

Two Pathways to Truthfulness: On the Intrinsic Encoding of LLM Hallucinations

Wen Luo[♡], Guangyue Peng[♡], Wei Li[♡], Shaohang Wei[♡], Feifan Song[♡],
Liang Wang[♣], Nan Yang[♣], Xingxing Zhang[♣], Jing Jin[♡], Furu Wei[♣], Houfeng Wang^{♡*}

[♡] State Key Laboratory of Multimedia Information Processing,

School of Computer Science, Peking University

[♣] Microsoft Research Asia

rowan.wenluo@gmail.com, wanghf@pku.edu.cn

Abstract

Despite their impressive capabilities, large language models (LLMs) frequently generate hallucinations. Previous work shows that their internal states encode rich signals of truthfulness, yet the origins and mechanisms of these signals remain unclear. In this paper, we demonstrate that truthfulness cues arise from two distinct information pathways: (1) a Question-Anchored pathway that depends on question-answer information flow, and (2) an Answer-Anchored pathway that derives self-contained evidence from the generated answer itself. First, we validate and disentangle these pathways through attention knockout and token patching. Afterwards, we uncover notable and intriguing properties of these two mechanisms. Further experiments reveal that (1) the two mechanisms are closely associated with LLM knowledge boundaries; and (2) internal representations are aware of their distinctions. Finally, building on these insightful findings, two applications are proposed to enhance hallucination detection performance. Overall, our work provides new insight into how LLMs internally encode truthfulness, offering directions for more reliable and self-aware generative systems.¹

1 Introduction

Despite their remarkable capabilities in natural language understanding and generation, large language models (LLMs) often produce *hallucinations*—outputs that appear plausible but are factually incorrect. This phenomenon poses a critical challenge for deploying LLMs in real-world applications where reliability and trustworthiness are paramount (Shi et al., 2024; Bai et al., 2024). One line of research tackles hallucination detection from an extrinsic perspective (Min et al., 2023; Hu et al., 2025; Huang et al., 2025), evaluating only

the model’s outputs while disregarding its internal dynamics. Although such approaches can identify surface-level textual inconsistencies, their extrinsic focus limits the insight they offer into the underlying causes of hallucinations. Complementing these efforts, another line of work investigates the intrinsic properties of LLMs, revealing that their internal representations encode rich truthfulness signals (Burns et al., 2023; Li et al., 2023; Chen et al., 2024; Orgad et al., 2025; Niu et al., 2025). These internal truthfulness signals can be exploited to detect an LLM’s own generative hallucinations by training a linear classifier (i.e., a probe) on its hidden representations. However, while prior work establishes the presence of such cues, the mechanisms by which they arise and operate remain largely unexplored. Recent studies indicate well-established mechanisms in LLMs that underpin complex capabilities such as in-context learning (Wang et al., 2023), long-context retrieval (Wu et al., 2025), and reasoning (Qian et al., 2025). This observation naturally leads to a key question: *how do truthfulness cues arise and function within LLMs?*

In this paper, we uncover that truthfulness signals in LLMs arise from **two distinct information pathways**: (1) a **Question-Anchored (Q-Anchored) pathway**, which depends on the flow of information from the input question to the generated answer, and (2) an **Answer-Anchored (A-Anchored) pathway**, which derives self-contained evidence directly from the model’s own outputs. We begin with a preliminary study using saliency analysis to quantify information flow potentially relevant to hallucination detection. Results reveal a bimodal distribution of dependency on question-answer interactions, suggesting heterogeneous truthfulness encoding mechanisms. To validate this hypothesis, we design two experiments across 4 diverse datasets using 12 models that vary in both architecture and scale, including base, instruction-tuned, and reasoning-oriented models.

*Corresponding author

¹Code is publicly available at <https://github.com/RowanWenLuo/llm-truthfulness-pathways>

By (i) blocking critical question-answer information flow through attention knockout (Geva et al., 2023; Fierro et al., 2025) and (ii) injecting hallucinatory cues into questions via token patching (Ghandeharioun et al., 2024; Todd et al., 2024), we disentangle these truthfulness pathways. Our analyses confirm that Q-Anchored signals rely heavily on question-derived cues, whereas A-Anchored signals are robust to their removal and primarily originate from the generated answer itself.

Building on this foundation, we further investigate emergent properties of these truthfulness pathways through large-scale experiments. Our findings highlight two intriguing characteristics: **(1) Association with knowledge boundaries:** Q-anchored encoding predominates for well-established facts that fall within the knowledge boundary, whereas A-anchored encoding is favored in long-tail cases. **(2) Self-awareness:** LLM internal states can distinguish which mechanism is being employed, suggesting intrinsic awareness of pathway distinctions.

Finally, these analyses not only deepen our mechanistic understanding of hallucinations but also enable practical applications. Specifically, by leveraging the fundamentally different dependencies of the truthfulness pathways and the model’s intrinsic awareness, we propose two pathway-aware strategies to enhance hallucination detection. **(1) Mixture-of-Probes (MoP):** Motivated by the specialization of internal pathways, MoP employs a set of expert probing classifiers, each tailored to capture distinct truthfulness encoding mechanisms. **(2) Pathway Reweighting (PR):** From the perspective of selectively emphasizing pathway-relevant internal cues, PR modulates information intensity to amplify signals that are most informative for hallucination detection, aligning internal activations with pathway-specific evidence. Experiments demonstrate that our proposed methods consistently outperform competing approaches, achieving up to a 10% AUC gain across various datasets and models.

Overall, our key contributions are summarized as follows:

- **(Mechanism)** We conduct a systematic investigation into how internal truthfulness signals emerge and operate within LLMs, revealing two distinct information pathways: a *Question-Anchored* pathway that relies on question-answer information flow, and an *Answer-Anchored* pathway that derives self-contained evidence from the generated output.

- **(Discovery)** Through large-scale experiments across multiple datasets and model families, we identify two key properties of these mechanisms: (i) association with knowledge boundaries, and (ii) intrinsic self-awareness of pathway distinctions.
- **(Application)** Building on these findings, we propose two pathway-aware detection methods that exploit the complementary nature of the two mechanisms to enhance hallucination detection, providing new insights for building more reliable generative systems.

2 Background

2.1 Hallucination Detection

Given an LLM f , we denote the dataset as $D = \{(q_i, \hat{y}_i^f, z_i^f)\}_{i=1}^N$, where q_i is the question, \hat{y}_i^f the model’s answer in open-ended generation, and $z_i^f \in \{0, 1\}$ indicates whether the answer is hallucinatory. The task is to predict z_i^f given the input $x_i^f = [q_i, \hat{y}_i^f]$ for each instance. Cases in which the model refuses to answer are excluded, as they are not genuine hallucinations and can be trivially classified. Methods based on internal signals assume access to the model’s hidden representations but no external resources (e.g., retrieval systems or fact-checking APIs) (Xue et al., 2025a). Within this paradigm, probing trains a lightweight linear classifier on hidden activations to discriminate between hallucinatory and factual outputs, and has been shown to be among the most effective approaches in this class of internal-signal-based methods (Orgad et al., 2025).

2.2 Exact Question and Answer Tokens

To analyze the origins and mechanisms of truthfulness signals in LLMs, we primarily focus on **exact tokens** in question-answer pairs. Not all tokens contribute equally to detecting factual errors: some carry core information essential to the meaning of the question or answer, while others provide peripheral details. We draw on *semantic frame* theory (Baker et al., 1998; Pagnoni et al., 2021), which represents a situation or event along with its participants and their roles. In the theory, frame elements are categorized as: (1) *Core frame elements*, which define the situation itself, and (2) *Non-core elements*, which provide additional, non-essential context.

As shown in Table 1, we define: (1) **Exact question tokens**: core frame elements in the question, typically including **the exact subject and property tokens** (i.e., *South Carolina* and *capital*). (2) **Exact answer tokens**: core frame elements in the answer that convey the critical information required to respond correctly (i.e., *Columbia*). Humans tend to rely more on core elements when detecting errors, as these tokens carry the most precise information. Consistent with this intuition, recent work (Orgad et al., 2025) shows that probing activations on the exact answer tokens offers the strongest signal for hallucination detection, outperforming all other token choices. Motivated by these findings, our analysis mainly centers on exact tokens to probe truthfulness signals in LLMs. Moreover, to validate the robustness of our conclusions, we also conduct comprehensive experiments using alternative, non-exact-token configurations (see Appendix B.2).

Question: What is the capital of South Carolina ?

Answer: It is Columbia, a hub for government, culture, and education that houses the South Carolina State House and the University of South Carolina.

Table 1: Example of exact question and answer tokens. Colors indicate token types: ■ - exact property, ■ - exact subject, and ■ - exact answer tokens.

3 Two Internal Truthfulness Pathways

We begin with a preliminary analysis using metrics based on saliency scores (§3.1). The quantitative results reveal **two distinct information pathways for truthfulness encoding**: (1) a **Question-Anchored (Q-Anchored) Pathway**, which relies heavily on exact question tokens (i.e., the questions), and (2) an **Answer-Anchored (A-Anchored) Pathway**, in which the truthfulness signal is largely independent of the question-to-answer information flow. Section 3.2 presents experiments validating this hypothesis. In particular, we show that Q-Anchored Pathway depends critically on information flowing from the question to the answer, whereas the signals along the A-Anchored Pathway are primarily derived from the LLM-generated answer itself.

3.1 Saliency-Driven Preliminary Study

This section investigates the intrinsic characteristics of LLM attention interactions and their potential role in truthfulness encoding. We employ

saliency analysis (Simonyan et al., 2014), a widely used interpretability method, to reveal how attention among tokens influences probe decisions. Following common practice (Michel et al., 2019; Wang et al., 2023), we compute the saliency score as:

$$S^l(i, j) = \left| A^l(i, j) \frac{\partial \mathcal{L}(x)}{\partial A^l(i, j)} \right|, \quad (1)$$

where S^l denotes the saliency score matrix of the l -th layer, A^l represents the attention weights of that layer, and \mathcal{L} is the loss function for hallucination detection (i.e., the binary cross-entropy loss). Scores are averaged over all attention heads within each layer. In particular, $S^l(i, j)$ quantifies the saliency of attention from query i to key j , capturing how strongly the information flow from j to i contributes to the detection. We study two types of information flow: (1) $S_{E_Q \rightarrow E_A}$, the saliency of direct information flow from the exact question tokens to the exact answer tokens, and (2) $S_{E_Q \rightarrow *}$, the saliency of the total information disseminated by the exact question tokens.




Figure 1: Kernel density estimates of saliency-score distributions for critical question-to-answer information flows. The bimodal pattern suggests two distinct information mechanisms.

Results We demonstrate Kernel Density Estimation results of the saliency scores on TriviaQA (Joshi et al., 2017) and Natural Questions (Kwiatkowski et al., 2019) datasets. As shown in Figure 1, probability densities reveal a clear bimodal distribution: for all examined information types originating from the question, the probability mass concentrates around two peaks, one near zero saliency and another at a substantially higher value. The near-zero peak suggests that, for a substantial subset of samples, the question-to-answer information flow contributes minimally to hallucination detection, whereas the higher peak reflects strong dependence on such flow.

Hypothesis These observations lead to the hypothesis that there are **two distinct mechanisms of**




Figure 2: ΔP under attention knockout. The layer axis indicates the Transformer layer on which the probe is trained. Shaded regions indicate 95% confidence intervals. Full results in Appendix C.

internal truthfulness encoding for hallucination detection: (1) one characterized by strong reliance on the key question-to-answer information from the exact question tokens, and (2) one in which truthfulness encoding is largely independent of the question. We validate the proposed hypothesis through further experiments in the next section.

3.2 Disentangling Information Mechanisms

We hypothesize that the internal truthfulness encoding operates through two distinct information flow mechanisms, driven by the attention modules within Transformer blocks. To validate the hypothesis, we first block information flows associated with the exact question tokens and analyze the resulting changes in the probe’s predictions. Subsequently, we apply a complementary technique, called token patching, to further substantiate the existence of these two mechanisms. Finally, we demonstrate that the self-contained information from the LLM-generated answer itself drives the truthfulness encoding for the A-Anchored type.

3.2.1 Experimental Setup

Our analysis covers a diverse collection of 12 LLMs that vary in both scale and architectural design. Specifically, we consider three categories: (1) *base models*, including Llama-3.2-1B (Grattafiori et al., 2024), Llama-3.2-3B, Llama-3-8B, Llama-3-70B, Mistral-7B-v0.1 (Jiang et al., 2023), and Mistral-7B-v0.3; (2) *instruction-tuned models*, including Llama-3.2-3B-Instruct, Llama-3-8B-Instruct, Mistral-7B-Instruct-v0.1, and Mistral-7B-Instruct-v0.3; and (3) *reasoning-oriented models*, namely Qwen3-8B (Yang et al., 2025) and Qwen3-32B. We conduct experiments on 4 widely used question-answering datasets: PopQA (Mallen et al., 2023), TriviaQA (Joshi et al., 2017), Hot-

potQA (Yang et al., 2018), and Natural Questions (Kwiatkowski et al., 2019). Additional implementation details are provided in Appendix B.

3.2.2 Identifying Anchored Modes via Attention Knockout

Experiment To investigate whether internal truthfulness encoding operates via distinct information mechanisms, we perform an attention knockout experiment targeting the exact question tokens. Specifically, for a probe trained on representations from the k -th layer, we set $A_l(i, E_Q) = 0$ for layers $l \in \{1, \dots, k\}$ and positions $i > E_Q$. This procedure blocks the information flow from question tokens to subsequent positions in the representation. We then examine how the probe’s predictions respond to this intervention. To provide a clearer picture, instances are categorized according to whether their prediction \hat{z} changes after the attention knockout:

$$\text{Mode}(x) = \begin{cases} \text{Q-Anchored}, & \text{if } \hat{z} \neq \tilde{z} \\ \text{A-Anchored}, & \text{otherwise} \end{cases} \quad (2)$$

where \hat{z} and \tilde{z} denote predictions before and after the attention knockout, respectively.

Results The results in Figure 2 and Appendix C reveal a clear bifurcation of behaviors: for one subset of instances, probabilities shift substantially, while for another subset, probabilities remain nearly unchanged across all layers. Shaded regions indicate 95% confidence intervals, confirming that this qualitative separation is statistically robust. This sharp divergence supports the hypothesis that internal truthfulness encoding operates via two distinct mechanisms with respect to question-answer information. In Appendix C, we conduct




Figure 3: Prediction flip rate under token patching. Q-Anchored samples demonstrate significantly higher sensitivity than the counterparts when hallucinatory cues are injected into exact questions. Full results in Appendix D.

a comprehensive analysis of alternative configurations for token selection, activation extraction, and various instruction- or reasoning-oriented models, and observe consistent patterns across all settings. Moreover, Figure 16 in Appendix C shows that blocking information from randomly selected question tokens yields negligible changes, in contrast to blocking exact question tokens, underscoring the nontrivial nature of the identified mechanisms.

3.2.3 Further Validation via Token Patching

Experiment To further validate our findings, we employ a critical token patching technique to investigate how the internal representations of the LLM respond to hallucinatory signals originating from exact question tokens under the two proposed mechanisms. Given a context sample d_c , we randomly select a patch sample d_p and replace the original question tokens E_Q^c in d_c with the exact question tokens E_Q^p from d_p . This operation introduces hallucinatory cues into the context sample, allowing us to assess whether the LLM’s internal states appropriately reflect the injected changes. We restrict our analysis to context instances where the original LLM answers are factual, ensuring that any observed changes can be attributed solely to the injected hallucinatory cues.

Results We measure the sensitivity of the truthfulness signals using the prediction flip rate, defined as the frequency with which the probe’s prediction changes after hallucinatory cues are introduced. Figure 3 and Appendix D present the results of the best-performing layer of each model on four datasets when patching the exact subject tokens. Across models and datasets, Q-Anchored mode exhibits significantly higher sensitivity compared to A-Anchored mode when exposed to hallucination cues from the questions. Furthermore, within each pathway, the flip rates where exact question

tokens are patched are substantially higher than those observed when random tokens are patched, ruling out the possibility that the observed effects are mainly due to general semantic disruption from token replacement. These consistent results provide further support for our hypothesis regarding distinct mechanisms of information pathways.

3.2.4 What Drives A-Anchored Encoding?

Experiment Since the A-Anchored mode operates largely independently of the question-to-answer information flow, it is important to investigate the source of information it uses to identify hallucinations. To this end, we remove the questions entirely from each sample and perform a separate forward pass using only the LLM-generated answers. This procedure yields answer-only hidden states, which are subsequently provided as input to the probe. We then evaluate how the probe’s predictions change under this “answer-only” condition. This setup enables us to assess whether A-Anchored predictions rely primarily on the generated answer itself rather than on the original question.

Results As shown in Figure 4 and Appendix E, Q-Anchored instances exhibit substantial changes in prediction probability when the question is removed, reflecting their dependence on question-to-answer information. In contrast, A-Anchored instances remain largely invariant, indicating that the probe continues to detect hallucinations using information encoded within the LLM-generated answer itself. These findings suggest that the A-Anchored mechanism primarily leverages self-contained answer information to build signals about truthfulness.




Figure 4: $-\Delta P$ with only the LLM-generated answer. Q-Anchored instances exhibit substantial shifts, whereas A-Anchored instances remain stable, confirming that A-Anchored truthfulness encoding relies on information in the LLM-generated answer itself. Full results in Appendix E.

4 Properties of Truthfulness Pathways

This section examines notable properties and distinct behaviors of intrinsic truthfulness encoding: **(1) Associations with knowledge boundaries:** samples within the LLM’s knowledge boundary tend to encode truthfulness via the Q-Anchored pathway, whereas samples beyond the boundary often rely on the A-Anchored signal; **(2) Self-awareness:** internal representations can be used to predict which mechanism is being employed, suggesting that LLMs possess intrinsic awareness of pathway distinctions.

4.1 Associations with Knowledge Boundaries

We find that distinct patterns of truthfulness encoding are closely associated with the knowledge boundaries of LLMs. To characterize these boundaries, three complementary metrics are employed: **(1) Answer accuracy**, the most direct indicator of an LLM’s factual competence; **(2) I-don’t-know rate** (shown in Appendix G), which reflects the model’s ability to recognize and express its own knowledge limitations; **(3) Entity popularity**, which is widely used to distinguish between common and long-tail factual knowledge (Mallen et al., 2023).

As shown in Figure 5 and Appendix F, Q-Anchored samples achieve significantly higher accuracy than those driven by the A-Anchored pathway. The results for the I-don’t-know rate, reported in Appendix G, exhibit trends consistent with answer accuracy, further indicating stronger knowledge handling in Q-Anchored samples. Moreover, entity popularity, shown in Figure 6, provides a more fine-grained perspective on knowledge boundaries. Specifically, Q-Anchored samples tend to involve more popular entities, whereas A-Anchored samples are more frequently associated with less

popular, long-tail factual knowledge. These findings suggest that truthfulness encoding is strongly aligned with the availability of stored knowledge: when LLMs possess the requisite knowledge, they predominantly rely on question-answer information flow (Q-Anchored); when knowledge is unavailable, they instead draw upon internal patterns within their own generated outputs (A-Anchored).

4.2 Self-Awareness of Pathway Distinctions

Given that LLMs encode truthfulness via two distinct mechanisms, this section investigates whether their internal representations contain discriminative information that can be used to distinguish between these mechanisms. To this end, we train probing classifiers on the models’ original internal states (i.e., without knockout interventions) to predict which mechanism is being utilized.

Table 2 reports the pathway classification results of the best-performing layers in hallucination detection across different models. Our findings demonstrate that different mechanisms can be reliably inferred from internal representations, suggesting that, in addition to encoding truthfulness, LLMs exhibit intrinsic awareness of pathway distinctions. These findings highlight a potential avenue for fine-grained improvements targeting specific truthfulness encoding mechanisms.

Datasets	Llama-3-8B	Llama-3-70B	Mistral-7B-v0.3
PopQA	87.80	92.66	87.64
TriviaQA	75.10	83.91	85.87
HotpotQA	86.31	87.34	92.13
NQ	78.31	84.14	84.83

Table 2: AUCs for encoding pathway classification. The predictability from internal representations indicates that LLMs possess intrinsic awareness of pathway distinctions.




Figure 5: Comparisons of answer accuracy between pathways. Q-Anchored samples show higher accuracy than A-Anchored ones, highlighting the association between truthfulness encoding and LLM knowledge boundaries. Full results in Appendix F and G.




Figure 6: Entity frequency distributions for both pathways on PopQA. Q-Anchored samples concentrate on more popular entities, whereas A-Anchored samples skew toward long-tail entities.

5 Pathway-Aware Detection

Building on the intriguing findings, we explore how the discovered pathway distinctions can be leveraged to improve hallucination detection. Specifically, two simple yet effective pathway-aware strategies are proposed: (1) **Mixture-of-Probes (MoP)** (§5.1), which allows expert probes to specialize in Q-Anchored and A-Anchored pathways respectively, and (2) **Pathway Reweighting (PR)** (§5.2), a plug-and-play approach that amplifies pathway-relevant cues salient for detection.

5.1 Mixture-of-Probes

Motivated by the fundamentally different dependencies of the two encoding pathways and the LLMs’ intrinsic awareness of them, we propose a **Mixture-of-Probes (MoP)** framework that explicitly captures this heterogeneity. Rather than training a single probe to handle all inputs, MoP employs two pathway-specialized experts and leverages the self-awareness probe (§4.2) as a gating network to combine their predictions. Let $\mathbf{h}^{l^*}(x) \in \mathbb{R}^d$ be the token hidden state from the best detection

layer l^* . Two expert probes $p_Q(\cdot)$ and $p_A(\cdot)$ are trained separately for two pathway samples, and the self-awareness probe provides a gating coefficient $\pi(\mathbf{h}^{l^*}(x)) \in [0, 1]$. The final prediction is a convex combination, **requiring no extra training**:

$$p_{\text{MoP}}(z=1 | \mathbf{h}^{l^*}(x)) = \pi_Q p_Q(z=1 | \mathbf{h}^{l^*}(x)) + (1 - \pi_Q) p_A(z=1 | \mathbf{h}^{l^*}(x)). \quad (3)$$

5.2 Pathway Reweighting

From the perspective of emphasizing pathway-relevant internal cues, we introduce a plug-and-play **Pathway Reweighting (PR)** method that directly modulates the question-answer information flow. The key idea is to adjust the attention from exact answer to question tokens according to the predicted pathway, amplifying the signals most salient for hallucination detection. For each layer $l \leq l^*$, two learnable scalars $\alpha_Q^l, \alpha_A^l > 0$ are introduced. Given self-awareness probability $\pi(\mathbf{h}^{l^*}(x))$, we rescale attention edges $i \in E_A, j \in E_Q$ to construct representations tailored for detection:

$$\tilde{A}^l(i, j) = \begin{cases} [1 + s(\mathbf{h}^{l^*}(x))] A^l(i, j), & i \in E_A, j \in E_Q, \\ A^l(i, j), & \text{otherwise,} \end{cases} \quad (4)$$

where

$$s(\mathbf{h}^{l^*}(x)) = \pi_Q \alpha_Q^l - (1 - \pi_Q) \alpha_A^l. \quad (5)$$

The extra parameters serve as a lightweight adapter, used only during detection to guide salient truthfulness cues and omitted during generation, leaving the generation capacity unaffected.

5.3 Experiments

Setup The experimental setup follows Section 3.2.1. We compare our method against several

Method	Llama-3-8B				Mistral-7B-v0.3			
	PopQA	TriviaQA	HotpotQA	NQ	PopQA	TriviaQA	HotpotQA	NQ
P(True)	55.85	49.92	52.14	53.27	45.49	47.61	57.87	52.79
Logits-mean	74.52	60.39	51.94	52.63	69.52	66.76	55.45	57.88
Logits-min	85.36	70.89	61.28	56.50	87.05	77.33	68.08	54.40
Probing Baseline	88.71	77.58	82.23	70.20	87.39	81.74	83.19	73.60
MoP-RandomGate	75.52	69.17	79.88	66.56	79.81	70.88	72.23	61.19
MoP-VanillaExperts	89.11	78.73	84.57	71.21	88.53	80.93	82.93	73.77
MoP	92.11	81.18	85.45	74.64	91.66	83.57	85.82	76.87
PR	94.01	83.13	87.81	79.10	93.09	84.36	89.03	79.09

Table 3: Comparison of hallucination detection performance (AUC). Full results in Appendix H.

internal-based baselines, including (1) P(True) (Kadavath et al., 2022), (2) uncertainty-based metrics (Aichberger et al., 2024; Xue et al., 2025a), and (3) probing classifiers (Chen et al., 2024; Orgad et al., 2025). Results are averaged over three random seeds. Additional implementation details are provided in Appendix B.5 and B.6.

