

Calibrated Deferral Routing for Cost-Efficient Guardrails

Giandomenico Cornacchia*, Inkit Padhi, Manish Nagireddy
Subhajit Chaudhury, Pierre Dognin, Tejaswini Pedapati, Ambrish Rawat
Mark Purcell, Kush R. Varshney, Prasanna Sattigeri
IBM Research

Abstract

Large Language Models (LLMs) have achieved strong performance across diverse NLP tasks, but deploying them in real-world scenarios raises concerns regarding safety, reliability, and computational cost. To mitigate the risk of harmful, unreliable, or unfaithful content, LLM pipelines increasingly rely on guardrails. Yet deploying multiple guardrails across different tasks and risk categories introduces significant computational and orchestration overhead. In this work, we propose a *calibrated deferral-routing* framework for deploying guardrails both accurately and cost-effectively. Our approach arranges guardrails as a cascade of increasing capability and cost: each input is first evaluated by a lightweight guardrail, while only uncertain instances are deferred to a larger, more capable, and more expensive model. We formalize this routing problem and study practical deferral-calibration strategies by applying post-hoc calibration, including isotonic regression, Platt scaling, and temperature scaling. We evaluate 18 guardrail models from four model families (i.e., Granite Guardian, Llama Guard, ShieldGemma, and GuardReasoner) across 10 harm-detection benchmarks. Our experiments show that: (i) all guardrail models are often poorly calibrated out of the box (ECE ranging from 6% to 28%); (ii) post-hoc calibration substantially improves reliability reducing ECE below 2.5% in most cases; (iii) cascade routing yields large cost savings with minimal accuracy loss, while monotone calibration primarily makes routing thresholds deployable rather than improving rank ordering.

1 Introduction

LLMs have become integral to real-world NLP applications, powering assistants, summarization tools, and content generators. Despite their capabilities, LLMs remain prone to generating hallucinated, biased, or malicious content, which poses

safety and trust challenges (Bender et al., 2021; Weidinger et al., 2021). To address this concern, the notion of *guardrails*, i.e., external mechanisms for input-output verification, moderation, and alignment, has gained attention.

Guardrails can be instantiated as rule-based filters (Cornacchia et al., 2024), fine-tuned classifiers (OpenAI, 2023), or transformer-based models (Inan et al., 2023; Padhi et al., 2025, 2024), each offering different trade-offs between interpretability, coverage, and inference cost (Markov et al., 2023; Ghosh et al., 2024). While effective in isolation, combining multiple guardrails to distinct risk dimensions (e.g., factuality, toxicity, compliance) rapidly inflates latency and resource demand.

To mitigate this cost, practitioners often fine-tune smaller guardrail variants, each specialised for a particular safety dimension. However, smaller models are typically associated with lower Natural Language Understanding capabilities than their larger counterparts under the same training regime. A principled remedy is to delegate the decisions of the small model to a more capable one when the uncertainty of the former is high.

The literature on deferral architectures has often considered AI systems that defer to a human expert, or vice versa (Beede et al., 2020). In the former setting, the focus lies on identifying inputs on which the AI system cannot provide a confident prediction; in the latter, the AI serves as a second opinion that informs, and may bias, the expert’s evaluation (Jacobs et al., 2021; Mozannar et al., 2022; Liu et al., 2021). In both cases, learning to defer is typically cast as the joint optimization of a surrogate model and the downstream predictor, so as to learn an optimal configuration of safe, low-confidence deferral (Mozannar et al., 2023; Madras et al., 2018; Mozannar and Sontag, 2020).

In the NLP setting, the scale of production data makes it impractical to keep a human expert in the loop for every prediction. A common alternative is

*Work done at IBM Research; now with AWS AI.

to calibrate model confidence estimates to obtain more reliable uncertainty scores. Calibration alone, however, does not eliminate misclassification; what it provides is a trustworthy signal that can drive downstream decisions such as deferral. Jointly calibrating a small and a large guardrail, and routing only uncertain inputs to the latter, can therefore reduce the overall cost of the system while preserving its accuracy.

In this work, we propose a deferral-based orchestration strategy for LLM guardrails (Figure 1). By arranging guardrails in a calibrated cascade, we route only uncertain cases to more expensive evaluators, analogously to decision-making under budget constraints (Teerapittayanon et al., 2016; Xin et al., 2020). In offline benchmark study, the resulting cascade maintains high overall accuracy while reducing average inference cost; we do not claim demonstrated production readiness. We introduce a formal routing framework, clarify that monotone calibration makes deferral thresholds deployable but does not create new ranking signal, and evaluate our approach on multiple safety-critical datasets.

2 Related Work

Guardrails and LLM safety. Existing approaches to LLM safety can be broadly categorised into training-time alignment and inference-time detection. Training-time methods, such as reinforcement learning from human feedback (Ouyang et al., 2022) and Constitutional AI (Bai et al., 2022), aim to reduce harmful outputs at the source, while red-teaming (Ganguli et al., 2022; Perez et al., 2022) probes for residual vulnerabilities. These methods, however, offer no guarantees at deployment time, which motivates inference-time guardrails for real time input-output filtering. Llama Guard (Inan et al., 2023) generates binary safe/unsafe labels under a conversational safety taxonomy; ShieldGemma (Zeng et al., 2024) evaluates content against explicit policy guidelines; Granite Guardian (Padhi et al., 2024) unifies harm and RAG-hallucination detection via a structured safety template. Rule-based frameworks such as NeMo Guardrails (Rebedea et al., 2023) lack probabilistic outputs and are unsuitable for confidence-based routing.

Deploying multiple guardrails across risk dimensions incurs substantial overhead (Cornacchia et al., 2024) and benchmarks show no single defence dominates across attack styles (Zizzo et al., 2024;

Rawat et al., 2024; Schoepf et al., 2025). Yet existing guardrails are typically deployed in isolation, with no principled mechanism for routing across models of different sizes. Our work fills this gap with a calibrated deferral framework that orchestrates guardrails in a cost-aware cascade.

Deferral and cascaded inference. The idea of abstaining from prediction when confidence is low has a long history in selective classification (Chow, 1970; El-Yaniv and Wiener, 2010; Geifman and El-Yaniv, 2017). In the deep learning era, budgeted prediction methods such as BranchyNet (Teerapittayanon et al., 2016) and DeeBERT (Xin et al., 2020) use early exits to reduce average computation, but they require architectural modifications and are therefore not applicable to pre-trained guardrail models.

The learning-to-defer framework formalises the problem of jointly optimising a classifier and a deferral rule. Madras et al. (2018) introduce a joint objective with fairness constraints, while Mozannar et al. (2023) prove that the problem is NP-hard in the linear setting and propose realisable-consistent surrogate losses. Jitkrittum et al. (2023) extend this analysis to cascades of pre-trained classifiers and establish conditions under which confidence-based deferral is Bayes-optimal and conditions under which it is not (e.g., when the heavy model is a specialist with non-uniform error). A central observation of that work, which we leverage throughout our framework, is that confidence-based deferral suffices when the heavy model has approximately constant error probability.

Most relevant to our setting, Nagireddy et al. (2025) propose a cascade of small detectors for social bias detection and show that a two-model pipeline of small models can approach GPT-4 performance at a fraction of the cost. Their approach, however, relies on deterministic binary classifiers without calibrated confidence scores, which limits the ability to tune the cost-accuracy trade-off. We extend the cascade paradigm to general-purpose guardrails by introducing post-hoc calibration as the mechanism for producing reliable deferral signals, so that practitioners control the coverage-accuracy trade-off through a single threshold.

Uncertainty calibration. Reliable deferral requires probabilities that reflect the likelihood of correctness. Post-hoc calibration methods such as temperature scaling (Guo et al., 2017), Platt scaling (Platt, 1999), and isotonic regression (Zadrozny and Elkan, 2002; Niculescu-Mizil and Caruana,




Figure 1: An input is first evaluated by the lightweight guardrail G_l , which produces a hard prediction $\hat{y}_l(x)$ and a raw confidence score $u_l(x) = \max\{\hat{p}_l(x), 1 - \hat{p}_l(x)\}$. A post-hoc calibrator f_l maps $u_l(x)$ to a calibrated routing score $s_l(x)$. If $s_l(x) \geq \tau$, the system accepts the lightweight prediction $\hat{y}_l(x)$ at cost c_l ; otherwise, it defers to G_h and uses $\hat{y}_h(x)$, incurring cost $c_l + c_h$. Higher coverage ϕ reduces the fraction of double-inference calls.

2005) are well studied for standard classifiers, but remain less explored for LLM-based guardrails. Recent work shows that LLM guard models are often overconfident and become more miscalibrated under adversarial inputs (Liu et al., 2025), while Zhang et al. (2025) propose self-evaluation for safety calibration. However, these works do not study how calibration affects deferral thresholds. We focus on this connection, showing that calibration improves threshold reliability while fixed-coverage deferral depends primarily on ranking quality (Proposition A.2).

