LEARNING THE COST OF RELIABLE INFERENCE

Anonymous authors
Paper under double-blind review

ABSTRACT

Benchmarking and routing platforms increasingly act as intermediaries connecting large language model providers with end-users. However, providers on these platforms typically use a fixed price per token, preventing users from achieving the most competitive price for their tasks. In this work, we design a procurement platform where token prices for each task are driven by provider competition, enabling users to secure competitive pricing for guaranteed quality levels. To this end, the platform sequentially routes queries via a reverse second-price auction that incentivizes model providers to truthfully bid their best estimate of the average cost to serve a user's query. As it routes queries, the platform learns the quality offered by each provider and progressively routes queries to the most cost-competitive provider among those meeting a desired quality threshold. To validate our design, we conduct experiments with multiple LLMs from the `Llama` and `Qwen` families on popular mathematical reasoning and question-answering benchmarks. The results show that the pricing margin of the most cost-competitive provider on our platform varies significantly—from 10% to 71%—depending on the task and quality threshold. This suggests a substantial inefficiency in the current fixed-price market, and it demonstrates that our platform may enable users to capture maximum savings whenever competitive market conditions permit.

1 INTRODUCTION

The rapid proliferation of large language models (LLMs) has led to a highly competitive market of online service providers offering models of varying capabilities at different prices. To help users navigate their choice of models, benchmarking and routing platforms have become popular, letting users compare model performance on well-established tasks and even automatically selecting a provider to serve their queries.[1]

However, existing platforms do not allow users to achieve the most competitive price for their tasks. This is because providers on these platforms typically charge users a fixed price per token,[2] which fails to reflect that a provider's cost to serve a model and the response quality the user obtains in return depend heavily on the specific task. As a concrete illustration, consider a user who wants to use an LLM to assist with medical diagnosis and can choose between two providers, each offering a different model. Suppose the providers charge $\$0.06$ and $\$0.07$ per token, while generating a token costs them $\$0.01$ and $\$0.04$, respectively. Then, if both models produce responses of similar quality and number of tokens for the user's task, the user—acting rationally—would select the cheapest model and pay $\$0.06$ per token. However, if the providers competed on price for the user's task, the user could in principle pay as little as $\$0.04$ per token: the first provider would lower its per-token price down to this point, since below this the second provider—which cannot profitably price below its own cost—would no longer be a competitive alternative (Milgrom & Roberts, 1982; Mahmood, 2024; Bergemann et al., 2025c).

In this work, we design a procurement platform that addresses this price inefficiency, enabling users to secure competitive pricing on their tasks for guaranteed quality levels. To this end, the platform sequentially routes queries via a reverse second-price auction that incentivizes providers to truthfully bid their best estimate of the average cost to serve a user's query. As it routes queries, the platform learns the quality offered by each provider and progressively routes queries to the most

[1] See, *e.g.*, `artificialanalysis.ai`, `openrouter.ai`, and `thegrid.ai`.

[2] `ai.google.dev/gemini-api/docs/pricing`, `openai.com/api/pricing/`, `claude.com/pricing`.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

cost-competitive provider among those meeting a desired quality threshold. More formally, we show that, with high probability, the platform routes all but a vanishing fraction of the user's queries to the most cost-competitive provider, whose (average) pricing margin converges to the difference between the second-lowest average cost among qualified providers and its own average cost.

To validate our platform, we conduct experiments with multiple LLMs from the `Llama` and `Qwen` families on popular mathematical reasoning and question-answering benchmarks. We find that the pricing margin of the most cost-competitive provider varies significantly—from 10% to 71%—depending on the task and quality threshold. This suggests a substantial inefficiency in the current fixed-price market, and it demonstrates that our platform may enable users to capture maximum savings whenever competitive market conditions permit.

Further related work. Our work contributes to a rapidly growing body of literature on the economic aspects of LLMs-as-a-service (Raghavan, 2024; Mahmood, 2024; La Malfa et al., 2024; Laufer et al., 2024; Bergemann et al., 2025b;c; Olmedo et al., 2026; Velasco et al., 2026b;c;a). Within this literature, the works most closely related to ours characterize how LLM providers price their models when serving users in a competitive market. Specifically, Mahmood (2024) study the price competition between providers serving models of varying quality across different tasks, while Bergemann et al. (2025c) characterize the menu of token prices that maximizes a provider's revenue across heterogeneous users. Our work also relates to a recent line that evaluates LLMs not only by their capabilities but also by the price users pay (Wang et al., 2024b; Zellinger & Thomson, 2025; Erol et al., 2026; Chen et al., 2026a). This alternative view has very recently materialized in already-deployed platforms: some platforms publish provider rankings that explicitly factor in their prices, while others act as routers that lower the price by letting providers compete on their rates.[1] In contrast to these solutions, our platform explicitly ensures that the user asymptotically pays the most competitive market price for their specific task at guaranteed quality levels.

Our work also builds on a large body of literature on mechanism and auction design (McAfee & McMillan, 1986; Ledyard, 1987; Che, 1993; Bar-Yossef et al., 2002; Roughgarden et al., 2017; Akbarpour & Li, 2020; Awaya et al., 2025; Cai & Liao, 2026). In this context, our sequential setting connects to a strand on dynamic auctions, where bidders sequentially learn private information, such as their type, that determines their utility, and the objective is to design mechanisms that ensure truthful bidding over multiple time steps (Hossain, 2008; Pavan et al., 2014; Weed et al., 2016; Feng et al., 2018; Bergemann & Välimäki, 2019; Bergemann et al., 2025a). To the best of our knowledge, however, we are the first to consider a setting where LLM providers act as bidders whose operational costs and serving qualities vary across tasks.

Lastly, at a technical level, our work relates to a strand of the multi-armed bandit literature in which the learner must choose among alternatives that are themselves strategic agents seeking to maximize their utility, requiring learning algorithms that explicitly account for these incentives (Auer et al., 2002; Even-Dar et al., 2006; Babaioff et al., 2009; Bubeck & Cesa-Bianchi, 2012; Jain et al., 2015; Braverman et al., 2019; Lattimore & Szepesvári, 2020; Freeman et al., 2020; Padmanabhan et al., 2022). Therein, the closest works to ours are those that study procurement auctions through a bandit lens (Singla & Krause, 2013; Chen et al., 2026b; Patra et al., 2026). Specifically, both Chen et al. (2026b) and Patra et al. (2026) formulate procurement as a contextual bandit problem, with the latter studying the selection of LLM providers through a reverse auction where providers report their generation costs and the user learns query-dependent model quality. Unlike Patra et al. (2026), however, we do not assume that providers know their model's cost on the user's task, and our payment rule asymptotically converges to the most competitive price, equal to the second-lowest average cost among qualified providers. In a similar LLM routing setting, Cao et al. (2026) instead design an allocation and payment mechanism that disincentivizes providers from misreporting their capabilities or manipulating the token bills charged to the user.

2 A Procurement Platform for Reliable Inference

In this section, we introduce a procurement platform that acts as an intermediary between a user who seeks to solve a specific task (*e.g.*, coding, mathematical reasoning, or question answering) and a set of N LLM providers. By incentivizing providers to compete on price, the platform ensures the user secures the most competitive pricing at a guaranteed quality level on all but a vanishing fraction of queries.

The platform operates sequentially over time steps $t = 1, 2, \ldots, T$, where the total number of queries T is unknown to the providers but assumed to be at least the number of providers N. At any given time step, each provider participates in the platform by maintaining a standing bid $b_{i,t} \in [0, c_{\max}]$, where $c_{\max} > 0$ is an upper bound on the cost of serving a single query. Given the current bid vector $\mathbf{b}_t = (b_{1,t}, \ldots, b_{N,t})$, the platform receives a query $x_t \in \mathcal{X}$ that the user wishes to solve, where $\mathcal{X}$ is the space of possible queries characterizing the user's task, and selects the provider $I_t \in [N]$ that serves it. Provider I_t generates a response $y_t \in \mathcal{Y}$ to the query, where $\mathcal{Y}$ is the space of possible responses,[3] and, by doing so, it incurs a generation cost $c_{I_t,t} \in [0, c_{\max}]$, which includes the cost of energy, compute, or hardware. The platform then grades the response using a score $r(x_t, y_t) \in [0, 1]$, ranging from unsatisfactory (0) to fully satisfactory (1), and sets the payment the user makes to provider I_t. Finally, provider I_t may update its standing bid to $b_{I_t,t+1}$, while the bid of every other provider remains unchanged: $b_{i,t+1} = b_{i,t}$ for all $i \neq I_t$. In what follows, we denote the distribution of the query as P^x, the distribution of the cost incurred by provider i as P_i^c, and the distribution of the score received by provider i's response as P_i^r, with a mean $q_i \in [0, 1]$ that we refer to as the (average) quality of the provider.

We now describe the rules used by the platform to allocate queries and pay the providers. Our key idea is to combine an allocation rule based on confidence bounds—which progressively narrows the search to the *qualified* provider offering the lowest price to the user—with a second-price payment rule that incentivizes providers to bid as low as possible without risking an unprofitable win. Specifically, the platform begins with an initialization phase spanning the first N time steps, during which it assigns one query to each provider in turn and pays them $c_{\max}$. This ensures every provider serves at least one query, rendering the estimates below well-defined from round $N + 1$ onward.

Thereafter, at each round $t > N$, the platform maintains a running estimate of each provider's quality,

$$\hat{q}_{i,t} = \frac{1}{m_{i,t}} \sum_{\tau=1}^{t-1} r(x_\tau, y_\tau)\, \mathbb{1}\{I_\tau = i\}, \tag{1}$$

where $m_{i,t} = \sum_{\tau=1}^{t-1} \mathbb{1}\{I_\tau = i\}$ counts how many times provider i has been selected before round t. The goal of these estimates is to identify the set of qualified providers $\mathcal{Q} = \{i \in [N] : q_i \geq q_{\min}\}$, meeting a user-defined quality threshold $q_{\min}$, which we assume contains at least two providers[4]. Since the true qualities q_i are unknown, the platform instead forms an optimistic estimate of $\mathcal{Q}$,

$$\hat{\mathcal{Q}}_t = \big\{i \in [N] : \hat{q}_{i,t} + \beta(m_{i,t}) \geq q_{\min}\big\}, \quad \text{where} \quad \beta(m_{i,t}) = \sqrt{\frac{1}{2m_{i,t}} \ln\left(\frac{2\pi^2 N\, m_{i,t}^2}{3\delta}\right)}. \tag{2}$$

Here, $\beta(m_{i,t})$ is a confidence radius at a level $\delta \in (0, 1)$ fixed by the platform that accounts for the error in the empirical estimate $\hat{q}_{i,t}$, so that $\hat{\mathcal{Q}}_t$ contains every provider whose upper confidence bound on quality exceeds the threshold $q_{\min}$. In Appendices B.1 and B.2, we show that, with probability at least $1 - \delta$, this optimistic estimate never excludes a qualified provider, *i.e.*, $\mathcal{Q} \subseteq \hat{\mathcal{Q}}_t$ holds simultaneously across all time steps $t > N$.

Using the set $\hat{\mathcal{Q}}_t$, the platform allocates the query by prioritizing providers with lower bids while giving an explicit advantage to those that have not yet been selected often. More concretely, let $g(t)$ be an increasing function specified by the platform that grows sublinearly in t, and $\mathcal{S}_t = \{\, i \in \hat{\mathcal{Q}}_t : m_{i,t} < g(t) \,\}$ denote the set of estimated-qualified providers selected fewer than $g(t)$ times so far. Then, at each round $t > N$, the platform selects:[5]

$$I_t = \begin{cases} \text{Uniform}\,(\mathcal{S}_t) & \text{if } \mathcal{S}_t \neq \emptyset, \\ \displaystyle\arg\min_{i \in \hat{\mathcal{Q}}_t} \big(b_{i,t} - c_{\max}\, \beta(m_{i,t})\big) & \text{otherwise,} \end{cases} \tag{3}$$

[3] In practice, the provider generates y_t by sampling a sequence of tokens from its model, and so $\mathcal{Y}$ is the set of finite-length token sequences.

[4] This assumption is needed only for the analysis. If a single provider is estimated to be qualified, the minimum in Eq. 4 is over an empty set and the provider is paid $c_{\max}$.

