

The Illusion of Specialization: Unveiling the Domain-Invariant "Standing Committee" in Mixture-of-Experts Models

Yan Wang^{1*†}, Yitao Xu^{2*}, Nanhan Shen^{2*},
Jinyan Su³, Jimin Huang^{5,1}, Zining Zhu⁴

¹The Fin AI ²Georgia Institute of Technology

³Cornell University ⁴Stevens Institute of Technology ⁵The University of Manchester

Emails: wy2266336@gmail.com, zzhu41@stevens.edu

Abstract

Mixture of Experts models are widely assumed to achieve domain specialization through sparse routing. In this work, we question this assumption by introducing COMMITTEEAUDIT, a post hoc framework that analyzes routing behavior at the level of expert groups rather than individual experts. Across three representative models and the MMLU benchmark, we uncover a domain invariant **Standing Committee**. This is a compact coalition of routed experts that consistently captures the majority of routing mass across domains, layers, and routing budgets, even when architectures already include shared experts. Qualitative analysis further shows that Standing Committees anchor reasoning structure and syntax, while peripheral experts handle domain-specific knowledge. These findings reveal a strong structural bias toward centralized computation, suggesting that specialization in Mixture of Experts models is far less pervasive than commonly believed. Crucially, this inherent bias indicates that current training objectives, such as load-balancing losses that enforce uniform expert utilization, may be working against the model’s natural optimization path, thereby limiting training efficiency and performance. The code is available at GitHub¹.


1 Introduction

Large Language Models (LLMs) have demonstrated remarkable capabilities in complex reasoning and understanding tasks (Qian et al., 2025; Wang et al., 2025a,b,c). To further scale these models without incurring proportional computational costs, the Mixture-of-Experts (MoE) architecture has emerged as a dominant approach. By activating only a sparse subset of parameters for each input token, MoE models promise to decouple model


* These authors contributed equally as co-first authors.

† Corresponding author

¹<https://github.com/The-FinAI/CommitteeAudit>



(a) Conventional domain-specific intuition.



(b) Empirically observed expert sharing.

Figure 1: From domain-specific intuition to empirically observed expert sharing in Mixture-of-Experts models. (a) **The Intuition:** The ideal “Divide-and-Conquer” strategy assumes disjoint sets of experts for different domains. (b) **The Observation:** Empirical routing patterns reveal a **Standing Committee** (e.g., Experts E4 and E5) that is consistently activated across disparate domains (Math, Legal, Biology), acting as a generalist core hidden within the routed experts.

capacity from inference latency. This conditional computation paradigm is particularly appealing for general-purpose LLMs because different domains exhibit heterogeneous computational patterns that theoretically benefit from expert specialization.

The design philosophy of MoE models often follows a divide-and-conquer principle where experts are expected to specialize by domain. Under this view, sparsity arises when the model routes inputs to distinct expert groups, such as a dedicated set for mathematics and another for legal tasks, as illustrated in Figure 1a. However, the optimization dynamics of sparse routing frequently contradict this ideal separation. Prior research on *Representation Collapse* (Chi et al., 2022; Do et al., 2025) warns of a pathological state where gating networks fail to optimize effectively, causing experts to become redundant or completely inactive. Recognizing that

Table 1: Comparison with existing MoE interpretability works. Prior studies mainly focus on individual expert specialization, internal representations, or frequency-based importance. In contrast, our work identifies stable, domain-invariant expert committees.

| Research Focus | Representative Works | Unit of Analysis | Dominant Assumption | Captures Co-activation |
|--|---|------------------|---------------------|------------------------|
| Routing & Behavior Analysis | A Closer Look into MoE (Lo et al., 2025) | Individual | General Polysemy | ✗ |
| | Probing Semantic Routing (Olson et al., 2025) | Individual | Domain-Specific | ✗ |
| | Context Faithfulness (Bai et al., 2025) | Individual | Context-Dependent | ✗ |
| → Focus: Which individual expert activates for a specific token? | | | | |
| Representation & Intrinsic Mechanisms | Secretly Embedding Model (Li and Zhou, 2024) | Global/Layer | Latent Space | ✗ |
| | MoE-X (Yang et al., 2025b) | Individual | Modular | ✗ |
| | Intrinsic User-Centric (Swamy et al., 2024) | Individual | User-Aligned | ✗ |
| → Focus: What information is encoded within experts? | | | | |
| Criticality Analysis | Unveiling Super Experts (Su et al., 2025) | Individual | Pareto Principle | ✗ |
| → Focus: Which single experts are dominant or critical? | | | | |
| Structural Organization | Standing Committee Analysis (Ours) | Committee | Domain-Invariant | ✓ |
| → Focus: How stable groups of experts dominate computation across domains? | | | | |

natural language relies heavily on high-frequency and domain-agnostic patterns, recent state-of-the-art (SOTA) architectures have moved to institutionalize a centralized processing unit. For instance, DeepSeek (Dai et al., 2024; DeepSeek-AI et al., 2024, 2025) introduces **Shared Experts** that are always activated to isolate common knowledge from the routed experts. The prevailing assumption is that by architecturally separating these generalists, the remaining routed experts are free to become true specialists. Yet, our empirical analysis reveals that this architectural fix does not fully suppress the drive toward centralization. As shown in Figure 1b, we observe substantial cross-domain sharing even among the experts explicitly designated for specialization. Unlike representation collapse, where experts die out due to optimization failure, these shared experts are highly active and functionally competent but simply refuse to specialize. This evidence suggests that the formation of a generalist core is not merely an architectural choice but an inevitable emergent property of sparse computation. However, this emergent structural bias poses a significant challenge to conventional MoE training: standard load-balancing auxiliary losses, which are designed to prevent expert idle-out by encouraging uniform selection, may inadvertently suppress this natural computational hierarchy. This conflict potentially leads to suboptimal convergence and wasted FLOPs on peripheral experts that lack fundamental reasoning capabilities.

Recent advances in MoE interpretability have largely focused on the properties of individual experts. Prior work has examined semantic routing patterns (Olson et al., 2025; Lo et al., 2025; Bai

et al., 2025), analyzed internal representations (Li and Zhou, 2024; Yang et al., 2025b), and most recently identified frequency-based “Super Experts” (Su et al., 2025). However, these studies predominantly treat experts as independent computational units, where importance is quantified through isolated activation statistics. Consequently, they overlook the potential for a higher-level structural organization within the routing mechanism. While the “Super Expert” phenomenon highlights the Pareto distribution of individual criticality, it fails to capture the relational stability between experts across varying contexts. This leaves a critical gap in our understanding: do experts function as isolated specialists whose prominence is a mere statistical byproduct, or do they spontaneously organize into stable, domain-invariant coalitions?

To resolve this discrepancy between domain-specific intuition and observed expert sharing, we propose COMMITTEEAUDIT, a post-hoc analytical framework designed to audit the group-level structural organization of pre-trained MoE models. Unlike prior works that focus on individual expert statistics, our framework quantifies the stability and intersection of expert coalitions across divergent tasks. We apply this auditing process to three representative models (OLMoE (Muennighoff et al., 2025), Qwen3-30B-A3B (Yang et al., 2025a), and DeepSeek-V2-Lite (DeepSeek-AI et al., 2024)), covering both standard routing and architectures with explicit shared experts.

Guided by this structural perspective, our study addresses three fundamental questions regarding the hidden organization of MoE computation: **Existence:** Do routed experts naturally self-organize into stable, domain-invariant groups that dominate computation, or do they remain specialized by task? **Dynamics:** How does this group-level organization evolve across network depths? Is the centralization of experts an inevitable emergent property of sparse routing? **Functionality:** What functional roles do these stable groups play? Specifically, does the model rely on them for general reasoning while relegating specific knowledge to a fringe of other experts?

Our contributions answer these questions and challenge the prevailing view of MoE specialization:

(1) We provide systematic evidence of a domain-invariant **Standing Committee**, a compact expert coalition that emerges regardless of shared-expert architectures, revealing a structural bias that chal-

lenges “fairness-oriented” load balancing.

(2) We introduce a model-agnostic framework, COMMITTEEAUDIT, that utilizes Pareto-optimal ranking and stability diagnostics to quantify group-level expert organization beyond individual activation statistics.

(3) Through qualitative analysis, we uncover a core-periphery organization where committee members anchor logical and syntactic structures, while peripheral experts manage domain-specific knowledge.

2 Related Works

Expert Specialization and Routing Analysis

MoE models are commonly motivated by a divide-and-conquer intuition: sparsity arises when tokens are routed to domain-specialized experts (Xue et al., 2024; Jiang et al., 2024; Zoph et al., 2022; ?; Fan et al., 2024). Early multilingual studies supported this view, reporting experts that preferentially served specific languages (Lepikhin et al., 2020; Zheng et al., 2024). However, recent analyses of general-purpose LLMs reveal a more nuanced picture. Experts often behave polysemously rather than strictly specializing (Lo et al., 2025), and routing only weakly aligns with human semantic domains (Olson et al., 2025). Other work shows that specialization is modulated by context rather than being an intrinsic property of an expert (Bai et al., 2025). In parallel, studies on “super experts” highlight a small set of disproportionately active experts (Su et al., 2025), shifting attention from specialization to expert criticality.

Internal Representations and Intrinsic Interpretability A complementary line of work examines the internal representations of MoE systems. Evidence suggests that experts contribute to a shared latent space rather than operating as isolated modules (Li and Zhou, 2024). To improve interpretability, architectural interventions have been proposed, including constraints that encourage interpretable expert roles (Yang et al., 2025b) and routing mechanisms designed to align usage with higher-level semantic concepts (Swamy et al., 2024).

From Individual Experts to Collective Structure

Despite these advances, most prior studies analyze experts as independent computational units, focusing on activation patterns or internal states (Ghandeharioun et al., 2024). What remains unclear is whether experts organize into stable, co-activated

groups that persist across tasks. Our work addresses this gap by shifting the lens from individual experts to structured collectives, “standing committees”, and shows that such committees emerge in a domain-invariant manner, challenging the conventional assumption of purely domain-specific routing.

