G_*}(x)$  holds true for almost every  $x$ , then it follows that  $G \equiv G_*$ .*

*Proof of Proposition 2.* Assume that  $f_G(x) = f_{G_*}(x)$  for almost every  $x$ , which gives:

$$\sum_{i=1}^N \frac{1}{1 + \exp(-x^\top A_i x - b_i^\top x - c_i)} \mathcal{E}(x, \eta_i) = \sum_{i=1}^{N_*} \frac{1}{1 + \exp(-x^\top A_i^* x - (b_i^*)^\top x - c_i^*)} \mathcal{E}(x, \eta_i^*). \quad (35)$$

Since  $\mathcal{E}(x, \eta)$  is identifiable, the set  $\{\mathcal{E}(x, \eta_i), i \in [N]\}$  is linearly independent, and  $\{\mathcal{E}(x, \eta_i^*), i \in [N_*]\}$  is distinct for some  $N_* \in \mathbb{N}$ . Therefore, if  $N \neq N_*$ , there must exist some index  $i \in [N]$  such that  $\eta_i \neq \eta_j^*$  for any  $j \in [N_*]$ . This leads to a contradiction as the coefficients cannot simultaneously satisfy equation (35). Consequently, we must have  $N = N_*$ .

Thus, the sets of weights and gating functions on both sides of (35) are equivalent:

$$\left\{ \frac{1}{1 + \exp(-x^\top A_i x - b_i^\top x - c_i)}, i \in [N] \right\} = \left\{ \frac{1}{1 + \exp(-x^\top A_i^* x - (b_i^*)^\top x - c_i^*)}, i \in [N_*] \right\}.$$

Without loss of generality, we can assume that the correspondence is such that

$$\frac{1}{1 + \exp(-x^\top A_i x - b_i^\top x - c_i)} = \frac{1}{1 + \exp(-x^\top A_i^* x - (b_i^*)^\top x - c_i^*)},$$

for all  $i \in [N]$  and for almost every  $x$ . This implies the sigmoid function's invariance to translations, leading to:

$$A_i = A_i^* + v_2, \quad b_i = b_i^* + v_1, \quad c_i = c_i^* + v_0,$$

for some  $v_2 \in \mathbb{R}^{d \times d}$ ,  $v_1 \in \mathbb{R}^d$  and  $v_0 \in \mathbb{R}$ . However, due to the assumption  $A_k = A_k^* = 0_{d \times d}$ ,  $b_k = b_k^* = 0_d$  and  $c_k = c_k^* = 0$ , it follows that  $v_2 = 0_{d \times d}$ ,  $v_1 = 0_d$  and  $v_0 = 0$ , leading to:

$$A_i = A_i^*, \quad b_i = b_i^*, \quad c_i = c_i^*, \quad \text{for any } i \in [N].$$

Substituting this back into equation (35), we have:

$$\sum_{i=1}^N \frac{1}{1 + \exp(-x^\top A_i x - b_i^\top x - c_i)} \mathcal{E}(x, \eta_i) = \sum_{i=1}^N \frac{1}{1 + \exp(-x^\top A_i^* x - (b_i^*)^\top x - c_i^*)} \mathcal{E}(x, \eta_i^*). \quad (36)$$

Next, let us partition the index set  $[N]$  into subsets  $J_1, J_2, \dots, J_m$ , where  $m \leq k$ , such that  $(A_i, b_i, c_i) = (A_{i'}^*, b_{i'}^*, c_{i'}^*)$  for any  $i, i' \in J_j$  and  $j \in [m]$ . For indices  $i$  and  $i'$  belonging to distinct subsets, it holds that  $(A_i, b_i, c_i) \neq (A_{i'}, b_{i'}, c_{i'})$ .

Using this partition, the equation can be reorganized as:

$$\sum_{j=1}^m \sum_{i \in J_j} \frac{1}{1 + \exp(-x^\top A_i x - b_i^\top x - c_i)} \mathcal{E}(x, \eta_i) = \sum_{j=1}^m \sum_{i \in J_j} \frac{1}{1 + \exp(-x^\top A_i^* x - (b_i^*)^\top x - c_i^*)} \mathcal{E}(x, \eta_i^*), \quad (37)$$

for almost every  $x$ .

Recall that  $A_i = A_i^*$ ,  $b_i = b_i^*$  and  $c_i = c_i^*$  for all  $i \in [N^*]$ . This implies that for any  $j \in [m]$ , we can identify the sets:

$$\{\eta_i : i \in J_j\} \equiv \{\eta_i^* : i \in J_j\}.$$

Consequently, we obtain:

$$G = \sum_{j=1}^m \sum_{i \in J_j} \frac{1}{1 + \exp(-c_i)} \delta(A_i, b_i, \eta_i) = \sum_{j=1}^m \sum_{i \in J_j} \frac{1}{1 + \exp(-c_i^*)} \delta(A_i^*, b_i^*, \eta_i^*) = G_*,$$

showing  $G \equiv G_*$ . Thus, the proof is complete.  $\square$

### B.3 Proof of Theorem 2

To prove the result, we first examine the Voronoi loss:

$$\begin{aligned}\mathcal{L}_{2,r}(G, G_*) &:= \sum_{j=1}^{\bar{N}} \left| \sum_{i \in \mathcal{A}_j} \frac{1}{1 + \exp(-c_i)} - \frac{1}{1 + \exp(-c_j^*)} \right| \\ &\quad + \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} [\|\Delta A_{ij}\|^r + \|\Delta b_{ij}\|^r + \|\Delta \alpha_{ij}\|^r + |\Delta \beta_{ij}|^r] \\ &\quad + \sum_{j=\bar{N}+1}^{N^*} \sum_{i \in \mathcal{A}_j} [\|\Delta A_{ij}\|^r + \|\Delta b_{ij}\|^r + |\Delta c_{ij}|^r + \|\Delta \alpha_{ij}\|^r + |\Delta \beta_{ij}|^r]\end{aligned}$$

where the difference terms  $\Delta A_{ij} := A_i - A_j^*$ ,  $\Delta b_{ij} := b_i - b_j^*$ ,  $\Delta c_{ij} := c_i - c_j^*$ ,  $\Delta \alpha_{ij} := \alpha_i - \alpha_j^*$ ,  $\Delta \beta_{ij} := \beta_i - \beta_j^*$ . These terms are asymmetric by definition because  $\Delta A_{ij} \neq \Delta A_{ji}$  in general (i.e.,  $A_i - A_j^* \neq A_j - A_i^*$  where  $i \in \mathcal{A}_j$ ). So  $\mathcal{L}_{2,r}(G, G_*)$  is not symmetric, but it still satisfies a weak triangle inequality. For experts meeting all the assumptions in Theorem 2, we derive the following results using Taylor expansion:

**Lemma 2.** *Given experts in Theorem 2, we achieve for any  $r \geq 1$  that*

$$\lim_{\epsilon \rightarrow 0} \inf_{G \in \mathcal{M}_N(\Theta)} \left\{ \frac{\|f_G - f_{G_*}\|_{L^2(\mu)}}{\mathcal{L}_{2,r}(G, G_*)} : \mathcal{L}_{2,r}(G, G_*) \leq \epsilon \right\} = 0.$$

We will prove this lemma later.

*Main Proof.* Recall that  $(X_1, Y_1), (X_2, Y_2), \dots, (X_n, Y_n) \in \mathbb{R}^d \times \mathbb{R}$  follow a standard regression model

$$Y_i = f_{G_*}(X_i) + \varepsilon_i, \quad i = 1, \dots, n,$$

where  $X_1, X_2, \dots, X_n$  are i.i.d. samples from a probability distribution  $\mu$  on  $\mathbb{R}^d$ , and  $\varepsilon_i$  are i.i.d. Gaussian noise variables with  $\mathbb{E}[\varepsilon_i | X_i] = 0$  and  $\text{Var}[\varepsilon_i | X_i] = \nu$ ,  $i \in [n]$ . Under the Gaussian assumption on the noise variables  $\varepsilon_i$ , we have that  $Y_i | X_i \sim \mathcal{N}(f_{G_*}(X_i), \nu)$ ,  $i = 1, \dots, n$ . Given Lemma 2, there exists a sufficiently small  $\epsilon > 0$  and a mixing measure  $G'_* \in \mathcal{M}_N(\Theta)$  such that  $\mathcal{L}_{2,r}(G'_*, G_*) = 2\epsilon$  and  $\|f_{G'_*} - f_{G_*}\| \leq \sqrt{C_1} \epsilon$  for a fixed constant  $C_1$ . Thus, applying Le Cam's lemma to address the  $\mathcal{L}_{2,r}$  loss, which satisfies the weak triangle inequality, we obtain:

$$\begin{aligned}&\inf_{\tilde{G}_n \in \mathcal{M}_N(\Theta)} \sup_{G \in \mathcal{M}_N(\Theta) \setminus \mathcal{M}_{N^*-1}(\Theta)} \mathbb{E}_{f_G} [\mathcal{L}_{2,r}(\tilde{G}_n, G)] \\ &\gtrsim \frac{\mathcal{L}_{2,r}(G'_*, G_*)}{8} \exp \left\{ -n \mathbb{E}_{X \sim \mu} [\text{KL}(\mathcal{N}(f_{G'_*}(X), \nu), \mathcal{N}(f_{G_*}(X), \nu))] \right\}\end{aligned}$$

Recall that the KL divergence between two Gaussian distributions with the same variance is

$$\text{KL}(\mathcal{N}(f_{G'_*}(X), \nu), \mathcal{N}(f_{G_*}(X), \nu)) = \frac{1}{2\nu} (f_{G'_*}(X) - f_{G_*}(X))^2.$$

Hence, we can deduce that:

$$\begin{aligned} \inf_{\tilde{G}_n \in \mathcal{M}_N(\Theta)} \sup_{G \in \mathcal{M}_N(\Theta) \setminus \mathcal{M}_{N^*-1}(\Theta)} \mathbb{E}_{f_G} [\mathcal{L}_{2,r}(\tilde{G}_n, G)] &\gtrsim \epsilon \exp \left( -n \|f_{G_*} - f_{G_*}\|_{L^2(\mu)}^2 \right) \\ &\gtrsim \epsilon \exp(-nC_1\epsilon^2). \end{aligned}$$

By setting  $\epsilon = n^{-1/2}$ , we obtain  $\epsilon \exp(-nC_1\epsilon^2) = n^{-1/2} \exp(-C_1)$ , which means that

$$\inf_{\tilde{G}_n \in \mathcal{M}_N(\Theta)} \sup_{G \in \mathcal{M}_N(\Theta) \setminus \mathcal{M}_{N^*-1}(\Theta)} \mathbb{E}_{f_G} [\mathcal{L}_{2,r}(\tilde{G}_n, G)] \gtrsim n^{-\frac{1}{2}},$$

Consequently, we establish the result for Theorem 2.  $\square$

*Proof of Lemma 2.* Now we want to demonstrate that the following limit holds true for any  $r \geq 1$ :

$$\lim_{\epsilon \rightarrow 0} \inf_{G \in \mathcal{M}_N(\Theta) : \mathcal{L}_{2,r}(G, G_*) \leq \epsilon} \frac{\|f_G - f_{G_*}\|_{L^2(\mu)}}{\mathcal{L}_{2,r}(G, G_*)} = 0. \quad (38)$$

To achieve this, we need to construct a sequence of mixing measures  $(G_n)$  that satisfies  $\mathcal{L}_{2,r}(G_n, G_*) \rightarrow 0$  and

$$\frac{\|f_G - f_{G_*}\|_{L^2(\mu)}}{\mathcal{L}_{2,r}(G, G_*)} \rightarrow 0,$$

as  $n \rightarrow \infty$ . Recall we consider that at least one expert parameter  $\alpha_i^*$ , where  $i \in \mathcal{A}_j, j \in [N^*]$  in the over-specified gating parameters Voronoi cell equals  $0_d$ . WLOG, assume  $\alpha_1^* = 0_d$ . Next, we consider the sequence  $(G_n)$  with  $k^* + 1$  atoms, in which for over-specified parameters  $i = 1, 2$ :  $A_i^n = A_1^* = 0_{d \times d}$ ,  $b_i^n = b_1^* = 0_d$ ,  $c_1^n = c_2^n$  such that

$$\sum_{i=1}^2 \frac{1}{1 + \exp(-c_i^n)} = \frac{1}{1 + \exp(-c_1^*)} + \frac{1}{n^{r+1}},$$

$\alpha_i^n = \alpha_1^* = 0_d$ ,  $\beta_1^n = \beta_1^* + 1/n$  and  $\beta_2^n = \beta_1^* - 1/n$ . And for exactly-specified parameters  $i = 3, \dots, N^* + 1$ :  $A_i^n = A_{i-1}^*$ ,  $b_i^n = b_{i-1}^*$ ,  $c_i^n = c_{i-1}^*$ ,  $\alpha_i^n = \alpha_{i-1}^*$ ,  $\beta_i^n = \beta_{i-1}^*$ . Recall the construction of  $\mathcal{L}_{2,r}$  loss, we will have

$$\mathcal{L}_{2,r}(G_n, G_*) = \frac{1}{n^{r+1}} + \frac{2}{n^r} = \mathcal{O}(n^{-r}).$$

Now recall the Taylor expansion in equation (19) for  $f_{G_n}(x) - f_{G_*}(x)$ :

$$\begin{aligned} f_{G_n}(x) - f_{G_*}(x) &= \sum_{i=1}^2 \left[ \frac{(\alpha_i^n)^\top x + \beta_i^n}{1 + \exp(-x^\top A_i^n x - (b_i^n)^\top x - c_i^n)} - \frac{(\alpha_1^*)^\top x + \beta_1^*}{1 + \exp(-c_1^*)} \right] := \text{I}_n \\ &+ \left[ \sum_{i=1}^2 \frac{1}{1 + \exp(-c_i^n)} - \frac{1}{1 + \exp(-c_1^*)} \right] \cdot \left[ (\alpha_1^*)^\top x + \beta_1^* \right] := \text{II}_n \\ &+ \sum_{i=3}^{N^*+1} \left[ \frac{(\alpha_i^n)^\top x + \beta_i^n}{1 + \exp(-x^\top A_i^n x - (b_i^n)^\top x - c_i^n)} - \frac{(\alpha_{i-1}^*)^\top x + \beta_{i-1}^*}{1 + \exp(-x^\top A_{i-1}^* x - (b_{i-1}^*)^\top x - c_{i-1}^*)} \right] := \text{III}_n. \end{aligned}$$

From our construction for the sequence  $(A_i^n, b_i^n, c_i^n, \alpha_i^n, \beta_i^n)_{i=1}^{N^*+1}$ , it's easy to verify that

$$f_{G_n}(x) - f_{G_*}(x) = \sum_i^2 \left[ \frac{\beta_i^n}{1 + \exp(-c_i^n)} - \frac{\beta_1^*}{1 + \exp(-c_i^n)} \right] + \frac{1}{n^{r+1}} \left[ (\alpha_1^*)^\top x + \beta_1^* \right] = \frac{\beta_1^*}{n^{r+1}}.$$

Based on the above result, we conclude that  $[f_{G_n}(x) - f_{G_*}(x)]/\mathcal{L}_{2,r}(G_n, G_*) \rightarrow 0$  for almost every  $x$ . As a result,  $\|f_{G_n} - f_{G_*}\|_{L^2(\mu)} \mathcal{L}_{2,r}(G_n, G_*) \rightarrow 0$  as  $n \rightarrow \infty$ . This establishes the claim stated in equation (38).  $\square$

## B.4 Proof of Theorem 3

*Proof.* Following from the result of Theorem 1, it is sufficient to show that the following inequality holds true:

$$\inf_{G \in \mathcal{M}_N(\Theta)} \frac{\|f_G - f_{\bar{G}}\|_{L^2(\mu)}}{\mathcal{L}_3(G, \bar{G})} > 0, \quad (39)$$

for any mixing measure  $\bar{G} \in \overline{\mathcal{M}}_N(\Theta)$ . To prove the above inequality, we follow a similar approach to the proof in Appendix B.2, dividing the analysis into a local part and a global part. However, since the arguments for the global part remain the same (up to some notational changes) in the over-specified setting, they are omitted.