[5] If $\hat{\mathcal{Q}}_t = \emptyset$, then I_t is chosen uniformly at random. However, this occurs with probability at most δ (see Appendices B.1 and B.2).

with ties broken arbitrarily. That is, if $\mathcal{S}_t \neq \emptyset$, which we refer to as an exploration step, the platform selects at random an estimated-qualified provider that has not been sampled sufficiently according to the function $g(t)$. Otherwise, the rule selects the provider minimizing an optimistic bidding index, where the term $c_{\max}\,\beta(m_{i,t})$ gives an additional advantage to providers whose confidence radius remains large. Altogether, this allocation rule ensures that qualified providers are selected often enough for the platform to reliably estimate their quality and, as we will see in Section 3, for the providers to learn their own costs.

To design the payment rule, we build on the classical second-price rule from the auction and mechanism design literature (Vickrey, 1961; Nisan et al., 2011). Specifically, for any provider i, we define the critical payment $P_{i,t}$ as

$$P_{i,t} = c_{\max}\,\beta(m_{i,t}) + \min_{j \in \hat{\mathcal{Q}}_t \setminus \{i\}} \big(b_{j,t} - c_{\max}\,\beta(m_{j,t})\big), \tag{4}$$

where the minimum over an empty set is taken to be $+\infty$ by convention. The actual payment made to the selected provider I_t is then set to

$$\pi_t = \begin{cases} c_{\max} & \text{if } \mathcal{S}_t \neq \emptyset, \\ \min\{P_{I_t,t}, c_{\max}\} & \text{otherwise.} \end{cases} \tag{5}$$

In words, in an exploration step, the provider selected uniformly at random from the set $\mathcal{S}_t$ is paid $c_{\max}$. Otherwise, the selected provider I_t receives a payment that depends on how often the providers have been selected in the past and on the other providers' bids $\mathbf{b}_{-I_t,t} = (b_{j,t})_{j \neq I_t}$, but not on its own bid $b_{I_t,t}$. Algorithm 1 in Appendix A summarizes the complete procedure.

3 Pricing Inference Through Competition

In this section, we study the incentives that providers face when participating in our procurement platform and characterize their optimal bidding strategy. As a first step, we analyze the conditions under which a provider is selected to serve a query as a function of their bid. Concretely, the following proposition shows that, when $\mathcal{S}_t = \emptyset$, a provider's assignment is determined by whether their standing bid falls below the critical payment $P_{i,t}$.[6]

Proposition 1. *Fix a step $t > N$ with $\mathcal{S}_t = \emptyset$, a provider $i \in \hat{\mathcal{Q}}_t$, and the other providers' bids $\mathbf{b}_{-i,t}$. Then, provider i is assigned the query x_t if they bid $b_{i,t} < P_{i,t}$, and is not assigned the query if $b_{i,t} > P_{i,t}$.*

The above proposition does not fully characterize how a provider *should* bid; however, it already offers valuable intuition. When $\mathcal{S}_t \neq \emptyset$, the query is assigned independently of the bids $\mathbf{b}_t$, meaning a provider's bid only influences the allocation when $\mathcal{S}_t = \emptyset$. In this case, a provider is assigned the query only if it bids below the threshold $P_{i,t}$. Because this threshold depends on the other providers' bids, it is unknown to provider i in advance, which incentivizes it to bid as low as possible. Conversely, once a provider bids low enough to secure the query, due to the payment rule in Eq. 5, their final payout remains independent of their own bid. Together, these properties suggest that providers are incentivized to bid truthfully according to the generation costs they incur. In the remainder of this section, we formalize this intuition and characterize the resulting equilibrium.

We cast the interaction among the N providers on the platform as a multi-stage game (Fudenberg & Tirole, 1991). In this game, each provider i is characterized by its score and cost distributions $\theta_i = (P_i^r, P_i^c)$ on the user's task, which need not be known even to the provider itself; we write $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_N)$.[7] Moreover, its action in the game corresponds to choosing a bidding strategy σ_i that determines its standing bid $b_{i,t}$ based on its observed history $H_{i,t}$, which collects the costs $c_{i,\tau}$ and scores $r(x_\tau, y_\tau)$ of every query that provider i has served before time step t, *i.e.*,[8]

$$b_{i,t} = \sigma_i(H_{i,t}) \quad \text{where} \quad H_{i,t} = \big(c_{i,\tau}, r(x_\tau, y_\tau)\big)_{\tau < t\,:\, I_\tau = i}. \tag{6}$$

[6] All proofs are deferred to Appendix B.

[7] In the game theory literature, this corresponds to an incomplete-information setting in which θ_i is the (private) *type* of each provider (Harsanyi, 2004; Hossain, 2008; Feng et al., 2018).

[8] Our theoretical results also hold if each provider i is allowed to use a time-varying bidding strategy $\sigma_{i,t}$. This is because, as we will show later, our notion of bidding strategy optimality applies to any arbitrary strategy at any time t. However, to lighten the notation, we focus on time-invariant bidding strategies mapping histories to bids.

In what follows, we denote the providers' overall history as H_t and say that provider i's history $H_{i,t}$ is compatible with H_t if it is contained in H_t.

Crucially, although a provider does not necessarily know its true mean cost $\bar{c}_i$ of serving a user query, it can use its observed history to compute an estimate of the conditional mean cost of serving the next query, $\mathbb{E}\big[c_{i,t} \mid H_{i,t}\big]$. For ease of exposition, we focus on providers that compute this estimate as the empirical average of past costs, given by

$$\hat{c}_{i,t} = \frac{1}{m_{i,t}} \sum_{\tau=1}^{t-1} c_{i,\tau} \, \mathbb{1}\{I_\tau = i\}.$$

In Appendix C, we show that Theorem 2 holds for any cost estimator and, if the platform widens the cost radius $c_{\max}\beta(\cdot)$ to $(1+\gamma)c_{\max}\beta(\cdot)$, Theorems 3 and 4 hold for providers that compute $\hat{c}_{i,t}$ through Bayesian estimators whose estimates deviates from the empirical mean by at most a constant multiple γ of the confidence radius.

Then, at time step $t > N$, the utility that a provider i expects to receive from participating in the platform under any possible history H_t that is compatible with its observed history $H_{i,t}$ is the conditional average of payments received minus the mean cost it estimates to incur by serving queries over the remaining horizon from t to T, *i.e.*,

$$U_{i|H_t}(\sigma_i, \boldsymbol{\sigma}_{-i}; \boldsymbol{\theta}_{-i}, T) = \mathbb{E}\left[\sum_{s=t}^{T} (\pi_s - \hat{c}_{i,s}) \cdot \mathbb{1}\{I_s = i\} \,|\, H_t\right], \tag{7}$$

where $\boldsymbol{\sigma}_{-i} = (\sigma_j)_{j \neq i}$ denotes the strategies of all providers other than i, and the expectation is over the randomness in the providers' generations conditional on the history H_t, which depends implicitly on $\boldsymbol{\theta}$.

Having formalized the providers' interaction as a game, we now turn to characterizing their equilibrium bidding strategies. Specifically, since a provider i knows neither the score and cost distributions $\boldsymbol{\theta}_{-i}$ characterizing the other providers nor the total number of user queries T, we focus on strategies that maximize its utility under this uncertainty using the framework of robust game theory (Aghassi & Bertsimas, 2006). Formally, the next theorem shows that the strategy σ_i^* that bids the provider's cost estimate for the next query, *i.e.*,

$$\sigma_i^*(H_{i,t}) = \hat{c}_{i,t}, \tag{8}$$

offers every provider the best worst-case utility over its uncertainty about $\boldsymbol{\theta}_{-i}$ and T for any given history H_t that is compatible with its observed history $H_{i,t}$.

Theorem 2. *For any provider $i \in [N]$, any strategy profile $\boldsymbol{\sigma}_{-i}$ of all other providers, and any history H_t with $t > N$ that is compatible with provider i's observed history $H_{i,t}$, the strategy σ_i^* satisfies:*

i) for every $T \geq t$ and $\boldsymbol{\theta}_{-i}$, $U_{i|H_t}(\sigma_i^*, \boldsymbol{\sigma}_{-i}; \boldsymbol{\theta}_{-i}, T) \geq 0$;
ii) for every σ_i, $\inf_{T \geq t,\, \boldsymbol{\theta}_{-i}} U_{i|H_t}(\sigma_i^*, \boldsymbol{\sigma}_{-i}; \boldsymbol{\theta}_{-i}, T) \geq \inf_{T \geq t,\, \boldsymbol{\theta}_{-i}} U_{i|H_t}(\sigma_i, \boldsymbol{\sigma}_{-i}; \boldsymbol{\theta}_{-i}, T)$;
iii) for every σ_i, if $\mathcal{S}_t = \emptyset$, $i \in \hat{\mathcal{Q}}_t$ and $(\sigma_i^*(H_{i,t}) - P_{i,t})(\sigma_i(H_{i,t}) - P_{i,t}) < 0$, the inequality in ii) is strict.

In the above theorem, claim i) states that σ_i^* never incurs an expected loss, and it holds because a provider bidding its cost estimate serves a query only when the payment covers that estimate. Claim ii) states that σ_i^* maximizes the worst-case utility, and it holds because a provider who does not know T cannot rule out that the current query is the last one; consequently, the worst-case utility of any strategy is bounded above by its net payment from the current query. Moreover, under σ_i^*, the worst-case utility matches this upper bound exactly because, by claim i), no later query can reduce it. Since this worst-case scenario effectively isolates the current query, and because the payment does not depend on the provider's own bid, the optimal bid is the lowest possible value that avoids an expected loss—which is precisely the provider's cost estimate. Claim iii) states that any deviation changing whether the provider is assigned the query yields a strictly lower worst-case utility, and it holds because, if the current query happens to be the last one, such a deviation either causes the provider to forgo a profitable query or forces it to serve an unprofitable one.

Furthermore, as an immediate consequence of Theorem 2, the choice of bidding strategies $\boldsymbol{\sigma}^* = (\sigma_1^*, \ldots, \sigma_N^*)$ is a robust equilibrium, *i.e.*, for every σ_i

$$\inf_{T \geq t,\, \boldsymbol{\theta}_{-i}} U_{i|H_t}(\sigma_i^*, \boldsymbol{\sigma}_{-i}^*; \boldsymbol{\theta}_{-i}, T) \;\geq\; \inf_{T \geq t,\, \boldsymbol{\theta}_{-i}} U_{i|H_t}(\sigma_i, \boldsymbol{\sigma}_{-i}^*; \boldsymbol{\theta}_{-i}, T). \tag{9}$$

In words, at every step $t > N$ and after every history, no provider i can improve its worst-case utility by deviating from σ_i^* when the other providers follow $\boldsymbol{\sigma}_{-i}^*$.

4 Finding the Lowest Competitive Inference Price

In this section, we show that, when all providers choose their equilibrium strategy σ_i^*, the user is served by a provider meeting their quality threshold $q_{\min}$ on all but a vanishing fraction of the queries, while the price they pay converges to the most competitive price the market can offer for their task. For concreteness, in the remainder of the section, we specify the function $g(t)$ to be of the form $g(t) = k(t/N)^\alpha$, where $\alpha \in (0, 1)$ and $k > 0$ are given parameters.

We begin by identifying the provider that serves most of the user's queries. More precisely, we show that, on all but a vanishing fraction of queries, the platform selects the qualified provider i^* with the lowest cost, *i.e.*,

$$i^* = \arg\min_{i \in \mathcal{Q}} \bar{c}_i, \tag{10}$$

which we refer to as the optimal provider and assume to be unique. The number of queries not served by the optimal provider depends on the specific parameters α and k chosen by the platform as well as on the difficulty of distinguishing each provider from i^*, which we characterize using two quantities for each provider. First, for an unqualified provider $i \notin \mathcal{Q}$, we let $\varepsilon_i = q_{\min} - q_i > 0$ measure how far its quality falls below the quality threshold. Second, for a qualified provider $i \in \mathcal{Q} \setminus \{i^*\}$, we let $\Delta_i = \bar{c}_i - \bar{c}_{i^*} > 0$ measure how far its mean cost exceeds that of i^*. Intuitively, the smaller ε_i or Δ_i, the harder the corresponding provider is to distinguish from i^*, and the more queries the platform may allocate to them, as formalized in the next theorem.

