

Agora: Enhancing LLM Agent Reasoning Via Auction-Based Task Allocation

Kaiji Zhou, Aleš Leonardis, and Yue Feng*

School of Computer Science, University of Birmingham
Birmingham, United Kingdom
kxz571@student.bham.ac.uk
{a.leonardis,y.feng,6}@bham.ac.uk

Abstract

Enhancing the reasoning capabilities of large language model (LLM) agents requires effective orchestration of diverse expert models and tools. However, existing frameworks typically call APIs, based on coarse-grained matching between tasks and the functions of expert models or tools, while overlooking critical factors such as performance variability and cost efficiency among functionally similar alternatives. To address this, we propose Agora, a framework that uses a confidence-calibrated auction to dynamically allocate tasks to expert models and tools. By treating reasoning steps as tradeable items, Agora bases allocation on calibrated competence rather than raw confidence. Across five main benchmarks, Agora improves or remains competitive with single-model, routing, and cascade baselines under matched candidate pools.¹

1 Introduction

Advancing the reasoning capabilities of Large Language Models (LLMs) requires moving beyond monolithic execution. While strategies like Chain-of-Thought (CoT) (Wei et al., 2022) provide a structural blueprint by decomposing queries into atomic steps, executing these complex chains often exceeds the reliable scope of any single generalist model. A critical bottleneck restricts current reasoning systems: *the mismatch between task difficulty and model capability*. Static routing ignores query-dependent differences in model suitability (Nguyen et al., 2024). These limitations motivate *dynamic competence discovery*: routing each reasoning step to a suitable agent.

However, implementing this fine-grained orchestration faces two hurdles: *structural alignment* and *trustworthy valuation*. First, coarse-grained routing (at the query level) fails to exploit the sub-task

structures generated by planners. Second, and more critically, establishing a reliable auction is plagued by *overconfidence* (Huang et al., 2025). Agents often hallucinate certainty, claiming high confidence on incorrect answers. Without a reliable measure of an agent’s *true* probability of success, dynamic allocation risks assigning critical logic nodes to overconfident but incompetent agents, causing the reasoning chain to collapse.

To address these challenges, we propose *Agora*, a framework that reformulates task allocation as a confidence-calibrated *auction*. Specifically, Agora operates in two phases: a *Planner* decomposes the query into atomic units, and an *Auction* treats these units as tradeable items. Agents compete to solve them by submitting “bids” derived from their execution cost and *calibrated confidence*. Crucially, Agora employs a hierarchical calibration strategy that combines static baselines with online adaptation to filter out hallucinated certainty. This lets routing adapt to calibrated, step-specific estimates of agent reliability.

In summary, our contributions are as follows:

- *Auction-Based Reasoning Framework*: We propose Agora, a framework that leverages an auction mechanism to dynamically route reasoning steps, enabling specialized agents to collaborate efficiently on complex tasks.
- *Competence-Driven Calibration*: We introduce a strategy combining embedding-based binning with online refinement to standardize confidence estimation, ensuring that allocation is based on reliability rather than hallucinated certainty.
- *Empirical Validation*: Across five main benchmarks, Agora improves or remains competitive under matched candidate pools while making the cost-quality trade-off directly tunable.

*Corresponding author.

¹Code: <https://github.com/KAIJUJO/Agora>




Figure 1: Overview of Agora’s *auction-based reasoning framework*. Given a complex query x , a planner decomposes it into a graph of dependencies, which are grouped into *task units*. Each unit is dynamically allocated to the highest-scoring agent via a *confidence-calibrated auction* that balances calibrated competence against execution cost. Finally, the unit outputs are synthesized into the final answer Y .

2 Related Work

2.1 Model Selection and Specialized Reasoning

Enhancing reasoning capabilities relies on coordinating diverse expertise. The prohibitive cost of generalist models has spurred research on *dynamic model selection* to access specialized capabilities efficiently. Early heuristic *cascades*, such as FrugalGPT, sequentially query models to cut costs but suffer from serial latency (Chen et al., 2024). To address this, recent work employs learned *routers*. Supervised classifiers like HybridLLM (Ding et al., 2024) and AdaptiveLLM (Cheng et al., 2025) optimize accuracy under budgets. More recent approaches learn routers from preference data (RouteLLM (Ong et al., 2025)) or use structural modeling (MetaLLM (Nguyen et al., 2024)) to adapt to query complexity. Building on this, TensorOpera Router achieves significant throughput gains by projecting queries into a learned latent space to map inputs to experts (Stripelis et al., 2024).

Parallel research pushes towards finer granularities. Techniques like Confidence-Token routing (Chuang et al., 2025) and Mixture-of-Depths (MoD) (Raposo et al., 2024) optimize compute at the *token* level by skipping unnecessary operations.

Parallel to system-level routing, architectural *Mixture-of-Experts (MoE)* methods have gained prominence for scaling model capacity. Archi-

tectures like Switch Transformers (Fedus et al., 2022) and Mixtral (Jiang et al., 2024) route tokens to specific internal expert layers to optimize computational efficiency. Recent innovations like DeepSeek-MoE (Dai et al., 2024) further refine this by employing fine-grained expert segmentation. However, while these methods share the philosophy of specialization, they require end-to-end pre-training. In contrast, our framework focuses on *auction-based orchestration*, enabling *diverse agents* (e.g., proprietary planners combined with open-weight executors) to collaborate on reasoning tasks without architectural modifications.

Crucially, existing systems operate at extremes: either routing the *entire query* (coarse-grained) or individual *tokens* (fine-grained). They lack the semantic granularity to decompose complex problems into distinct *sub-tasks*. Our framework addresses this gap by auctioning *task units* using calibrated competence and execution cost.

2.2 Calibration for Reliable Valuation

In an auction-based system, reliable routing hinges on accurate self-valuation. Early “Ask-for-calibration” methods prompted models to verbalize certainty (Tian et al., 2023), and follow-up studies refined this via prompting strategies and aggregation (Xiong et al., 2024; Yang et al., 2024). Post-hoc methods like QA-Calibration improved reliability via embedding-based grouping (Mang-

gala et al., 2025), yet they typically rely on static calibration sets.