Results As shown in Table 3 and Appendix H, both MoP and PR consistently outperform competing approaches across different datasets and model scales. Specifically, for MoP, we further examine two ablated variants: (1) MoP-RandomGate, which randomly routes the two pathway experts without leveraging the self-awareness probe; and (2) MoP-VanillaExperts, which replaces the expert probes with two vanilla probes to serve as a simple ensemble strategy. Both ablated variants exhibit substantially degraded performance compared to MoP, underscoring the roles of pathway specialization and self-awareness gating. For PR, the method proves particularly effective in improving performance by dynamically adjusting the focus on salient truthfulness cues. These results demonstrate that explicitly modeling truthfulness encoding heterogeneity can effectively translate the insights of our analysis into practical gains for hallucination detection.

6 Related Work

Hallucination detection in LLMs has received increasing attention because of its critical role in building reliable and trustworthy generative systems (Tian et al., 2024; Shi et al., 2024; Bai et al., 2024). Existing approaches can be broadly grouped by whether they rely on external resources (e.g., retrieval systems or fact-checking APIs). Externally assisted methods cross-verify output texts against external knowledge bases (Min et al., 2023; Hu et al., 2025; Huang et al., 2025) or specialized LLM judges (Luo et al., 2024; Bouchard and Chauhan,

2025; Zhang et al., 2025). Resource-free methods avoid external data and instead exploit the model’s own intermediate computations. Some leverage the model’s self-awareness of knowledge boundaries (Kadavath et al., 2022; Luo et al., 2025), while others use uncertainty-based measures (Aichberger et al., 2024; Xue et al., 2025a), treating confidence as a proxy for truthfulness. These techniques analyze output distributions (e.g., logits) (Aichberger et al., 2024), variance across multiple samples (e.g., consistency) (Min et al., 2023; Aichberger et al., 2025), or other statistical indicators of prediction uncertainty (Xue et al., 2025b). Another line of work trains linear probing classifiers on hidden representations to capture intrinsic truthfulness signals. Prior work (Burns et al., 2023; Li et al., 2023; Chen et al., 2024; Orgad et al., 2025) shows that LLMs encode rich latent features correlated with factual accuracy, enabling efficient detection with minimal overhead. Yet the mechanisms behind these internal truthfulness encoding remain poorly understood. Compared to previous approaches, our work addresses this gap by dissecting how such intrinsic signals emerge and operate, revealing distinct information pathways that not only yield explanatory insights but also enhance detection performance.

7 Conclusion

We investigate how LLMs encode truthfulness, revealing two complementary pathways: a **Question-Anchored** pathway relying on question-answer flow, and an **Answer-Anchored** pathway extracting self-contained evidence from generated outputs. Analyses across datasets and models highlight their ties to knowledge boundaries and intrinsic self-awareness. Building on these insights, we further propose two applications to improve hallucination detection. Overall, our findings not only advance mechanistic understanding of intrinsic truthfulness

encoding but also offer practical applications for building more reliable generative systems.

Limitations

While this work provides a systematic analysis of intrinsic truthfulness encoding mechanisms in LLMs and demonstrates their utility for hallucination detection, one limitation is that, similar to prior work on mechanistic interpretability, our analyses and pathway-aware applications assume access to internal model representations. Such access may not always be available in strictly black-box settings. In these scenarios, additional engineering or alternative approximations may be required for practical deployment, which we leave for future work.

Ethics Statement

Our work presents minimal potential for negative societal impact, primarily due to the use of publicly available datasets and models. This accessibility inherently reduces the risk of adverse effects on individuals or society.

Acknowledgments

This work was supported by Beijing Natural Science Foundation (L253020) & the Academic Research Projects of Beijing Union University (NO.ZK10202405).

References

Lukas Aichberger, Kajetan Schweighofer, Mykyta Ielanskyi, and Sepp Hochreiter. 2024. Semantically diverse language generation for uncertainty estimation in language models. *arXiv preprint arXiv:2406.04306*.

Lukas Aichberger, Kajetan Schweighofer, Mykyta Ielanskyi, and Sepp Hochreiter. 2025. [Improving uncertainty estimation through semantically diverse language generation](#). In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net.

Ge Bai, Jie Liu, Xingyuan Bu, Yancheng He, Jiaheng Liu, Zhanhui Zhou, Zhuoran Lin, Wenbo Su, Tiezheng Ge, Bo Zheng, and Wanli Ouyang. 2024. [Mt-bench-101: A fine-grained benchmark for evaluating large language models in multi-turn dialogues](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 7421–7454. Association for Computational Linguistics.

Collin F Baker, Charles J Fillmore, and John B Lowe. 1998. The berkeley framenet project. In *36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics, Volume 1*, pages 86–90.

Dylan Bouchard and Mohit Singh Chauhan. 2025. Uncertainty quantification for language models: A suite of black-box, white-box, llm judge, and ensemble scorers. *arXiv preprint arXiv:2504.19254*.

Collin Burns, Haotian Ye, Dan Klein, and Jacob Steinhardt. 2023. [Discovering latent knowledge in language models without supervision](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.

Chao Chen, Kai Liu, Ze Chen, Yi Gu, Yue Wu, Mingyuan Tao, Zhihang Fu, and Jieping Ye. 2024. [INSIDE: llms' internal states retain the power of hallucination detection](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.

Constanza Fierro, Negar Foroutan, Desmond Elliott, and Anders Søgaard. 2025. [How do multilingual language models remember facts?](#) In *Findings of the Association for Computational Linguistics, ACL 2025, Vienna, Austria, July 27 - August 1, 2025*, pages 16052–16106. Association for Computational Linguistics.

Mor Geva, Jasmijn Bastings, Katja Filippova, and Amir Globerson. 2023. [Dissecting recall of factual associations in auto-regressive language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 12216–12235. Association for Computational Linguistics.

Asma Ghandeharioun, Avi Caciularu, Adam Pearce, Lucas Dixon, and Mor Geva. 2024. [Patchscopes: A unifying framework for inspecting hidden representations of language models](#). In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.

Wentao Hu, Wengyu Zhang, Yiyang Jiang, Chen Jason Zhang, Xiaoyong Wei, and Qing Li. 2025. [Removal of hallucination on hallucination: Debate-augmented RAG](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2025, Vienna, Austria, July 27 - August 1, 2025*, pages 15839–15853. Association for Computational Linguistics.

- Lei Huang, Xiaocheng Feng, Weitao Ma, Yuchun Fan, Xiachong Feng, Yuxuan Gu, Yangfan Ye, Liang Zhao, Weihong Zhong, Baoxin Wang, Dayong Wu, Guoping Hu, Lingpeng Kong, Tong Xiao, Ting Liu, and Bing Qin. 2025. [Alleviating hallucinations from knowledge misalignment in large language models via selective abstention learning](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2025, Vienna, Austria, July 27 - August 1, 2025, pages 24564–24579. Association for Computational Linguistics.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *Preprint*, arXiv:2310.06825.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. [TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611, Vancouver, Canada. Association for Computational Linguistics.
- Saurav Kadavath, Tom Conerly, Amanda Askell, T. J. Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zachary Dodds, Nova Dassarma, Eli Tran-Johnson, Scott Johnston, Sheer El-Showk, Andy Jones, Nelson Elhage, Tristan Hume, Anna Chen, Yuntao Bai, Sam Bowman, Stanislav Fort, and 17 others. 2022. [Language models \(mostly\) know what they know](#). *ArXiv*, abs/2207.05221.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural questions: A benchmark for question answering research](#). *Transactions of the Association for Computational Linguistics*, 7:452–466.
- Kenneth Li, Oam Patel, Fernanda B. Vi  gas, Hanspeter Pfister, and Martin Wattenberg. 2023. [Inference-time intervention: Eliciting truthful answers from a language model](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Wen Luo, Tianshu Shen, Wei Li, Guangyue Peng, Richeng Xuan, Houfeng Wang, and Xi Yang. 2024. [Halludial: A large-scale benchmark for automatic dialogue-level hallucination evaluation](#). *Preprint*, arXiv:2406.07070.
- Wen Luo, Feifan Song, Wei Li, Guangyue Peng, Shao-hang Wei, and Houfeng Wang. 2025. [Odysseus navigates the sirens’ song: Dynamic focus decoding for factual and diverse open-ended text generation](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 27200–27218, Vienna, Austria. Association for Computational Linguistics.
- Alex Mullen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. [When not to trust language models: Investigating effectiveness of parametric and non-parametric memories](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9802–9822, Toronto, Canada. Association for Computational Linguistics.
- Paul Michel, Omer Levy, and Graham Neubig. 2019. Are sixteen heads really better than one? *Advances in neural information processing systems*, 32.
- Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen-tau Yih, Pang Wei Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. [Factscore: Fine-grained atomic evaluation of factual precision in long form text generation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 12076–12100. Association for Computational Linguistics.
- Mengjia Niu, Hamed Haddadi, and Guansong Pang. 2025. Robust hallucination detection in llms via adaptive token selection. *arXiv preprint arXiv:2504.07863*.
- Hadas Orgad, Michael Toker, Zorik Gekhman, Roi Reichart, Idan Szpektor, Hadas Kotek, and Yonatan Belinkov. 2025. [LLMs know more than they show: On the intrinsic representation of LLM hallucinations](#). In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net.
- Artidoro Pagnoni, Vidhisha Balachandran, and Yulia Tsvetkov. 2021. Understanding factuality in abstractive summarization with frank: A benchmark for factuality metrics. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4812–4829.
- Chen Qian, Dongrui Liu, Haochen Wen, Zhen Bai, Yong Liu, and Jing Shao. 2025. Demystifying reasoning dynamics with mutual information: Thinking tokens are information peaks in llm reasoning. *arXiv preprint arXiv:2506.02867*.
- Zhengliang Shi, Shuo Zhang, Weiwei Sun, Shen Gao, Pengjie Ren, Zhumin Chen, and Zhaochun Ren. 2024. [Generate-then-ground in retrieval-augmented generation for multi-hop question answering](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2024, Bangkok, Thailand, August 11-16,

- 2024, pages 7339–7353. Association for Computational Linguistics.
- Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. 2014. [Deep inside convolutional networks: Visualising image classification models and saliency maps](#). In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Workshop Track Proceedings*.
- Yuanhe Tian, Ruyi Gan, Yan Song, Jiaying Zhang, and Yongdong Zhang. 2024. [Chimed-gpt: A chinese medical large language model with full training regime and better alignment to human preferences](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 7156–7173. Association for Computational Linguistics.
- Eric Todd, Millicent L. Li, Arnab Sen Sharma, Aaron Mueller, Byron C. Wallace, and David Bau. 2024. [Function vectors in large language models](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Lean Wang, Lei Li, Damai Dai, Deli Chen, Hao Zhou, Fandong Meng, Jie Zhou, and Xu Sun. 2023. Label words are anchors: An information flow perspective for understanding in-context learning. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9840–9855.
- Wenhao Wu, Yizhong Wang, Guangxuan Xiao, Hao Peng, and Yao Fu. 2025. [Retrieval head mechanistically explains long-context factuality](#). In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net.
- Boyang Xue, Fei Mi, Qi Zhu, Hongru Wang, Rui Wang, Sheng Wang, Erxin Yu, Xuming Hu, and Kam-Fai Wong. 2025a. [UAlign: Leveraging uncertainty estimations for factuality alignment on large language models](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6002–6024, Vienna, Austria. Association for Computational Linguistics.
- Yihao Xue, Kristjan Greenewald, Youssef Mroueh, and Baharan Mirzasoleiman. 2025b. Verify when uncertain: Beyond self-consistency in black box hallucination detection. *arXiv preprint arXiv:2502.15845*.
- An Yang, Anfeng 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. 2025. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [HotpotQA: A dataset for diverse, explainable multi-hop question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.
- Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, and 1 others. 2025. Siren’s song in the ai ocean: A survey on hallucination in large language models. *Computational Linguistics*, pages 1–46.

A LLM Usage

In this work, we employ LLMs solely for language refinement to enhance clarity and explanatory quality. All content has been carefully verified for factual accuracy, and the authors take full responsibility for the entire manuscript. The core ideas, experimental design, and methodological framework are conceived and developed independently by the authors, without the use of LLMs.

B Implementation Details

B.1 Identifying Exact Question and Answer Tokens

To locate the exact question and answer tokens within a QA pair, we prompt GPT-4o (version gpt-4o_2024-11-20) to identify the precise positions of the core frame elements. The instruction templates are presented in Tables 5 and 6. A token is considered an *exact question* or *exact answer* if and only if it constitutes a valid substring of the corresponding question or answer. To mitigate potential biases, each example is prompted at most five times, and only successfully extracted instances are retained for downstream analysis. Prior work (Orgad et al., 2025) has shown that LLMs can accurately identify exact answer tokens, typically achieving over 95% accuracy. In addition, we manually verified GPT-4o’s identification quality in our setting. Specifically, it achieves 99.92%, 95.83%, and 96.62% accuracy on exact subject tokens, exact property tokens, and exact answer tokens, respectively. Furthermore, we also explore alternative configurations without the use of exact tokens to ensure the robustness of our findings (see Section B.2).

B.2 Probing Implementation Details

We investigate multiple probing configurations. For token selection, we consider three types of tokens: (1) the final token of the answer, which is the most commonly adopted choice in prior work due to its global receptive field under attention (Chen et al., 2024); (2) the token immediately preceding the exact answer span; and (3) the final token within the exact answer span. For activation extraction, we obtain representations from either (1) the output of each attention sublayer or (2) the output of the final multi-layer perceptron (MLP) in each transformer layer. Across all configurations, our experimental results exhibit consistent trends, indicating that the observed findings are robust to

these design choices. For the probing classifier, we follow standard practice (Chen et al., 2024; Orgad et al., 2025) and employ a logistic regression model implemented in `scikit-learn`.

B.3 Models

Our analysis covers a diverse collection of 12 LLMs that vary in both scale and architectural design. Specifically, we consider three categories: (1) *base models*, including Llama-3.2-1B (Grattafiori et al., 2024), Llama-3.2-3B, Llama-3-8B, Llama-3-70B, Mistral-7B-v0.1 (Jiang et al., 2023), and Mistral-7B-v0.3; (2) *instruction-tuned models*, including Llama-3.2-3B-Instruct, Llama-3-8B-Instruct, Mistral-7B-Instruct-v0.1, and Mistral-7B-Instruct-v0.3; and (3) *reasoning-oriented models*, namely Qwen3-8B (Yang et al., 2025) and Qwen3-32B.

B.4 Datasets

We consider four widely used question-answering datasets: PopQA (Mallen et al., 2023), TriviaQA (Joshi et al., 2017), HotpotQA (Yang et al., 2018), and Natural Questions (Kwiatkowski et al., 2019).

PopQA is an open-domain question-answering dataset that emphasizes entity-centric factual knowledge with a long-tail distribution. It is designed to probe LLMs’ ability to memorize less frequent facts, highlighting limitations in parametric knowledge.

TriviaQA is a reading comprehension dataset constructed by pairing trivia questions authored independently of evidence documents. The questions are often complex, requiring multi-sentence reasoning, and exhibit substantial lexical and syntactic variability.

HotpotQA is a challenging multi-hop question-answering dataset that requires reasoning. It includes diverse question types—span extraction, yes/no, and novel comparison questions—along with sentence-level supporting fact annotations, promoting the development of explainable QA systems.

Natural Questions is an open-domain dataset consisting of real, anonymized questions from Google search queries. Each question is annotated with both a long answer (paragraph or section) and a short answer (span or yes/no), or marked as null when no answer is available. Due to computational constraints, we randomly sample 2,000 training samples and 2,000 test samples for each dataset.

B.5 Implementation Details of Baselines

In our experiments regarding applications, we compare our proposed methods against several internal-based baselines for hallucination detection. These baselines leverage the LLM’s internal signals, such as output probabilities, logits, and hidden representations, without relying on external resources. Below, we detail the implementation of each baseline.

P(True) P(True) (Kadavath et al., 2022) exploits the LLM’s self-awareness of its knowledge boundaries by prompting the model to assess the correctness of its own generated answer. Specifically, for each question-answer pair (q_i, \hat{y}_i^f) , we prompt the LLM with a template that asks it to evaluate whether its answer is factually correct. Following Kadavath et al. (2022), the prompt template is shown in Table 4.

Question: {Here is the question}
Possible answer: {Here is the answer}
Is the possible answer:
(A) True
(B) False
The possible answer is:

Table 4: Prompt template used for the P(True) baseline.

Logits-based Baselines The logits-based baselines utilize the raw logits produced by the LLM during the generation of the exact answer tokens. Let $\hat{y}_{i,E_A}^f = [t_1, t_2, \dots, t_m]$ represent the sequence of exact answer tokens for a given question-answer pair, where m is the number of exact answer tokens. For each token t_j (where $j \in \{1, \dots, m\}$), the LLM produces a logit vector $L_j \in \mathbb{R}^V$, where V is the vocabulary size, and the logit for the generated token t_j is denoted $L_j[t_j]$. The logits-based metrics are defined as follows:

- **Logits-mean:** The average of the logits across all exact answer tokens:

$$\text{Logits-mean} = \frac{1}{m} \sum_{j=1}^m L_j[t_j] \quad (6)$$

- **Logits-max:** The maximum logit value among the exact answer tokens:

$$\text{Logits-max} = \max_{j \in \{1, \dots, m\}} L_j[t_j] \quad (7)$$

- **Logits-min:** The minimum logit value among the exact answer tokens:

$$\text{Logits-min} = \min_{j \in \{1, \dots, m\}} L_j[t_j] \quad (8)$$

These metrics serve as proxies for the model’s confidence in the generated answer, with lower logit values potentially indicating uncertainty or hallucination.

Scores-based Baselines The scores-based baselines are derived from the softmax probabilities of the exact answer tokens. Using the same notation as above, for each exact answer token t_j , the softmax probability is computed as:

$$p_j[t_j] = \frac{\exp(L_j[t_j])}{\sum_{k=1}^V \exp(L_j[k])} \quad (9)$$

where $L_j[k]$ is the logit for the k -th token in the vocabulary. The scores-based metrics are defined as follows:

- **Scores-mean:** The average of the softmax probabilities across all exact answer tokens:

$$\text{Scores-mean} = \frac{1}{m} \sum_{j=1}^m p_j[t_j] \quad (10)$$

- **Scores-max:** The maximum softmax probability among the exact answer tokens:

$$\text{Scores-max} = \max_{j \in \{1, \dots, m\}} p_j[t_j] \quad (11)$$

- **Scores-min:** The minimum softmax probability among the exact answer tokens:

$$\text{Scores-min} = \min_{j \in \{1, \dots, m\}} p_j[t_j] \quad (12)$$

These probabilities provide a normalized measure of the model’s confidence, bounded between 0 and 1, with lower values potentially indicating a higher likelihood of hallucination.

Probing Baseline The probing baseline follows the standard approach described in Chen et al. (2024); Orgad et al. (2025). A linear classifier is trained on the hidden representations of the last exact answer token from the best-performing layer. The training and evaluation data for the probing classifier are constructed following the procedure described in Appendix B.4. The classifier is implemented using scikit-learn with default hyperparameters, consistent with the probing setup

described in Appendix B.2. The probing baseline serves as a direct comparison to our proposed applications, as it relies on the same type of internal signals but does not account for the heterogeneity of truthfulness encoding pathways.

B.6 Implementation Details of MoP and PR

Model Backbone and Hidden Representations

All experiments use the same base LLM as in the main paper. Hidden representations $\mathbf{h}^{l^*}(x)$ are extracted from the best-performing layer l^* determined on a held-out validation split.

Mixture-of-Probes (MoP) Similar to Appendix B.5, the two expert probes p_Q and p_A are implemented using scikit-learn with default hyperparameters, consistent with the probing setup described in Appendix B.2. The gating network is directly from the self-awareness probe described in Section 4.2. The training and evaluation data for the probing classifier are the same as Appendix B.5. The proposed MoP framework requires no additional retraining: we directly combine the two expert probes with the pathway-discrimination classifier described in Section 4.2 and perform inference without further parameter updates.

Pathway Reweighting (PR) The training and evaluation data used for the probing classifier are identical to those described in Appendix B.5. For each Transformer layer $l \leq l^*$, we introduce two learnable scalars α_Q^l and α_A^l for every attention head. These parameters, together with the probe parameters, are optimized using the Adam optimizer with a learning rate of 1×10^{-2} , $\beta_1 = 0.9$, and $\beta_2 = 0.999$. Training is conducted with a batch size of 512 for 10 epochs, while all original LLM parameters remain frozen.

You are given a factual open-domain question-answer pair.

Your task is to identify:

1. Core Entity (c) - the known specific entity in the question that the answer is about (a person, place, organization, or other proper noun).
2. Relation (r) - the minimal phrase in the question that expresses what is being asked about the core entity, using only words from the question.

Guidelines:

The core entity must be a concrete, known entity mentioned in the question, not a general category. If multiple entities appear, choose the one most central to the question—the entity the answer primarily concerns.

The relation should be the smallest meaningful span that directly connects the core entity to the answer.

Use only words from the question; do not paraphrase or add new words.

Exclude extra context, modifiers, or descriptive phrases that are not essential to defining the relationship.

For complex questions with long modifiers or embedded clauses, focus on the words that directly express the property, action, or attribute of the core entity relevant to the answer.

If you cannot confidently identify the core entity or the relation, output NO ANSWER.

Output format:

Core Entity: exact text

Relation: exact text

Example 1

Question: Who was the director of Finale?

Answer: Ken Kwapis

Core Entity: Finale

Relation: director

Example 2

Question: What film, in production between 2007 and 2009, is directed by James Cameron ("Titanic")?

Answer: Avatāra

Core Entity: James Cameron

Relation: film directed by

Example 3

Question: Which novel, written in 1925 and often cited as a classic of American literature, was authored by F. Scott Fitzgerald?

Answer: The Great Gatsby

Core Entity: F. Scott Fitzgerald

Relation: novel authored by

Question: {Here is the question}

Answer: {Hese is the answer}

Table 5: Prompt template used to locate the exact question tokens.

You are given a factual open-domain question-answer pair.
Your task is to extract the short exact answer from the long answer.

Guidelines:

Retain only the relevant tokens from the answer.
Use only words from the answer; do not paraphrase or add new words.
Extract the short answer regardless of whether the long answer is correct or incorrect.
If the short answer cannot be confidently determined, or if the long answer does not address the question, output NO ANSWER.

Example 1

Q: Which musical featured the song The Street Where You Live?
A: The song "The Street Where You Live" is from the Lerner and Loewe musical "My Fair Lady." It is one of the most famous songs from the show, and it is sung by Professor Henry Higgins as he reflects on the transformation of Eliza Doolittle and the memories they have shared together.
Exact answer: My Fair Lady

Example 2

Q: Which Swedish actress won the Best Supporting Actress Oscar for Murder on the Orient Express?
A: I'm glad you asked about a Swedish actress who won an Oscar for "Murder on the Orient Express," but I must clarify that there seems to be a misunderstanding here. No Swedish actress has won an Oscar for Best Supporting Actress for that film. The 1974 "Murder on the Orient Express" was an American production, and the cast was predominantly British and American. If you have any other questions or if there's another
Exact answer: NO ANSWER

Example 3

Q: Who invented the telephone?
A: The telephone was invented by Thomas Edison in the 19th century.
Exact answer: Thomas Edison

Q: {Here is the question}
A: {Here is the answer}
Exact answer:

Table 6: Prompt template used to locate the exact answer tokens.

You are given one factual question. Interpret it literally and think carefully.
Your task is to decide whether you can answer it correctly with high confidence based only on your internal knowledge (no tools or web). If yes, output exactly: YES. If not or uncertain, output exactly: NO. You should output one word only.

Question: {Here is the question}
Your Output:

Table 7: Prompt template used to obtain the i-don't-know rate.