3 Background

3.1 Problem Setting

Let $\mathcal{X}$ denote the input space and $Y \in \{0, 1\}$ a binary label, where $Y=1$ indicates a safety violation. We consider a two-stage cascade with a lightweight guardrail G_l and a heavier guardrail G_h at raw inference costs $c_l < c_h$. Each model produces a raw score $\hat{p}_i(x) \in [0, 1]$, interpreted as an estimate of $P(Y=1 | X=x)$, and a hard prediction $\hat{y}_i(x) = \mathbb{I}[\hat{p}_i(x) \geq 0.5]$ that is held fixed throughout. A deferral rule $r : \mathcal{X} \rightarrow \{0, 1\}$ accepts G_l ($r=0$) or defers to G_h ($r=1$); under 0–1 loss with deferral penalty $c > 0$, the Bayes-optimal rule (Jitkritum et al., 2023) is $r^*(x) = \mathbb{I}[\Delta(x) > c]$ with $\Delta(x) = P(Y \neq \hat{y}_l(x) | x) - P(Y \neq \hat{y}_h(x) | x)$, generalizing Chow’s rejection rule (Chow, 1970) and the learning-to-defer setup of Mozannar et al. (2023). The full risk decomposition and excess-risk bound are in Appendix A. Since $\Delta(x)$ is unobserved, we use confidence-based deferral as a plug-in approximation; the risk-coverage curve is

determined by how well the routing score ranks examples by the benefit of deferral.

3.2 Confidence-Based Deferral via a Routing Score

In practice, the conditional errors in $\Delta(x)$ are unknown. We therefore use a scalar routing score derived from the lightweight model’s confidence in its own hard prediction. Let $\hat{p}_l(x)$ denote the raw estimate of $\Pr(Y = 1 | X = x)$, and let

$$\hat{y}_l(x) = \mathbb{I}[\hat{p}_l(x) \geq 1/2]. \quad (1)$$

The raw confidence assigned by G_l to this hard prediction is

$$u_l(x) = \max\{\hat{p}_l(x), 1 - \hat{p}_l(x)\}. \quad (2)$$

A post-hoc calibrator maps this raw confidence to a calibrated routing score:

$$s_l(x) = f_l(u_l(x)). \quad (3)$$

In the no-calibration case, f_l is the identity map. Calibration is used only to produce a reliable routing score; the hard prediction $\hat{y}_l(x)$ is unchanged.

The confidence-based deferral rule is

$$\hat{r}_\tau(x) = \mathbb{I}[s_l(x) < \tau], \quad \phi(\tau) = \Pr(s_l(X) \geq \tau), \quad (4)$$

where $\phi(\tau)$ is the coverage, i.e., fraction of examples handled by the lightweight model. If $s_l(x) \geq \tau$, cascade accepts $\hat{y}_l(x)$; otherwise, deferred to the heavier guardrail for hard prediction.

This routing rule is a plug-in approximation to the Bayes-optimal deferral rule: it is effective

when inputs assigned low calibrated routing confidence by G_l are also those on which G_h provides a meaningful conditional improvement. Because any monotone calibrator f preserves the ordering induced by u_l (Proposition A.4, Appendix A), calibration improves the reliability and interpretability of thresholds but does not by itself create new ranking information; ECE improvements therefore need not translate into improvements of the exact fixed-coverage risk-coverage curve. Since $\hat{y}_l$ is held fixed, calibration does not change the standalone accuracy, F1, precision, or recall of G_l ; cascade-level changes arise only through routing.

At coverage ϕ , the expected cost is $\bar{C}(\phi) = c_l + (1 - \phi) \cdot c_h$, with positive savings over G_h alone when $\phi > c_l/c_h$ (full derivation in Appendix H). While we focus on 0–1 loss for clarity, the framework extends to general loss functions with appropriate redefinition of $\Delta(x)$.

4 Calibration Methodology

The deferral policy in Eq. (4) requires calibrated routing score for the lightweight guardrail. Let $\hat{y}_i(x)$ denote the hard prediction of guardrail G_i . We write $u_i(x) \in [1/2, 1]$ for the uncalibrated confidence assigned by G_i to its own hard prediction. In the binary case,

$$u_i(x) = \max\{\hat{p}_i(x), 1 - \hat{p}_i(x)\}, \quad (5)$$

i.e., $u_i(x) = \hat{p}_i(x)$ when $\hat{y}_i(x) = 1$ and $u_i(x) = 1 - \hat{p}_i(x)$ otherwise.

A post-hoc calibrator $f_i : [0, 1] \rightarrow [0, 1]$ maps this raw confidence to a calibrated routing score:

$$s_i(x) = f_i(u_i(x)). \quad (6)$$

The cascade accepts the lightweight prediction when $s_l(x) \geq \tau$ and defers otherwise:

$$\hat{r}_\tau(x) = \mathbb{I}[s_l(x) < \tau], \quad \phi(\tau) = \mathbb{P}(s_l(X) \geq \tau). \quad (7)$$

Thus calibration is used only to produce a reliable routing score; the hard prediction $\hat{y}_l$ itself is unchanged.

We compare four calibration cases for the routing score, depending on the supervision available. *Case 1 (No Calibration)* sets $s_l(x) = u_l(x)$ and thresholds the raw confidence. *Case 2 (Deferral Annotation)* fits f_l to predict agreement $a_j = \mathbb{I}[\hat{y}_l(x_j) = \hat{y}_h(x_j)]$, requiring no ground truth; since this target encodes exactly where G_l and G_h disagree, closeness to Case 3 should

be read as “agreement is a sufficient surrogate.” *Case 3 (Annotation)* fits f_l to empirical correctness $z_j = \mathbb{I}[\hat{y}_l(x_j) = y_j]$ (the standard post-hoc setting). *Case 4 (Ideal)* additionally calibrates G_h on ground truth as a fully supervised reference. Under a deferral-aligned-confidence assumption (formal statement and discussion in Appendix F), confidence thresholding is a plug-in approximation to the Bayes-optimal rule.

Evaluation Protocol. We split the evaluation data 80/20 into calibration and test sets. The calibration set is used to fit post-hoc maps f_l , and when applicable, f_h . All metrics are reported on the held-out test set. For each calibration case, we sweep threshold τ to trace the coverage-accuracy Pareto frontier and report the Area Under the Risk-Rejection Rate Curve (AURRRC) as a threshold-free summary metric.

5 Experimental Setup

Datasets. For harm detection, we aggregate 10 public benchmarks spanning prompt harmfulness (ToxicChat, OpenAI Moderation, AegisSafetyTest, SimpleSafetyTests, HarmBench) and response harmfulness (BeaverTails, SafeRLHF, and three XSTest variants: XSTest-RR, XSTest-RH, XSTest-RR(h)); The statistics of the full data set are reminded in Appendix D.

Models. We evaluate 18 guardrail models from four families spanning 1B to 27B parameters. The Granite Guardian (GG) family includes eight models across versions 3.0, 3.1, 3.2, 3.3, and 4.1, ranging from 2B to 8B parameters. Llama Guard (LG) comprises four models: the original 7B, version 2 at 8B, and version 3 at both 1B and 8B. ShieldGemma (SG) offers three sizes: 2B, 9B, and 27B. GuardReasoner (GR) provides 1B, 3B, and 8B variants. Each model outputs a binary safety label and a confidence score. For generative models (GG, LG, GR), we extract the probability of the positive label from token-level logprobs. For ShieldGemma, we use the softmax over “Yes”/“No” logits. Full details on probability extraction are in Appendix K.

Calibration methods. We fit three post-hoc calibrators on the 80% calibration split. All calibrators are trained on the scalar routing confidence $u_i(x)$: with correctness labels in the ground-truth setting and agreement labels in the deferral-annotation setting. Isotonic regression fits a non-parametric monotone map. Platt scaling fits $\sigma(a \cdot \text{logit}(u_i(x))) +$

b). Temperature scaling applies a one-parameter monotone rescaling before mapping to the routing confidence. In all cases the hard prediction $\hat{y}_i$ is held fixed. Calibration does not by itself change a model’s standalone accuracy, F1, precision, or recall. Changes in cascade accuracy/F1 arise from the routing decision, because low-confidence examples are evaluated by G_h instead of G_l .

Deferral protocol. For each model pair (G_l, G_h) , we sweep the confidence threshold $\tau \in [0, 1]$ and compute coverage $\phi(\tau) = \mathbb{P}(s_l(x) \geq \tau)$. At each threshold, inputs with $s_l(x) < \tau$ are deferred to G_h . We evaluate four calibration cases (Section 4).

Metrics. For detection quality: AUC, F1, Recall. For calibration quality: Expected Calibration Error (ECE, 15 bins) and Area Under the Risk-Rejection Rate Curve (AURRRC). For deferral: accuracy and F1 as a function of coverage and cost savings under the formula $\bar{C} = c_l + (1 - \phi) \cdot c_h$. We report the *aggregate deferral risk*, defined as the mean error rate of the cascade with coverage levels (Eq. 12).

6 Results

6.1 Benchmark and Calibration Results

Table 1 presents detection and calibration results for four model families. Granite Guardian models achieve the strongest AUC (0.90 for GG-3.1-8B and GG-4.1-8B), while GuardReasoner-1B delivers the highest F1 (0.84) with balanced precision and recall. ShieldGemma-27B underperforms in recall (0.27), yielding low F1 despite high precision. Additional model families (Qwen3Guard, Nemotron) are evaluated in Appendix Table 3.

