

Beyond the Final Layer: Intermediate Representations for Better Multilingual Calibration in Large Language Models

Ej Zhou* Caiqi Zhang* Tiancheng Hu Chengzu Li

Nigel Collier Ivan Vulić Anna Korhonen

Language Technology Lab, University of Cambridge

{yz926,cz391,th656,c1917,nhc30,iv250,alk23}@cam.ac.uk

Abstract

Confidence calibration, the alignment of a model’s predicted confidence with its actual accuracy, is critical for the reliable deployment of Large Language Models (LLMs), yet remains under-explored in multilingual contexts. We conduct the first large-scale systematic study of multilingual calibration across eight model families and over 100 languages, finding that non-English languages are worse-calibrated than English (e.g., LLaMA3’s average Expected Calibration Error is 23.1% on non-English against 4.6% on English). To diagnose this, we investigate the model’s internal representations and find that the final layer, shaped by English-centric training, provides a poor signal for multilingual confidence, whereas late-intermediate layers offer a better-calibrated signal. Building on this, we introduce a suite of training-free methods, including Language-Aware Confidence Ensemble (LACE), which adaptively selects an ensemble of layers for each language. Both the multilingual calibration gap and the late-intermediate peak persist beyond MCQA, in open-ended QA, chain-of-thought reasoning, and alternative MCQA scoring protocols. Our study highlights the calibration costs associated with English-centric alignment and offers a path toward more globally equitable LLMs.

1 Introduction

Calibration in machine learning denotes the alignment between a model’s predicted confidence and the empirical probability that its predictions are correct (Guo et al., 2017; Geng et al., 2024; Zhang et al., 2025a, 2026).¹ A model is perfectly calibrated if predictions assigned 80% confidence are correct approximately 80% of the time. Calibration is particularly critical for large language

models (LLMs) in high-stakes applications such as medical diagnosis, legal advice, and decision support, where miscalibration can amplify harm (Zhang et al., 2024b; Yang et al., 2025b,c; Hu et al., 2026). Well-calibrated LLMs make their reliability explicit and interpretable, improving downstream trust and safety.

Despite its importance, most calibration work on LLMs has focused on English (Xue et al., 2024). This reflects what Ruder et al. (2022) termed *Square One Bias*: research often advances along a single axis (e.g., English alignment or multilingual coverage), while their intersection (multilingual calibration) remains under-explored. Prior studies also have practical limitations: they cover only a few languages, often rely on machine-translated data, and evaluate a limited set of models (Xue et al., 2024; Yang et al., 2023).

To address this gap, we conduct a large-scale systematic study of multilingual calibration. Our analysis spans eight major model families and over 100 languages, using human-curated benchmarks (MMMLU and Belebele) with $\sim 10^5$ instances. Non-English languages consistently exhibit lower accuracy and worse calibration: LLaMA3’s average Expected Calibration Error (ECE) is 23.1% on non-English against 4.6% on English, a five-fold gap. The miscalibration pattern itself depends on training priorities. Prior work has largely framed LLM miscalibration as overconfidence (Zhang et al., 2024c; Chhikara, 2025); we observe two regimes. English-aligned models like LLaMA3 maintain calibration only for English; multilingual-first models like Aya are systematically over-confident across all languages, including English. Current alignment strategies thus fail to generalise to non-English contexts.

To identify the architectural source of this miscalibration, we examine the common practice of deriving confidence scores only from the final layer. Following work showing intermediate layers en-

* Equal contribution.

¹In this paper, we distinguish between calibration as a *property* and as a *process*. We use the term *calibration* to refer to the property of a model’s confidence being well-aligned with its accuracy. In contrast, the methods used to achieve this alignment are referred to as *calibration methods*.

code more language-agnostic representations (Bandarkar et al., 2025; Wendler et al., 2024), we hypothesise that the final layer, shaped by English-dominated training, provides a sub-optimal signal for multilingual confidence. Our layer-wise analysis reveals a dichotomy: English calibration improves *monotonically* with model depth and peaks at the final layer, whereas for nearly all non-English languages, *late-intermediate layers* provide better-calibrated confidence estimates than the final layer.

This insight motivates our methodological contribution: training-free calibration methods that leverage these intermediate representations. We compare three confidence elicitation strategies: the *best layer* method selects the single most calibrated intermediate layer; the *good-layers ensemble* aggregates signals from multiple layers; and *LACE* (Language-Aware Confidence Ensemble) adaptively tailors layer selection to each target language. The three methods yield consistent multilingual calibration improvements across all models, and combining *LACE* with post-hoc methods such as Temperature Scaling (Guo et al., 2017) yields further gains. Finally, we verify that both the multilingual gap and the late-intermediate peak hold in free-form generation (open-ended QA, chain-of-thought) and under alternative MCQA scoring protocols (§6).

Our contributions are fourfold: (1) An empirical analysis of multilingual LLM calibration across eight model families and over 100 languages, revealing systematic disparities between English and other languages. (2) The first layer-wise investigation of multilingual calibration, showing that intermediate layers offer a more reliable calibration signal than the final layer. (3) Training-free calibration methods that leverage intermediate representations and close part of the cross-lingual calibration gap. (4) Evidence that both findings hold beyond MCQA, in open-ended QA, chain-of-thought reasoning, and alternative MCQA scoring protocols.

2 Related Work

Multilingual Calibration Recent work has highlighted that modern LLMs, despite their strong performance, frequently struggle with calibration in their predictions (Xiong et al., 2024; Zhang et al., 2025b). Parallel studies document language-specific biases in LLMs (Zhang et al., 2024a; Qin et al., 2025; Zhou and Lu, 2025; Zhou et al., 2026a) and a broader linguistic hierarchy in the global

AI ecosystem (Occhini et al., 2026). Yet calibration in multilingual settings remains underexplored. Ahuja et al. (2022) first established that multilingual models like mBERT and XLM-R are poorly calibrated, especially for low-resource languages like Swahili. Xue et al. (2024) conducted a confidence estimation study across various models, covering both language-agnostic and language-specific tasks, but datasets in their study included only 5 languages and were machine-translated which can potentially import bias (Vanmassenhove et al., 2021; Choenni et al., 2024). We extend prior research in two key dimensions. First, we provide more comprehensive evaluation using human-curated datasets across 100+ languages and eight LLM families. Second, while existing studies focus exclusively on final-layer outputs, we investigate the internal architectural origins of calibration through a novel layer-wise analysis.

Layer-wise Representations Recent work studies how layers specialise in multilingual transformers. Intermediate layers often encode cross-lingual semantics in a largely language-agnostic shared space (Bandarkar et al., 2025; Zhou et al., 2026b). By contrast, later layers become more language-specific, adapting these representations to surface features such as syntax and word order. For predominantly English-trained LLMs (e.g., LLaMA), studies further suggest that multilingual inputs are mapped in middle layers to shared internal representations that often resemble English-like structures (Wendler et al., 2024; Kojima et al., 2024; Alabi et al., 2024). This view helps explain why prompting the model to “think in English” can improve cross-lingual generation: it matches the model’s internal pathway before decoding back into the target language (Shi et al., 2023; Zhang et al., 2024d). Building on these findings, we examine how layer-wise specialisation, especially the language-neutrality of intermediate layers, affects calibration across languages.

3 Benchmarking Multilingual Calibration

3.1 Experimental Setup

Task and Confidence Elicitation. We focus on multilingual Multiple-Choice Question Answering (MCQA) for two reasons. First, MCQA provides *unambiguous correctness labels*: in open-ended generation, correctness is approximated via metrics such as ROUGE, BLEU, or LLM-as-a-judge,

Language	LLaMA3				Aya			
	ECE↓	BRIER↓	AUROC↑	Accuracy	ECE↓	BRIER↓	AUROC↑	Accuracy
Arabic	33.06	24.37	61.00	38.20	28.41	33.79	71.49	45.20
Bengali	24.93	23.39	58.44	35.20	29.01	31.48	60.01	31.30
German	25.81	24.92	65.36	44.40	26.54	33.51	69.70	53.00
Spanish	18.21	21.89	71.65	52.00	28.17	31.86	71.12	51.10
French	13.87	22.75	71.39	51.30	23.80	32.72	70.69	53.40
Hindi	28.31	24.28	62.07	39.90	30.21	34.98	70.08	42.30
Indonesian	19.67	23.76	66.25	45.00	27.88	31.54	70.85	51.20
Italian	21.19	22.74	71.57	51.80	26.65	30.33	71.76	52.70
Japanese	28.36	27.27	61.73	43.00	16.30	26.26	69.92	46.70
Korean	30.86	25.06	62.59	42.50	32.07	37.09	72.06	45.00
Portuguese	10.51	21.76	71.37	50.40	27.33	31.42	70.71	53.50
Swahili	23.84	21.45	61.10	32.20	32.01	36.72	58.23	31.30
Yoruba	8.18	19.43	58.00	27.40	30.11	28.56	60.73	26.40
Chinese	41.94	19.56	50.63	23.10	17.12	28.75	67.35	52.20
English	4.61	17.63	80.36	61.20	20.66	25.30	74.65	57.40
Avg. Non-English	23.12	22.95	64.06	41.47	26.77	31.97	68.26	45.49
Avg. Low-Resource	23.00	22.78	61.14	36.32	29.60	32.84	65.23	37.95
Avg. High-Resource	21.71	22.62	67.41	46.63	24.29	30.80	70.88	51.67
Avg. Non-Latin-Script	27.44	23.10	59.44	35.19	26.90	32.20	66.23	40.05
Avg. Latin-Script	16.27	22.21	71.14	50.87	25.86	30.95	71.35	53.19
Average (All)	22.22	22.68	64.90	42.51	26.42	31.62	68.62	46.18

Table 1: Multilingual performance of **LLaMA3** (left) and **Aya** (right) on the MMMLU dataset. Metrics include ECE, Brier Score, AUROC, and Accuracy. All numbers are in percentages.

which often disagree and can flip conclusions depending on the metric (Ielanskyi et al., 2025). Second, MCQA admits a *clear definition of confidence* as the probability assigned to the selected option, whereas open-ended confidence estimation is method-dependent and unstable. Concretely, for each instance x with C choices we take the model’s prediction $\hat{y} = \arg \max_i p_i$ at the final layer and elicit confidence as $\text{Conf}(x) = \max_{i \in \{1, \dots, C\}} p_i$. We later validate that our findings generalise beyond MCQA on generative tasks (§6) and across alternative MCQA scoring protocols (Appendix F.4).

Datasets and Models. We use two high-quality, human-translated benchmarks: (1) **MMMLU** (Hendrycks et al., 2021) (15 languages) and (2) **Belebele** (Bandarkar et al., 2024) (122 languages). Together they cover a much larger range of languages and instances ($\sim 10^5$) than prior multilingual calibration studies. All experiments use eight-shot prompting in the respective language. We evaluate eight model families spanning 7–14B parameters: **LLaMA3**, **Qwen2.5**, **Qwen3**, **Mistral**, **Aya**, **DeepSeek**, **Phi**, and **Olmo3**. Full model identifiers and additional setup details are in Appendix A.

Metrics We evaluate calibration using expected calibration error (ECE; Guo et al., 2017) and the Brier score (Brier, 1950). To measure model’s ability to discriminate between correct and incorrect predictions, we also report AUROC (Fawcett,

Model	ECE↓		Δ	AUROC↑		Δ
	EN	Non-EN		EN	Non-EN	
LLaMA3	4.61	23.12	-18.51	80.36	64.06	+16.30
Aya	20.66	26.77	-6.11	74.65	68.26	+6.39
Mistral	23.92	34.24	-10.32	73.75	63.67	+10.08
Qwen2.5	15.77	21.22	-5.45	78.23	70.18	+8.05
Phi	20.48	27.05	-6.57	71.13	57.41	+13.72
DeepSeek	9.10	26.21	-17.11	66.21	57.14	+9.07
Qwen3	12.09	13.65	-1.56	82.26	73.41	+8.85
Olmo3	28.65	27.24	+1.41	64.70	56.37	+8.33

Table 2: Per-model MMMLU calibration. $\Delta = \text{EN} - \text{Non-EN}$ for both metrics. Per-language results are in Table 1 and Tables 9–12.