Therefore, for an arbitrary mixing measure  $\bar{G} := \sum_{i=1}^k \frac{1}{1+\exp(-\bar{c}_i)} \delta_{(\bar{A}_i, \bar{b}_i, \bar{\eta}_i)} \in \overline{\mathcal{M}}_N(\Theta)$ , we focus exclusively on demonstrating that:

$$\lim_{\varepsilon \rightarrow 0} \inf_{G \in \mathcal{M}_N(\Theta): \mathcal{L}_3(G, \bar{G}) \leq \varepsilon} \frac{\|f_G - f_{\bar{G}}\|_{L^2(\mu)}}{\mathcal{L}_3(G, \bar{G})} > 0. \quad (40)$$

Assume, for contradiction, that the above claim does not hold. Then there exists a sequence of mixing measures  $G_n = \sum_{i=1}^k \frac{1}{1+\exp(-c_i^n)} \delta_{(A_i^n, b_i^n, \eta_i^n)}$  in  $\mathcal{M}_N(\Theta)$  such that as  $n \rightarrow \infty$ , we get

$$\begin{cases} \mathcal{L}_{3n} := \mathcal{L}_3(G_n, \bar{G}) \rightarrow 0, \\ \|f_{G_n} - f_{\bar{G}}\|_{L^2(\mu)} / \mathcal{L}_{3n} \rightarrow 0. \end{cases} \quad (41)$$

Let us denote by  $\mathcal{A}_i^n := \mathcal{A}_i(G_n)$  a Voronoi cell of  $G_n$  generated by the  $j$ -th components of  $\bar{G}$ . Since our analysis is asymptotic, we can assume that the Voronoi cells are independent of the sample size, i.e.  $\mathcal{A}_j = \mathcal{A}_j^n$ .

In Dense Regime, since  $G_n$  and  $\bar{G}$  have the same number of atoms  $N$ , and  $\mathcal{L}_{3n} \rightarrow 0$ , it follows that each Voronoi cell  $\mathcal{A}_i$  contains precisely one element for all  $i \in [N]$ . Without loss of generality, we assume  $\mathcal{A}_i = \{i\}$  for all  $i \in [N]$ . This ensures that  $(A_i^n, b_i^n, c_i^n, \eta_i^n) \rightarrow (\bar{A}_i, \bar{b}_i, \bar{c}_i, \bar{\eta}_i)$  as  $n \rightarrow \infty$  for every  $i \in [N]$ .

Consequently, the Voronoi loss  $\mathcal{L}_{3n}$  can be expressed as:

$$\mathcal{L}_{3n} := \sum_{i=1}^N \left( \|\Delta \bar{A}_i^n\| + \|\Delta \bar{b}_i^n\| + |\Delta \bar{c}_i^n| + \|\bar{\eta}_i^n\| \right), \quad (42)$$

where the increments are given by:  $\Delta \bar{A}_i^n = A_i^n - \bar{A}_i$ ,  $\Delta \bar{b}_i^n = b_i^n - \bar{b}_i$ ,  $\Delta \bar{c}_i^n = c_i^n - \bar{c}_i$ ,  $\Delta \bar{\eta}_i^n = \eta_i^n - \bar{\eta}_i$ .

We now break the proof of the local part into the following three steps:

**Step 1 - Taylor expansion:** In this step, we decompose the term  $f_{G_n}(x) - f_{\bar{G}}(x)$  using a Taylor expansion. First, let us denote  $\sigma(x, A, b, c) := \frac{1}{1 + \exp(-x^\top A x - b^\top x - c)}$ , then we have that

$$\begin{aligned}
& f_{G_n}(x) - f_{\bar{G}}(x) \\
&= \sum_{i=1}^N \left[ \frac{1}{1 + \exp(-x^\top \bar{A}_i^n x - (b_i^n)^\top x - c_i^n)} \cdot \mathcal{E}(x, \eta_i^n) - \frac{1}{1 + \exp(-x^\top \bar{A}_i x - (b_i)^\top x - \bar{c}_i)} \cdot \mathcal{E}(x, \bar{\eta}_i) \right] \\
&= \sum_{i=1}^N \sum_{|\alpha|=1} \frac{1}{\alpha!} (\Delta \bar{A}_i^n)^{\alpha_1} (\Delta \bar{b}_i^n)^{\alpha_2} (\Delta \bar{c}_i^n)^{\alpha_3} (\Delta \bar{\eta}_i^n)^{\alpha_4} \frac{\partial^{|\alpha_1|+|\alpha_2|+|\alpha_3|} \sigma}{\partial A^{\alpha_1} \partial b^{\alpha_2} \partial c^{\alpha_3}}(x, \bar{A}_i, \bar{b}_i, \bar{c}_i) \frac{\partial^{|\alpha_4|} \mathcal{E}}{\partial \eta^{\alpha_4}}(x, \bar{\eta}_i) + R_1(x) \\
&= \sum_{i=1}^N \sum_{|\alpha|=1} S_{i,\alpha_1,\alpha_2,\alpha_3,\alpha_4}^n \frac{\partial^{|\alpha_1|+|\alpha_2|+|\alpha_3|} \sigma}{\partial A^{\alpha_1} \partial b^{\alpha_2} \partial c^{\alpha_3}}(x, \bar{A}_i, \bar{b}_i, \bar{c}_i) \frac{\partial^{|\alpha_4|} \mathcal{E}}{\partial \eta^{\alpha_4}}(x, \bar{\eta}_i) + R_1(x), \tag{43}
\end{aligned}$$

where  $R_1(x)$  is a Taylor remainder such that  $R_1(x)/\mathcal{L}_{3n} \rightarrow 0$  as  $n \rightarrow \infty$ . Now we could denote

$$S_{i,\alpha_1:4}^n = \frac{1}{\alpha!} (\Delta \bar{A}_i^n)^{\alpha_1} (\Delta \bar{b}_i^n)^{\alpha_2} (\Delta \bar{c}_i^n)^{\alpha_3} (\Delta \bar{\eta}_i^n)^{\alpha_4}, \quad i \in [N], \alpha \in \mathbb{R}^{d \times d} \times \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^q \tag{44}$$

where  $\sum_{i=1}^4 |\alpha_i| = 1$ .

**Step 2 - Non-vanishing coefficients:** Now we claim that at least one among the ratios  $S_{i,\alpha_1:4}^n / \mathcal{L}_{3n}$  will not vanish as  $n$  goes to infinity. We prove by contradiction that all of them converge to zero when  $n \rightarrow 0$ :

$$\frac{S_{i,\alpha_1:4}^n}{\mathcal{L}_{3n}} \rightarrow 0.$$

for any  $i \in [N], \alpha \in \mathbb{R}^{d \times d} \times \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^q$  such that  $\sum_{i=1}^4 |\alpha_i| = 1$ .

Now, consider  $1 \leq i \leq N$ , for arbitrary  $u, v \in [d]$ , let  $\alpha_1 = e_{d \times d, uv}, \alpha_2 = 0_d, \alpha_3 = 0$  and  $\alpha_4 = 0_q$ , we will have

$$\frac{1}{\mathcal{L}_{3n}} \sum_{i=1}^N \left| \Delta(\bar{A}_i^n)^{(uv)} \right| = \frac{1}{\mathcal{L}_{3n}} \left| S_{i,e_{d \times d,uv},0_d,0,0_q}^n \right| \rightarrow 0, \quad n \rightarrow \infty.$$

Then by taking the summation of the term with  $u, v \in [d]$ , we will have

$$\frac{1}{\mathcal{L}_{3n}} \sum_{i=1}^N \|\Delta \bar{A}_i^n\|_1 = \frac{1}{\mathcal{L}_{3n}} \sum_{i=1}^N \sum_{u=1}^d \sum_{v=1}^d |S_{i,e_{d \times d,uv},0_d,0,0_q}^n| \rightarrow 0.$$

Recall the topological equivalence between  $L_1$ -norm and  $L_2$ -norm on finite-dimensional vector space over  $\mathbb{R}$ , we will have

$$\frac{1}{\mathcal{L}_{3n}} \sum_{i=1}^N \|\Delta \bar{A}_i^n\| \rightarrow 0. \tag{45}$$

Following a similar argument, since

$$\begin{aligned} \frac{1}{\mathcal{L}_{3n}} \sum_{i=1}^N \|\Delta \bar{b}_i^n\|_1 &= \frac{1}{\mathcal{L}_{3n}} \sum_{i=1}^N \sum_{u=1}^d |S_{i,0_{d \times d},e_{d,u},0,0_q}^n| \rightarrow 0, \\ \frac{1}{\mathcal{L}_{3n}} \sum_{i=1}^N |\Delta \bar{c}_i^n| &= \frac{1}{\mathcal{L}_{3n}} \sum_{i=1}^N |S_{i,0_{d \times d},0_d,1,0_q}^n| \rightarrow 0, \\ \frac{1}{\mathcal{L}_{3n}} \sum_{i=1}^N \|\Delta \bar{\eta}_i^n\|_1 &= \frac{1}{\mathcal{L}_{3n}} \sum_{i=1}^N \sum_{w=1}^q |S_{i,0_{d \times d},0_d,0,e_{q,w}}^n| \rightarrow 0, \end{aligned}$$

we obtain that

$$\frac{1}{\mathcal{L}_{3n}} \sum_{i=1}^N \|\Delta \bar{b}_i^n\| \rightarrow 0, \quad \frac{1}{\mathcal{L}_{3n}} \sum_{i=1}^N |\Delta \bar{c}_i^n| \rightarrow 0, \quad \frac{1}{\mathcal{L}_{3n}} \sum_{i=1}^N \|\Delta \bar{\eta}_i^n\| \rightarrow 0. \quad (46)$$

Now taking the summation of limits in equations (45) - (46), we could deduce that

$$1 = \frac{\mathcal{L}_{3n}}{\mathcal{L}_{3n}} = \frac{1}{\mathcal{L}_{3n}} \cdot \sum_{i=1}^N (\|\Delta \bar{A}_i^n\| + \|\Delta \bar{b}_i^n\| + |\Delta \bar{c}_i^n| + \|\Delta \bar{\eta}_i^n\|) \rightarrow 0,$$

as  $n \rightarrow \infty$ , which is a contradiction. Thus, at least one among the ratios  $S_{i,\alpha_{1:4}}^n / \mathcal{L}_{3n}$  must not approach zero as  $n \rightarrow \infty$ . Let us denote by  $m_n$  the maximum of the absolute values of those elements. It follows from the previous result that  $1/m_n \not\rightarrow \infty$  as  $n \rightarrow \infty$ .

**Step 3 - Application of Fatou's lemma:** In this step, we apply Fatou's lemma to obtain the desired inequality in equation (40). Recall in the beginning we have assumed that  $\|f_{G_n} - f_{\bar{G}}\|_{L^2(\mu)} / \mathcal{L}_{3n} \rightarrow 0$  and the topological equivalence between  $L_1$ -norm and  $L_2$ -norm on finite-dimensional vector space over  $\mathbb{R}$ , we obtain  $\|f_{G_n} - f_{\bar{G}}\|_{L^1(\mu)} / \mathcal{L}_{3n} \rightarrow 0$ . By applying the Fatou's lemma, we get

$$0 = \lim_{n \rightarrow \infty} \frac{\|f_{G_n} - f_{\bar{G}}\|_{L^1(\mu)}}{m_n \mathcal{L}_{3n}} \geq \int \liminf_{n \rightarrow \infty} \frac{|f_{G_n}(x) - f_{\bar{G}}(x)|}{m_n \mathcal{L}_{3n}} d\mu(x) \geq 0.$$

This result suggests that for almost every  $x$ ,

$$\frac{f_{G_n}(x) - f_{\bar{G}}(x)}{m_n \mathcal{L}_{3n}} \rightarrow 0. \quad (47)$$

Let us denote

$$\frac{S_{i,\alpha_{1:4}}^n}{m_n \mathcal{L}_{3n}} \rightarrow s_{i,\alpha_{1:4}},$$

as  $n \rightarrow \infty$  with a note that at least one among the limits  $s_{i,\alpha_{1:4}}$  is non-zero. Then from equation (47), we will have

$$\sum_{i=1}^N \sum_{|\alpha|=1} s_{i,\alpha_{1,\alpha_2,\alpha_3,\alpha_4}}^n \frac{\partial^{|\alpha_1|+|\alpha_2|+|\alpha_3|} \sigma}{\partial A^{\alpha_1} \partial b^{\alpha_2} \partial c^{\alpha_3}}(x, \bar{A}_i, \bar{b}_i, \bar{c}_i) \frac{\partial^{|\alpha_4|} \mathcal{E}}{\partial \eta^{\alpha_4}}(x, \bar{\eta}_i) = 0,$$

for almost every  $x$ . Note that the expert function  $\mathcal{E}(\cdot, \eta)$  is weakly identifiable, then the above equation implies that

$$s_{i,\alpha_{1:4}} = 0,$$

for any  $i \in [N]$ ,  $(\alpha_1, \alpha_2, \alpha_3, \alpha_4) \in \mathbb{N}^{d \times d} \times \mathbb{N}^d \times \mathbb{N} \times \mathbb{N}^q$  such that  $\sum_{i=1}^4 |\alpha_i| = 1$ . This violates that at least one among the limits in the set  $\{s_{i,\alpha_{1:4}}\}$  is different from zero.