All models are substantially miscalibrated out of the box, with ECE ranging from 6.2% (GG-3.1-3B) to 27.5% (SG-27B). We compare three post-hoc calibration methods (Table 1, right columns). **Isotonic regression** achieves the lowest ECE for most models (below 2.5% for 17 out of 21), confirming its effectiveness as a flexible non-parametric calibrator. **Platt scaling** yields competitive ECE with a smoother parametric map. Since hard predictions are held fixed in our protocol, standalone F1, precision, recall, and accuracy are unchanged by calibration; differences in cascade-level metrics arise only through the routing decision. **Temperature scaling** is effective for moderately overconfident models (e.g., Granite Guardian, $T \approx 2$ to 5) but fails for highly overconfident ones (e.g., GuardReasoner, $T \approx 8$ to 10; Nemotron, ECE = 13.8%

after scaling). The learned T values, reported in Appendix Table 3, serve as a diagnostic of model overconfidence.

For the deferral experiments that follow, we use isotonic regression as the default calibrator because it achieves the lowest ECE for most models. Since isotonic regression is monotone, it primarily improves the reliability and interpretability of routing thresholds rather than introducing new ranking information. At exact fixed coverage, differences between calibrated and uncalibrated curves should be attributed to ties, finite threshold discretization, or differences in the underlying routing score, rather than to monotone calibration alone (Proposition A.4). Appendix G discusses four calibration strategies (No Calibration / Case Zero, Deferral Annotation, Annotation, Ideal Annotation; Appendix F) for the routing decision; we adopt the Annotation (ground-truth) case throughout.

6.2 Deferral Across Model Families

We compare the deferral trade-off across the four core model families using isotonic calibration on ground truth (Appendix Figure 4).

Granite Guardian exhibits the strongest deferral dynamics: at 80% coverage, the 2B $\rightarrow$ 8B cascade retains 83.2% accuracy (only 0.7 pp below the 8B, 95% CI [0.13, 1.63] pp) while reducing cost by 55% (accounting for overhead of always running G_l). Llama Guard shows a different pattern: its small model (1B) is weaker than the large, so accuracy degrades more steeply with coverage. For GuardReasoner, the 1B and 8B are of comparable quality (GR-1B leads on AUC and F1, GR-8B on recall), so the modest cascade movement reflects complementary error patterns rather than a clear quality ordering. ShieldGemma shows no deferral benefit due to the negligible accuracy gap between 2B and 27B (0.001).

Cross-Family Deferral. A key question for deployment is whether deferral works across model families, i.e., routing from a cheap model of one architecture to a stronger model of another. Table 2 shows the *aggregate deferral risk* (Eq. 12), the cascade’s mean error rate across coverage levels, for within-family and cross-family pairs under the No-Calibration (Case Zero) and Ideal Annotation cases. The Case Zero vs Ideal Annotation gap on this aggregate metric reflects coverage-grid bracketing and tie-breaking, not new ranking information (Proposition A.4); per-coverage

Table 1: Detection and calibration results on the OOD harm benchmark (20% held-out test set). ECE and AURRRC in %. Three post-hoc calibration methods compared. Best per column in **bold**. Additional models in Appendix Table 3.

Model	Detection			ECE ↓				AURRRC ↓			
	AUC	F1	Rec	Unc.	Iso.	Sig.	Temp.	Unc.	Iso.	Sig.	Temp.
GG-3.0-2B	0.79	0.68	0.76	10.6	1.1	4.3	2.9	17.2	15.1	16.4	17.2
GG-3.0-8B	0.87	0.76	0.74	6.8	2.4	4.4	2.5	9.1	9.0	9.5	9.1
GG-3.1-2B	0.86	0.75	0.76	6.2	2.3	4.1	4.5	10.5	9.5	11.3	10.5
GG-3.1-8B	0.90	0.80	0.78	6.7	1.2	2.4	3.5	7.4	7.1	8.4	7.4
GG-3.2-3B	0.85	0.75	0.79	8.5	1.2	2.9	3.7	12.2	11.0	11.8	12.2
GG-3.2-5B	0.88	0.78	0.77	9.9	1.7	3.5	3.5	9.4	8.1	10.1	9.4
GG-3.3-8B	0.89	0.80	0.82	13.9	1.8	3.1	2.1	7.8	7.3	9.5	7.8
GG-4.1-8B	0.90	0.80	0.79	12.7	1.9	1.8	2.3	7.2	6.8	8.9	7.2
LG-7B	0.83	0.66	0.55	14.2	3.1	3.9	2.5	13.3	12.0	13.5	13.4
LG-2-8B	0.85	0.74	0.65	15.4	2.6	3.0	3.0	10.9	10.6	11.3	10.9
LG-3-1B	0.80	0.65	0.56	17.0	1.4	4.6	2.0	14.2	13.2	13.7	14.2
LG-3-8B	0.83	0.72	0.63	14.8	2.5	2.5	4.7	12.3	10.6	11.7	12.3
SG-2B	0.74	0.39	0.25	26.5	2.2	5.6	6.5	22.0	18.5	19.2	22.0
SG-9B	0.73	0.38	0.24	26.4	2.5	5.8	7.1	22.2	18.6	19.2	22.2
SG-27B	0.76	0.41	0.27	27.5	2.2	2.7	4.9	20.0	16.7	18.7	20.0
GR-1B	0.86	0.84	0.83	17.2	1.7	1.1	0.7	11.1	11.4	11.1	10.4
GR-3B	0.83	0.83	0.86	18.3	1.9	1.5	1.9	16.4	14.0	14.7	16.4
GR-8B	0.82	0.82	0.86	18.7	1.0	0.4	1.9	15.5	14.7	14.9	17.9

curves are in Appendix Figure 4. Cross-family deferral to Granite-Guardian-3.1-8B achieves the lowest aggregate deferral risk in our evaluation (2.82% for GR-1B→GG-8B, under the Ideal Annotation case). However, we do not assess statistical significance and do not claim statistical separation from the strongest within-family pair (3.45%, GG-3.1-2B→GG-3.1-8B), and results may vary across datasets. This demonstrates that, under both-model ground-truth (Ideal) calibration, deferral is not limited to same-architecture pairs and that complementary strengths across families can be exploited; whether this advantage survives the label-free regimes is left open.

Table 2: Average risk across coverage levels for within-family and cross-family deferral pairs, under the No-Calibration (Case Zero) and Ideal Annotation cases

Small Model	Large Model	Case Zero	Ideal Ann.
<i>Within-family</i>			
GG-3.1-2B	GG-3.1-8B	11.04%	3.45%
LG-3-1B	LG-3-8B	8.40%	4.05%
GR-1B	GR-8B	9.07%	4.19%
SG-2B	SG-27B	6.28%	5.94%
<i>Cross-family</i>			
GR-1B	GG-3.1-8B	15.60%	2.82%
LG-3-1B	GG-3.1-8B	9.29%	3.68%

6.3 Cost Reduction Analysis

Every input is first evaluated by G_l and only deferred inputs invoke G_h (per-model cost breakdown in Appendix E.4).

Under two deployment scenarios (fixed 80% coverage and the maximum coverage preserving $\geq 99\%$ of G_h accuracy), the GG-2B→GG-8B cascade attains $1.8\times$ cost (55% saving) and the LG-1B→LG-8B cascade reaches $1.3\times$ (68% saving) at 80% coverage; for GG, the 99%-accuracy operating point coincides with $\phi=0.8$ (Appendix H, Table 6). The cost-accuracy Pareto frontier (Appendix H, Figure 6) shows the cascade strictly dominating both standalone baselines: cheaper than G_h at matched accuracy, more accurate than G_l at matched cost.

7 Conclusion

We presented a calibration-aware deferral routing framework for LLM guardrails, where a lightweight detector handles most inputs and defers uncertain cases to a more capable but costlier model, with post-hoc calibration (isotonic, Platt, temperature) making routing thresholds deployable rather than improving rank ordering.

Across four guardrail families and ten harm-detection benchmarks, three findings emerge. First, all guardrails are substantially miscalibrated out of the box (ECE 6%–25%); post-hoc calibration reduces ECE below 2.5% in most cases. Sec-

ond, confidence-based deferral yields favorable accuracy-cost trade-offs when G_l 's routing confidence aligns with the benefit of deferral: the GG-3.1-2B→GG-3.1-8B cascade saves 55% cost at 80% coverage, losing <1 pp accuracy vs. the 8B alone. Third, effectiveness varies sharply by family and depends on the *shape* of the accuracy curve, not merely on the presence of an accuracy gap. The framework is model-agnostic, requires no retraining, and bridges learning-to-defer with cost-efficient guardrail orchestration.

Limitations

This paper focuses on harm-detection guardrails and studies calibrated deferral in a single two-model cascade. The setting is deliberately simple, reflecting a practical deployment scenario and several extensions remain open. Multi-stage cascades with more than two models and routing policies that exploit complementary strengths across model families, are natural generalizations of our framework. Cross-family deferral, where routing decisions span architectures, is a promising direction; a systematic study with statistical significance estimation is left to future work. The calibration data requirements of isotonic regression deserve further investigation, since non-parametric methods may overfit with very small calibration sets. Finally, integrating calibrated deferral with latency-aware scheduling and batching strategies would complete the path from calibration theory to deployable guardrail orchestration. Beyond harm detection, evaluating calibrated deferral across other risk definitions, such as groundedness, hallucination detection, factual consistency, privacy leakage and policy compliance, may reveal different scaling behavior and motivate different cascade architectures.