2006). We report ECE with number of bins $K = 10$; for bin size robustness, see Appendix F.2.

3.2 Results

Table 1 reports per-language MMMLU results for **LLaMA3** and **Aya**; Table 2 extends the comparison to the English-vs.-Non-English gap across all eight model families.

Not only are LLMs more accurate but also more calibrated in English. Table 1 shows that non-English languages underperform English in both accuracy and calibration: LLaMA3’s average ECE is 23.1% on non-English against 4.6% on English. The same pattern appears in Brier Score and AUROC and holds across all eight model families: Table 2 reports a positive AUROC Δ for every model, and a negative ECE Δ for every model except




Figure 1: Brier score vs. resource level for LLaMA3 on Belebele. Each point is one language; colour indicates writing system. Correlations: Spearman $\rho = -0.59$, $p < 10^{-8}$; Kendall $\tau = -0.43$, $p < 10^{-8}$; Pearson $r = -0.39$, $p < 0.001$.

Olmo3, which is poorly calibrated in English as well. The size of the ECE gap tracks multilingual pre-training coverage: Qwen3, pre-trained on 36T tokens spanning 119 languages (Yang et al., 2025a), shows the smallest English-favouring gap (1.56), while Qwen2.5, whose 18T-token corpus covers 29 languages (Qwen Team, 2024), shows 5.45; this within-family comparison partially controls for general model quality. English-first LLaMA3 shows the largest gap (18.51). Bootstrap confidence intervals for Table 2 are in Appendix F.1.

Calibration correlates with language resource availability. The calibration gap also tracks resource availability: in Table 1, Hindi and Swahili are worse calibrated than the high-resource group. Plotting resource level² against calibration on Belebele (Figure 1; per-language tables in Appendix A.4) yields strong negative correlations with Brier (Spearman $\rho = -0.59$, $p < 10^{-8}$) and positive correlations with AUROC ($\rho = 0.66$, $p < 10^{-12}$), suggesting that calibration is associated with the representation of a language in the pre-training corpus.

Miscalibration patterns reflect training and alignment priorities. Confidence distributions reveal two regimes. For LLaMA3 (Figure 2a), the English setting is well-calibrated (ECE = 4.6) while the non-English setting shows underconfidence. Aya’s distributions in English and non-English are similar in shape (Figure 2b) but right-skewed (overconfident) in both. These patterns track training priorities. LLaMA3’s alignment pipeline (supervised

fine-tuning, preference ranking, safety) (Grattafiori et al., 2024) is primarily English-focused. Aya prioritises multilingual coverage (Üstün et al., 2024; Dang et al., 2024) with less attention to calibration. The two regimes echo the *Square One Bias* (Ruder et al., 2022): progress along either axis (English alignment or multilingual coverage) leaves the orthogonal dimension under-served. Confidence behaviour of other models is in Appendix A.3.

4 Mid-Layers Reveal Better Calibration

Inspired by recent insights in layer-wise multilingual representations (Bandarkar et al., 2025; Wendler et al., 2024), we examine how confidence evolves throughout the model’s depth to understand the source of the poor calibration observed in §3. Specifically, we ask: *is it possible to identify intermediate representations that elicit better calibrated confidence scores w.r.t. final-layer accuracy?*

4.1 Early-Decoded Confidence Estimation

We adopt a layer-wise probing technique inspired by the early exiting paradigm (Elbayad et al., 2020) to offer a new way of confidence estimation. Instead of applying the modelling head only to the final hidden state, we attach it to each intermediate transformer layer. This allows us to extract logits and compute prediction confidence from every layer, providing a granular view of the model’s decision-making process.

Formally, let $\mathbf{h}_\ell \in \mathbb{R}^d$ denote the hidden representation at layer ℓ , where $\ell = 1, \dots, L$, and d is the dimensionality of the hidden state. We apply the original language modelling head, with weight matrix $W \in \mathbb{R}^{V \times d}$, to compute the logits at each

²LLaMA3 does not release its exact training data; we use Common Crawl (CC-MAIN-2025-30; Common Crawl Foundation, 2025) per-language page percentages as a proxy.




Figure 2: Density of model confidence scores for **English** and **Non-English** samples. Dashed lines mark mean confidence; solid lines mark accuracy.

layer:

$$\mathbf{z}_\ell = W\mathbf{h}_\ell$$

where $\mathbf{z}_\ell \in \mathbb{R}^V$ are the unnormalised token logits over the vocabulary of size V . These logits are then converted into probabilities using the softmax function:

$$\mathbf{p}_\ell = \text{softmax}(\mathbf{z}_\ell), \quad \sum_{v=1}^V [\mathbf{p}_\ell]_v = 1.$$

With the $\mathbf{p}_\ell$, we *trace* the token ultimately predicted at the final layer, $\hat{y}_L = \arg \max_v [\mathbf{p}_L]_v$, back through the intermediate layers. At each layer ℓ , we then define the confidence score as the probability mass that this layer assigns to the final prediction $\text{Conf}_\ell(x) = [\mathbf{p}_\ell]_{\hat{y}_L}$. Calibration is then evaluated by comparing these $\text{Conf}_\ell(x)$ with the prediction accuracy determined at the final layer $\hat{y}_L$:

$$\text{ECE}_\ell = \text{ECE}\left(\left\{(\text{Conf}_\ell(x), \mathbf{1}\{\hat{y}_L = y\})\right\}\right).$$

where $\mathbf{1}\{\hat{y}_L = y\}$ is the indicator function of whether the final-layer prediction is correct.

4.2 Layer-wise Multilingual Calibration Results

Multilingual settings reveal best calibration in late intermediate layers. The conventional understanding (Tenney et al., 2019) holds that intermediate representations become progressively more refined and task-specific, so calibration should improve with depth and peak at the final layer; we observe this pattern in the English-only setting (Appendix B.1). However, the multilingual picture differs: as Figure 3 shows, many languages achieve their best ECE in the late-intermediate layers (from layer 24 for LLaMA3, and from layer 26 for Aya, both 32-layer models), with calibration worsening again towards the final output layer. The turning point aligns with the entropy trend: as entropy drops sharply in these intermediate layers, ECE also decreases. Reliability diagrams are in Appendix B.3.

Why are late-intermediate layers better calibrated? Three observations point to a consistent interpretation. First, we see English calibration improves monotonically with depth while non-English calibration peaks before the final layer; Second, the region of late-intermediate layers is where prior work locates the most language-agnostic representations (Bandarkar et al., 2025; Wendler et al., 2024), and where models develop alignable geometry (Zhou et al., 2026b); a shared representation space supports a confidence signal that transfers across languages. Third, alignment updates act most directly on the output distribution, and the models with the most English-centric alignment show the largest gaps (§3). Together these are consistent with the final layer re-specialising towards English output distributions at the cost of multilingual confidence.

The mid-layer calibration peak is robust across model sizes and families. The same pattern holds across all eight model families we analyse layer-wise (Appendix B.2) and on LLaMA-3.1-70B (Appendix D): for most non-English languages, a late-intermediate layer yields lower ECE than the decoding layer. This questions the common practice of using final-layer probabilities as confidence, and motivates the layer-aware calibration strategies developed next.




Figure 3: ECE vs. entropy across layers on the MMMLU subset for LLaMA3 and Aya. See Figure 8 for the English-only curve.

5 Improving Multilingual Calibration

5.1 Multilingual Calibration Methods

Our observations in Section 4 motivate confidence estimators that exploit intermediate representations. We explore to extract confidence in three different ways:

- **Final layer (baseline).** We follow the standard practice in prior work by using the model’s final layer to derive probabilities.
- **Best layer.** We identify the *best layer* as the one that achieves the lowest average calibration error across languages. From this layer, we extract probability estimates following the procedure in Section 4.1. The best layer ℓ^* is selected using a separate validation set and defined as:

$$\ell^* = \arg \min_{\ell \in \{1, \dots, L\}} \text{ECE}_\ell^{\text{avg}}.$$

- **Good layers ensemble.** We consider the set of layers (*good layers*) whose multilingual calibration outperforms the final layer. To obtain confidence estimates, we average the predictive distributions across these layers $\mathcal{G}$, this reduces layer-specific noise while pooling calibration-aware signals:

$$\mathcal{G} = \{\ell : \text{ECE}_\ell^{\text{avg}} < \text{ECE}_L^{\text{avg}}\}, \mathbf{p}_{\text{ensemble}} = \frac{1}{|\mathcal{G}|} \sum_{\ell \in \mathcal{G}} \mathbf{p}_\ell.$$

Combining with Classical Post-hoc Calibration. Since confidence elicitation and calibration are orthogonal approaches, we further test whether the

proposed elicitation methods can be enhanced by standard post-hoc calibration techniques such as Temperature Scaling (Guo et al., 2017) and Isotonic Regression (Zadrozny and Elkan, 2002), which operate independently of how probabilities were obtained (Kadavath et al., 2022; Minderer et al., 2021). We adopt a two-stage pipeline:

$$\mathbf{p}_{\text{final}} = \text{Calibrate}(\mathbf{p}_{\text{intermediate}}),$$

where $\mathbf{p}_{\text{intermediate}}$ comes from the confidence scores discussed above. See Appendix C.1 for details.

Language-Aware Confidence Ensemble. The methods above optimise jointly across languages. *LACE* (Language-Aware Confidence Ensemble) instead tailors layer selection to each language. For each language k , we select the layers that are better calibrated than the final layer,

$$\mathcal{G}^{(k)} = \{\ell : \text{ECE}_\ell^{(k)} < \text{ECE}_L^{(k)}\},$$

$$\mathbf{p}_{\text{ensemble}}^{(k)} = \frac{1}{|\mathcal{G}^{(k)}|} \sum_{\ell \in \mathcal{G}^{(k)}} \mathbf{p}_\ell^{(k)},$$

and apply a language-specific post-hoc calibrator:

$$\mathbf{p}_{\text{final}}^{(k)} = \text{Calibrate}^{(k)}(\mathbf{p}_{\text{ensemble}}^{(k)}).$$

Per-language layer selection avoids negative transfer from layers that are miscalibrated for the target language, and ensembling reduces variance while preserving language-relevant signals. *LACE*




Figure 4: Forest plot of average ECE in MMMLU, with means and 95% CIs.

is training-free and adds negligible latency: it reuses intermediate logits from the single forward pass already required for generation, with no re-training or weight modification.

Experiment Setup We use the data from Section 4 as a held-out validation set and evaluate on a separate MMMLU test split. Both splits are balanced across languages, with a total of 30K examples. For Belebele, we construct a comparable evaluation set with 24K examples overall with a similar validation/test split. We report macro-averaged results across languages.

5.2 LACE Results

Intermediate-layer confidence shows better calibration than final-layer confidence. Figure 4 shows that moving from Best Layer to Good-Layers Ensemble to LACE yields progressive improvements; full numbers for MMMLU and Belebele are in Table 3. Occasional drops in AUROC arise because discrimination and calibration are not necessarily correlated (Gao et al., 2022; Carriero et al., 2025) and our layer selection optimises ECE. By tailoring layer selection and calibration to each language, LACE achieves the best overall results, with ECE as low as 5.96% on LLaMA3.

Classical post-hoc methods offer improvements but face challenges in multilingual settings. Temperature scaling and isotonic regression noticeably reduce calibration error on Aya, which is systematically overconfident (avg. confidence 76.9% on MMMLU, 82.5% on Belebele; Figure 2). However, they offer marginal benefit or even degradation for LLaMA3 (avg. confidence 36.0% / 58.3%): a globally fit temperature gives limited per-language gains under distributional heterogeneity (cf. Simpson, 1951).




Figure 5: Jaccard Similarity between different language’s layer-sets in LLaMA-3.1 70B.