3 COMMITTEEAUDIT

3.1 Preliminaries

Mixture-of-Experts Architecture. Mixture-of-Experts (MoE) models extend the Transformer by replacing the feed-forward network with a set of E parallel experts $\{E_i\}_{i=1}^E$ (Shazeer et al., 2017; Lepikhin et al., 2020; Fedus et al., 2022). For a token $x \in \mathbf{R}^d$ at layer ℓ , a gating network produces a routing vector $G^{(\ell)}(x)$. Under Top- k routing, the layer output is the weighted sum of k activated experts:

$$y = \sum_{i \in \text{Top-}k} G^{(\ell)}(x)_i E_i(x). \quad (1)$$

While token-level routing is sparse, aggregating decisions over a corpus reveals structural regularities.

Expert Contribution Index (ECI). To quantify expert importance at the domain task level, we define the *Expert Contribution Index* (ECI). Given a corpus \mathcal{D} partitioned into domain tasks $\mathcal{T} = \{\tau\}$, we denote \mathcal{D}_τ as the subset for a domain task τ . For expert i at layer ℓ , the ECI is the expected routing weight:

$$c_{i,\tau}^{(\ell)} = \mathbf{E}_{x \in \mathcal{D}_\tau} [G^{(\ell)}(x)_i] \quad (2)$$

Unlike activation frequency, ECI preserves the *magnitude* of router preference, providing a more informative signal for ranking. ECI serves as the building block for analyzing cross-task invariants.

3.2 Framework Description

COMMITTEEAUDIT, as shown in Figure 2, is a domain-conditioned routing analysis framework that (i) extracts domain-level routing profiles, (ii) quantifies inter-domain routing divergence, and (iii) explores Standing Committees. In high-capacity MoEs ($E \geq 64$), while single experts may occasionally dominate, activation is generally too distributed for individual-centric analysis. We hypothesize that specialization is expressed through a structured distribution over a subset of experts, referred to as a *committee*.

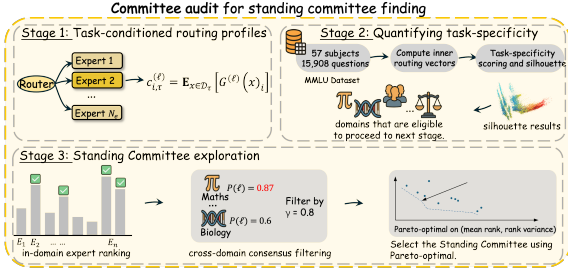


Figure 2: Overview of the COMMITTEEAUDIT framework.

Stage I: Task-conditioned routing profiles. Before constructing committees, we first extract routing representations from the MoE model. For every sample $x \in \mathcal{D}_\tau$ and MoE layer ℓ , we run the model and record the full routing vector $G^{(\ell)}(x)$ taken at the last token unless otherwise specified:

$$G^{(\ell)}(x) = \text{softmax}(z^{(\ell)}(x)) \in \Delta^{E-1}, \quad (3)$$

where Δ^{E-1} denotes the probability simplex over E experts, that is $\Delta^{E-1} = \{p \in \mathbf{R}^E : p_i \geq 0, \sum_i p_i = 1\}$.

We use the full routing distribution (rather than discrete Top- k activations) because it preserves the complete preference structure over experts.

We then aggregate routing behavior at the domain-task level. For each expert i , we compute its ECI $c_{i,\tau}^{(\ell)}$ used Eq (2) and collect all expert contributions into a task-conditioned profile

$$\bar{G}_{\ell,\tau} = [c_{1,\tau}^{(\ell)}, \dots, c_{i,\tau}^{(\ell)}, \dots, c_{E,\tau}^{(\ell)}]^\top, \quad (4)$$

This profile serves as the building block for both task-specificity analysis (Stage II) and expert contribution estimation (Stage III).

Stage II: Quantifying task-specificity. We next assess the degree of routing specialization per domain task τ using a silhouette-based score. Stage II determines whether a task’s routing is sufficiently distinctive. Let $d(\cdot, \cdot)$ denote the cosine distance between routing vectors (as defined in Eq (3)). For each $x_i \in \mathcal{D}_\tau$, define

$$a_i = \frac{1}{|\mathcal{D}_\tau| - 1} \sum_{\substack{x_j \in \mathcal{D}_\tau \\ j \neq i}} d(G^{(\ell)}(x_i), G^{(\ell)}(x_j)),$$

$$b_i = \min_{\tau' \neq \tau} \frac{1}{|\mathcal{D}_{\tau'}|} \sum_{x_j \in \mathcal{D}_{\tau'}} d(G^{(\ell)}(x_i), G^{(\ell)}(x_j)),$$
(5)

and compute the sample-level silhouette

$$s_i = \frac{b_i - a_i}{\max(a_i, b_i)} \in [-1, 1]. \quad (6)$$

The task-specificity score is the mean silhouette:

$$S_\ell(\tau) = \frac{1}{|\mathcal{D}_\tau|} \sum_{x_i \in \mathcal{D}_\tau} s_i. \quad (7)$$

where high $S_\ell(\tau)$ indicates that routing vectors within domain task τ form a compact and clearly separated cluster. Only tasks with sufficiently high $S_\ell(\tau)$ proceed to Standing Committee construction in Stage III.

Stage III: Standing Committee exploration.

We finally identify the *domain-invariant backbone* for well-structured tasks. Given $c_{i,\tau}^{(\ell)}$ computed in Stage I, which corresponds to the average routing weight assigned to expert i when processing samples from domain task τ . Experts are then ranked within each domain task, and we denote the resulting rank by

$$R(i, \tau)^{(\ell)}, \quad (8)$$

assigning a penalty rank $k+1$ to experts that do not appear in the Top- k . Here, k corresponds to the model’s routing sparsity.

To distinguish experts that are globally useful from those that are merely task-specific, we measure how often each expert appears among the Top- k across domain tasks:

$$P_i^{(\ell)} = \frac{1}{|\mathcal{T}|} \sum_{\tau \in \mathcal{T}} \mathbf{I}(R(i, \tau)^{(\ell)} \leq k). \quad (9)$$

Experts that occur sufficiently frequently form the set of consensus candidates \mathcal{E} :

$$\mathcal{E}^{(\ell)} = \left\{ i : P_i^{(\ell)} \geq \gamma \right\}, \quad (10)$$

where γ controls the threshold for cross-domain agreement. To ensure committee members are active in nearly all observed domain tasks, in this paper, we use $\gamma > 0.8$ to set up our experiment.

For each candidate expert, we further characterize stability across domain tasks:

$$\mu_i^{(\ell)} = \mathbf{E}_\tau [R(i, \tau)^{(\ell)}],$$

$$\sigma_i^{(\ell)} = \text{Var}_\tau [R(i, \tau)^{(\ell)}].$$
(11)

Here, $\mu_i^{(\ell)}$ captures how highly the expert is ranked on average, while $\sigma_i^{(\ell)}$ quantifies how consistently it holds that position.

Finally, the *Standing Committee* at layer ℓ is defined as the Pareto-optimal set:

$$\mathcal{C}^{(\ell)} = \text{Pareto} \left(\{(\mu_i^{(\ell)}, \sigma_i^{(\ell)})\}_{i \in \mathcal{E}^{(\ell)}} \right), \quad (12)$$

Table 2: Aggregation of MMLU subjects into nine domain tasks.

| Domain | Representative Subjects |
|------------------|--|
| STEM–Math | algebra, geometry, probability, statistics, college/high-school mathematics |
| STEM–Physics | high-school/college physics, astronomy, conceptual physics |
| STEM–Chemistry | high-school and college chemistry |
| STEM–BioMed | biology, anatomy, genetics, clinical knowledge, virology, nutrition, aging, medicine |
| CS–Eng SocSci | computer science, security, operating systems, ML, electrical engineering economics, econometrics, sociology, psychology, political science, public relations |
| Humanities | history, philosophy, ethics, religion, art history, world facts |
| Lang–Ling | English, literature, linguistics |
| Biz–Law | business, management, accounting, marketing, law |

Table 3: MoE configurations of evaluated models.

| Model | Experts (E) | Top- k | Shared | Size |
|------------------|-----------------|----------|--------|------|
| DeepSeek-V2-Lite | 64 | 6 | 2 | 16B |
| Qwen3-30B-A3B | 128 | 8 | 0 | 30B |
| OLMoE-1B-7B | 64 | 8 | 0 | 7B |

This selection favors experts which achieve a favorable trade-off between high average rank and low cross-domain variability.

3.3 Experimental Setup

3.3.1 Dataset

We evaluate COMMITTEEAUDIT on the Massive Multitask Language Understanding (MMLU) benchmark (Hendrycks et al., 2020), which contains 57 multiple-choice subjects spanning science, humanities, social science, and professional domains. To study domain-conditioned routing rather than per-subject idiosyncrasies, we reorganize all subjects into nine semantically coherent domains (Table 2). To provide a more comprehensive evaluation of our method, we further extend the experimental benchmark to Colossal Clean Crawled Corpus (Raffel et al., 2020), with additional dataset details deferred to Appendix B.

Formally, each subject is mapped to a domain task $\tau \in \mathcal{T}$, yielding domain-specific subsets $\{\mathcal{D}_\tau\}$. All routing analyses in this paper, including task-specificity and Standing Committee extraction, are conducted at the domain level.

3.3.2 Model

We evaluate COMMITTEEAUDIT on three representative MoE language models that differ in expert-pool size and routing configuration. All models are used in inference-only mode, and we extract routing weights from every MoE layer for analysis.

As shown in Table 3, DeepSeek-V2-Lite (DeepSeek-AI et al., 2024) includes two shared experts that are always active, forming a centralized processing path, whereas

Qwen3-30B-A3B (Yang et al., 2025a) and OLMoE-1B-7B (Muennighoff et al., 2025) rely purely on routed experts. The variation in (E, k) and shared-expert usage allows us to probe whether Standing Committees are an architectural artifact or a persistent phenomenon across MoE designs (details are in Appendix A).

3.4 Metrics

3.4.1 Jaccard Similarity (Cross-Domain Expert Sharing)

To quantify how much different domains reuse the same experts, we compute the Jaccard similarity between domain-level Top- k expert sets. For layer ℓ and domains τ_1, τ_2 , let $\mathcal{E}_{\ell, \tau}$ denote the Top- k experts; then

$$\text{Jaccard}_\ell(\tau_1, \tau_2) = \frac{|\mathcal{E}_{\ell, \tau_1} \cap \mathcal{E}_{\ell, \tau_2}|}{|\mathcal{E}_{\ell, \tau_1} \cup \mathcal{E}_{\ell, \tau_2}|}. \quad (13)$$

Values near 0 indicate domain-specific routing, whereas larger values reflect substantial cross-domain expert sharing.