References

- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, and 1 others. 2022. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073*.
- Emma Beede, Elizabeth Baylor, Fred Hersch, Anna Lurchenko, Lauren Wilcox, Paisan Ruamviboonsuk, and Laura M. Vardoulakis. 2020. [A human-centered evaluation of a deep learning system deployed in clinics for the detection of diabetic retinopathy](#). In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI '20, page 1–12, New York, NY, USA. Association for Computing Machinery.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the dangers of stochastic parrots: Can language models be too big?](#) In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 610–623. Association for Computing Machinery.
- Chi-Keung Chow. 1970. On optimum recognition error and reject tradeoff. *IEEE Transactions on Information Theory*, 16(1):41–46.
- Giandomenico Cornacchia, Giulio Zizzo, Kieran Fraser, Muhammad Zaid Hameed, Ambrish Rawat, and Mark Purcell. 2024. Moje: Mixture of jailbreak experts, naive tabular classifiers as guard for prompt attacks. In *AIES (1)*, pages 304–315. AAAI Press.
- Josef Dai, Xuehai Pan, Ruiyang Sun, Jiaming Ji, Xinbo Xu, Mickel Liu, Yizhou Wang, and Yaodong Yang. 2024. [Safe RLHF: safe reinforcement learning from human feedback](#). In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net.
- Ran El-Yaniv and Yair Wiener. 2010. On the foundations of noise-free selective classification. *JMLR*, 11:1605–1641.
- Deep Ganguli, Liane Lovitt, Jackson Kernion, Amanda Askell, Yuntao Bai, Saurav Kadavath, Ben Mann, Ethan Perez, Nicholas Schiefer, Kamal Ndousse, and 1 others. 2022. Red teaming language models to reduce harms: Methods, scaling behaviors, and lessons learned. *arXiv preprint arXiv:2209.07858*.
- Yonatan Geifman and Ran El-Yaniv. 2017. Selective classification for deep neural networks. In *Advances in Neural Information Processing Systems*, volume 30, pages 4878–4887.
- Shaona Ghosh, Prasoon Varshney, Erick Galinkin, and Christopher Parisien. 2024. Aegis: Online adaptive ai content safety moderation with ensemble of llm experts. *arXiv preprint arXiv:2404.05993*.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. 2017. On calibration of modern neural networks. In *ICML*.
- Hakan Inan, Kartikeya Upasani, Jianfeng Chi, Rashi Rungta, Krithika Iyer, Yuning Mao, Michael Tontchev, Qing Hu, Brian Fuller, Davide Testuggine, and Madian Khabisa. 2023. [Llama guard: Llm-based input-output safeguard for human-ai conversations](#). *CoRR*, abs/2312.06674.
- Maia Jacobs, Melanie F Pradier, Thomas H McCoy Jr, Roy H Perlis, Finale Doshi-Velez, and Krzysztof Z Gajos. 2021. How machine-learning recommendations influence clinician treatment selections: the example of antidepressant selection. *Translational psychiatry*, 11(1):108.

- Jiaming Ji, Mickel Liu, Josef Dai, Xuehai Pan, Chi Zhang, Ce Bian, Boyuan Chen, Ruiyang Sun, Yizhou Wang, and Yaodong Yang. 2023. Beavertails: Towards improved safety alignment of LLM via a human-preference dataset. In *NeurIPS*.
- Wittawat Jitkrittum, Neha Gupta, Aditya Krishna Menon, Harikrishna Narasimhan, Ankit Singh Rawat, and Sanjiv Kumar. 2023. When does confidence-based cascade deferral suffice? In *NeurIPS*.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *SOSP*.
- Zi Lin, Zihan Wang, Yongqi Tong, Yangkun Wang, Yuxin Guo, Yujia Wang, and Jingbo Shang. 2023. Toxicchat: Unveiling hidden challenges of toxicity detection in real-world user-ai conversation. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 4694–4702.
- Han Liu, Vivian Lai, and Chenhao Tan. 2021. Understanding the effect of out-of-distribution examples and interactive explanations on human-ai decision making. *Proc. ACM Hum. Comput. Interact.*, 5(CSCW2):408:1–408:45.
- Hongfu Liu, Hengguan Huang, Xiangming Gu, Hao Wang, and Ye Wang. 2025. [On calibration of llm-based guard models for reliable content moderation](#). In *International Conference on Learning Representations*.
- David Madras, Toniann Pitassi, and Richard S. Zemel. 2018. Predict responsibly: Improving fairness and accuracy by learning to defer. In *NeurIPS*, pages 6150–6160.
- Todor Markov, Chong Zhang, Sandhini Agarwal, Florentine Eloundou Nékoul, Theodore Lee, Steven Adler, Angela Jiang, and Lilian Weng. 2023. [A holistic approach to undesired content detection in the real world](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 15009–15018.
- Mantas Mazeika, Long Phan, Xuwang Yin, Andy Zou, Zifan Wang, Norman Mu, Elham Sakhae, Nathaniel Li, Steven Basart, Bo Li, David Forsyth, and Dan Hendrycks. 2024. Harmbench: A standardized evaluation framework for automated red teaming and robust refusal. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 35181–35224. PMLR.
- Hussein Mozannar, Hunter Lang, Dennis Wei, Prasanna Sattigeri, Subhro Das, and David A. Sontag. 2023. Who should predict? exact algorithms for learning to defer to humans. In *AISTATS*, volume 206 of *Proceedings of Machine Learning Research*, pages 10520–10545. PMLR.
- Hussein Mozannar, Arvind Satyanarayan, and David A. Sontag. 2022. Teaching humans when to defer to a classifier via exemplars. In *AAAI*, pages 5323–5331. AAAI Press.
- Hussein Mozannar and David A. Sontag. 2020. Consistent estimators for learning to defer to an expert. In *ICML*, volume 119 of *Proceedings of Machine Learning Research*, pages 7076–7087. PMLR.
- Manish Nagireddy, Inkit Padhi, Soumya Ghosh, and Prasanna Sattigeri. 2025. [When in doubt, cascade: Towards building efficient and capable guardrails](#). *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 8(2):1812–1821.
- Alexandru Niculescu-Mizil and Rich Caruana. 2005. Predicting good probabilities with supervised learning. In *Proceedings of the 22nd International Conference on Machine Learning (ICML)*, pages 625–632.
- OpenAI. 2023. [Openai moderation api](#).
- Long Ouyang, Jeffrey Wu, Xu Jiang, and 1 others. 2022. Training language models to follow instructions with human feedback. In *NeurIPS*.
- Inkit Padhi, Manish Nagireddy, Giandomenico Cornacchia, Subhajit Chaudhury, Tejaswini Pedapati, Pierre L. Dognin, Keerthiram Murugesan, Erik Miehl, Martin Santillan Cooper, Kieran Fraser, Giulio Zizzo, Muhammad Zaid Hameed, Mark Purcell, Michael Desmond, Qian Pan, Zahra Ashktorab, Inge Vejsbjerg, Elizabeth M. Daly, Michael Hind, and 4 others. 2024. Granite guardian. *CoRR*, abs/2412.07724.
- Inkit Padhi, Manish Nagireddy, Giandomenico Cornacchia, Subhajit Chaudhury, Tejaswini Pedapati, Pierre L. Dognin, Keerthiram Murugesan, Erik Miehl, Martín Santillán Cooper, Kieran Fraser, Giulio Zizzo, Muhammad Zaid Hameed, Mark Purcell, Michael Desmond, Qian Pan, Inge Vejsbjerg, Elizabeth M. Daly, Michael Hind, Werner Geyer, and 3 others. 2025. Granite guardian: Comprehensive LLM safeguarding. In *NAACL (Industry Track)*, pages 607–615. Association for Computational Linguistics.
- Ethan Perez, Saffron Huang, Francis Song, Trevor Cai, Roman Ring, John Aslanides, Amelia Glaese, Nat McAleese, and Geoffrey Irving. 2022. Red teaming language models with language models. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*.
- John C Platt. 1999. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Advances in Large Margin Classifiers*, pages 61–74.
- Amrith Rawat, Stefan Schoepf, Giulio Zizzo, Giandomenico Cornacchia, Muhammad Zaid Hameed, Kieran Fraser, Erik Miehl, Beat Buesser, Elizabeth M Daly, Mark Purcell, Prasanna Sattigeri,