Theorem 3. *When all providers choose their equilibrium strategy $\boldsymbol{\sigma}^*$, with probability at least $1-\delta$, the number of time steps in which the platform does not assign the query to the lowest-cost qualified provider,* i.e.*, $I_t \neq i^*$, is at most $N - 1 + B_{\mathrm{id}}(T)$, where*

$$B_{\mathrm{id}}(T) \coloneqq \sum_{i \notin \mathcal{Q}} M(\varepsilon_i) + \sum_{i \in \mathcal{Q} \setminus \{i^*\}} \max\left\{ M\left(\frac{\Delta_i}{c_{\max}}\right),\ k \cdot \left(\frac{T}{N}\right)^\alpha \right\}, \tag{11}$$

where, for any $x > 0$,

$$M(x) = \left\lceil \frac{4}{x^2}\ \ln\left(\frac{57N}{\delta x^4}\right) \right\rceil.$$

In the above theorem, the bound consists of the $N-1$ initialization steps not served by i^* alongside a sum of per-provider terms. For an unqualified provider, the term grows as $1/\varepsilon_i^2$ up to logarithmic factors. Conversely, for a qualified provider other than i^*, the term is determined by the maximum of two terms: a term growing as $1/\Delta_i^2$ that bounds the steps in which it outbids i^* and a term $k(T/N)^\alpha$ that bounds the exploration steps in which it is selected. Crucially, this bound depends sublinearly on the total number of queries T. Hence, the platform allocates the query to the optimal provider on all but a vanishing fraction of the user's queries.

We now turn to the price the user pays on the platform for a query, *i.e.*, the payment π_t made to the provider serving it. More concretely, we focus on comparing this payment with the second-lowest mean cost among qualified providers, $\bar{c}_{(2)} = \min_{i \in \mathcal{Q} \setminus \{i^*\}} \bar{c}_i$. This cost represents the most competitive (average) price the user could hope to pay for a query: it is the lowest price the optimal provider i^* would need to offer to remain cheaper than every other qualified provider, since the next-cheapest cannot profitably offer a price below its own mean cost $\bar{c}_{(2)}$. The next theorem shows that the average payment converges to the second-lowest mean cost among qualified providers.

Theorem 4. *When all providers choose their equilibrium strategy $\boldsymbol{\sigma}^*$ and the user submits a minimum number T_0 of queries, with probability at least $1 - \delta$ the payments made by the user satisfy*

$$\begin{aligned}\frac{1}{T}\sum_{t=1}^{T}\left|\pi_t - \bar{c}_{(2)}\right| \ \le\ & \frac{T_0}{T}\, c_{\max} + k\left(\frac{N}{T}\right)^{1-\alpha}\left(c_{\max} - \bar{c}_{(2)}\right) \\ & + \frac{6\, c_{\max}}{2-\alpha}\sqrt{\frac{1}{2k}\left(\frac{N}{T}\right)^{\alpha} \ln\left(\frac{2\pi^2 N k^2}{3\delta}\left(\frac{T}{N}\right)^{2\alpha}\right)},\end{aligned} \tag{12}$$

Benchmark	N	$q_{\min}$	$\lvert\mathcal{Q}\rvert$	i^*	$\bar{c}_{i^*}$	$\bar{v}_{i^*}$	$\bar{c}_{(2)}$	Payment π_t	Margin	Share of i^*
GSM8K	9	0.764	3	Qwen2.5-3B	16.1	20.1	34.4	30.3 (27.5)	71%	0.96
GPQA	4	0.145	2	Llama-3.1-8B	43.0	53.8	50.7	48.3 (47.4)	10%	0.97

Table 1: **Our procurement platform on queries from GSM8K and GPQA.** The table reports the number of simulated providers N, the quality threshold $q_{\min}$, the number of qualified providers $|\mathcal{Q}|$, the most cost-competitive provider i^* and its mean cost $\bar{c}_{i^*}$ and price $\bar{v}_{i^*}$, the mean cost of the second most cost-competitive provider $\bar{c}_{(2)}$, the average payment π_t over all steps and, in parentheses, over non-exploration steps, the average pricing margin of the winner over non-exploration steps, and the share of queries routed to i^*. The last three quantities are computed over the last K queries, averaged across 30 independent runs. Monetary quantities are in units of 10^{-6} USD.

where $T_0 = N \max\{1, (\bar{M}/k)^{1/\alpha}\}$ *and* $\bar{M} = \max\{1, \max_{i \notin \mathcal{Q}} M(\varepsilon_i)\}$. *Hence,* $\frac{1}{T}\sum_{t=1}^{T} \pi_t \to \bar{c}_{(2)}$ *as* $T \to \infty$.

In the above theorem, the first term accounts for the first T_0 time steps, which include the N initialization steps and the non-exploration steps during which an unqualified provider may still belong to $\hat{\mathcal{Q}}_t$; the second term accounts for the exploration steps after T_0, in which the selected provider is paid $c_{\max}$; and the third term accounts for the remaining steps. As T grows, the average difference between the payment and the second-lowest mean cost $\bar{c}_{(2)}$ of a qualified provider vanishes at a rate $\max\{T^{-(1-\alpha)}, T^{-\alpha/2}\}$, up to logarithmic factors. In Corollary 12 in Appendix B.9, we show that this rate is maximized when the platform sets $\alpha = 2/3$, giving a rate of $T^{-1/3}$.

5 Experiments

In this section, we evaluate our platform on simulated providers serving real LLMs from the Llama and Qwen families. Specifically, we evaluate whether the platform learns the set of qualified providers and progressively routes queries to the optimal provider i^*. We then compare the price paid on our platform against the average per-token price listed by inference providers on Hugging Face[9]. Appendix E and Appendices F–H contain further details on our experimental setup and additional experimental results, respectively.

Experimental setup. Each of the N simulated providers serves one LLM, which we use to identify them. We consider two settings: (i) $N = 9$ providers, each serving a model from the Llama family (Llama-{3-8B, 3.1-8B, 3.2-1B, 3.2-3B}-Instruct) or the Qwen family (Qwen-{2-0.5B, 2-1.5B, 2-7B, 2.5-3B, 2.5-7B}-Instruct), and (ii) $N = 4$ providers serving the newest of the largest models and the smallest model of each family, *i.e.*, Llama-3.1-8B, Llama-3.2-1B, Qwen2.5-7B and Qwen2-0.5B.

In both settings, the user's queries are drawn from one of two widely used benchmarks. In the setting with $N = 9$ providers, we draw queries from the GSM8K dataset (Cobbe et al., 2021), which consists of $K = 1{,}319$ grade-school mathematical reasoning problems. In the setting with $N = 4$ providers, we draw queries from the GPQA (Rein et al., 2024) dataset, which consists of $K = 546$ graduate-level science questions. Both benchmarks feature verifiable ground-truth answers, and the score $r(x_t, y_t)$ represents whether the response is correct (1) or incorrect (0). Consequently, the quality q_i of a provider corresponds to the probability that the model generates a correct answer. We rely on a public dataset of recorded generations,[10] which contains 128 generations per model for every question across both benchmarks. In each experiment, we simulate $T = 70{,}000$ queries by taking 53.07 and 128.21 passes over all queries in GSM8K and GPQA, respectively, sampling without replacement. At each step, once a provider is selected, its response is sampled uniformly at random from its 128 recorded generations for that query. We repeat each experiment across 30 independent runs to report averaged results.

Since the public dataset we use lacks generation costs, we construct them using real-world market signals. In particular, the cost of each recorded generation is determined by its token count and the model's average per-token price on the Hugging Face inference-provider list—on the order of

[9] huggingface.co/docs/inference-providers/index, consulted on December 30, 2025.

[10] https://huggingface.co/datasets/Human-Centric-Machine-Learning/strategic-ttc-data

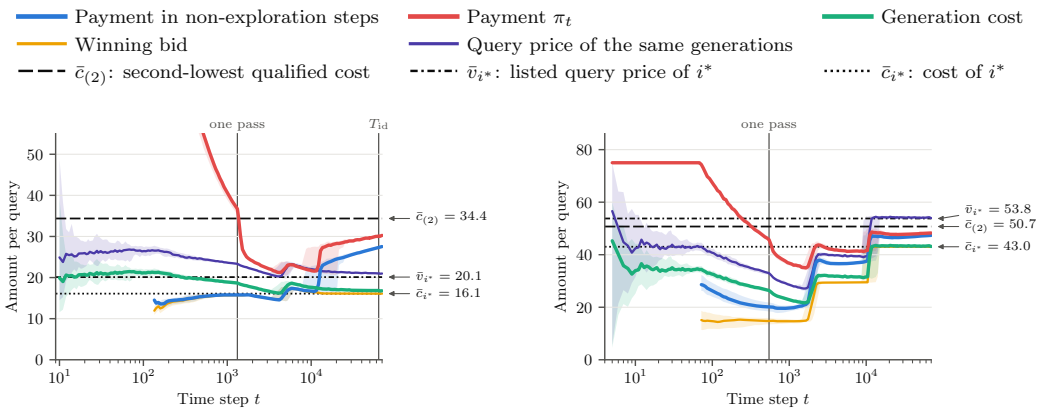

(a) GSM8K, weak competition

(b) GPQA, strong competition

Figure 1: **Temporal evolution of the winning bids, payments, and costs.** For each reported quantity, every point corresponds to the median over 30 independent runs of a moving average whose window size equals the benchmark size K, and the shaded regions span the 10th to 90th percentiles across runs. The vertical lines mark one complete pass over the benchmark queries and, where applicable within the time horizon, the time step T_{id} beyond which Theorem 3 guarantees that i^* becomes the unique most selected provider. Amount per query is in units of 10^{-6} USD.

\$0.1 per million tokens—discounted by a fixed 25% margin (*i.e.*, divided by 1.25), common to all providers and unknown to the platform (see Appendix E, Table 2). We define $\bar{v}_i = 1.25\,\bar{c}_i$ as the mean query price of provider i, representing the average amount a user pays per query under publicly listed prices. Further, we place the quality threshold $q_{\min}$ at the midpoint between two adjacent provider qualities, so that three providers qualify on GSM8K and two on GPQA.

We purposely choose these two settings to vary the intensity of competition among qualified providers, which fundamentally determines the user's final payment. Specifically, since the average payment converges to the second-lowest mean cost $\bar{c}_{(2)}$ among qualified providers by Theorem 4, we call competition strong when $\bar{c}_{(2)}$ is smaller than the listed query price of the optimal provider $\bar{v}_{i^*}$, and weak otherwise. On GSM8K, the qualified competitors of i^* are considerably more expensive ($\bar{c}_{(2)} > \bar{v}_{i^*}$, weak), whereas on GPQA, the sole competitor is only slightly more expensive ($\bar{c}_{(2)} < \bar{v}_{i^*}$, strong). Table 1 summarizes both settings.

The platform runs Algorithm 1 with confidence level $\delta = 0.05$ and exploration function $g(t) = k(t/N)^{\alpha}$, with $k = 2$ and $\alpha = 3/4$. We choose this value over the rate-optimal $\alpha = 2/3$ from Corollary 12 to induce greater early exploration at the expense of a slower asymptotic rate. Under this setup, the simulated providers follow their equilibrium strategy $\boldsymbol{\sigma}^*$.

Results. Table 1 summarizes our main results. We find that the platform successfully identifies the optimal provider: over the final pass of the benchmark, it routes at least 96% of all queries to i^*, and i^* becomes the most selected provider after roughly 10,000 and 30,000 queries on GSM8K and GPQA respectively. In both cases, this occurs well before the time step $T_{\mathrm{id}} = \min\{T : T > N + 2B_{\mathrm{id}}(T)\}$, *i.e.*, the time step beyond which Theorem 3 guarantees, with probability at least $1 - \delta$, that i^* is the unique most-selected provider (63,740 and 153,697 queries, respectively; see Appendices F.1 and H.1). This gap demonstrates that our sample complexity bounds are highly conservative in practice. The delayed convergence on GPQA stems from the presence of the unqualified `Llama-3.2-1B` model; because its quality falls a mere 0.037 below $q_{\min}$, the platform requires approximately 11,000 queries to statistically exclude it from $\hat{\mathcal{Q}}_t$.