C Attention Knockout

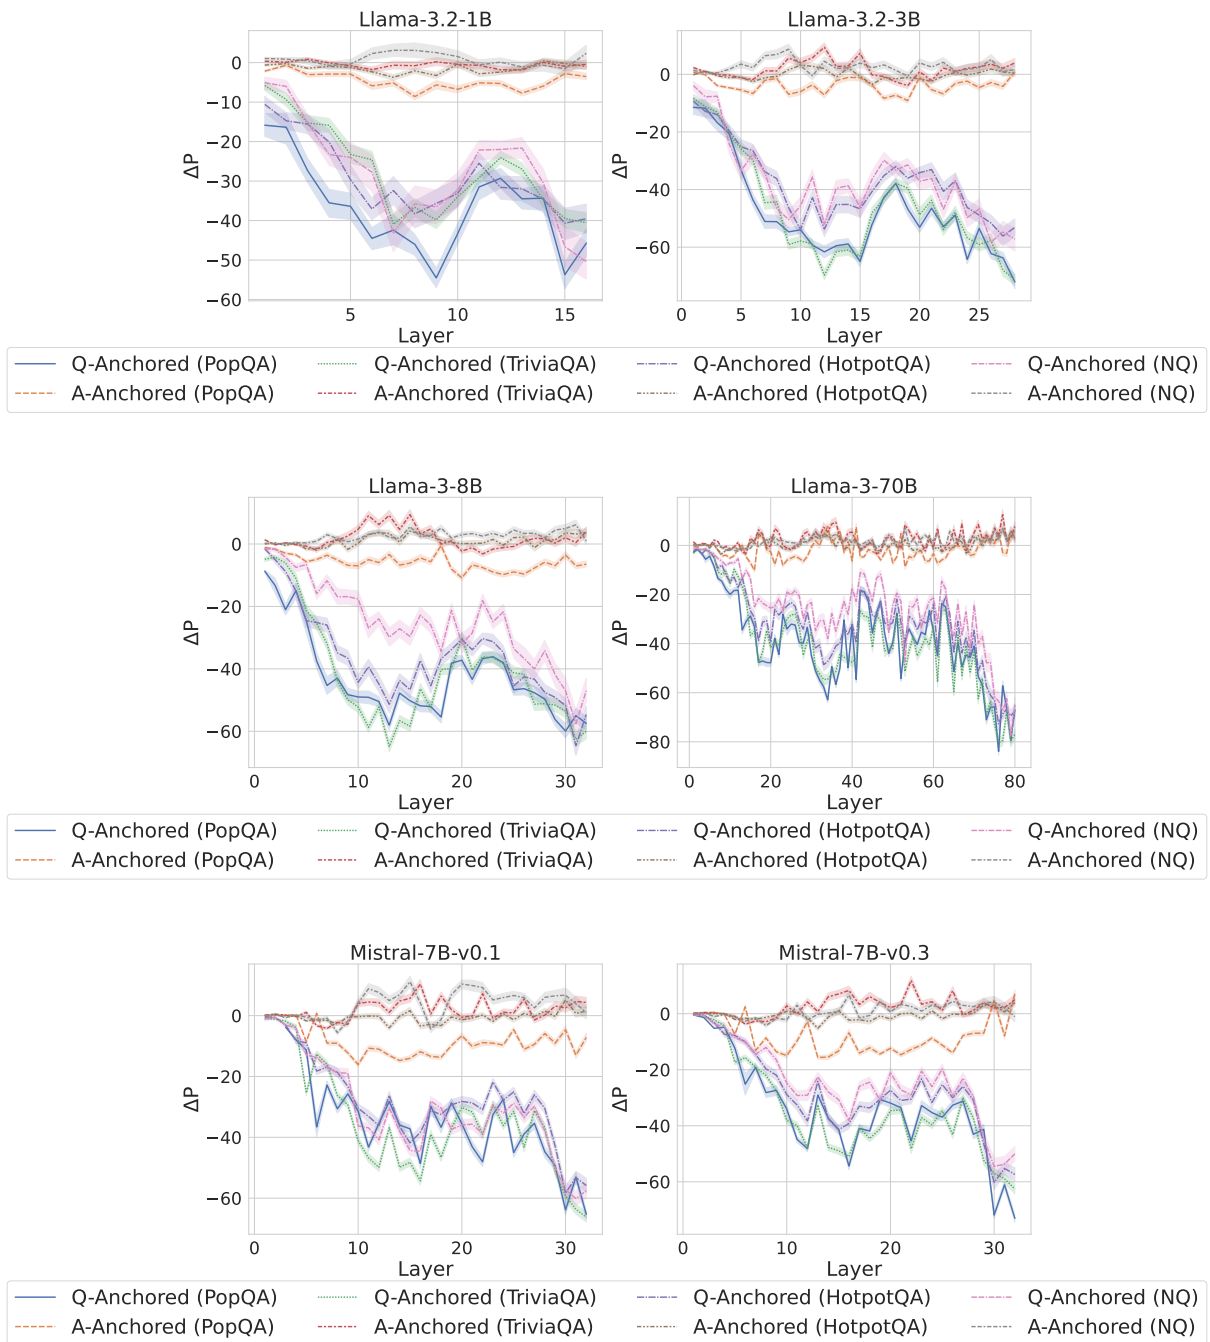


Figure 7: ΔP under attention knockout, probing attention activations of the final token.

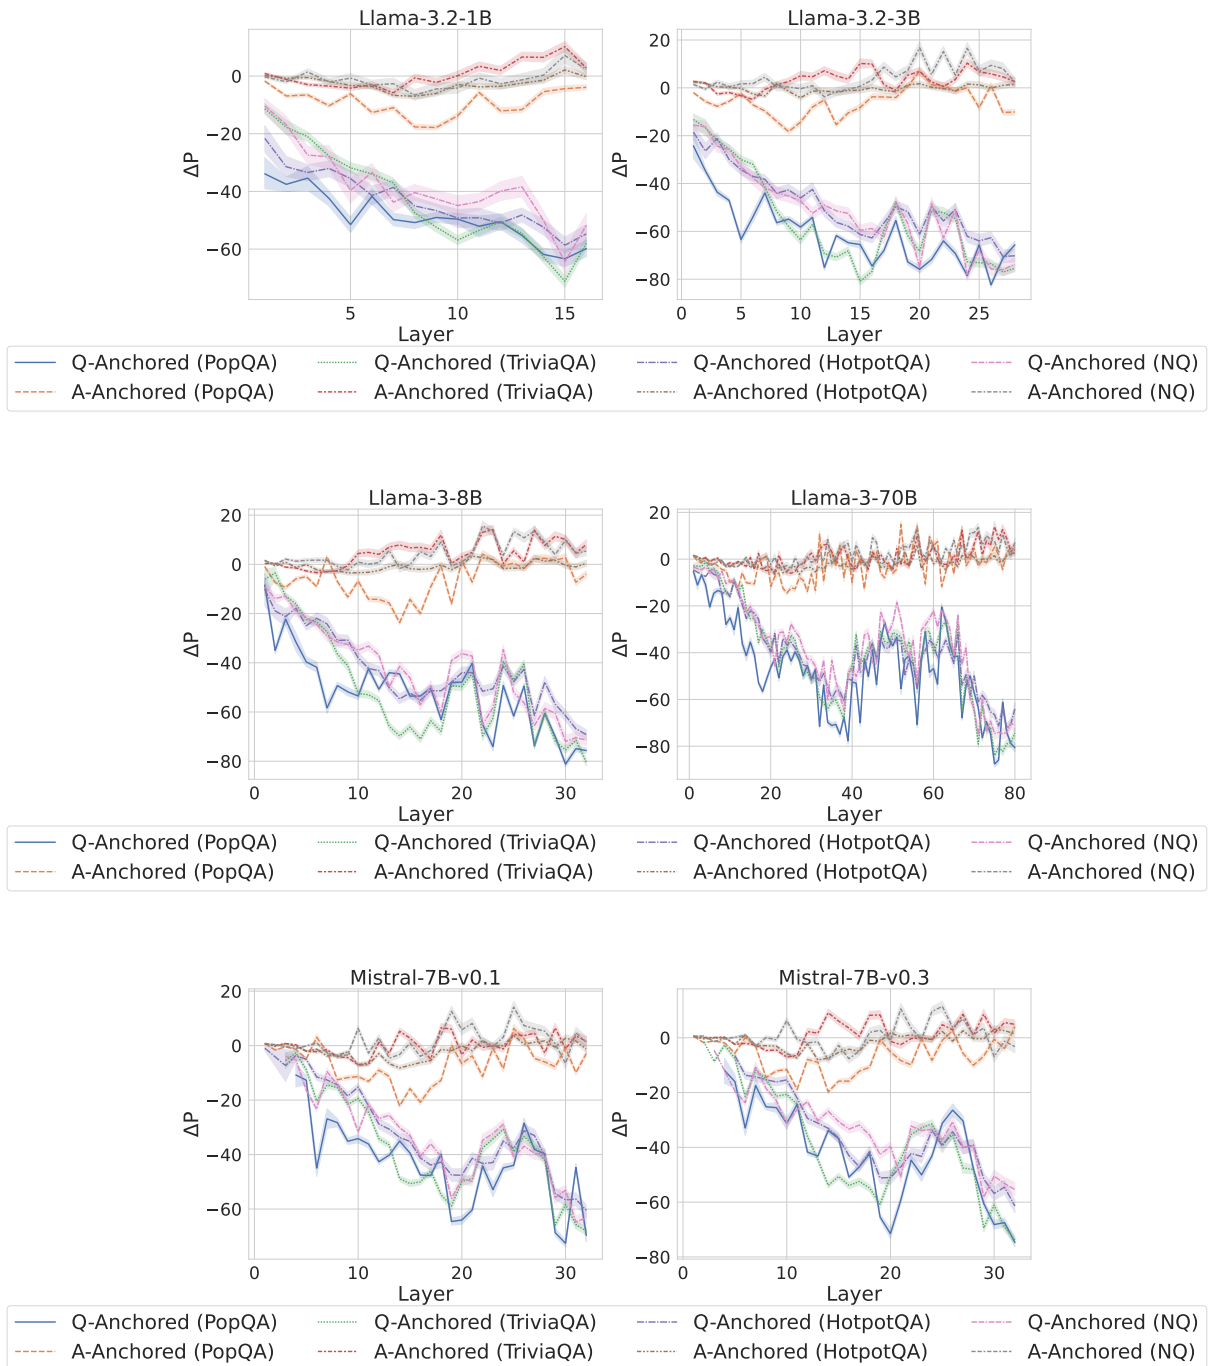


Figure 8: ΔP under attention knockout, probing attention activations of the token immediately preceding the exact answer tokens.

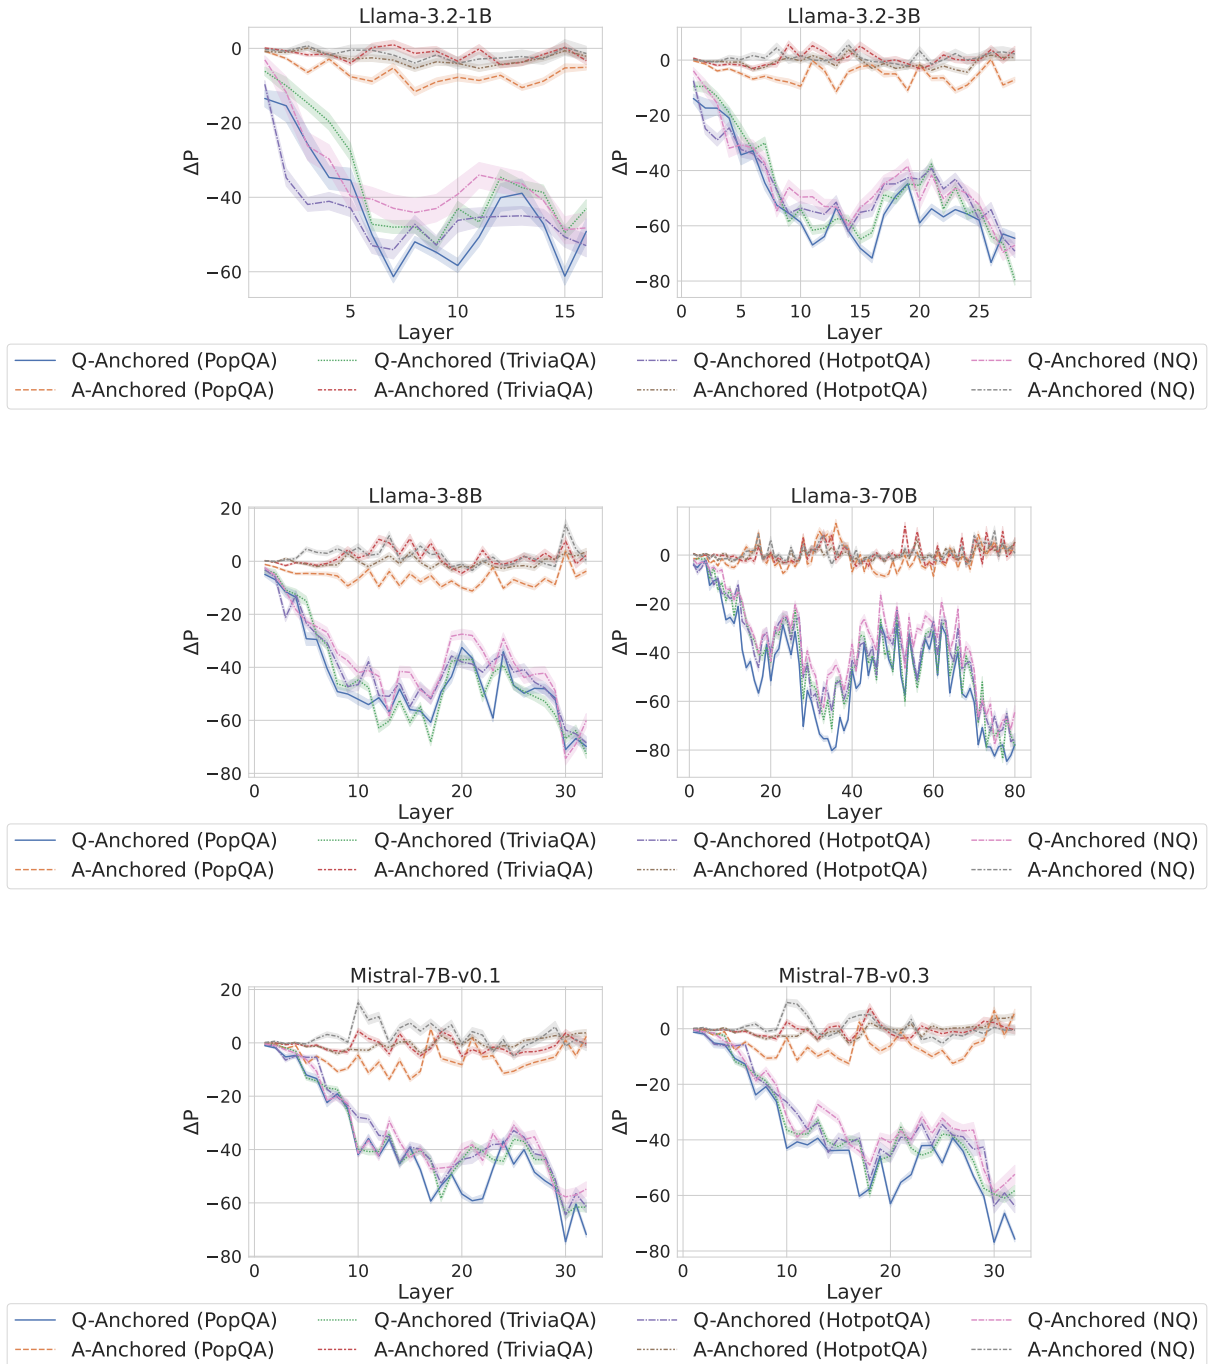


Figure 9: ΔP under attention knockout, probing attention activations of the last exact answer token.

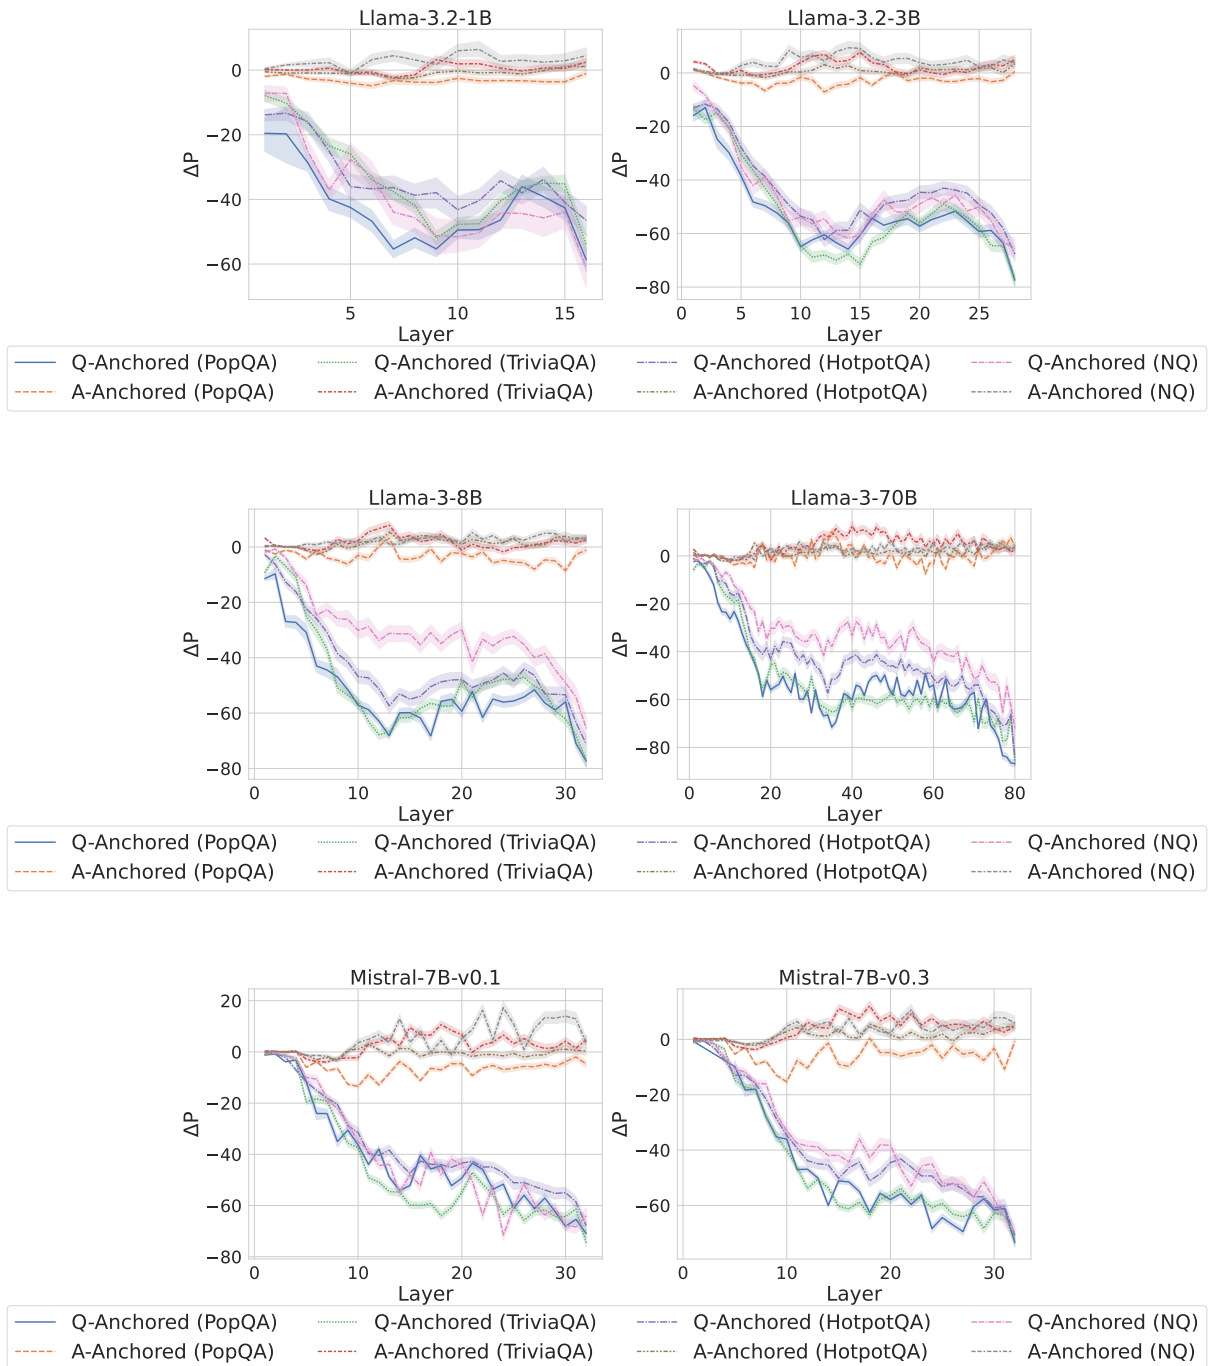


Figure 10: ΔP under attention knockout, probing mlp activations of the final token.

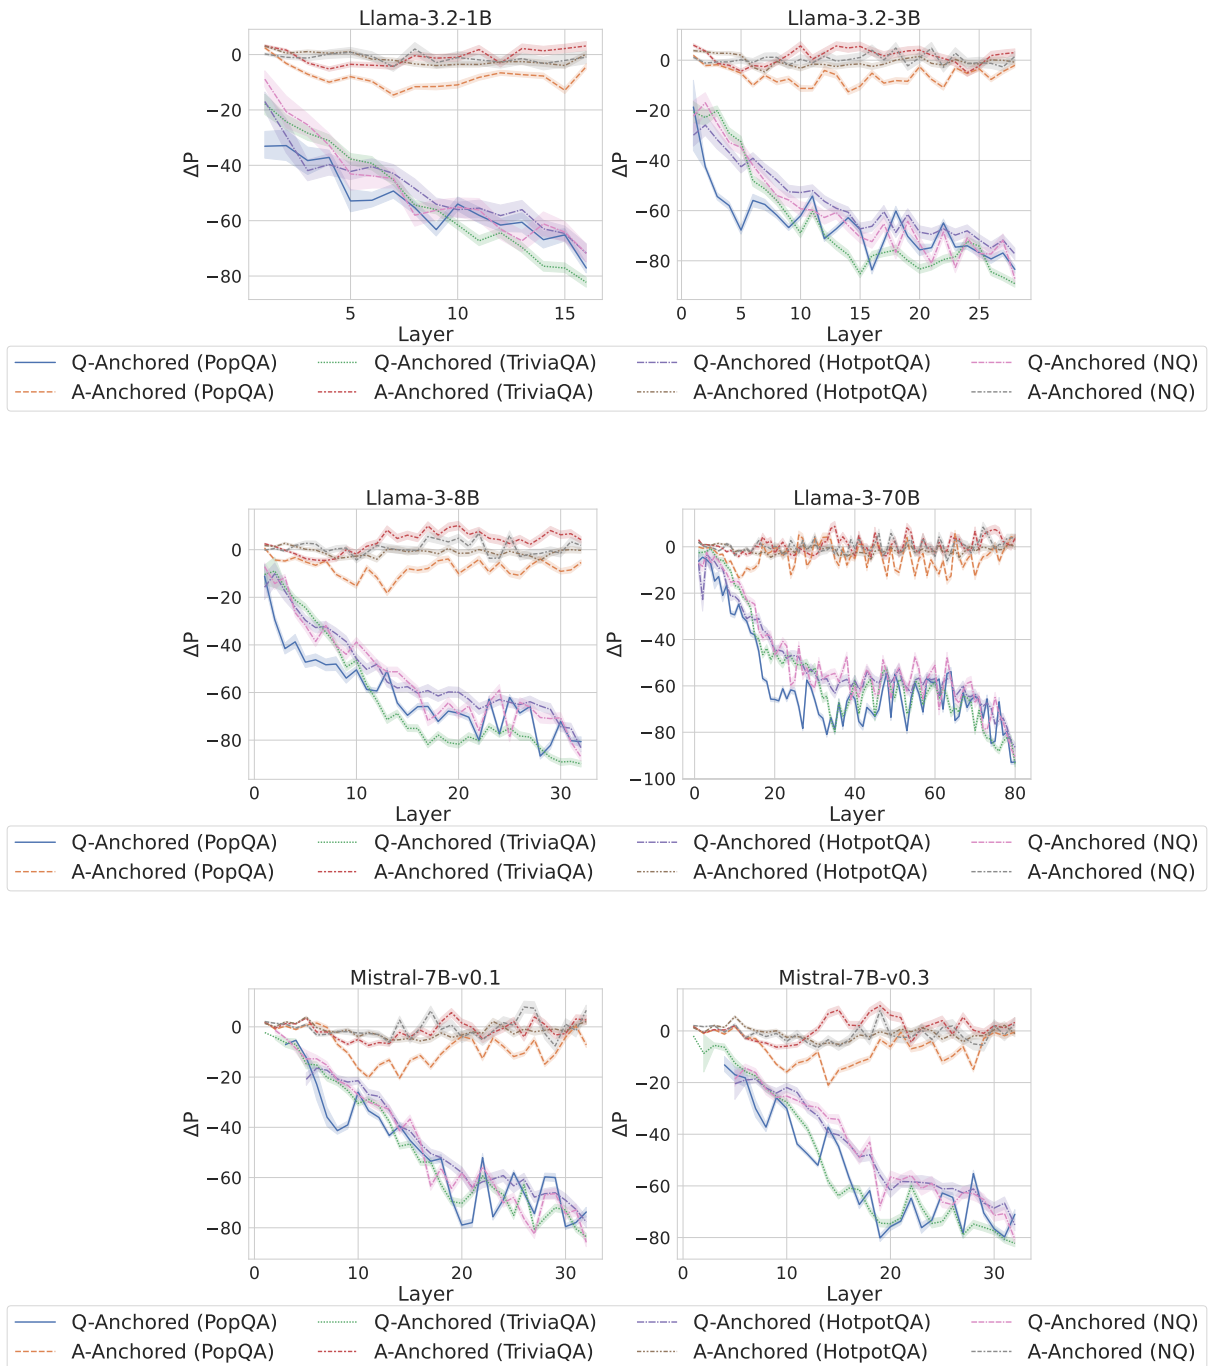


Figure 11: ΔP under attention knockout, probing mlp activations of the token immediately preceding the exact answer tokens.