Recent advancements leverage continuous semantic spaces for greater precision. Kernel Language Entropy (KLE) calculates uncertainty via semantic similarity kernels (Nikitin et al., 2024), while Semantic Nearest Neighbor Entropy (SNNE) aggregates pairwise similarities to handle long-form generation (Nguyen et al., 2025). These methods provide complementary semantic uncertainty estimates. Agora instead applies static and online calibration to self-reported confidence before allocation.

2.3 Mechanism-Driven Reasoning

Before modern LLM agents, task-oriented dialogue systems tracked structured states (Feng et al., 2021); Chain-of-Thought (Wei et al., 2022) and Least-to-Most prompting (Zhou et al., 2023) later made reasoning decomposition explicit. Hugging-GPT (Shen et al., 2023), Chameleon (Lu et al., 2023), and Gradientsys (Song et al., 2025) dispatch sub-tasks to specialists; Med-VRAgent combines visual guidance, self-reward, and tree search for medical reasoning (Guo et al., 2025). However, they rarely optimize competence against cost.

To ground collaborative reasoning in economic principles, we draw on *mechanism design*. Dütting et al. study incentive properties for token-level LLM aggregation under explicit preference and payment rules (Dütting et al., 2024). Inspired by this perspective, Agora treats sub-tasks as auctionable items and allocates them by balancing calibrated competence against execution cost.

Multi-agent debate can improve reasoning and factuality through social interaction (Du et al., 2024). However, game-theoretic evaluations reveal substantial limits in LLM rationality (Fan et al., 2024). This motivates calibrating self-reported confidence before allocation.

3 Methodology

Our framework follows a closed-loop pipeline: *Planning* $\rightarrow$ *Calibration* $\rightarrow$ *Auction* $\rightarrow$ *Execution* $\rightarrow$ *Refinement* (see Fig. 1). Given a user request x , the *Planning* module decomposes it into a directed graph of atomic task units. In the *Calibration* phase, potential agents estimate their success probabilities using a nested mechanism: a pre-trained calibrator S for base accuracy, and a dynamic post-hoc calibrator S' that adapts to distribution shifts. This

gives the auction a reliability-adjusted basis for allocation. The *Auction* module assigns each task unit using a score that balances calibrated confidence against execution cost. Selected agents *Execute* the subtasks, and a composer synthesizes the *Final Answer*. Finally, the *Refinement* module can use evaluator feedback to update S' over time.

3.1 Task Planning and Unit Extraction

The *LLM Planner* (f_1) first interprets the input x to generate a directed task graph $G = (V, E)$, where nodes represent atomic steps with defined skill requirements and data dependencies (see Appendix A for detailed prompts).

Subsequently, the *Task-unit Extractor* optimizes this graph by merging tightly coupled nodes into coarser executable *task units* T .

Candidate groups are identified via coupling signals, specifically shared I/O, prompt similarity, or explicit control dependencies.

To ensure feasibility, a merge is finalized only if it satisfies a decision gate: (i) the resulting unit must be supported by agent capabilities and (ii) maintain structural validity within the graph.

3.2 Confidence Calibration

We define the calibrated confidence $\hat{p}_{ij}$ for agent a_j on task t_i using a hierarchical composition of static and dynamic transformations.

Static Calibration (S). To improve calibration across diverse domains, we employ a static calibrator S trained on a *broad-spectrum corpus of diverse public benchmarks* (including math, coding, and multimodal datasets; details in Appendix B). Instead of fitting a single global scaler, S applies a group-specific scaling followed by histogram binning:

$$\hat{p}_0 = S(p_{\text{raw}}) = \mathcal{S}_{\text{bin}} \left(\sigma(\mathbf{w}_g^T \cdot \phi(p_{\text{raw}}) + b_g) \right), \quad (1)$$

where p_{raw} is the raw confidence, and parameters $(\mathbf{w}_g, b_g)$ are specific to task group g , determined by embedding clustering (KD-tree) on the heterogeneous training data. This allows the system to retrieve appropriate calibration parameters even for unseen task types based on semantic similarity.

Dynamic Calibration (S'). To handle distribution shifts (e.g., specific scientific workflows not present in the static corpus), we introduce an online dynamic calibrator S' . It refines the static estimate

$\hat{p}_0$ via a time-variant transformation:

$$\hat{p}_{\text{final}} = S'(\hat{p}_0; \theta_t) = \sigma(\alpha_t \cdot \text{logit}(\hat{p}_0) + \beta_t). \quad (2)$$

When correctness feedback is available, the parameters $\theta_t = \{\alpha_t, \beta_t\}$ are updated online via gradient descent to minimize the negative log-likelihood of recent auction outcomes.

3.3 Auction and Execution

Bid Construction. For each task unit t_i and candidate agent a_j , we compute a scalar bid b_{ij} based on two components. First, we apply a concave power-law transformation to the calibrated confidence $\hat{p}_{ij}$:

$$v(\hat{p}_{ij}) = (\hat{p}_{ij})^\gamma, \quad \text{with } \gamma \in (0, 1]. \quad (3)$$

This transformation compresses the high-confidence regime, making the mechanism less sensitive to minor fluctuations in high-probability estimates (e.g., distinguishing 0.99 from 0.95 becomes less critical than 0.60 from 0.50). Second, we calculate a *Normalized Cost* $C_{\text{norm},j} \in [0, 1]$, defined as a weighted sum of monetary expense and latency:

$$C_{\text{norm},j} = w_p \cdot \min\left(\frac{c_j^{\text{in}} + c_j^{\text{out}}}{C_{\text{ref}}}, 1\right) + w_l \cdot \min\left(\frac{L_{\text{ref}}}{l_j}, 1\right). \quad (4)$$

Here, c_j^{in} and c_j^{out} represent the agent’s unit execution costs (e.g., token prices for LLMs), while l_j denotes its throughput (tokens/sec). These are normalized against a domain-specific reference price ceiling (C_{ref}) and reference throughput (L_{ref}). The weights w_p and w_l govern the trade-off between budget and speed.