Figure 1 provides additional insight into the temporal evolution of the winning bids, payments and costs. In agreement with Theorem 4, we find that, for the initial queries, the payment on non-exploration steps lies below $\bar{c}_{i^*}$; this occurs because cheaper, unqualified providers that the platform has not yet excluded from $\hat{\mathcal{Q}}_t$ are still selected and influence the pricing mechanism. However, as the platform routes more queries, the payment converges towards $\bar{c}_{(2)}$. Importantly, we also find that the user's payment for a guaranteed quality level is fundamentally determined by the intensity of com-

petition among qualified providers on each benchmark. On the final pass, under strong competition on GPQA, the average payment per query on non-exploration steps lies 10% above the generation cost $\bar{c}_{i^*}$ and 12% below the listed price $\bar{v}_{i^*}$ that the user would pay in the current fixed-price market. Conversely, under weak competition on GSM8K, the average payment lies 71% above $\bar{c}_{i^*}$ and 37% above $\bar{v}_{i^*}$, due to the higher cost of the nearest qualified competitor. Altogether, the results suggest that our platform allows users to pay a price that approaches the one supported by provider competition on their specific task, capturing substantial savings whenever the market-clearing price falls below public fixed listings.

6 Discussion and Limitations

In this section, we highlight several limitations of our work and discuss avenues for future research.

Methodology. We assume that a provider maximizes its worst-case utility by bidding its current cost estimate, since it knows neither the total number of user queries nor the score and cost distributions of its competitors. However, because providers may violate this assumption in practice, we extend our theoretical results in Appendix D to a setting where bids can deviate above or below the cost estimate by bounded amounts. In this context, analyzing our platform using other game-theoretic solution concepts remains an interesting direction for future work. Further, whenever the cost difference across providers is small, or the providers' quality sits right on the threshold, our platform requires a relatively large number of queries to secure the best price for the user. In such cases, future work could lower query requirements by leveraging quality and cost estimates across tasks. Finally, in our platform, users carry the cost of exploration. However, because the most cost-competitive provider directly benefits from this exploration, one may argue that this burden should be shifted to, or at least shared with, the providers, who typically possess greater financial resources. Exploring alternative payment rules that redistribute such a burden is an interesting direction for future work.

Evaluation. We evaluated our platform using simulated providers who compete to serve queries from benchmarking datasets with verifiable ground truth. An important next step would be to evaluate our platform in settings where providers compete to serve open-ended queries and quality is measured automatically without relying on predefined correct outputs. Further, in our experiments, we estimated the generation costs by discounting publicly listed token prices by a fixed 25% profit margin. Future research should leverage specialized profiling tools to measure these true costs precisely. Ultimately, deploying and evaluating the proposed platform in a real-world production environment remains crucial.

Practical considerations. We assume that our platform can accurately estimate the quality of the responses generated by the providers. In many domains where LLMs are routinely used, such as coding or mathematical reasoning, automated evaluation procedures make this feasible with minimal overhead. Moving forward, the viability of our platform in settings with less well-established quality metrics and benchmarks will depend on continued progress in the LLM evaluation literature. Lastly, we note that, while second-price auctions possess many desirable properties, they also make a provider's payment dependent on the bids of its competitors. This dependence leaves the mechanism potentially vulnerable to manipulation or collusion (Hendricks & Porter, 1989; Mailath & Zemsky, 1991; Che et al., 2018).

7 Conclusions

In this work, we introduced a procurement platform that connects large language model providers with end-users, where token prices are driven by provider competition via reverse second-price auctions. More concretely, this mechanism incentivizes providers to truthfully bid their best estimate of the average cost to serve a user's query. We demonstrated that, with high probability, the platform routes all but a vanishing fraction of the user's queries to the most cost-competitive provider meeting a desired quality threshold. Moreover, this provider's average pricing margin converges to the difference between the second-lowest average cost among such providers and its own average cost. Empirically, we demonstrated that the pricing margin on our platform varies significantly—from 10% to 71%—depending on the task and quality threshold, revealing a substantial inefficiency in the current fixed-price market. Ultimately, we hope that our work inspires users, providers, and routing platforms to explore alternative, more efficient economic paradigms for the market of LLM-as-a-service.

AI USE STATEMENT

In this work, we used generative AI tools for limited language and style refinement, proofreading, and minor assistance with code. We did not use generative AI tools to generate research ideas, develop the methodology, or derive theoretical results or proofs. All AI-assisted text and code were reviewed and verified by the authors. We take responsibility for the final content of this work, including text, claims, and artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

We have taken several steps to make our results reproducible. On the theoretical side, Algorithm 1 specifies the platform completely, every assumption is stated where it is used, and complete proofs of all results in Sections 3 and 4 are given in Appendix B, with the extensions to Bayesian cost estimation and to deviations from cost-estimate bidding in Appendices C and D. On the empirical side, all experiments use a publicly available dataset of recorded LLM generations, and Appendix E describes how providers, qualities, generation costs and the quality threshold are constructed from it, including the listed prices we use (Table 2), the platform's parameters, the seeding of the runs, and every metric we report; Appendices F–H report the complete results, of which Section 5 shows a subset. The supplementary material contains the source code of the simulator, the configuration file defining every environment, the scripts that download and build the datasets, and the executed notebooks that produce all figures and tables in the paper, together with the per-run results and the seeds of every run.

REFERENCES

Michele Aghassi and Dimitris Bertsimas. Robust game theory. *Math. Program.*, 107(1–2):231–273, June 2006. ISSN 0025-5610. doi: 10.1007/s10107-005-0686-0. URL https://doi.org/10.1007/s10107-005-0686-0.

Mohammad Akbarpour and Shengwu Li. Credible auctions: A trilemma. *Econometrica*, 88(2): 425–467, 2020.

Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Mach. Learn.*, 47(2-3):235–256, May 2002.

Yu Awaya, Naoki Fujiwara, and Marton Szabo. Quality and price in scoring auctions. *Journal of Mathematical Economics*, 116:103083, 2025. ISSN 0304-4068. doi: https://doi.org/10.1016/j.jmateco.2024.103083. URL https://www.sciencedirect.com/science/article/pii/S0304406824001435.

Moshe Babaioff, Yogeshwer Sharma, and Aleksandrs Slivkins. Characterizing truthful multi-armed bandit mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce*, pp. 79–88, New York, NY, USA, July 2009. ACM.

Ziv Bar-Yossef, Kirsten Hildrum, and Felix Wu. Incentive-compatible online auctions for digital goods. In *SODA*, volume 2, pp. 964–970, 2002.

Dirk Bergemann and Juuso Välimäki. Dynamic mechanism design: An introduction. *J. Econ. Lit.*, 57(2):235–274, June 2019.

Dirk Bergemann, Marek Bojko, Paul Duetting, Renato Paes Leme, Haifeng Xu, and Song Zuo. Data-driven mechanism design: Jointly eliciting preferences and information. In *Proceedings of the 26th ACM Conference on Economics and Computation*, EC '25, pp. 507, New York, NY, USA, 2025a. Association for Computing Machinery. ISBN 9798400719431. doi: 10.1145/3736252.3742578. URL https://doi.org/10.1145/3736252.3742578.

Dirk Bergemann, Alessandro Bonatti, and Alex Smolin. The economics of large language models: Token allocation, fine-tuning, and optimal pricing. In *Proceedings of the 26th ACM Conference on Economics and Computation*, EC '25, pp. 786, New York, NY, USA, 2025b. Association for Computing Machinery. ISBN 9798400719431. doi: 10.1145/3736252.3742625. URL https://doi.org/10.1145/3736252.3742625.

Dirk Bergemann, Alessandro Bonatti, and Alex Smolin. Menu pricing of large language models. *arXiv preprint arXiv:2502.07736*, 2025c.

Mark Braverman, Jieming Mao, Jon Schneider, and S. Matthew Weinberg. Multi-armed bandit problems with strategic arms. In Alina Beygelzimer and Daniel Hsu (eds.), *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pp. 383–416. PMLR, 25–28 Jun 2019. URL https://proceedings.mlr.press/v99/braverman19b.html.

Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Found. Trends® Mach. Learn.*, 5(1):1–122, December 2012.

Yuhang Cai and Huchang Liao. An overview of multi-attribute auctions: bibliometrics, methodologies, applications and future directions. *Expert Systems with Applications*, 299:130083, 2026. ISSN 0957-4174. doi: https://doi.org/10.1016/j.eswa.2025.130083. URL https://www.sciencedirect.com/science/article/pii/S0957417425036991.

Yuhan Cao, Yu Wang, Sitong Liu, Miao Li, Yixin Tao, and Tianxing He. Pay for the second-best service: A game-theoretic approach against dishonest llm providers. In *Proceedings of the ACM Web Conference 2026*, WWW '26, pp. 249–260, New York, NY, USA, 2026. Association for Computing Machinery. ISBN 9798400723070. doi: 10.1145/3774904.3792446. URL https://doi.org/10.1145/3774904.3792446.

Yeon-Koo Che. Design competition through multidimensional auctions. *The RAND Journal of Economics*, 24(4):668–680, 1993. ISSN 07416261. URL http://www.jstor.org/stable/2555752.

Yeon-Koo Che, Daniele Condorelli, and Jinwoo Kim. Weak cartels and collusion-proof auctions. *Journal of Economic Theory*, 178:398–435, 2018. ISSN 0022-0531. doi: https://doi.org/10.1016/j.jet.2018.09.005. URL https://www.sciencedirect.com/science/article/pii/S0022053118305842.

Lingjiao Chen, Chi Zhang, Yeye He, Ion Stoica, Matei Zaharia, and James Zou. The price reversal phenomenon: When cheaper reasoning models cost more. *arXiv preprint arXiv:2603.23971*, 2026a.

Yiling Chen, Shi Feng, and Sadie Zhao. Contextual procurement auctions with bandit learning. *arXiv preprint arXiv:2607.05813*, 2026b.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems, 2021. URL https://arxiv.org/abs/2110.14168.

Mehmet Hamza Erol, Batu El, Mirac Suzgun, Mert Yuksekgonul, and James Zou. Cost-of-pass: An economic framework for evaluating language models. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=vC9S20zsgN.

Eyal Even-Dar, Shie Mannor, and Yishay Mansour. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *Journal of Machine Learning Research*, 7(39):1079–1105, 2006. URL http://jmlr.org/papers/v7/evendar06a.html.

Zhe Feng, Chara Podimata, and Vasilis Syrgkanis. Learning to bid without knowing your value. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, EC '18, pp. 505–522, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450358293. doi: 10.1145/3219166.3219208. URL https://doi.org/10.1145/3219166.3219208.

Rupert Freeman, David Pennock, Chara Podimata, and Jennifer Wortman Vaughan. No-regret and incentive-compatible online learning. In Hal Daumé III and Aarti Singh (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 3270–3279. PMLR, 13–18 Jul 2020. URL https://proceedings.mlr.press/v119/freeman20a.html.

Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, 1991.

John C. Harsanyi. Games with incomplete information played by "bayesian" players, i-iii: Part i. the basic model. *Management Science*, 50(12):1804–1817, 2004. ISSN 00251909, 15265501. URL http://www.jstor.org/stable/30046151.

Kenneth Hendricks and Robert H. Porter. Collusion in auctions. *Annales d'Économie et de Statistique*, (15/16):217–230, 1989. ISSN 0769489X, 22726497. URL http://www.jstor.org/stable/20075758.

Tanjim Hossain. Learning by bidding. *The RAND Journal of Economics*, 39(2):509–529, 2008. ISSN 07416261. URL http://www.jstor.org/stable/25474380.

Shweta Jain, Sujit Gujar, Satyanath Bhat, Onno Zoeter, and Y. Narahari. An incentive compatible multi-armed-bandit crowdsourcing mechanism with quality assurance, 2015. URL https://arxiv.org/abs/1406.7157.

Emanuele La Malfa, Aleksandar Petrov, Simon Frieder, Christoph Weinhuber, Ryan Burnell, Raza Nazar, Anthony Cohn, Nigel Shadbolt, and Michael Wooldridge. Language-models-as-a-service: Overview of a new paradigm and its challenges. *Journal of Artificial Intelligence Research*, 80: 1497–1523, 2024.

Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, Cambridge, England, July 2020.

Benjamin Laufer, Jon Kleinberg, and Hoda Heidari. Fine-Tuning Games: Bargaining and Adaptation for General-Purpose Models. In *Proceedings of the ACM Web Conference 2024*, WWW '24, pp. 66–76, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400701719. doi: 10.1145/3589334.3645366. URL https://doi.org/10.1145/3589334.3645366.

John O Ledyard. Incentive compatibility. In *The New Palgrave dictionary of economics*, pp. 1–9. Springer, 1987.