Thus, we obtain the local inequality in equation (40).  $\square$

## B.5 Convergence to a Single Sigmoid under Vanishing Gating

In this subsection, we formally justify the statement that the limit of a sum of two sigmoid functions can reduce to a single sigmoid function only when one of the experts effectively vanishes. This result is used in Section 3 to explain why the regression function learned by the over-specified model converges to a single-expert form. The key condition for this collapse is that the gating parameters of the redundant expert must vanish asymptotically. We state and prove this result below.

**Proposition 3.** *Let  $(\hat{A}_i^n, \hat{b}_i^n, \hat{c}_i^n) \rightarrow (A_i^*, b_i^*, c_i^*)$  in probability for  $i = 1, 2$ . Suppose that for  $\mu$ -almost every  $x \in \mathbb{R}^d$ ,*

$$\sum_{i=1}^2 \frac{1}{1 + \exp(-x^\top \hat{A}_i^n x - (\hat{b}_i^n)^\top x - \hat{c}_i^n)} \xrightarrow{P} \frac{1}{1 + \exp(-x^\top A_1^* x - (b_1^*)^\top x - c_1^*)}.$$

*Then it must hold that  $A_2^* = 0_{d \times d}$ ,  $b_2^* = 0_d$ , and  $c_2^* = +\infty$  (or tends to  $+\infty$  in probability), i.e., the second sigmoid term vanishes in the limit.*

*Proof.* Let us denote the sigmoid function as  $\sigma(z) = 1/(1 + \exp(-z))$ , and consider the limiting expressions:

$$\begin{aligned} \sigma_1(x) &= \frac{1}{1 + \exp(-x^\top A_1^* x - (b_1^*)^\top x - c_1^*)}, \\ \sigma_2(x) &= \frac{1}{1 + \exp(-x^\top A_2^* x - (b_2^*)^\top x - c_2^*)}. \end{aligned}$$

Suppose toward a contradiction that  $(A_2^*, b_2^*) \neq (0, 0)$  or  $c_2^*$  is finite. Then there exists a set of positive  $\mu$ -measure where  $\sigma_2(x)$  is bounded away from 0. In that case, the sum  $\sigma_1(x) + \sigma_2(x)$  must strictly exceed  $\sigma_1(x)$  for such  $x$ , and hence cannot equal  $\sigma_1(x)$ .

But the convergence in the hypothesis implies that for  $\mu$ -almost every  $x$ ,

$$\sigma_1^n(x) + \sigma_2^n(x) \xrightarrow{P} \sigma_1(x),$$

which can only hold if  $\sigma_2^n(x) \rightarrow 0$  in probability. This, in turn, requires that the argument of the exponential in  $\sigma_2^n(x)$  tends to  $+\infty$  in probability for almost every  $x$ , i.e.,

$$x^\top A_2^* x + (b_2^*)^\top x + c_2^* \rightarrow +\infty.$$

This is only possible if  $A_2^* = 0$ ,  $b_2^* = 0$ , and  $c_2^* \rightarrow +\infty$ . Hence, the second sigmoid term vanishes in the limit.  $\square$

## C Additional Results

### C.1 Sample Complexity under the Partially Quadratic Score Function

In this section, we proceed the analysis of the sigmoid gating MoE with partially quadratic affinity score function based on the regression frame work in equation (5), the corresponding regression function is redefined as follows:

$$\tilde{f}_{G_*}(x) := \sum_{i=1}^{N^*} \frac{1}{1 + \exp(-x^\top A_i^* x - c_i^*)} \cdot \mathcal{E}(x, \eta_i^*) \quad (48)$$

where  $\{(A_i, c_i, \eta_i) \in \tilde{\Theta}, i = 1, \dots, N\}$  is the set of learnable parameters,  $\tilde{\Theta} \subseteq \mathbb{R}^{d \times d} \times \mathbb{R} \times \mathbb{R}^q$  stands for the parameter space and  $\tilde{\mathcal{M}}_N(\tilde{\Theta}) := \{\tilde{G} = \sum_{i=1}^{N'} \frac{1}{1 + \exp(-c_i)} \delta_{(A_i, \eta_i)} : 1 \leq N' \leq N, (A_i, c_i, \eta_i) \in \tilde{\Theta}\}$  is the set of all mixing measures with at most  $N$  atoms. We assume  $N > N^*$ . And for simplicity, we still denote  $\tilde{\Theta}$  as  $\Theta$ ,  $\tilde{\mathcal{M}}_N(\tilde{\Theta})$  as  $\mathcal{M}_N(\Theta)$  and  $\tilde{G}$  as  $G$  in the following.

In contrast to the fully quadratic score function, the first-degree monomial term  $b^\top x$  has been omitted from the scoring function. Consequently, the least squares estimator in this context is modified as follows:

$$\tilde{G}_n := \arg \min_{G \in \mathcal{M}_N(\Theta)} \sum_{i=1}^n \left( Y_i - \tilde{f}_G(X_i) \right)^2. \quad (49)$$

#### C.1.1 Convergence of the Regression Function Estimator

Similar with the polynomial case, we fit the ground-truth MoE model with a mixture of  $N > N^*$ , there must be some true atoms  $(A_i^*, \eta_i^*)$  fitted by more than one component. We over-specify the true MoE model by a mixture of  $N$  experts where  $N > N^*$ . There exist some atoms  $(A_i^*, \eta_i^*)$  approximated by at least two fitted components, the over-specified atoms. We assume that  $(\hat{A}_i^n, \hat{\eta}_i^n) \rightarrow (A_1^*, \eta_1^*)$  for  $i = 1, 2$ , in probability. Then, the term  $\|\tilde{f}_{\tilde{G}_n} - \tilde{f}_{G_*}\|_{L^2(\mu)} \rightarrow 0$  only when

$$\sum_{i=1}^2 \frac{1}{1 + \exp(-x^\top \hat{A}_i^n x - \hat{c}_i^n)} \rightarrow \frac{1}{1 + \exp(-x^\top A_1^* x - c_1^*)},$$

as  $n \rightarrow \infty$  for  $\mu$ -almost every  $x$ , which occurs only when  $A_1^* = 0_d$ . Following the approach outlined in Section 3, we will partition our analysis into two complementary regimes of the gating parameters:

- (i) *Sparse Regime*: all the over-specified gating parameters are zero:  $A_i^* = 0_{d \times d}$ ;
- (ii) *Dense Regime*: at least one among the over-specified gating parameters is non-zero:  $A_i^* \neq 0_{d \times d}$ .

We now present the convergence behavior of the regression function estimator under each of the two regimes, respectively.

**Proposition 4.** *Under the sparse regime and with the least squares estimator  $\tilde{G}_n$  defined in equation (49), the regression estimator  $\tilde{f}_{\tilde{G}_n}$  admits the following rate of convergence to  $\tilde{f}_{G_*}$ :*

$$\|\tilde{f}_{\tilde{G}_n} - \tilde{f}_{G_*}\|_{L^2(\mu)} = \mathcal{O}_P \left( \sqrt{\log(n)/n} \right). \quad (50)$$

The proof of Proposition 4 can be found in Appendix C.2.1. The bound (50) reveals that the convergence rate of the regression function estimator  $\tilde{f}_{\tilde{G}_n}$  to the ground-truth regression function is of parametric order  $\mathcal{O}_P(\sqrt{\log(n)/n})$  under the sparse regime. However, under the dense regime, such convergence does not happen as previously mentioned. Instead, the regression function estimator  $\tilde{f}_{\tilde{G}_n}$  converges to the closest MoE with more than  $N^*$  experts to  $\tilde{f}_{G_*}$ , that is,  $\tilde{f}_{\bar{G}}$  where  $\bar{G} \in \overline{\mathcal{M}}_N(\Theta) := \arg \min_{G \in \mathcal{M}_N(\Theta) \setminus \mathcal{M}_{N^*}(\Theta)} \|\tilde{f}_G - \tilde{f}_{G_*}\|_{L^2(\mu)}$ . Using similar arguments to Proposition 4, we can also determine the convergence behavior of the regression function estimator in the following corollary:

**Corollary 2.** *Under the dense regime and with the least squares estimator  $\tilde{G}_n$  defined in equation (49), the regression estimator  $\tilde{f}_{\tilde{G}_n}$  admits the following rate of convergence to  $\tilde{f}_{\bar{G}}$ :*

$$\inf_{\bar{G} \in \overline{\mathcal{M}}_N(\Theta)} \|\tilde{f}_{\tilde{G}_n} - \tilde{f}_{\bar{G}}\|_{L^2(\mu)} = \mathcal{O}_P\left(\sqrt{\log(n)/n}\right).$$

Similar with what we have discussed in Section 3, given the results of Propositions 4 and 2 and a loss function  $\mathcal{L}(G, G_*)$  among parameters satisfying  $\mathcal{L}(\tilde{G}_n, G_*) \lesssim \|\tilde{f}_{\tilde{G}_n} - \tilde{f}_{G_*}\|_{L^2(\mu)}$ , we will have that  $\mathcal{L}(\tilde{G}_n, G_*) = \mathcal{O}_P(\sqrt{\log(n)/n})$ . From this, we can derive the expert convergence rate, which will be analyzed in the following.

### C.1.2 Sparse Regime of Gating Parameters

Recall that under the sparse regime, all the over-specified parameters  $A_i^* = 0_{d \times d}$ . Assume  $\{A_i^*\}_{i=1}^{\bar{N}}$  are over-specified parameters, i.e., those fitted by at least two estimators, where  $1 \leq \bar{N} \leq N^*$ . And the remaining parameters fitted by exactly one estimator  $\{A_i^*\}_{\bar{N}+1}^{N^*}$ :

$$\underbrace{A_1^*, \dots, A_{\bar{N}}^*}_{\text{over-specified}}, \underbrace{A_{\bar{N}+1}^*, \dots, A_{N^*}^*}_{\text{exactly-specified}}.$$

As discussed in Section 4, to derive the expert convergence rate, we decompose the difference  $\tilde{f}_{\tilde{G}_n} - \tilde{f}_{G_*}$  into a sum of linearly independent terms. This requires considering the Taylor expansion of the product of the sigmoid gating function and the expert function, given by

$$F(x; A, c, \eta) := \sigma(x^\top A x + c) \cdot \mathcal{E}(x, \eta) = \frac{1}{1 + \exp(-x^\top A x - c)} \cdot \mathcal{E}(x, \eta).$$

To ensure that the parameter discrepancies converge to zero—thereby guaranteeing parameter and expert convergence—along with  $\|\tilde{f}_{\tilde{G}_n} - \tilde{f}_{G_*}\|_{L^2(\mu)} \rightarrow 0$ , it is crucial to clarify the linear independence properties of  $F(x; A, c, \eta)$ , a fundamental expert component with a partially quadratic score in the mixture of experts model. Therefore, we impose a *partial-strong identifiability* condition, defined as follows:

**Definition 3** (Partial-strong identifiability). *An expert function  $\mathcal{E}(x, \eta)$  is partial-strongly identifiable if it is twice differentiable w.r.t its parameter  $\eta$  for  $\mu$ -almost all  $x$  and, for any positive integer  $\ell$  and any pair-wise distinct choices of parameters  $\{(A_i, c_i, \eta_i)\}_{i=1}^\ell$ , the functions in the classes*

$$\left\{ \frac{\partial^{|\gamma_1|+|\gamma_2|} F}{\partial A^{\gamma_1} \partial \eta^{\gamma_2}}(x, 0_{d \times d}, c_i, \eta_i) : i \in [\ell], \sum_{i=1}^2 |\gamma_i| \in [2] \right\} \text{ and } \left\{ \frac{\partial F}{\partial A^{\tau_1} \partial c^{\tau_2} \partial \eta^{\tau_3}}(x, A_i, c_i, \eta_i) : i \in [\ell], \sum_{i=1}^3 |\tau_i| = 1 \right\}$$

are linearly independent, for  $\mu$ -almost all  $x$ , where  $(\gamma_1, \gamma_2) \in \mathbb{N}^{d \times d} \times \mathbb{N}^q$  and  $(\tau_1, \tau_2, \tau_3) \in \mathbb{N}^{d \times d} \times \mathbb{N} \times \mathbb{N}^q$ .

**Examples.** Similar to Section 4.1, we consider two-layer neural networks of the form  $\mathcal{E}(x; (\alpha, \beta, \lambda)) = \lambda\phi(\alpha^\top x + \beta)$ , where  $\phi$  is an activation function and  $(\alpha, \beta, \lambda) \in \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}$ . It can be verified that if  $\phi$  is the ReLU or GELU function, and if  $\alpha \neq 0_d$  and  $\lambda \neq 0$ , then the function  $x \mapsto \mathcal{E}(x; (\alpha, \beta, \lambda))$  is partial-strongly identifiable.

However, unlike the fully quadratic score case, where the term  $b^\top x$  is present, its removal allows the partial-strong identifiability condition in Definition 3 to remain valid even when the expert function follows a polynomial form, given by  $\mathcal{E}(x; (\alpha, \beta)) = (\alpha^\top x + \beta)^p$  for some  $p \in \mathbb{N}$ . Now, the function takes the form  $F(x; A, c, \alpha, \beta) = 1/[1 + \exp(-x^\top Ax - c)] \cdot (\alpha^\top x + \beta)^p$ . Notably, the sigmoid function contributes only a second-order term  $-x^\top Ax$  along with a constant, while the expert function contains only a first-order term  $\alpha^\top x$  and a constant. This distinct structure ensures that parameter interactions disappear as PDEs in equation (8).

**Voronoi loss.** For a mixing measure  $G$  with  $1 \leq N' \leq N$  atoms, we allocate its atoms across the Voronoi cells  $\{\mathcal{A}_j \equiv \mathcal{A}_j(G), j \in [N^*]\}$  generated by the atoms of  $G_*$ , where

$$\mathcal{A}_j := \{i \in [N'] : \|\theta_i - \theta_j^*\| \leq \|\theta_i - \theta_\ell^*\|, \forall \ell \neq j\}, \quad (51)$$

with  $\theta_i := (A_i, \eta_i)$  and  $\theta_j^* := (A_j^*, \eta_j^*)$  for all  $j \in [N^*]$ . Similar to the fully quadratic score case, we define the Voronoi loss function as

$$\begin{aligned} \mathcal{L}_4(G, G_*) := & \sum_{j=1}^{\bar{N}} \left| \sum_{i \in \mathcal{A}_j} \frac{1}{1 + \exp(-c_i)} - \frac{1}{1 + \exp(-c_j^*)} \right| + \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} [\|\Delta A_{ij}\|^2 + \|\Delta \eta_{ij}\|^2] \\ & + \sum_{j=\bar{N}+1}^{N^*} \sum_{i \in \mathcal{A}_j} [\|\Delta A_{ij}\| + |\Delta c_{ij}| + \|\Delta \eta_{ij}\|] \end{aligned} \quad (52)$$

where we denote  $\Delta A_{ij} := A_i - A_j^*$ ,  $\Delta c_{ij} := c_i - c_j^*$  and  $\Delta \eta_{ij} := \eta_i - \eta_j^*$ . In the aforementioned statement, if the Voronoi cell  $\mathcal{A}_j$  is empty, the associated summation term is conventionally set to zero. Additionally, the computation of the Voronoi loss function  $\mathcal{L}_4$  is efficient, with a computational complexity of  $\mathcal{O}(N \times N^*)$ .