Winner Selection. The final bid is derived by penalizing the transformed confidence with the weighted cost:

$$b_{ij} = (\hat{p}_{ij})^\gamma - \beta \cdot C_{\text{norm},j}. \quad (5)$$

The parameter $\beta \geq 0$ acts as the global *Cost Sensitivity*. A higher β biases the system toward cost-efficiency, selecting cheaper or faster agents unless a premium model offers a substantial confidence gain. The winning agent is identified via $j^* = \arg \max_j b_{ij}$.

Auction Overhead. The auction adds one confidence query per eligible task-unit-agent pair, for $\sum_{t \in T} |A_t|$ bid calls, where A_t is the eligible agent set for unit t ; this reduces to $|T| \times |A|$ when all agents are eligible. These independent queries request scalar confidences and can be parallelized. For larger pools, a preliminary top- k filter could reduce bid traffic.

3.4 Refinement Loop

To handle distribution shifts in long-horizon deployment, the final stage *optionally* closes the loop by updating the dynamic calibrator S' . This module is designed to be adaptive: it activates when a feedback mechanism (e.g., unit tests for code or a lightweight verifier) provides a binary correctness label $y_{\text{label}} \in \{0, 1\}$. We use ground-truth feedback only for upper-bound analysis; without reliable feedback, Agora defaults to the static calibrator S . This label is used to update the parameters θ_t of S' via gradient descent:

$$\theta_{t+1} \leftarrow \theta_t - \eta \nabla_{\theta} \mathcal{L}_{\text{BCE}}(S'(\hat{p}_0; \theta_t), y_{\text{label}}). \quad (6)$$

This update lets S' adapt when correctness feedback is available.

4 Experiments

4.1 Datasets

AIME 2025: We utilize the complete set of 30 problems from the 2025 American Invitational Mathematics Examination to test advanced reasoning, evaluated via Pass@1. Each answer changes accuracy by 3.33 percentage points, so we treat AIME as a small-sample stress test.

MMLU-Pro: To evaluate robustness, we use MMLU-Pro (Wang et al., 2024b). Due to computational constraints, we evaluate on a stratified random subset of 2,000 examples, reporting the overall accuracy.

SciCode: This benchmark simulates scientific workflows (Tian et al., 2024). We report Pass Rates on the full test split. Following official recommendations, we utilize the `with_background` setting to strictly test knowledge retrieval.

SPIQA: For multimodal context, we use the Test-A split of SPIQA (Pramanick et al., 2024), reporting Retrieval Accuracy and L3 Score.

MathVision: We assess visual mathematical reasoning using the *testmini* split of MathVision (Wang et al., 2024a), reporting Pass@1.

Table 1: Overview of baseline methods and their compatibility with text-based versus multimodal inference tasks (✓: Supported; ×: Not supported/evaluated).

Type	Method	Text	Vision
Static	Single Strong / Weak	✓	✓
	Random Router	✓	✓
Heuristic	INN Router (Text Embed)	✓	×
Cascade	Consistency Cascade	✓	✓
	FrugalGPT	✓	✓
Learned	AdaptiveLLM	✓	×
	Hybrid LLM (Best-Route)	✓	×

4.2 Baselines

We compare our proposed framework against a diverse set of baselines, ranging from static single-model executions to advanced routing and cascading strategies. Since not all baselines support multimodal inputs, we categorize them based on their compatibility with text-based and visual inference tasks (see Table 1).

Static and Heuristic Baselines: To establish performance bounds, we evaluate *Single Strong/Weak* executions, which exclusively utilize the most capable or cost-effective model for all queries. We also include a *Random Router* as a stochastic lower bound. For retrieval-based routing, we employ a *INN Router* that selects agents based on text embedding similarity for language tasks.

Cascading Strategies: To address cost-efficiency, we evaluate sequential frameworks. *Consistency Cascade* (Wang et al., 2023) starts with a smaller model and escalates to a stronger one only if the consistency of sampled outputs falls below a threshold. *FrugalGPT* (Chen et al., 2024) queries a chain of models ordered by cost, stopping early if the answer is deemed reliable; we adapt this for visual tasks by chaining vision-language models.

Learned Routers: Finally, we compare against representative predictors trained to optimize routing. Our adapted implementation of *AdaptiveLLM* (Cheng et al., 2025) uses MiniLM query embeddings with a lightweight benchmark-specific classifier. *Hybrid LLM* (Ding et al., 2024) (Best-Route) explicitly models query difficulty and agent expertise to maximize expected accuracy under specific budget constraints.

4.3 Implementation

Model Selection. To test the framework’s adaptability, we employed distinct model pairs

for each benchmark. Within each benchmark, all routing and cascade methods use the same candidate pool:

- (1) *AIME 2025*: xai/grok-4-1 (Strong) paired with gpt-oss-20b (Weak);
- (2) *MMLU-Pro*: Mistral-Small-3.2-24B (Strong) paired with qwen3-14b (Weak);
- (3) *SciCode*: xai/grok-4-1 (Strong) paired with openai/gpt-5-mini (Weak).
- (4) For multimodal tasks (*SPIQA*, *MathVision*), we consistently used Qwen/Qwen3-VL-Thinking and xai/grok-4-1-fast-vision.

System Configuration. To demonstrate the controllability of Agora, we define three operating modes based on the cost sensitivity β : (1) *Quality-First* ($\beta = 0.001$): Prioritizes accuracy, using cost only as a tie-breaker. This is our default setting for matched-pool accuracy comparisons. (2) *Balanced* ($\beta = 0.1$): Seeks a trade-off between performance and budget. (3) *Cost-Efficient* ($\beta = 0.5$): Aggressively optimizes for savings, comparable to frugal baselines. All operating modes use the same static calibrators; changing β only changes the cost penalty and requires no retraining. In our main results (Table 2), we report the *Quality-First* performance to compare against strong baselines, while the trade-off characteristics are analyzed in Sec. 5.4.