- Pin-Yu Chen, and Kush R Varshney. 2024. Attack atlas: A practitioner’s perspective on challenges and pitfalls in red teaming GenAI. *arXiv preprint arXiv:2409.15398*.
- Traian Rebedea, Razvan Dinu, Makesh Narsimhan Sreedhar, Christopher Parisien, and Jonathan Cohen. 2023. **NeMo guardrails: A toolkit for controllable and safe LLM applications with programmable rails**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 431–445, Singapore. Association for Computational Linguistics.
- Paul Röttger, Hannah Rose Kirk, Bertie Vidgen, Giuseppe Attanasio, Federico Bianchi, and Dirk Hovy. 2024. Xstest: A test suite for identifying exaggerated safety behaviours in large language models. In *NAACL*.
- Stefan Schoepf, Muhammad Zaid Hameed, Ambrish Rawat, Kieran Fraser, Giulio Zizzo, Giandomenico Cornacchia, and Mark Purcell. 2025. **MAD-MAX: Modular and diverse malicious attack mixtures for automated LLM red teaming**. In *ICML 2025 Workshop on Data in Generative Models, Benchmarking, Understanding, and Governance*.
- Surat Teerapittayanon, Bradley McDanel, and H. T. Kung. 2016. **Branchynet: Fast inference via early exiting from deep neural networks**. In *2016 23rd International Conference on Pattern Recognition (ICPR)*, pages 2464–2469. IEEE.
- Bertie Vidgen, Nino Scherrer, Hannah Rose Kirk, Rebecca Qian, Anand Kannappan, Scott A. Hale, and Paul Röttger. 2024. **Simplesafetytests: a test suite for identifying critical safety risks in large language models**. *Preprint*, arXiv:2311.08370.
- Laura Weidinger, John Mellor, Maribeth Rauh, Conor Griffin, Jonathan Uesato, Po-Sen Huang, Myra Cheng, Mia Glaese, Borja Balle, Atoosa Kasirzadeh, and 1 others. 2021. Ethical and social risks of harm from language models. *arXiv preprint arXiv:2112.04359*.
- Ji Xin, Raphael Tang, Jaejun Lee, Yaoliang Yu, and Jimmy Lin. 2020. **Deebert: Dynamic early exiting for accelerating bert inference**. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2246–2251. Association for Computational Linguistics.
- Bianca Zadrozny and Charles Elkan. 2002. **Transforming classifier scores into accurate multiclass probability estimates**. In *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 694–699.
- Wenjun Zeng, Yuchi Liu, Ryan Mullins, Ludovic Peran, Joe Fernandez, Hamza Harkous, Karthik Narasimhan, Drew Proud, Piyush Kumar, Bhaktipriya Radharapu, Olivia Sturman, and Oscar Wahltinez. 2024. **Shieldgemma: Generative ai content moderation based on gemma**. *arXiv preprint arXiv:2407.21772*.
- Bo Zhang, Cong Gao, Linkang Yang, Bingxu Han, Minghao Hu, Zhunchen Luo, Guotong Geng, Xiaoying Bai, Jun Zhang, Wen Yao, and Zhong Wang. 2025. **SafeConf: A confidence-calibrated safety self-evaluation method for large language models**. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 3483–3495, Suzhou, China. Association for Computational Linguistics.
- Giulio Zizzo, Giandomenico Cornacchia, Kieran Fraser, Muhammad Zaid Hameed, Ambrish Rawat, Beat Buesser, Mark Purcell, Pin-Yu Chen, Prasanna Sattigeri, and Kush R Varshney. 2024. Adversarial prompt evaluation: Systematic benchmarking of guardrails against prompt input attacks on LLMs. In *NeurIPS Safe Generative AI Workshop*.

A Proofs and Theoretical Details

We provide detailed derivations for the propositions stated in Section 3. The population risk of a deferral rule r under 0–1 loss with deferral penalty $c > 0$ is

$$R(r) = P(Y \neq \hat{y}_l(X), r(X)=0) + P(Y \neq \hat{y}_h(X), r(X)=1) + c \cdot P(r(X)=1), \quad (8)$$

where c is a normalized loss-equivalent cost of invoking G_h , distinct from the raw inference costs c_l, c_h used for deployment cost accounting.

A.1 Bayes-Optimal Deferral

Proposition A.1 (Bayes-optimal deferral; [Jitkritum et al., 2023](#)). *The Bayes-optimal deferral rule minimizing (8) is $r^*(x) = \mathbb{I}[\Delta(x) > c]$, where $\Delta(x) = P(Y \neq \hat{y}_l(x) \mid X=x) - P(Y \neq \hat{y}_h(x) \mid X=x)$ is the expected error reduction from deferring to G_h .*

Proof. The population risk of a deferral rule r is

$$R(r) = \mathbb{E}_{x,y} [(1 - r(x)) \cdot \mathbb{I}[y \neq \hat{y}_l(x)] + r(x) \cdot \mathbb{I}[y \neq \hat{y}_h(x)] + c \cdot r(x)] = \mathbb{E}_x [\text{err}_l(x) - r(x) \cdot \alpha(x)] + \text{const}, \quad (9)$$

where $\text{err}_i(x) = P(y \neq \hat{y}_i(x) \mid x)$ is the expected error of guardrail G_i on input x , $\Delta(x) = \text{err}_l(x) - \text{err}_h(x)$ is the expected error reduction, and $\alpha(x) = \Delta(x) - c$ is the net benefit of deferral. The first term does not depend on r . To minimise $R(r)$, we maximise $r(x) \cdot \alpha(x)$ pointwise. Since $r(x) \in \{0, 1\}$, the optimal choice is $r^*(x) = \mathbb{I}[\alpha(x) > 0]$, which yields $r^*(x) = \mathbb{I}[\Delta(x) > c]$. $\square$

A.2 Excess Risk Decomposition

Proposition A.2 (Excess risk; adapted from Jitkrit-tum et al., 2023). *With $\alpha(x) = \Delta(x) - c$, the excess risk of any deferral rule r relative to r^* is $R(r) - R(r^*) = \mathbb{E}_x[|\alpha(x)| \cdot \mathbb{I}\{r(x) \neq r^*(x)\}]$.*

Proof. From the derivation above, $R(r) = \mathbb{E}_x[\text{err}_l(x) - r(x) \cdot \alpha(x)]$ and $R(r^*) = \mathbb{E}_x[\text{err}_l(x) - r^*(x) \cdot \alpha(x)]$. Subtracting:

$$\begin{aligned} R(r) - R(r^*) &= \mathbb{E}_x[(r^*(x) - r(x)) \cdot \alpha(x)] \\ &= \mathbb{E}_x[(\mathbb{I}[\alpha(x) > 0] - \mathbb{I}[r(x) = 1]) \cdot \alpha(x)] \end{aligned} \quad (10)$$

This quantity is non-negative, since: when $r(x)$ and $r^*(x)$ agree, the contribution is zero; when $\alpha(x) > 0$ but $r(x) = 0$ (missed deferral), the contribution is $\alpha(x) > 0$; when $\alpha(x) \leq 0$ but $r(x) = 1$ (unnecessary deferral), the contribution is $-\alpha(x) \geq 0$. Equivalently, $R(r) - R(r^*) = \mathbb{E}_x[|\alpha(x)| \cdot \mathbb{I}\{r(x) \neq r^*(x)\}]$. $\square$

A.3 Confidence-Based Deferral as a Plug-in Approximation

Proof. The Bayes-optimal deferral rule from Proposition A.1 defers on an input x whenever the expected reduction in classification error exceeds the deferral cost:

$$r^*(x) = \mathbb{I}[\Delta(x) > c],$$

where

$$\Delta(x) = \Pr(Y \neq \hat{y}_l(X) \mid X = x) - \Pr(Y \neq \hat{y}_h(X) \mid X = x).$$

Equivalently, defining $\alpha(x) = \Delta(x) - c$, the optimal rule defers exactly when $\alpha(x) > 0$.

In practice, $\Delta(x)$ is unknown because it depends on the conditional errors of both the lightweight and heavy guardrails. The confidence-based rule

$$\hat{r}_\tau(x) = \mathbb{I}[s_l(x) < \tau]$$

therefore acts as a plug-in approximation to r^* , using the calibrated routing score $s_l(x)$ as a proxy for the benefit of deferral. This approximation is accurate when the ordering induced by low $s_l(x)$ aligns with the ordering induced by large $\Delta(x)$. In other words, inputs assigned lower routing confidence by G_l should be precisely those for which deferring to G_h provides larger expected improvement.

A sufficient condition for this alignment is that the heavy model’s conditional error varies slowly

across inputs relative to the lightweight model’s conditional error. Writing

$$\text{err}_i(x) = \Pr(Y \neq \hat{y}_i(X) \mid X = x),$$

we have

$$\Delta(x) = \text{err}_l(x) - \text{err}_h(x).$$

If $\text{err}_h(x)$ is approximately constant, or more generally does not increase with $\text{err}_l(x)$ fast enough to offset changes in $\text{err}_l(x)$, then $\Delta(x)$ is larger on inputs where G_l is more likely to be wrong. In that case, thresholding low $s_l(x)$ approximates thresholding high $\Delta(x)$.

Thus, confidence-based deferral is not optimal merely because a larger model is more accurate on average. It is effective when the lightweight model’s calibrated routing confidence ranks examples in a way that is aligned with the conditional benefit of deferral. $\square$

A.4 Scalar Confidence Calibration Preserves Ranking

Proposition A.3 (Monotone calibration preserves ranking). *For monotone non-decreasing f applied to the scalar routing confidence u_l , $u_l(x_i) \leq u_l(x_j) \Rightarrow s_l(x_i) \leq s_l(x_j)$. If f is strictly increasing, thresholding s_l induces the same acceptance sets as thresholding u_l up to a reparameterization of τ .*

Proof. Let $u_l : X \rightarrow [1/2, 1]$ denote the uncalibrated scalar confidence assigned by the lightweight model to its own hard prediction, and let $x = x$.

$$s_l(x) = f(u_l(x))$$

be the calibrated routing score, where f is monotone non-decreasing.

For any two inputs x_i, x_j , if

$$u_l(x_i) \leq u_l(x_j),$$

then monotonicity of f implies

$$s_l(x_i) = f(u_l(x_i)) \leq f(u_l(x_j)) = s_l(x_j).$$

Therefore, the calibrated score preserves the ordering induced by the raw routing confidence.