With the Voronoi loss function established, we are now prepared to present the parameter convergence rate in Theorem 4.

**Theorem 4.** *If the expert function  $x \mapsto \mathcal{E}(x, \eta)$  is partial-strongly identifiable, then the lower bound  $\|\tilde{f}_G - \tilde{f}_{G_*}\|_{L^2(\mu)} \gtrsim \mathcal{L}_4(G, G_*)$  holds true for any  $G \in \mathcal{M}_N(\Theta)$ . As a result, when combined with the regression estimation rate in Proposition 4, this suggests that*

$$\mathcal{L}_4(\tilde{G}_n, G_*) = \mathcal{O}_P\left(\sqrt{\log(n)/n}\right).$$

Proof of Theorem 4 is in Appendix C.2.2. A few comments regarding the above result are in order.

(i) *Parameter convergence rates:* From the construction of the Voronoi loss  $\mathcal{L}_4$ , the estimation of over-specified parameters  $A_i^*, \eta_i^*$  for  $i \in [\bar{N}]$  converges at a rate of  $\mathcal{O}_P([\log(n)/n]^{1/4})$ . In the mean time, the exactly-specified parameters  $A_i^*, \eta_i^*$  for  $\bar{N} + 1 \leq i \leq N^*$  exhibit a faster convergence rate of  $\mathcal{O}_P([\log(n)/n]^{1/2})$ .

(ii) *Expert convergence rates:* Because  $\mathcal{E}(x, \eta)$  is twice differentiable (and thus Lipschitz) in  $\eta$  over a bounded domain, we have, for any  $i \in \mathcal{A}_j$ ,  $\sup_x |\mathcal{E}(x, \tilde{\eta}_i^n) - \mathcal{E}(x, \eta_j^*)| \leq \tilde{L} \|\tilde{\eta}_i^n - \eta_j^*\|$ , where  $\tilde{L}$  is a

Lipschitz constant and  $\tilde{G}_n = \sum_{i=1}^{\tilde{N}_n} \frac{1}{1+\exp(-\tilde{c}_i^n)} \delta_{(\tilde{A}_i^n, \tilde{b}_i^n, \tilde{\eta}_i^n)}$ . This implies that exactly-specified experts converge at a rate of  $\mathcal{O}_P((\log(n)/n)^{1/2})$ , while over-specified experts converge at  $\mathcal{O}_P((\log(n)/n)^{1/4})$ . Consequently, exactly-specified experts require  $\mathcal{O}(\epsilon^{-2})$  samples to achieve an approximation error  $\epsilon$ , whereas over-specified experts need  $\mathcal{O}(\epsilon^{-4})$  samples for the same error.

**Sample complexity comparison under the sparse regime:** Recall from Section 4.1 that polynomial experts, although not strongly identifiable, are partial-strongly identifiable. Hence, while they might require  $\mathcal{O}(\exp(\epsilon^{-1/\tau}))$  samples to achieve an approximation error  $\epsilon$  under a fully quadratic score function (for some constant  $\tau > 0$ ), they only need  $\mathcal{O}(\epsilon^{-4})$  samples under a partially quadratic score function.

According to [1], linear experts with quadratic monomial gates require an exponential number of samples. But the linear expert  $\mathcal{E}(x, (\alpha, \beta)) = \alpha^\top x + \beta$  satisfies partial-strong identifiability when  $\beta \neq 0$ , so it need only a polynomial number  $\mathcal{O}(\epsilon^{-4})$  data points to attain the approximation error  $\epsilon$ .

Consequently, for partially quadratic score functions with linear experts, the sigmoid self-attention is more sample-efficient than the softmax self-attention under the sparse gating regime.

### C.1.3 Dense Regime of Gating Parameters

In this section, we focus on the dense regime. Under dense regime, for the overspecified parameters, there exists  $i \in [N^*]$ , s.t.  $A_i^* \neq 0_{d \times d}$ . In this situation, the least squares regression estimator  $f_{\tilde{G}_n} \rightarrow f_{\bar{G}}$ , a regression function, where the parameters

$$\bar{G} \in \overline{\mathcal{M}}_N(\Theta) := \arg \min_{G \in \mathcal{M}_N(\Theta) \setminus \mathcal{M}_{N^*}(\Theta)} \|\tilde{f}_G - \tilde{f}_{G^*}\|_{L^2(\mu)}.$$

In other words, the estimators of the parameters defining  $\tilde{f}_{\tilde{G}_n}$  converge to the parameters of  $\tilde{f}_{\bar{G}}$ . WLOG, we assume that

$$\bar{G} := \sum_{i=1}^N \frac{1}{1 + \exp(-\bar{c}_i)} \delta_{(\bar{A}_i, \bar{\eta}_i)}.$$

Now we introduce a definition for *partial-weak identifiability* to specify the experts that has a similar fast convergence rate as the partial-strongly identifiable experts in Definition 3 under sparse regime. Partial-weakly identifiable experts need to satisfy only a subset of the conditions required for partial-strongly identifiable experts.

**Definition 4** (Partial-weak identifiability). *An expert function  $\mathcal{E}(x, \eta)$  is partial-weakly identifiable if it is differentiable w.r.t its parameter  $\eta$  for  $\mu$ -almost all  $x$  and, for any positive integer  $\ell$  and any pair-wise distinct choices of parameters  $\{(A_i, c_i, \eta_i)\}_{i=1}^\ell$ , the functions in the class*

$$\left\{ \frac{\partial F}{\partial A^{\tau_1} \partial c^{\tau_2} \partial \eta^{\tau_3}}(x, A_i, c_i, \eta_i) : i \in [\ell], \sum_{i=1}^3 |\tau_i| = 1 \right\}$$

are linearly independent, for  $\mu$ -almost all  $x$ , where  $(\tau_1, \tau_2, \tau_3) \in \mathbb{N}^{d \times d} \times \mathbb{N} \times \mathbb{N}^q$ .

**Examples.** It is worth noting that partial-strongly identifiable experts also meet the partial-weak identifiability condition. In particular, the previously mentioned two-layer neural networks  $\mathcal{E}(x; (\alpha, \beta, \lambda)) = \lambda \phi(\alpha^\top x + \beta)$ , where  $\phi$  is either ReLU or GELU,  $\alpha \neq 0_d$ , and  $\lambda \neq 0$ , are weakly

identifiable. Likewise, polynomial experts of the form  $\mathcal{E}(x; (\alpha, \beta)) = (\alpha^\top x + \beta)^p, p \in \mathbb{N}$ , also satisfy the weak identifiability condition.

We now establish, in Theorem 5, the convergence behavior of partial-weakly identifiable experts in the dense regime. Specifically, let  $\mathcal{L}_5$  denote the Voronoi loss function defined by

$$\mathcal{L}_5(G, \bar{\bar{G}}) := \sum_{j=1}^N \sum_{i \in \mathcal{A}_j} [\|A_i - \bar{\bar{A}}_i\| + |c_i - \bar{\bar{c}}_i| + \|\eta_i - \bar{\bar{\eta}}_i\|]. \quad (53)$$

**Theorem 5.** *If the function  $x \mapsto \mathcal{E}(x, \eta)$  is partial-weakly identifiable, then the lower bound  $\inf_{\bar{\bar{G}} \in \bar{\bar{\mathcal{M}}}_N(\Theta)} \|\tilde{f}_G - \tilde{f}_{\bar{\bar{G}}}\|_{L^2(\mu)} \gtrsim \mathcal{L}_5(G, \bar{\bar{G}})$  holds true for any mixing measure  $G \in \mathcal{M}_N(\Theta)$ . As a consequence, we obtain that*

$$\inf_{\bar{\bar{G}} \in \bar{\bar{\mathcal{M}}}_N(\Theta)} \mathcal{L}_5(\tilde{G}_n, \bar{\bar{G}}) = \mathcal{O}_P\left(\sqrt{\log(n)/n}\right).$$

Proof of Theorem 5 is in Appendix C.2.3. Since  $\tilde{\eta}_i^n$  converges at a rate of  $\mathcal{O}_P(\sqrt{\log(n)/n})$ , the partial-weakly identifiable expert estimators  $\mathcal{E}(x, \tilde{\eta}_i^n)$  enjoy the same convergence order. Consequently, these experts require a polynomial number of samples,  $\mathcal{O}(\epsilon^{-2})$ , to achieve an approximation error of  $\epsilon$ .

**Sample complexity comparison under the dense regime:** Recall that, under softmax self-attention, partial-strongly identifiable and polynomial “experts” require  $\mathcal{O}(\epsilon^{-4})$  and  $\mathcal{O}(\exp(\epsilon^{-1/\tau}))$  samples, respectively, to achieve an approximation error  $\epsilon$ . However, because these experts all satisfy partial-weak identifiability under sigmoid self-attention, they need only  $\mathcal{O}(\epsilon^{-2})$  samples in the dense regime for partially quadratic scores. We therefore conclude that, in practice, where the dense regime is more common, sigmoid self-attention is more sample-efficient than its softmax counterpart.

## C.2 Proofs of the results for Partially Quadratic Scores in Appendix C.1

### C.2.1 Proof of Proposition 4

*Proof.* The proof of Proposition 4 can be done in a similar fashion to that of Proposition 1 in Appendix B.1.  $\square$

### C.2.2 Proof of Theorem 4

*Proof.* In order to prove  $\|\tilde{f}_G - \tilde{f}_{G_*}\|_{L^2(\mu)} \gtrsim \mathcal{L}_4(G, G_*)$  for any  $G \in \mathcal{M}_N(\Theta)$ , it is sufficient to show that

$$\inf_{G \in \mathcal{M}_N(\Theta)} \frac{\|\tilde{f}_G - \tilde{f}_{G_*}\|_{L^2(\mu)}}{\mathcal{L}_4(G, G_*)} > 0. \quad (54)$$

To prove the above inequality, we consider two cases for the denominator  $\mathcal{L}_4(G, G_*)$ : either it lies within a ball  $B(0, \epsilon)$  where the loss is sufficiently small, or it falls outside this region, where  $\mathcal{L}_4(G, G_*)$  will not vanish.

**Local part:** At first, we focus on that

$$\lim_{\epsilon \rightarrow 0} \inf_{G \in \mathcal{M}_N(\Theta): \mathcal{L}_4(G, G_*) \leq \epsilon} \frac{\|\tilde{f}_G - \tilde{f}_{G_*}\|_{L^2(\mu)}}{\mathcal{L}_4(G, G_*)} > 0. \quad (55)$$

Assume, for contradiction, that the above claim does not hold. Then there exists a sequence of mixing measures  $\tilde{G}_n = \sum_{i=1}^{N^*} \delta_{(A_i^n, c_i^n, \eta_i^n)}$  in  $\mathcal{M}_N(\Theta)$  such that as  $n \rightarrow \infty$ , we get

$$\begin{cases} \mathcal{L}_{4n} := \mathcal{L}_4(\tilde{G}_n, G_*) \rightarrow 0, \\ \|\tilde{f}_{\tilde{G}_n} - \tilde{f}_{G_*}\|_{L^2(\mu)} / \mathcal{L}_{4n} \rightarrow 0. \end{cases} \quad (56)$$

Let us recall that

$$\begin{aligned} \mathcal{L}_{4n} &:= \sum_{j=1}^{\bar{N}} \left| \sum_{i \in \mathcal{A}_j} \frac{1}{1 + \exp(-c_i^n)} - \frac{1}{1 + \exp(-c_j^*)} \right| + \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} [\|\Delta A_{ij}^n\|^2 + \|\Delta \eta_{ij}^n\|^2] \\ &+ \sum_{j=\bar{N}+1}^{N^*} \sum_{i \in \mathcal{A}_j} [\|\Delta A_{ij}^n\| + |\Delta c_{ij}^n| + \|\Delta \eta_{ij}^n\|], \end{aligned}$$

where  $\Delta A_{ij}^n := A_i^n - A_j^*$ ,  $\Delta c_{ij}^n := c_i^n - c_j^*$ ,  $\Delta \eta_{ij}^n := \eta_i^n - \eta_j^*$ . Since  $\mathcal{L}_{4n} \rightarrow 0$  as  $n \rightarrow 0$ , there are two different situation for parameter convergence:

- For  $j = 1, \dots, \bar{N}$ , parameters are over-fitted:  $\sum_{i \in \mathcal{A}_j} \frac{1}{1 + \exp(-c_i^n)} \rightarrow \frac{1}{1 + \exp(-c_j^*)}$  and  $(A_i^n, \eta_i^n) \rightarrow (A_j^*, \eta_j^*)$ ,  $\forall i \in \mathcal{A}_j$ ;
- For  $j = \bar{N} + 1, \dots, N^*$ , parameters are exact-fitted:  $(A_i^n, c_i^n, \eta_i^n) \rightarrow (A_j^*, c_j^*, \eta_j^*)$ ,  $\forall i \in \mathcal{A}_j$ .