LACE is complementary to post-hoc approaches and delivers the best calibration performance. Importantly, our method *reshapes*, rather than *rescales*, the calibration signal. Therefore, our method does not compete with post-hoc calibrators but provides orthogonal improvements. As shown in Table 3, all methods combined with temperature scaling or isotonic regression yields further incremental gains, and notably LACE is further boosted to consistently deliver the lowest calibration error across benchmarks: ECE to 3.09% on LLaMA3 and 3.45% on Aya.

5.3 Analysis

We confirm that LACE scales to larger models, and analyse robustness.

LACE scales effectively to larger models. Consistent with our findings at the 7–14B scales, our layer-based methods continue to outperform final-layer confidence on LLaMA-3.1-70B (Table 4; full results in Appendix D). LACE provides the largest overall improvement at this scale, reducing ECE from 23.51% (baseline) to as low as 2.81% when combined with Isotonic Regression.

Layer-set overlap and robustness across languages and datasets. To assess whether LACE’s layer sets generalise beyond specific samples or datasets, we examine their consistency across languages and benchmarks. Similar languages share overlapping layer subsets: Figure 5 shows clusters of high overlap between Western European languages and CJK (Chinese, Japanese, Korean) languages (red boxes). Layer sets derived from Belebele also retain strong calibration when transferred to MMMLU. The overlapping structure makes LACE robust to language-ID noise: in a controlled corruption experiment that randomly replaces the language tag of test examples, LACE degrades only

Method	MMMLU						Belebele					
	LLaMA3 (Acc. = 43.2%)			Aya (Acc. = 48.8%)			LLaMA3 (Acc. = 68.6%)			Aya (Acc. = 67.8%)		
	ECE	Brier	AUROC	ECE	Brier	AUROC	ECE	Brier	AUROC	ECE	Brier	AUROC
FINAL LAYER BASELINE [†]	22.44	23.03	64.05	24.39	30.45	68.31	17.68	20.38	69.01	15.73	19.08	72.92
▷ (+Temperature Scaling)	23.35	22.59	64.05	15.40	23.76	68.31	17.63	20.39	69.01	10.66	17.19	72.92
▷ (+Isotonic Regression)	20.23	22.55	63.47	9.15	22.02	68.07	11.09	19.07	68.85	8.33	16.44	72.78
<i>Intermediate Layer Calibration (Global Selection)</i>												
BEST LAYER [†]	14.28	22.78	71.44	17.57	27.08	66.68	13.67	18.28	71.33	15.40	20.06	66.79
▷ (+Temperature Scaling)	13.71	20.60	71.44	9.34	22.52	66.68	13.12	17.52	71.34	14.99	19.16	66.79
▷ (+Isotonic Regression)	13.12	20.80	71.28	10.66	22.46	66.26	12.40	17.39	71.48	14.02	19.00	66.80
GOOD LAYERS ENSEMBLE [†]	11.84	20.23	73.91	13.10	23.78	68.62	10.78	15.59	76.26	11.47	18.00	70.24
▷ (+Temperature Scaling)	11.30	20.01	73.91	10.23	22.21	68.62	10.49	15.62	76.26	11.81	17.57	70.24
▷ (+Isotonic Regression)	9.60	19.90	73.49	7.71	21.82	68.25	10.16	15.54	76.07	10.42	17.48	70.31
<i>Intermediate Layer Calibration (LACE)</i>												
LANGUAGE-AWARE ENSEMBLE [‡]	5.96	20.51	72.94	11.42	22.70	68.38	7.05	14.35	75.61	10.22	17.77	69.79
▷ (+Temperature Scaling)	4.34	19.73	73.40	4.88	21.87	68.49	6.05	14.47	74.98	5.46	16.36	70.45
▷ (+Isotonic Regression)	3.09	20.51	69.13	3.45	21.46	67.10	5.79	14.53	73.70	4.80	16.73	68.40

Table 3: Calibration results for **LLaMA3** and **Aya** on MMMLU (left) and the Belebele subset (right). [†] denote methods that do not assume access to language identity. [‡] denote methods with given language identity. Indented italic rows correspond to post-hoc calibration adjustments.

Method	ECE	Brier	AUROC
ORIG. PROB.	23.51	22.10	65.19
(+Temp. Scal.)	22.56	21.97	65.19
(+Isotonic)	19.01	21.34	65.07
LACE	6.14	15.76	79.42
(+Temp. Scal.)	4.93	15.60	79.92
(+Isotonic)	2.81	15.38	78.79

Table 4: Calibration for LLaMA-3.1 70B on MMMLU.

minimally across all metrics. Full discussion is in Appendix C.

6 Generalisation Beyond MCQA

Our main analysis focuses on MCQA because it offers unambiguous correctness labels and a well-defined notion of confidence (§3). To probe whether our two core findings (the multilingual calibration gap and the late-intermediate sweet spot) generalise to free-form generation, we extend the analysis to two generative benchmarks.

Setup. We use **MKQA** (Longpre et al., 2021) (open-ended multilingual QA) and **MGSM** (Shi et al., 2023) (multilingual chain-of-thought math) on **LLaMA3**. Correctness is scored against gold answers: for MKQA we apply Unicode-aware normalisation and require gold-answer containment; for MGSM we extract the final numeric token and require an exact match. Following Lewis-Lim et al. (2025), we evaluate four standard generative-confidence elicitation methods: *Perplexity* (geometric mean of next-token probabilities), *Margin* (top-1 vs. top-2 softmax gap, averaged across generated to-

Dataset	Group	Perp.	Margin	P(True)	Verb.
MKQA	English	13.5	10.6	19.2	26.9
	Avg. Non-Eng.	33.9	34.8	31.7	41.1
	Δ	-20.4	-24.2	-12.5	-14.2
MGSM	English	18.9	15.9	11.0	12.8
	Avg. Non-Eng.	42.9	41.1	12.6	23.8
	Δ	-24.0	-25.2	+1.6	-11.0

Table 5: Generative-task calibration on **LLaMA3**: ECE (%), $\downarrow$ for four confidence elicitation methods on MKQA and MGSM. $\Delta = \text{EN} - \text{Non-EN}$; negative values mean English is better-calibrated. Per-language numbers are in Appendix E.

kens), *P(True)* (Kadavath et al., 2022) (re-prompted true/false probability), and *Verbalised* (Tian et al., 2023) (expected confidence on a 1–9 scale, mapped to $[0, 1]$). Further details are in Appendix E.1.

The multilingual calibration gap persists in generation. Table 5 reports group-averaged ECE for each confidence method on the two datasets. English is the best-calibrated language for most cells. Averaging across the four methods, English ECE is 17.6% on MKQA and 14.7% on MGSM, while non-English ECE is 35.4% and 30.1%, roughly $2\times$ worse on both datasets. The gap is widest on the lowest-resource languages (Japanese on MKQA, avg. ECE 47.8%; Bengali on MGSM, avg. ECE 46.7%), mirroring the resource correlation we report on MCQA (§3).

Late-intermediate layers remain better-calibrated. We apply the early-decoded

Method	Lang	Acc.	Last	Best (ℓ^*)	ΔECE
Margin	en	64.0	18.0	11.9 ($\ell=29$)	-6.1
	de	55.0	26.2	9.9 ($\ell=29$)	-16.3
	es	55.0	26.7	6.0 ($\ell=29$)	-20.7
	fr	41.0	42.5	5.1 ($\ell=27$)	-37.4
	ru	53.0	30.3	8.2 ($\ell=30$)	-22.1
	ja	45.0	36.9	6.2 ($\ell=28$)	-30.7
	th	41.0	42.7	9.0 ($\ell=27$)	-33.7
	bn	18.0	66.0	1.6 ($\ell=23$)	-64.4
Perplexity	en	64.0	15.4	15.4 ($\ell=32$)	0.0
	de	55.0	29.1	10.7 ($\ell=31$)	-18.4
	es	55.0	28.7	11.3 ($\ell=31$)	-17.5
	fr	41.0	43.2	23.1 ($\ell=31$)	-20.1
	ru	53.0	29.5	10.2 ($\ell=31$)	-19.3
	ja	45.0	39.4	15.7 ($\ell=31$)	-23.6
	th	41.0	42.9	33.2 ($\ell=31$)	-9.6
	bn	18.0	69.1	8.6 ($\ell=30$)	-60.5

Table 6: Layer-wise calibration on MGSM with **LLaMA3** (32 layers). ECE in percent. Last = final-layer ECE; Best = minimum ECE over all layers (best-layer index in parentheses); ΔECE = Best - Last.




Figure 6: Layer-wise ECE on MGSM (LLaMA3) for en, fr, ru, and th. Left: margin confidence. Right: perplexity confidence. x -axis is layer index (0-32); y -axis is ECE.

confidence-estimation procedure of §4.1 to the two probability-space methods that admit a per-layer formulation (margin and perplexity; $P(\text{True})$ and Verbalised require re-prompting and therefore have no natural layer analogue). Table 6 reports the result for the eight MGSM languages: for every non-English (language $\times$ method) pair, a late-intermediate layer is better calibrated than the final layer, with ΔECE from -9.6 to -64.4 percentage points and the largest gains on Bengali. English follows its MCQA pattern (Appendix B.1): the margin gain is small (-6.1) and the final layer is already best for perplexity. For non-English languages the best layers concentrate in $\ell=23$ -31 on a 32-layer model, the same region we identified for MCQA. Figure 6 shows the per-layer ECE curves.

Layer selection without language identity. LACE requires a language-ID signal at inference

Method	MGSM			MKQA		
	EN	Non-EN	Avg	EN	Non-EN	Avg
Final layer [†]	14.75	38.41	35.45	15.52	25.72	24.44
Global layer [†]	17.57	25.13	24.18	25.02	24.03	24.15
Learned selector [†]	20.07	19.18	19.29	22.04	23.73	23.51
LACE [‡]	18.09	19.28	19.13	18.24	21.14	20.78

Table 7: Test ECE (%) on MGSM and MKQA with LLaMA3, averaged over 20 random validation/test splits. Global layer = a single layer tuned on validation data; Learned selector = a logistic regression over per-layer confidences (Appendix E.4). [†]: without language identity; [‡]: given language identity.

time. To test whether the layer choice can instead be inferred from the input alone, we train a logistic-regression selector that routes each example to a layer using only its per-layer confidence signature, with no access to language identity (Appendix E.4). Table 7 compares it with the final layer, a single language-agnostic layer tuned on validation data, and LACE, averaged over 20 random validation/test splits. On MGSM the selector matches LACE (average ECE 19.29 vs. 19.13) and halves the non-English ECE of the final layer (38.41 $\rightarrow$ 19.18); on MKQA it improves over both language-agnostic baselines, while LACE remains ahead (20.78 vs. 23.51).

Both findings also hold under two alternative MCQA scoring protocols (full-option-likelihood and text-generation; Appendix F.4).

7 Conclusion

This paper provides a full-scale evaluation of multilingual calibration in LLMs over 100 languages, uncovering disparities between English and other languages in both accuracy and reliability. Our layer-wise analysis identifies a mid-layer calibration peak: for non-English languages, late-intermediate layers are better-calibrated than the final layer. Leveraging this insight, we introduce LACE, a training-free, layer-aware ensemble that complements post-hoc techniques such as Temperature Scaling. Because these methods reuse intermediate logits without retraining, they can apply directly to deployed models to extract more reliable answering. We further show that the multilingual gap and the mid-layer effect generalise beyond MCQA to generative settings. We hope these results encourage further work on multilingual calibration to support more equitable language technologies.

Limitations

Our primary analysis is on MCQA. §6 extends the findings to MKQA and MGSIM, but a larger study across more models, languages, and generative confidence-aggregation schemes remains future work. LACE is post-hoc and requires a small per-language validation set (Appendix C.5). A natural next step is to address the gap at training time, e.g. via per-language ECE penalties in RLHF/DPO, or via mechanistic interventions (Templeton et al., 2024) that localise the final-layer English bias. LACE also selects layers at the language level; a per-instance scoring mechanism could further improve robustness for code-mixed inputs. Finally, our evaluation is statistical (ECE, Brier, AUROC); we do not measure whether improved calibration translates into better human decision-making under deployment, which would require a targeted user study.