Rafid Mahmood. Pricing and Competition for Generative AI. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL https://openreview.net/forum?id=8LbJfEjIrT.

George J Mailath and Peter Zemsky. Collusion in second price auctions with heterogeneous bidders. *Games and Economic Behavior*, 3(4):467–486, 1991. ISSN 0899-8256. doi: https://doi.org/10.1016/0899-8256(91)90016-8. URL https://www.sciencedirect.com/science/article/pii/0899825691900168.

R Preston McAfee and John McMillan. Bidding for contracts: A principal-agent analysis. *Rand J. Econ.*, 17(3):326, 1986.

Paul Milgrom and John Roberts. Predation, reputation, and entry deterrence. *Journal of Economic Theory*, 27(2):280–312, 1982. ISSN 0022-0531. doi: https://doi.org/10.1016/0022-0531(82)90031-X. URL https://www.sciencedirect.com/science/article/pii/002205318290031X.

Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani (eds.). *Algorithmic game theory*. Cambridge University Press, Cambridge, England, January 2011.

Ricardo Olmedo, Bernhard Schölkopf, and Moritz Hardt. Computational Arbitrage in AI Model Markets. In *Forty-third International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=0xUvKZSQxV.

Divya Padmanabhan, Satyanath Bhat, K J Prabuchandran, Shirish Shevade, and Y Narahari. Dominant strategy truthful, deterministic multi-armed bandit mechanisms with logarithmic regret for sponsored search auctions. *Appl. Intell.*, 52(3):3209–3226, February 2022.

Pronoy Patra, Sankarshan Damle, Manisha Padala, and Sujit Gujar. Truthful reverse auctions for adaptive selection via contextual multi-armed bandits. In *Proceedings of the 25th International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '26, pp. 578–586, Richland, SC, 2026. International Foundation for Autonomous Agents and Multiagent Systems. ISBN 9798400723179. doi: 10.65109/MBRQ7564. URL https://doi.org/10.65109/MBRQ7564.

Alessandro Pavan, Ilya Segal, and Juuso Toikka. Dynamic mechanism design: A myersonian approach. *Econometrica*, 82(2):601–653, 2014. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/24029270.

Manish Raghavan. Competition and Diversity in Generative AI. *arXiv preprint arXiv:2412.08610*, 2024.

David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R. Bowman. GPQA: A graduate-level google-proof q&a benchmark. In *First Conference on Language Modeling*, 2024. URL https://openreview.net/forum?id=Ti67584b98.

Tim Roughgarden, Vasilis Syrgkanis, and Eva Tardos. The price of anarchy in auctions. *Journal of Artificial Intelligence Research*, 59:59–101, 2017.

Adish Singla and Andreas Krause. Truthful incentives in crowdsourcing tasks using regret minimization mechanisms. In *Proceedings of the 22nd international conference on World Wide Web*, New York, NY, USA, May 2013. ACM.

Ander Artola Velasco, Dimitrios Rontogiannis, Stratis Tsirtsis, and Manuel Gomez-Rodriguez. Test-Time Compute Games, 2026a. URL https://arxiv.org/abs/2601.21839.

Ander Artola Velasco, Stratis Tsirtsis, and Manuel Gomez-Rodriguez. Auditing pay-per-token in large language models. In *The 29th International Conference on Artificial Intelligence and Statistics*, 2026b.

Ander Artola Velasco, Stratis Tsirtsis, Nastaran Okati, and Manuel Gomez Rodriguez. Is Your LLM Overcharging You? Tokenization, Transparency, and Incentives. In *Forty-third International Conference on Machine Learning*, 2026c. URL https://openreview.net/forum?id=aGaBFE9pta.

William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961. ISSN 00221082, 15406261. URL http://www.jstor.org/stable/2977633.

Haoxiang Wang, Wei Xiong, Tengyang Xie, Han Zhao, and Tong Zhang. Interpretable preferences via multi-objective reward modeling and mixture-of-experts, 2024a. URL https://arxiv.org/abs/2406.12845.

Junlin Wang, Siddhartha Jain, Dejiao Zhang, Baishakhi Ray, Varun Kumar, and Ben Athiwaratkun. Reasoning in token economies: Budget-aware evaluation of LLM reasoning strategies. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 19916–19939, Miami, Florida, USA, November 2024b. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.1112. URL https://aclanthology.org/2024.emnlp-main.1112/.

Jonathan Weed, Vianney Perchet, and Philippe Rigollet. Online learning in repeated auctions. In Vitaly Feldman, Alexander Rakhlin, and Ohad Shamir (eds.), *29th Annual Conference on Learning Theory*, volume 49 of *Proceedings of Machine Learning Research*, pp. 1562–1583, Columbia University, New York, New York, USA, 23–26 Jun 2016. PMLR. URL https://proceedings.mlr.press/v49/weed16.html.

Michael J. Zellinger and Matt Thomson. Economic evaluation of llms, 2025. URL https://arxiv.org/abs/2507.03834.

A ALGORITHM

Algorithm 1 It routes the user's queries to a qualified provider at the lowest competitive price.

Input: providers $[N]$, quality threshold $q_{\min}$, confidence level δ, allocation parameters k, α, user-defined score r, cost bound $c_{\max}$
Initialization: assign one query to each provider $i \in [N]$, pay them $c_{\max}$, and receive their first standing bid b_i
Set $m_i \leftarrow 1$ and $\hat{q}_i \leftarrow r(x_i, y_i)$ for all $i \in [N]$
for $t = N+1, N+2, \ldots$ **do**
  Receive user query $x_t \sim P^x$
  $\beta_i \leftarrow \sqrt{\frac{1}{2m_i} \ln\left(\frac{2\pi^2 N m_i^2}{3\delta}\right)}$ for all $i \in [N]$ ▷ confidence radius
  $\hat{\mathcal{Q}} \leftarrow \{\, i \in [N] : \hat{q}_i + \beta_i \geq q_{\min} \,\}$ ▷ optimistic qualified set
  $g \leftarrow k\,(t/N)^{\alpha}$ ▷ target selection count
  if $\hat{\mathcal{Q}} = \emptyset$ **then**
    $I_t \leftarrow$ a provider chosen uniformly at random from $[N]$; $\pi_t \leftarrow c_{\max}$
  else if $m_i < g$ for some $i \in \hat{\mathcal{Q}}$ **then**
    $I_t \leftarrow$ a provider chosen uniformly at random from $\{\, i \in \hat{\mathcal{Q}} : m_i < g \,\}$; $\pi_t \leftarrow c_{\max}$
  else
    $I_t \leftarrow \arg\min_{i \in \hat{\mathcal{Q}}} \left(b_i - c_{\max}\,\beta_i\right)$ ▷ allocation rule, ties broken arbitrarily
    $P \leftarrow c_{\max}\,\beta_{I_t} + \min_{j \in \hat{\mathcal{Q}} \setminus \{I_t\}} \left(b_j - c_{\max}\,\beta_j\right)$
    $\pi_t \leftarrow \min\{P, c_{\max}\}$ ▷ second-price payment
  end if
  The user pays provider I_t the amount π_t
  Provider I_t serves the query, returning a response y_t
  Observe the score $r(x_t, y_t)$ and set $\hat{q}_{I_t} \leftarrow \frac{m_{I_t}\,\hat{q}_{I_t} + r(x_t, y_t)}{m_{I_t}+1}$ and $m_{I_t} \leftarrow m_{I_t} + 1$
  Provider I_t may revise its standing bid b_{I_t}; all other bids are unchanged
end for

B PROOFS

B.1 CONCENTRATION INEQUALITIES FOR THE QUALITIES AND COSTS

In this section, we characterize the relationship between the estimated confidence sets $\hat{\mathcal{Q}}_t$, introduced in Eq. 2, and the set $\mathcal{Q}$ of qualified providers. To this end, we first define a high-probability event under which, for every provider and every selection count, the empirical quality and cost averages are simultaneously close to their corresponding population averages. For each provider $i \in [N]$ and $m \geq 1$, let $\hat{q}_{i,m}$ and $\hat{c}_{i,m}$ denote the empirical means of the scores and generation costs observed over the first m queries served by provider i, respectively. Recall that

$$\beta(m) = \sqrt{\frac{1}{2m} \ln\left(\frac{2\pi^2 N m^2}{3\delta}\right)}.$$

Lemma 1. *Let $\mathcal{E}$ denote the event that, simultaneously for every provider $i \in [N]$ and every $m \geq 1$,*

$$|\hat{q}_{i,m} - q_i| \leq \beta(m), \qquad |\hat{c}_{i,m} - \bar{c}_i| \leq c_{\max}\beta(m). \tag{13}$$

Then, $\Pr(\mathcal{E}) \geq 1 - \delta$.

Proof. Fix a provider $i \in [N]$, and let $\tau_i(1) < \tau_i(2) < \cdots$ denote the steps in which provider i is selected. Since the platform selects a provider before the score and generation cost of that step are realized, the successive observations

$$r(x_{\tau_i(1)}, y_{\tau_i(1)}), r(x_{\tau_i(2)}, y_{\tau_i(2)}), \ldots$$

are i.i.d. draws from $\mathcal{P}_i^r$ with mean q_i, and the successive realized costs

$$c_{i,\tau_i(1)}, c_{i,\tau_i(2)}, \ldots$$

are i.i.d. draws from $\mathcal{P}_i^c$ with mean $\bar{c}_i$. Thus, for every fixed selection count $m \geq 1$, Hoeffding's inequality and the fact that scores lie in $[0, 1]$ give

$$\Pr(|\hat{q}_{i,m} - q_i| > \beta(m)) \leq 2\exp\left(-2m\beta(m)^2\right).$$

Similarly, since generation costs lie in $[0, c_{\max}]$,

$$\Pr(|\hat{c}_{i,m} - \bar{c}_i| > c_{\max}\beta(m)) \leq 2\exp\left(-2m\beta(m)^2\right).$$

By the definition of $\beta(m)$,

$$2\exp\left(-2m\beta(m)^2\right) = \frac{3\delta}{\pi^2 N m^2}.$$

A union bound over the two estimates, all providers, and all selection counts therefore gives

$$\begin{aligned}\Pr(\mathcal{E}^c) &\leq \sum_{i=1}^{N}\sum_{m=1}^{\infty} 2\frac{3\delta}{\pi^2 N m^2} \\ &= \frac{6\delta}{\pi^2}\sum_{m=1}^{\infty}\frac{1}{m^2} = \delta,\end{aligned}$$

where $\sum_{m=1}^{\infty} m^{-2} = \pi^2/6$. Hence,

$$\Pr(\mathcal{E}) \geq 1 - \delta.$$

□

We next prove a technical result showing that a lower bound on the confidence radius $\beta(m)$ implies an upper bound on the selection count m. We will use this result in the proof of Lemma 7 and Lemma 8.

Lemma 2. *For any $a > 0$ and $m \geq 1$, if*

$$a \leq 2\beta(m), \tag{14}$$

then

$$m \leq \frac{4}{a^2}\ln\left(\frac{57N}{\delta a^4}\right). \tag{15}$$

Proof. By equation 14 and the definition of $\beta(m)$,

$$\frac{1}{2m}\ln\left(\frac{2\pi^2 N m^2}{3\delta}\right) \geq \frac{a^2}{4}.$$

Equivalently,

$$m \leq \frac{2}{a^2}\ln\left(\frac{2\pi^2 N m^2}{3\delta}\right) = \frac{2}{a^2}\ln\left(\frac{2\pi^2 N}{3\delta}\right) + \frac{4}{a^2}\ln m.$$

Set

$$A = \frac{2}{a^2}\ln\left(\frac{2\pi^2 N}{3\delta}\right), \qquad B = \frac{4}{a^2},$$

so that $m \leq A + B\ln m$. For every $r > 0$, the inequality $\ln x \leq x - 1$ gives

$$\ln m = \ln\left(\frac{m}{r}\right) + \ln r \leq \frac{m}{r} + \ln r - 1.$$

Hence,

$$\left(1 - \frac{B}{r}\right) m \leq A + B(\ln r - 1).$$

Choosing $r = 2B = 8/a^2$ yields

$$\frac{m}{2} \leq A + B\left(\ln\left(\frac{8}{a^2}\right) - 1\right).$$

Therefore,

$$\begin{aligned} m &\leq 2A + 2B\left(\ln\left(\frac{8}{a^2}\right) - 1\right) \\ &= \frac{4}{a^2}\left[\ln\left(\frac{2\pi^2 N}{3\delta}\right) + 2\ln\left(\frac{8}{a^2}\right) - 2\right] \\ &= \frac{4}{a^2}\ln\left(\frac{128\pi^2 N}{3e^2\delta a^4}\right) \\ &\leq \frac{4}{a^2}\ln\left(\frac{57N}{\delta a^4}\right), \end{aligned}$$

where the final inequality follows from $128\pi^2/(3e^2) < 57$. □

B.2 Properties of exploration steps

We next establish several basic properties of the allocation rule that will be used when proving the theoretical results stated in Section 3 and Section 4.