3.4.2 Gini Coefficient (Expert Concentration)

To quantify the inequality of expert contributions, we compute the Gini coefficient over the distribution of the Expert Contribution Index (ECI) at each layer ℓ . Let $\bar{\mathbf{c}}^{(\ell)} = (\bar{c}_1^{(\ell)}, \dots, \bar{c}_E^{(\ell)})$ denote the global contribution vector, where $\bar{c}_i^{(\ell)} = \mathbf{E}_{\tau \in \mathcal{T}}[c_{i, \tau}^{(\ell)}]$ represents the average ECI for expert i across all tasks. The Gini coefficient is defined as:

$$\text{Gini}(\bar{\mathbf{c}}^{(\ell)}) = \frac{\sum_{i=1}^E \sum_{j=1}^E |\bar{c}_i^{(\ell)} - \bar{c}_j^{(\ell)}|}{2E \sum_{i=1}^E \bar{c}_i^{(\ell)}}. \quad (14)$$

In this context, a Gini coefficient approaching 0 indicates a uniform utilization of experts, where each expert provides an equal contribution to the model’s computation. Conversely, a high Gini coefficient (approaching 1) signals extreme contribution inequality, where the total routing mass is monopolized by a small subset of experts, providing macroscopic evidence for the existence of a standing committee.

4 Experiment

In this section, we present a series of experiments to address the three questions introduced in Section 1. These experiments allow us to determine whether standing committees actually emerge, how they evolve across depth, and what functional role they play in model behavior.

Table 4: Cross-domain sharing (Jaccard) and expert concentration (Gini) across models.

| Metric | Statistic | OLMoE | DeepSeek-V2-Lite | Qwen3-30B-A3B |
|--------------------|-----------|--------|------------------|---------------|
| Jaccard Similarity | Max | 1.0000 | 1.0000 | 1.0000 |
| | Min | 0.7963 | 0.7103 | 0.5300 |
| | Overall | 0.8735 | 0.8670 | 0.8670 |
| Gini Coefficient | Max | 0.9082 | 0.9360 | 0.9605 |
| | Min | 0.8814 | 0.9092 | 0.9405 |
| | Overall | 0.8957 | 0.9207 | 0.9465 |

4.1 Existence and Stability: The “Standing Committee” Phenomenon

Question 1: Do routed experts converge into stable, domain-invariant groups?

4.1.1 Jaccard–Gini Analysis of Expert Sharing and Concentration

Table 4 evaluates the standing-committee hypothesis from two complementary perspectives: (i) whether the same experts tend to reappear across domains, and (ii) how unevenly routing mass is distributed among them. Despite substantial differences in expert capacity (E) and routing sparsity (k), all three models display high overlap and concentration, indicating that MoE routing tends to self-organize into “standing committees” rather than task-specific specialization.

Membership Stability (Jaccard Similarity).

The Jaccard index captures whether the same experts repeatedly appear among the top- k routed set across domains. OLMoE ($E = 64, k = 8$) achieves the highest mean overlap (0.8735) and the strongest minimum stability (0.7963), suggesting that the model frequently reuses a common subset of experts. Qwen3 ($E = 128, k = 8$) shows greater local variability (Min: 0.5300), yet its high global average (0.8670) indicates that such deviations occur on top of a largely stable routing structure rather than replacing it entirely.

Contribution Concentration (Gini Coefficient).

The Gini coefficient quantifies the inequality of ECI across the expert population. All models exhibit extreme values (> 0.88), meaning that a small fraction of experts absorbs most of the routing mass. Interestingly, concentration correlates with expert capacity: Qwen3 ($E = 128$) attains the highest overall Gini (0.9465). Rather than distributing computation more broadly, larger pools appear to amplify the dominance of a compact set of frequently selected experts.

Figure 3 links these statistics to routing behavior. In Panel (a), the mean normalized weight assigned

Table 5: An audit of representative Standing Committees (\mathcal{C}) across network phases, showing shallow, middle, and deep layers for each model. **Avg. μ** and **Avg. σ^2** represent the mean and the variance of ranks of the committee members across domains. **ECI Cov.** is the cumulative contribution, and **Ratio** indicates the influence density vs. a uniform baseline.

| Model | Phase | Layer (l) | Committee Members (\mathcal{C}) | $ \mathcal{C} $ | Avg. $\mu \downarrow$ | Avg. $\sigma^2 \downarrow$ | ECI Cov. | Ratio (\times) |
|------------------|---------|---------------|-------------------------------------|-----------------|-----------------------|----------------------------|----------|--------------------|
| DeepSeek-V2-Lite | Shallow | 3 | {7, 13, 22, 42} | 4 | 3.36 | 1.81 | 66.3% | 29.5 \times |
| | Middle | 11 | {28, 43, 54} | 3 | 3.15 | 1.98 | 60.7% | 31.4 \times |
| | Deep | 19 | {4, 14, 47, 61} | 4 | 3.11 | 0.76 | 70.5% | 35.8 \times |
| OLMoE | Shallow | 2 | {30, 58, 63} | 3 | 3.41 | 2.15 | 43.9% | 15.9 \times |
| | Middle | 8 | {13, 45} | 2 | 3.28 | 0.49 | 29.7% | 13.1 \times |
| | Deep | 16 | {17, 52, 60} | 3 | 3.19 | 1.52 | 44.0% | 16.0 \times |
| Qwen3-30B-A3B | Shallow | 3 | {38, 40, 80, 93} | 4 | 3.61 | 3.44 | 54.0% | 36.4 \times |
| | Middle | 33 | {16, 26, 57, 116, 121} | 5 | 3.82 | 2.16 | 67.0% | 49.9 \times |
| | Deep | 46 | {94, 101, 107} | 3 | 3.15 | 1.59 | 50.9% | 43.3 \times |


to the routed top- k experts remains both high and stable across layers. If experts were mainly task-specialized, different domains would activate different experts, and the variance bands would widen. Instead, we observe persistent dominance by the same routed subset, indicating a domain-invariant backbone. Panels (b–d) arrive at the same conclusion from a distributional view: Lorenz curves show that only a tiny fraction of experts accounts for most ECI, confirming a strongly centralized allocation of computation.

Together, the two views support the Standing Committee hypothesis: MoE models concentrate routing mass onto a small, persistent core, while most experts operate only peripherally.


4.1.2 Analyzing the Stability and Contribution of Standing Committees

Table 5 summarizes representative standing committees based on the Pareto-optimal set across depth. Committee members consistently occupy very high routing positions (Avg. $\mu \approx 3.1$ –3.8) with low rank variability ($\sigma^2 \leq 3.44$). For example, the middle-layer committee in OLMoE exhibits a variance of only 0.49, indicating that these experts remain near the top of the routing hierarchy regardless of domain. Rather than transient specialists, they function as a de-facto backbone that the model repeatedly relies on. For additional transparency, Appendix C provides an illustrative expert-level table showing domain coverage and ECI statistics for representative DeepSeek-V2-Lite layers. Complete layer-wise MMLU results are provided in Appendix D, Section D.1, while matched supplementary results on the C4-based evaluation subset are reported in Appendix D, Section D.2.


Although $|\mathcal{C}|$ remains small (2-5) across all models, these groups capture up to 70.5% of total routing mass. Importantly, the size of \mathcal{C} remains stable even as capacity increases from $E = 64$ to




(a) Standing Committee Stability Across Layers.



(b) OLMoE expert concentration.



(c) DeepSeek expert concentration.



(d) QWen expert concentration.

Figure 3: Evidence of standing committees in MoE models. (a) Layer-wise concentration of top- k experts across tasks. For each model, the solid line shows the mean normalized weight assigned to the top- k experts at each layer, and the shaded region denotes one standard deviation. High and stable values indicate a small subset of experts (“standing committees”) consistently absorbs most routing mass. (b–d) Lorenz curves reveal that only a small subset of experts accounts for most contributions (other Lorenz curves are shown in Appendix E), showing that these committees are highly centralized rather than uniformly shared.

$E = 128$. This suggests that committee-like behavior is not an artifact of a particular architecture, but an emergent pattern of sparse routing optimization.

4.2 Variation and Sensitivity: Structural Dynamics

Question 2: How does group structure evolve with depth, and is centralization inevitable under sparse routing?

4.2.1 Robustness of Standing Committees

We begin by examining whether routed experts actually form persistent groups. Figure 4 reports the Jaccard similarity of top- k expert sets across domains for each layer. Across all three MoE models, the similarity remains consistently high (often ≥ 0.85), showing that the same experts are repeatedly activated across tasks and depths. Rather than rotating specialists, the routing network converges to a shared backbone of experts that is largely invariant to both input domain and layer position.

Having established the existence of a persistent backbone, we next ask how it evolves with depth






Figure 4: Cross-layer stability of routed experts across models, measured by Jaccard similarity between top- k expert sets over domains. All three MoE models maintain high overlap (≥ 0.8 for most layers), showing that the same experts are repeatedly selected despite changes in input domain and network depth.

and routing sparsity. Figure 5 analyzes OLMoE under different routing budgets k .

In Figure 5(a), the Influence Density Ratio remains consistently above 1 across layers and routing budgets. This shows that, on a per-expert basis, committee members absorb substantially more routing mass than non-committee experts. Importantly, increasing k does not flatten this imbalance. Instead, additional routing capacity mainly intro-



(a) Influence Density Ratio across layers and routing budgets k .



(b) Committee size $|\mathcal{C}|$ across layers and routing budgets k .

Figure 5: Dynamics of Standing Committees in OLMoE under different routing budgets. We show all 16 layers. (a) The Influence Density Ratio measures the average per-expert contribution of committee members relative to that of non-committee experts. Values greater than 1 indicate that committee members absorb disproportionately more routing mass per expert. Increasing k does not eliminate this concentration. (b) The size of the identified committees stays small and changes only mildly with depth and k , indicating a compact but persistent core of experts.

duces peripheral experts with relatively low contribution density, while the committee continues to dominate the allocation of routing mass.

Figure 5(b) shows that the committee size itself remains small, typically around 1-4 experts, and changes only mildly across depth and routing budgets. Thus, MoE models do not substantially expand the committee as depth increases or routing capacity grows; they instead rely on a compact, persistent core whose influence remains structurally robust across different values of k . Together with Figure 4, these findings indicate that centralization is not accidental but an emergent structural property of sparse routing.