**Step 1 - Taylor expansion:** In this step, we decompose the term  $\tilde{f}_{\tilde{G}_n}(x) - \tilde{f}_{G_*}(x)$  using a Taylor expansion. First, let us denote

$$\begin{aligned} \tilde{f}_{\tilde{G}_n}(x) - \tilde{f}_{G_*}(x) &= \sum_{j=1}^{\bar{N}} \left[ \sum_{i \in \mathcal{A}_j} \frac{1}{1 + \exp(-x^\top A_i^n x - c_i^n)} \cdot \mathcal{E}(x, \eta_i^n) - \frac{1}{1 + \exp(-x^\top A_j^* x - c_j^*)} \cdot \mathcal{E}(x, \eta_j^*) \right] \\ &+ \sum_{j=\bar{N}+1}^{N^*} \left[ \sum_{i \in \mathcal{A}_j} \frac{1}{1 + \exp(-x^\top A_i^n x - c_i^n)} \cdot \mathcal{E}(x, \eta_i^n) - \frac{1}{1 + \exp(-x^\top A_j^* x - c_j^*)} \cdot \mathcal{E}(x, \eta_j^*) \right] \quad (57) \\ &= \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} \left[ \frac{1}{1 + \exp(-x^\top A_i^n x - c_i^n)} \cdot \mathcal{E}(x, \eta_i^n) - \frac{1}{1 + \exp(-c_i^n)} \cdot \mathcal{E}(x, \eta_j^*) \right] := \text{I}_n \\ &+ \sum_{j=1}^{\bar{N}} \left[ \sum_{i \in \mathcal{A}_j} \frac{1}{1 + \exp(-c_i^n)} - \frac{1}{1 + \exp(-c_j^*)} \right] \cdot \mathcal{E}(x, \eta_j^*) := \text{II}_n \\ &+ \sum_{j=\bar{N}+1}^{N^*} \left[ \sum_{i \in \mathcal{A}_j} \frac{1}{1 + \exp(-x^\top A_i^n x - c_i^n)} \cdot \mathcal{E}(x, \eta_i^n) - \frac{1}{1 + \exp(-x^\top A_j^* x - c_j^*)} \cdot \mathcal{E}(x, \eta_j^*) \right] := \text{III}_n \end{aligned}$$

Let us denote  $\sigma(x, A, c) = \frac{1}{1 + \exp(-x^\top A x - c)}$ , then  $I_n, II_n$  and  $III_n$  could be denoted as

$$\begin{aligned} I_n &= \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} \left[ \frac{1}{1 + \exp(-x^\top A_i^n x - c_i^n)} \cdot \mathcal{E}(x, \eta_i^n) - \frac{1}{1 + \exp(-c_i^n)} \cdot \mathcal{E}(x, \eta_j^*) \right] \\ &= \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} [\sigma(x, A_i^n, c_i^n) \mathcal{E}(x, \eta_i^n) - \sigma(x, 0_{d \times d}, c_i^n) \mathcal{E}(x, \eta_j^*)] \\ &= \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} \sum_{|\gamma|=1}^2 \frac{1}{\gamma!} (\Delta A_{ij}^n)^{\gamma_1} (\Delta \eta_{ij}^n)^{\gamma_2} \frac{\partial^{|\gamma_1|} \sigma}{\partial A^{\gamma_1}}(x, 0_{d \times d}, c_i^n) \frac{\partial^{|\gamma_2|} \mathcal{E}}{\partial \eta^{\gamma_2}}(x, \eta_j^*) + R_1(x), \end{aligned} \quad (58)$$

$$II_n = \sum_{j=1}^{\bar{N}} \left[ \sum_{i \in \mathcal{A}_j} \frac{1}{1 + \exp(-c_i^n)} - \frac{1}{1 + \exp(-c_j^*)} \right] \cdot \mathcal{E}(x, \eta_j^*), \quad (59)$$

$$\begin{aligned} III_n &= \sum_{j=\bar{N}+1}^{N^*} \left[ \sum_{i \in \mathcal{A}_j} \frac{1}{1 + \exp(-x^\top A_i^n x - c_i^n)} \cdot \mathcal{E}(x, \eta_i^n) - \frac{1}{1 + \exp(-x^\top A_j^* x - c_j^*)} \cdot \mathcal{E}(x, \eta_j^*) \right] \\ &= \sum_{j=\bar{N}+1}^{N^*} \left[ \sum_{i \in \mathcal{A}_j} [\sigma(x, A_i^n, c_i^n) \mathcal{E}(x, \eta_i^n) - \sigma(x, A_j^*, c_j^*) \mathcal{E}(x, \eta_j^*)] \right] \\ &= \sum_{j=\bar{N}+1}^{N^*} \sum_{|\tau|=1} \left[ \sum_{i \in \mathcal{A}_j} \frac{1}{\tau!} (\Delta A_{ij}^n)^{\tau_1} (\Delta c_{ij}^n)^{\tau_2} (\Delta \eta_{ij}^n)^{\tau_3} \right] \frac{\partial^{|\tau_1|+|\tau_2|} \sigma}{\partial A^{\tau_1} \partial c^{\tau_2}}(x, A_j^*, c_j^*) \frac{\partial^{|\tau_3|} \mathcal{E}}{\partial \eta^{\tau_3}}(x, \eta_j^*) + R_2(x), \end{aligned} \quad (60)$$

where  $R_i(x), i = 1, 2$  are Taylor remainder such that  $R_i(x)/\mathcal{L}_{4n} \rightarrow 0$  as  $n \rightarrow \infty$  for  $i = [2]$ .

Now we could denote

$$K_{j,i,\gamma_{1:2}}^n = \frac{1}{\gamma!} (\Delta A_{ij}^n)^{\gamma_1} (\Delta \eta_{ij}^n)^{\gamma_2}, \quad j \in [\bar{N}], i \in \mathcal{A}_j, \gamma \in \mathbb{R}^{d \times d} \times \mathbb{R}^q \quad (61)$$

$$L_j^n = \sum_{i \in \mathcal{A}_j} \frac{1}{1 + \exp(-c_i^n)} - \frac{1}{1 + \exp(-c_j^*)}, \quad j \in [\bar{N}] \quad (62)$$

$$T_{j,\tau_{1:3}}^n = \sum_{i \in \mathcal{A}_j} \frac{1}{\tau!} (\Delta A_{ij}^n)^{\tau_1} (\Delta c_{ij}^n)^{\tau_2} (\Delta \eta_{ij}^n)^{\tau_3}, \quad j \in \{\bar{N}+1, \dots, N^*\}, \tau \in \mathbb{R}^{d \times d} \times \mathbb{R} \times \mathbb{R}^q \quad (63)$$

where  $1 \leq \sum_{i=1}^2 |\gamma_i| \leq 2$  and  $\sum_{i=1}^3 |\tau_i| = 1$ . Also recall that for the last term,  $i$  was settled directly by  $j$  since all the parameters are exact-fitted when  $\bar{N}+1 \leq j \leq N^*$ , i.e.  $|\mathcal{A}_j| = 1$ .

Using these notations (61) - (63) we can now rewrite the difference  $\tilde{f}_{\tilde{G}_n}(x) - \tilde{f}_{G_*}(x)$  as

$$\begin{aligned} \tilde{f}_{\tilde{G}_n}(x) - \tilde{f}_{G_*}(x) &= \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} \sum_{|\gamma|=1}^2 K_{j,i,\gamma_{1:2}}^n \frac{\partial^{|\gamma_1|} \sigma}{\partial A^{\gamma_1}}(x, 0_{d \times d}, c_i^n) \frac{\partial^{|\gamma_2|} \mathcal{E}}{\partial \eta^{\gamma_2}}(x, \eta_j^*) + \sum_{j=1}^{\bar{N}} L_j^n \cdot \mathcal{E}(x, \eta_j^*) \\ &\quad + \sum_{j=\bar{N}+1}^{N^*} \sum_{|\tau|=1} T_{j,\tau_{1:3}}^n \frac{\partial^{|\tau_1|+|\tau_2|} \sigma}{\partial A^{\tau_1} \partial c^{\tau_2}}(x, A_j^*, c_j^*) \frac{\partial^{|\tau_3|} \mathcal{E}}{\partial \eta^{\tau_3}}(x, \eta_j^*) + R_1(x) + R_2(x). \end{aligned} \quad (64)$$

**Step 2 - Non-vanishing coefficients:** Now we claim that at least one in the set

$$\mathcal{S} = \left\{ \frac{K_{j,i,\gamma_{1:2}}^n}{\mathcal{L}_{4n}}, \frac{L_j^n}{\mathcal{L}_{4n}}, \frac{T_{j,\tau_{1:3}}^n}{\mathcal{L}_{4n}} \right\}$$

will not vanish as  $n$  goes to infinity. We prove by contradiction that all of them converge to zero when  $n \rightarrow 0$ :

$$\frac{K_{j,i,\gamma_{1:2}}^n}{\mathcal{L}_{4n}} \rightarrow 0, \quad \frac{L_j^n}{\mathcal{L}_{4n}} \rightarrow 0, \quad \frac{T_{j,\tau_{1:3}}^n}{\mathcal{L}_{4n}} \rightarrow 0.$$

Then follows directly from  $L_j^n/\mathcal{L}_{4n} \rightarrow 0$  we could conclude that

$$\frac{1}{\mathcal{L}_{4n}} \sum_{j=1}^{\bar{N}} \left| \sum_{i \in \mathcal{A}_j} \frac{1}{1 + \exp(-c_i^n)} - \frac{1}{1 + \exp(-c_j^*)} \right| = \frac{1}{\mathcal{L}_{4n}} \sum_{j=1}^{\bar{N}} |L_j^n| \rightarrow 0. \quad (65)$$

Before consider other coefficients, for simplicity, we denote  $e_{d,u} = \overbrace{(0, \dots, 0, \underbrace{1}_{u\text{-th}}, 0, \dots, 0)}^{d\text{-tuple}} \in \mathbb{R}^d$  as a  $d$ -tuple with all components equal to 0, except the  $u$ -th, which is 1; and  $e_{d \times d, uv}$  as a  $d \times d$  matrix with all components equal to 0, except the element in the  $u$ -th row and  $v$ -th column, which is 1, i.e.  $e_{d \times d, uv} = \overbrace{(0_d^\top, \dots, 0_d^\top, e_{d,u}^\top, 0_d^\top, \dots, 0_d^\top)}^{d\text{-column}} \in \mathbb{R}^{d \times d}$ .

Now consider for arbitrary  $u, v \in [d]$ , let  $\gamma_1 = 2e_{d \times d, uv}$  and  $\gamma_2 = 0_q$ , we will have

$$\frac{1}{\mathcal{L}_{4n}} \left| \Delta(A_{ij}^n)^{(uv)} \right|^2 = \frac{1}{2\mathcal{L}_{4n}} \left| K_{j,i,2e_{d \times d, uv}, 0_q}^n \right| \rightarrow 0, \quad n \rightarrow \infty.$$

Then by taking the summation of the term with  $u, v \in [d]$ , we will have

$$\frac{1}{\mathcal{L}_{4n}} \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} \|\Delta A_{ij}^n\|^2 = \frac{1}{2\mathcal{L}_{4n}} \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} \sum_{u=1}^d \sum_{v=1}^d |K_{j,i,2e_{d \times d, uv}, 0_q}^n| \rightarrow 0. \quad (66)$$

Similarly we will have

$$\frac{1}{\mathcal{L}_{4n}} \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} \|\Delta \eta_{ij}^n\|^2 = \frac{1}{2\mathcal{L}_{4n}} \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} \sum_{w=1}^q |K_{j,i,0_{d \times d}, 2e_{q,w}}^n| \rightarrow 0. \quad (67)$$

Now, consider  $\bar{N} + 1 \leq j \leq N^*$ , such that its corresponding Voronoi cell has only one element, i.e.  $|\mathcal{A}_j| = 1$ . For arbitrary  $u, v \in [d]$ , let  $\tau_1 = e_{d \times d, uv}$ ,  $\tau_2 = 0$  and  $\tau_3 = 0_q$ , we will have

$$\frac{1}{\mathcal{L}_{4n}} \sum_{i \in \mathcal{A}_j} \left| \Delta(A_{ij}^n)^{(uv)} \right| = \frac{1}{\mathcal{L}_{4n}} \left| T_{j,e_{d \times d, uv}, 0, 0_q}^n \right| \rightarrow 0, \quad n \rightarrow \infty.$$

Then by taking the summation of the term with  $u, v \in [d]$ , we will have

$$\frac{1}{\mathcal{L}_{4n}} \sum_{j=\bar{N}+1}^{N^*} \sum_{i \in \mathcal{A}_j} \|\Delta A_{ij}^n\|_1 = \frac{1}{\mathcal{L}_{4n}} \sum_{j=\bar{N}+1}^{N^*} \sum_{u=1}^d \sum_{v=1}^d |T_{j,e_{d \times d,uv},0,0_q}^n| \rightarrow 0.$$

Recall the topological equivalence between  $L_1$ -norm and  $L_2$ -norm on finite-dimensional vector space over  $\mathbb{R}$ , we will have

$$\frac{1}{\mathcal{L}_{4n}} \sum_{j=\bar{N}+1}^{N^*} \sum_{i \in \mathcal{A}_j} \|\Delta A_{ij}^n\| \rightarrow 0. \quad (68)$$

Following a similar argument, since

$$\begin{aligned} \frac{1}{\mathcal{L}_{4n}} \sum_{j=\bar{N}+1}^{N^*} \sum_{i \in \mathcal{A}_j} |\Delta c_{ij}^n| &= \frac{1}{\mathcal{L}_{4n}} \sum_{j=\bar{N}+1}^{N^*} |T_{j,0_{d \times d},1,0_q}^n| \rightarrow 0, \\ \frac{1}{\mathcal{L}_{4n}} \sum_{j=\bar{N}+1}^{N^*} \sum_{i \in \mathcal{A}_j} \|\Delta \eta_{ij}^n\|_1 &= \frac{1}{\mathcal{L}_{4n}} \sum_{j=\bar{N}+1}^{N^*} \sum_{w=1}^q |T_{j,0_{d \times d},0,e_{q,w}}^n| \rightarrow 0, \end{aligned}$$

we obtain that

$$\frac{1}{\mathcal{L}_{4n}} \sum_{j=\bar{N}+1}^{N^*} \sum_{i \in \mathcal{A}_j} |\Delta c_{ij}^n| \rightarrow 0, \quad \frac{1}{\mathcal{L}_{4n}} \sum_{j=\bar{N}+1}^{N^*} \sum_{i \in \mathcal{A}_j} \|\Delta \eta_{ij}^n\| \rightarrow 0. \quad (69)$$

Now taking the summation of limits in equations (65) - (69), we could deduce that

$$\begin{aligned} 1 &= \frac{\mathcal{L}_{4n}}{\mathcal{L}_{4n}} = \frac{1}{\mathcal{L}_{4n}} \sum_{j=1}^{\bar{N}} \left| \sum_{i \in \mathcal{A}_j} \frac{1}{1 + \exp(-c_i^n)} - \frac{1}{1 + \exp(-c_j^*)} \right| + \frac{1}{\mathcal{L}_{4n}} \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} [\|\Delta A_{ij}^n\|^2 + \|\Delta \eta_{ij}^n\|^2] \\ &\quad + \frac{1}{\mathcal{L}_{4n}} \sum_{j=\bar{N}+1}^{N^*} \sum_{i \in \mathcal{A}_j} [\|\Delta A_{ij}^n\| + |\Delta c_{ij}^n| + \|\Delta \eta_{ij}^n\|] \rightarrow 0, \end{aligned}$$

as  $n \rightarrow \infty$ , which is a contradiction. Thus, not all the coefficients of elements in the set

$$\mathcal{S} = \left\{ \frac{K_{j,i,\gamma_{1:2}}^n}{\mathcal{L}_{4n}}, \frac{L_j^n}{\mathcal{L}_{4n}}, \frac{T_{j,\tau_{1:3}}^n}{\mathcal{L}_{4n}} \right\}$$

tend to 0 as  $n \rightarrow \infty$ . Let us denote by  $m_n$  the maximum of the absolute values of those elements. It follows from the previous result that  $1/m_n \not\rightarrow \infty$  as  $n \rightarrow \infty$ .