Ethics Statement

Our research adheres to strict ethical guidelines. We verified the licenses of all software and datasets used in this study to ensure full compliance with their terms. No privacy concerns have been identified. We have conducted a thorough assessment of the project and do not anticipate any further risks. We only use LLMs for grammar checking during the paper writing.

Acknowledgements

This work was supported by the UKRI Frontier Research Grant EP/Y031350/1 (EQUATE). T.H. was supported by the Gates Cambridge Trust (grant OPP1144 from the Bill & Melinda Gates Foundation).

References

Marah Abidin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J. Hewett, Mojan Javaheripi, Piero Kauffmann, James R. Lee, Yin Tat Lee, Yuanzhi Li, Weishung Liu, Caio C. T. Mendes, Anh Nguyen, Eric Price, Gustavo de Rosa, Olli Saarikivi, and 8 others. 2024. [Phi-4 technical report](#).

Kabir Ahuja, Sunayana Sitaram, Sandipan Dandapat, and Monojit Choudhury. 2022. [On the calibration of massively multilingual language models](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 4310–4323, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Jesujoba Alabi, Marius Mosbach, Matan Eyal, Dietrich Klakow, and Mor Geva. 2024. [The hidden space of transformer language adapters](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6588–6607, Bangkok, Thailand. Association for Computational Linguistics.

Lucas Bandarkar, Davis Liang, Benjamin Muller, Mikel Artetxe, Satya Narayan Shukla, Donald Husa, Naman Goyal, Abhinandan Krishnan, Luke Zettlemoyer, and Madian Khabsa. 2024. [The belebele benchmark: a parallel reading comprehension dataset in 122 language variants](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, page 749–775. Association for Computational Linguistics.

Lucas Bandarkar, Benjamin Muller, Pritish Yuvraj, Rui Hou, Nayan Singhal, Hongjiang Lv, and Bing Liu. 2025. [Layer swapping for zero-shot cross-lingual transfer in large language models](#). In *International Conference on Learning Representations*.

Glenn W Brier. 1950. Verification of forecasts expressed in terms of probability. *Monthly weather review*, 78(1):1–3.

Alex Carriero, Kim Luijken, Anne de Hond, Karel G. M. Moons, Ben van Calster, and Maarten van Smeden. 2025. [The harms of class imbalance corrections for machine learning based prediction models: A simulation study](#). *Statistics in Medicine*, 44(3-4):e10320.

Prateek Chhikara. 2025. [Mind the confidence gap: Overconfidence, calibration, and distractor effects in large language models](#).

Rochelle Choenni, Sara Rajaei, Christof Monz, and Ekaterina Shutova. 2024. [On the evaluation practices in multilingual nlp: Can machine translation offer an alternative to human translations?](#) *Preprint*, arXiv:2406.14267.

Common Crawl Foundation. 2025. Common Crawl. <https://commoncrawl.org>. Accessed: 2025-07-26.

John Dang, Shivalika Singh, Daniel D’souza, Arash Ahmadian, Alejandro Salamanca, Madeline Smith, Aidan Peppin, Sungjin Hong, Manoj Govindassamy, Terrence Zhao, Sandra Kublik, Meor Amer, Viraat Aryabumi, Jon Ander Campos, Yi-Chern Tan, Tom Kocmi, Florian Strub, Nathan Grinsztajn, Yannis Flet-Berliac, and 26 others. 2024. [Aya expand: Combining research breakthroughs for a new multilingual frontier](#).

DeepSeek-AI. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#).

Maha Elbayad, Jiatao Gu, Edouard Grave, and Michael Auli. 2020. [Depth-adaptive transformer](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

- Tom Fawcett. 2006. [An introduction to roc analysis](#). *Pattern Recognition Letters*, 27(8):861–874. ROC Analysis in Pattern Recognition.
- Riqiang Gao, Thomas Li, Yucheng Tang, Zhoubing Xu, Michael Kammer, Sanja L. Antic, Kim Sandler, Fabien Moldonado, Thomas A. Lasko, and Bennett Landman. 2022. [A comparative study of confidence calibration in deep learning: From computer vision to medical imaging](#).
- Jiahui Geng, Fengyu Cai, Yuxia Wang, Heinz Koepl, Preslav Nakov, and Iryna Gurevych. 2024. [A survey of confidence estimation and calibration in large language models](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6577–6595, Mexico City, Mexico. Association for Computational Linguistics.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. [The llama 3 herd of models](#). *ArXiv preprint*, abs/2407.21783.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. 2017. [On calibration of modern neural networks](#). In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 1321–1330. PMLR.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Songbo Hu, Yinhong Liu, Ej Zhou, Evgeniia Razu-movskaia, Xiaobin Wang, Alexander Fraser, Ivan Vulić, and Anna Korhonen. 2026. [Dial HEALTH-DIAL for advice: A multilingual and multi-parallel spoken dialogue dataset for knowledge-grounded information seeking](#). In *Findings of the Association for Computational Linguistics: ACL 2026*, pages 25529–25552, San Diego, California, United States. Association for Computational Linguistics.
- Mykyta Ielanskyi, Kajetan Schweighofer, Lukas Aichberger, and Sepp Hochreiter. 2025. Addressing pitfalls in the evaluation of uncertainty estimation methods for natural language generation. In *ICLR Workshop: Quantify Uncertainty and Hallucination in Foundation Models: The Next Frontier in Reliable AI*.
- 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](#).
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-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](#). *Preprint*, arXiv:2207.05221.
- Takeshi Kojima, Itsuki Okimura, Yusuke Iwasawa, Hitomi Yanaka, and Yutaka Matsuo. 2024. [On the multilingual ability of decoder-based pre-trained language models: Finding and controlling language-specific neurons](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6919–6971, Mexico City, Mexico. Association for Computational Linguistics.
- Samuel Lewis-Lim, Xingwei Tan, Zhixue Zhao, and Nikolaos Aletras. 2025. [Analysing chain of thought dynamics: Active guidance or unfaithful post-hoc rationalisation?](#) In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 29838–29853, Suzhou, China. Association for Computational Linguistics.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. [Teaching models to express their uncertainty in words](#). *Preprint*, arXiv:2205.14334.
- Shayne Longpre, Yi Lu, and Joachim Daiber. 2021. [MKQA: A linguistically diverse benchmark for multilingual open domain question answering](#). *Transactions of the Association for Computational Linguistics*, 9:1389–1406.
- Matthias Minderer, Josip Djolonga, Rob Romijnders, Frances Hubis, Xiaohua Zhai, Neil Houlsby, Dustin Tran, and Mario Lucic. 2021. [Revisiting the calibration of modern neural networks](#). In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 15682–15694.
- Giulia Occhini, Kumiko Tanaka-Ishii, Anna Barford, Refael Tikochinski, Songbo Hu, Roi Reichart, Yijie Zhou, Hannah Claus, Ulla Petti, Ivan Vulić, Ramit Debnath, and Anna Korhonen. 2026. [Artificial intelligence is creating a new global linguistic hierarchy](#). *Preprint*, arXiv:2602.12018.
- Team Olmo, Allyson Ettinger, Amanda Bertsch, Bailey Kuehl, David Graham, David Heineman, Dirk Groeneveld, Faeze Brahman, Finbarr Timbers, Hamish Ivison, Jacob Morrison, Jake Poznanski, Kyle Lo, Luca Soldaini, Matt Jordan, Mayee Chen, Michael Noukhovitch, Nathan Lambert, Pete Walsh, and 49 others. 2026. [Olmo 3](#). *Preprint*, arXiv:2512.13961.

- Libo Qin, Qiguang Chen, Yuhang Zhou, Zhi Chen, Yinghui Li, Lizi Liao, Min Li, Wanxiang Che, and Philip S Yu. 2025. A survey of multilingual large language models. *Patterns*, 6(1).
- Qwen Team. 2024. [Qwen2.5 technical report](#). *Preprint*, arXiv:2412.15115.
- Sebastian Ruder, Ivan Vulić, and Anders Søgaard. 2022. [Square one bias in NLP: Towards a multi-dimensional exploration of the research manifold](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2340–2354, Dublin, Ireland. Association for Computational Linguistics.
- Freda Shi, Mirac Suzgun, Markus Freitag, Xuezhi Wang, Suraj Srivats, Soroush Vosoughi, Hyung Won Chung, Yi Tay, Sebastian Ruder, Denny Zhou, Dipanjan Das, and Jason Wei. 2023. [Language models are multi-lingual chain-of-thought reasoners](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- E. H. Simpson. 1951. [The interpretation of interaction in contingency tables](#). *Journal of the Royal Statistical Society: Series B (Methodological)*, 13(2):238–241.
- Pedro Ortiz Suarez, Laurie Burchell, Catherine Arnett, Rafael Mosquera, Sara Hincapié Monsalve, Thom Vaughan, Damian Stewart, Malte Ostendorff, Idris Abdulmumin, Vukosi Marivate, Shamsuddeen Hassan Muhammad, Atnafu Lambebo Tonja, Hend Al-Khalifa, Nadia Ghezaiel Hammouda, Verrah Akinyi Otiende, Tack Hwa Wong, Jakhongir Saydaliev, Melika Nobakhtian, Muhammad Ravi Shulthan Habibi, and 78 others. 2026. [CommonLID: Re-evaluating state-of-the-art language identification performance on web data](#). In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 33063–33080, San Diego, California, United States. Association for Computational Linguistics.
- Adly Templeton, Tom Conerly, Jonathan Marcus, Jack Lindsey, Trenton Bricken, Brian Chen, Adam Pearce, Craig Citro, Emmanuel Ameisen, Andy Jermyn, Catherine Olsson, Tristan Hume, Jared Kaplan, Tom Henighan, Sam McCandlish, and Chris Olah. 2024. Scaling monosemanticity: Extracting interpretable features from Claude 3 Sonnet. <https://transformer-circuits.pub/2024/scaling-monosemanticity/>. Anthropic Research.
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. [BERT rediscovers the classical NLP pipeline](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601, Florence, Italy. Association for Computational Linguistics.
- Katherine Tian, Eric Mitchell, Allan Zhou, Archit Sharma, Rafael Rafailov, Huaxiu Yao, Chelsea Finn, and Christopher Manning. 2023. [Just ask for calibration: Strategies for eliciting calibrated confidence scores from language models fine-tuned with human feedback](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5433–5442, Singapore. Association for Computational Linguistics.
- Ahmet Üstün, Viraat Aryabumi, Zheng Yong, Wei-Yin Ko, Daniel D’souza, Gbemileke Onilude, Neel Bhandari, Shivalika Singh, Hui-Lee Ooi, Amr Kayid, Freddie Vargus, Phil Blunsom, Shayne Longpre, Niklas Muennighoff, Marzieh Fadaee, Julia Kreutzer, and Sara Hooker. 2024. [Aya model: An instruction fine-tuned open-access multilingual language model](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15894–15939, Bangkok, Thailand. Association for Computational Linguistics.
- Eva Vanmassenhove, Dimitar Shterionov, and Matthew Gwilliam. 2021. [Machine translationese: Effects of algorithmic bias on linguistic complexity in machine translation](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2203–2213, Online. Association for Computational Linguistics.
- Chris Wendler, Veniamin Veselovsky, Giovanni Monea, and Robert West. 2024. [Do llamas work in English? on the latent language of multilingual transformers](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15366–15394, Bangkok, Thailand. Association for Computational Linguistics.
- Miao Xiong, Zhiyuan Hu, Xinyang Lu, Yifei Li, Jie Fu, Junxian He, and Bryan Hooi. 2024. [Can llms express their uncertainty? an empirical evaluation of confidence elicitation in llms](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Boyang Xue, Hongru Wang, Rui Wang, Sheng Wang, Zezhong Wang, Yiming Du, Bin Liang, and Kam-Fai Wong. 2024. [Mlingconf: A comprehensive study of multilingual confidence estimation on large language models](#). *ArXiv preprint*, abs/2410.12478.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zhang, 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.
- Ruihan Yang, Caiqi Zhang, Zhisong Zhang, Xinting Huang, Sen Yang, Nigel Collier, Dong Yu, and Deqing Yang. 2025b. [LoGU: Long-form generation with uncertainty expressions](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 18947–18968, Vienna, Austria. Association for Computational Linguistics.
- Ruihan Yang, Caiqi Zhang, Zhisong Zhang, Xinting Huang, Dong Yu, Nigel Collier, and Deqing Yang.