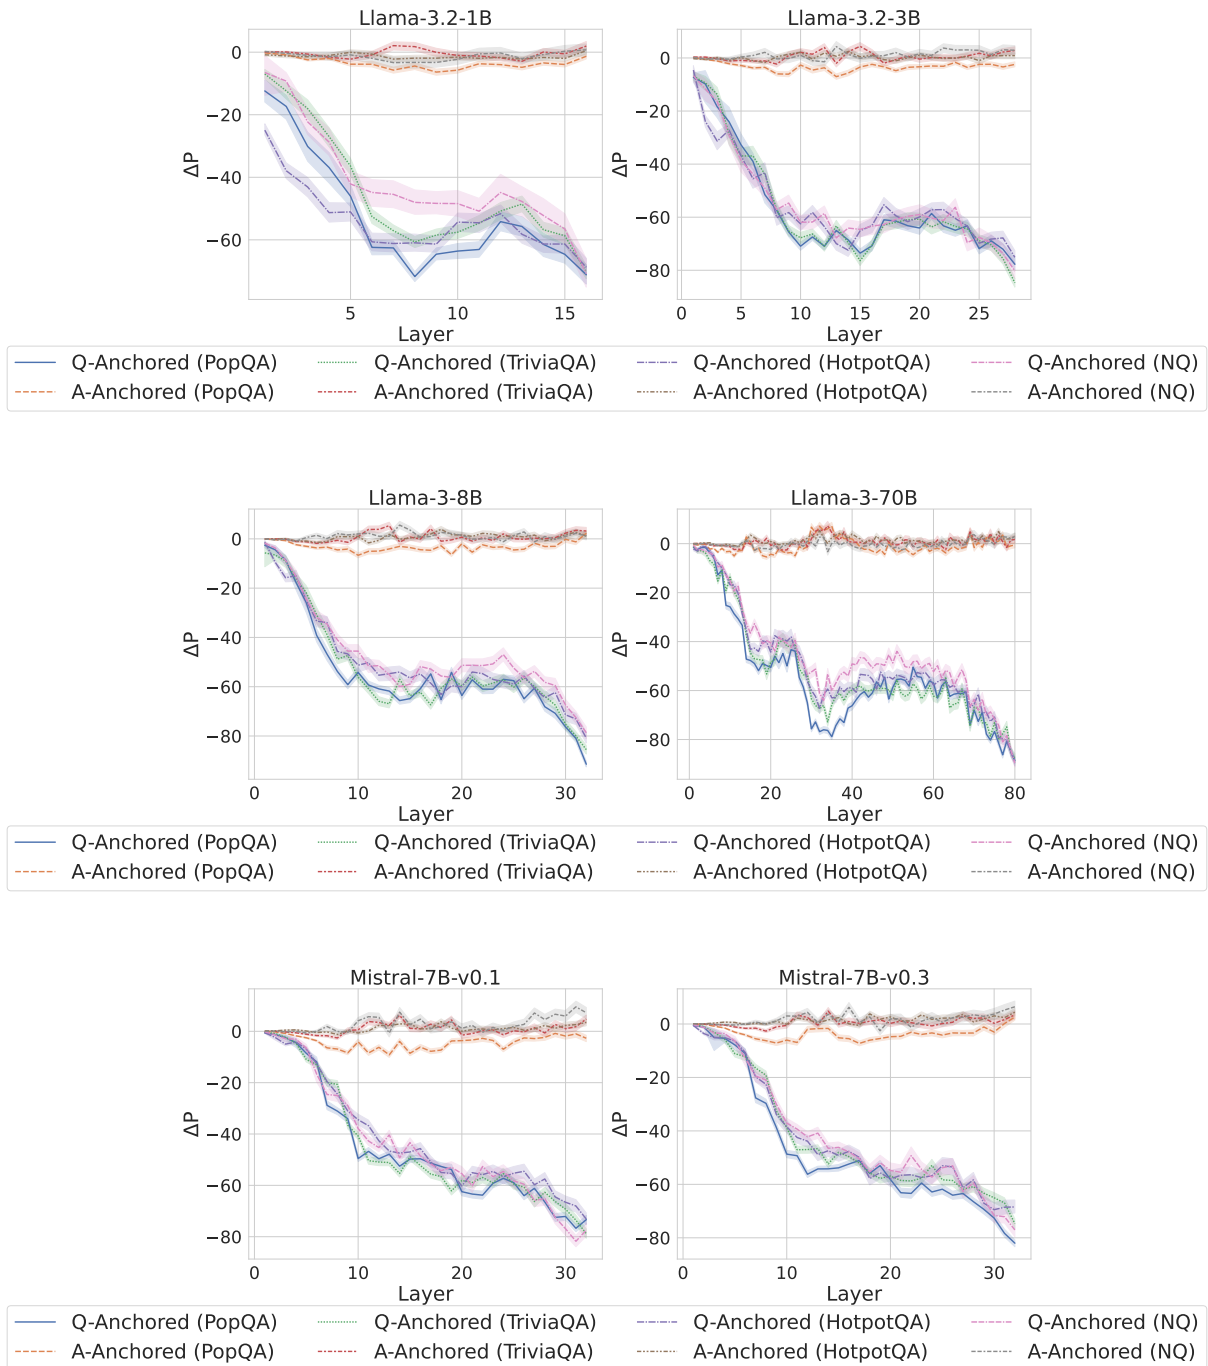


Figure 12: ΔP under attention knockout, probing mlp activations of the last exact answer token.

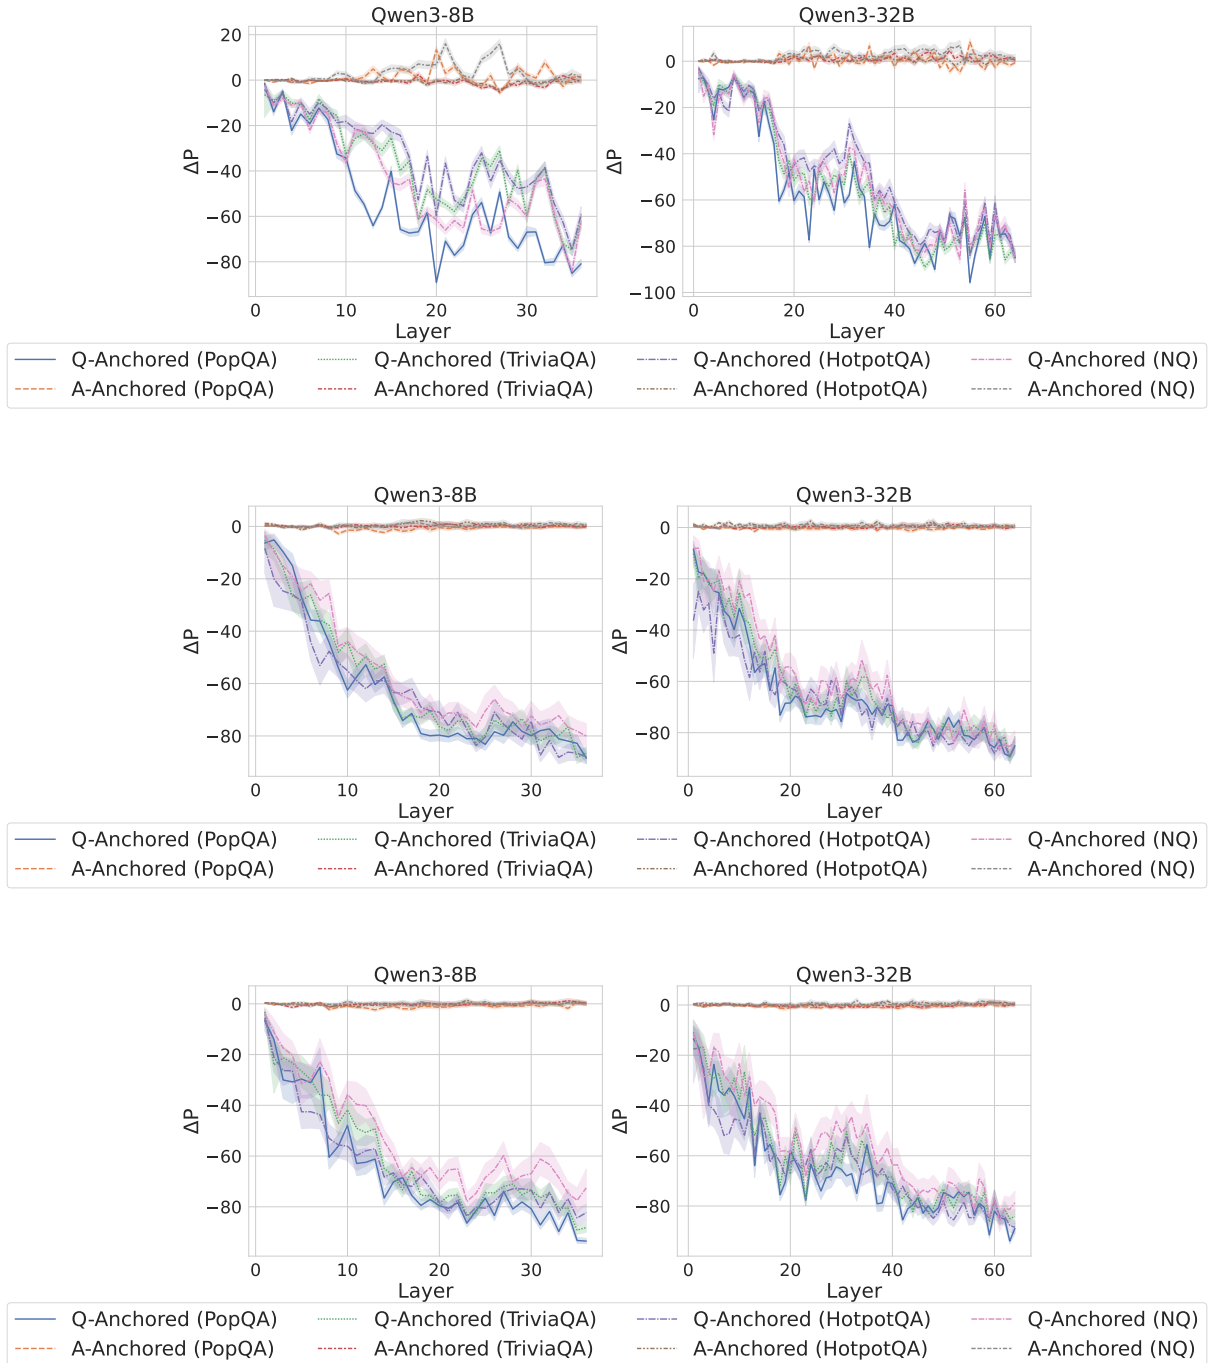


Figure 13: ΔP under attention knockout for reasoning models. Probing attention activations for the final token (top), the token immediately preceding the exact answer tokens (middle), and the last exact answer token (bottom).

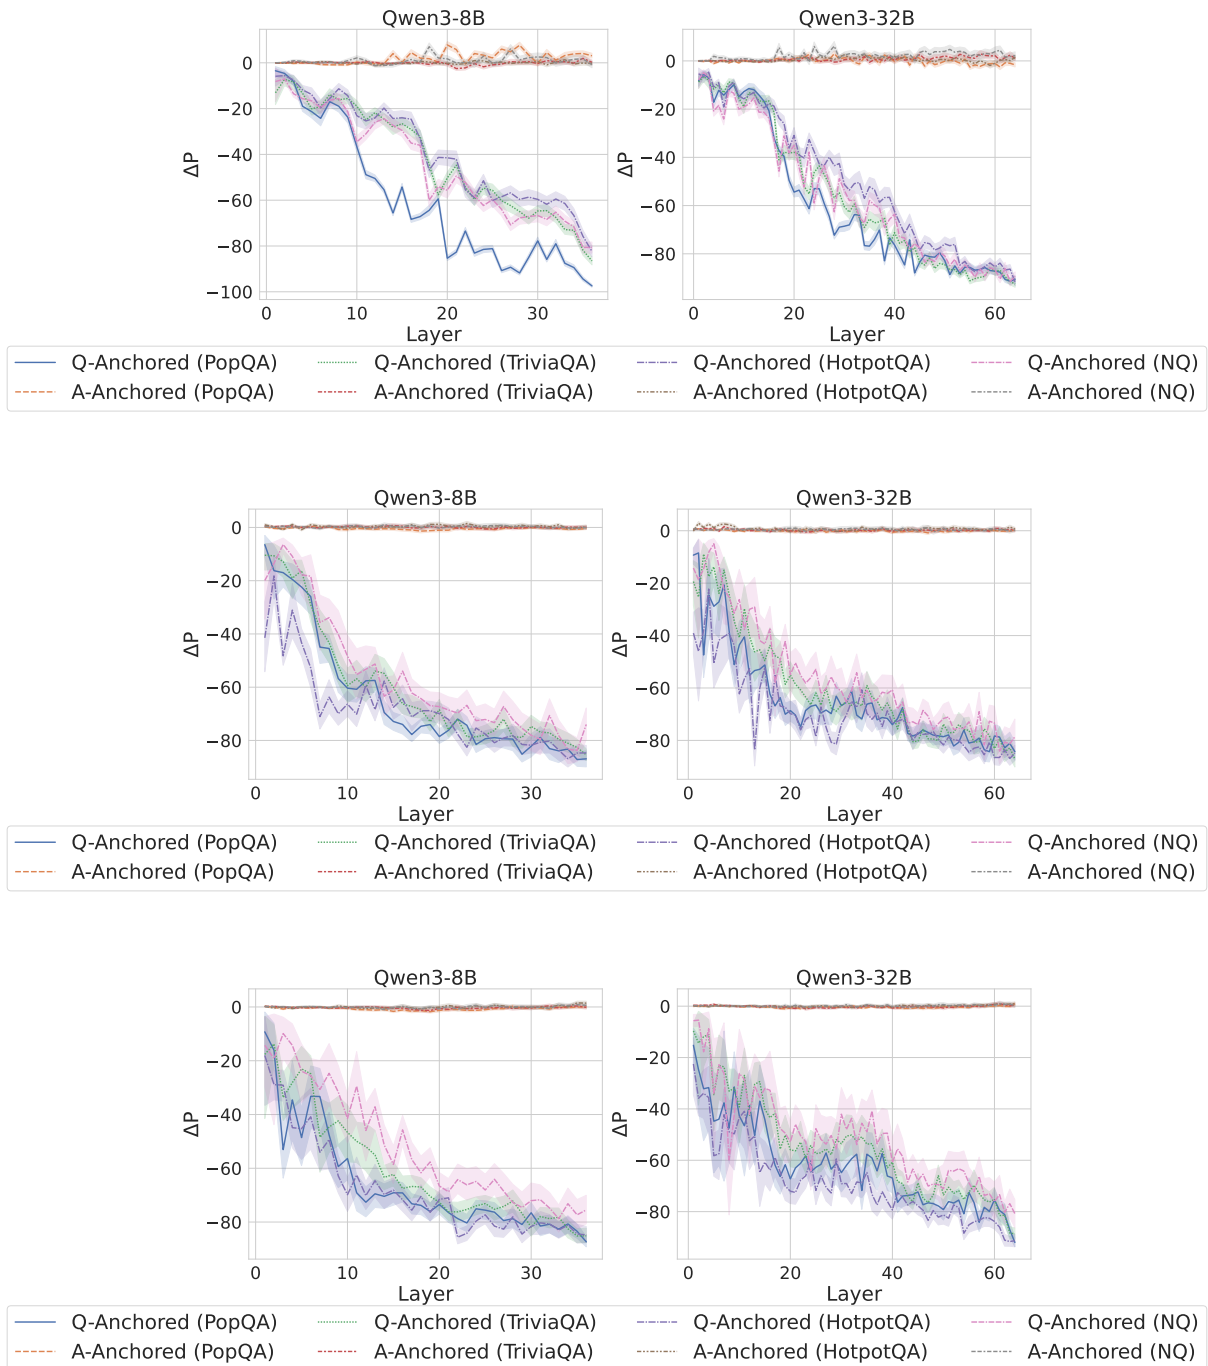


Figure 14: ΔP under attention knockout for reasoning models. Probing mlp activations for the final token (top), the token immediately preceding the exact answer tokens (middle), and the last exact answer token (bottom).

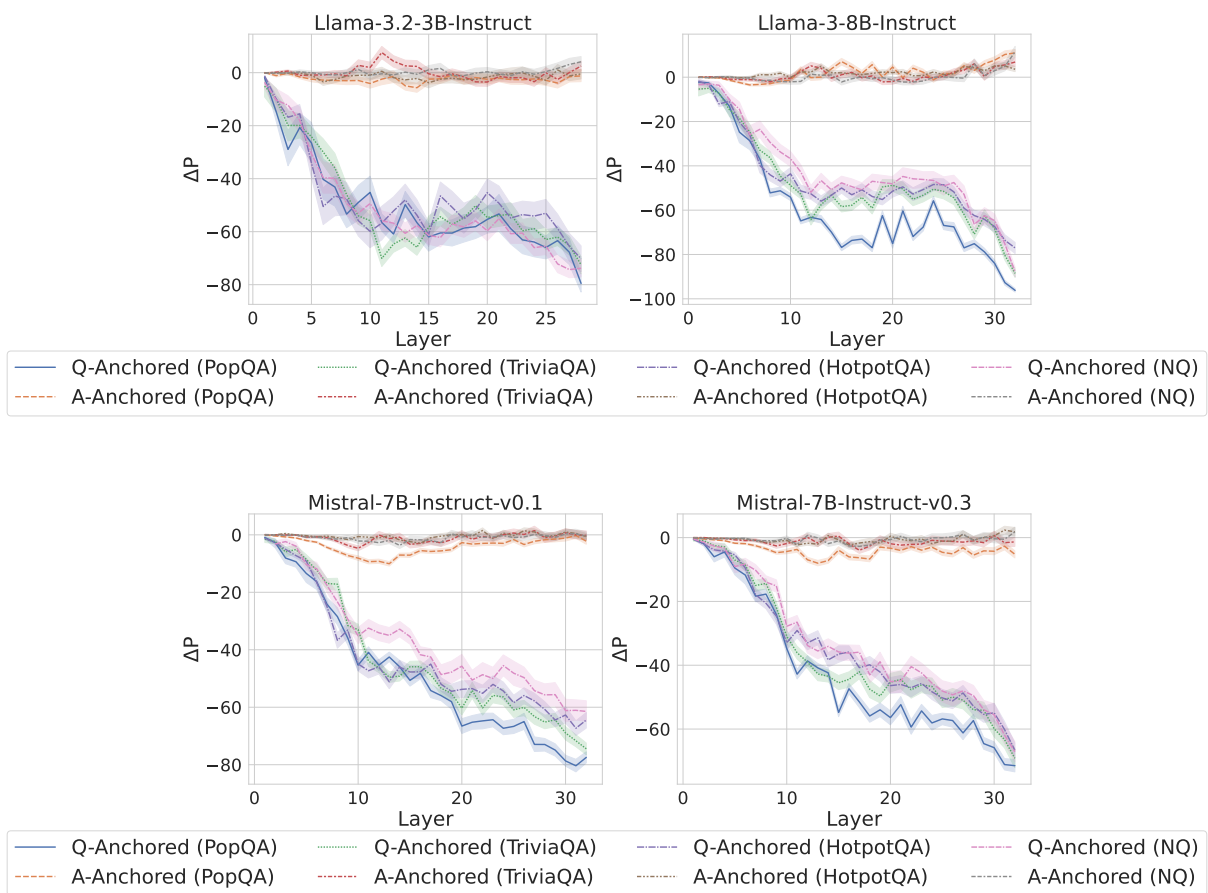


Figure 15: ΔP under attention knockout for instruct models.

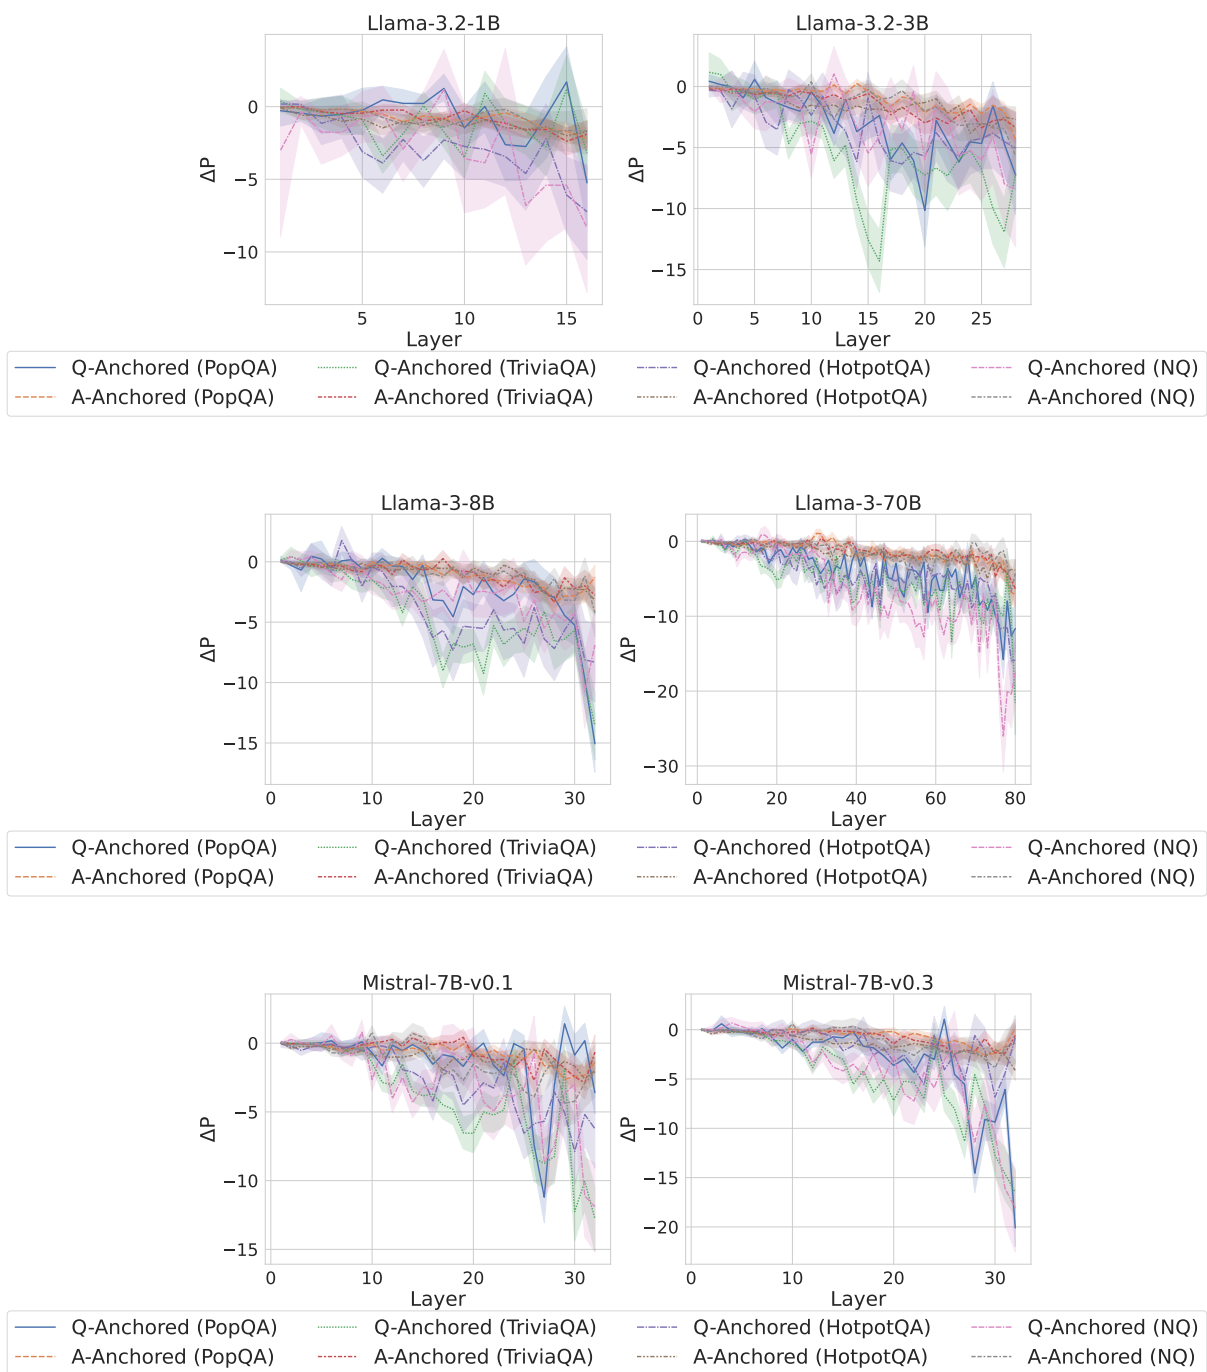


Figure 16: ΔP under attention knockout with randomly masked question tokens. Unlike selectively blocking the exact question tokens, both Q-Anchored and A-Anchored samples exhibit similar patterns, with substantially smaller probability changes when question tokens are masked at random. This suggests that exact question tokens play a critical role in conveying the semantic information of core frame elements.