If f is strictly increasing, then thresholding s_l induces the same acceptance sets as thresholding u_l , up to a reparameterization of the threshold. Specifically, for any calibrated threshold τ , there exists an uncalibrated threshold $\tau' = f^{-1}(\tau)$ such that

$$\{x : s_l(x) \geq \tau\} = \{x : u_l(x) \geq \tau'\}.$$

If f is monotone non-decreasing but not strictly increasing, then flat regions of f may map distinct raw confidence values to the same calibrated score. In this case, the ordering is preserved up to ties, and finite-sample differences may arise because of tie-breaking or discrete threshold grids.

Isotonic regression, temperature scaling with $T > 0$, and Platt scaling with nonnegative slope are monotone non-decreasing transformations of the scalar confidence score. Hence, in the setting considered here, these calibrators improve the reliability and interpretability of routing thresholds without introducing new ranking information. $\square$

A.5 Extension to K -Model Cascades

For K guardrails with standalone inference costs $c^{(1)} \leq \dots \leq c^{(K)}$, the Bayes-optimal selector (which picks one guardrail per input) generalises Proposition A.1 (Jitkrittum et al., 2023):

$$s^*(x) = \arg \min_{k \in [K]} (1 - \eta_{\hat{y}_k(x)}(x) + c^{(k)}). \quad (11)$$

This selects the guardrail that minimises the sum of misclassification risk and standalone invocation cost. In a *sequential* cascade, however, the cost to reach G_k is cumulative: every input pays at least $c^{(1)}$; inputs routed to the second stage pay $c^{(1)} + c^{(2)}$; more generally, inputs reaching the k -th stage pay $\sum_{j \leq k} c^{(j)}$. Extending calibrated routing to this sequential multi-stage setting is left for future work.

B Full Benchmark Results

Table 3 presents detection and calibration results for all 21 model variants evaluated in this study. The 18 models from the four core families (Granite Guardian, Llama Guard, ShieldGemma, GuardReasoner) are identical to those reported in the main paper (Table 1); this table additionally includes Qwen3Guard and Nemotron, and reports the learned temperature T for each model as a diagnostic of overconfidence.

Figure 2 shows reliability diagrams for the model variants not included in the main paper.

Figure 3 shows cross-family coverage-risk trade-off.

C Per-Family Deferral Analysis

Figure 4 shows the unified comparison across families (joint accuracy vs. coverage and vs. average cost). Figure 5 then presents the full four-panel

deferral analysis (Accuracy, F1, Threshold, ECE vs. Coverage) for each of the five model families under four calibration cases.

Granite Guardian (Figure 5a). Both Accuracy and F1 improve steadily as coverage decreases (i.e., more inputs are deferred to G_h). The “Ideal (Both GT)” case achieves the lowest ECE across all coverage levels, confirming that calibrating both models on ground truth provides the best routing signal. The threshold curve declines smoothly, indicating well-separated confidence distributions.

Llama Guard (Figure 5b). The “No Calibration” baseline shows limited improvement with deferral, as the uncalibrated confidence scores provide poor routing signal. Calibration on ground truth substantially improves F1, particularly at intermediate coverage levels (50 to 70%). The threshold curve is steeper than Granite Guardian, reflecting less reliable confidence estimates.

GuardReasoner (Figure 5c). GuardReasoner exhibits a distinctive pattern: the small model (1B) already achieves high F1 (0.84), so deferral to the 8B model provides modest accuracy gains. However, calibration significantly reduces ECE, improving the reliability of the confidence estimates even when the accuracy improvement is small.

ShieldGemma (Figure 5d). The accuracy gap between SG-2B and SG-27B is negligible (0.001), rendering deferral ineffective regardless of calibration strategy. All four cases produce nearly identical curves, confirming that deferral requires a meaningful performance gap between G_l and G_h .

Qwen3Guard (Figure 5e). The within-family gap is small (0.014), so deferral provides limited benefit. However, QG-0.6B achieves the highest AUC among all models (0.91), making it an excellent candidate for cross-family deferral to stronger models such as GG-3.1-8B (see Section 6).

D Dataset Details

We evaluate on out-of-distribution (OOD) harm detection. All datasets are publicly available. Following the evaluation protocol of Padhi et al. (2024), we assign a positive (harmful) label to any instance marked as unsafe under its original taxonomy. Table 4 reports per-benchmark statistics for the 10 evaluation subsets.

Table 3: Full detection and calibration results across all 21 model variants evaluated in this study. The main paper (Table 1) reports the same 18 models from the four core families (Granite Guardian, Llama Guard, ShieldGemma, GuardReasoner); this table additionally includes Qwen3Guard and Nemotron, and reports the learned temperature T for each model (higher T indicates greater overconfidence). ECE in %.

Model	AUC	F1	ECE _{unc}	ECE _{iso}	ECE _{sig}	ECE _{temp}	T
GG-3.0-2B	0.79	0.68	10.6	1.1	4.3	2.9	2.4
GG-3.0-8B	0.87	0.76	6.8	2.4	4.4	2.5	1.7
GG-3.1-2B	0.86	0.75	6.2	2.3	4.1	4.5	1.8
GG-3.1-8B	0.90	0.80	6.7	1.2	2.4	3.5	2.0
GG-3.2-3B	0.85	0.75	8.5	1.2	2.9	3.7	2.2
GG-3.2-5B	0.88	0.78	9.9	1.7	3.5	3.5	1.8
GG-3.3-8B	0.89	0.80	13.9	1.8	3.1	2.1	4.7
GG-4.1-8B	0.90	0.80	12.7	1.9	1.8	2.3	3.7
LG-7B	0.83	0.66	14.2	3.1	3.9	2.5	2.9
LG-2-8B	0.85	0.74	15.4	2.6	3.0	3.0	4.6
LG-3-1B	0.80	0.65	17.0	1.4	4.6	2.0	3.7
LG-3-8B	0.83	0.72	14.8	2.5	2.5	4.7	4.1
SG-2B	0.74	0.39	26.5	2.2	5.6	6.5	5.8
SG-9B	0.73	0.38	26.4	2.5	5.8	7.1	6.6
SG-27B	0.76	0.41	27.5	2.2	2.7	4.9	6.1
GR-1B	0.86	0.84	17.2	1.7	1.1	0.7	8.3
GR-3B	0.83	0.83	18.3	1.9	1.5	1.9	10.0
GR-8B	0.82	0.82	18.7	1.0	0.4	1.9	10.0
QG-0.6B	0.90	0.81	11.8	1.6	3.9	3.2	3.2
QG-8B	0.90	0.84	11.3	1.6	6.3	4.5	3.7
Nemotron-8B	0.77	0.72	17.7	1.2	3.8	13.8	5.5




Figure 2: Reliability diagrams for all remaining model variants (smaller models within each family). Red curves show uncalibrated probabilities; green curves show isotonic-calibrated probabilities. The dashed diagonal indicates perfect calibration.

From raw datasets to evaluation subsets. The sample counts listed below (e.g., ToxicChat 2,853; SafeRLHF 1,000; TRUE datasets) are the raw sizes of each source benchmark. The counts reported in

Table 4 of the main paper are the filtered evaluation subsets used after preprocessing, i.e., (i) keeping only human-annotated entries where multiple annotation types exist, (ii) removing instances with




Figure 3: Full deferral risk heatmap. Rows above the black separator are within-family cascades; rows below are cross-family cascades. Risk is the average 1–Accuracy across coverage levels; lower (green) is better. Cross-family deferral to GG-3.1-8B (LG-3-1B, GR-1B, QG-0.6B) achieves the lowest residual risk under the Ideal Annotation case; under the other three calibration cases the cross-family pairs do not uniformly beat their within-family counterparts.

Table 4: Benchmark datasets used in our evaluation.

Dataset	Task	N	%Pos
ToxicChat	Prompt harm	1,427	7.2
OpenAI Mod.	Prompt harm	1,680	42.3
AegisSafetyTest	Prompt harm	363	50.1
SimpleSafetyTests	Prompt harm	100	100
HarmBench	Prompt harm	320	100
BeaverTails	Response harm	3,021	47.8
SafeRLHF	Response harm	2,000	50.0
XSTest-RR	Response harm	449	50.0
XSTest-RH	Response harm	446	50.0
XSTest-RR(h)	Response harm	200	50.0

missing labels or invalid model outputs, (iii) class-balancing within each benchmark when applicable, and (iv) sub-sampling the largest benchmarks to avoid overweighting them in the aggregate. The 80/20 calibration-evaluation split (Section D, last subsection) is applied to these filtered subsets.

D.1 Harm Detection Benchmarks

Prompt Harmfulness.

- **ToxicChat** (Lin et al., 2023): 2,853 human-annotated samples from real user-chatbot interactions. We use version 0124 and select only human-annotated samples, assigning the harmful label if either the toxicity or jailbreak flag is positive.
- **OpenAI Moderation** (Markov et al., 2023):

1,680 prompts labeled across eight safety categories (sexual, hate, violence, harassment, self-harm, sexual/minors, hate/threatening, violence/graphic).

- **AegisSafetyTest** (Ghosh et al., 2024): 359 prompt-only entries from Anthropic’s HH-RLHF, manually annotated across 13 risk categories.
- **SimpleSafetyTests** (Vidgen et al., 2024): 100 manually crafted harmful prompts targeting sensitive topics including child abuse, suicide, self-harm, and fraud.
- **HarmBench Prompt** (Mazeika et al., 2024): 239 harmful prompts designed to test robustness against jailbreak attacks.

Response Harmfulness.