Lemma 3. *On $\mathcal{E}$, every qualified provider remains in the estimated-qualified set at every post-initialization step:*

$$\mathcal{Q} \subseteq \hat{\mathcal{Q}}_t, \qquad t > N. \tag{16}$$

In particular, $\hat{\mathcal{Q}}_t \neq \emptyset$ for every $t > N$.

Proof. Fix $i \in \mathcal{Q}$ and $t > N$. Since i is qualified,

$$q_i \geq q_{\min}.$$

On $\mathcal{E}$, Lemma 1 gives, for any provider i and time step $t > N$,

$$\hat{q}_{i,t} \geq q_i - \beta(m_{i,t}).$$

Therefore,

$$\hat{q}_{i,t} + \beta(m_{i,t}) \geq q_i \geq q_{\min},$$

and hence $i \in \hat{\mathcal{Q}}_t$ by the definition of the estimated-qualified set. Since this holds for every $i \in \mathcal{Q}$, we obtain $\mathcal{Q} \subseteq \hat{\mathcal{Q}}_t$. Finally, $\mathcal{Q}$ contains at least two providers by assumption, so $\hat{\mathcal{Q}}_t$ is nonempty under the event $\mathcal{E}$. □

Lemma 4. *Suppose that $g\colon \mathbb{N} \to \mathbb{R}$ is non-decreasing, and fix a total number of user queries $T > N$. If provider i is selected in an exploration step $t \le T$, then*

$$m_{i,t} < g(t) \le g(T). \tag{17}$$

Consequently, provider i can be selected in an exploration step at most

$$\lceil g(T) \rceil - 1 \tag{18}$$

times, and the total number of post-initialization exploration steps up to T queries is at most

$$N\left(\lceil g(T) \rceil - 1\right). \tag{19}$$

Proof. If provider i is selected in an exploration step t, then $i \in \mathcal{S}_t$. By definition of $\mathcal{S}_t$,

$$m_{i,t} < g(t).$$

Since g is non-decreasing and $t \le T$,

$$m_{i,t} < g(t) \le g(T),$$

which proves equation 17.

Now suppose that t is an exploration step in which provider i is selected. Since $m_{i,t}$ is a positive integer after initialization, equation 17 implies

$$m_{i,t} \le \lceil g(T) \rceil - 1.$$

Thus, provider i can therefore be selected in an exploration step at most $\lceil g(T) \rceil - 1$ when the total number of user queries is T. Summing this bound over all N providers gives equation 19. □

B.3 Properties of non-exploration steps

We next establish properties of the allocation rule on non-exploration steps, that is, when $\mathcal{S}_t = \emptyset$ and the selection of provider I_t depends on the bids $\mathbf{b}_t$. We begin by proving Proposition 1.

Proof of Proposition 1.

Proof of Proposition 1. Fix a time step $t > N$ with $\mathcal{S}_t = \emptyset$, a provider $i \in \hat{\mathcal{Q}}_t$, and the other providers' bids $\mathbf{b}_{-i,t}$. Since t is not an exploration step, the allocation rule (Eq. 3) selects a provider minimizing the adjusted bid

$$b_{j,t} - c_{\max}\beta(m_{j,t}) \tag{20}$$

over $j \in \hat{\mathcal{Q}}_t$.

Provider i is therefore selected whenever

$$b_{i,t} - c_{\max}\beta(m_{i,t}) < \min_{j \in \hat{\mathcal{Q}}_t \setminus \{i\}} \left(b_{j,t} - c_{\max}\beta(m_{j,t})\right).$$

Rearranging gives

$$b_{i,t} < c_{\max}\beta(m_{i,t}) + \min_{j \in \hat{\mathcal{Q}}_t \setminus \{i\}} \left(b_{j,t} - c_{\max}\beta(m_{j,t})\right) = P_{i,t},$$

where the equality follows from the definition of $P_{i,t}$. Thus $b_{i,t} < P_{i,t}$ implies that provider i is selected.

Similarly, if

$$b_{i,t} > P_{i,t},$$

then

$$b_{i,t} - c_{\max}\beta(m_{i,t}) > \min_{j \in \hat{\mathcal{Q}}_t \setminus \{i\}} \left(b_{j,t} - c_{\max}\beta(m_{j,t})\right).$$

so provider i does not minimize the adjusted bid and is therefore not selected. This proves the claim. □

Lemma 5. *Suppose that all providers use their equilibrium strategy $\boldsymbol{\sigma}^*$. On $\mathcal{E}$, for every provider $i \in [N]$ and every step $t > N$,*

$$\bar{c}_i - 2c_{\max}\beta(m_{i,t}) \le b_{i,t} - c_{\max}\beta(m_{i,t}) \le \bar{c}_i. \tag{21}$$

Proof. Fix a provider $i \in [N]$ and a step $t > N$. Under the equilibrium strategy σ^*,

$$b_{i,t} = \hat{c}_{i,t}.$$

Thus, on $\mathcal{E}$,

$$\bar{c}_i - c_{\max}\beta(m_{i,t}) \le \hat{c}_{i,t} \le \bar{c}_i + c_{\max}\beta(m_{i,t}). \tag{22}$$

Substituting $b_{i,t} = \hat{c}_{i,t}$ and subtracting $c_{\max}\beta(m_{i,t})$ yields

$$\bar{c}_i - 2c_{\max}\beta(m_{i,t}) \le b_{i,t} - c_{\max}\beta(m_{i,t}) \le \bar{c}_i.$$

which proves the claim. □

Lemma 6. *Suppose that all providers use their equilibrium strategy $\boldsymbol{\sigma}^*$. On $\mathcal{E}$, if a provider $i \in \mathcal{Q} \setminus \{i^*\}$ is selected at a non-exploration step $t > N$, then*

$$\Delta_i \le 2c_{\max}\beta(m_{i,t}), \tag{23}$$

where $\Delta_i = \bar{c}_i - \bar{c}_{i^*} > 0$.

Proof. Fix a non-exploration step $t > N$ in which $i \in \mathcal{Q} \setminus \{i^*\}$ is selected. Under the event $\mathcal{E}$, Lemma 3 implies that $i^* \in \hat{\mathcal{Q}}_t$. Since t is a non-exploration step, the allocation rule therefore implies

$$b_{i,t} - c_{\max}\beta(m_{i,t}) \le b_{i^*,t} - c_{\max}\beta(m_{i^*,t}).$$

Applying Lemma 5 to both providers gives

$$\bar{c}_i - 2c_{\max}\beta(m_{i,t}) \le b_{i,t} - c_{\max}\beta(m_{i,t}) \le b_{i^*,t} - c_{\max}\beta(m_{i^*,t}) \le \bar{c}_{i^*}.$$

Hence,

$$\bar{c}_i - \bar{c}_{i^*} \le 2c_{\max}\beta(m_{i,t}).$$

Since $\Delta_i = \bar{c}_i - \bar{c}_{i^*}$, this proves equation 23. □

Lemma 7. *Suppose that all providers use their equilibrium strategy $\boldsymbol{\sigma}^*$. On $\mathcal{E}$, if a provider $i \in \mathcal{Q} \setminus \{i^*\}$ is selected at a non-exploration step $t > N$, then*

$$m_{i,t} \le L_i := \left\lceil \frac{4c_{\max}^2}{\Delta_i^2} \ln\left(\frac{57Nc_{\max}^4}{\delta\Delta_i^4}\right) \right\rceil. \tag{24}$$

Consequently, once $m_{i,t} > L_i$, provider i cannot be selected at a non-exploration step

Proof. By Lemma 6, if provider $i \in \mathcal{Q} \setminus \{i^*\}$ is selected at a non-exploration step $t > N$, then

$$\Delta_i \le 2c_{\max}\beta(m_{i,t}).$$

Dividing by $c_{\max}$ gives

$$\frac{\Delta_i}{c_{\max}} \le 2\beta(m_{i,t}).$$

Applying Lemma 2 with

$$a = \frac{\Delta_i}{c_{\max}}$$

therefore yields

$$m_{i,t} \le \frac{4c_{\max}^2}{\Delta_i^2} \ln\left(\frac{57Nc_{\max}^4}{\delta\Delta_i^4}\right) \le L_i.$$

Thus, if $m_{i,t} > L_i$, provider i cannot be selected by the allocation rule at a non-exploration step. □

B.4 Elimination of unqualified providers

We next bound the number of times an unqualified provider can be selected on the platform, including both non-exploration and exploration steps.

Lemma 8. *On $\mathcal{E}$, every unqualified provider $i \notin \mathcal{Q}$, with quality gap*

$$\varepsilon_i := q_{\min} - q_i > 0,$$

is selected at most

$$M_i := \left\lceil \frac{4}{\varepsilon_i^2} \ln\left(\frac{57N}{\delta \varepsilon_i^4}\right) \right\rceil \tag{25}$$

times in post-initialization steps.

Proof. Fix an unqualified provider $i \notin \mathcal{Q}$ and a post-initialization step $t > N$ in which $I_t = i$. Whether t is a non-exploration or exploration step, provider i can be selected only if $i \in \hat{\mathcal{Q}}_t$. Hence,

$$\hat{q}_{i,t} + \beta(m_{i,t}) \geq q_{\min}.$$

On $\mathcal{E}$, Lemma 1 gives

$$\hat{q}_{i,t} \leq q_i + \beta(m_{i,t}).$$

Combining the two inequalities yields

$$q_{\min} \leq q_i + 2\beta(m_{i,t}),$$

and therefore

$$\varepsilon_i \leq 2\beta(m_{i,t}). \tag{26}$$

Applying Lemma 2 with $a = \varepsilon_i$ gives

$$m_{i,t} \leq \frac{4}{\varepsilon_i^2} \ln\left(\frac{57N}{\delta \varepsilon_i^4}\right) \leq M_i.$$

Since this holds whenever provider i is selected in a post-initialization step, provider i is selected at most M_i times after initialization. Consequently, once $m_{i,t} > M_i$, provider i cannot be selected at any step $t > N$. □

B.5 Selection and identification bounds

We now bound the total number of times a qualified suboptimal provider can be selected.

Lemma 9. *Suppose that all providers use their equilibrium strategy $\boldsymbol{\sigma}^*$. On $\mathcal{E}$, for every provider $i \in \mathcal{Q} \setminus \{i^*\}$, the number of post-initialization steps $t \leq T$ in which $I_t = i$ is at most*

$$\max\left\{L_i,\, \lceil g(T) \rceil - 1\right\}, \tag{27}$$

where L_i is defined in Lemma 7.

Proof. Fix $T > N$ and a provider $i \in \mathcal{Q} \setminus \{i^*\}$, and consider any post-initialization step $t \leq T$ in which $I_t = i$. If t is a non-exploration step, Lemma 7 gives

$$m_{i,t} \leq L_i.$$

If instead t is an exploration step, Lemma 4 gives

$$m_{i,t} < g(t) \leq g(T).$$

Since $m_{i,t}$ is an integer,

$$m_{i,t} \leq \lceil g(T) \rceil - 1.$$

Thus, in either case,

$$m_{i,t} \leq \max\left\{L_i,\, \lceil g(T) \rceil - 1\right\}.$$

Since provider i is selected once during initialization and $m_{i,t}$ increases by one at each of its selections, its ℓ-th post-initialization selection occurs at a step t with $m_{i,t} = \ell$. Hence, provider i is selected at most $\max\{L_i, \lceil g(T) \rceil - 1\}$ times in post-initialization steps $t \leq T$. □

Proof of Theorem 3.

Proof of Theorem 3. We work under the event $\mathcal{E}$, which by Lemma 1 occurs with probability at least $1-\delta$. In the first N initialization steps, each provider is selected exactly once. Hence, at most $N-1$ initialization steps are not assigned to i^*.