4.2.2 Top- k Sensitivity Sweep

To probe how sensitive the standing committee is to the routing budget, we perform a top- k sensitivity sweep on OLMoE. For each setting of $k \in \{4, 6, 8, 12, 16\}$, we re-identify the standing committee based on the Pareto-optimal set and

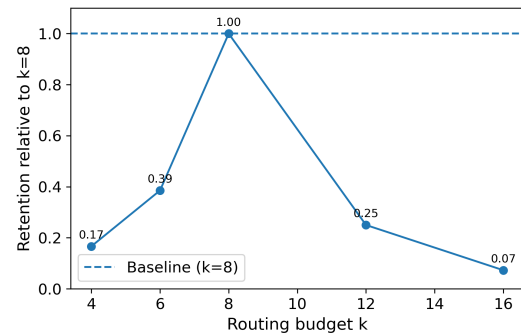


Figure 6: Top- k sensitivity sweep for OLMoE. We take the standing committee identified at $k = 8$ as the reference core and measure, for each routing budget k , the fraction of its members that are still present in the new committee (averaged over layers). The dashed line marks the $k = 8$ baseline (retention = 1.0).

compute the retention rate with respect to the reference committee at $k=8$, i.e., the fraction of $k=8$ core members that remain in the new committee.

Figure 6 shows that retention peaks at the nominal configuration $k=8$ and drops on both sides: it falls to 0.39 at $k=6$, 0.17 at $k=4$, and below 0.3 once k is expanded to 12 or 16. This pattern suggests that the core experts are not an artifact of a single k choice, but they are also not completely rigid. When k is too small, the gate is forced to exclude part of the original core; when k is too large, the gate dilutes its attention and recruits additional experts, replacing some core members.


Overall, the sweep indicates that sparse routing induces a centralized committee around $k=8$, but that committee can partially reorganize as the routing budget becomes substantially more aggressive or more constrained.

4.3 Interaction and Behavior: Functional Interpretation


Question 3: What roles do stable expert groups play in reasoning versus domain knowledge?

We illustrate the functional roles of Standing Committees using a qualitative case study. Figure 7 shows an activation matrix between committee experts (columns) and functional tokens (rows). A cell is marked when a token repeatedly activates an expert in at least three domains. Two consistent behaviors emerge.

Anchor 1: Logical framing and reasoning control. Across OLMoE and Qwen3-30B-A3B, abstract reasoning triggers such as *Which*, *What*, *Suppose*, and question marks are routed to the same subset of committee experts. These tokens define



(a) OLMoE: functional tokens repeatedly map to a fixed subset of committee experts.



(b) Qwen3-30B-A3B shows similar convergence despite larger expert capacity.

Figure 7: Case study of token-level routing. Rows denote functional tokens and columns denote members of the Standing Committee. A cell is marked when a token reliably activates an expert across domains.

the logical scaffolding of the prompt, suggesting that the committee acts as a reasoning controller.

Anchor 2: Domain-invariant syntactic backbone. High-frequency structural tokens, including *the*, *a*, and *in*, also converge to overlapping committee members across domains, indicating that the committee maintains a stable syntactic layer independent of content.

Peripheral specialization. By contrast, domain-specific terminology rarely stabilizes: chemical symbols, biomedical identifiers, and financial jargon are distributed across many experts depending on context. This pattern supports a core-periphery organization in which the committee anchors reasoning and syntax, while peripheral experts are recruited on demand for specialized knowledge.

Taken together, Standing Committees function as a domain-invariant control layer, coordinating logical structure and grammar while delegating domain knowledge to peripheral experts.

4.4 Masking-Based Evidence of Functional Importance

To complement our observational analysis, we conduct a masking-based intervention on DeepSeek-

Table 6: Representative masking-based intervention results on DeepSeek-V2-Lite. For each selected layer, we suppress the routing weights of the identified Standing Committee experts, re-normalize the remaining routing distribution, and evaluate answer outcomes on MMLU.

| Phase | Mask Layer | Correct | Wrong | No Answer |
|----------|------------|---------|-------|-----------|
| Baseline | None | 0.39 | 0.58 | 0.03 |
| Shallow | 2 | 0.12 | 0.52 | 0.36 |
| Middle | 10 | 0.09 | 0.55 | 0.36 |
| Deep | 26 | 0.03 | 0.59 | 0.38 |

V2-Lite to test whether identified Standing Committee experts are functionally important for answer generation. For each selected layer, we suppress the routing weights of the Standing Committee experts, re-normalize the remaining routing distribution, and then evaluate answer outcomes on MMLU. Table 6 reports representative shallow, middle, and deep layers, while the full results are provided in Appendix F.

Masking committee experts consistently degrades performance. The unmasked baseline achieves a correct rate of 0.39, which drops to 0.12, 0.09, and 0.03 when masking representative shallow, middle, and deep committee layers, respectively. At the same time, the proportion of *No Answer* responses rises sharply from 0.03 to 0.36–0.38. These results provide intervention-based evidence that Standing Committee experts are not merely statistically frequent routing destinations, but functionally important components of the model’s computational backbone.

5 Conclusion

This work introduces COMMITTEEAUDIT and shows that MoE models rely on a domain-invariant **Standing Committee** that anchors reasoning and syntax, while peripheral experts handle domain knowledge. Our cross-model analysis suggests that this centralized structure arises from sparse routing itself rather than from specific architectural choices. This structural bias also reveals a tension with load-balancing objectives that encourage uniform expert usage, motivating function-aware routing and architectures that explicitly support a core-periphery organization of expertise.

Limitations

This work introduces COMMITTEEAUDIT and reports evidence for a domain-invariant Standing Committee in Mixture-of-Experts models. However, several limitations remain.

First, our analysis covers only a small set of representative MoE architectures and settings, and does not span hybrid, hierarchical, or dynamically adaptive routing designs. Whether similar organizational patterns persist in broader systems remains an open question.

Second, although we supplement our observational analysis with masking-based interventions on representative layers, our study still falls short of a fully controlled causal analysis. In particular, we do not systematically compare committee masking against alternative controls such as random or frequency-matched non-committee masking across models and tasks.

Third, our evaluation is centered primarily on domain-level analyses over MMLU, with supplementary evidence from C4 and intervention-based experiments. While this reduces subject-level noise and broadens the empirical picture, it may not fully capture behavior in conversational, multi-step reasoning, coding, or tool-augmented scenarios.

Finally, COMMITTEEAUDIT focuses on routing statistics rather than training dynamics. Understanding when and how Standing Committees emerge during optimization remains an important direction for future work.

Potential Risks

Potential Positive Impacts. By revealing group-level routing structure, this work may inform more transparent and efficient MoE design, support diagnosis of routing failures, and encourage principled interpretability research for sparse models.

Potential Negative Impacts. However, several risks remain. First, insights into centralized computation could be misinterpreted as evidence of inherent safety or robustness, leading to overreliance in deployment. Second, identifying persistent expert coalitions may enable adversarial targeting of critical routing pathways. Third, benchmarks may become over-optimized toward interpretability metrics without improving real-world safety.

These findings should therefore be used as analytical tools rather than deployment guarantees.

Ethical Considerations and Usage Disclaimer

All experiments use publicly available models and datasets without personal or sensitive information. This work is intended for academic and educational

purposes only and does not constitute guidance for production deployment.

The framework exposes structural properties of MoE systems, but does not certify fairness, safety, robustness, or regulatory compliance. The authors make no warranties regarding completeness or suitability for downstream use, and any application to high-stakes settings should involve domain experts, risk assessment, and human oversight.

We acknowledge that AI-assisted tools were used during writing and editing (e.g., grammar checking, phrasing refinement, and formatting suggestions). These tools were not used to generate research ideas, experimental results, model outputs, or claims, and all technical content, analyses, and conclusions were designed, verified, and interpreted by the authors.

Licenses and Terms of Use

All datasets and pre-trained models used in this work are publicly available and redistributed under their respective licenses. We respect the original terms of use for each artifact. MMLU is used under its public research license, and all evaluated MoE models (OLMoE, Qwen3-30B-A3B, and DeepSeek-V2-Lite) are accessed and used in compliance with their published licenses. We do not redistribute any third-party artifacts.

Acknowledgments

The authors acknowledge *The Fin AI* community for its research support, feedback, and collaborative environment that contributed to this work.

References

- Jun Bai, Minghao Tong, Yang Liu, Zixia Jia, and Zilong Zheng. 2025. Understanding and leveraging the expert specialization of context faithfulness in mixture-of-experts llms. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 21938–21953.
- Zewen Chi, Li Dong, Shaohan Huang, Damai Dai, Shuming Ma, Barun Patra, Saksham Singhal, Payal Bajaj, Xia Song, Xian-Ling Mao, and 1 others. 2022. On the representation collapse of sparse mixture of experts. *Advances in Neural Information Processing Systems*, 35:34600–34613.
- Damai Dai, Chengqi Deng, Chenggang Zhao, RX Xu, Huazuo Gao, Deli Chen, Jiashi Li, Wangding Zeng, Xingkai Yu, Yu Wu, and 1 others. 2024. Deepseek-moe: Towards ultimate expert specialization in