- **BeaverTails** (Ji et al., 2023): 3,021 prompt-response pairs from the test set, focusing on response harmfulness with responses generated by Alpaca-7B.
- **SafeRLHF** (Dai et al., 2024): 1,000 prompt-response pairs with manually annotated preference comparisons between safe and unsafe responses.
- **XSTest** (Röttger et al., 2024): Three variants testing the boundary between safe and unsafe content. XSTest-RH (446 samples) tests refusal of harmless prompts; XSTest-RR (449 samples) and XSTest-RR(h) (200 samples) test compliance with risky prompts.

D.2 Data Splitting Protocol

All datasets are split 80/20 into calibration and evaluation partitions. The split is performed at the instance level with a fixed random seed (42) to ensure reproducibility. The calibration partition is used exclusively for fitting post-hoc calibrators (isotonic regression, Platt scaling, or temperature scaling). All metrics reported in the paper are computed on the held-out 20% evaluation partition. This protocol ensures that calibration and evaluation are performed on disjoint data, preventing information leakage.

D.3 Aggregate Deferral Risk

The heatmaps in Figures 3 summarise cascade performance with a single scalar we refer to as the *aggregate deferral risk*. For a cascade (G_l, G_h) with




Figure 4: Deferral performance across model families (isotonic calibration, Calibrated GT case). Left: joint accuracy vs coverage. Right: joint accuracy vs average inference cost.

a confidence-based deferral rule, let $\text{Acc}(\phi)$ denote the cascade’s accuracy when a fraction $\phi \in [0, 1]$ of inputs is handled by G_l (the rest is deferred to G_h). The aggregate risk is the mean error rate across a coverage grid $\Phi = \{0.5, 0.6, 0.7, 0.8, 0.9, 1.0\}$:

$$\overline{\text{Risk}}(G_l, G_h) = \frac{1}{|\Phi|} \sum_{\phi \in \Phi} (1 - \text{Acc}(\phi)). \quad (12)$$

This is a threshold-free scalar summary: each coverage level gives a different operating trade-off between cost and accuracy, and $\overline{\text{Risk}}$ averages the error rate across these operating points. It is analogous to AURRRC but without the cost axis; lower is better. We report the deployment-relevant range $\phi \in [0.5, 1.0]$: below $\phi = 0.5$ the cheap model handles only a minority of traffic, a regime rarely cost-effective in deployment because the cascade pays the double-inference overhead on the majority of inputs. The underlying risk-coverage curves (Figures 4 and 5) extend to lower coverage.

E Calibration Methods: Extended Discussion

E.1 Isotonic Regression

Isotonic regression (Zadrozny and Elkan, 2002; Niculescu-Mizil and Caruana, 2005) fits a non-parametric, monotone piecewise-constant function f by solving:

$$\min_{f \in \mathcal{F}_{\text{iso}}} \sum_{j=1}^n (f(\hat{p}(x_j)) - y_j)^2 \quad \text{subject to } f \text{ non-decreasing}, \quad (13)$$

where $\mathcal{F}_{\text{iso}}$ is the class of non-decreasing step functions. The solution is computed efficiently via the pool adjacent violators algorithm in $O(n \log n)$

time. For out-of-sample predictions, values outside the training range are clipped to the nearest boundary value.

Strengths: Can capture arbitrary monotone probability mappings; achieves the lowest ECE in most experiments. **Weaknesses:** May overfit with small calibration sets (< 500 samples); the step-function output can produce ties in confidence scores.

E.2 Platt Scaling (Sigmoid Calibration)

Platt scaling (Platt, 1999) fits a two-parameter logistic function:

$$f(z) = \sigma(az + b) = \frac{1}{1 + \exp(-(az + b))}, \quad (14)$$

where $a, b \in \mathbb{R}$ are learned by minimizing the negative log-likelihood on the calibration set. The optimization is convex and solved via L-BFGS.


Strengths: Robust with small calibration sets; smooth output probabilities. **Weaknesses:** Assumes a sigmoidal relationship between raw and calibrated probabilities, which may not hold for all model architectures.

E.3 Temperature Scaling


Temperature scaling (Guo et al., 2017) is a special case of Platt scaling with $a = 1/T$ and $b = 0$:

$$f(z) = \sigma(\text{logit}(z)/T), \quad \text{where } \text{logit}(z) = \log \frac{z}{1 - z}. \quad (15)$$


The single parameter $T > 0$ is learned by minimizing the negative log-likelihood. Values $T > 1$ indicate overconfidence (the model’s raw probabilities are too extreme); $T < 1$ indicates underconfidence.



(a) Granite Guardian: $G_l = \text{GG-3.1-2B}$, $G_h = \text{GG-3.1-8B}$. Accuracy gap: 4.2%.




(b) Llama Guard: $G_l = \text{LG-3-1B}$, $G_h = \text{LG-3-8B}$. Accuracy gap: 4.5%.



(c) GuardReasoner: $G_l = \text{GR-1B}$, $G_h = \text{GR-8B}$. Accuracy gap: 0.4%.



(d) ShieldGemma: $G_l = \text{SG-2B}$, $G_h = \text{SG-27B}$. Accuracy gap: 0.1%.



(e) Qwen3Guard: $G_l = \text{QG-0.6B}$, $G_h = \text{QG-8B}$. Accuracy gap: 1.4%.

Figure 5: Per-family deferral analysis. Each row shows four panels: Accuracy, F1, Threshold, and ECE as a function of coverage under four calibration cases (No Calibration, Calibrated on GT, Calibrated on Pseudo-labels, Ideal with Both GT). Deferral is most effective when the accuracy gap between G_l and G_h is large (Granite Guardian, Llama Guard) and least effective when the gap is negligible (ShieldGemma).

Strengths: Preserves the ranking of predictions (and thus AUC and the deferral curve per Proposition A.4); single parameter, minimal overfitting risk. **Weaknesses:** Cannot correct non-linear miscalibration patterns.

E.4 Model Inference Cost

Table 5 reports per-model cost accounting based on parameters and output tokens. Granite Guardian 3.0/3.1 models are the most token-efficient, generating a single Yes/No token, whereas GuardReasoner generates ~ 300 reasoning tokens, making its 1B variant ($0.8\times$) comparable in cost to GG-3.1-2B ($1.0\times$) despite having half the parameters.

Table 5: Per-model inference cost. Cost is proportional to parameters $\times$ (input + output tokens), normalized to GG-3.1-2B. Output tokens vary by architecture.

Model	Params	Out. tok.	Cost	Output format
GG-3.1-2B	2B	1	$1.0\times$	Yes/No
GG-3.1-8B	8B	1	$4.0\times$	Yes/No
GG-3.3-8B	8B	16	$4.1\times$	$\langle \text{score} \rangle$ tags
LG-3-1B	1B	3	$0.5\times$	safe/unsafe
LG-3-8B	8B	3	$4.0\times$	safe/unsafe
SG-2B	2B	1	$1.0\times$	Yes/No
SG-27B	27B	1	$13.5\times$	Yes/No
GR-1B	1B	~ 300	$0.8\times$	Reasoning chain
GR-8B	8B	~ 300	$6.4\times$	Reasoning chain
QG-0.6B	0.6B	1	$0.3\times$	Yes/No
QG-8B	8B	1	$4.0\times$	Yes/No

E.5 Empirical Comparison

In our experiments, the learned temperature values reveal the degree of overconfidence across model families:

- **Granite Guardian:** $T \approx 2$ to 5, indicating moderate overconfidence.
- **Llama Guard:** $T \approx 3$ to 5, similar to Granite Guardian.
- **ShieldGemma:** $T \approx 6$ to 7, indicating substantial overconfidence consistent with the high uncalibrated ECE (26 to 28%).
- **GuardReasoner:** $T \approx 8$ to 10, the most overconfident family. This is consistent with the reasoning-based architecture, which may produce highly peaked probability distributions.
- **Qwen3Guard:** $T \approx 3$ to 4, moderate overconfidence.
- **Nemotron:** $T \approx 5.5$, but temperature scaling fails to reduce ECE effectively (13.8% after scaling vs. 17.7% uncalibrated), suggesting non-sigmoidal miscalibration.

F Calibration Cases: Formal Definitions and Deferral-Aligned Assumption

Case 1: No Calibration (Case Zero). No post-hoc calibration is applied. We set $s_l(x) = u_l(x)$ and sweep the threshold τ directly on the raw confidence.

Case 2: Deferral Annotation. No ground-truth labels are available, but both G_l and G_h can be evaluated on unlabeled production inputs. We calibrate the lightweight model’s confidence to predict agreement with the heavier model:

$$f_l^{\text{agr}} = \arg \min_{f \in \mathcal{F}} \sum_{j=1}^n \ell(f(u_l(x_j)), a_j), \quad (16)$$

$$a_j = \mathbb{I}[\hat{y}_l(x_j) = \hat{y}_h(x_j)].$$

The resulting score estimates $\mathbb{P}(\hat{Y}_l = \hat{Y}_h \mid u_l)$, not $\mathbb{P}(\hat{Y}_l = Y \mid u_l)$, and is therefore a teacher-agreement proxy for routing. This calibration affects only the routing score; the hard predictions $\hat{y}_l$ and $\hat{y}_h$ are kept fixed.

Case 3: Annotation. Ground-truth labels are available. We calibrate the lightweight model’s confidence to its empirical correctness:

$$f_l^{\text{gt}} = \arg \min_{f \in \mathcal{F}} \sum_{j=1}^n \ell(f(u_l(x_j)), z_j), \quad (17)$$

$$z_j = \mathbb{I}[\hat{y}_l(x_j) = y_j].$$

This is the standard post-hoc confidence calibration setting for selective prediction and deferral.