Planning Configuration. While our framework generally employs an LLM Planner, we adapted this for specific benchmarks to leverage intrinsic structures. For *AIME 2025* and *MMLU-Pro*, the planner decomposes queries into explicit reasoning steps. However, for *SciCode* and *SPIQA*, we bypassed the planner to utilize their pre-defined sub-problem and retrieval-reasoning structures, respectively. This isolation allows us to evaluate the auction mechanism’s efficacy purely on task allocation, decoupling it from planning quality.

Baseline Implementation Details. To ensure a rigorous comparison, we categorize baselines into *training-free* strategies and *learned* routers. For *Learned Baselines* (AdaptiveLLM, Hybrid LLM, INN), we strictly isolated training data to avoid leakage: (1) For *AIME 2025*, we used historical AIME problems from 1983–2024; (2) For *MMLU-Pro*, we utilized the remaining examples excluding our stratified test subset; (3) For *SciCode*, we used the official validation split. For *Cascading Strategies*, we implemented inference-time variants with-

Table 2: *Text-based reasoning results.* *Vanilla (Strong/Weak)* denotes single-model execution with the candidates in Sec. 4.3. AIME Vanilla scores are from Artificial Analysis’s independent AIME 2025 evaluation (Artificial Analysis, 2025). *Agora (S/W)* uses the strong/weak model as planner. Results use the *quality-first setting* ($\beta = 0.001$); Sec. 5.4 reports cost trade-offs.

Methods	AIME	MMLU	SciCode	
	2025	Pro	Sub	Main
Vanilla (Strong)	89.5	68.1	44.2	13.6
Vanilla (Weak)	71.7	64.4	39.2	12.4
Random Router	76.7	66.3	42.3	12.6
Consist. Cascade	89.5	67.0	45.4	14.1
1NN Router	76.7	68.3	43.0	12.3
AdaptiveLLM	86.7	<u>69.3</u>	43.2	13.6
Hybrid LLM	80.0	64.5	43.2	11.2
FrugalGPT	86.7	66.7	<u>44.4</u>	13.6
Agora (S)	93.3	71.9	<u>44.4</u>	<u>13.8</u>
Agora (W)	<u>90.0</u>	69.1		

Table 3: *Multimodal Results.* Comparison on SPIQA and MathVision (MV). For SPIQA, *Ret* denotes Retrieval Accuracy; Avg, ≥ 0.6 , and ≥ 0.8 refer to the GPT-4o-evaluated L3 Reasoning Scores.

Methods	SPIQA				Math Vision
	Ret	Avg	≥ 0.6	≥ 0.8	
Vanilla (Grok)	85.4	60.5	<u>65.6</u>	43.4	50.3
Vanilla (Qwen)	67.3	55.2	59.2	<u>48.2</u>	53.3
Random Router	72.0	57.6	62.2	44.4	50.0
Consist. Cascade	72.6	55.5	60.0	45.6	51.4
FrugalGPT	75.2	<u>61.7</u>	64.4	45.1	<u>54.6</u>
Agora	<u>84.9</u>	65.0	72.4	56.9	55.3

out parameter updates: *Consistency Cascade* generates $k = 3$ samples and escalates if agreement is low, while *FrugalGPT* employs an LLM-based verifier (confidence threshold ≥ 0.7) to substitute the original trained scoring function. Detailed hyperparameters (e.g., consistency thresholds) are provided in Appendix E.

4.4 Results

We analyze performance across text-based and multimodal benchmarks (Tables 2 and 3).

Text-based Reasoning. Agora improves matched-pool performance on logic-intensive tasks. On *AIME 2025*, *Agora (S)* reaches 93.3%, versus 89.5% for the strong single-model baseline. On *MMLU-Pro*, it reaches 71.9%, the best result among the evaluated methods. On *SciCode*, Agora obtains a 44.4% subproblem pass rate, above




Figure 2: Ablation study on AIME 2025 and MMLU-Pro. The AIME series uses the strong model as planner (*Agora (S)*). Task planning and auction routing yield progressive gains over Vanilla.

the strong single-model baseline (44.2%), tied with FrugalGPT, and below Consistency Cascade (45.4%). This three-run mean suggests that distribution shift limits routing gains: the calibrator is trained on competitive programming rather than scientific code.

Multimodal Reasoning.

Table 3 reports results in visual domains. On *SPIQA*, Agora achieves the best Average L3 score among the evaluated methods (65.05%), exceeding FrugalGPT by over 3 percentage points. On the strict metric (≥ 0.8), Agora achieves 56.9%, an 8.7-point improvement over the best single model (Qwen: 48.2%). This result reflects functional complementarity: Grok is stronger at retrieval, whereas Qwen is stronger at strict reasoning. Together, the results delimit when routing helps: gains are largest when candidate strengths are separable, whereas domain-misaligned calibration yields competitive rather than dominant performance. For Agora, *SciCode* and *SPIQA* use predefined units, so their outcomes isolate allocation from planner quality.

5 Discussions

5.1 Ablation Study

As shown in Figure 2, both decomposition and dynamic routing drive consistent gains. Focusing on *MMLU-Pro*, which tests robust reasoning across diverse subjects, we observe a step-wise enhancement. *Task Planning* first lifts the single-model baseline by +1.90% (from 68.10% to 70.00%), demonstrating that structural decomposition reduces problem complexity even for strong generalist models. Building on this, the *Auction Mechanism* contributes an equal, additional +1.90% boost,

reaching a peak accuracy of 71.90%. This cumulative improvement shows that planning provides the reasoning blueprint, while the auction assigns each sub-task to the highest-scoring agent.