**Step 3 - Application of Fatou's lemma:** In this step, we apply Fatou's lemma to obtain the desired inequality in equation (55). Recall in the beginning we have assumed that  $\|\tilde{f}_{\tilde{G}_n} - \tilde{f}_{G_*}\|_{L^2(\mu)} / \mathcal{L}_{4n} \rightarrow 0$  and the topological equivalence between  $L_1$ -norm and  $L_2$ -norm on finite-dimensional vector space over  $\mathbb{R}$ , we obtain  $\|\tilde{f}_{\tilde{G}_n} - \tilde{f}_{G_*}\|_{L^1(\mu)} / \mathcal{L}_{4n} \rightarrow 0$ . By applying the Fatou's lemma, we get

$$0 = \lim_{n \rightarrow \infty} \frac{\|\tilde{f}_{\tilde{G}_n} - \tilde{f}_{G_*}\|_{L^1(\mu)}}{m_n \mathcal{L}_{4n}} \geq \int \liminf_{n \rightarrow \infty} \frac{|\tilde{f}_{\tilde{G}_n}(x) - \tilde{f}_{G_*}(x)|}{m_n \mathcal{L}_{4n}} d\mu(x) \geq 0.$$

This result suggests that for almost every  $x$ ,

$$\frac{\tilde{f}_{\tilde{G}_n}(x) - \tilde{f}_{G_*}(x)}{m_n \mathcal{L}_{4n}} \rightarrow 0, \quad (70)$$

recall equation (64), we deduce that

$$\begin{aligned} \frac{\tilde{f}_{\tilde{G}_n}(x) - \tilde{f}_{G_*}(x)}{m_n \mathcal{L}_{4n}} &= \sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} \sum_{|\gamma|=1}^2 \frac{K_{j,i,\gamma_{1:2}}^n}{m_n \mathcal{L}_{4n}} \frac{\partial^{|\gamma_1|} \sigma}{\partial A^{\gamma_1}}(x, 0_{d \times d}, c_i^n) \frac{\partial^{|\gamma_2|} \mathcal{E}}{\partial \eta^{\gamma_2}}(x, \eta_j^*) + \sum_{j=1}^{\bar{N}} \frac{L_j^n}{m_n \mathcal{L}_{4n}} \cdot \mathcal{E}(x, \eta_j^*) \\ &+ \sum_{j=\bar{N}+1}^{N^*} \sum_{|\tau|=1}^3 \frac{T_{j,\tau_{1:3}}^n}{m_n \mathcal{L}_{4n}} \frac{\partial^{|\tau_1|+|\tau_2|} \sigma}{\partial A^{\tau_1} \partial c^{\tau_2}}(x, A_j^*, c_j^*) \frac{\partial^{|\tau_3|} \mathcal{E}}{\partial \eta^{\tau_3}}(x, \eta_j^*) + \frac{R_1(x)}{m_n \mathcal{L}_{4n}} + \frac{R_2(x)}{m_n \mathcal{L}_{4n}}. \end{aligned}$$

Let us denote

$$\frac{K_{j,i,\gamma_{1:2}}^n}{m_n \mathcal{L}_{4n}} \rightarrow k_{j,i,\gamma_{1:2}}, \quad \frac{L_j^n}{m_n \mathcal{L}_{4n}} \rightarrow l_j, \quad \frac{T_{j,\tau_{1:3}}^n}{m_n \mathcal{L}_{4n}} \rightarrow t_{j,\tau_{1:3}},$$

then from equation (70), we will have

$$\begin{aligned} &\sum_{j=1}^{\bar{N}} \sum_{i \in \mathcal{A}_j} \sum_{|\gamma|=1}^2 k_{j,i,\gamma_{1:2}} \frac{\partial^{|\gamma_1|} \sigma}{\partial A^{\gamma_1}}(x, 0_{d \times d}, c_i^n) \frac{\partial^{|\gamma_2|} \mathcal{E}}{\partial \eta^{\gamma_2}}(x, \eta_j^*) + \sum_{j=1}^{\bar{N}} l_j^n \cdot \mathcal{E}(x, \eta_j^*) \\ &+ \sum_{j=\bar{N}+1}^{N^*} \sum_{|\tau|=1}^3 t_{j,\tau_{1:3}}^n \frac{\partial^{|\tau_1|+|\tau_2|} \sigma}{\partial A^{\tau_1} \partial c^{\tau_2}}(x, A_j^*, c_j^*) \frac{\partial^{|\tau_3|} \mathcal{E}}{\partial \eta^{\tau_3}}(x, \eta_j^*) = 0, \end{aligned}$$

for almost every  $x$ . Note that the expert function  $\mathcal{E}(\cdot, \eta)$  is strongly identifiable, then the above equation implies that

$$k_{j,i,\gamma_{1:2}} = l_j = t_{j,\tau_{1:3}} = 0,$$

for any  $j \in [N^*]$ ,  $(\gamma_1, \gamma_2) \in \mathbb{N}^{d \times d} \times \mathbb{N}^q$  and  $(\tau_1, \tau_2, \tau_3) \in \mathbb{N}^{d \times d} \times \mathbb{N} \times \mathbb{N}^q$  such that  $1 \leq \sum_{i=1}^2 |\gamma_i| \leq 2$  and  $\sum_{i=1}^3 |\tau_i| = 1$ . This violates that at least one among the limits in the set  $\{k_{j,i,\gamma_{1:2}}, l_j, t_{j,\tau_{1:3}}\}$  is different from zero.

Thus, we obtain the local inequality in equation (55). Consequently, there exists some  $\varepsilon' > 0$  such that

$$\inf_{G \in \mathcal{M}_N(\Theta): \mathcal{L}_4(G, G_*) \leq \varepsilon'} \frac{\|\tilde{f}_G - \tilde{f}_{G_*}\|_{L^2(\mu)}}{\mathcal{L}_4(G, G_*)} > 0.$$

**Global part:** We now proceed to demonstrate equation (54) for the case where the denominator does not vanish, i.e.

$$\inf_{G \in \mathcal{M}_N(\Theta): \mathcal{L}_4(G, G_*) > \varepsilon'} \frac{\|\tilde{f}_G - \tilde{f}_{G_*}\|_{L^2(\mu)}}{\mathcal{L}_4(G, G_*)} > 0. \quad (71)$$

Suppose, for contradiction, that inequality (71) does not hold. Then there exists a sequence of mixing measures  $G'_n \in \mathcal{M}_N(\Theta)$  such that  $\mathcal{L}_4(G'_n, G_*) > \varepsilon'$  and

$$\lim_{n \rightarrow \infty} \frac{\|\tilde{f}_{G'_n} - \tilde{f}_{G_*}\|_{L^2(\mu)}}{\mathcal{L}_4(G'_n, G_*)} = 0.$$

Under this situation, we could deduce that  $\|\tilde{f}_{G'_n} - \tilde{f}_{G_*}\|_{L^2(\mu)} \rightarrow 0$  as  $n \rightarrow \infty$ . Since  $\Theta$  is a compact set, we can replace the sequence  $G'_n$  with a convergent subsequence, which approaches a mixing measure  $G' \in \mathcal{M}_N(\Theta)$ . Given that  $\mathcal{L}_4(G'_n, G_*) > \varepsilon'$ , we conclude that  $\mathcal{L}_4(G', G_*) > \varepsilon'$ . Then, using Fatou's lemma, we deduce:

$$0 = \lim_{n \rightarrow \infty} \|\tilde{f}_{G'_n} - \tilde{f}_{G_*}\|_{L^2(\mu)}^2 \geq \int \liminf_{n \rightarrow \infty} \left| \tilde{f}_{G'_n}(x) - \tilde{f}_{G_*}(x) \right|^2 d\mu(x),$$

which indicates that  $\tilde{f}_{G'}(x) = \tilde{f}_{G_*}(x)$  for almost every  $x$ . Based on the Proposition 5 followed, we will have that  $G' \equiv G_*$ .  $\square$

**Proposition 5.** *If the equation  $\tilde{f}_G(x) = \tilde{f}_{G_*}(x)$  holds true for almost every  $x$ , then it follows that  $G \equiv G_*$ .*

*Proof.* The proof of Proposition 5 can be done in a similar fashion to that of Proposition 2.  $\square$

### C.2.3 Proof of Theorem 5

*Proof.* Following from the result of Corollary 2, it is sufficient to show that the following inequality holds true:

$$\inf_{G \in \mathcal{M}_N(\Theta)} \frac{\|\tilde{f}_G - \tilde{f}_{\bar{G}}\|_{L^2(\mu)}}{\mathcal{L}_5(G, \bar{G})} > 0, \quad (72)$$

for any mixing measure  $\bar{G} \in \overline{\mathcal{M}}_N(\Theta)$ . To prove the above inequality, we follow a similar approach to the proof in Appendix B.4, dividing the analysis into a local part and a global part and omitting the global part.

Therefore, for an arbitrary mixing measure  $\bar{G} := \sum_{i=1}^N \frac{1}{1+\exp(-\bar{c}_i)} \delta_{(\bar{A}_i, \bar{\eta}_i)} \in \overline{\mathcal{M}}_N(\Theta)$ , we focus exclusively on demonstrating that:

$$\lim_{\varepsilon \rightarrow 0} \inf_{G \in \mathcal{M}_N(\Theta): \mathcal{L}_5(G, \bar{G}) \leq \varepsilon} \frac{\|\tilde{f}_G - \tilde{f}_{\bar{G}}\|_{L^2(\mu)}}{\mathcal{L}_5(G, \bar{G})} > 0. \quad (73)$$

Assume, for contradiction, that the above claim does not hold. Then there exists a sequence of mixing measures  $G_n = \sum_{i=1}^N \frac{1}{1+\exp(-c_i^n)} \delta_{(A_i^n, \eta_i^n)}$  in  $\mathcal{M}_N(\Theta)$  such that as  $n \rightarrow \infty$ , we get

$$\begin{cases} \mathcal{L}_{5n} := \mathcal{L}_5(\tilde{G}_n, \bar{G}) \rightarrow 0, \\ \|\tilde{f}_{\tilde{G}_n} - \tilde{f}_{\bar{G}}\|_{L^2(\mu)} / \mathcal{L}_{5n} \rightarrow 0. \end{cases} \quad (74)$$

Let us denote by  $\mathcal{A}_i^n := \mathcal{A}_i(\tilde{G}_n)$  a Voronoi cell of  $\tilde{G}_n$  generated by the  $j$ -th components of  $\overline{\overline{G}}$ . Since our analysis is asymptotic, we can assume that the Voronoi cells are independent of the sample size, i.e.  $\mathcal{A}_j = \mathcal{A}_i^n$ .

In Dense Regime, since  $\tilde{G}_n$  and  $\overline{\overline{G}}$  have the same number of atoms  $k$ , and  $\mathcal{L}_{5n} \rightarrow 0$ , it follows that each Voronoi cell  $\mathcal{A}_i$  contains precisely one element for all  $i \in [N]$ . Without loss of generality, we assume  $\mathcal{A}_i = \{i\}$  for all  $i \in [N]$ . This ensures that  $(A_i^n, c_i^n, \eta_i^n) \rightarrow (\overline{\overline{A}}_i, \overline{\overline{c}}_i, \overline{\overline{\eta}}_i)$  as  $n \rightarrow \infty$  for every  $i \in [N]$ .

Consequently, the Voronoi loss  $\mathcal{L}_{5n}$  can be expressed as:

$$\mathcal{L}_{5n} := \sum_{i=1}^N \left( \|\Delta \overline{\overline{A}}_i^n\| + |\Delta \overline{\overline{c}}_i^n| + \|\overline{\overline{\eta}}_i^n\| \right), \quad (75)$$

where the increments are given by:  $\Delta \overline{\overline{A}}_i^n = A_i^n - \overline{\overline{A}}_i$ ,  $\Delta \overline{\overline{c}}_i^n = c_i^n - \overline{\overline{c}}_i$ ,  $\Delta \overline{\overline{\eta}}_i^n = \eta_i^n - \overline{\overline{\eta}}_i$ .

We now break the proof of the local part into the following three steps:

**Step 1 - Taylor expansion:** In this step, we decompose the term  $\tilde{f}_{\tilde{G}_n}(x) - \tilde{f}_{\overline{\overline{G}}}(x)$  using a Taylor expansion. First, let us denote  $\sigma(x, A, c) := \frac{1}{1 + \exp(-x^\top A x - c)}$ , then we have that

$$\begin{aligned} \tilde{f}_{\tilde{G}_n}(x) - \tilde{f}_{\overline{\overline{G}}}(x) &= \sum_{i=1}^N \left[ \frac{1}{1 + \exp(-x^\top A_i^n x - c_i^n)} \cdot \mathcal{E}(x, \eta_i^n) - \frac{1}{1 + \exp(-x^\top \overline{\overline{A}}_i x - \overline{\overline{c}}_i)} \cdot \mathcal{E}(x, \overline{\overline{\eta}}_i) \right] \\ &= \sum_{i=1}^N \sum_{|\alpha|=1} \frac{1}{\alpha!} (\Delta \overline{\overline{A}}_i^n)^{\alpha_1} (\Delta \overline{\overline{c}}_i^n)^{\alpha_2} (\Delta \overline{\overline{\eta}}_i^n)^{\alpha_3} \frac{\partial^{|\alpha_1|+|\alpha_2|} \sigma}{\partial A^{\alpha_1} \partial c^{\alpha_2}}(x, \overline{\overline{A}}_i, \overline{\overline{c}}_i) \frac{\partial^{|\alpha_3|} \mathcal{E}}{\partial \eta^{\alpha_3}}(x, \overline{\overline{\eta}}_i) + R_1(x) \\ &= \sum_{i=1}^N \sum_{|\alpha|=1} S_{i,\alpha_1,\alpha_2,\alpha_3}^n \frac{\partial^{|\alpha_1|+|\alpha_2|} \sigma}{\partial A^{\alpha_1} \partial c^{\alpha_2}}(x, \overline{\overline{A}}_i, \overline{\overline{c}}_i) \frac{\partial^{|\alpha_3|} \mathcal{E}}{\partial \eta^{\alpha_3}}(x, \overline{\overline{\eta}}_i) + R_1(x), \end{aligned} \quad (76)$$

where  $R_1(x)$  is a Taylor remainder such that  $R_1(x)/\mathcal{L}_{5n} \rightarrow 0$  as  $n \rightarrow \infty$ . Now we could denote

$$S_{i,\alpha_1:3}^n = \frac{1}{\alpha!} (\Delta \overline{\overline{A}}_i^n)^{\alpha_1} (\Delta \overline{\overline{c}}_i^n)^{\alpha_2} (\Delta \overline{\overline{\eta}}_i^n)^{\alpha_3}, \quad i \in [N], \alpha \in \mathbb{R}^{d \times d} \times \mathbb{R} \times \mathbb{R}^q \quad (77)$$

where  $\sum_{i=1}^3 |\alpha_i| = 1$ .