- mixture-of-experts language models. *arXiv preprint arXiv:2401.06066*.
- DeepSeek-AI, Aixin Liu, Bei Feng, Bin Wang, Bingxuan Wang, Bo Liu, Chenggang Zhao, Chengqi Deng, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fuli Luo, Guangbo Hao, Guanting Chen, and 138 others. 2024. [Deepseek-v2: A strong, economical, and efficient mixture-of-experts language model](#). *Preprint*, arXiv:2405.04434.
- DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, and 181 others. 2025. [Deepseek-v3 technical report](#). *Preprint*, arXiv:2412.19437.
- Giang Do, Hung Le, and Truyen Tran. 2025. Simsmoe: Toward efficient training mixture of experts via solving representational collapse. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 2012–2025.
- Dongyang Fan, Bettina Messmer, and Martin Jaggi. 2024. [Towards an empirical understanding of moe design choices](#). *Preprint*, arXiv:2402.13089.
- William Fedus, Barret Zoph, and Noam Shazeer. 2022. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120):1–39.
- Asma Ghandeharioun, Avi Caciularu, Adam Pearce, Lucas Dixon, and Mor Geva. 2024. Patchscopes: A unifying framework for inspecting hidden representations of language models. *arXiv preprint arXiv:2401.06102*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*.
- Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, and 1 others. 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- Dmitry Lepikhin, HyoukJoong Lee, Yanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam Shazeer, and Zhifeng Chen. 2020. Gshard: Scaling giant models with conditional computation and automatic sharding. *arXiv preprint arXiv:2006.16668*.
- Ziyue Li and Tianyi Zhou. 2024. Your mixture-of-experts llm is secretly an embedding model for free. *arXiv preprint arXiv:2410.10814*.
- Ka Man Lo, Zeyu Huang, Zihan Qiu, Zili Wang, and Jie Fu. 2025. A closer look into mixture-of-experts in large language models. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 4427–4447.
- Niklas Muennighoff, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Jacob Morrison, Sewon Min, Weijia Shi, Pete Walsh, Oyvind Tafjord, Nathan Lambert, Yuling Gu, Shane Arora, Akshita Bhagia, Dustin Schwenk, David Wadden, Alexander Wettig, Binyuan Hui, Tim Dettmers, Douwe Kiela, and 5 others. 2025. [Olmoe: Open mixture-of-experts language models](#). *Preprint*, arXiv:2409.02060.
- Matthew Lyle Olson, Neale Ratzlaff, Musashi Hinck, Man Luo, Sungduk Yu, Chendi Xue, and Vasudev Lal. 2025. Probing semantic routing in large mixture-of-expert models. *arXiv preprint arXiv:2502.10928*.
- Lingfei Qian, Weipeng Zhou, Yan Wang, Xueqing Peng, Han Yi, Yilun Zhao, Jimin Huang, Qianqian Xie, and Jian yun Nie. 2025. [Fino1: On the transferability of reasoning-enhanced llms and reinforcement learning to finance](#). *Preprint*, arXiv:2502.08127.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67.
- Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. 2017. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*.
- Zunhai Su, Qingyuan Li, Hao Zhang, Weihao Ye, Qibo Xue, YuLei Qian, Yuchen Xie, Ngai Wong, and Kehong Yuan. 2025. [Unveiling super experts in mixture-of-experts large language models](#). *Preprint*, arXiv:2507.23279.
- Vinitra Swamy, Syrielle Montariol, Julian Blackwell, Jibril Frej, Martin Jaggi, and Tanja Käser. 2024. Intrinsic user-centric interpretability through global mixture of experts. *arXiv preprint arXiv:2402.02933*.
- Yan Wang, Yueru He, Ruoyu Xiang, and Jeff Zhao. 2025a. Rkefino1: A regulation knowledge-enhanced large language model. In *2025 IEEE 11th International Conference on Intelligent Data and Security (IDS)*, pages 49–51. IEEE.
- Yan Wang, Yang Ren, Lingfei Qian, Xueqing Peng, Keyi Wang, Yi Han, Dongji Feng, Fengran Mo, Shengyuan Lin, Qinchuan Zhang, Kaiwen He, Chenri Luo, Jianxing Chen, Junwei Wu, Jimin Huang, Guojun Xiong, Xiao-Yang Liu, Qianqian Xie, and Jian-Yun Nie. 2025b. [Fintagging: Benchmarking llms for extracting and structuring financial information](#). *Preprint*, arXiv:2505.20650.

Yan Wang, Keyi Wang, Shanshan Yang, Jaisal Patel, Jeff Zhao, Fengran Mo, Xueqing Peng, Lingfei Qian, Jimin Huang, Guojun Xiong, Xiao-Yang Liu, and Jian-Yun Nie. 2025c. [Finauditing: A financial taxonomy-structured multi-document benchmark for evaluating llms](#). *Preprint*, arXiv:2510.08886.

Fuzhao Xue, Zian Zheng, Yao Fu, Jinjie Ni, Zangwei Zheng, Wangchunshu Zhou, and Yang You. 2024. [Openmoe: An early effort on open mixture-of-experts language models](#). *arXiv preprint arXiv:2402.01739*.

An Yang, Anpeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025a. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.

Xingyi Yang, Constantin Venhoff, Ashkan Khakzar, Christian Schroeder de Witt, Puneet K Dokania, Adel Bibi, and Philip Torr. 2025b. [Mixture of experts made intrinsically interpretable](#). *arXiv preprint arXiv:2503.07639*.

Guorui Zheng, Xidong Wang, Juhao Liang, Nuo Chen, Yuping Zheng, and Benyou Wang. 2024. [Efficiently democratizing medical llms for 50 languages via a mixture of language family experts](#). *arXiv preprint arXiv:2410.10626*.

Barret Zoph, Irwan Bello, Sameer Kumar, Nan Du, Yanping Huang, Jeff Dean, Noam Shazeer, and William Fedus. 2022. [St-moe: Designing stable and transferable sparse expert models](#). *arXiv preprint arXiv:2202.08906*.

A Computational Budget and Infrastructure

All experiments are implemented in Python using PyTorch and the HuggingFace Transformers library, and are conducted in inference-only mode. We evaluate each model on the full MMLU benchmark using two NVIDIA A100 (80 GB) GPUs. OLMoE-1B-7B, DeepSeek-V2-Lite, and Qwen3-30B-A3B contain approximately 7B, 16B, and 30B parameters, respectively. The total computational cost of the routing analyses is about 40 GPU-hours, including forward passes and the collection of routing statistics.

B C4 Dataset

To complement the domain-level analysis on MMLU, we additionally construct a structured evaluation subset from C4 (Colossal Clean Crawled Corpus) (Raffel et al., 2020) for multilingual routing analysis. While MMLU provides subject annotations that can be aggregated into higher-level

| Domain Group | Representative Languages | # Samples | Avg. Seq. Len. |
|----------------|-------------------------------------|-----------|----------------|
| European | German, Spanish, French, Portuguese | 400 | 402.7 |
| Asian | Chinese, Hindi, Japanese | 300 | 402.9 |
| Middle Eastern | Arabic | 100 | 467.9 |
| Other | English, Russian | 200 | 348.8 |

Table 7: Aggregation of the C4-based evaluation subset into four language-oriented domain groups. The subset contains 1,000 samples in total, drawn from 10 language-specific sources with 100 samples per language.

semantic domains, C4 does not contain analogous task labels. We therefore organize the selected C4 subsets by language-oriented groups, yielding four higher-level domains: **European** (de, es, fr, pt), **Asian** (zh, hi, ja), **Middle Eastern** (ar), and **Other** (en, ru).

This C4-based evaluation set contains **1,000 samples** in total, drawn from **10 language-specific sources**, with **100 samples per language**. The resulting domain distribution is moderately imbalanced by construction, including 400 European samples, 300 Asian samples, 100 Middle Eastern samples, and 200 Other samples. This aggregation mirrors the design principle of Table 2: rather than analyzing routing behavior at the level of individual sources, we group related sources into broader domains in order to examine domain-conditioned routing patterns at a more stable level of abstraction.

We also observe non-trivial variation in input length across these groups. The average sequence length ranges from 348.8 in the Other group to 467.9 in the Middle Eastern group, while all groups contain samples reaching the truncation limit of 512 tokens. These differences make the C4 subset a useful complement to MMLU, allowing us to test whether Standing Committee behavior persists beyond supervised benchmark subjects and extends to a broader multilingual web corpus.

C Illustrative Expert-Level Statistics for Committee Selection

To make the committee selection process more transparent, we provide an illustrative expert-level view for representative DeepSeek-V2-Lite layers, as shown in Table 8. For each activated expert, we report its domain coverage, mean ECI, standard deviation of ECI, and an auxiliary stability score defined as the product of domain coverage and mean ECI. This score is introduced only for interpretive presentation, to provide a more tangible view of how highly covered and high-contribution experts

differ from weaker candidates. It is **not** part of the formal committee selection criterion. The formal method remains exactly as defined in Section 3.2: experts must first satisfy the cross-domain agreement threshold $\gamma > 0.8$, and committee members are then determined based on the Pareto-optimal set of rank statistics.

The table also clarifies an important point: domain coverage alone is not sufficient for committee membership. Some experts are activated in all or nearly all domains, but their mean ECI remains substantially lower than that of the final committee members. As a result, they do not belong to the committee, even though they are broadly active. This provides a more direct illustration that Standing Committees are not defined by activation frequency alone, but by stable and substantial contribution across domains.

D Standing Committees across all layers

D.1 MMLU Dataset Results

Full-layer analysis of standing committees. To move beyond representative snapshots, we perform a layer-by-layer audit of all identified standing committees in OLMoE, DeepSeek-V2-Lite, and Qwen3-30B-A3B, as shown in Table 9, 10, and 11. The resulting tables reveal a consistent organizational pattern that is obscured when only a few layers are examined. Across models, standing committees emerge early, consolidate in the middle layers, and persist into the deepest layers, while their composition changes only gradually.

Small, persistent coalitions. Despite large expert pools, the size of each committee remains compact: typically $|C| \in [1, 4]$ for OLMoE and DeepSeek-V2-Lite, and occasionally up to five members in Qwen. Increasing expert capacity does not diversify routing. Instead, optimization repeatedly converges onto a small coalition of experts that are selected across domains and prompts. This suggests that sparse routing does not primarily allocate experts by domain; rather, it reinforces a stable computational core.

Early centralization. A striking finding is that centralization appears already in shallow layers. Even in the first few layers, standing committees capture a non-trivial proportion of routing mass (often 20–40%). This indicates that MoE models commit to shared processing pathways almost immediately, likely encoding high-frequency patterns

such as token normalization, shallow syntactic cues, and generic lexical regularities. Contrary to the intuition that specialization gradually emerges with depth, the router begins consolidating computation from the outset.

Middle-layer consolidation. The middle layers display the clearest standing-committee behavior. Committees often grow slightly larger while their rank variance decreases, and their cumulative contribution increases sharply (frequently exceeding 50–65% in DeepSeek-V2-Lite and Qwen). These layers appear to implement domain-agnostic abstractions, reasoning templates, discourse structure, and general semantic scaffolding, that are shared across inputs. The router does not allocate different domains to distinct experts; instead, it repeatedly routes through the same committee.