Impact of System Architecture. Figure 2 reports the *Agora (S)* AIME chain: 89.50% $\rightarrow$ 90.00% $\rightarrow$ 93.30%. For *Agora (W)*, task planning raises AIME accuracy from 71.69% to 83.30% (11.61 percentage points), and auction routing further raises it to 90.00% (6.70 points). The two AIME chains separate these effects: decomposition contributes more with the weak planner, whereas auction routing improves both planner settings. Given AIME’s 30-item size, we treat the exact increments as supporting rather than high-precision component estimates.

Impact of Online Refinement (S'). The +1.2-point comparison (70.7% $\rightarrow$ 71.9%) uses ground-truth correctness feedback and therefore estimates the potential value of online adaptation, not the default black-box setting. S' is enabled only when reliable labels are available, such as unit-test or tool-based verification; otherwise *Agora* retains static S .

5.2 Calibration Efficacy Analysis

A central premise of our auction mechanism is that an agent’s bid must accurately reflect its probability of success. We evaluate this from two perspectives: intrinsic reliability (ECE metrics) and extrinsic impact (downstream accuracy).

5.2.1 Intrinsic Reliability

Figure 3 visualizes the reliability diagrams for six representative agents. The curves consistently deviate below the diagonal ($y = x$), indicating that when models predict high confidence (e.g., > 0.8), their empirical accuracy is substantially lower. For instance, *grok-4-1-fast-reasoning* and *qwen3-14b* show high Expected Calibration Errors (ECE) of 0.222 and 0.357, respectively, rendering their raw logits unreliable for bidding.

Applying our hierarchical calibrator (green curves) improves these probabilities. ECE decreases across all six models, with an order-of-magnitude reduction for *grok-4-1* (0.222 $\rightarrow$ 0.023). Crucially, this efficacy extends to challenging multimodal domains: the calibrator rectifies the vision-language model *Qwen3-VL* (ECE 0.182 $\rightarrow$ 0.100), enabling the auction to fairly compare visual and textual reasoning bids.




Figure 3: Reliability diagrams comparing raw (red) versus calibrated (green) confidence scores across six representative models. The grey dashed line represents perfect calibration. Metrics shown are Expected Calibration Error (ECE) and Maximum Calibration Error (MCE) before and after calibration.

Table 4: Ablation study of calibration impact on downstream accuracy (%). We compare the Vanilla strong baseline against our Auction mechanism with and without calibration. Values in parentheses denote the absolute performance gain/loss (Δ) relative to the Vanilla baseline. *Note:* For SPIQA, accuracy is reported on the strict subset where L3 Score ≥ 0.8 .

Configuration	AIME 2025	MMLU Pro	SciCode Sub	SPIQA (L3 $\geq$ 0.8)	Math Vision
Vanilla	89.5	68.1	44.2	48.2	53.3
Auction ((No Calibration))	90.0 (+0.5)	67.5 (-0.6)	42.3 (-1.9)	46.9 (-1.3)	49.3 (-4.0)
Auction (With Calibration)	93.3 (+3.8)	71.9 (+3.8)	44.4 (+0.2)	56.9 (+8.7)	55.3 (+2.0)

5.2.2 Extrinsic Impact: Preventing the Winner’s Curse

Calibration is central to reliable bidding. Without it, overconfident agents can win the auction. Table 4 measures this effect.

While text-based tasks show minor volatility without calibration, *multimodal tasks suffer a performance collapse*. For instance, on MathVision, the uncalibrated auction drops 4.0% below the single-model baseline, as the system mistakenly prioritizes agents hallucinating high confidence on incorrect answers. Incorporating calibration reverses this trend, yielding $\Delta + 2.0\%$ on MathVision and $+8.7\%$ on SPIQA. Appendix D shows the same pattern on MuSiQue. F1 rises from 44.1 for Single Mistral to 51.4 without calibration and 54.3 with calibration. Across the main suite, the uncalibrated




Figure 4: *Mechanism Analysis on SPIQA*. (a) Functional complementarity: Grok excels in retrieval, Qwen in reasoning. (b) Dynamic orchestration: The system strategically shifts logic-intensive sub-tasks to Qwen.

auction falls below the single-model baseline on MMLU-Pro, SciCode, SPIQA, and MathVision. Auction structure alone is therefore insufficient; calibrated bids are needed for reliable winner selection. MuSiQue uses dataset-provided decomposition, so its calibration gain is independent of planner quality.

5.3 Mechanism of Complementarity

We attribute Agora’s gains on SPIQA to the framework’s exploitation of *functional orthogonality*. As visualized in Figure 4(a), the candidate models possess complementary capability frontiers: grok-4-1 acts as a “Retrieval Specialist” (85.4% retrieval accuracy vs. 67.3% for Qwen), whereas Qwen3-VL-Thinking serves as a “Reasoning Specialist” (48.2% strict reasoning rate vs. 43.4% for Grok).

Consequently, the auction mechanism orchestrates these distinct strengths (Figure 4b). In the *Retrieval phase*, the system routes 81.1% of sub-tasks to Grok, shielding the pipeline from Qwen’s weaker visual grounding. Conversely, in the *Reasoning phase*, allocation shifts dynamically: while Grok handles easier instances, Qwen is activated for the 32% of harder tasks. On these winning bids, Qwen scores 64.2, compared with 53.2 for Grok. This shows that Agora combines complementary strengths rather than defaulting to one model. Because SPIQA supplies its retrieval–reasoning structure, the stage-specific shift reflects conditional allocation rather than planner quality or a global backend preference.

5.4 The Cost-Quality Trade-off Analysis

To quantify the controllability of our auction mechanism, we conducted a sensitivity analysis on the cost sensitivity parameter β using the MathVision benchmark. As shown in Figure 5, adjusting β acts




Figure 5: Impact of cost sensitivity β on MathVision. Increasing β shifts allocation toward the cost-efficient agent (Grok), reducing normalized cost at the expense of marginal accuracy drops.

as a direct lever for governing the trade-off between economic efficiency and reasoning accuracy.