D Token Patching

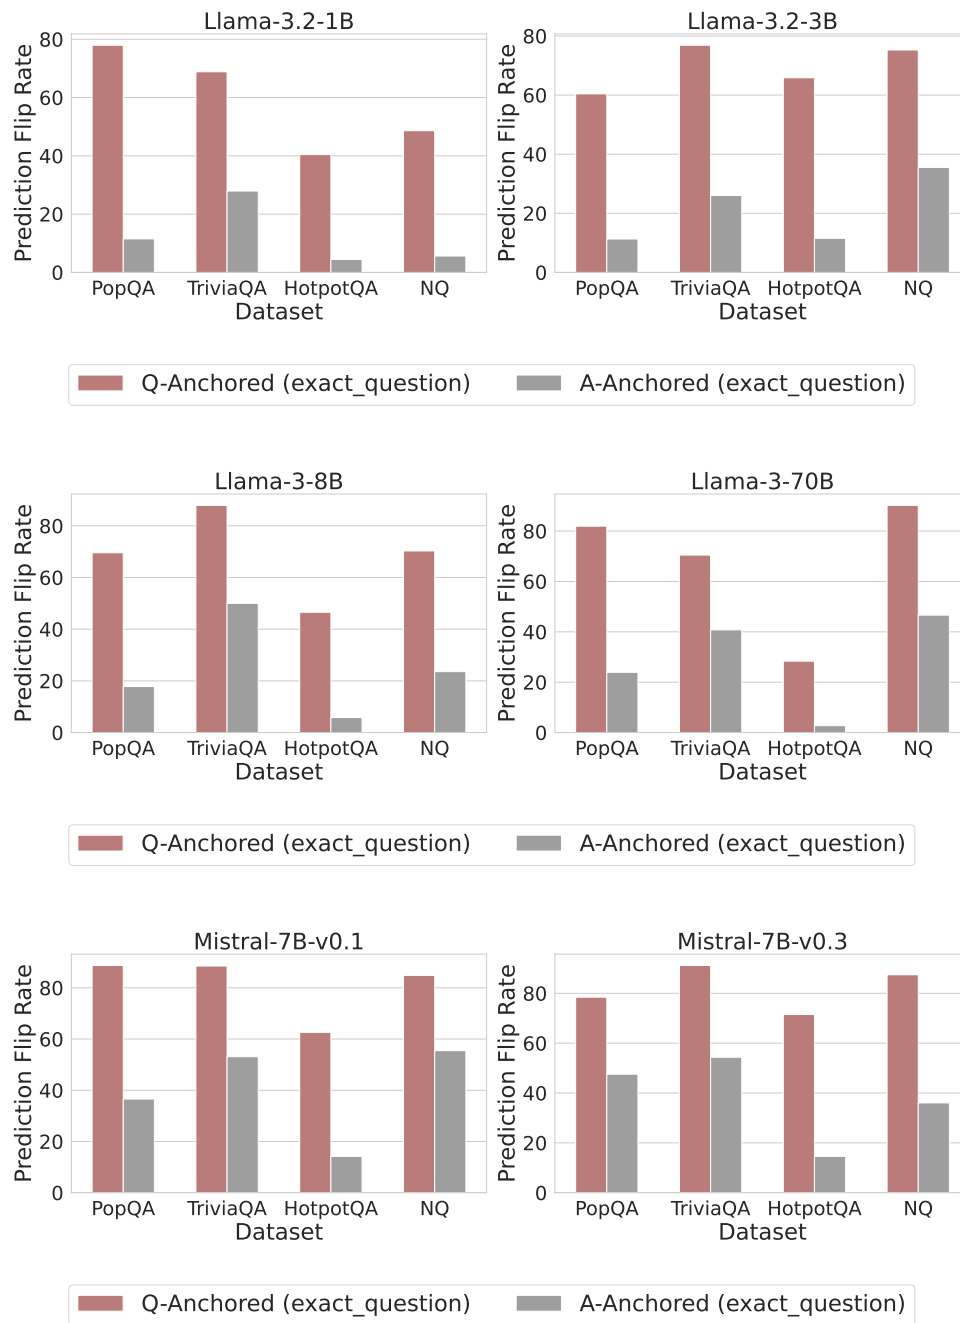


Figure 17: Prediction flip rate under token patching, probing attention activations of the final token.

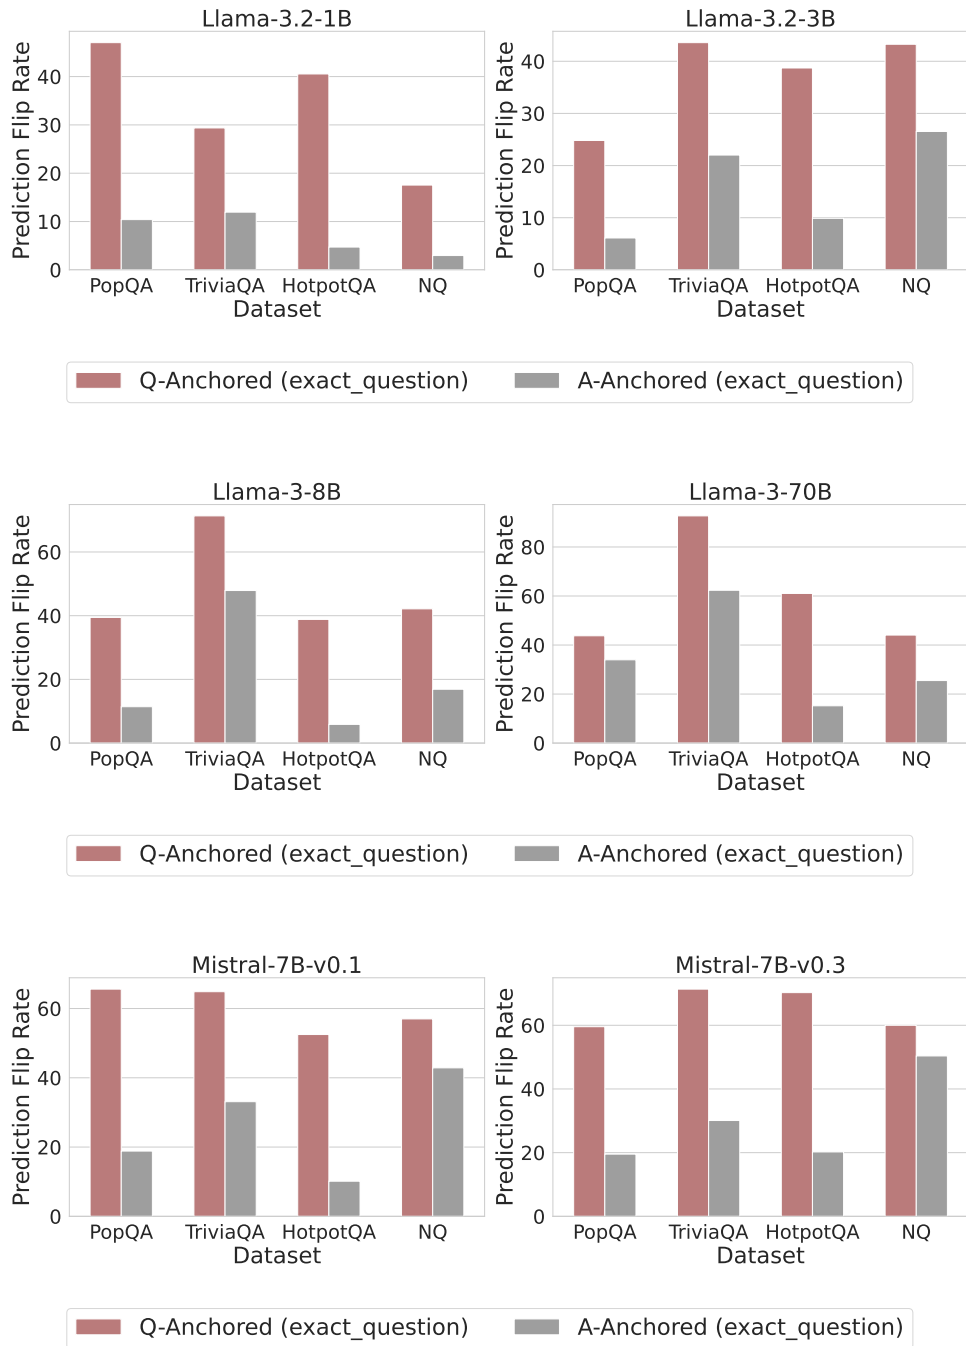


Figure 18: Prediction flip rate under token patching, probing attention activations of the token immediately preceding the exact answer tokens.

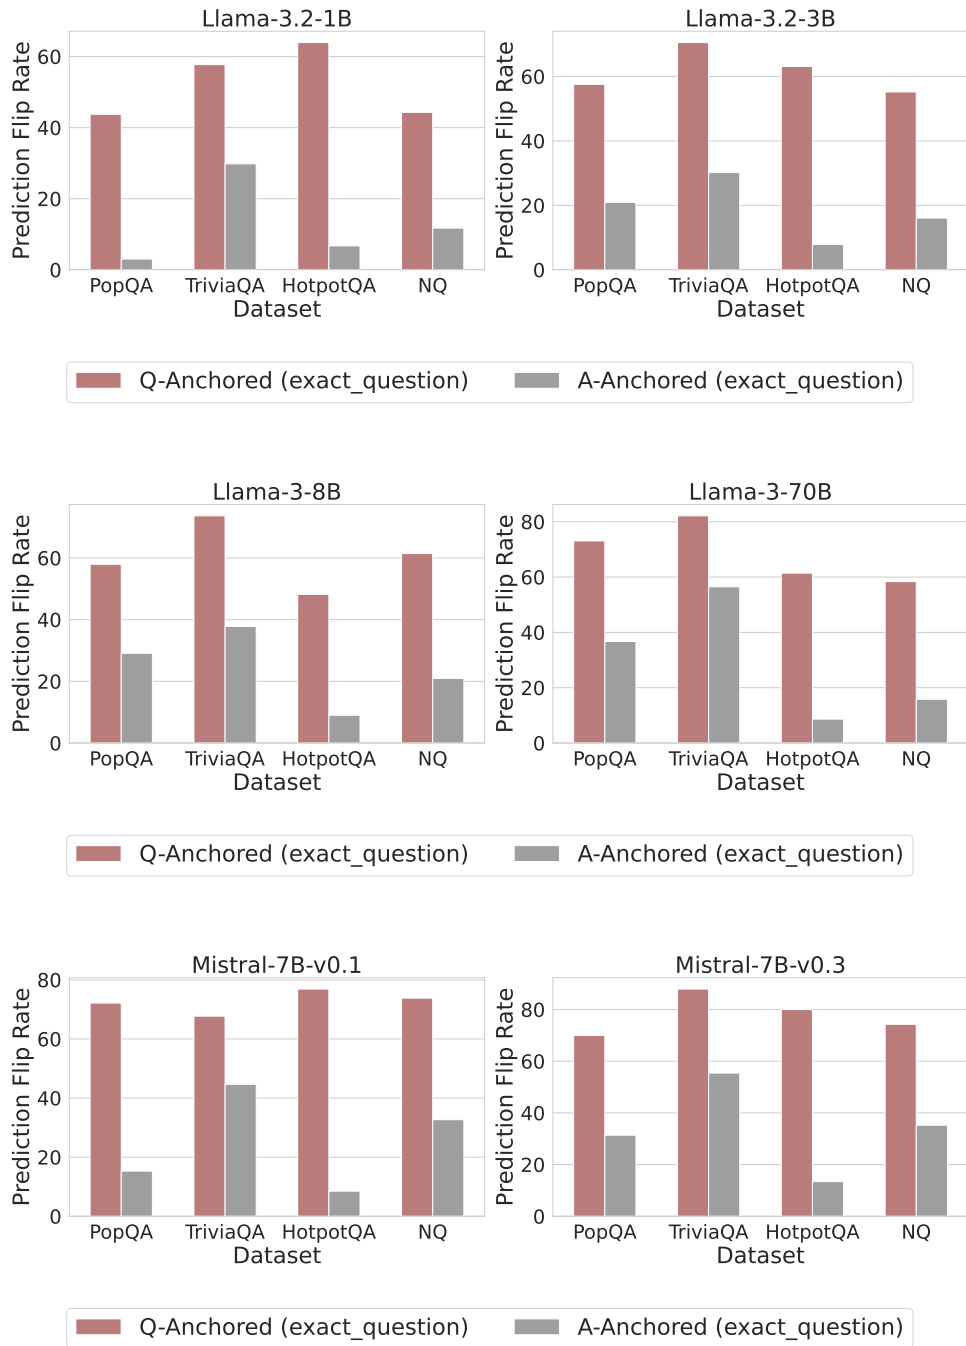


Figure 19: Prediction flip rate under token patching, probing attention activations of the last exact answer token.

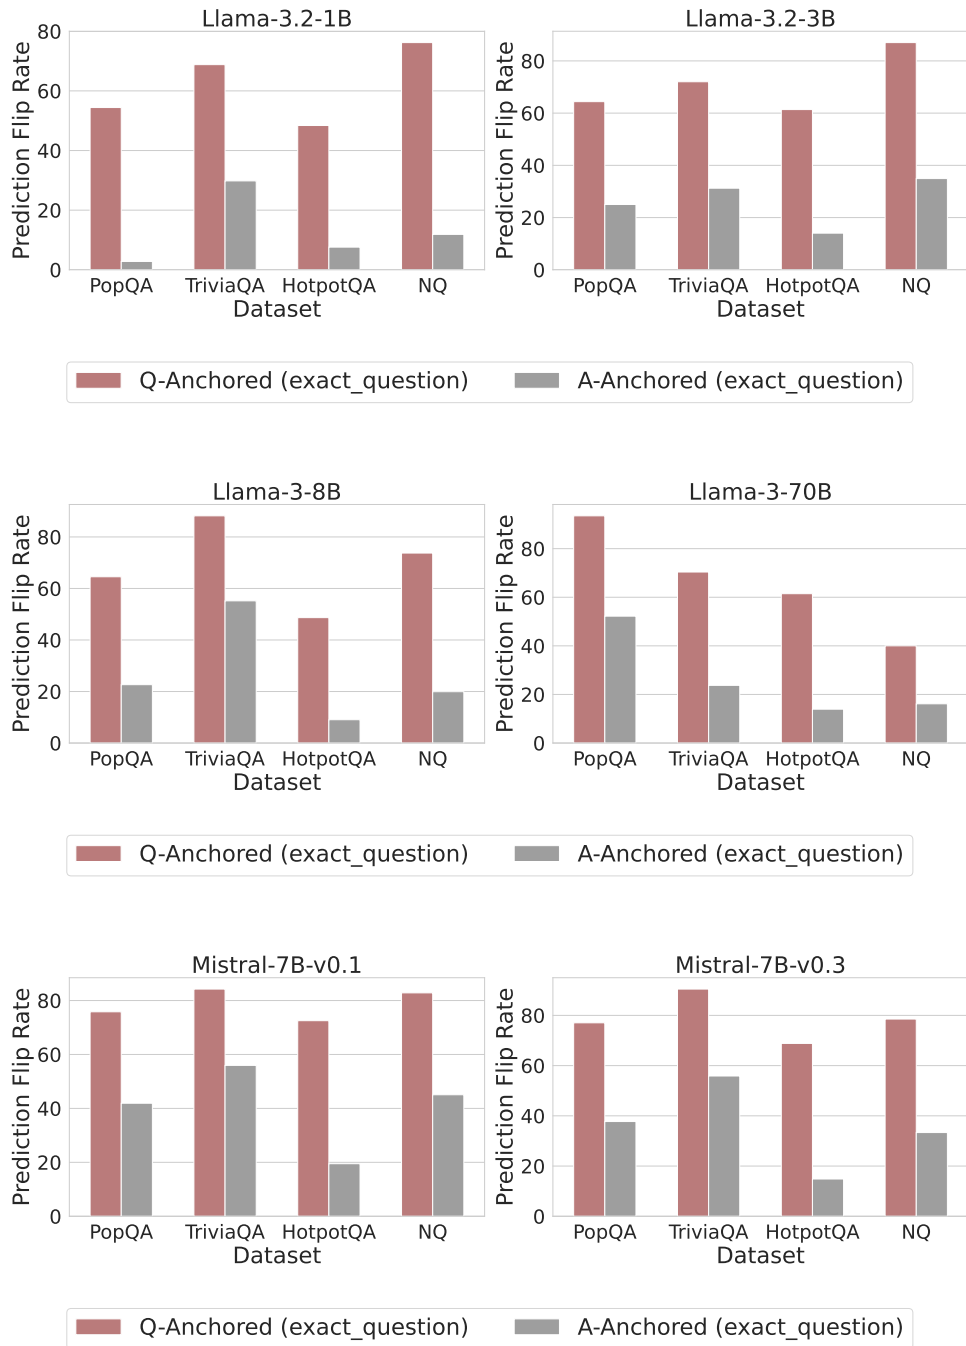


Figure 20: Prediction flip rate under token patching, probing mlp activations of the final token.

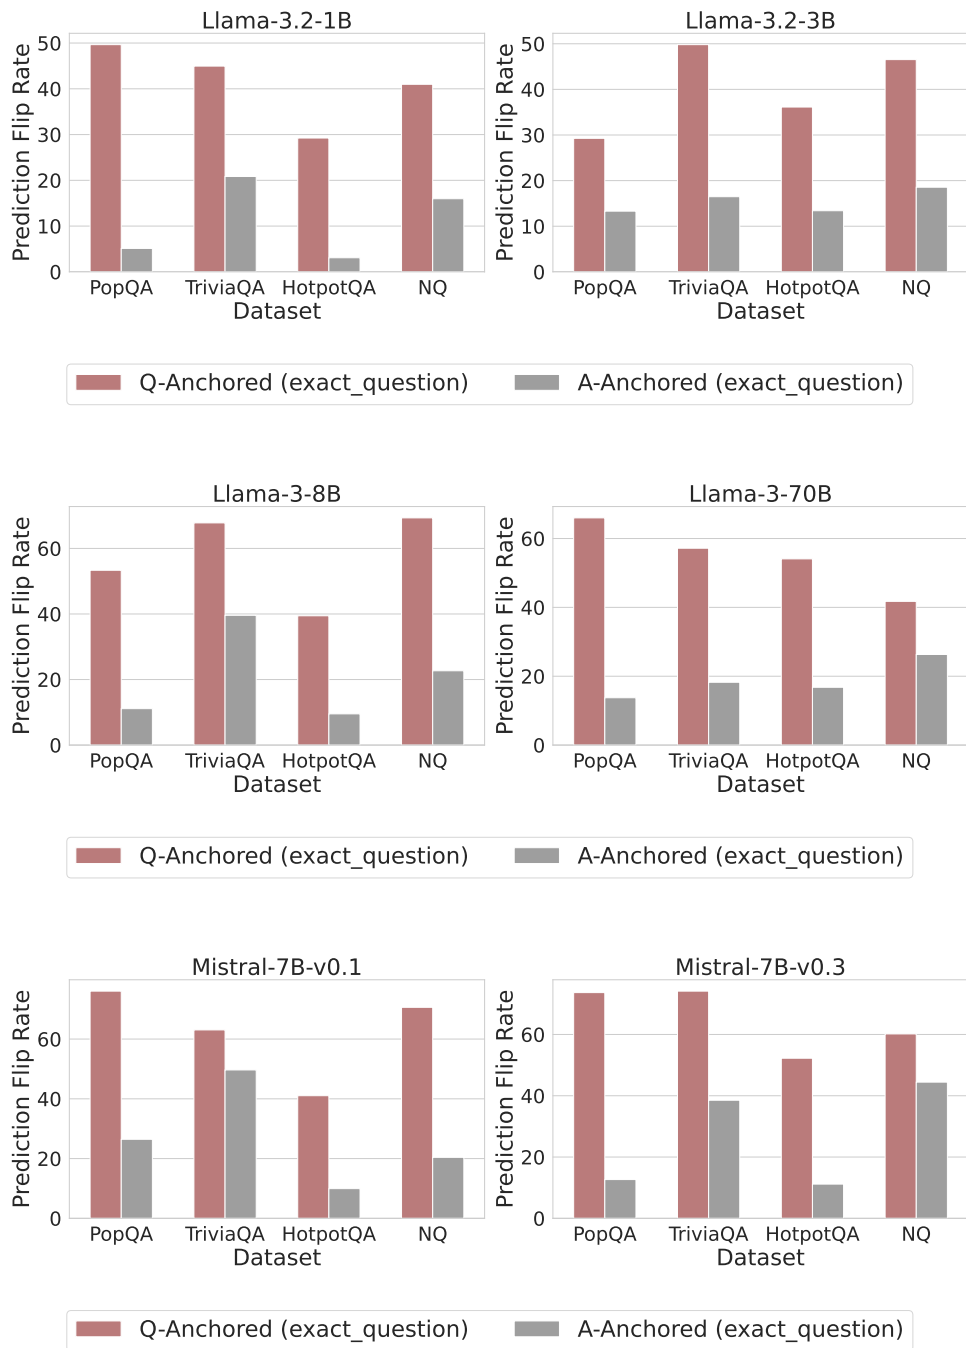


Figure 21: Prediction flip rate under token patching, probing mlp activations of the token immediately preceding the exact answer tokens.

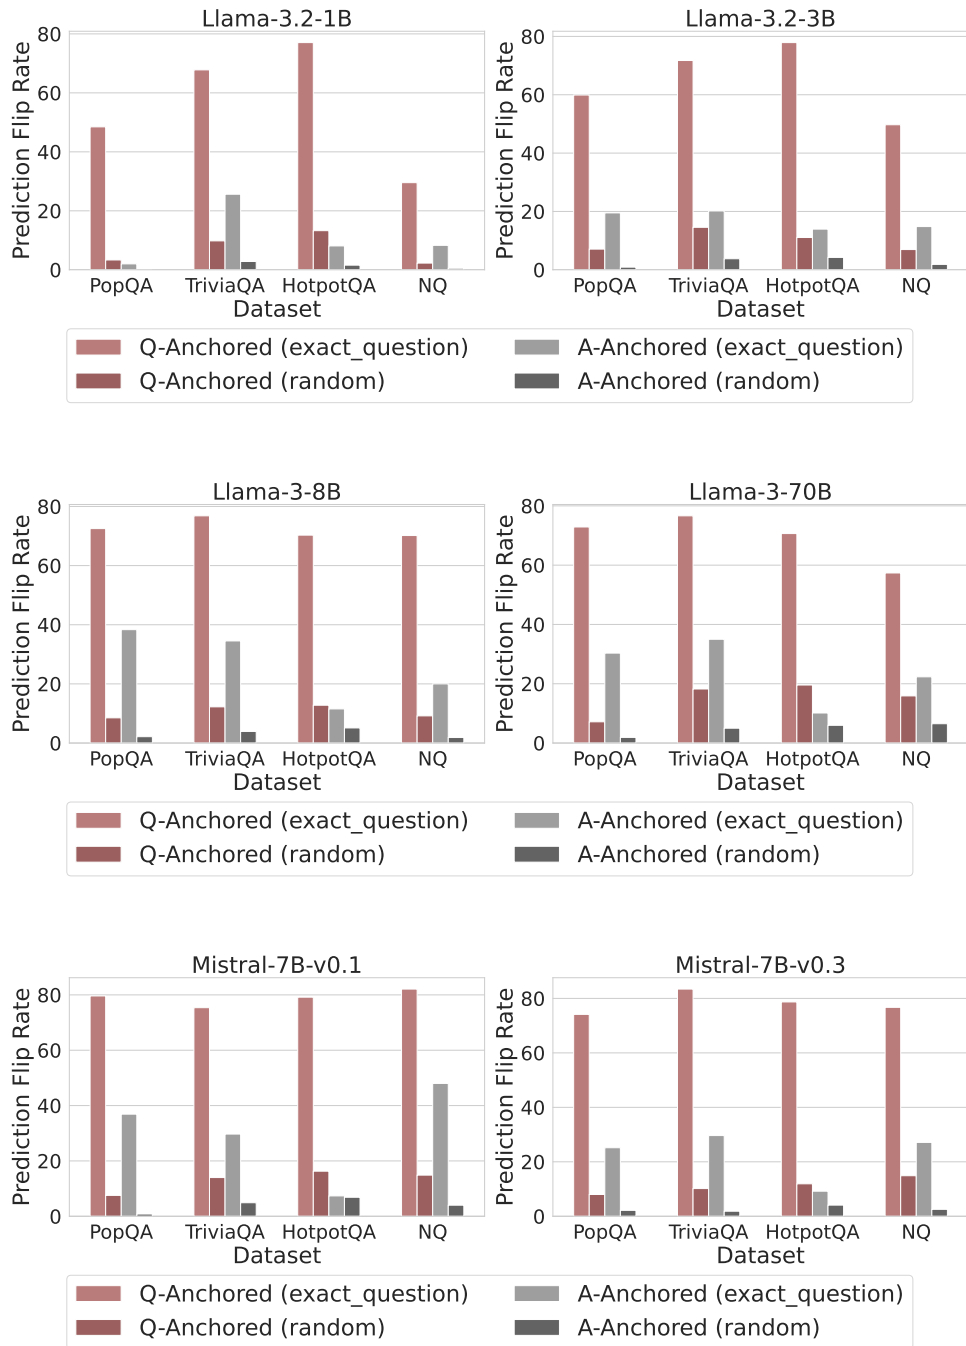


Figure 22: Prediction flip rate under token patching, probing mlp activations of the last exact answer token.