- 2025c. [UNCLE: Benchmarking uncertainty expressions in long-form generation](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 30340–30356, Suzhou, China. Association for Computational Linguistics.
- Yahan Yang, Soham Dan, Dan Roth, and Insup Lee. 2023. [On the calibration of multilingual question answering llms](#). *Preprint*, arXiv:2311.08669.
- Bianca Zadrozny and Charles Elkan. 2002. [Transforming classifier scores into accurate multiclass probability estimates](#). *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- Caiqi Zhang, Zhijiang Guo, and Andreas Vlachos. 2024a. [Do we need language-specific fact-checking models? the case of Chinese](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 1899–1914, Miami, Florida, USA. Association for Computational Linguistics.
- Caiqi Zhang, Fangyu Liu, Marco Basaldella, and Nigel Collier. 2024b. [LUQ: Long-text uncertainty quantification for LLMs](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 5244–5262, Miami, Florida, USA. Association for Computational Linguistics.
- Caiqi Zhang, Chang Shu, Ehsan Shareghi, and Nigel Collier. 2025a. [All roads lead to Rome: Graph-based confidence estimation for large language model reasoning](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 31814–31824, Suzhou, China. Association for Computational Linguistics.
- Caiqi Zhang, Ruihan Yang, Zhisong Zhang, Xinting Huang, Sen Yang, Dong Yu, and Nigel Collier. 2025b. [Atomic calibration of LLMs in long-form generations](#). In *Proceedings of the 14th International Joint Conference on Natural Language Processing and the 4th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics*, pages 148–169, Mumbai, India. The Asian Federation of Natural Language Processing and The Association for Computational Linguistics.
- Caiqi Zhang, Xiaochen Zhu, Chengzu Li, Nigel Collier, and Andreas Vlachos. 2026. [LoVeC: Reinforcement learning for better verbalized confidence in long-form generation](#). In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 33336–33363, San Diego, California, United States. Association for Computational Linguistics.
- Mozhi Zhang, Mianqiu Huang, Rundong Shi, Linsen Guo, Chong Peng, Peng Yan, Yaqian Zhou, and Xipeng Qiu. 2024c. [Calibrating the confidence of large language models by eliciting fidelity](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 2959–2979, Miami, Florida, USA. Association for Computational Linguistics.
- Zhihan Zhang, Dong-Ho Lee, Yuwei Fang, Wenhao Yu, Mengzhao Jia, Meng Jiang, and Francesco Barbieri. 2024d. [PLUG: Leveraging pivot language in cross-lingual instruction tuning](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7025–7046, Bangkok, Thailand. Association for Computational Linguistics.
- Ej Zhou and Weiming Lu. 2025. Bias beyond english: Evaluating social bias and debiasing methods in a low-resource setting. In *Natural Language Processing and Chinese Computing*, pages 214–227, Singapore. Springer Nature Singapore.
- Ej Zhou, Lucas Resck, Zheng Hui, and Anna Korhonen. 2026a. [Lower-resource, higher scores: Language bias in LLM evaluators](#). In *The Third Annual Conference on Language Modeling*.
- Ej Zhou, Suchir Salhan, Catherine Arnett, and Anna Korhonen. 2026b. [Cross-lingual alignment without joint training: Do monolingual language models converge on universal representations?](#) In *Proceedings of the 2026 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.

A Benchmarking Multilingual Calibration

In this section, we present the detailed multilingual evaluation results for the models and benchmarks discussed in the main text.

A.1 Experimental Details

A.1.1 Task Motivation and Confidence Elicitation

In this study, we focus on MCQA for two key reasons: first, MCQA provides unambiguous correctness labels for calibration. In open-ended generation, correctness is typically approximated using metrics such as ROUGE, BLEU, or LLM-as-a-judge, which often disagree and can even lead to flipped conclusions depending on the choice of metric (Ielanskyi et al., 2025). Second, MCQA enables a clear definition of confidence as the probability assigned to the selected option. In contrast, confidence estimation for open-ended generation remains unstable and method-dependent (e.g., verbalised confidence, perplexity-based methods, or sampling-based consistency). By using MCQA, we can isolate the effect of layer-wise signals on calibration without additional noise from generation-based confidence heuristics.

Correctness Calculation and Confidence Elicitation We take the output of the final layer as the model answer when evaluating correctness. We do not compute calibration using each layer’s own prediction, because our goal is to provide a reliability signal for the deployed model in a standard inference setting, where the user only observes the final-layer answer. Therefore, the confidence score should approximate *the probability that this final-layer answer is correct*. We follow the standard MCQA practice and use the output probability of the selected answer choice: $\text{Conf}(x) = \max_{i \in \{1, \dots, C\}} p_i$, where $\sum_{i=1}^C p_i = 1$, where C is the number of choices.

A.1.2 Models

Our experiments evaluate eight model families, plus a larger LLaMA-3.1-70B variant for scale analysis:

- **LLaMA3** (Grattafiori et al., 2024) (if **not** specified: Llama-3.1-8BInstruct)
- **Qwen2.5** (Qwen Team, 2024) (Qwen2.5-7B-Instruct)
- **Qwen3** (Yang et al., 2025a) (Qwen3-8B)

- **Mistral** (Jiang et al., 2023) (Mistral-7B-Instruct-v0.3)
- **Aya** (Dang et al., 2024) (aya-expanse-8b)
- **DeepSeek** (DeepSeek-AI, 2025) (DeepSeek-R1-DistillQwen-7B)
- **Phi** (Abdin et al., 2024) (phi-4)
- **Olmo3** (Olmo et al., 2026) (Olmo-3-7B)
- **LLaMA-3.1 70B** (Grattafiori et al., 2024) (Llama-3.1-70BInstruct)

A.1.3 Language Group Definitions

We randomly sampled 1,000 examples per language for MMMLU. We group languages in the MMMLU dataset according to resource availability and script type, as summarised in Table 8.

Group	Languages
Low-Resource	Arabic, Bengali, Swahili, Yoruba, Hindi, Indonesian
High-Resource	German, French, English, Spanish, Chinese, Italian, Japanese, Korean, Portuguese
Latin-Script	German, English, Spanish, French, Indonesian, Italian, Portuguese, Swahili, Yoruba
Non-Latin-Script	Arabic, Bengali, Hindi, Japanese, Korean, Chinese

Table 8: MMMLU language group definitions by resource availability and script type.

Language Codes The FLORES-200 language codes (3-letter form) for all evaluated languages are provided in the Belebele dataset documentation (Bandarkar et al., 2024).

A.2 MMMLU Results

Results for each model are reported in the following tables: Mistral (Table 9), Qwen2.5 (Table 10), Phi (Table 11), and DeepSeek (Table 12).

A.3 Model Confidence Behaviours

Figure 7 illustrates the distribution of confidence scores and accuracies across English and non-English settings for Qwen2.5, DeepSeek, Mistral, and Phi. Solid vertical lines indicate mean accuracies, while dashed vertical lines indicate mean confidences. The divergence between confidence and accuracy highlights calibration behaviour: underconfidence when the dashed line falls left of the solid line, and overconfidence when it falls to the right.

Language	AUROC	ECE	BRIER	Accuracy
Arabic	64.91	41.18	11.87	4.50
Bengali	64.56	49.70	11.72	0.10
German	70.84	24.14	29.32	43.00
English	73.75	23.92	27.95	54.00
Spanish	71.33	21.64	26.79	42.90
French	71.25	22.20	28.36	46.40
Hindi	75.08	39.77	6.23	1.60
Indonesian	69.48	26.98	29.69	38.80
Italian	74.08	25.24	28.25	44.50
Japanese	56.09	44.15	15.48	6.50
Korean	39.78	46.62	16.25	5.50
Portuguese	71.11	29.25	27.59	47.10
Swahili	56.02	30.81	27.34	26.30
Yoruba	44.79	44.18	21.99	16.10
Chinese	62.12	33.55	24.58	16.70
<i>Avg. Low-Resource</i>	62.47	38.77	18.14	14.57
<i>Avg. High-Resource</i>	65.59	30.08	24.95	34.07
<i>Avg. Latin-Script</i>	71.69	24.77	28.28	45.24
<i>Avg. Non-Latin-Script</i>	57.92	41.24	16.93	9.66
Average (All Languages)	64.35	33.56	22.23	26.27

Table 9: Performance comparison across languages for AUROC, ECE, BRIER score, and Accuracy in **Mistral**, evaluated on the MMMLU dataset.

Language	AUROC	ECE	BRIER	Accuracy
Arabic	67.15	14.30	26.67	54.90
Bengali	64.10	26.68	31.98	33.20
German	76.94	21.59	25.08	55.60
English	78.23	15.77	19.25	65.60
Spanish	76.95	19.26	23.98	61.10
French	75.65	16.92	22.88	62.20
Hindi	72.01	28.73	28.86	33.90
Indonesian	75.69	15.83	23.53	54.30
Italian	75.32	21.07	24.46	58.70
Japanese	80.03	6.71	17.10	33.10
Korean	74.15	17.60	25.75	52.20
Portuguese	75.85	18.86	23.61	58.40
Swahili	59.93	30.12	33.09	32.30
Yoruba	23.49	46.99	36.11	2.00
Chinese	85.31	12.47	17.42	47.00
<i>Avg. Low-Resource</i>	60.40	27.11	30.04	35.10
<i>Avg. High-Resource</i>	77.60	16.69	22.17	54.88
<i>Avg. Latin-Script</i>	76.38	18.47	23.26	59.41
<i>Avg. Non-Latin-Script</i>	65.77	22.95	27.12	36.08
Average (All Languages)	70.72	20.86	25.32	46.97

Table 10: Performance comparison across languages for AUROC, ECE, BRIER score, and Accuracy in **Qwen 2.5**, evaluated on the MMMLU dataset.

Table 13 reports detailed calibration and confidence statistics for LLaMA3. In addition to standard accuracy, we provide the model’s average confidence, the confidence gap (accuracy minus confidence), the proportion of correct predictions made with low confidence (“Underconf”), and the mean confidence levels assigned to correct vs. incorrect predictions. We further include the difference between these two distributions (“Corr–Inc Gap”), which captures how well the model separates correct from incorrect responses. English shows a relatively small confidence gap (2.5%), with strong separation between correct and incorrect predictions (23.8% Corr–Inc Gap). In contrast, most non-English languages show lower accuracy,

Language	AUROC	ECE	BRIER	Accuracy
Arabic	52.66	30.21	25.35	36.50
Bengali	52.62	34.13	24.73	27.20
German	63.47	22.86	22.86	65.60
English	71.13	20.48	17.92	73.10
Spanish	61.29	27.15	25.32	56.40
French	71.57	17.07	20.21	68.90
Hindi	37.74	46.43	26.16	15.70
Indonesian	42.89	32.36	30.63	30.70
Italian	72.25	10.51	19.13	67.50
Japanese	30.62	46.69	17.59	8.30
Korean	66.95	29.00	24.50	50.00
Portuguese	73.79	13.24	18.77	66.60
Swahili	64.42	16.18	23.61	40.50
Yoruba	53.76	20.83	21.01	27.60
Chinese	59.73	31.98	26.17	44.60
<i>Avg. Low-Resource</i>	50.68	30.02	25.25	29.70
<i>Avg. High-Resource</i>	63.42	24.33	21.39	55.67
<i>Avg. Latin-Script</i>	65.20	20.52	22.12	61.26
<i>Avg. Non-Latin-Script</i>	52.31	31.93	23.64	31.30
Average (All Languages)	58.33	26.61	22.93	45.28

Table 11: Performance comparison across languages for AUROC, ECE, BRIER score, and Accuracy in **Phi**, evaluated on the MMMLU dataset.