Increasing β from 0.001 to 0.5 triggers a strategic shift in allocation: the system progressively favors the cost-efficient agent (grok-4-1), with its usage ratio surging from 21.5% to 78.5%. This aggressive offloading demonstrates the mechanism’s responsiveness to cost constraints. However, this efficiency comes at a price: accuracy decreases from 55.26% to 51.06%. Notably, the transition from $\beta = 0.001$ to $\beta = 0.1$ reveals a balanced “sweet spot,” where Grok usage nearly doubles (to 36.5%) with only a negligible accuracy drop ($\approx 0.9\%$). In contrast, higher β values (≥ 0.25) prioritize budget over performance, pushing accuracy toward the lower bound of the standalone weak model. Because all β settings reuse the same static calibrators, the curve isolates allocation changes induced by the cost penalty rather than gains from retuning.

Bid collection is distinct from executor cost: scalar bid traffic depends on eligible task-unit–agent pairs, whereas β controls executor selection. Parallel queries and preliminary candidate filtering could limit traffic for larger pools.

6 Conclusion

We introduced *Agora*, a framework for enhancing LLM reasoning through confidence-calibrated, auction-based task allocation. By reformulating task allocation as a dynamic marketplace, Agora treats reasoning steps as tradeable items, enabling agents to bid based on verified competence rather than raw confidence. Our hierarchical calibration reduces the effect of overconfident bids and routes tasks using calibrated competence. Across five main benchmarks, Agora improves or remains competitive with matched-pool routing baselines, especially when candidate agents have complementary strengths. It also exposes a controllable cost–quality trade-off without retraining.

Limitations

Our framework has three limitations.

Calibration Generalization. Our mechanism relies on calibration quality. Developing a universal calibrator is data-hungry, and decomposed sub-tasks often present out-of-distribution challenges. Under distribution shift, the system can fall back to the strongest agent or require a larger bid margin before rerouting.

Dependency on Planning. Decomposition strategies are context-dependent and lack a “one-size-fits-all” solution. Tightly coupled tasks may violate the independent auction assumption and can instead be merged into a single task unit. Decomposition can also increase sub-task semantic variance, destabilizing the post-hoc calibrator.

Model Composition. Efficiency hinges on agent complementarity. If one model dominates all dimensions, candidate preselection or single-model execution avoids unnecessary bidding.

References

- Artificial Analysis. 2025. [LLM leaderboard](#).
- Lingjiao Chen, Matei Zaharia, and James Zou. 2024. [FrugalGPT: How to use large language models while reducing cost and improving performance](#). *Transactions on Machine Learning Research*.
- Junhang Cheng, Fang Liu, Chengru Wu, and Li Zhang. 2025. [AdaptiveLLM: A framework for selecting optimal cost-efficient LLM for code-generation based on CoT length](#). In *Proceedings of the 16th International Conference on Internetware*, Internetwork 2025, pages 461–473. ACM.
- Yu-Neng Chuang, Prathusha Kameswara Sarma, Parikshit Gopalan, John Boccio, Sara Bolouki, Xia Hu, and Helen Zhou. 2025. [Learning to route LLMs with confidence tokens](#). In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 10859–10878. PMLR.
- Damai Dai, Chengqi Deng, Chenggang Zhao, R. X. Xu, Huazuo Gao, Deli Chen, Jiashi Li, Wangding Zeng, Xingkai Yu, Y. Wu, Zhenda Xie, Y. K. Li, Panpan Huang, Fuli Luo, Chong Ruan, Zhifang Sui, and Wenfeng Liang. 2024. [DeepSeekMoE: Towards ultimate expert specialization in mixture-of-experts language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1280–1297. Association for Computational Linguistics.
- Dujian Ding, Ankur Mallick, Chi Wang, Robert Sim, Subhabrata Mukherjee, Victor Rühle, Laks V. S. Lakshmanan, and Ahmed Hassan Awadallah. 2024. [Hybrid LLM: Cost-efficient and quality-aware query routing](#). In *The Twelfth International Conference on Learning Representations*.
- Yilun Du, Shuang Li, Antonio Torralba, Joshua B. Tenenbaum, and Igor Mordatch. 2024. [Improving factuality and reasoning in language models through multiagent debate](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 11733–11763. PMLR.
- Paul Dütting, Vahab Mirrokni, Renato Paes Leme, Haifeng Xu, and Song Zuo. 2024. [Mechanism design for large language models](#). In *Proceedings of the ACM Web Conference 2024*, pages 144–155. ACM.
- Caoyun Fan, Jindou Chen, Yaohui Jin, and Hao He. 2024. [Can large language models serve as rational players in game theory? a systematic analysis](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(16):17960–17967.
- William Fedus, Barret Zoph, and Noam Shazeer. 2022. [Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity](#). *Journal of Machine Learning Research*, 23(120):1–39.
- Yue Feng, Yang Wang, and Hang Li. 2021. [A sequence-to-sequence approach to dialogue state tracking](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1714–1725, Online. Association for Computational Linguistics.
- Guangfu Guo, Xiaoqian Lu, and Yue Feng. 2025. [MedVRagent: A framework for medical visual reasoning-enhanced agents](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 18602–18616, Suzhou, China. Association for Computational Linguistics.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. 2025. [A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions](#). *ACM Transactions on Information Systems*, 43(2):1–55.
- Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, Guillaume Lample, L  lio Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, and 7 others. 2024. [Mixtral of experts](#). *Preprint*, arXiv:2401.04088.