E Answer-Only Input

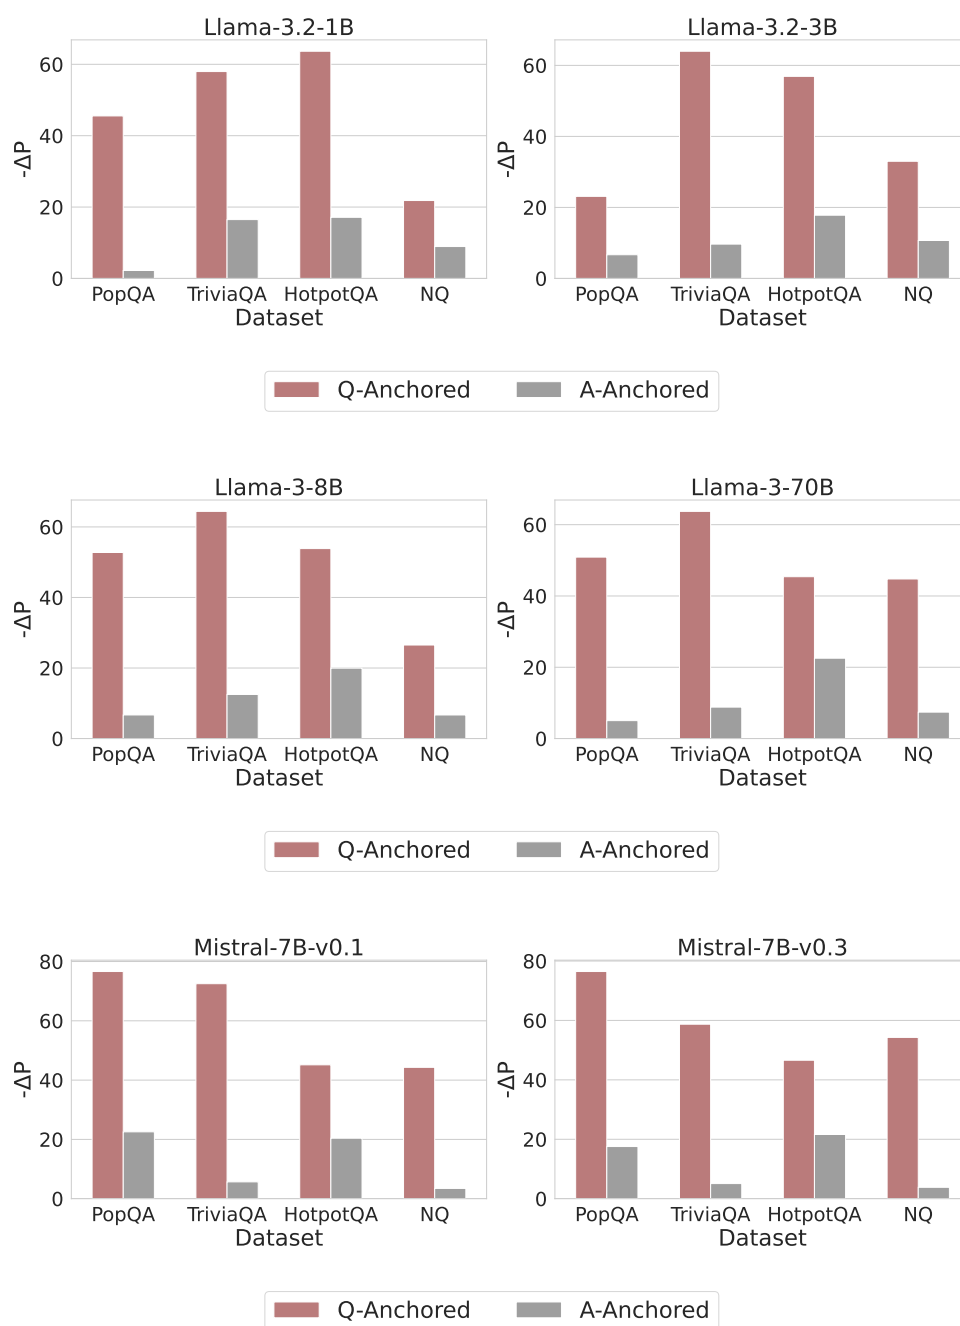


Figure 23: $-\Delta P$ with only the LLM-generated answer. Q-Anchored instances exhibit substantial shifts, whereas A-Anchored instances remain stable, confirming that A-Anchored truthfulness encoding relies on information in the LLM-generated answer itself.

F Answer Accuracy

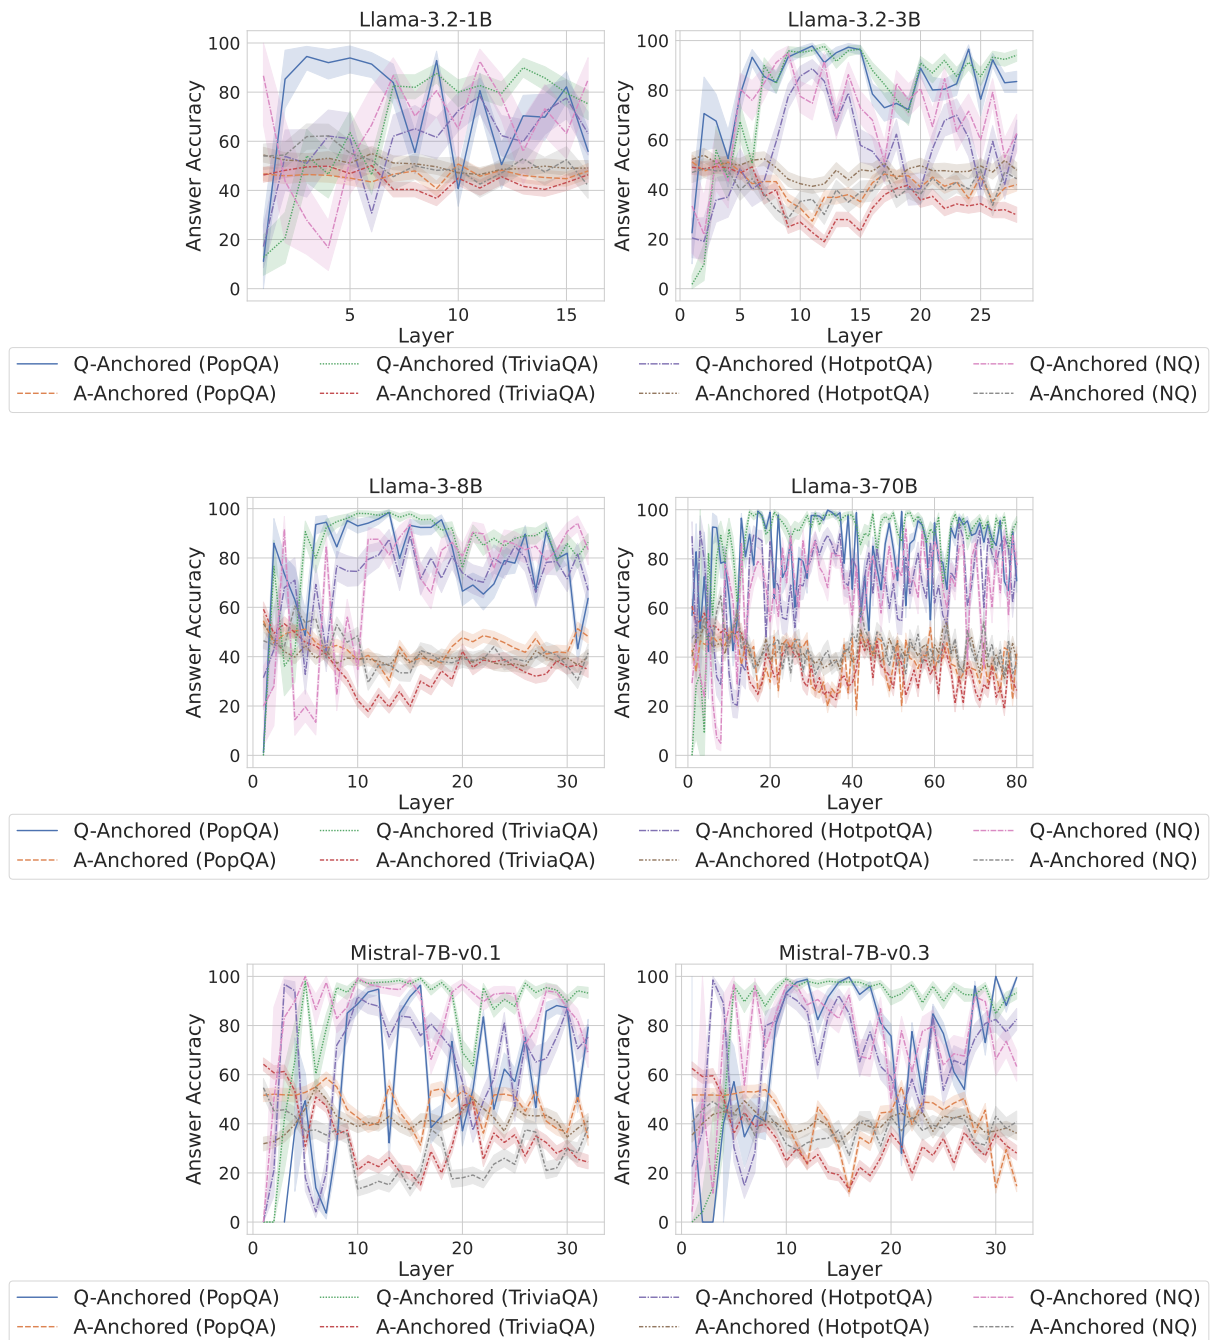


Figure 24: Comparisons of answer accuracy between pathways, probing attention activations of the final token.

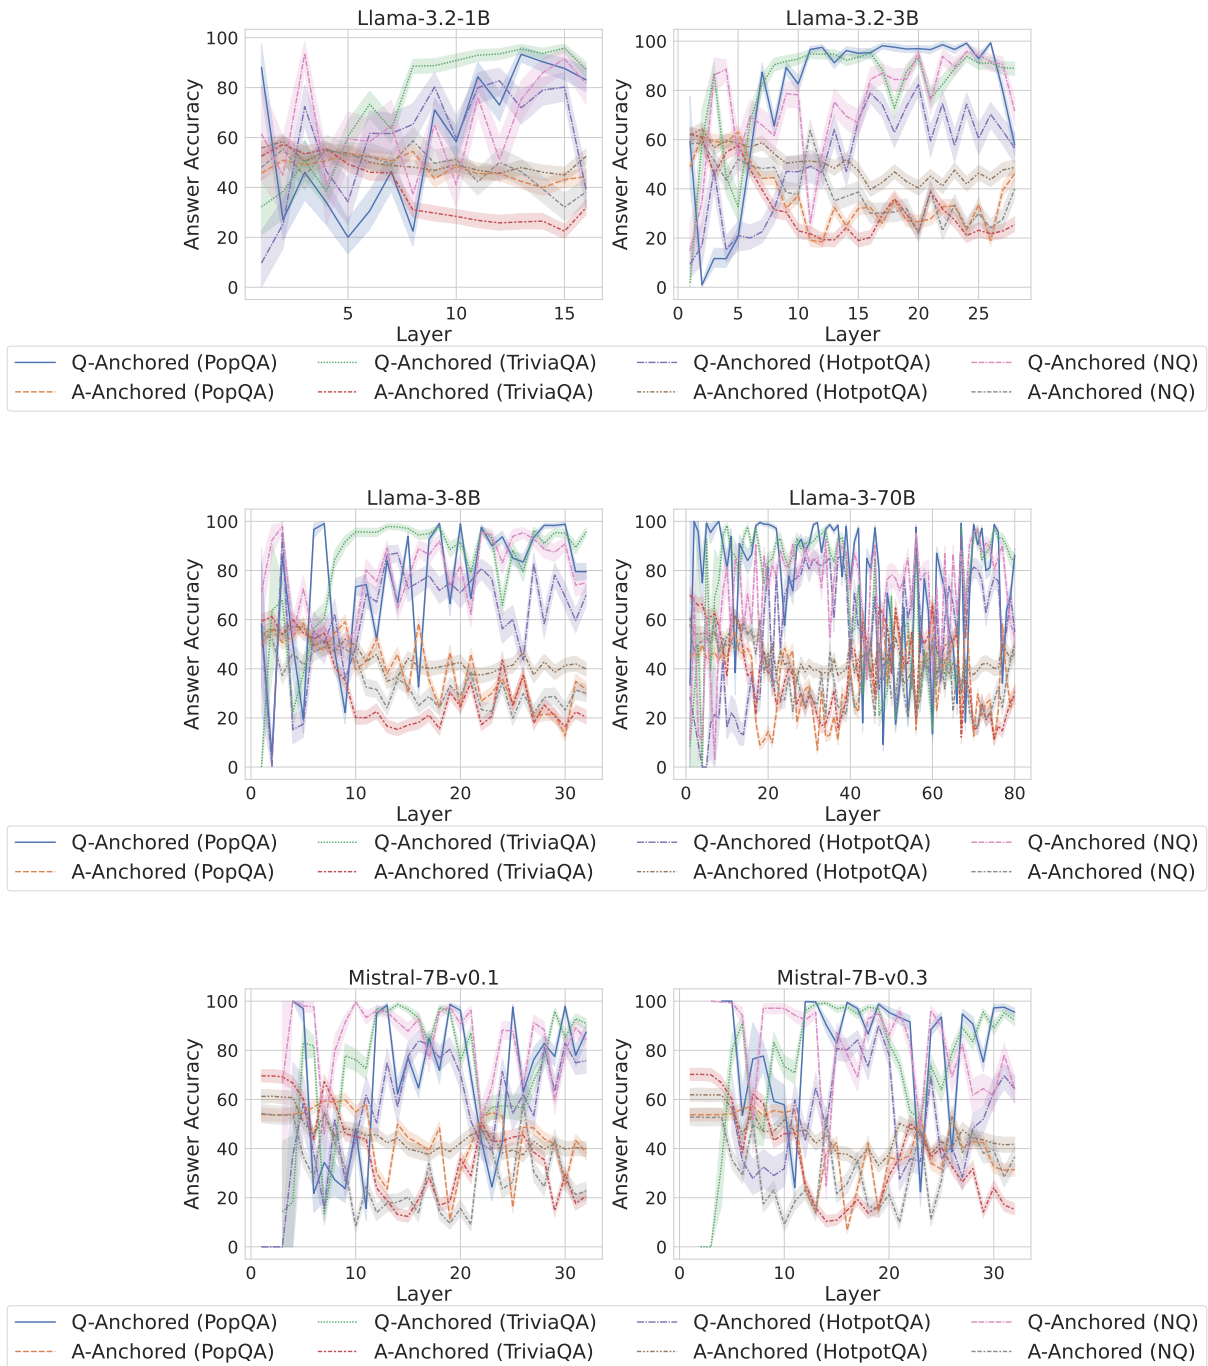


Figure 25: Comparisons of answer accuracy between pathways, probing attention activations of the token immediately preceding the exact answer tokens.

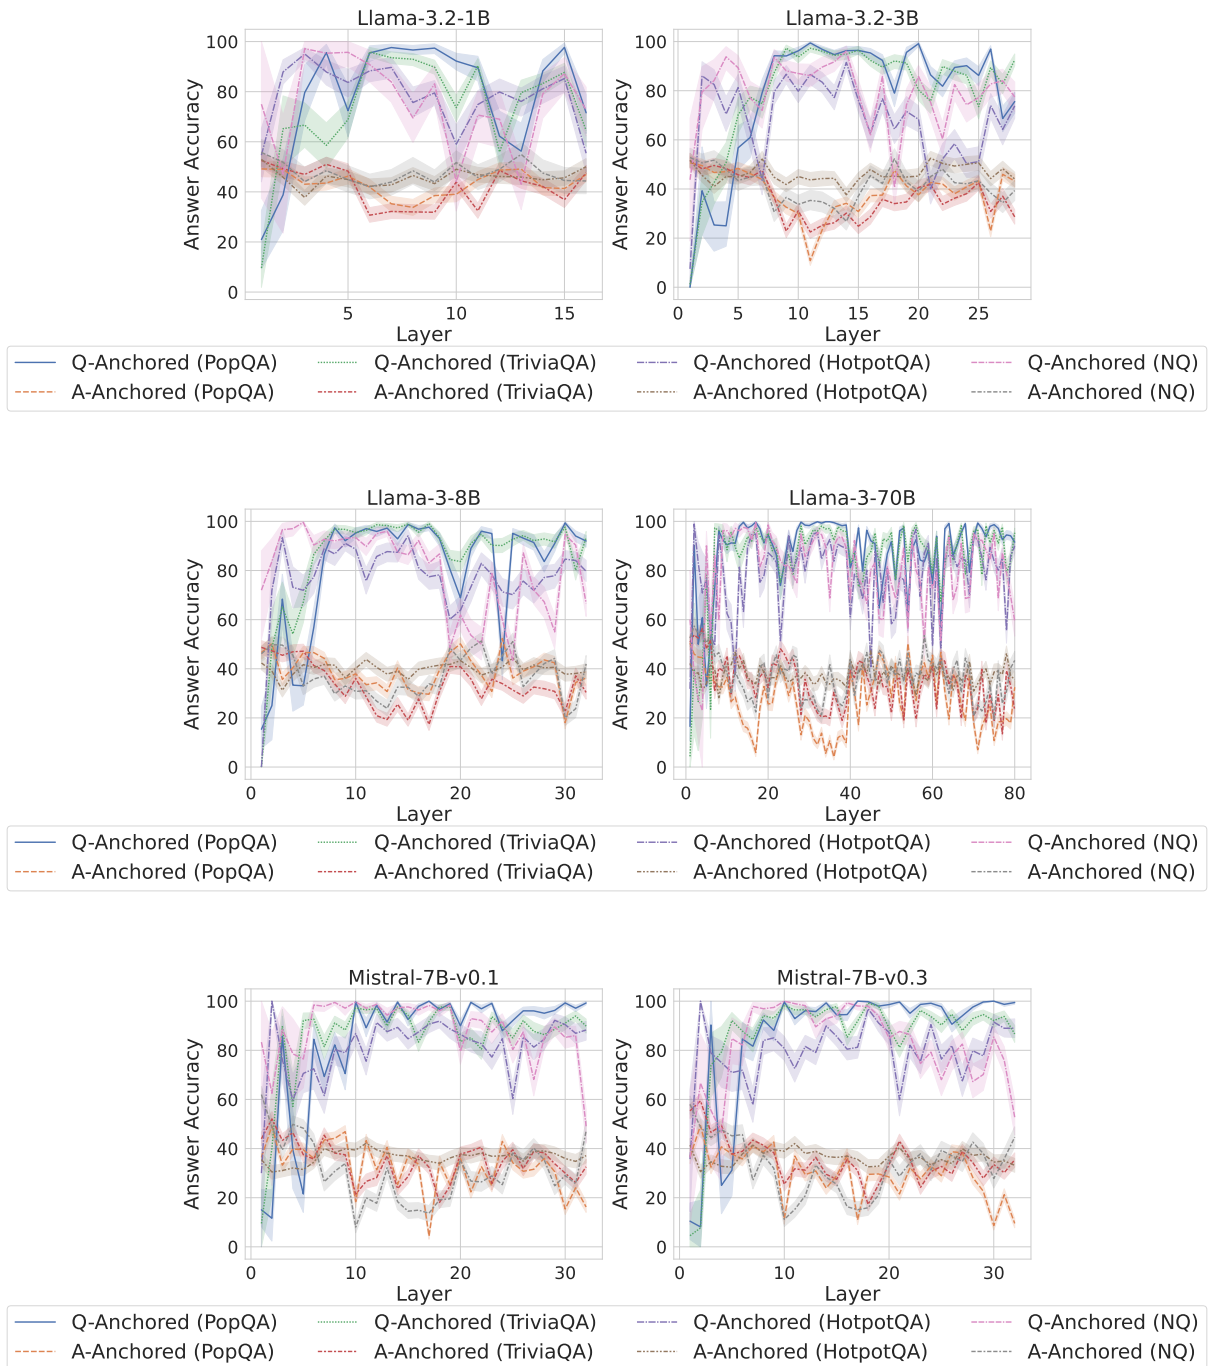


Figure 26: Comparisons of answer accuracy between pathways, probing attention activations of the last exact answer token.

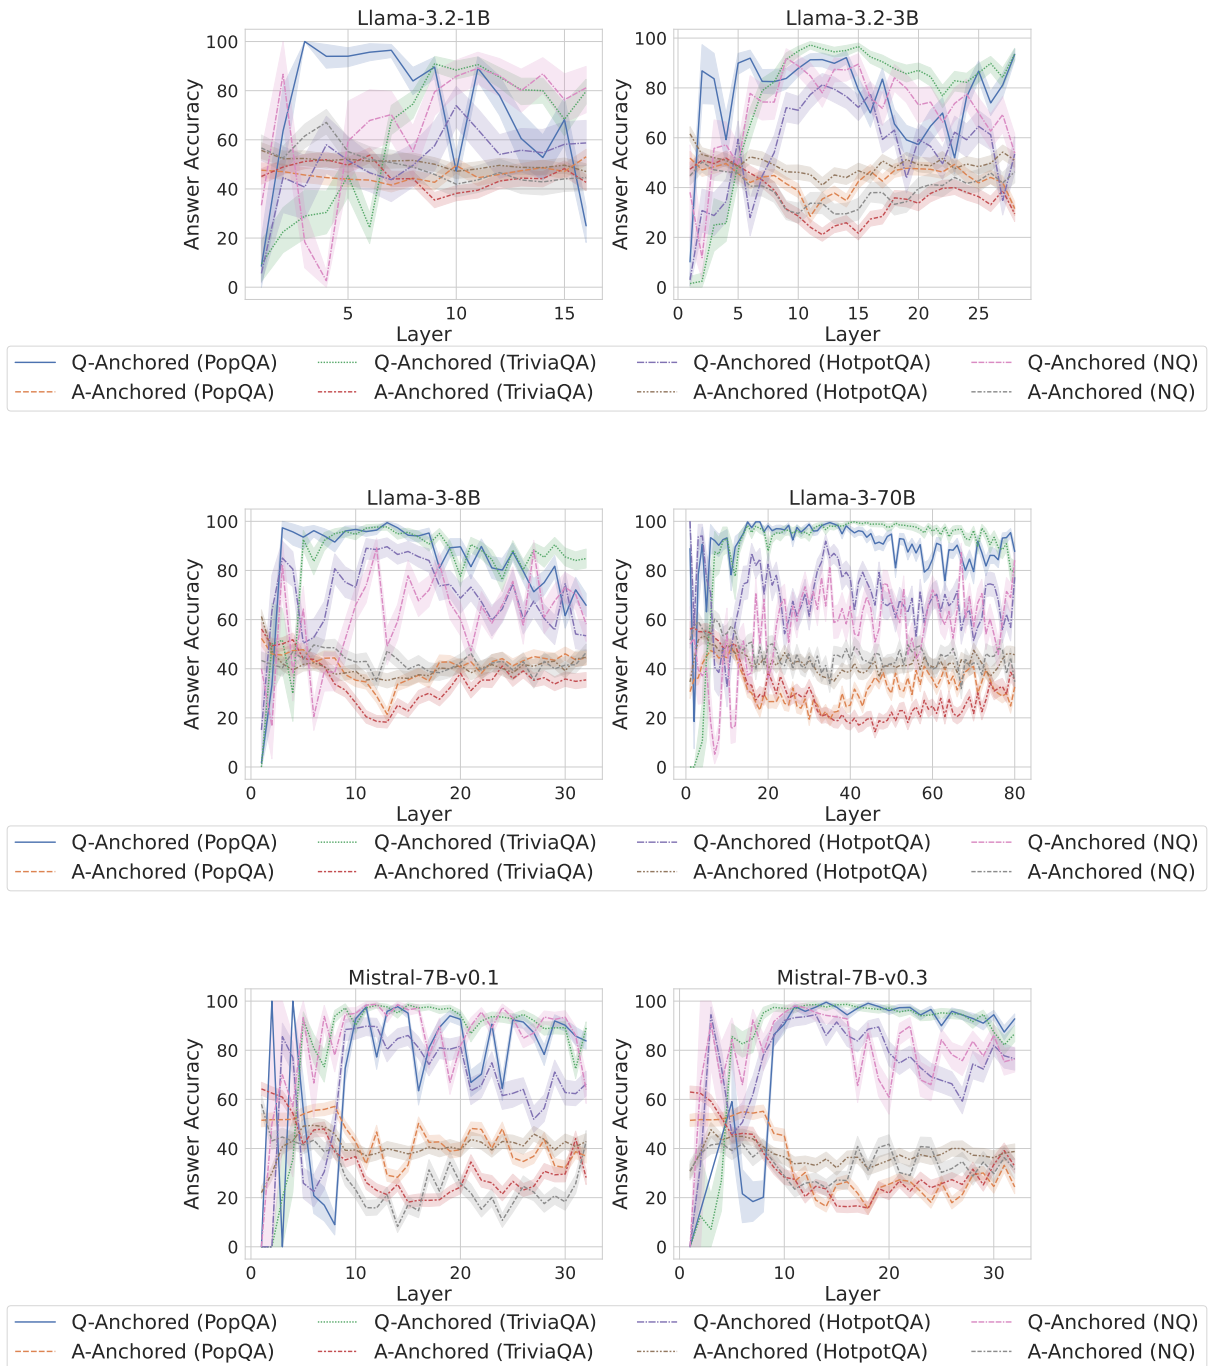


Figure 27: Comparisons of answer accuracy between pathways, probing mlp activations of the final token.