We next consider the post-initialization steps. Every provider $i \neq i^*$ is either unqualified or qualified and suboptimal. By Lemma 8, each provider $i \notin \mathcal{Q}$ is selected at most

$$M_i = \left\lceil \frac{4}{\varepsilon_i^2} \ln\left(\frac{57N}{\delta \varepsilon_i^4}\right) \right\rceil = M(\varepsilon_i)$$

times after initialization.

For each provider $i \in \mathcal{Q} \setminus \{i^*\}$, Lemma 9 gives at most

$$\max\{L_i,\ \lceil g(T) \rceil - 1\}$$

post-initialization selections, where $L_i = M(\Delta_i / c_{\max})$ by Lemma 7 and, for

$$g(t) = k\left(\frac{t}{N}\right)^{\alpha},$$

we have $\lceil g(T) \rceil - 1 = \lceil k(T/N)^{\alpha} \rceil - 1$.

Since every step with $I_t \neq i^*$ selects exactly one provider other than i^*, summing these bounds and adding the $N-1$ initialization steps gives

$$\begin{aligned}
\sum_{t=1}^{T} \mathbb{1}\{I_t \neq i^*\} &\leq N - 1 + \sum_{i \notin \mathcal{Q}} M_i + \sum_{i \in \mathcal{Q} \setminus \{i^*\}} \max\{L_i,\ \lceil g(T) \rceil - 1\} \\
&= N - 1 + \sum_{i \notin \mathcal{Q}} M(\varepsilon_i) + \sum_{i \in \mathcal{Q} \setminus \{i^*\}} \max\left\{ M\left(\frac{\Delta_i}{c_{\max}}\right),\ \left\lceil k\left(\frac{T}{N}\right)^{\alpha} \right\rceil - 1 \right\} \\
&\leq N - 1 + B_{\mathrm{id}}(T),
\end{aligned}$$

which proves the claimed bound. □

Corollary 5. *Suppose that all providers use their equilibrium strategy $\boldsymbol{\sigma}^*$. With probability at least $1-\delta$,*

$$\frac{1}{T} \sum_{t=1}^{T} \mathbb{1}\{I_t \neq i^*\} \longrightarrow 0 \qquad \text{as } T \to \infty. \tag{28}$$

In particular, the platform assigns a fraction tending to one of the user's queries to the lowest-cost qualified provider i^*.

Proof. On the event $\mathcal{E}$, Theorem 3 gives

$$\sum_{t=1}^{T} \mathbb{1}\{I_t \neq i^*\} \leq N - 1 + \sum_{i \notin \mathcal{Q}} M_i + \sum_{i \in \mathcal{Q} \setminus \{i^*\}} \max\left\{ L_i,\ k\left(\frac{T}{N}\right)^{\alpha} \right\}.$$

Dividing by T yields

$$\frac{1}{T} \sum_{t=1}^{T} \mathbb{1}\{I_t \neq i^*\} \leq \frac{N - 1 + \sum_{i \notin \mathcal{Q}} M_i}{T} + \frac{1}{T} \sum_{i \in \mathcal{Q} \setminus \{i^*\}} \max\left\{ L_i,\ k\left(\frac{T}{N}\right)^{\alpha} \right\}.$$

The quantities M_i and L_i are independent of T. Moreover, since $\alpha \in (0,1)$,

$$\frac{k(T/N)^{\alpha}}{T} = \frac{k}{N^{\alpha}} T^{\alpha - 1} \xrightarrow{T \to \infty} 0.$$

Hence every term on the right-hand side converges to zero, which proves equation 28. Since $\Pr(\mathcal{E}) \geq 1 - \delta$ by Lemma 1, the result holds with probability at least $1 - \delta$. □

Corollary 6. *Suppose that all providers use their equilibrium strategy $\boldsymbol{\sigma}^*$. On $\mathcal{E}$, if at some step $t > N$,*

$$m_{i,t} > M_i \quad \textit{for every } i \notin \mathcal{Q}, \qquad \textit{and} \qquad m_{i,t} > L_i \quad \textit{for every } i \in \mathcal{Q} \setminus \{i^*\}, \tag{29}$$

then every non-exploration step $\tau \geq t$ is assigned to the optimal provider i^*.

Proof. Fix an ordinary step $\tau \geq t$. Since selection counts are non-decreasing,

$$m_{i,\tau} > M_i \quad \text{for every } i \notin \mathcal{Q}, \qquad m_{i,\tau} > L_i \quad \text{for every } i \in \mathcal{Q} \setminus \{i^*\}.$$

By Lemma 8, no unqualified provider can be selected at step τ. By Lemma 7, no qualified provider $i \in \mathcal{Q} \setminus \{i^*\}$ can be selected at non-exploration step τ. Finally, Lemma 3 guarantees that $i^* \in \hat{\mathcal{Q}}_\tau$. Hence, in a non-exploration step, the allocation rule must select i^*. □

B.6 Generic cumulative counting and confidence-radius bounds

In this section, we establish a bound on sums of confidence radii along the sequence of selected providers. This counting argument is independent of whether a step is exploratory or non-exploratory and will be used in the proof of Theorem 7.

Lemma 10 (Radius sum)**.** *Fix $T \geq N$ and let $(a_t)_{t=N+1}^T$ be nonnegative numbers satisfying*

$$a_t \leq \kappa\, \beta(m_{I_t,t}), \qquad t = N+1, \ldots, T,$$

for some $\kappa \geq 0$. Then,

$$\sum_{t=N+1}^{T} a_t \leq \kappa \sqrt{2N(T-N) \ln\!\left(\frac{2\pi^2 N T^2}{3\delta}\right)}. \tag{30}$$

Proof. For each provider $i \in [N]$, let

$$D_i = m_{i,T+1} - 1$$

denote the number of times provider i is selected after initialization up to and including time step T. Since exactly one provider is selected at every post-initialization step,

$$\sum_{i=1}^{N} D_i = T - N.$$

Therefore,

$$\begin{aligned}\sum_{t=N+1}^{T} a_t &\leq \kappa \sum_{i=1}^{N} \sum_{\ell=1}^{D_i} \beta(\ell) \\ &= \kappa \sum_{i=1}^{N} \sum_{\ell=1}^{D_i} \sqrt{\frac{1}{2\ell} \ln\!\left(\frac{2\pi^2 N \ell^2}{3\delta}\right)}.\end{aligned}$$

Since $\ell \leq D_i \leq T$,

$$\beta(\ell) \leq \sqrt{\frac{1}{2\ell} \ln\!\left(\frac{2\pi^2 N T^2}{3\delta}\right)}.$$

Hence,

$$\begin{aligned}\sum_{t=N+1}^{T} a_t &\leq \kappa \sqrt{\frac{1}{2} \ln\!\left(\frac{2\pi^2 N T^2}{3\delta}\right)} \sum_{i=1}^{N} \sum_{\ell=1}^{D_i} \frac{1}{\sqrt{\ell}} \\ &\leq \kappa \sqrt{2 \ln\!\left(\frac{2\pi^2 N T^2}{3\delta}\right)} \sum_{i=1}^{N} \sqrt{D_i},\end{aligned}$$

where we used $\sum_{\ell=1}^{D} \ell^{-1/2} \le 2\sqrt{D}$. Finally, using the Cauchy–Schwarz inequality,

$$\sum_{i=1}^{N} \sqrt{D_i} \le \sqrt{N \sum_{i=1}^{N} D_i} = \sqrt{N(T-N)}.$$

Combining the preceding inequalities gives

$$\sum_{t=N+1}^{T} a_t \le \kappa \sqrt{2N(T-N)\ln\left(\frac{2\pi^2 NT^2}{3\delta}\right)},$$

as claimed. □

B.7 User performance guarantees

We now derive the quality guarantees that the user obtains on the platform.

Theorem 7. *On the event $\mathcal{E}$, the sum*

$$R_T^{\mathrm{qual}} = \sum_{t=1}^{T} \big(q_{\min} - q_{I_t}\big)^+,$$

which measures the cumulative quality under $q_{\min}$ obtained by the user, satisfies,

$$R_T^{\mathrm{qual}} \le \min\left\{ N + 2\sqrt{2N(T-N)\ln\left(\frac{2\pi^2 NT^2}{3\delta}\right)},\ \sum_{i \notin \mathcal{Q}} (M_i + 1)\varepsilon_i \right\}. \tag{31}$$

Moreover, the number of queries assigned to unqualified providers satisfies,

$$\sum_{t=1}^{T} \mathbb{1}\{I_t \notin \mathcal{Q}\} \le \sum_{i \notin \mathcal{Q}} (M_i + 1). \tag{32}$$

Consequently,

$$\frac{R_T^{\mathrm{qual}}}{T} = O\left(\frac{1}{T}\right) \longrightarrow 0, \qquad \frac{1}{T}\sum_{t=1}^{T} \mathbb{1}\{I_t \notin \mathcal{Q}\} = O\left(\frac{1}{T}\right) \longrightarrow 0 \tag{33}$$

Proof. In the first N initialization steps, every provider is selected once. Since $(q_{\min} - q_i)^+ \le 1$, the contribution of the first N steps to the sum R_T^{qual} is at most N.

Consider now a post-initialization step $t > N$. If $I_t \in \mathcal{Q}$, then $q_{I_t} \ge q_{\min}$ and the step does not contribute to the sum R_T^{qual}. If $I_t \notin \mathcal{Q}$, then the selected provider must belong to $\hat{\mathcal{Q}}_t$. In the proof of Lemma 8 (Eq. 26), we showed that this implies that:

$$\varepsilon_{I_t} = \big(q_{\min} - q_{I_t}\big)^+ \le 2\beta(m_{I_t,t}).$$

Applying Lemma 10 with $\kappa = 2$ therefore gives

$$\sum_{t=N+1}^{T} \big(q_{\min} - q_{I_t}\big)^+ \le 2\sqrt{2N(T-N)\ln\left(\frac{2\pi^2 NT^2}{3\delta}\right)}.$$

Adding the initialization contribution proves the first bound in equation 31.

Alternatively, Lemma 8, the elimination bound for unqualified providers, shows that every $i \notin \mathcal{Q}$ is selected at most M_i times after initialization. Including its initialization selection, provider i therefore contributes at most

$$(M_i + 1)\varepsilon_i$$

to the cumulative quality regret. Summing over all $i \notin \mathcal{Q}$ proves the second bound in equation 31.

The same selection bound immediately gives

$$\sum_{t=1}^{T} \mathbb{1}\{I_t \notin \mathcal{Q}\} \le \sum_{i \notin \mathcal{Q}} (M_i + 1).$$

proving equation 32.

Finally, the right-hand sides of the second bound in equation 31 and of equation 32 are independent of T. Dividing both inequalities by T therefore gives equation 33. □

We next bound the cumulative excess mean generation cost of the selected providers relative to the optimal provider i^*, i.e., how much more the user's queries cost to serve than they would if every query were served by i^*. We refer to this quantity as the generation-cost regret.

Theorem 8. *Suppose that all providers use their equilibrium strategy $\boldsymbol{\sigma}^*$. On the event $\mathcal{E}$, the sum*

$$R_T^{\mathrm{gen}} = \sum_{t=1}^{T} \big(\bar{c}_{I_t} - \bar{c}_{i^*}\big)^+$$

satisfies

$$R_T^{\mathrm{gen}} \le \min\Bigg\{ (N-1)\lceil g(T)\rceil c_{\max} + 2c_{\max}\sqrt{2N(T-N)\ln\!\left(\frac{2\pi^2 NT^2}{3\delta}\right)},$$
$$\sum_{i \notin \mathcal{Q}} (M_i + 1)\big(\bar{c}_i - \bar{c}_{i^*}\big)^+ + \sum_{i \in \mathcal{Q}\setminus\{i^*\}} \big(\max\{L_i, \lceil g(T)\rceil - 1\} + 1\big)\Delta_i \Bigg\}. \tag{34}$$

Proof. Since every mean cost lies in $[0, c_{\max}]$, each term of R_T^{gen} is at most $c_{\max}$, and it vanishes whenever $I_t = i^*$.