- Pan Lu, Baolin Peng, Hao Cheng, Michel Galley, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, and Jianfeng Gao. 2023. [Chameleon: Plug-and-play compositional reasoning with large language models](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 43447–43478.
- Putra Manggala, Atalanti A. Mastakouri, Elke Kirschbaum, Shiva Prasad Kasiviswanathan, and Aaditya Ramdas. 2025. [QA-Calibration of language model confidence scores](#). In *International Conference on Learning Representations*, pages 82631–82654.
- Dang Nguyen, Ali Payani, and Baharan Mirzasoleiman. 2025. [Beyond semantic entropy: Boosting LLM uncertainty quantification with pairwise semantic similarity](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 4530–4540. Association for Computational Linguistics.
- Quang H. Nguyen, Thanh Dao, Duy C. Hoang, Juliette Decugis, Saurav Manchanda, Nitesh V. Chawla, and Khoa D. Doan. 2024. [MetaLLM: A high-performant and cost-efficient dynamic framework for wrapping LLMs](#). *Preprint*, arXiv:2407.10834.
- Alexander Nikitin, Jannik Kossen, Yarin Gal, and Pekka Marttinen. 2024. [Kernel language entropy: Fine-grained uncertainty quantification for LLMs from semantic similarities](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 8901–8929.
- Isaac Ong, Amjad Almahairi, Vincent Wu, Wei-Lin Chiang, Tianhao Wu, Joseph E. Gonzalez, M. Waleed Kadous, and Ion Stoica. 2025. [RouteLLM: Learning to route LLMs from preference data](#). In *International Conference on Learning Representations*, pages 34433–34448.
- Shraman Pramanick, Rama Chellappa, and Subhashini Venugopalan. 2024. [SPIQA: A dataset for multi-modal question answering on scientific papers](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 118807–118833.
- David Raposo, Sam Ritter, Blake Richards, Timothy Lillcrap, Peter Conway Humphreys, and Adam Santoro. 2024. [Mixture-of-Depths: Dynamically allocating compute in transformer-based language models](#). *Preprint*, arXiv:2404.02258.
- Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. 2023. [Hugging-GPT: Solving AI tasks with ChatGPT and its friends in hugging face](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 38154–38180.
- Xinyuan Song, Zeyu Wang, Siyi Wu, Tianyu Shi, and Lynn Ai. 2025. [Gradientsys: A multi-agent LLM scheduler with ReAct orchestration](#). *Preprint*, arXiv:2507.06520.
- Dimitris Stripelis, Zhaozhuo Xu, Zijian Hu, Alay Dilipbhai Shah, Han Jin, Yuhang Yao, Jipeng Zhang, Tong Zhang, Salman Avestimehr, and Chaoyang He. 2024. [TensorOpera router: A multi-model router for efficient LLM inference](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 452–462. Association for Computational Linguistics.
- Katherine Tian, Eric Mitchell, Allan Zhou, Archit Sharma, Rafael Rafailov, Huaxiu Yao, Chelsea Finn, and Christopher D. 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. Association for Computational Linguistics.
- Minyang Tian, Luyu Gao, Shizhuo Dylan Zhang, Xinnan Chen, Cunwei Fan, Xuefei Guo, Roland Haas, Pan Ji, Kittithat Krongchon, Yao Li, Shengyan Liu, Di Luo, Yutao Ma, Hao Tong, Kha Trinh, Chenyu Tian, Zihan Wang, Bohao Wu, Yanyu Xiong, and 11 others. 2024. [SciCode: A research coding benchmark curated by scientists](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 30624–30650.
- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2022. [MuSiQue: Multi-hop questions via single-hop question composition](#). *Transactions of the Association for Computational Linguistics*, 10:539–554.
- Ke Wang, Juntao Pan, Weikang Shi, Zimu Lu, Houxing Ren, Aojun Zhou, Mingjie Zhan, and Hongsheng Li. 2024a. [Measuring multimodal mathematical reasoning with MATH-Vision dataset](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 95095–95169.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V. Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. [Self-consistency improves chain of thought reasoning in language models](#). In *The Eleventh International Conference on Learning Representations*.
- Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, Tianle Li, Max Ku, Kai Wang, Alex Zhuang, Rongqi Fan, Xiang Yue, and Wenhui Chen. 2024b. [MMLU-Pro: A more robust and challenging multi-task language understanding benchmark](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 95266–95290.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. [Chain-of-Thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837.
- 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*.

Daniel Yang, Yao-Hung Hubert Tsai, and Makoto Yamada. 2024. [On verbalized confidence scores for LLMs](#). *Preprint*, arXiv:2412.14737.

Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc V. Le, and Ed H. Chi. 2023. [Least-to-Most prompting enables complex reasoning in large language models](#). In *The Eleventh International Conference on Learning Representations*.

A Prompt Templates

To ensure reproducibility, we provide the exact system prompts used in the Agora pipeline. These templates are stored as plain text within the codebase.

A.1 Planner (Plan Graph Generation)

The planner is implemented in `aucteam/planner_agent/`. The prompt is assembled by concatenating the `SYSTEM_PRIMER`, few-shot exemplars, and `OUTPUT_INSTRUCTIONS`.

Rationale for Plan Constraints. We explicitly instruct the planner to limit output to a maximum of 4 steps. Empirical observation suggests that while longer plans offer finer granularity, they significantly increase the probability of *cascading failures*. The 4-step limit serves as a heuristic regularization.

System Primer.

You are an expert research planner.
Break complex requests into modular steps with explicit data/control dependencies.
Represent every plan as a single JSON object that conforms to this schema:

- Root object:
 - * steps (list): ordered execution steps.
 - * edges (list): dependencies between IDs.
 - * notes (str): clarifications.
- Step object template:


```
{
  "id": "T1", "description": "Outcome",
  "skills": ["capability"],
  "reads": ["art_id"],
  "writes": ["art_id"]}
      
```

 - * id: unique string starting with 'T'.
 - * description: what to accomplish.
 - * skills/reads/writes: always lists.
- Edge object template:


```
{
  "from": "T1", "to": "T2",
  "type": "data", "why": "Rationale"}
      
```

 - * type: "data" or "control".

Ensure the JSON is syntactically valid.

Output Instructions.

Respond with exactly one JSON object, followed by `<END_JSON>` on its own line. Do not include Markdown fences.

CRITICAL - Plan Requirements:

1. MAXIMUM 4 STEPS.
2. Keep step descriptions HIGH-LEVEL.
3. Do NOT embed:
 - Specific numerical calculations
 - Concrete formulas or equations
 - Intermediate computation results
4. Describe the GOAL, not the METHOD.

A.2 Executor (Task-Unit Execution)

Each task unit is executed with a structured prompt that isolates the unit's context.