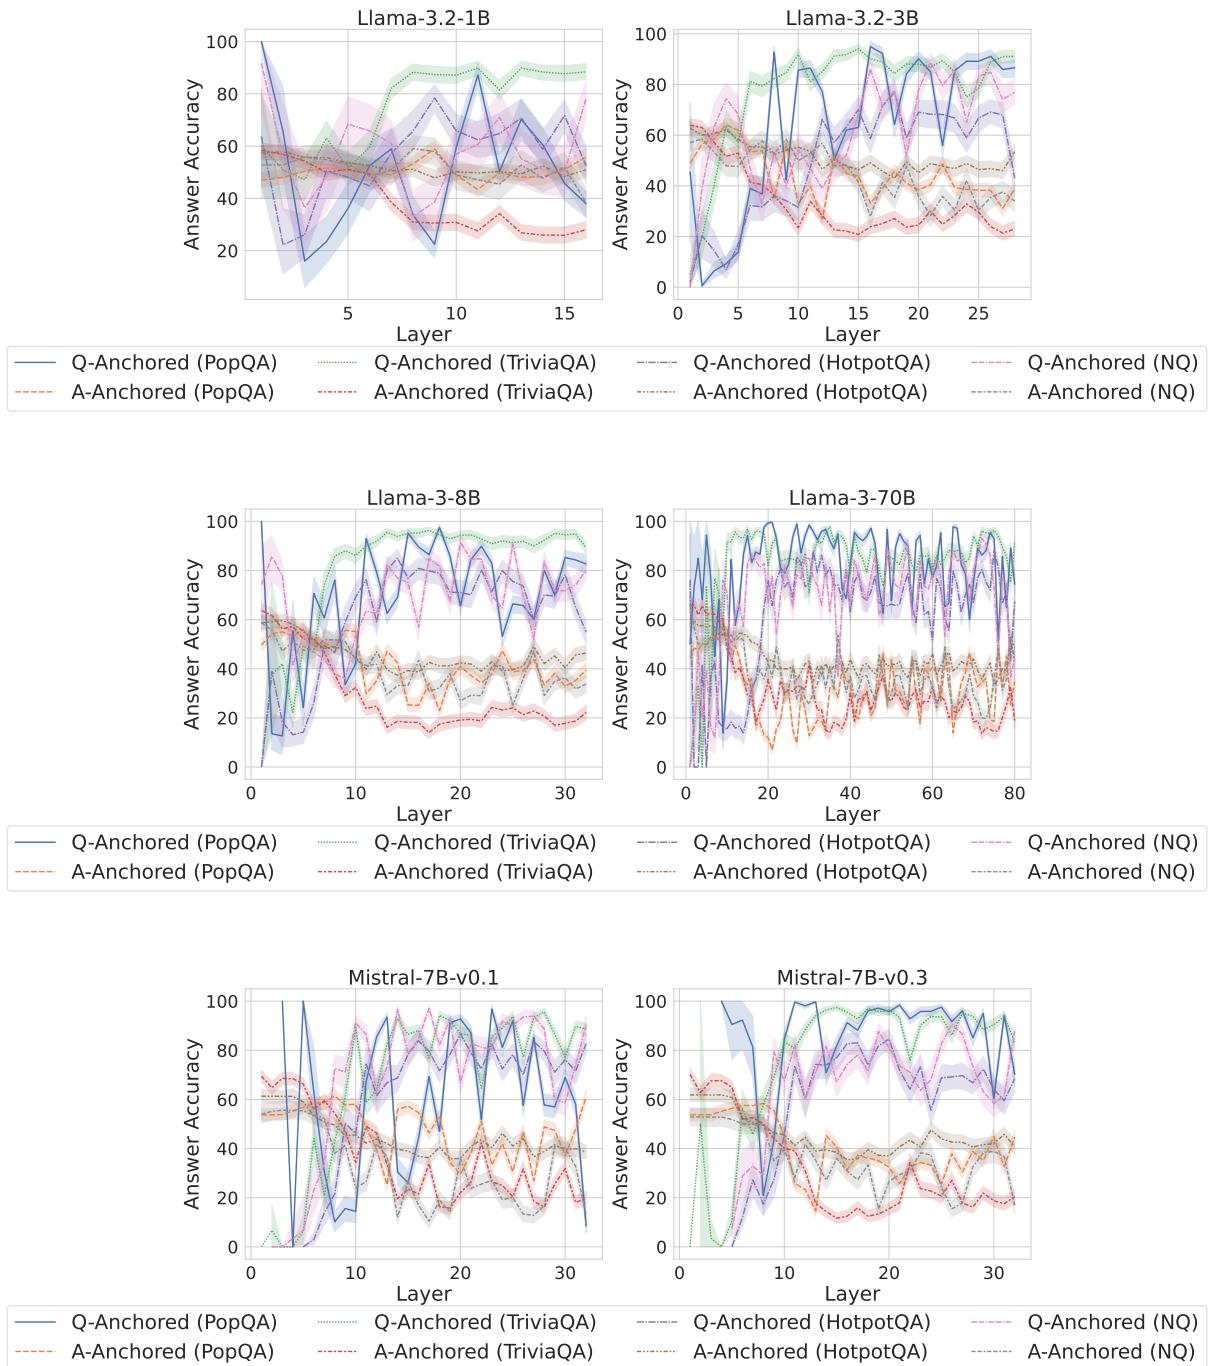


Figure 28: Comparisons of answer accuracy between pathways, probing mlp activations of the token immediately preceding the exact answer tokens.

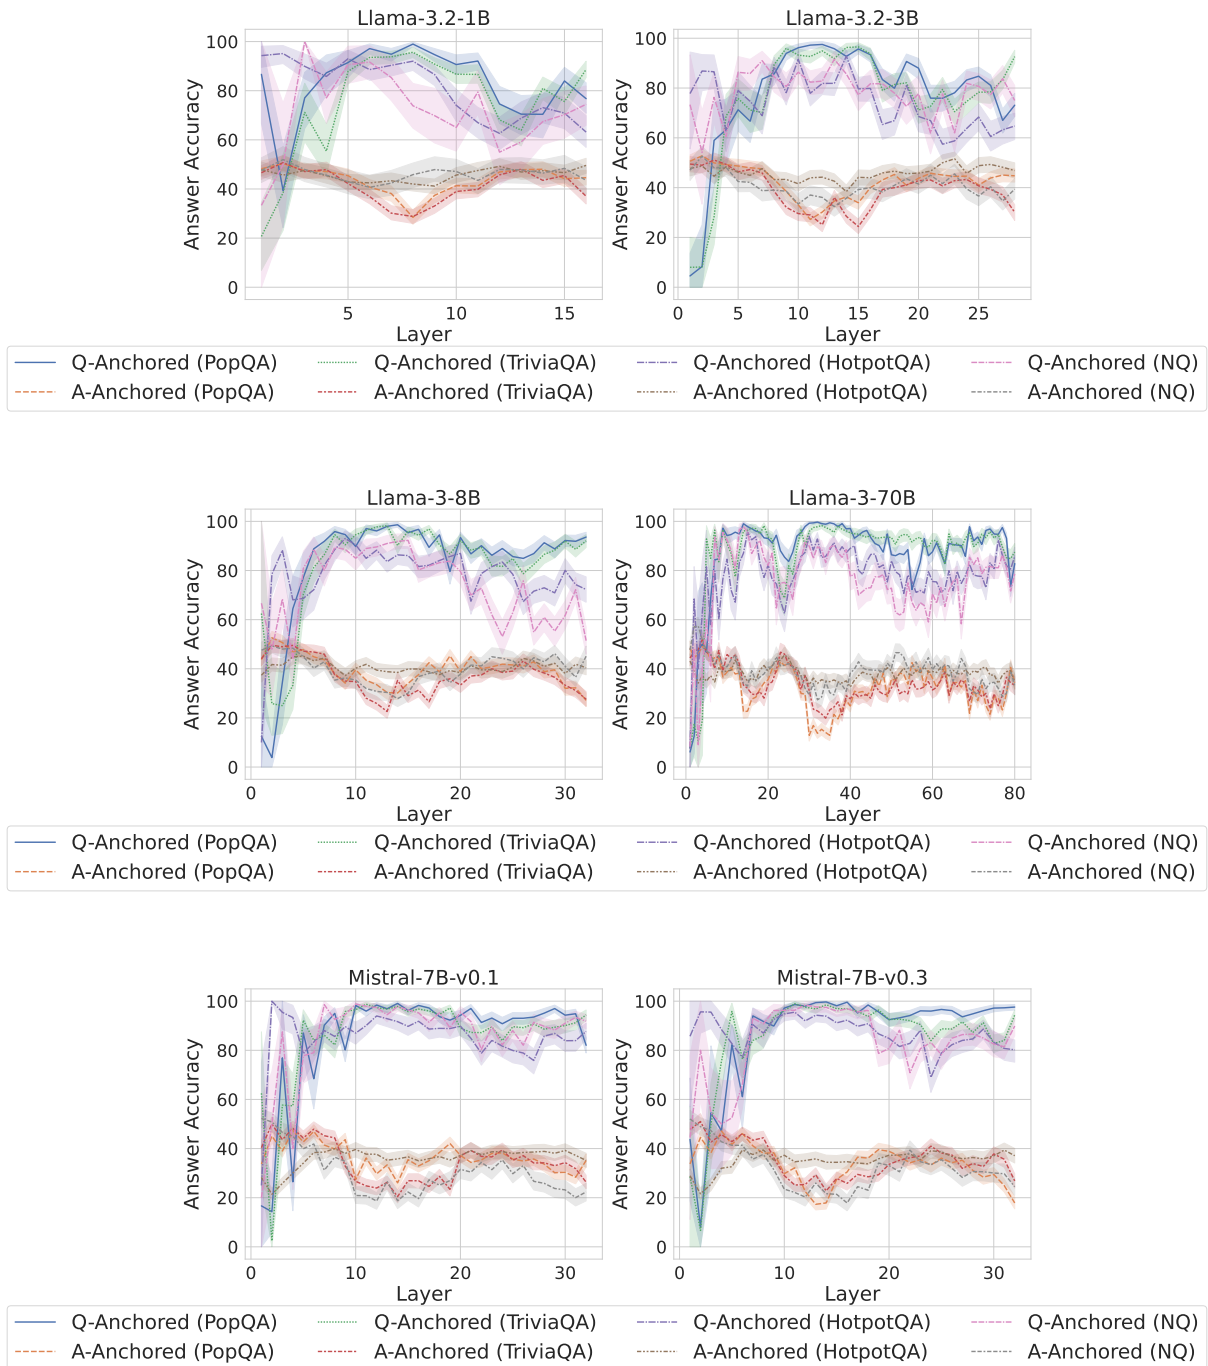


Figure 29: Comparisons of answer accuracy between pathways, probing mlp activations of the last exact answer token.

G I-Don't-Know Rate

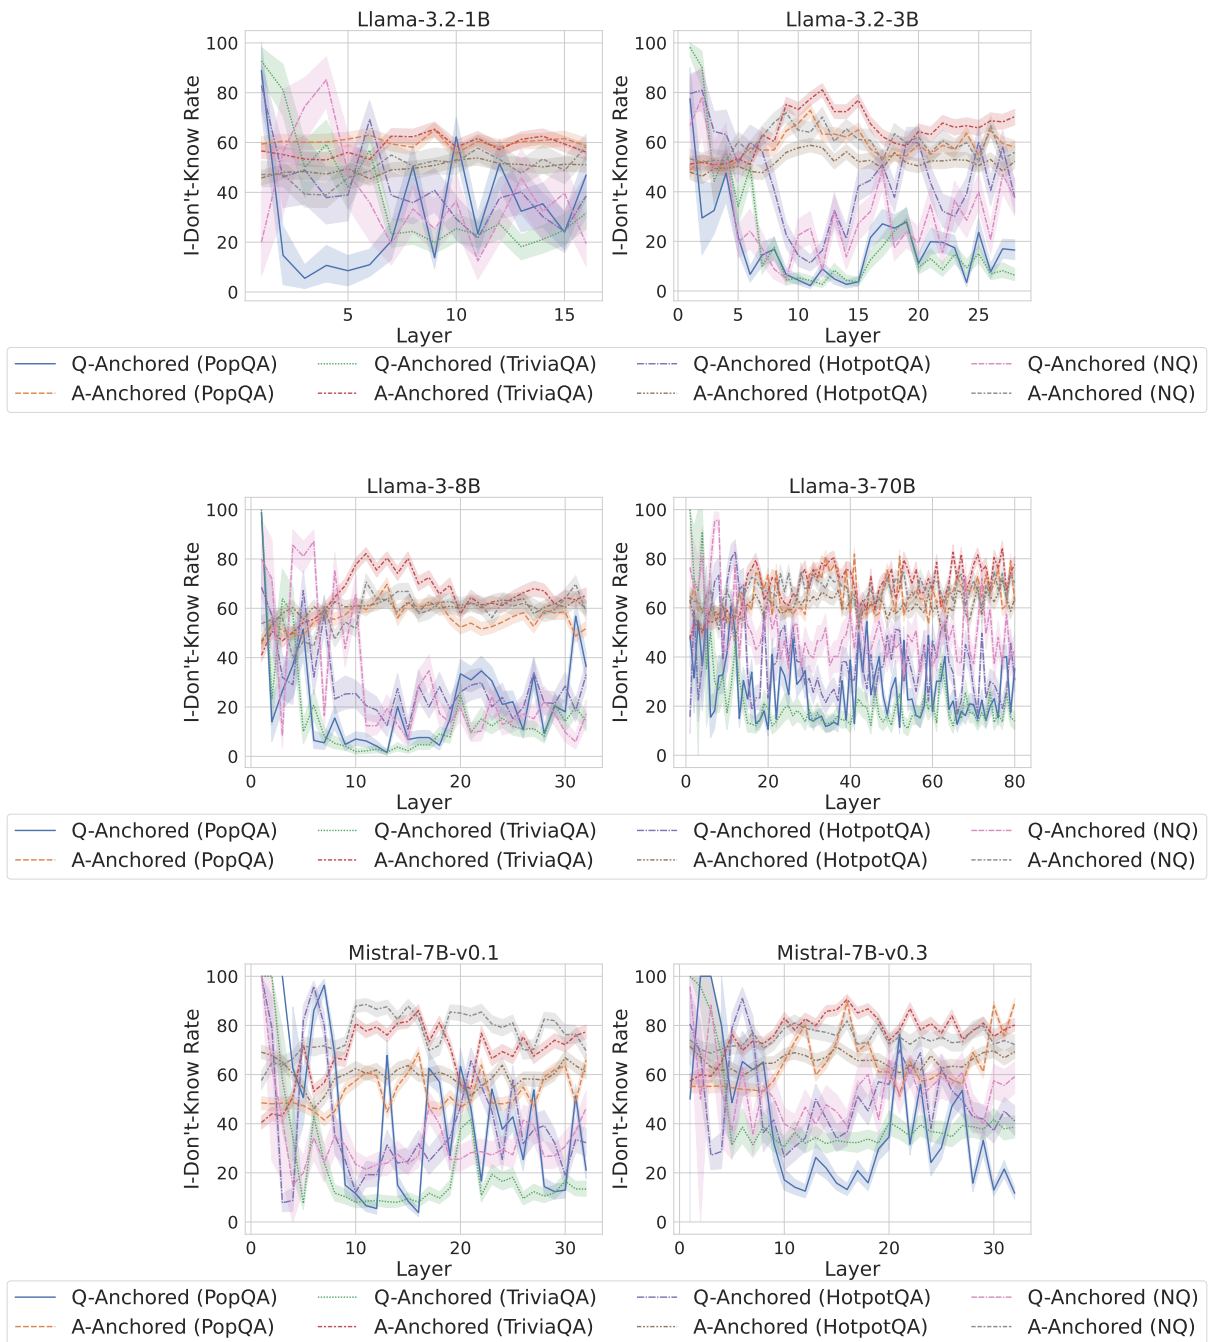


Figure 30: Comparisons of i-don't-know rate between pathways, probing attention activations of the final token.

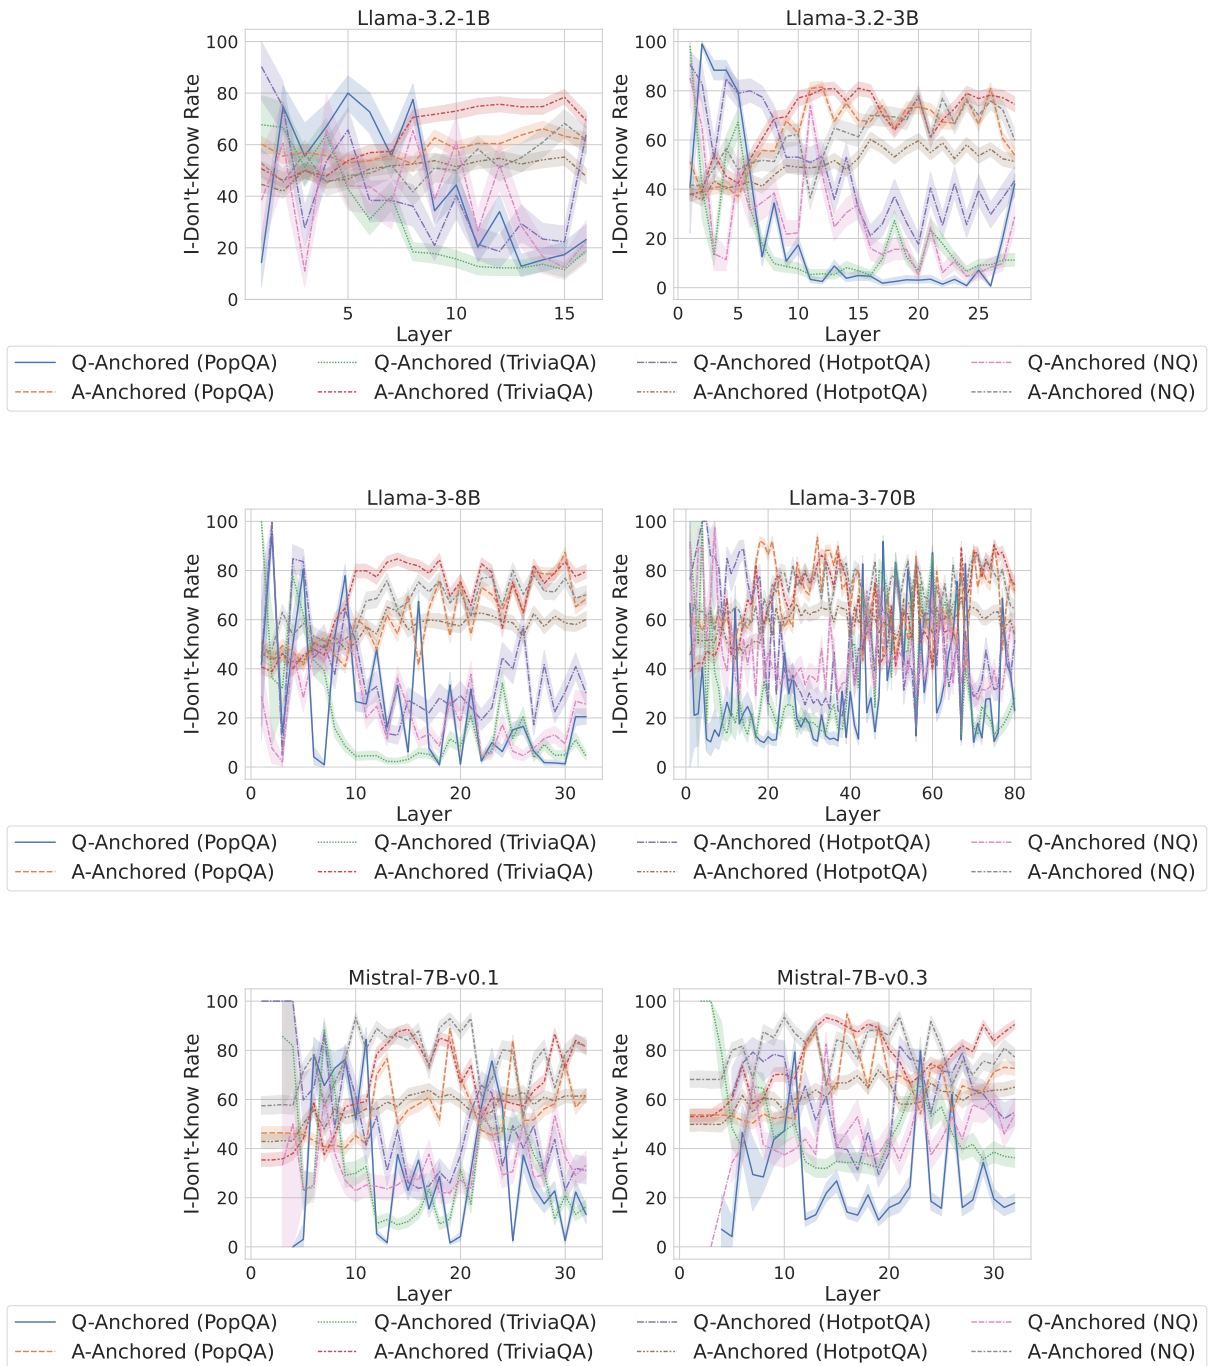


Figure 31: Comparisons of i-don't-know rate between pathways, probing attention activations of the token immediately preceding the exact answer tokens.