Deep-layer bottlenecks. In deep layers, both DeepSeek-V2-Lite and Qwen exhibit strong bottleneck effects: a small committee controls 50–70% of routing mass, often with high influence density. Rather than distributing final computation across diverse experts, the network funnels decision-making through a narrow coalition. This pattern challenges the traditional “divide-and-conquer” view of MoE systems, suggesting that final reasoning is centralized rather than decomposed.

Architectural variability, consistent behavior. Although OLMoE shows weaker committees than Qwen and DeepSeek-V2-Lite, the qualitative trend is remarkably stable across architectures. Even with different routing designs and training recipes, all three models converge toward small, domain-invariant committees that repeatedly dominate computation. Taken together, these findings indicate that standing committees are not an artifact of any particular implementation. Instead, they appear to be an emergent consequence of sparse routing optimization, reflecting a strong inductive bias toward centralization in modern MoE language models.

D.2 C4 Dataset Results

Matched layer-wise analysis on free-form web text. Based on the C4-based domain construction described in Appendix B, we further conduct a matched supplementary analysis to test whether the Standing Committee phenomenon depends on the multiple-choice structure of MMLU. Specifically, we apply the same committee-identification procedure to the C4-based evaluation subset using

Table 8: Illustrative expert-level statistics for committee selection in representative DeepSeek-V2-Lite layers. Domain Coverage is measured out of the nine aggregated MMLU domains. The Stability Score is an auxiliary presentation-only quantity defined as Domain Coverage \times Mean ECI, and is included only to improve interpretability; it is not part of the formal committee selection rule.

| Layer | Phase | Activated Expert | Domain Coverage | Mean ECI | Std ECI | Stability Score | Committee |
|-------|---------|------------------|-----------------|----------|---------|-----------------|-----------|
| 3 | Shallow | 7 | 9 | 0.8832 | 0.03 | 7.9488 | Y |
| 3 | Shallow | 13 | 9 | 0.8253 | 0.02 | 7.4277 | Y |
| 3 | Shallow | 22 | 9 | 0.8043 | 0.03 | 7.2387 | Y |
| 3 | Shallow | 42 | 9 | 0.7825 | 0.02 | 7.0425 | Y |
| 3 | Shallow | 1 | 9 | 0.5310 | 0.01 | 4.7790 | N |
| 3 | Shallow | 9 | 9 | 0.4928 | 0.02 | 4.4352 | N |
| 11 | Middle | 28 | 9 | 0.9259 | 0.03 | 8.3331 | Y |
| 11 | Middle | 43 | 9 | 0.9215 | 0.01 | 8.2935 | Y |
| 11 | Middle | 54 | 9 | 0.9005 | 0.02 | 8.1045 | Y |
| 11 | Middle | 47 | 7 | 0.8642 | 0.05 | 6.0494 | N |
| 11 | Middle | 16 | 6 | 0.6329 | 0.09 | 3.7974 | N |
| 11 | Middle | 62 | 5 | 0.5903 | 0.04 | 2.9515 | N |
| 11 | Middle | 35 | 3 | 0.3178 | 0.04 | 0.9534 | N |
| 11 | Middle | 8 | 3 | 0.2721 | 0.02 | 0.8163 | N |
| 11 | Middle | 1 | 1 | 0.1919 | 0.00 | 0.1919 | N |
| 19 | Deep | 4 | 9 | 0.9386 | 0.03 | 8.4474 | Y |
| 19 | Deep | 14 | 9 | 0.9298 | 0.03 | 8.3682 | Y |
| 19 | Deep | 17 | 9 | 0.9055 | 0.04 | 8.1495 | Y |
| 19 | Deep | 61 | 9 | 0.9013 | 0.02 | 8.1117 | Y |
| 19 | Deep | 53 | 9 | 0.7524 | 0.02 | 6.7716 | N |
| 19 | Deep | 12 | 7 | 0.3111 | 0.04 | 2.1777 | N |
| 19 | Deep | 32 | 3 | 0.1767 | 0.04 | 0.5301 | N |

Table 9: Comprehensive audit of Standing Committees for OLMoE.

| Layer (ℓ) | Committee Members (C) | C | Avg. μ \downarrow | Avg. σ^2 \downarrow | ECI Cov. | Ratio (\times) |
|------------------|-----------------------|---|-------------------------|------------------------------|----------|--------------------|
| 1 | 49 | 1 | 1.00 | 0.00 | 21.5% | 17.25 |
| 2 | 58, 63, 30 | 3 | 3.41 | 2.15 | 43.9% | 15.94 |
| 3 | 60, 14 | 2 | 3.17 | 1.17 | 30.9% | 13.84 |
| 4 | 9, 56 | 2 | 1.67 | 0.33 | 34.4% | 16.24 |
| 5 | 27 | 1 | 3.11 | 0.99 | 14.5% | 10.68 |
| 6 | 53 | 1 | 1.78 | 1.06 | 16.3% | 12.27 |
| 7 | 1 | 1 | 1.89 | 0.54 | 16.4% | 12.38 |
| 8 | 13, 45 | 2 | 3.28 | 0.49 | 29.7% | 13.10 |
| 9 | 27 | 1 | 1.33 | 0.22 | 17.3% | 13.15 |
| 10 | 8, 12 | 2 | 3.78 | 0.49 | 28.3% | 12.27 |
| 11 | 30, 6 | 2 | 4.28 | 1.05 | 26.8% | 11.35 |
| 12 | 33 | 1 | 2.00 | 2.00 | 17.1% | 12.96 |
| 13 | 30, 55 | 2 | 2.50 | 1.95 | 30.6% | 13.69 |
| 14 | 46, 4 | 2 | 2.22 | 0.73 | 31.3% | 14.14 |
| 15 | 1 | 1 | 1.78 | 0.62 | 16.9% | 12.78 |
| 16 | 60, 52, 17 | 3 | 3.19 | 1.52 | 44.0% | 15.99 |

the identical DeepSeek-V2-Lite configuration as in the main experiments. We keep the routing sparsity fixed at $k = 6$ and exclude shared experts from committee construction, so that the analysis isolates the effect of input distribution while holding the routing protocol constant. Due to space constraints, we report representative shallow, middle, and deep layers rather than the full layer-wise C4 results.

Compact committees persist beyond MMLU. The resulting pattern remains consistent with our main findings on MMLU. Across the reported representative layers, the identified committees remain compact, typically containing only 1–4 experts. This indicates that the emergence of a small dominant expert coalition is not specific to subject-

Table 10: Comprehensive audit of Standing Committees for DeepSeek-V2-Lite.

| Layer (ℓ) | Committee Members (C) | C | Avg. μ \downarrow | Avg. σ^2 \downarrow | ECI Cov. | Ratio (\times) |
|------------------|-----------------------|---|-------------------------|------------------------------|----------|--------------------|
| 1 | 25, 57 | 2 | 3.33 | 0.58 | 34.8% | 16.58 |
| 2 | 19, 51, 46 | 3 | 3.00 | 1.98 | 52.0% | 22.02 |
| 3 | 42, 7, 13, 22 | 4 | 3.36 | 1.81 | 66.3% | 29.46 |
| 4 | 25, 59, 13 | 3 | 2.44 | 1.72 | 57.3% | 27.25 |
| 5 | 38 | 1 | 1.22 | 0.17 | 22.0% | 17.74 |
| 6 | 35, 46 | 2 | 3.83 | 1.00 | 31.5% | 14.25 |
| 7 | 50, 17 | 2 | 3.33 | 0.33 | 34.5% | 16.36 |
| 8 | 45, 46 | 2 | 2.33 | 1.51 | 37.9% | 18.88 |
| 9 | 38, 46, 41 | 3 | 2.70 | 1.43 | 56.7% | 26.66 |
| 10 | 60 | 1 | 1.22 | 0.40 | 23.6% | 19.41 |
| 11 | 54, 43, 28 | 3 | 2.59 | 0.35 | 60.7% | 31.36 |
| 12 | 30 | 1 | 1.33 | 0.44 | 21.9% | 17.63 |
| 13 | 29, 6 | 2 | 2.39 | 0.46 | 40.6% | 21.17 |
| 14 | 5, 28, 33 | 3 | 2.37 | 1.43 | 58.7% | 28.93 |
| 15 | 8, 6, 0 | 3 | 3.04 | 1.22 | 56.1% | 26.00 |
| 16 | 24 | 1 | 1.33 | 0.22 | 20.4% | 16.18 |
| 17 | 40, 25, 31 | 3 | 3.89 | 0.91 | 44.3% | 16.18 |
| 18 | 51, 46, 53, 0 | 4 | 3.67 | 0.65 | 65.5% | 28.43 |
| 19 | 61, 14, 47, 4 | 4 | 3.11 | 0.76 | 70.5% | 35.78 |
| 20 | 44, 7 | 2 | 3.33 | 0.47 | 34.3% | 16.17 |
| 21 | 48 | 1 | 1.67 | 0.67 | 22.1% | 17.85 |
| 22 | 40, 21 | 2 | 3.33 | 1.32 | 35.4% | 17.00 |
| 23 | 23, 6, 38 | 3 | 3.30 | 0.58 | 52.1% | 22.14 |
| 24 | 60, 61 | 2 | 2.00 | 1.06 | 42.7% | 23.08 |
| 25 | 44, 1 | 2 | 1.83 | 0.32 | 46.3% | 26.74 |
| 26 | 36, 56 | 2 | 3.17 | 0.56 | 34.6% | 16.41 |

structured benchmark inputs, but also appears under free-form web-text inputs drawn from C4.