**Step 2 - Non-vanishing coefficients:** Now we claim that at least one among the ratios  $S_{i,\alpha_1:3}^n / \mathcal{L}_{5n}$  will not vanish as  $n$  goes to infinity. We prove by contradiction that all of them converge to zero when  $n \rightarrow 0$ :

$$\frac{S_{i,\alpha_1:3}^n}{\mathcal{L}_{5n}} \rightarrow 0.$$

for any  $i \in [N], \alpha \in \mathbb{R}^{d \times d} \times \mathbb{R} \times \mathbb{R}^q$  such that  $\sum_{i=1}^3 |\alpha_i| = 1$ .

Now, consider  $1 \leq i \leq k$ , for arbitrary  $u, v \in [d]$ , let  $\alpha_1 = e_{d \times d, uv}, \alpha_2 = 0$  and  $\alpha_3 = 0_q$ , we will have

$$\frac{1}{\mathcal{L}_{5n}} \sum_{i=1}^N \left| \Delta(\overline{\overline{A}}_i^n)^{(uv)} \right| = \frac{1}{\mathcal{L}_{5n}} \left| S_{i,e_{d \times d,uv},0,0_q}^n \right| \rightarrow 0, \quad n \rightarrow \infty.$$

Then by taking the summation of the term with  $u, v \in [d]$ , we will have

$$\frac{1}{\mathcal{L}_{5n}} \sum_{i=1}^N \|\Delta \bar{\bar{A}}_i^n\|_1 = \frac{1}{\mathcal{L}_{5n}} \sum_{i=1}^N \sum_{u=1}^d \sum_{v=1}^d |S_{i, e_{d \times d, u, v}, 0, 0_q}^n| \rightarrow 0.$$

Recall the topological equivalence between  $L_1$ -norm and  $L_2$ -norm on finite-dimensional vector space over  $\mathbb{R}$ , we will have

$$\frac{1}{\mathcal{L}_{5n}} \sum_{i=1}^N \|\Delta \bar{\bar{A}}_i^n\| \rightarrow 0. \quad (78)$$

Following a similar argument, since

$$\begin{aligned} \frac{1}{\mathcal{L}_{5n}} \sum_{i=1}^N |\Delta \bar{\bar{c}}_i^n| &= \frac{1}{\mathcal{L}_{5n}} \sum_{i=1}^N |S_{i, 0_{d \times d}, 1, 0_q}^n| \rightarrow 0, \\ \frac{1}{\mathcal{L}_{5n}} \sum_{i=1}^N \|\Delta \bar{\bar{\eta}}_i^n\|_1 &= \frac{1}{\mathcal{L}_{1n}} \sum_{i=1}^N \sum_{w=1}^q |S_{i, 0_{d \times d}, 0, e_{q, w}}^n| \rightarrow 0, \end{aligned}$$

we obtain that

$$\frac{1}{\mathcal{L}_{5n}} \sum_{i=1}^N |\Delta \bar{\bar{c}}_i^n| \rightarrow 0, \quad \frac{1}{\mathcal{L}_{5n}} \sum_{i=1}^N \|\Delta \bar{\bar{\eta}}_i^n\| \rightarrow 0. \quad (79)$$

Now taking the summation of limits in equations (78) - (79), we could deduce that

$$1 = \frac{\mathcal{L}_{5n}}{\mathcal{L}_{5n}} = \frac{1}{\mathcal{L}_{5n}} \cdot \sum_{i=1}^N \left( \|\Delta \bar{\bar{A}}_i^n\| + |\Delta \bar{\bar{c}}_i^n| + \|\Delta \bar{\bar{\eta}}_i^n\| \right) \rightarrow 0,$$

as  $n \rightarrow \infty$ , which is a contradiction. Thus, at least one among the ratios  $S_{i, \alpha_{1:3}}^n / \mathcal{L}_{5n}$  must not approach zero as  $n \rightarrow \infty$ . Let us denote by  $m_n$  the maximum of the absolute values of those elements. It follows from the previous result that  $1/m_n \not\rightarrow \infty$  as  $n \rightarrow \infty$ .

**Step 3 - Application of Fatou's lemma:** In this step, we apply Fatou's lemma to obtain the desired inequality in equation (73). Recall in the beginning we have assumed that  $\|\tilde{f}_{\tilde{G}_n} - \tilde{f}_{\bar{\bar{G}}}\|_{L^2(\mu)} / \mathcal{L}_{5n} \rightarrow 0$  and the topological equivalence between  $L_1$ -norm and  $L_2$ -norm on finite-dimensional vector space over  $\mathbb{R}$ , we obtain  $\|\tilde{f}_{\tilde{G}_n} - \tilde{f}_{\bar{\bar{G}}}\|_{L^1(\mu)} / \mathcal{L}_{5n} \rightarrow 0$ . By applying the Fatou's lemma, we get

$$0 = \lim_{n \rightarrow \infty} \frac{\|\tilde{f}_{\tilde{G}_n} - \tilde{f}_{\bar{\bar{G}}}\|_{L^1(\mu)}}{m_n \mathcal{L}_{5n}} \geq \int \liminf_{n \rightarrow \infty} \frac{|\tilde{f}_{\tilde{G}_n}(x) - \tilde{f}_{\bar{\bar{G}}}(x)|}{m_n \mathcal{L}_{5n}} d\mu(x) \geq 0.$$

This result suggests that for almost every  $x$ ,

$$\frac{\tilde{f}_{\tilde{G}_n}(x) - \tilde{f}_{\bar{\bar{G}}}(x)}{m_n \mathcal{L}_{5n}} \rightarrow 0. \quad (80)$$

Let us denote

$$\frac{S_{i,\alpha_{1:3}}^n}{m_n \mathcal{L}_{5n}} \rightarrow s_{i,\alpha_{1:3}},$$

as  $n \rightarrow \infty$  with a note that at least one among the limits  $s_{i,\alpha_{1:3}}$  is non-zero. Then from equation (80), we will have

$$\sum_{i=1}^N \sum_{|\alpha|=1} s_{i,\alpha_1,\alpha_2,\alpha_3}^n \frac{\partial^{|\alpha_1|+|\alpha_2|} \sigma}{\partial A^{\alpha_1} \partial c^{\alpha_2}}(x, \bar{A}_i, \bar{c}_i) \frac{\partial^{|\alpha_3|} \mathcal{E}}{\partial \eta^{\alpha_3}}(x, \bar{\eta}_i) = 0,$$

for almost every  $x$ . Note that the expert function  $\mathcal{E}(\cdot, \eta)$  is weakly identifiable, then the above equation implies that

$$s_{i,\alpha_{1:3}} = 0,$$

for any  $i \in [N]$ ,  $(\alpha_1, \alpha_2, \alpha_3) \in \mathbb{N}^{d \times d} \times \mathbb{N} \times \mathbb{N}^q$  such that  $\sum_{i=1}^3 |\alpha_i| = 1$ . This violates that at least one among the limits in the set  $\{s_{i,\alpha_{1:3}}\}$  is different from zero.

Thus, we obtain the local inequality in equation (73). Completes the proof.  $\square$

## D Additional Experiments

**Sigmoid gating convergence rate.** To empirically validate the theoretical findings presented in Section 4, we analyze the parameter estimation rates of the sigmoid quadratic gating MoE in a well-specified setting. Our evaluation spans various expert configurations, considering both dense and sparse regimes, to assess the sample complexity of the model.

**Setup.** Synthetic data is generated using a sigmoid quadratic gating MoE model with both ReLU and linear experts. Models are then fitted to the data for varying sample sizes, and the empirical parameter estimation rates are evaluated. Comparisons are made for ReLU and linear experts under both dense and sparse regimes to verify the theoretical findings. Please refer to Appendix E for a detailed description of synthetic data generation and ground-truth parameters.

**Results.** Figure 2 illustrates the empirical convergence rates of Voronoi loss for different expert configurations in the sigmoid quadratic gating MoE model under both sparse and dense regimes. Error bars indicate three standard deviations to account for variability across different runs. Specifically, Figure 2a compares the empirical Voronoi loss for MoE models with ReLU and linear experts in the sparse regime. The sigmoid quadratic gating MoE with linear experts exhibits a slower convergence rate of  $\mathcal{O}(n^{-0.07})$ , whereas the model with ReLU experts achieves a significantly faster rate of  $\mathcal{O}(n^{-0.54})$ .

Similarly, Figure 2b presents results for the dense regime, where both expert types achieve fast convergence rates. Notably, the MoE model with ReLU experts attains a convergence rate of  $\mathcal{O}(n^{-0.53})$ , while the model with linear experts follows with a rate of  $\mathcal{O}(n^{-0.44})$ . These results empirically validate our theoretical predictions, demonstrating that the sigmoid quadratic gating MoE yields superior parameter estimation rates across different expert configurations, particularly benefiting from the use of ReLU experts.

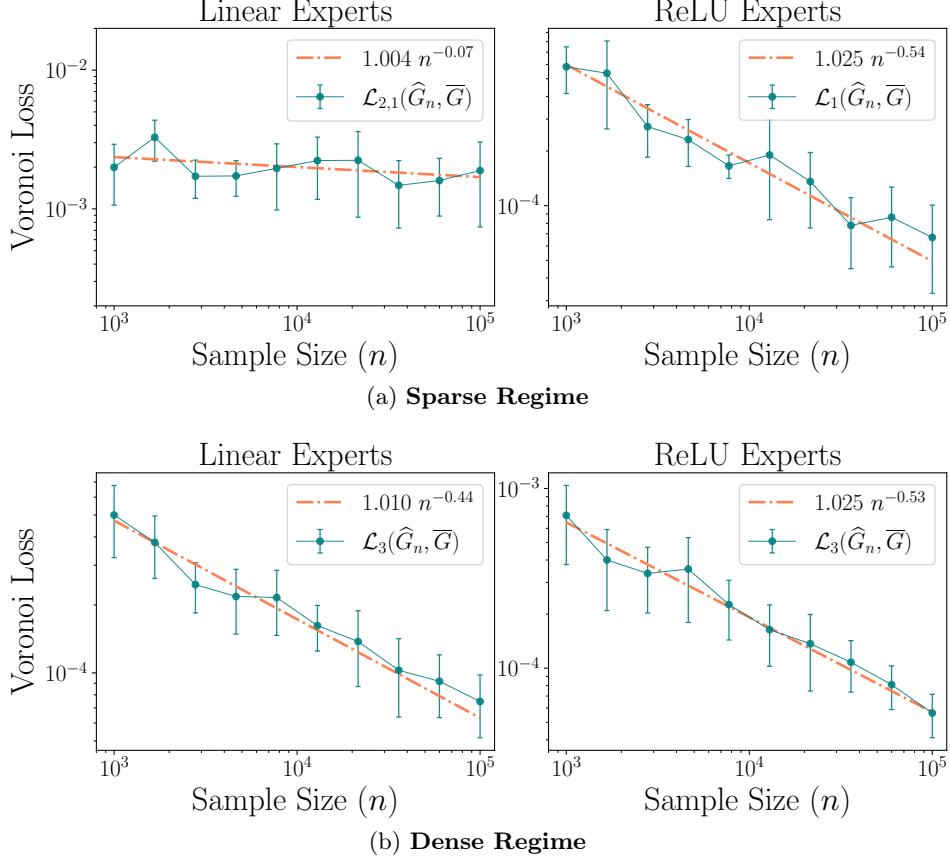


Figure 2: Log-log plots of empirical convergence rates of Voronoi losses for softmax and sigmoid quadratic gating mechanisms. 2a Comparison between sigmoid quadratic gating MoE with ReLU and linear experts under a sparse regime for ground-truth parameters. 2b Comparison between sigmoid quadratic gating with ReLU and linear experts under a dense regime for ground-truth parameters. Each plot illustrates the empirical Voronoi loss convergence rates, with solid lines representing the Voronoi losses and dash-dotted lines showing fitted trends to emphasize the empirical rates.

## E Experimental Details

In this appendix, we provide a comprehensive overview of the generation process of synthetic data in our numerical experiments.

For each experiment, we generate synthetic data using the corresponding MoE model. Specifically, we construct a dataset  $\{(X_i, Y_i)\}_{i=1}^n \subset \mathbb{R}^d \times \mathbb{R}$  as follows. First, we sample the input features independently from a uniform distribution:

$$X_i \sim \text{Uniform}([-1, 1]^d), \quad \text{for } i = 1, \dots, n. \quad (81)$$

The response variable  $Y_i$  is then generated according to the following model:

$$Y_i = f_{G_*}(X_i) + \varepsilon_i, \quad \text{for } i = 1, \dots, n, \quad (82)$$

where  $\varepsilon_i$  represents independent Gaussian noise with variance  $\nu = 0.01$ . The underlying regression

function  $f_{G_*}(\cdot)$  is defined as:

$$f_{G_*}(x) = \sum_{i=1}^k \mathcal{G}_i(x) \cdot \phi((\alpha_i^*)^\top x + \beta_i^*), \quad (83)$$

where  $\mathcal{G}_i(x)$  represents the gating function, which is modeled using either a softmax or a sigmoid quadratic transformation, and  $\phi(\cdot)$  denotes the expert activation function, which can be either a ReLU or a linear transformation. The input dimension is set to  $d = 8$ , and we employ  $N^* = 8$  experts. The gating and expert parameters are drawn from pre-specified distributions, as described in the following sections.

**Ground-truth gating parameters.** In the dense regime, for both sigmoid and softmax quadratic gating functions, the ground-truth parameters  $(A_i^*, b_i^*, c_i^*)$  are independently drawn from an isotropic Gaussian distribution with variance  $\nu_g = 1/d$  for all  $1 \leq i \leq 8$ .

In the sparse regime, we enforce sparsity by setting  $(A_i^*, b_i^*, c_i^*) = 0$  for  $i = 7, 8$ , while the remaining parameters are sampled as in the dense regime.

**Ground-truth expert parameters.** All ground-truth expert parameters are similarly drawn independently from an isotropic Gaussian distribution with zero mean and variance  $\nu_e = 1/d$ . These parameters are sampled once before running the experiments and remain unchanged throughout all trials.