Language	AUROC	ECE	BRIER	Accuracy
Arabic	55.33	32.74	21.54	26.40
Bengali	58.50	40.80	14.41	13.70
German	60.28	18.50	23.91	39.80
English	66.21	9.10	22.92	47.10
Spanish	62.24	12.47	23.51	40.80
French	62.93	10.84	23.12	41.40
Hindi	56.08	30.42	20.62	26.40
Indonesian	61.00	31.61	21.11	27.30
Italian	63.14	5.65	22.85	40.40
Japanese	55.56	18.05	23.14	32.10
Korean	21.56	49.09	18.66	1.10
Portuguese	62.37	16.78	23.26	39.10
Swahili	51.67	45.76	12.45	12.00
Yoruba	60.35	38.16	4.94	2.80
Chinese	69.00	16.13	23.97	43.10
<i>Avg. Low-Resource</i>	57.16	36.58	15.84	18.10
<i>Avg. High-Resource</i>	58.14	17.40	22.82	36.10
<i>Avg. Latin-Script</i>	62.60	14.99	22.95	39.41
<i>Avg. Non-Latin-Script</i>	53.51	33.89	17.47	19.70
Average (All Languages)	57.75	25.07	20.03	28.90

Table 12: Performance comparison across languages for AUROC, ECE, BRIER score, and Accuracy in **DeepSeek**, evaluated on the MMMLU dataset.

larger underconfidence rates, and much smaller separation between correct and incorrect predictions (average Corr–Inc Gap of 6.3%).

A.4 Belebele Results

Belebele (Bandarkar et al., 2024) is a multiple-choice dataset covering 122 language variants, enabling robust evaluation of NLU across high-, medium-, and low-resource languages. The dataset is fully parallel, allowing for direct cross-linguistic comparison of model performance. In our experiments, we sample 400 examples per language. Per-language calibration metrics on Belebele are reported in Table 14 for LLaMA3.




Figure 7: Distribution of confidence scores (predicted probabilities) versus accuracies in English (blue) and non-English (orange) for four models: (a) Qwen 2.5, (b) DeepSeek, (c) Mistral, and (d) Phi.

Language	Acc. (%)	Avg Conf (%)	Conf Gap (%)	Underconf (%)	Corr. Conf (%)	Inc. Conf (%)	Corr-Inc Gap (%)
Arabic	38.2	29.5	8.8	89.0	31.7	28.0	3.7
Chinese	23.1	23.2	-0.1	99.6	22.2	23.6	-1.4
Korean	42.5	31.2	11.3	92.0	33.5	29.5	4.0
Japanese	43.0	24.0	19.0	92.3	26.8	21.8	5.0
Swahili	32.2	31.5	0.7	88.8	34.2	30.3	3.9
Italian	51.8	41.8	10.0	58.3	47.4	35.8	11.6
Bengali	35.2	34.6	0.6	88.4	36.1	33.8	2.3
Spanish	52.0	46.1	5.9	50.0	52.6	39.0	13.6
Portuguese	50.4	47.1	3.3	47.0	53.6	40.6	13.0
Indonesian	45.0	37.0	8.0	75.3	41.1	33.7	7.4
Hindi	39.9	31.7	8.2	88.2	34.1	30.1	4.0
German	44.4	33.8	10.6	80.4	37.1	31.2	5.9
Yoruba	27.4	29.4	-2.0	93.8	31.5	28.7	2.8
French	51.3	41.5	9.8	59.5	47.4	35.3	12.1
English	61.2	58.8	2.5	25.7	68.0	44.2	23.8
Non-English	41.2	34.5	6.7	78.8	37.8	31.5	6.3

Table 13: LLaMA3 Calibration and underconfidence analysis across languages. Metrics include accuracy, average confidence, confidence gap (accuracy minus confidence), proportion of underconfident correct predictions (confidence < 0.5 when correct), average confidence for correct vs. incorrect predictions, and their difference (Corr-Inc Gap).

B Layer-Wise Calibration Analysis

B.1 English Calibration Improves as Layer Deepens

Figure 8 shows the layer-wise Expected Calibration Error (ECE) for English in LLaMA3, illustrating how calibration improves progressively in deeper layers.

B.2 Multilingual Calibration is Best at Late-Intermediate Layers

We visualise calibration performance across layers by plotting metrics against entropy on the MMMLU dataset for LLaMA3 (Figure 9a), Aya (Figure 9b), Mistral (Figure 9c), Phi (Figure 10a), DeepSeek (Figure 10b), and Qwen2.5 (Figure 10c).




Figure 8: Layer-wise predicted confidence ECE for English in LLaMA3.

B.3 Reliability Diagrams

Per-language calibration trends reveal that late-intermediate layer improves calibration for most non-English languages, with a slight trade-off for English. To further explore these dynamics, Figure 11 presents per-language reliability diagrams for nine languages, comparing the selected intermediate layer (Layer 29) against the final layer of LLaMA3. Nearly all non-English languages benefit greatly from moving away from the final layer: their reliability curves align more closely with the diagonal, and ECE decreases substantially. For example, German and Hindi show reductions in ECE of more than 13% and 16%, respectively. In contrast, calibration for English appears to degrade slightly at intermediate layers. Since the final layer already exhibits strong calibration for English, earlier layers offer no additional benefit. This highlights a potential bias introduced during pretraining, where the model overfits to English patterns or introduces noise during the final stages of adaptation. While the degree of improvement varies across languages, the trend remains consistent across diverse linguistic families and scripts. This suggests that the effect is not language-specific but rather a systematic property of multilingual calibration.

C LACE — Supporting Details & Robustness

This section provides implementation details for the post-hoc calibration baselines used alongside LACE, followed by a series of robustness and generalisation analyses.

C.1 Post-hoc Calibration Baselines

We include two widely used post-hoc calibration methods as baselines: **Temperature Scaling** (Guo et al., 2017) and **Isotonic Regression** (Zadrozny and Elkan, 2002). Both are trained on a held-out validation set (15k examples from MMMLU and 12k examples from Belebele, non-overlapping with the evaluation sets). The fitted calibrators are then

Set 1					Set 2					Set 3				
Lang	Acc	AUR.	ECE	Brier	Lang	Acc	AUR.	ECE	Brier	Lang	Acc	AUR.	ECE	Brier
acm	58.2	78.4	7.1	19.2	arz	69.8	77.8	12.2	18.2	ceb	65.8	75.6	11.1	19.7
fin	80.8	75.6	6.0	14.1	hin	67.0	72.9	14.0	20.6	ita	86.0	79.6	7.4	10.8
khm	4.0	5.5	66.2	51.4	lvs	74.8	79.7	11.3	16.2	npi	59.8	74.8	5.1	20.0
pol	80.0	78.9	6.3	13.3	slv	81.0	76.6	8.7	13.8	swe	79.2	81.6	7.3	12.8
tso	34.0	62.5	3.4	21.5	xho	37.0	60.0	5.0	22.8	afr	80.8	81.2	8.3	12.7
asm	45.8	68.0	6.1	22.8	ces	82.2	80.2	7.8	11.8	fra	86.2	77.4	9.9	11.4
hin	57.0	72.1	4.7	21.1	jav	68.0	78.4	11.1	18.4	kin	34.8	63.3	4.7	21.3
mal	60.5	74.0	12.3	21.3	npi	32.5	59.9	6.3	21.4	por	86.2	79.8	5.6	10.2
sna	35.8	58.6	6.2	23.0	swl	67.0	75.0	8.1	19.1	tur	78.2	78.8	8.5	14.6
yor	31.2	61.1	4.4	20.4	als	73.5	78.0	7.2	16.2	azj	66.8	71.8	14.3	21.4
ckb	46.0	71.8	8.3	22.2	fuv	28.0	51.9	5.2	20.4	hrv	79.8	78.1	8.2	14.3
jpn	66.5	73.3	29.9	28.0	kir	63.0	73.8	21.9	24.7	mar	67.5	73.5	9.2	19.4
nso	37.8	60.3	3.2	22.7	snd	17.5	50.5	30.0	25.3	tam	65.5	73.4	14.5	21.5
ukr	84.2	77.4	12.1	12.8	zho	76.5	71.2	30.6	25.2	amh	34.8	63.2	5.0	21.5
bam	31.2	60.6	4.8	20.8	dan	79.8	78.8	8.2	13.9	gaz	31.8	53.1	2.4	21.7
hun	82.5	84.0	9.7	11.8	kac	30.2	61.8	3.5	20.5	kor	77.8	77.9	14.5	16.3
mkd	77.8	79.1	7.8	14.4	nya	32.0	60.7	3.3	21.1	ron	80.0	80.4	6.3	13.0
som	35.2	59.0	3.9	22.2	tel	59.5	73.7	7.2	20.4	urd	59.5	67.3	19.7	25.8
zho	81.2	68.9	24.4	21.1	apc	65.0	78.3	9.7	18.4	ben	65.5	72.6	10.1	20.3
deu	86.8	72.3	13.4	12.9	grn	39.8	65.5	6.6	22.4	hye	0.2	1.5	63.8	42.1
kan	58.5	72.1	5.2	20.9	lao	32.5	59.5	2.1	21.5	mlt	69.8	76.4	9.9	18.4
ory	55.8	70.0	16.4	24.7	rus	81.2	81.0	10.4	13.0	sot	32.8	57.5	2.6	21.8
tgk	63.8	70.1	12.5	22.1	urd	41.2	64.8	4.3	22.5	zsm	82.5	82.9	9.5	11.8
arb	79.5	74.2	11.8	15.5	ben	35.2	59.9	4.9	22.4	ell	80.5	80.6	10.1	13.7
guj	58.0	68.8	8.2	22.0	ibo	40.2	62.0	3.7	22.6	kat	1.5	6.0	68.4	51.6
lin	34.2	59.3	3.5	21.9	mri	35.5	63.3	4.3	21.5	pan	58.0	74.4	11.4	21.6
shn	16.8	47.9	15.1	16.9	spa	84.0	79.3	5.8	11.7	tgl	75.2	80.1	7.8	15.3
uzn	69.0	77.0	13.4	19.2	zul	36.5	59.3	5.1	22.8	arb	29.8	56.4	3.6	20.8
bod	29.0	60.2	6.8	20.3	sun	65.5	74.3	13.5	20.9	hat	55.8	72.9	4.6	20.9
ilo	54.0	69.8	9.2	22.6	kaz	63.5	75.6	19.3	23.0	lit	73.8	82.3	11.0	15.7
mya	0.8	7.2	72.5	55.7	pbt	47.5	67.7	3.1	22.6	sin	32.2	59.2	3.8	21.5
srp	83.2	77.5	13.0	13.7	tha	71.8	75.5	19.8	21.2	vie	83.5	78.4	8.7	12.4
ars	62.0	78.5	9.6	18.8	bul	80.8	77.7	13.6	14.6	est	71.8	78.2	4.2	16.4
hau	45.2	67.1	11.7	23.9	ind	81.8	75.6	6.4	13.4	kea	48.8	73.0	8.0	21.1
lug	35.5	57.6	2.5	22.5	nld	83.0	78.2	5.9	12.0	pes	79.2	77.8	8.5	14.4
sin	58.8	72.7	10.3	21.6	ssw	31.8	61.5	3.0	20.8	tir	28.0	57.7	4.9	19.7
war	62.2	74.7	12.1	21.0	ary	58.5	72.0	7.7	21.2	cat	86.2	84.8	9.4	10.2
eus	69.5	75.7	7.0	18.1	heb	77.2	76.9	17.2	17.8	isl	67.0	78.3	5.2	17.4
khk	48.0	70.2	10.0	22.9	luo	31.8	55.7	5.7	21.5	nob	79.2	77.8	9.5	14.7
plt	44.2	65.6	5.8	23.1	slk	82.0	76.9	4.7	13.0	eng	87.8	87.9	4.0	8.1
tsn	31.8	62.7	4.6	20.8	wol	33.0	53.2	5.7	22.3	Avg.	57.6	68.9	10.9	19.8

Table 14: Per-language performance on the belebele test set for the **LLaMA3** model, reporting AUROC, ECE, and Brier score. Each row is color-coded by language category, based on resource availability (high, medium, low) and script type (Latin vs. non-Latin). The categories are shaded with soft pastel colors: high-resource Latin (light blue), high-resource non-Latin (light pink), medium-resource Latin (light green), medium-resource non-Latin (lavender), low-resource Latin (cream), and low-resource non-Latin (tan). English line and the Average line is **bolded**.