Substantial routing coverage under free-form inputs. Although the committees are small, they still capture a substantial fraction of routing mass. In the reported layers, committee coverage ranges from 12.43% to 66.32%, with several middle and deep layers exceeding 30% and the strongest layers surpassing 50%. In particular, Layer 0 reaches 66.32% coverage, while Layer 24 reaches 62.81%, showing that centralized routing persists even when

Table 11: Comprehensive audit of Standing Committees for Qwen3-30B-A3B.

| Layer (ℓ) | Committee Members (C) | $ C $ | Avg. $\mu \downarrow$ | Avg. $\sigma^2 \downarrow$ | ECI Cov. | Ratio (\times) |
|------------------|---------------------------|-------|-----------------------|----------------------------|----------|--------------------|
| 1 | 114 | 1 | 2.56 | 0.91 | 14.9% | 22.26 |
| 2 | 119 | 1 | 2.78 | 2.17 | 14.4% | 21.40 |
| 3 | 40, 93, 80, 38 | 4 | 3.61 | 3.44 | 54.0% | 36.43 |
| 4 | 34, 120, 84 | 3 | 3.63 | 3.78 | 40.2% | 28.02 |
| 5 | 104, 63, 81 | 3 | 3.22 | 2.49 | 43.6% | 32.16 |
| 6 | 68, 1, 37, 66 | 4 | 4.25 | 2.95 | 52.6% | 34.39 |
| 7 | 56, 71, 78 | 3 | 3.37 | 4.05 | 41.9% | 30.03 |
| 8 | 112, 101 | 2 | 3.67 | 0.77 | 27.1% | 23.36 |
| 9 | 26 | 1 | 1.33 | 0.44 | 17.6% | 27.20 |
| 10 | 114, 84 | 2 | 4.39 | 0.68 | 26.7% | 22.96 |
| 11 | 98, 53, 112 | 3 | 4.19 | 1.28 | 40.1% | 27.89 |
| 12 | 60, 125, 7 | 3 | 3.15 | 1.98 | 43.3% | 31.83 |
| 13 | 65, 78, 56 | 3 | 3.59 | 1.67 | 43.2% | 31.64 |
| 14 | 122 | 1 | 1.78 | 0.62 | 17.1% | 26.25 |
| 15 | 20, 90, 17 | 3 | 4.04 | 3.21 | 38.9% | 26.47 |
| 16 | 116 | 1 | 3.33 | 1.33 | 20.6% | 24.64 |
| 17 | 70, 34, 83 | 3 | 3.48 | 1.74 | 39.4% | 27.05 |
| 18 | 61, 31, 77 | 3 | 3.59 | 2.44 | 46.7% | 32.06 |
| 19 | 6, 105, 21 | 3 | 3.04 | 1.54 | 41.2% | 30.50 |
| 20 | 121, 92, 17 | 3 | 3.89 | 1.83 | 42.1% | 28.01 |
| 21 | 93, 106, 63 | 3 | 3.59 | 2.65 | 41.9% | 27.91 |
| 22 | 99, 106 | 2 | 2.39 | 1.21 | 27.7% | 25.04 |
| 23 | 37, 44, 65 | 3 | 3.78 | 2.71 | 39.2% | 27.69 |
| 24 | 121, 86, 36 | 3 | 3.63 | 2.46 | 42.4% | 29.03 |
| 25 | 52, 35, 24 | 3 | 3.22 | 2.27 | 39.7% | 27.28 |
| 26 | 113, 109, 33 | 3 | 3.81 | 2.13 | 44.5% | 31.03 |
| 27 | 31, 123 | 2 | 2.78 | 1.38 | 28.4% | 24.77 |
| 28 | 78, 73, 62 | 3 | 3.78 | 2.61 | 41.3% | 28.77 |
| 29 | 17, 47, 49 | 3 | 3.74 | 1.87 | 39.4% | 27.54 |
| 30 | 116, 65, 40 | 3 | 3.81 | 2.24 | 41.7% | 28.96 |
| 31 | 57, 24, 92 | 3 | 3.56 | 2.32 | 41.1% | 29.18 |
| 32 | 83, 9, 32 | 3 | 3.63 | 2.49 | 43.6% | 30.14 |
| 33 | 57, 121, 16, 26, 116 | 5 | 3.82 | 2.16 | 67.0% | 49.88 |
| 34 | 9, 96, 110, 64 | 4 | 3.86 | 2.83 | 54.0% | 36.40 |
| 35 | 105, 56 | 2 | 3.11 | 1.53 | 29.0% | 25.75 |
| 36 | 63, 23 | 2 | 3.00 | 2.99 | 28.7% | 25.35 |
| 37 | 96 | 1 | 2.33 | 0.89 | 14.7% | 21.84 |
| 38 | 0 | 1 | 1.11 | 0.10 | 17.2% | 26.40 |
| 39 | 20, 48, 86 | 3 | 4.15 | 1.96 | 40.0% | 27.80 |
| 40 | 85, 49 | 2 | 2.33 | 1.32 | 30.9% | 28.16 |
| 41 | 81, 17, 87 | 3 | 4.30 | 1.07 | 38.9% | 26.51 |
| 42 | 55, 21 | 2 | 2.17 | 2.14 | 30.6% | 27.80 |
| 43 | 31, 6, 56 | 3 | 3.74 | 2.46 | 41.8% | 29.91 |
| 44 | 71, 31 | 2 | 3.33 | 0.84 | 29.6% | 26.53 |
| 45 | 90 | 1 | 2.11 | 0.54 | 16.1% | 24.40 |
| 46 | 107, 94, 101 | 3 | 3.15 | 1.59 | 50.9% | 43.26 |
| 47 | 38, 34 | 2 | 2.39 | 1.21 | 36.8% | 36.65 |
| 48 | 101 | 1 | 2.89 | 0.99 | 14.2% | 20.99 |

the input format differs substantially from MMLU.

High influence density and persistent centralization. We also observe consistently high influence density across layers. The ratio between the average contribution of committee members and that of non-members ranges from $13.48\times$ to $45.22\times$. This means that even when more experts remain nominally available, the effective routing mass is still disproportionately absorbed by a very small committee. The strongest concentration appears in deeper layers, where the ratio exceeds $40\times$, further supporting the view that sparse routing naturally converges toward a centralized computational core.

Implications for robustness across datasets. Taken together, these results suggest that the Standing Committee phenomenon is not an artifact of MMLU’s multiple-choice format. Instead, it persists under a qualitatively different input distribution based on free-form web text. This cross-dataset consistency strengthens our central claim that Standing Committees reflect a general structural tendency of sparse MoE routing, rather than a benchmark-specific effect.

Table 12: Representative Standing Committees on the C4-based evaluation subset using DeepSeek-V2-Lite under the same routing configuration as in the main experiments. We report representative shallow, middle, and deep layers rather than the full layer-wise results. The committees remain compact while capturing substantial routing mass, showing that the Standing Committee phenomenon persists beyond MMLU’s multiple-choice format.

| Layer (ℓ) | Committee Members (C) | $ C $ | Avg. $\mu \downarrow$ | Avg. $\sigma^2 \downarrow$ | ECI Cov. | Ratio (\times) |
|------------------|---------------------------|-------|-----------------------|----------------------------|----------|--------------------|
| 1 | 4, 49, 50, 25 | 4 | 2.41 | 0.55 | 66.32% | 13.78 |
| 2 | 0 | 1 | 1.25 | 0.86 | 12.43% | 18.94 |
| 8 | 32, 6 | 2 | 2.75 | 0.62 | 35.59% | 27.13 |
| 9 | 45, 41 | 2 | 2.50 | 1.19 | 33.78% | 25.81 |
| 10 | 60 | 1 | 2.25 | 1.66 | 17.62% | 13.48 |
| 11 | 16 | 1 | 2.00 | 0.00 | 16.74% | 22.67 |
| 19 | 22, 6, 25, 4 | 4 | 3.31 | 1.64 | 52.92% | 27.36 |
| 20 | 6, 44 | 2 | 4.62 | 1.34 | 21.50% | 18.49 |
| 21 | 56 | 1 | 2.50 | 2.23 | 22.00% | 18.59 |
| 22 | 11, 2 | 2 | 3.12 | 5.22 | 30.56% | 43.64 |
| 23 | 22, 38 | 2 | 3.88 | 4.72 | 27.13% | 31.54 |
| 24 | 48, 61 | 2 | 2.12 | 2.72 | 38.05% | 32.09 |
| 25 | 39, 25, 5 | 3 | 3.75 | 5.06 | 62.81% | 45.22 |
| 26 | 56 | 1 | 1.25 | 0.19 | 19.92% | 15.67 |

E Contribution Concentration Analysis

E.1 Lorenz Curve for OLMoE Model

The Lorenz curves for OLMoE, as shown in Figure 8, reveal a similarly concentrated contribution pattern, despite its smaller scale and more conservative routing design. Across layers, the curves bend sharply away from the equality baseline, with Gini coefficients consistently around 0.88–0.90. This indicates that only a small subset of experts receives the majority of effective routing mass. Even when the nominal expert pool is relatively modest, the allocation of computation remains far from uniform.

Layer-wise inspection shows that this concentration is remarkably stable. Early layers, middle layers, and deep layers all display nearly identical Lorenz profiles, suggesting that specialization does not gradually diversify as representations become more abstract. Instead, OLMoE repeatedly falls back on the same compact subset of experts, while the remaining experts contribute minimally.

Interestingly, the fraction of “used” experts typically lies between 12% and 20%, even though the router is free to assign mass more broadly. This implies that sparsity is driven less by necessity and more by the optimization dynamics of the gating network. Rather than distributing computation in a balanced way, the router converges toward a persistent core of high-traffic experts that dominate inference across inputs.

Taken together, the Lorenz curves demonstrate that contribution inequality is not merely a byproduct of model scale. Even in OLMoE, contribution is highly centralized, reinforcing the broader Standing Committee pattern: most experts exist on the

periphery, while a small, repeatedly selected core absorbs the majority of computational responsibility.

E.2 Lorenz Curve for DeepSeek-V2-Lite Model

DeepSeek-V2-Lite exhibits an even sharper form of contribution concentration, as shown in Figure 9. Across layers, the Lorenz curves bend aggressively toward the lower-right corner, with Gini coefficients consistently around 0.91–0.92. This indicates that routing mass is dominated by a very small subset of experts. In several layers, fewer than 15% of experts account for nearly all effective contributions, while the remaining experts receive negligible traffic.

Unlike what one might expect from a lightweight architecture optimized for efficiency, the inequality pattern does not relax as depth increases. Early, middle, and late layers display almost indistinguishable Lorenz shapes. The router repeatedly converges to the same compact set of high-traffic experts, rather than distributing load adaptively as representations evolve.

A notable pattern is the oscillation in the proportion of “used” experts. Some layers activate roughly 20% of the pool, while others rely on as little as 9%. However, even in layers with broader activation, the cumulative share of contribution remains steeply skewed. This suggests that the increase in participation does not meaningfully change who dominates, but merely introduces additional peripheral experts who play marginal roles.

Taken together, these curves reinforce the central observation: contribution concentration is not mitigated by architectural simplification. DeepSeek-V2-Lite still organizes computation around a small, persistent Standing Committee, while most experts remain structurally available but functionally underutilized.