**Computational Resources.** All numerical experiments were conducted on a MacBook Air equipped with an Apple M4 chip.

## References

- [1] P. Akbarian, H. Nguyen, X. Han, and N. Ho. Quadratic gating functions in mixture of experts: A statistical insight. *arXiv preprint arXiv:2410.11222*, 2024. (Cited on pages 2, 5, 9, 10, 12, and 33.)
- [2] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. (Cited on page 1.)
- [3] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. (Cited on page 1.)
- [4] L. Chen, K. Lu, A. Rajeswaran, K. Lee, A. Grover, M. Laskin, P. Abbeel, A. Srinivas, and I. Mordatch. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34:15084–15097, 2021. (Cited on page 1.)
- [5] Z. Chen, Y. Deng, Y. Wu, Q. Gu, and Y. Li. Towards understanding the mixture-of-experts layer in deep learning. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 23049–23062. Curran Associates, Inc., 2022. (Cited on page 3.)
- [6] Z. Chen, Y. Shen, M. Ding, Z. Chen, H. Zhao, E. G. Learned-Miller, and C. Gan. Mod-squad: Designing mixtures of experts as modular multi-task learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11828–11837, 2023. (Cited on page 3.)

- [7] Z. Chi, L. Dong, S. Huang, D. Dai, S. Ma, B. Patra, S. Singhal, P. Bajaj, X. Song, X.-L. Mao, H. Huang, and F. Wei. On the representation collapse of sparse mixture of experts. In A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, editors, *Advances in Neural Information Processing Systems*, 2022. (Cited on page 3.)
- [8] R. Csordás, K. Irie, and J. Schmidhuber. Approximating two-layer feedforward networks for efficient transformers. In H. Bouamor, J. Pino, and K. Bali, editors, *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 674–692, Singapore, Dec. 2023. Association for Computational Linguistics. (Cited on page 3.)
- [9] R. Csordás, P. Piékko, K. Irie, and J. Schmidhuber. Switchhead: Accelerating transformers with mixture-of-experts attention. *arXiv preprint arXiv:2312.07987*, 2023. (Cited on pages 2 and 12.)
- [10] T. Dao. Flashattention-2: Faster attention with better parallelism and work partitioning. In *The Twelfth International Conference on Learning Representations*, 2024. (Cited on page 2.)
- [11] T. Dao, D. Y. Fu, S. Ermon, A. Rudra, and C. Re. Flashattention: Fast and memory-efficient exact attention with IO-awareness. In *Advances in Neural Information Processing Systems*, 2022. (Cited on page 2.)
- [12] DeepSeek-AI, A. Liu, B. Feng, B. Xue, B. Wang, B. Wu, C. Lu, C. Zhao, C. Deng, C. Zhang, C. Ruan, D. Dai, D. Guo, D. Yang, D. Chen, D. Ji, E. Li, F. Lin, F. Dai, F. Luo, G. Hao, G. Chen, G. Li, H. Zhang, H. Bao, H. Xu, H. Wang, H. Zhang, H. Ding, H. Xin, H. Gao, H. Li, H. Qu, J. L. Cai, J. Liang, J. Guo, J. Ni, J. Li, J. Wang, J. Chen, J. Chen, J. Yuan, J. Qiu, J. Li, J. Song, K. Dong, K. Hu, K. Gao, K. Guan, K. Huang, K. Yu, L. Wang, L. Zhang, L. Xu, L. Xia, L. Zhao, L. Wang, L. Zhang, M. Li, M. Wang, M. Zhang, M. Zhang, M. Tang, M. Li, N. Tian, P. Huang, P. Wang, P. Zhang, Q. Wang, Q. Zhu, Q. Chen, Q. Du, R. J. Chen, R. L. Jin, R. Ge, R. Zhang, R. Pan, R. Wang, R. Xu, R. Zhang, R. Chen, S. S. Li, S. Lu, S. Zhou, S. Chen, S. Wu, S. Ye, S. Ye, S. Ma, S. Wang, S. Zhou, S. Yu, S. Zhou, S. Pan, T. Wang, T. Yun, T. Pei, T. Sun, W. L. Xiao, W. Zeng, W. Zhao, W. An, W. Liu, W. Liang, W. Gao, W. Yu, W. Zhang, X. Q. Li, X. Jin, X. Wang, X. Bi, X. Liu, X. Wang, X. Shen, X. Chen, X. Zhang, X. Chen, X. Nie, X. Sun, X. Wang, X. Cheng, X. Liu, X. Xie, X. Liu, X. Yu, X. Song, X. Shan, X. Zhou, X. Yang, X. Li, X. Su, X. Lin, Y. K. Li, Y. Q. Wang, Y. X. Wei, Y. X. Zhu, Y. Zhang, Y. Xu, Y. Xu, Y. Huang, Y. Li, Y. Zhao, Y. Sun, Y. Li, Y. Wang, Y. Yu, Y. Zheng, Y. Zhang, Y. Shi, Y. Xiong, Y. He, Y. Tang, Y. Piao, Y. Wang, Y. Tan, Y. Ma, Y. Liu, Y. Guo, Y. Wu, Y. Ou, Y. Zhu, Y. Wang, Y. Gong, Y. Zou, Y. He, Y. Zha, Y. Xiong, Y. Ma, Y. Yan, Y. Luo, Y. You, Y. Liu, Y. Zhou, Z. F. Wu, Z. Z. Ren, Z. Ren, Z. Sha, Z. Fu, Z. Xu, Z. Huang, Z. Zhang, Z. Xie, Z. Zhang, Z. Hao, Z. Gou, Z. Ma, Z. Yan, Z. Shao, Z. Xu, Z. Wu, Z. Zhang, Z. Li, Z. Gu, Z. Zhu, Z. Liu, Z. Li, Z. Xie, Z. Song, Z. Gao, and Z. Pan. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*, 2024. (Cited on page 3.)
- [13] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *North American Chapter of the Association for Computational Linguistics*, 2019. (Cited on page 1.)
- [14] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby. An image is worth 16x16

- words: Transformers for image recognition at scale. *ArXiv*, abs/2010.11929, 2020. (Cited on page 1.)
- [15] N. Du, Y. Huang, A. M. Dai, S. Tong, D. Lepikhin, Y. Xu, M. Krikun, Y. Zhou, A. W. Yu, O. Firat, B. Zoph, L. Fedus, M. Bosma, Z. Zhou, T. Wang, Y. E. Wang, K. Webster, M. Pellat, K. Robinson, K. S. Meier-Hellstern, T. Duke, L. Dixon, K. Zhang, Q. V. Le, Y. Wu, Z. Chen, and C. Cui. Glam: Efficient scaling of language models with mixture-of-experts. In *International Conference on Machine Learning*, 2021. (Cited on page 3.)
- [16] S. Faria and G. Soromenho. Fitting mixtures of linear regressions. *Journal of Statistical Computation and Simulation*, 80(2):201–225, 2010. (Cited on page 3.)
- [17] W. Fedus, B. Zoph, and N. Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23:1–39, 2022. (Cited on page 3.)
- [18] N. Gaur, B. Farris, P. Haghani, I. Leal, P. J. Moreno, M. Prasad, B. Ramabhadran, and Y. Zhu. Mixture of informed experts for multilingual speech recognition. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6234–6238. IEEE, 2021. (Cited on page 3.)
- [19] X. Gu, T. Pang, C. Du, Q. Liu, F. Zhang, C. Du, Y. Wang, and M. Lin. When attention sink emerges in language models: An empirical view. *arXiv preprint arXiv:2410.10781*, 2024. (Cited on page 1.)
- [20] H. Hazimeh, Z. Zhao, A. Chowdhery, M. Sathiamoorthy, Y. Chen, R. Mazumder, L. Hong, and E. Chi. Dselect-k: Differentiable selection in the mixture of experts with applications to multi-task learning. *Advances in Neural Information Processing Systems*, 34:29335–29347, 2021. (Cited on page 3.)
- [21] N. Ho, C.-Y. Yang, and M. I. Jordan. Convergence rates for Gaussian mixtures of experts. *Journal of Machine Learning Research*, 23(323):1–81, 2022. (Cited on page 13.)
- [22] E. S. Hu, K. Ahn, Q. Liu, H. Xu, M. Tomar, A. Langford, D. Jayaraman, A. Lamb, and J. Langford. Learning to achieve goals with belief state transformers. *ArXiv*, abs/2410.23506, 2024. (Cited on page 1.)
- [23] R. A. Jacobs, M. I. Jordan, S. J. Nowlan, and G. E. Hinton. Adaptive mixtures of local experts. *Neural Computation*, 3, 1991. (Cited on pages 2 and 3.)
- [24] A. Jain, V. P. Singh, and S. P. Rath. A multi-accent acoustic model using mixture of experts for speech recognition. In *Interspeech*, pages 779–783, 2019. (Cited on page 3.)
- [25] A. Q. Jiang, A. Sablayrolles, A. Roux, A. Mensch, B. Savary, C. Bamford, D. S. Chaplot, D. de las Casas, E. B. Hanna, F. Bressand, G. Lengyel, G. Bour, G. Lample, L. R. Lavaud, L. Saulnier, M.-A. Lachaux, P. Stock, S. Subramanian, S. Yang, S. Antoniak, T. L. Scao, T. Gervet, T. Lavril, T. Wang, T. Lacroix, and W. E. Sayed. Mixtral of experts. *arxiv preprint arxiv 2401.04088*, 2024. (Cited on page 3.)

- [26] P. Jin, B. Zhu, L. Yuan, and S. Yan. Moh: Multi-head attention as mixture-of-head attention. *arXiv preprint arXiv:2410.11842*, 2024. (Cited on page 12.)
- [27] M. I. Jordan and R. A. Jacobs. Hierarchical mixtures of experts and the EM algorithm. *Neural Computation*, 6:181–214, 1994. (Cited on pages 3 and 12.)
- [28] C. Kim, J. Park, J. Shin, H. Lee, P. Abbeel, and K. Lee. Preference transformer: Modeling human preferences using transformers for rl. *arXiv preprint arXiv:2303.00957*, 2023. (Cited on page 1.)
- [29] Y. Kwon and S.-W. Chung. Mole: Mixture of language experts for multi-lingual automatic speech recognition. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2023. (Cited on page 3.)
- [30] B. Lindsay. *Mixture models: Theory, geometry and applications*. In NSF-CBMS Regional Conference Series in Probability and Statistics. IMS, Hayward, CA., 1995. (Cited on page 3.)
- [31] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. (Cited on page 1.)
- [32] J. Ma, Z. Zhao, X. Yi, J. Chen, L. Hong, and E. H. Chi. Modeling task relationships in multi-task learning with multi-gate mixture-of-experts. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 1930–1939, 2018. (Cited on page 3.)
- [33] T. Manole and N. Ho. Refined convergence rates for maximum likelihood estimation under finite mixture models. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 14979–15006. PMLR, 17–23 Jul 2022. (Cited on page 6.)
- [34] E. F. Mendes and W. Jiang. Convergence rates for mixture-of-experts. *arXiv preprint arXiv:1110.2058*, 2011. (Cited on page 13.)
- [35] N. Muennighoff, L. Soldaini, D. Groeneveld, K. Lo, J. D. Morrison, S. Min, W. Shi, P. Walsh, O. Tafjord, N. Lambert, Y. Gu, S. Arora, A. Bhagia, D. Schwenk, D. Wadden, A. Wettig, B. Hui, T. Dettmers, D. Kiela, A. Farhadi, N. A. Smith, P. W. Koh, A. Singh, and H. Hajishirzi. Olmoe: Open mixture-of-experts language models. *ArXiv*, abs/2409.02060, 2024. (Cited on page 3.)
- [36] H. Nguyen, P. Akbarian, F. Yan, and N. Ho. Statistical perspective of top-k sparse softmax gating mixture of experts. In *International Conference on Learning Representations*, 2024. (Cited on page 13.)
- [37] H. Nguyen, X. Han, C. W. Harris, S. Saria, and N. Ho. On expert estimation in hierarchical mixture of experts: Beyond softmax gating functions. *arXiv preprint arXiv:2410.02935*, 2024. (Cited on page 12.)
- [38] H. Nguyen, N. Ho, and A. Rinaldo. On least square estimation in softmax gating mixture of experts. In *Proceedings of the ICML*, 2024. (Cited on pages 6 and 12.)

- [39] H. Nguyen, T. Nguyen, and N. Ho. Demystifying softmax gating function in Gaussian mixture of experts. In *Advances in Neural Information Processing Systems*, 2023. (Cited on page 13.)
- [40] J. Puigcerver, C. Riquelme, B. Mustafa, and N. Houlsby. From sparse to soft mixtures of experts. In *The Twelfth International Conference on Learning Representations*, 2024. (Cited on page 3.)
- [41] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. (Cited on page 1.)
- [42] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67, 2020. (Cited on page 1.)
- [43] J. Ramapuram, F. Danieli, E. Dhekane, F. Weers, D. Busbridge, P. Ablin, T. Likhomanenko, J. Digani, Z. Gu, A. Shidani, and R. Webb. Theory, analysis, and best practices for sigmoid self-attention. *arXiv preprint arXiv:2409.04431*, 2024. (Cited on pages 1, 2, and 12.)
- [44] N. Shazeer, A. Mirhoseini, K. Maziarz, A. Davis, Q. Le, G. Hinton, and J. Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In *In International Conference on Learning Representations*, 2017. (Cited on page 3.)
- [45] H. Touvron, T. Lavig, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. (Cited on page 1.)
- [46] S. van de Geer. *Empirical Processes in M-estimation*. Cambridge University Press, 2000. (Cited on pages 5 and 13.)
- [47] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. (Cited on pages 1, 3, and 12.)
- [48] X. Wu, S. Huang, W. Wang, and F. Wei. Multi-head mixture-of-experts. *arXiv preprint arXiv:2404.15045*, 2024. (Cited on pages 2 and 12.)
- [49] F. Yan, H. Nguyen, D. Le, P. Akbarian, and N. Ho. Understanding expert structures on minimax parameter estimation in contaminated mixture of experts. In *Proceedings of The 28th International Conference on Artificial Intelligence and Statistics*, 2025. (Cited on page 13.)
- [50] Z. You, S. Feng, D. Su, and D. Yu. Speechmoe: Scaling to large acoustic models with dynamic routing mixture of experts. *arXiv preprint arXiv:2105.03036*, 2021. (Cited on page 3.)
- [51] Y. Zhou, N. Du, Y. Huang, D. Peng, C. Lan, D. Huang, S. Shakeri, D. So, A. Dai, Y. Lu, Z. Chen, Q. Le, C. Cui, J. Laudon, and J. Dean. Brainformers: Trading simplicity for efficiency. In *International Conference on Machine Learning*, pages 42531–42542. PMLR, 2023. (Cited on page 3.)