Case 4: Ideal Annotation. Both G_l and G_h are calibrated on ground truth. This setting is a fully supervised calibration reference; it is not a Bayes-optimal oracle, since $\Delta(x)$ is not observed directly.

Assumption F.1 (Deferral-aligned confidence). For a model pair (G_l, G_h) , the expected benefit of deferral

$$\Delta(x) = \mathbb{P}(Y \neq \hat{y}_l(X) \mid X=x) - \mathbb{P}(Y \neq \hat{y}_h(X) \mid X=x)$$

is larger on inputs assigned lower calibrated confidence by G_l , i.e., the ordering induced by $s_l(x)$ is

a useful proxy for the ordering induced by $-\Delta(x)$. For Case 2, we additionally require that agreement with G_h be positively correlated with correctness of G_l .

Under this assumption, confidence thresholding is a plug-in approximation to the Bayes-optimal deferral rule. The assumption is weaker than pointwise dominance: G_l may outperform G_h on some instances, but low-confidence regions of G_l must be regions where deferral is beneficial on average.

G Calibration Case Ablation

Using the four calibration cases formally defined in Appendix F, we report the empirical comparison on the Granite Guardian 3.1 family (2B $\rightarrow$ 8B) as the primary case study. Because the calibrators are monotone, we interpret the accuracy-coverage curves as evidence about the quality of the routing score and its threshold mapping rather than as evidence of new ranking information.

H Cost Scenarios and Pareto Frontier

Expected cost of the cascade. Every input is first evaluated by G_l . Confident inputs ($r(x)=0$) incur cost c_l only; deferred inputs ($r(x)=1$) incur $c_l + c_h$. At coverage $\phi = P(r(x)=0)$, the expected cost is:

$$\bar{C}(\phi) = c_l + (1 - \phi) \cdot c_h. \quad (18)$$

This expression assumes fixed per-example costs c_l and c_h ; with input- or output-dependent costs, the corresponding quantity is $\mathbb{E}[c_l(X) + r(X)c_h(X)]$. We compare against the baseline of using G_h alone, with cost c_h . The fractional saving is

$$\text{Saving} = 1 - \frac{c_l + (1 - \phi) \cdot c_h}{c_h} = \phi - \frac{c_l}{c_h}. \quad (19)$$

Positive savings require $\phi > c_l/c_h$. For a 2B/8B cascade with $c_l/c_h = 0.25$, break-even occurs at $\phi = 0.25$; at $\phi = 0.8$, the saving is 55%.

Table 6 reports cascade costs under two deployment scenarios. In *Scenario 1: Fixed 80% coverage*, we require G_l to handle 80% of inputs, so only the remaining 20% are deferred to G_h . In *Scenario 2: Accuracy $\geq 99\%$ of G_h* , we choose the maximum coverage that preserves at least 99% of the accuracy of G_h alone. For the GG-2B $\rightarrow$ GG-8B cascade, Scenario 2 is satisfied at 80% coverage, yielding the same cost and savings as Scenario 1, indicating that the cascade maintains near-optimal accuracy up to this operating point before degradation begins.

Table 6: Cascade cost under two deployment scenarios. c_l, c_h account for model parameters and output tokens (Table 5).

G_l	G_h	Cov.	Acc	Cost	Save
<i>Scenario 1: Fixed 80% coverage</i>					
GG-3.1-2B	GG-3.1-8B	80%	.832	$1.8\times$	55%
LG-3-1B	LG-3-8B	80%	.751	$1.3\times$	68%
<i>Scenario 2: Accuracy $\geq 99\%$ of G_h</i>					
GG-3.1-2B	GG-3.1-8B	80%	.832	$1.8\times$	55%




Figure 6: Cost-accuracy Pareto frontier for Granite Guardian cascades (isotonic calibration). Square markers indicate small-model-only performance while diamonds large-model-only. The curves trace the trade-off as coverage varies from 30% to 100%.

I Calibration Data Size Ablation

A practical question for deployment is: how much calibration data is needed? We vary the calibration/test split from 50/50 to 90/10 across five model families and two calibration methods.

Key findings:

- **Isotonic regression** shows a clear downward trend in ECE as the calibration fraction increases, particularly for Granite Guardian (ECE drops from 2.8% at 50% calibration to 1.2% at 90%). This confirms that the non-parametric method benefits from more data.
- **Sigmoid scaling** is remarkably stable, with ECE varying by less than 0.5 percentage points across all split ratios. This makes it the safer choice when calibration data is limited.
- The **overfitting gap** (difference between calibration-set ECE and test-set ECE) is largest for isotonic regression at small calibration fractions, confirming the overfitting risk.




Figure 7: ECE on the test set as a function of calibration set fraction. Isotonic regression benefits more from additional data, while sigmoid scaling is stable across all split ratios.

- For practical deployment, we recommend at least 2,000 calibration samples for isotonic regression and 500 for sigmoid scaling.

J Deferral Threshold Selection

The deferral threshold τ in Equation (4) controls the coverage-accuracy trade-off. We recommend the following deployment procedure:

1. **Calibrate** the small model G_l on a held-out calibration set using isotonic regression (or sigmoid scaling if the calibration set is small).
2. **Compute the coverage-accuracy curve** by sweeping τ from 0.5 to 1.0 on the calibration set.
3. **Select** τ to achieve the desired coverage target. For example, $\tau = 0.85$ typically yields 70 to 80% coverage for well-calibrated Granite Guardian models.
4. **Monitor** the deferral rate in production. If the rate deviates significantly from the calibration-set estimate, this indicates distribution shift and the threshold should be re-tuned.

The threshold plots in Figure 5 show that well-calibrated models (Granite Guardian, GuardReasoner) exhibit smooth, monotonically decreasing thresholds as coverage increases, making threshold selection straightforward. Poorly calibrated models (ShieldGemma) show noisy threshold curves, indicating that the confidence scores are less reliable for deferral decisions.

K Model Details and Probability Extraction

K.1 Probability Extraction

All models are evaluated using vLLM (Kwon et al., 2023) for efficient batched inference. The probability of risk (i.e., $\hat{p}(x) = P(y=1 | x)$) is extracted as follows:

Table 7: Guardrail models evaluated in this study.

Family	Model	Params	Prob. Method
Granite Guardian	GG-3.0-2B	2B	Top- k logprob (Yes/No)
	GG-3.0-8B	8B	Top- k logprob (Yes/No)
	GG-3.1-2B	2B	Top- k logprob (Yes/No)
	GG-3.1-8B	8B	Top- k logprob (Yes/No)
	GG-3.2-3B	3B	Top- k logprob (Yes/No)
	GG-3.2-5B	5B	Top- k logprob (Yes/No)
	GG-3.3-8B	8B	$\langle \text{score} \rangle$ tag logprob
	GG-4.1-8B	8B	Guardian block logprob
Llama Guard	LG-7B	7B	First-token logit (safe/unsafe)
	LG-2-8B	8B	First-token logit
	LG-3-1B	1B	First-token logit
	LG-3-8B	8B	First-token logit
ShieldGemma	SG-2B	2B	Vocab logit (Yes/No)
	SG-9B	9B	Vocab logit
	SG-27B	27B	Vocab logit
GuardReasoner	GR-1B	1B	Decision-token logit
	GR-3B	3B	Decision-token logit
	GR-8B	8B	Decision-token logit
Qwen3Guard	QG-0.6B, 8B	0.6/8B	Safety-token logit
Nemotron	Nemotron-8B	8B	JSON-token logit

Granite Guardian 3.0/3.1/3.2. Following Padhi et al. (2024), we compute the detection scores by aggregating the exponentiated log-likelihoods of all tokens containing the substring “Yes” or “No” within the top- k ($k=20$) tokens at the first generated position. The scores are normalized via softmax to produce the probability of risk.

Granite Guardian 3.3/4.1. These models output $\langle \text{score} \rangle \text{yes/no} / \langle \text{score} \rangle$ tags. We extract logprobs at the token position immediately following the $\langle \text{score} \rangle$ tag and apply the same softmax normalization over Yes/No logprobs.

Llama Guard. The model generates “safe” or “unsafe” as the first token. We extract the logits for these two tokens at the last input position and apply softmax.

ShieldGemma. The model is prompted to answer “Yes” or “No”. We extract the logits for the vocabulary tokens “Yes” and “No” at the last input position and apply softmax.

GuardReasoner. The model generates a reasoning chain followed by a label. We scan the gen-

erated tokens for the position where “harmful” or “unharmful” appears after “Request:” and extract logprobs at that position.

Qwen3Guard. The model generates “Safety: Safe/Unsafe/Controversial”. We extract logprobs at the token position where the safety label appears and compute $P(\text{unsafe}) = P(\text{Unsafe}) + P(\text{Controversial})$.

Nemotron. The model generates a JSON response with “User Safety”: “safe”/“unsafe”. We extract logprobs at the token position where the safety value appears.

When neither the safe nor unsafe token appears in the top- k logprobs, we treat the missing token’s logprob as -50 (effectively $-\infty$) and use the generated label to determine the direction. This ensures that highly confident predictions (where the alternative token falls outside top- k) receive probabilities close to 0 or 1 rather than the uninformative value of 0.5.