Figure 9: Calibration metrics (ECE, Brier score, AUROC) vs. entropy across layers on the MMMLU dataset for LLaMA3, Aya, and Mistral.

applied to test-set predictions, and metrics (ECE, AUROC, Brier score, Accuracy) are computed using the calibrated probabilities.

Temperature Scaling Temperature scaling applies a single scalar parameter $T > 0$ to rescale logits before computing probabilities. The parameter is optimised by minimising the **negative log-likelihood (NLL)** on the validation set. In practice, we perform a coarse-to-fine grid search: first over a wide range ($T \in [0.05, 5.0]$ with 60 candidates), then locally refining around the best value with a denser grid. This procedure provides stable estimates across languages and avoids degenerate minima. The resulting optimal temperature is then used to rescale logits of all models prior to evaluation.

Isotonic Regression Isotonic regression learns a non-parametric, monotone mapping from predicted probabilities to calibrated probabilities in $[0, 1]$. We use the `scikit-learn` implementation with out-of-bounds clipping and monotonicity constraints. The model is fitted on the validation set and then applied to test-set predictions.

C.2 Cross-Lingual Layer-Sharing and Transfer

To better understand whether calibration improvements transfer across languages, we analyse the

extent to which languages share similar “good layer” sets under our intermediate-layer calibration framework. If two languages select highly overlapping sets of good layers, we would expect stronger cross-lingual transfer: optimising for one language could benefit the other. Conversely, disjoint sets would suggest that calibration signals are language-specific.

Measuring cross-lingual similarity. For each language ℓ , let G_ℓ denote its selected set of good layers. We compute a language–language similarity matrix using the Jaccard index:

$$\text{Jaccard}(\ell_i, \ell_j) = \frac{|G_{\ell_i} \cap G_{\ell_j}|}{|G_{\ell_i} \cup G_{\ell_j}|}.$$

A score of 1.0 indicates identical layer sets, while 0.0 means no overlap. Higher similarity reflects stronger potential cross-lingual transfer.

Findings. Figure 12 shows the cross-lingual similarity matrices for LLaMA-3 8B and 70B. While both models show some alignment among typologically related languages, the 70B model exhibits *far clearer and more coherent clustering*, reflecting stronger structural organisation in its intermediate representations. In particular:

- **Western European languages** (e.g., German,




Figure 10: Calibration metrics (ECE, Brier score, AUROC) vs. entropy across layers on the MMMLU dataset for Phi, Deepseek, and Qwen 2.5.

Spanish, Portuguese, French, Italian) share overlapping good-layer subsets.

- **CJK languages** (Japanese, Chinese, Korean) also form a consistent cluster.
- Languages across distant families (e.g., English vs. CJK; Romance vs. Indic or African languages) show much lower overlap.

Example good-layer sets within Romance languages in the 70B model illustrate this structured overlap:

- de: [72, 75, 76, 77, 78, 79, 80]
- es: [72, 73, 75, 76, 77, 78, 79, 80]
- pt: [71, 72, 73, 75, 76, 77, 78, 79, 80]
- fr: [77, 78, 79, 80]
- it: [72, 77, 78, 79, 80]

Overall, these results confirm that **similar languages tend to share partially overlapping calibration-relevant layer patterns**, while more distant languages exhibit far less similarity.

C.3 Cross-Dataset Consistency of Good-Layer Selection

We assess whether “good layers” for calibration are consistent across different tasks by performing a cross-dataset layer-overlap analysis. For each language ℓ common to both MMMLU and Belebele, we extract the good-layer sets G^{MMMLU} and G^{Belebele} and compute their Jaccard similarity:

$$\text{Jaccard}(m) = \frac{1}{|\mathcal{L}_m|} \sum_{\ell \in \mathcal{L}_m} \frac{|G^{\text{MMMLU}}_{m,\ell} \cap G^{\text{Belebele}}_{m,\ell}|}{|G^{\text{MMMLU}}_{m,\ell} \cup G^{\text{Belebele}}_{m,\ell}|}.$$

This evaluation was conducted across the 12 languages common to both datasets. The results show moderate cross-dataset consistency, with an average Jaccard similarity of **0.5227** for LLaMA-3 and **0.4704** for Aya. This indicates that despite the different task formulations, the benchmarks identify partially overlapping sets of calibration-relevant layers. This suggests that certain structural properties of intermediate representations support robust calibration across tasks, though task-specific effects remain.

Transfer evaluation on MMMLU using Belebele-derived layers. To further test generalisation, we applied the good-layer sets identified on Belebele




Figure 11: **Per-language calibration reliability diagrams for LLaMA3.** Each panel shows a reliability histogram with evenly spaced confidence bins. **Blue** bars correspond to the chosen intermediate layer (Layer 29), and **orange** bars correspond to the original final layer. The dashed diagonal is the perfectly calibrated line ($y=x$). Hatched overlays indicate the absolute calibration gap within each bin. The inset reports ECE (%) for both layers and the change $\Delta\text{ECE} = \text{ECE}_{\text{Final}} - \text{ECE}_{29}$ (positive values denote improved calibration at Layer 29).




Figure 12: Cross-lingual Jaccard similarity between languages' good-layer sets. Clear correlation between Western European and CJK language groups indicate stronger organisation of calibration-relevant information in deeper layers.

directly to MMMLU without any post-hoc calibration. The results, shown in Table 15, confirm that Belebele-derived layers retain meaningful calibration performance on MMMLU, demonstrating clear cross-task transferability.

C.4 Robustness to Language-ID Noise

To quantify how sensitive LACE is to noisy language labels, we run a controlled corruption exper-

Model	Setting	ECE ↓	AUROC ↑	Brier ↓
LLaMA-3	baseline	22.2	66.1	23.3
	<i>transferred-LACE</i>	15.6	70.6	21.8
Aya	baseline	27.4	68.9	32.3
	<i>transferred-LACE</i>	19.8	69.3	25.2

Table 15: Calibration transfer results when applying **Belebele-derived good layers** directly to MMMLU (without post-hoc calibration).

iment on MMMLU for two base models (Aya and LLaMA-3).

We start from our setup in the main paper where language-specific layer sets are selected on the validation split using gold language labels. Then at evaluation time, we simulate language-ID mistakes by randomly corrupting the language tag for a proportion $p \in \{0.0, 0.1, 0.2\}$ of test examples: for each selected example, we replace its true language label with a *different* language sampled from the remaining languages. LACE uses these (possibly incorrect) language tags to choose which language-specific layer set to average, while the underlying

model logits remain unchanged.

Therefore, $p = 0$ corresponds to perfect language identification, whereas $p = 0.1$ and $p = 0.2$ approximate moderate error rates in an external language-ID system; state-of-the-art LID systems achieve above 90% accuracy on sentence-level text (Suarez et al., 2026), so $p = 0.2$ is a conservative upper bound on realistic misidentification rates. We compute ECE, AUROC, and Brier score for each corruption level. The results are shown in Table 16.

Corruption	Model	ECE (%)	AUROC (%)	Brier (%)
0%	LLaMA-3	5.96	72.94	19.91
	Aya	11.42	68.61	23.16
10%	LLaMA-3	5.97	72.80	19.97
	Aya	11.60	68.38	23.31
20%	LLaMA-3	6.29	72.39	20.15
	Aya	12.07	68.25	23.39

Table 16: Robustness of LACE to language-ID corruption on MMMLU.

The impact of language-ID noise is mild. For Aya, corrupting 10% of tags increases ECE by only 0.18 (about 1.5% relative) and Brier by 0.14, while AUROC decreases by 0.23. Even at 20% corruption, ECE rises by only 0.65 compared to the clean setting, and Brier changes by 0.23. For LLaMA-3, the effects are even smaller: ECE changes by +0.01 at 10% and +0.33 at 20%, with AUROC decreasing by less than 0.60 and Brier increasing by at most 0.24. These variations are negligible relative to the large cross-lingual calibration gaps reported in the main paper.

Intuitively, LACE remains robust because it averages over multiple “good” layers per language rather than relying on a single, highly language-specific layer: misrouting an example to another language’s layer set still yields a reasonable ensemble, so moderate language-ID noise only weakly perturbs the final predictions.

C.5 Sensitivity to Validation Set Size

A potential limitation of LACE is that it requires a labelled validation set for each language in order to identify the corresponding Good Layers. This requirement may appear burdensome for low-resource languages, which are often the settings where improved calibration is most needed. However, in our main experiments, we use only 1,000 labelled validation examples per language, which we believe is already a practically attainable amount of data for many languages.

To further examine the data requirement of LACE, we conduct a sensitivity analysis by reducing the validation set size from the full 1,000 examples to 500, 200, and 100 examples. For each reduced setting, we recompute the selected Good Layers and compare them against those obtained from the full 1,000-example validation set. As a measure of agreement, we report the Jaccard similarity between the two selected layer sets.

Table 17 shows that the selected Good Layers remain highly consistent when using only 500 validation examples, with an average Jaccard similarity of 0.88 relative to the 1,000-example baseline. When the validation size is further reduced to 200 examples, the agreement remains moderate (0.65), while at 100 examples the consistency drops to 0.51. These results suggest that LACE is robust to substantial reductions in validation data, and that in practice as few as 500 labelled examples are sufficient to recover nearly the same layer selection as the full validation set.

Comparison	Avg. Jaccard	Interpretation
1000 vs. 500	0.88	High consistency
1000 vs. 200	0.65	Moderate consistency
1000 vs. 100	0.51	Lower consistency

Table 17: Stability of Good Layer selection under reduced validation set sizes, measured by Jaccard similarity against the full 1,000-example validation set.

Overall, this analysis indicates that the validation data requirement of LACE is lower than it may initially seem. While performance naturally degrades as the validation set becomes very small, the method remains highly stable at 500 examples, substantially lowering the barrier to applying LACE in lower-resource settings.

D Experiments on the 70B Model

To answer whether the calibration patterns observed in 7B–8B models persist in much larger models, we conducted a layer-wise calibration study on the **LLaMA-3.1 70B Instruct** model, using the same multilingual MMMLU setup.

Experimental Setup. We evaluated the 70B model under identical conditions as the 8B model: 15 languages and 15,000 test examples. To manage computational demands, the model was loaded using 4-bit quantization, which reduces memory usage while preserving performance. All experiments were performed on NVIDIA H100 80GB

GPUs. LLaMA-3.1 70B has **80 transformer layers** (compared to 32 in the smaller models).

Layer-wise Calibration in a Large Model. The results, shown in Figure 13, demonstrate that the phenomenon observed in smaller models is also present in the 70B model: calibration improves substantially in the late-intermediate layers, well before the final layer.




Figure 13: Expected Calibration Error (ECE) across the 80 layers of LLaMA-3.1 70B. A clear calibration minimum is visible in the intermediate layers across all language groups.

Quantifying the Improvement. To evaluate whether our calibration methods continue to be effective at larger scales, we also applied BEST LAYER, GOOD LAYER, and LACE to the 70B model. The results are reported in Table 18. Following the same format as Table 3 in the main text, this table shows that the intermediate-layer signal remains strong in the 70B model: using probabilities from the best-performing intermediate layer yields substantially better calibration than relying on the final layer. Moreover, LACE continues to provide the largest overall improvement, achieving the best calibration among all methods. These results confirm that our proposed layer-based calibration strategies, including LACE, scale effectively to much larger models.

E Generalisation Beyond MCQA: Full Details

This appendix supplements the generative-task analysis presented in §6.