E.3 Lorenz Curve for Qwen3 Model

As shown in Figure 10, Qwen-30B-A3B shows one of the strongest forms of contribution inequality across all models we study. The Lorenz curves are almost vertical near the right edge, yielding Gini coefficients around 0.94 across layers. This indicates that routing mass is funneled into an extremely small subset of experts. In several layers, fewer than 10% of experts account for nearly all effective contribution, leaving the majority effectively idle.

The pattern is also highly consistent across depth. Early layers, mid-depth layers, and late layers exhibit nearly identical Lorenz profiles, suggesting that the model does not gradually diversify expert usage as representations become more abstract. Instead, the router repeatedly returns to the same high-traffic subset, which acts as a default computational pathway for most inputs.

A striking observation is that the proportion of “active” experts fluctuates between 6% and 12%, yet the inequality curve barely changes. Even when more experts are nominally activated, the cumulative contribution remains concentrated in a tiny elite group. Additional experts merely contribute marginal amounts, without altering the dominance structure.

These results indicate that contribution centralization intensifies as model capacity increases. In Qwen-30B-A3B, a large pool of experts does not translate into broader participation. Rather, the gating dynamics amplify the emergence of a persistent Standing Committee, while most experts remain structurally available but functionally peripheral.

E.4 Cross-Model Synthesis.

Across architectures of very different sizes and routing designs, we observe a consistent pattern of extreme contribution concentration. OLMoE, DeepSeek-V2-Lite, and Qwen-30B-A3B all display Lorenz curves that deviate sharply from the equality baseline, with Gini coefficients typically above 0.88 and often exceeding 0.94. In every case, only a small fraction of experts accounts for the vast majority of effective routing mass, while the remaining experts make negligible contributions.

Importantly, this phenomenon persists across depth. Early layers, middle layers, and late layers show nearly identical inequality profiles, indicating that expert participation does not broaden as representations become more abstract. Instead, the gating networks repeatedly allocate computation to a compact, stable subset of experts that serve as default processing routes, regardless of layer position or domain.

At the same time, fluctuations in the proportion of “used” experts do not materially change this distribution. Even when more experts are nominally activated, the cumulative contribution remains dominated by the same small core. Additional experts tend to act as low-impact auxiliaries rather than genuine participants in computation.

Taken together, these results suggest that con-

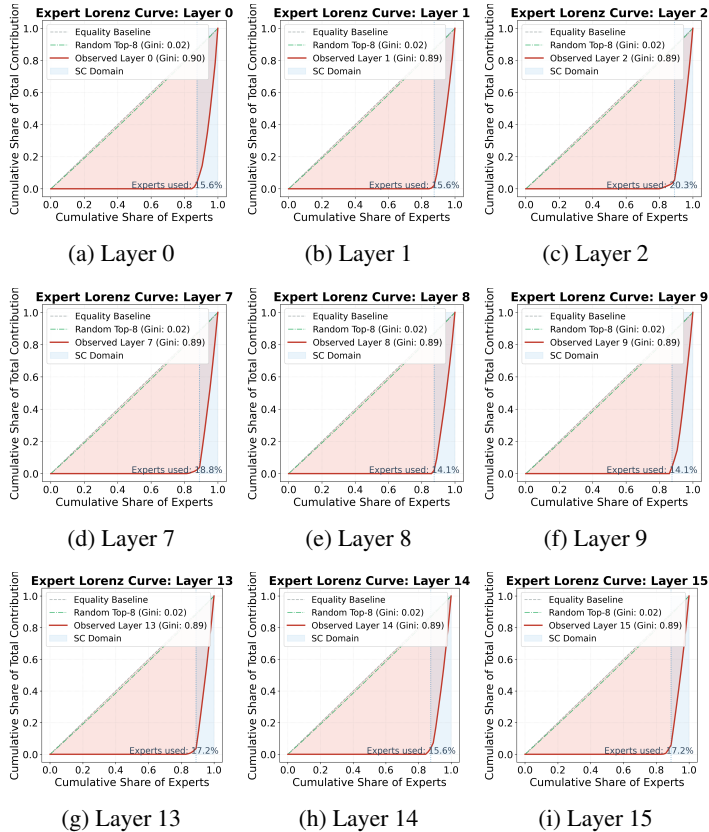


Figure 8: Expert Lorenz Curves across layers for OLMoE model.

tribution concentration is not merely an artifact of scale, architecture, or routing hyperparameters. Rather, it reflects a robust inductive tendency of sparse MoE optimization. The models converge toward a Standing Committee structure, in which a persistent core of experts monopolizes computation while most experts operate peripherally.

F Full Masking-Based Intervention Results

To further examine the functional importance of Standing Committee experts, we conduct a masking-based intervention on DeepSeek-V2-Lite. For each selected layer, we suppress the routing weights of the identified Standing Committee experts, re-normalize the remaining routing distribution, and then evaluate answer outcomes on MMLU. While the main text reports representative shallow, middle, and deep layers, Table 13 provides the complete set of evaluated layers.

The full results reveal a consistent degradation pattern across depth. Relative to the unmasked baseline, masking committee experts sharply reduces the correct answer rate and substantially increases the proportion of *No Answer* responses. The effect becomes especially severe in deeper lay-

ers, where correct responses nearly collapse.

| Phase | Mask Layer | Correct | Wrong | No Answer |
|----------|------------|---------|-------|-----------|
| Baseline | None | 0.39 | 0.58 | 0.03 |
| Shallow | 1 | 0.19 | 0.46 | 0.35 |
| Shallow | 2 | 0.12 | 0.52 | 0.36 |
| Shallow | 3 | 0.17 | 0.50 | 0.33 |
| Middle | 9 | 0.11 | 0.53 | 0.36 |
| Middle | 10 | 0.09 | 0.55 | 0.36 |
| Middle | 11 | 0.10 | 0.54 | 0.36 |
| Deep | 19 | 0.07 | 0.56 | 0.37 |
| Deep | 20 | 0.06 | 0.57 | 0.37 |
| Deep | 21 | 0.05 | 0.58 | 0.37 |
| Deep | 25 | 0.04 | 0.58 | 0.38 |
| Deep | 26 | 0.03 | 0.59 | 0.38 |

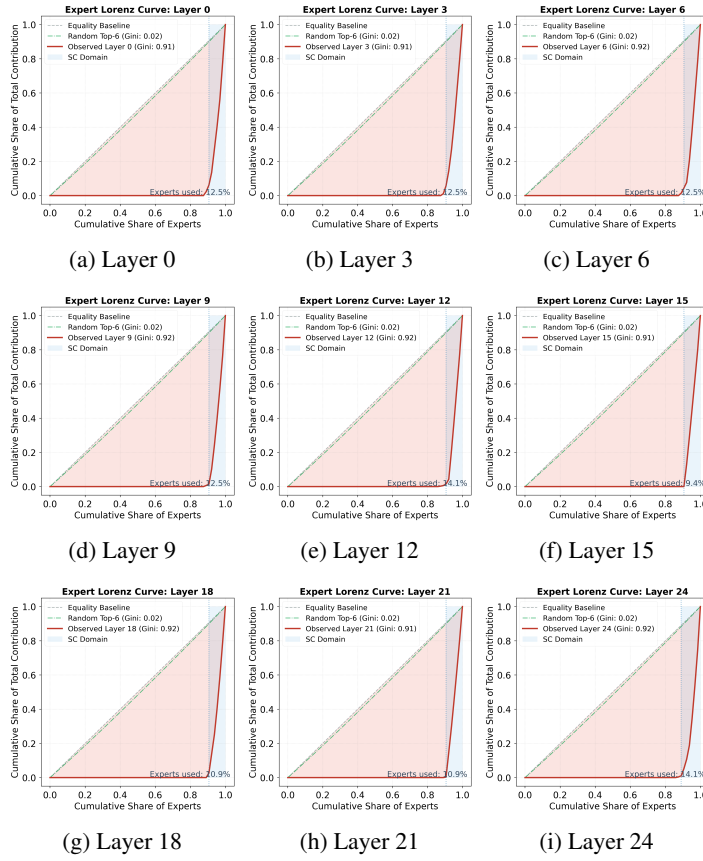


Figure 9: Expert Lorenz Curves across layers for DeekSeek-V2-Lite model.

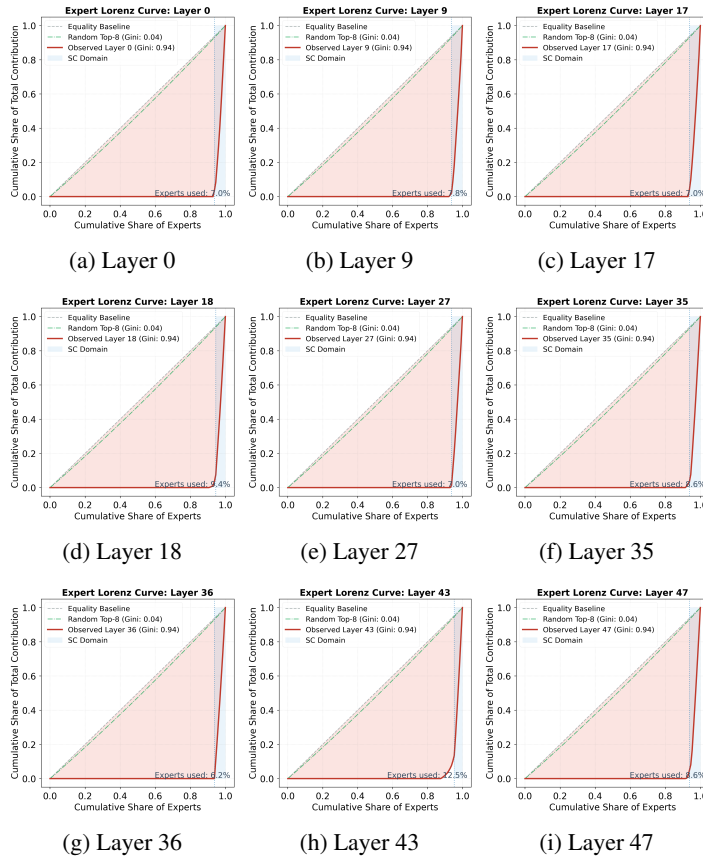


Figure 10: Expert Lorenz Curves across layers for QWen3-30B-A3B model.