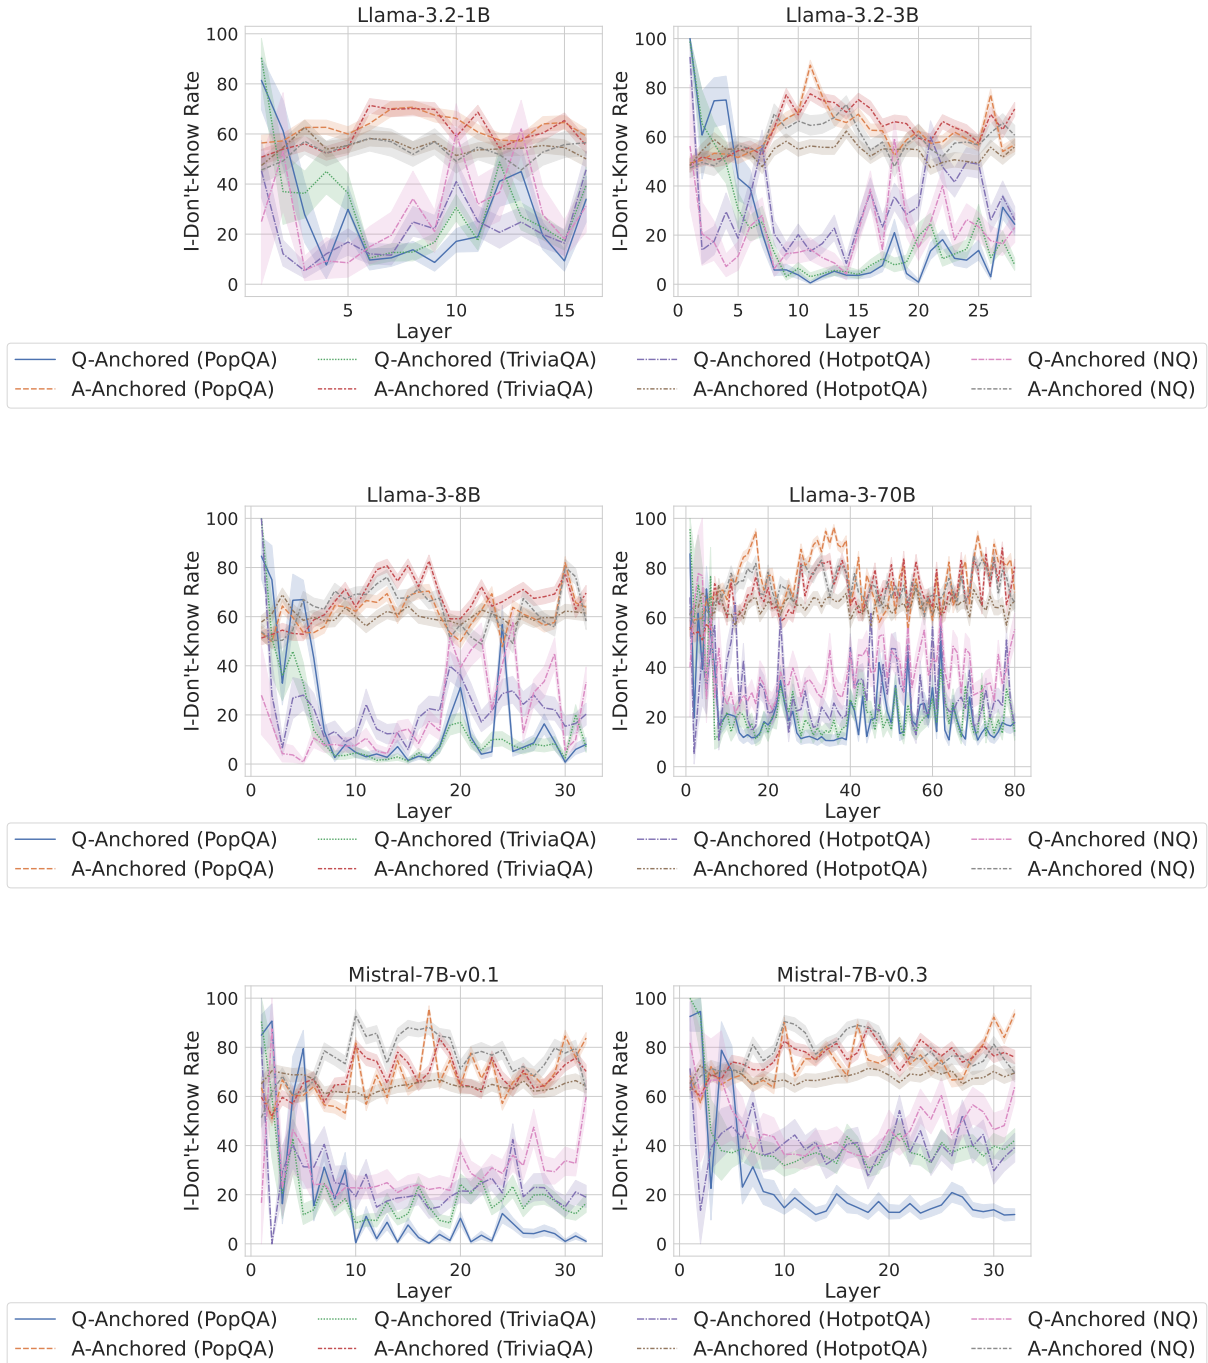


Figure 32: Comparisons of i-don't-know rate between pathways, probing attention activations of the last exact answer token.

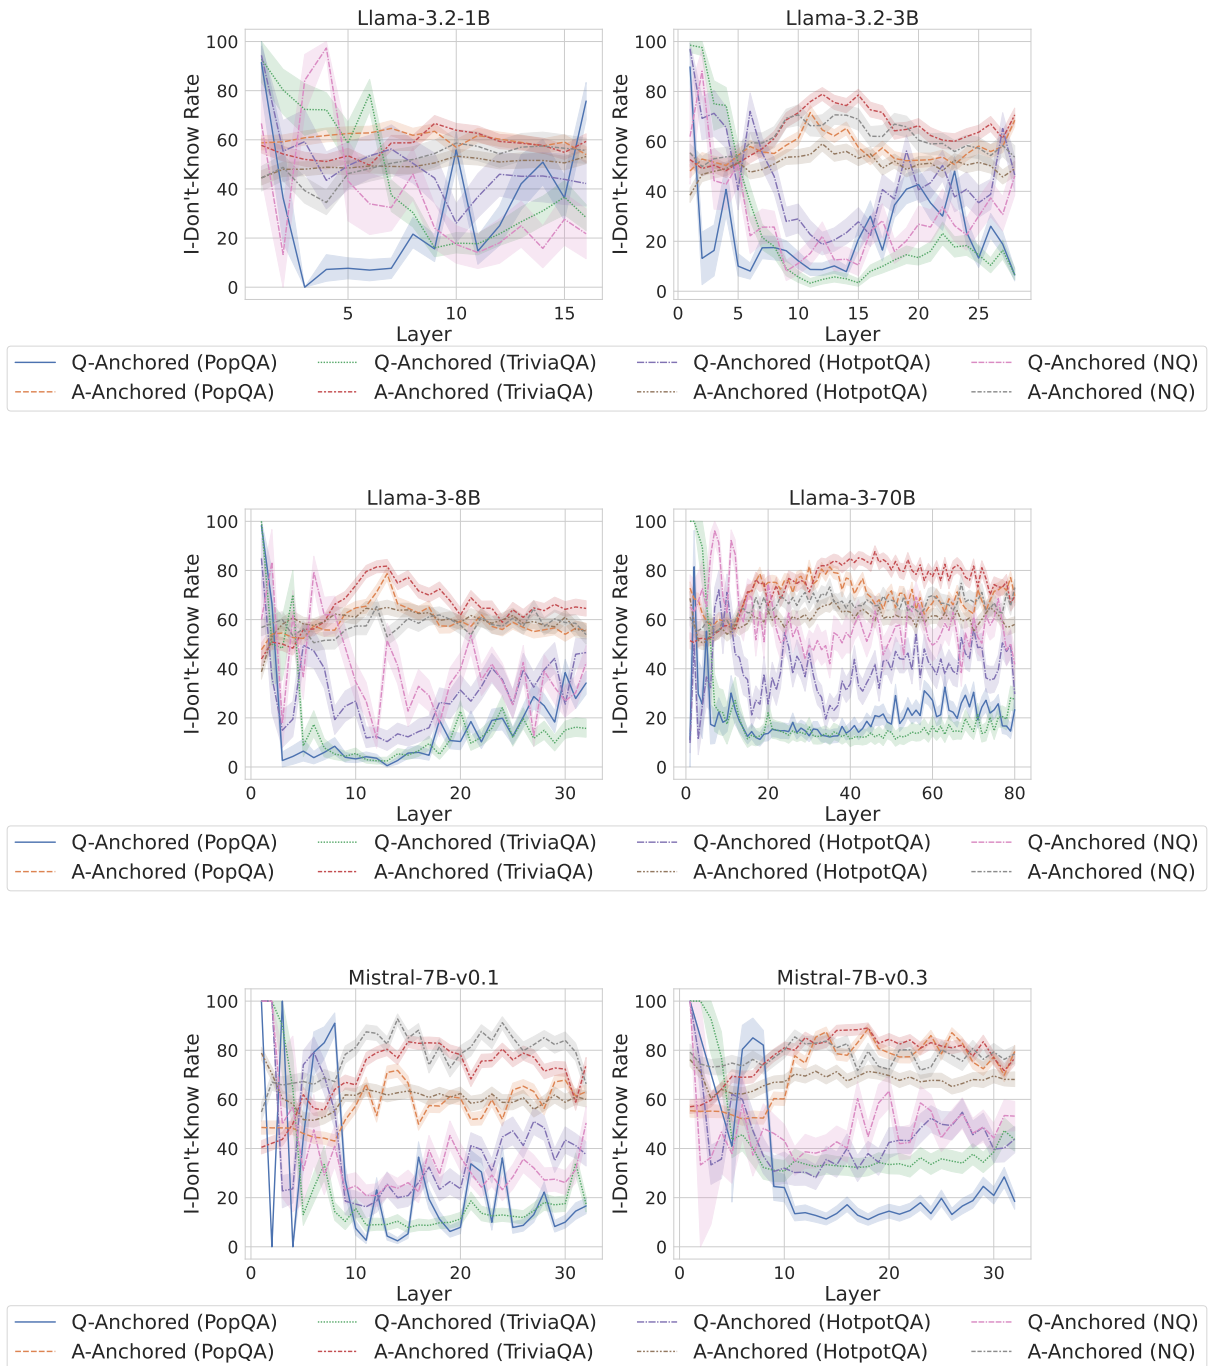


Figure 33: Comparisons of i-don't-know rate between pathways, probing mlp activations of the final token.

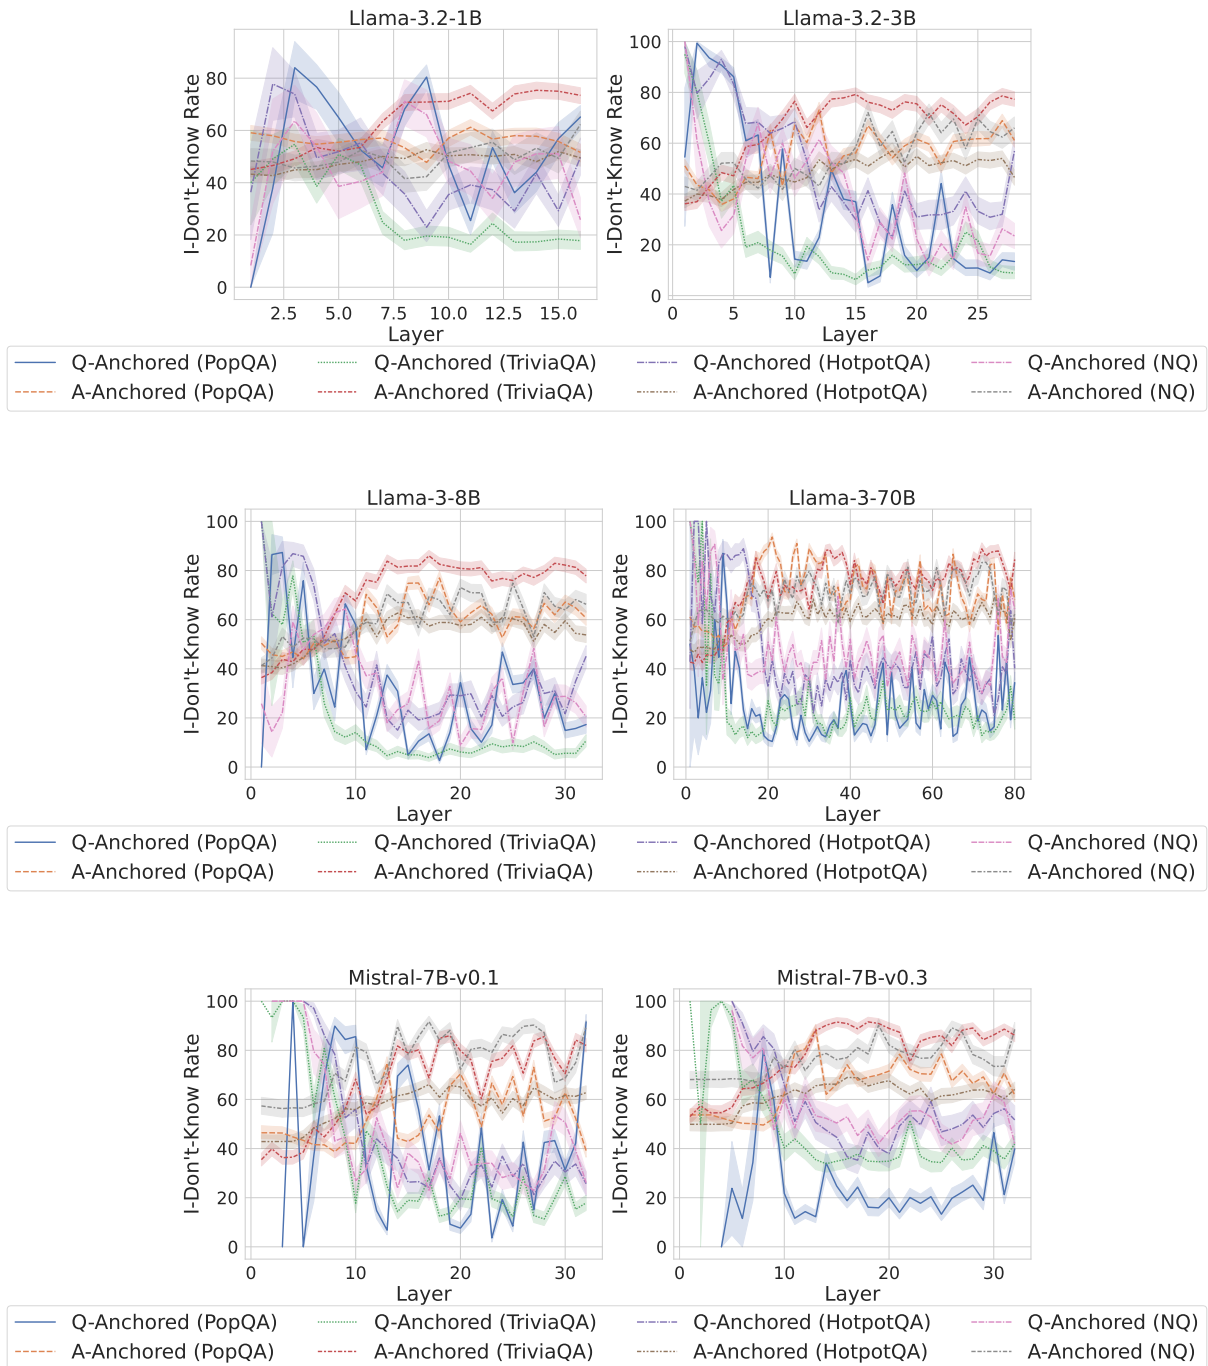


Figure 34: Comparisons of i-don't-know rate between pathways, probing mlp activations of the token immediately preceding the exact answer tokens.

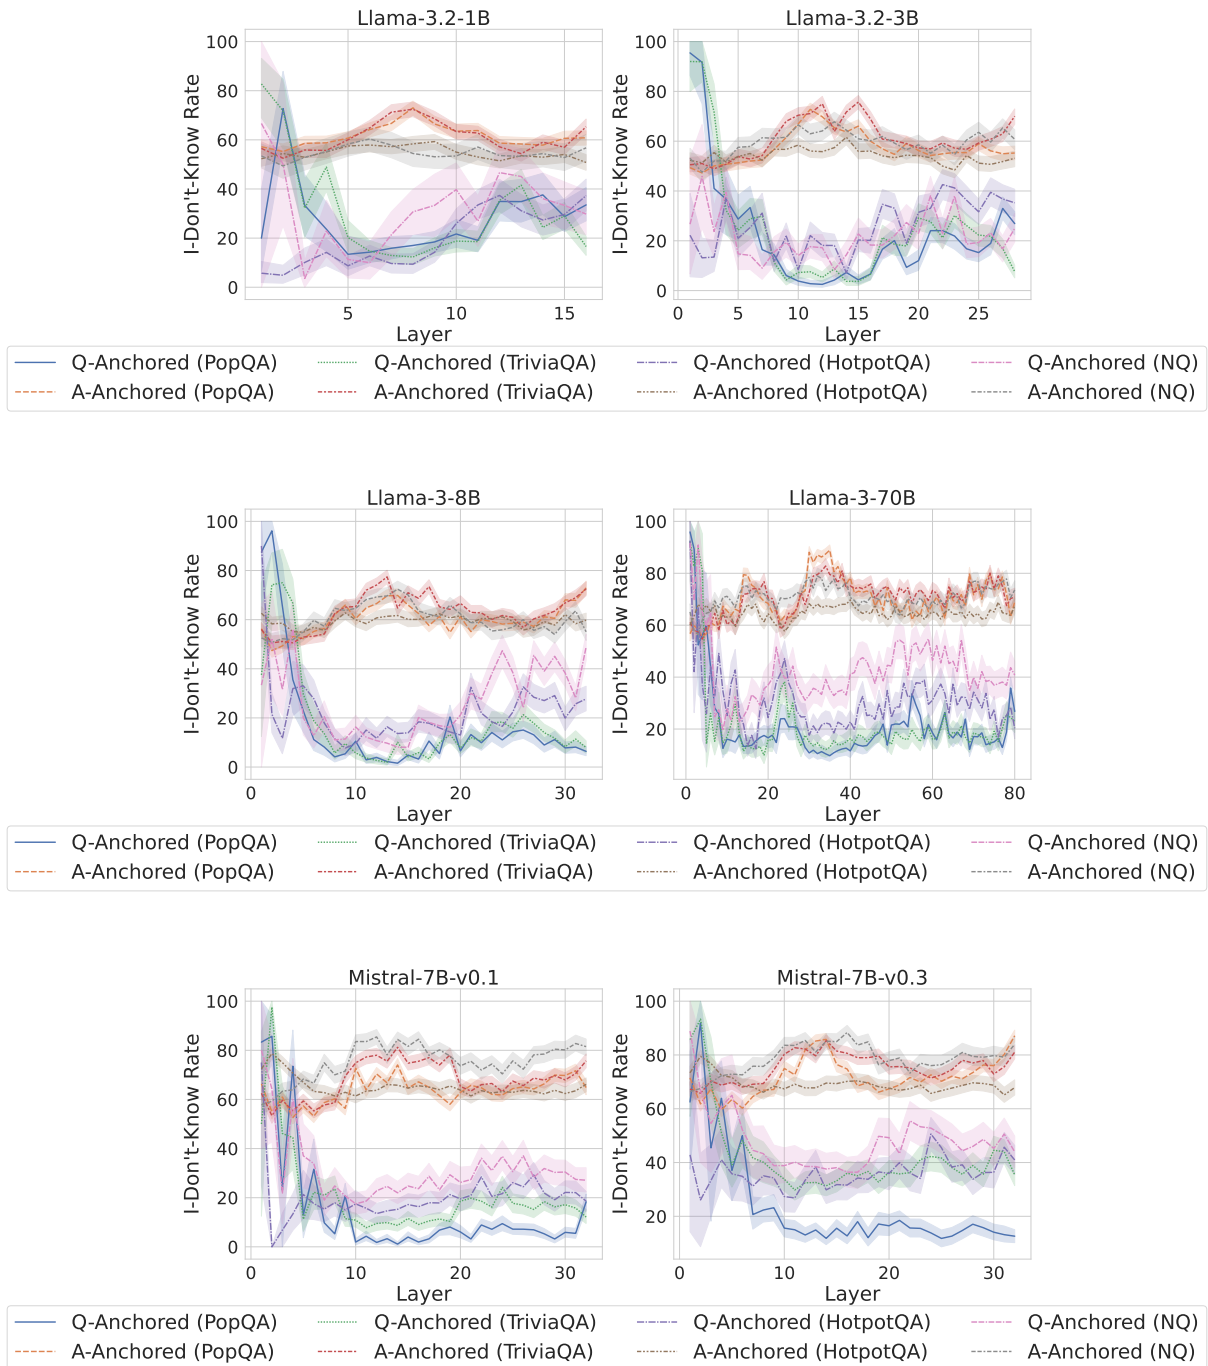


Figure 35: Comparisons of i-don't-know rate between pathways, probing mlp activations of the last exact answer token.

H Pathway-Aware Detection

Method	LLama-3.2-1B				LLama-3.2-3B			
	PopQA	TriviaQA	HotpotQA	NQ	PopQA	TriviaQA	HotpotQA	NQ
P(True)	60.00	49.65	43.34	52.83	54.58	51.76	47.73	53.78
Logits-mean	74.89	60.24	60.18	49.92	73.47	63.46	60.35	54.89
Logits-max	58.56	52.37	52.29	46.19	56.03	54.33	48.65	48.88
Logits-min	78.66	62.37	67.14	51.20	80.92	69.60	71.11	58.24
Scores-mean	72.91	61.13	62.16	64.67	67.99	61.96	64.91	61.71
Scores-max	69.33	59.74	61.29	64.08	63.34	61.92	61.09	57.56
Scores-min	64.84	55.93	59.28	55.81	61.51	56.76	63.95	57.43
Probing Baseline	94.25	77.17	90.25	74.83	90.96	76.61	86.54	74.20
MoP-RandomGate	83.69	69.20	84.11	68.76	79.69	72.38	75.13	67.11
MoP-VanillaExperts	93.86	78.63	90.91	75.73	90.98	77.68	86.41	75.30
MoP	95.85	80.07	91.51	79.19	92.74	78.72	88.16	78.14
PR	96.18	84.22	92.80	86.45	95.70	80.66	90.66	81.91

Table 8: Comparison of hallucination detection performance (AUC) on LLama-3.2-1B and LLama-3.2-3B.

Method	LLama-3-8B				LLama-3-70B			
	PopQA	TriviaQA	HotpotQA	NQ	PopQA	TriviaQA	HotpotQA	NQ
P(True)	55.85	49.92	52.14	53.27	54.83	50.96	49.39	51.18
Logits-mean	74.52	60.39	51.94	52.63	67.81	52.40	50.45	48.28
Logits-max	58.08	52.20	46.40	47.89	56.21	48.16	43.42	45.33
Logits-min	85.36	70.89	61.28	56.50	79.96	61.53	62.63	52.16
Scores-mean	62.87	62.09	62.06	60.32	56.81	60.70	60.91	58.05
Scores-max	56.62	60.24	59.85	56.06	55.15	59.60	57.32	51.93
Scores-min	60.99	58.27	60.33	57.68	58.77	58.22	64.06	58.05
Probing Baseline	88.71	77.58	82.23	70.20	86.88	81.59	84.45	74.39
MoP-RandomGate	75.52	69.17	79.88	66.56	67.96	70.56	72.16	66.28
MoP-VanillaExperts	89.11	78.73	84.57	71.21	86.04	82.47	82.48	73.85
MoP	92.11	81.18	85.45	74.64	88.54	84.12	86.65	76.12
PR	94.01	83.13	87.81	79.10	90.08	84.21	87.69	78.24

Table 9: Comparison of hallucination detection performance (AUC) on LLama-3-8B and LLama-3-70B.

Method	Mistral-7B-v0.1				Mistral-7B-v0.3			
	PopQA	TriviaQA	HotpotQA	NQ	PopQA	TriviaQA	HotpotQA	NQ
P(True)	48.78	50.43	51.94	55.52	45.49	47.61	57.87	52.79
Logits-mean	69.09	64.95	54.47	59.41	69.52	66.76	55.45	57.88
Logits-max	54.37	54.76	46.74	56.45	54.34	55.24	48.39	54.37
Logits-min	86.02	76.56	68.06	53.73	87.05	77.33	68.08	54.40
Scores-mean	59.00	59.61	64.18	57.60	58.84	60.22	63.28	60.05
Scores-max	51.71	56.58	63.29	55.82	53.00	55.55	63.13	57.73
Scores-min	60.00	57.48	61.17	48.51	60.59	57.84	59.85	50.76
Probing Baseline	89.61	78.43	83.76	74.10	87.39	81.74	83.19	73.60
MoP-RandomGate	80.50	68.27	74.51	68.05	79.81	70.88	72.23	61.19
MoP-VanillaExperts	89.82	79.51	83.54	74.78	88.53	80.93	82.93	73.77
MoP	92.44	84.03	84.63	76.38	91.66	83.57	85.82	76.87
PR	94.72	84.66	89.04	80.92	93.09	84.36	89.03	79.09

Table 10: Comparison of hallucination detection performance (AUC) on Mistral-7B-v0.1 and Mistral-7B-v0.3.