E.1 Experimental Setup

Datasets. We use two human-translated multilingual generative benchmarks. **MKQA** (Longpre et al., 2021) is an open-ended multilingual question-answering benchmark. We sample 500 examples per language for seven languages (en, de, es, fr, ru, ja, th), spanning Latin, Cyrillic, Japanese, and

Method	ECE	Brier	AUROC
ORIGINAL PROBABILITY [†]	23.51	22.10	65.19
(+Temperature Scaling)	22.56	21.97	65.19
(+Isotonic Regression)	19.01	21.34	65.07
BEST LAYER [†]	16.31	16.14	84.57
(+Temperature Scaling)	13.81	14.94	84.57
(+Isotonic Regression)	7.62	13.77	84.40
GOOD LAYER [†]	13.90	16.43	79.30
(+Temperature Scaling)	15.00	16.42	79.30
(+Isotonic Regression)	13.44	16.25	78.98
LACE [‡]	6.14	15.76	79.42
(+Temperature Scaling)	4.93	15.60	79.92
(+Isotonic Regression)	2.81	15.38	78.79

Table 18: Calibration performance of different layer-based methods on the LLaMA-3.1 70B model using MMMLU. Intermediate-layer methods continue to yield substantial calibration improvements even at large model scale.

Thai scripts. **MGSM** (Shi et al., 2023) is a multilingual chain-of-thought math benchmark derived from GSM8K. We use 250 examples per language for eight languages (en, de, es, fr, ru, ja, th, bn).

Correctness measure. For MKQA, we follow the standard open-domain QA convention: a generation is scored correct if any gold answer (or alias) is contained in the normalised prediction, or vice versa. Normalisation uses Unicode-aware case folding, punctuation stripping, and whitespace normalisation (consistent with the reference implementation released by the dataset authors). For MGSM, we extract the final numeric token from the generated chain-of-thought and require an exact match to the gold answer up to a tolerance of 10^{-4} .

Confidence methods. We evaluate four standard generative-confidence elicitation methods, following the taxonomy of Lewis-Lim et al. (2025):

- **Perplexity:** $\exp\left(\frac{1}{T} \sum_{t=1}^T \log p(y_t | y_{<t}, x)\right)$, the geometric mean of next-token probabilities over the generated answer of length T .
- **Margin:** the per-token softmax margin between the top-1 and top-2 candidate tokens, averaged across generated tokens. The output lies in $[0, 1]$.
- **P(True)** (Kadavath et al., 2022): we re-prompt the model with “Question: . . . Proposed answer: . . . Is the proposed answer correct?” and read the softmax probability assigned to the “True” token at the next position.
- **Verbalised** (Tian et al., 2023): we append “How confident are you on a scale of 1–9?” and decode an expected confidence $\hat{c} = \mathbb{E}[s | s \in \{1, \dots, 9\}]$,

mapped to $[0, 1]$.

E.2 Per-Language Results

Tables 19 and 20 report per-language accuracy and ECE for each of the four confidence methods. The full numerical layer-wise results, including AUROC and Brier scores at every layer, are released alongside the code (layerwise_mgsm_llama3.csv).

Lang.	Acc.	Perp.	Margin	P(True)	Verb.
English	52.0	13.5	10.6	19.2	26.9
German	43.0	23.4	23.7	26.3	31.2
Spanish	42.8	25.2	24.1	25.7	31.0
French	43.2	25.1	24.7	24.2	32.3
Russian	26.4	39.8	41.2	37.4	49.0
Japanese	18.4	45.3	47.7	42.4	55.7
Thai	24.0	44.5	47.6	34.1	47.2
Avg. Non-Eng.	33.0	33.9	34.8	31.7	41.1

Table 19: Per-language MKQA results (LLaMA3, 500 examples per language). ECE (%), $\downarrow$ for four confidence methods. Bold marks the column-minimum.

Lang.	Acc.	Perp.	Margin	P(True)	Verb.
English	65.6	18.9	15.9	11.0	12.8
German	51.6	32.5	30.3	10.9	21.7
Spanish	56.4	27.8	26.0	13.0	16.3
French	44.8	40.1	38.7	9.5	24.4
Russian	47.2	37.9	36.4	13.5	19.2
Japanese	44.0	40.0	37.9	17.1	25.6
Thai	36.4	48.4	47.6	12.2	28.8
Bengali	14.0	73.3	70.7	12.2	30.4
Avg. Non-Eng.	42.1	42.9	41.1	12.6	23.8

Table 20: Per-language MGSM results (LLaMA3, 250 examples per language). ECE (%), $\downarrow$ for four confidence methods. Bold marks the column-minimum.

E.3 Discussion of the P(True) Outlier on MGSM

On MGSM, P(True) reports nearly identical ECE for every language (~ 10 – 17%), in contrast to the wide spread observed for perplexity, margin, and verbalised. This is consistent with the documented brittleness of self-reprompting probes: Kadavath et al. (2022) and Lin et al. (2022) both note that the “Is the proposed answer correct?” template is read by the model in the language of the template (here, English), independently of the language of the original question. The resulting probabilities are dominated by the English-language meta-question rather than by the multilingual content of

the answer, so the protocol effectively short-circuits the language-conditional confidence signal. We retain the method to faithfully report the benchmark and note that the other three methods, which read confidence directly from the generated tokens, all preserve the multilingual gap.

E.4 Learned Layer Selector

The selector’s features are the example’s per-layer confidences (33 values for LLaMA3); a logistic regression trained on the validation split maps this signature to a choice of layer. Experiments use the MGSM and MKQA subsets of §6 (eight languages, 100 examples per language); the language-agnostic global-layer baseline is tuned on the same validation split in each run.

F Robustness of the Multilingual Calibration Gap

The English/Non-English calibration gap reported in §3 is the central empirical claim of the paper. This appendix collects the robustness checks that establish the gap is preserved under statistical re-sampling of the evaluation instances, different ECE bin counts, and alternative MCQA scoring protocols, and examines one case in which ECE and Brier disagree.

F.1 Bootstrap Confidence Intervals

Model	ECE $\downarrow$		AUROC $\uparrow$	
	EN	Non-EN	EN	Non-EN
LLaMA3	[3.93, 8.81]	[20.72, 24.80]	[77.67, 83.01]	[62.85, 64.71]
Aya	[17.23, 24.58]	[26.74, 29.06]	[71.57, 77.65]	[67.22, 69.11]
Mistral	[21.72, 30.28]	[33.12, 35.19]	[70.78, 76.73]	[62.09, 65.07]
Qwen2.5	[12.71, 19.09]	[21.13, 22.93]	[75.10, 81.36]	[69.17, 71.25]
Phi	[16.25, 28.71]	[26.97, 29.27]	[66.94, 75.18]	[56.38, 58.46]
DeepSeek	[5.02, 16.16]	[24.62, 27.91]	[62.79, 69.60]	[55.70, 58.58]
Qwen3	[8.19, 17.14]	[13.17, 15.56]	[77.38, 86.99]	[71.91, 74.94]
Olmo3	[24.07, 34.66]	[26.19, 28.87]	[58.05, 70.96]	[54.47, 58.21]

Table 21: 95% bootstrap confidence intervals for the per-model MMMLU entries of Table 2.

Table 21 reports 95% confidence intervals for the per-model entries of Table 2, computed with a non-parametric bootstrap: instances are resampled with replacement within each language (2,000 resamples; ECE with $K=10$ bins), per-language metrics are recomputed, and the non-English value macro-averages the 14 non-English languages of each resample. The English intervals are wider because they rest on a single language split. The English and non-English intervals are disjoint for ECE in five of the eight models (LLaMA3, Aya,

Mistral, Qwen2.5, DeepSeek) and for AUROC in every model except Olmo3.

F.2 ECE Binning Sensitivity

In the main paper, following Guo et al. (2017), we report ECE using $K = 10$ uniformly spaced bins over $[0, 1]$, applying the identical binning scheme across all models, layers, languages, and calibration methods.

To evaluate robustness with respect to this hyperparameter, we recomputed ECE for MMMLU predictions for both LLaMA-3 and Aya using $K \in \{5, 10, 15, 20, 25\}$. Table 22 presents the complete results for both LLaMA-3 and Aya.

Bins (K)	LLaMA-3 ECE (%)		Aya ECE (%)	
	English	Non-Eng.	English	Non-Eng.
5	3.21	17.87	19.90	25.93
10	4.61	23.12	20.66	26.77
15	5.04	23.27	20.34	27.45
20	5.58	22.83	21.72	28.26
25	6.07	23.13	21.34	28.27

Table 22: ECE binning sensitivity for LLaMA-3 and Aya on MMMLU. English consistently shows lower calibration error than non-English across all binning configurations.

Although absolute values shift somewhat under different bin counts, the *qualitative pattern remains unchanged*: non-English languages remain substantially less calibrated than English for every choice of K , and non-English calibration is highly stable once $K \geq 10$ (LLaMA-3 within 22.83–23.27%, Aya within 26.77–28.27%). Brier score and AUROC (both bin-free) exhibit the same cross-lingual trends, so our conclusions are robust to reasonable variations in the ECE computation scheme.

F.3 ECE and Brier Disagreement under Capability Collapse

In Table 1, LLaMA3’s Chinese entry has the worst ECE among all languages (41.94) yet nearly the best Brier score among non-English languages (19.56, second only to Yoruba’s 19.43). The cause is that LLaMA3’s Chinese accuracy collapses to 23.1%, below the 25% random baseline for four-option MCQA, with confidences clustered at a similar level (average 23.2%; Table 13). In this regime the two metrics diverge by construction: Brier rewards uniformly low confidence on mostly-wrong predictions, while ECE captures the bin-level mismatch between confidence and accuracy.

The anomaly does not appear in models with functional Chinese performance (Tables 10 and 12). Calibration metrics should therefore not be interpreted for languages where accuracy is at or below chance.

F.4 Alternative MCQA Scoring Protocols

The main paper elicits MCQA confidence as the softmax probability assigned to the option-letter token at the choice position (Appendix A). This is a standard convention in MCQA evaluation, but not the only one. To verify that the multilingual calibration gap holds under alternative scoring protocols, we evaluate LLaMA3 under three protocols:

- **P1 – Letter probability** (paper baseline): softmax over the four option-letter tokens (A, B, C, D) at the choice position; confidence is the max-class probability. This is the protocol used throughout the main paper, reported here at full MMMLU scale (1000 examples per language; numbers reproduced from Table 1).
- **P2 – Option likelihood**: each full option text is scored under the model (sum of token log-probabilities), then softmax is taken across the four options; the model’s prediction is $\arg \max$, the confidence is the max softmax value. Evaluated on a balanced 14-language $\times$ 200-example MMMLU subset.
- **P3 – Text generation**: greedy-decode up to 8 new tokens, regex-extract the predicted letter, and use the softmax probability of that decoded letter at the first generated position as confidence. Evaluated on the same 14-language $\times$ 200-example subset.

Protocol	Group	Acc.	ECE↓	AUROC↑	Brier↓
P1: Letter prob. [†]	EN	61.20	4.61	80.36	17.63
	Non-EN avg.	41.47	23.12	64.06	22.95
P2: Option lik.	EN	39.50	20.10	59.70	29.10
	Non-EN avg.	31.40	27.00	54.10	30.90
P3: Text gen.	EN	40.50	6.70	76.10	18.30
	Non-EN avg.	27.00	12.20	49.60	19.00
Avg. Δ (EN – Non-EN)		+13.78	-10.30	+16.13	-2.61

Table 23: MMMLU calibration under three MCQA scoring protocols on LLaMA3. P1 is reported at full MMMLU scale (1000 examples per language; same numbers as Table 1). P2 and P3 are evaluated on a controlled 14-language $\times$ 200-example MMMLU subset. The bottom row averages the EN – Non-EN gap across protocols. [†]P1 values come from the main paper’s full-scale evaluation.

Finding. As Table 23 shows, the EN — Non-EN ordering documented in the main paper is preserved on *every* metric under all three scoring protocols. The averaged gap across protocols is +13.78 points of accuracy and +16.13 points of AUROC in favour of English, with ECE and Brier both reduced for English. Reading down the table, the P1 row reproduces the main-paper gap (full scale); P2 and P3 verify on a controlled subset that the gap persists even when confidence is derived from full-option likelihoods (P2) or from text-generation outputs (P3) rather than letter-token probabilities. The multilingual calibration gap therefore holds across these three scoring protocols.