Task-Unit Prompt Template.

Overall task: {task_description}

Executing unit {unit_id}

Steps in this unit:
- {node_id}: {step_description}
- ...

Skills: {skills}

=====

INPUTS FROM PRIOR STEPS:

=====

[{artifact_name}]:
{artifact_content}
truncated to 8000 chars per artifact

=====

Instructions:

- If final unit:
 - This is the FINAL step.
 - {final_answer_hint_1}
- Else:
 - Focus only on this unit.
 - Provide intermediate findings.
 - Do not give the final answer yet.
 - Use INPUTS FROM PRIOR STEPS.

Context: {optional_context}

A.3 Self-Reported Confidence Prompt

Before the auction, agents self-report a scalar confidence p_{raw} .

Rate your confidence (0-1) that you can correctly implement this function:

{step_description_prompt}
truncated to 800 chars

Function signature:
{function_header}

Respond with ONLY a number between 0 and 1 (e.g., 0.85).

B Calibration Implementation Details

Our calibration framework is designed to ensure robustness across diverse task types. We explicitly distinguish the training data sources for the static and dynamic components to balance generalization with adaptability.

Static Calibrator (S): Broad-Spectrum Training. To enable the static calibrator to adapt to unseen task types (e.g., domain-specific scientific problems) and minimize the risk of distribution

shift, we trained it on a *comprehensive, large-scale corpus* aggregated from diverse public benchmarks. Instead of relying on simple binning which fails to capture the complexity of broad-domain data, we utilized the ground-truth labels from these datasets to train the Hierarchical Scaling (HS-QAB) parameters. Crucially, we prioritized using validation or test splits from source benchmarks where possible to minimize overlap with the pre-training corpora of large language models.

The training corpus is categorized by modality:

- *Text-based Reasoning Tasks*: We aggregated datasets covering mathematics, coding, and commonsense reasoning, including: *ACEReason (Math)*, *AGIEval*, *AIME (Training split)*, *ARC-Challenge/Easy (Train/Val/Test)*, *Codeforces*, *CSQA*, *GSM8K*, *HumanEval*, *MMLU (Subset)*, and *TACO*.
- *Multimodal Tasks*: We utilized visual reasoning datasets including *MathVerse (Test)*, *SPIQA (Validation)*, *ChatQA*, and *Geometry3K*.

This training is intended to provide S with a broad baseline estimate, including for tasks outside the calibration corpus.

Dynamic Calibrator (S'): Historical Adaptation.

While the static calibrator targets broad generalization, the dynamic calibrator S' is trained on the *historical interaction data* accumulated during the current inference session. This allows the system to correct residual errors and adapt to the specific distribution of the target task in real-time.

C Algorithm Details

Task-Unit Extraction Algorithm. Pseudocode for the heuristic rules used to merge tightly coupled nodes.

Algorithm: Task-Unit Extraction

Input: Task Graph $G = (V, E)$
Params: Bandwidth=2, Jaccard_Thresh=0.6

1. Initialize Candidate Groups $S = \{\}$
2. Rule A: Data Reuse Band
For each pair (u, v) in Bandwidth:
 # Check intersection
 If $\text{Intersect}(\text{Writes}(u), \text{Reads}(v))$:
 If path $u \rightarrow v$ is clean:
 Add $\{u, v\}$ to S
3. Rule B: Prompt Similarity
For each pair (u, v) :
 If $\text{Jaccard}(u.\text{txt}, v.\text{txt}) \geq \text{Thresh}$:
 Add $\{u, v\}$ to S
4. Rule C: Control Gates

For each control edge $u \rightarrow v$:
 Add $\{u, v\}$ to S

5. Merge & Validate
 Union overlapping groups in S .
 For each Group g in S :
 If $\text{Valid}(g)$ AND $\text{Capable}(g)$:
 If $\text{Savings} > \text{Threshold}$:
 Merge nodes in g

D Additional MuSiQue Evaluation

We additionally evaluate Agora on MuSiQue-Ans (Trivedi et al., 2022) using the dataset-provided question decompositions as task units. We use 200 training examples for a two-bin calibrator and a disjoint, fixed 500-example development subset for evaluation. The candidate pool contains Mistral-Small-3.2-24B and Qwen3-14B; outputs are generated greedily with seed 9552 and replayed across routing methods. We report Exact Match (EM) and token-level F1. Agora obtains 43.0 EM and 54.3 F1; the F1 95% bootstrap confidence interval is [50.6, 58.2] over 1,000 resamples.

Table 5: MuSiQue-Ans results on the fixed 500-example development subset. Agora uses the quality-first setting ($\beta = 0.001$).

Method	EM	F1
Single Mistral	33.6	44.1
Single Qwen	25.2	35.5
Random Router	39.2	50.4
1NN Router	<u>42.4</u>	<u>54.1</u>
Auction w/o Calibration	39.6	51.4
Agora	43.0	54.3

E Baseline Implementation Details

E.1 Training Data Construction

For baselines requiring supervision (1NN Router, AdaptiveLLM, Hybrid LLM), we constructed training sets distinct from the evaluation splits:

- *AIME 2025*: We used 933 unique historical AIME problems from 1983–2024 for training.
- *MMLU-Pro*: The original dataset contains 12k+ examples. We reserved our stratified 2,000-sample test set and used the remaining ~ 10 k examples for training learned routers.
- *SciCode*: Following standard protocol, we utilized the provided validation split (featuring distinct problems from the test set) for training.

E.2 Configuration of Cascading Baselines

Consistency Cascade: This method operates without training. We configured it to first query the cost-efficient model. We generate $k = 3$ reasoning paths; if the self-consistency agreement (vote ratio) is below $\tau = 0.6$, the system escalates to the strong model.

FrugalGPT: We implement a training-free variant of FrugalGPT (Chen et al., 2024): a strong LLM judges the weak model’s answer, which is accepted at confidence ≥ 0.7 and otherwise escalated. This approximates the original learned scorer without a benchmark-specific classifier.