For the first bound, we partition the time steps up to T into initialization steps, exploration steps, and non-exploration steps. During initialization, at most $N-1$ steps are not assigned to i^*, which contribute at most $(N-1)c_{\max}$. By Lemma 4, every provider is selected in at most $\lceil g(T)\rceil - 1$ exploration steps, so the exploration steps assigned to a provider other than i^* contribute at most $(N-1)(\lceil g(T)\rceil - 1)c_{\max}$. Finally, consider a non-exploration step $t > N$ with $I_t = i \neq i^*$. On $\mathcal{E}$, Lemma 3 gives $i^* \in \hat{\mathcal{Q}}_t$, so the allocation rule (Eq. 3) implies

$$b_{i,t} - c_{\max}\beta(m_{i,t}) \le b_{i^*,t} - c_{\max}\beta(m_{i^*,t}),$$

and Lemma 5, applied to both providers, yields $\bar{c}_i - 2c_{\max}\beta(m_{i,t}) \le \bar{c}_{i^*}$. Hence, letting $a_t = (\bar{c}_{I_t} - \bar{c}_{i^*})^+$ on non-exploration steps and $a_t = 0$ on exploration steps, we have $a_t \le 2c_{\max}\beta(m_{I_t,t})$ for every $t = N+1, \dots, T$, and Lemma 10 with $\kappa = 2c_{\max}$ gives

$$\sum_{t=N+1}^{T} a_t \le 2c_{\max}\sqrt{2N(T-N)\ln\!\left(\frac{2\pi^2 NT^2}{3\delta}\right)}.$$

Summing the three contributions proves the first bound.

For the second bound, Lemma 8 shows that every $i \notin \mathcal{Q}$ is selected at most M_i times after initialization, and Lemma 9 shows that every $i \in \mathcal{Q}\setminus\{i^*\}$ is selected at most $\max\{L_i, \lceil g(T)\rceil - 1\}$ times after initialization. Including its initialization selection, each provider $i \neq i^*$ therefore contributes at most its number of selections times $(\bar{c}_i - \bar{c}_{i^*})^+$ to R_T^{gen}, where $(\bar{c}_i - \bar{c}_{i^*})^+ = \Delta_i$ for $i \in \mathcal{Q}\setminus\{i^*\}$. Summing over $i \neq i^*$ proves the second bound. □

B.8 Payments on exploration and non-exploration steps

We next bound the payments made by the user. On an exploration step, the selected provider is paid $c_{\max}$, whereas on an non-exploration step the payment is determined by the adjusted bid of the cheapest competing eligible provider.

Lemma 11. *On $\mathcal{E}$, suppose that all providers use their equilibrium strategy $\boldsymbol{\sigma}^*$. For every non-exploration step $t > N$ with $I_t = i$,*

$$\pi_t \le \min\left\{c_{\max}, \bar{c}_{(2)} + c_{\max}\beta(m_{i,t})\right\}. \tag{35}$$

Consequently,

$$\left(\pi_t - \bar{c}_{(2)}\right)^+ \le c_{\max}\beta(m_{i,t}). \tag{36}$$

Proof. Fix a non-exploration step $t > N$ with $I_t = i$. Then, the payment rule (Eq. 5) gives

$$\pi_t = \min\left\{c_{\max}\beta(m_{i,t}) + \min_{j\in\hat{\mathcal{Q}}_t\setminus\{i\}}\left(b_{j,t} - c_{\max}\beta(m_{j,t})\right), c_{\max}\right\}.$$

If $i = i^*$, let j be a second-cheapest qualified provider, so that $\bar{c}_j = \bar{c}_{(2)}$. If $i \ne i^*$, let $j = i^*$, in which case $\bar{c}_j = \bar{c}_{i^*} \le \bar{c}_{(2)}$. In either case, $j \in \mathcal{Q}$ and $j \ne i$ and $j \in \hat{\mathcal{Q}}_t$ on the event $\mathcal{E}$. Moreover, using Lemma 5,

$$b_{j,t} - c_{\max}\beta(m_{j,t}) \le \bar{c}_j \le \bar{c}_{(2)}.$$

Therefore,

$$\min_{j'\in\hat{\mathcal{Q}}_t\setminus\{i\}}\left(b_{j',t} - c_{\max}\beta(m_{j',t})\right) \le \bar{c}_{(2)},$$

and hence

$$\pi_t \le \min\left\{c_{\max}, \bar{c}_{(2)} + c_{\max}\beta(m_{i,t})\right\}.$$

The second claim follows immediately. □

Theorem 9. *On $\mathcal{E}$, when all providers use their equilibrium strategy $\boldsymbol{\sigma}^*$,*

$$\sum_{t=1}^{T}\left(\pi_t - \bar{c}_{(2)}\right)^+ \le \left[N + N\left(\lceil g(T)\rceil - 1\right)\right]\left(c_{\max} - \bar{c}_{(2)}\right) + c_{\max}\sqrt{2N(T-N)\ln\left(\frac{2\pi^2 NT^2}{3\delta}\right)}. \tag{37}$$

Proof. Partition the time steps up to T into initialization steps, exploration steps, and non-exploration steps.

During initialization, each of the N providers is paid $c_{\max}$, so the total excess payment contributed by these steps is

$$N\left(c_{\max} - \bar{c}_{(2)}\right).$$

On every exploration step, the selected provider is also paid $c_{\max}$. By Lemma 4, there are at most $N(\lceil g(T)\rceil - 1)$ such steps. Their total contribution is therefore at most

$$N\left(\lceil g(T)\rceil - 1\right)\left(c_{\max} - \bar{c}_{(2)}\right).$$

Finally, on every non-exploration step $t > N$, Lemma 11 gives

$$\left(\pi_t - \bar{c}_{(2)}\right)^+ \le c_{\max}\beta(m_{I_t,t}).$$

Hence, for the non-exploration steps, applying Lemma 10,

$$\sum_{\substack{t=N+1\\ t\text{non-exploration}}}^{T}\left(\pi_t - \bar{c}_{(2)}\right)^+ \le c_{\max}\sum_{t=N+1}^{T}\beta(m_{I_t,t}) \le c_{\max}\sqrt{2N(T-N)\ln\left(\frac{2\pi^2 NT^2}{3\delta}\right)}.$$

Summing the three contributions proves equation 37. □

Corollary 10. *Suppose*

$$g(t) = k\left(\frac{t}{N}\right)^{\alpha}, \qquad k > 0, \quad \alpha \in (0,1).$$

On $\mathcal{E}$, when all providers use their equilibrium strategy $\boldsymbol{\sigma}^*$,

$$\frac{1}{T}\sum_{t=1}^{T}\left(\pi_t - \bar{c}_{(2)}\right)^+ \le \left(\frac{N}{T} + k\left(\frac{N}{T}\right)^{1-\alpha}\right)\left(c_{\max} - \bar{c}_{(2)}\right) + c_{\max}\sqrt{\frac{2N}{T}\ln\left(\frac{2\pi^2 NT^2}{3\delta}\right)}. \tag{38}$$

Proof. Since

$$g(T) = k\left(\frac{T}{N}\right)^{\alpha},$$

the forced exploration term in Theorem 9 satisfies

$$N\left(\lceil g(T)\rceil - 1\right) \le Ng(T) = kN\left(\frac{T}{N}\right)^{\alpha} \tag{39}$$

Substituting this into equation 37 and dividing by T gives

$$\frac{Ng(T)}{T} = k\left(\frac{N}{T}\right)^{1-\alpha}$$

which proves the result. □

B.9 Payment convergence

We finally show that, on non-exploration steps, the payment converges to $\bar{c}_{(2)}$. Throughout, we write

$$\bar{M} = \max\{1, \max_{i\notin\mathcal{Q}} M_i\}, \qquad T_0 = N\max\{1, (\bar{M}/k)^{1/\alpha}\}$$

so that every time step $t > T_0$ comes after initialization and satisfies $g(t) > \bar{M}$

Lemma 12. *For every real $x \ge 1$, let*

$$\beta(x) = \sqrt{\frac{1}{2x}\ln\left(\frac{2\pi^2 Nx^2}{3\delta}\right)}. \tag{40}$$

Then β is decreasing on $[1,\infty)$. Consequently, $\beta(m) \le \beta(x)$ for every $m \ge x \ge 1$.

Proof. Let $C = 2\pi^2 N/(3\delta)$ and $h(x) = \beta(x)^2 = \ln(Cx^2)/(2x) \ge 0$. Since $\mathcal{Q} \subseteq [N]$ contains at least two providers, $N \ge 2$. Since moreover $\delta < 1$, for every $x \ge 1$,

$$Cx^2 \ge C > \frac{4\pi^2}{3} > e^2.$$

Differentiating gives

$$h'(x) = \frac{2 - \ln(Cx^2)}{2x^2} < 0,$$

so h is decreasing on $[1,\infty)$, and so is $\beta = \sqrt{h}$. □

Lemma 13. *On $\mathcal{E}$, for every unqualified provider $i \notin \mathcal{Q}$ and every step $t > N$, if $i \in \hat{\mathcal{Q}}_t$, then*

$$m_{i,t} \le M_i.$$

Proof. Fix $i \notin \mathcal{Q}$ and a step $t > N$ with $i \in \hat{\mathcal{Q}}_t$. By the definition of estimated-qualified set and Lemma 1,

$$q_{\min} \le \hat{q}_{i,t} + \beta(m_{i,t}) \le q_i + 2\beta(m_{i,t}),$$

and therefore $\varepsilon_i \le 2\beta(m_{i,t})$. Applying Lemma 2 with $a = \varepsilon_i$ gives

$$m_{i,t} \le \frac{4}{\varepsilon_i^2} \ln\left(\frac{57N}{\delta\varepsilon_i^4}\right) \le M_i,$$

which proves the claim. □

Lemma 14. *For every non-exploration step $t > N$,*

$$m_{j,t} \ge g(t) \qquad \textit{for every } j \in \hat{\mathcal{Q}}_t. \tag{41}$$

If, moreover, $\mathcal{E}$ holds and $t > T_0$, then $\hat{\mathcal{Q}}_t = \mathcal{Q}$.

Proof. Fix an ordinary step $t > N$, so that $\mathcal{S}_t = \emptyset$. By the definition of $\mathcal{S}_t$, no provider $j \in \hat{\mathcal{Q}}_t$ satisfies $m_{j,t} < g(t)$, which proves equation 41.

Now suppose that $\mathcal{E}$ holds and $t > T_0$, so that $g(t) > \bar{M}$. By Lemma 3, $\mathcal{Q} \subseteq \hat{\mathcal{Q}}_t$. Conversely, suppose that some unqualified provider $i \notin \mathcal{Q}$ belongs to $\hat{\mathcal{Q}}_t$. Then equation 41 gives

$$m_{i,t} \ge g(t) > \bar{M} \ge M_i,$$

whereas Lemma 13 gives $m_{i,t} \le M_i$, a contradiction. Hence, $\hat{\mathcal{Q}}_t \subseteq \mathcal{Q}$, and therefore $\hat{\mathcal{Q}}_t = \mathcal{Q}$. □

Lemma 15. *For every step $t \ge 1$, $0 \le \pi_t \le c_{\max}$. Consequently,*

$$\left|\pi_t - \bar{c}_{(2)}\right| \le c_{\max}, \qquad t \ge 1. \tag{42}$$

Proof. By the payment rule, $\pi_t \le c_{\max}$, with equality during initialization, on exploration steps, and whenever $\hat{\mathcal{Q}}_t = \emptyset$. On an ordinary round, Proposition 1 implies that the selected provider bids $b_{I_t,t} \le P_{I_t,t}$, and therefore

$$\pi_t = \min\{P_{I_t,t}, c_{\max}\} \ge \min\{b_{I_t,t}, c_{\max}\} = b_{I_t,t} \ge 0.$$

Since also $\bar{c}_{(2)} \in [0, c_{\max}]$, equation 42 follows. □

Lemma 16. *Suppose that all providers use their equilibrium strategy $\boldsymbol{\sigma}^*$. On $\mathcal{E}$, for every non-exploration step $t > T_0$,*

$$\bar{c}_{(2)} - 3c_{\max}\beta(g(t)) \le \pi_t \le \bar{c}_{(2)} + c_{\max}\beta(g(t)). \tag{43}$$

Proof. Fix a non-exploration step $t > T_0$ and let $i = I_t$. By Lemma 14, $\hat{\mathcal{Q}}_t = \mathcal{Q}$ and $m_{j,t} \ge g(t)$ for every $j \in \mathcal{Q}$. Moreover, $g(t) > \bar{M} \ge 1$, so Lemma 12 gives

$$\beta(m_{j,t}) \le \beta(g(t)), \qquad j \in \mathcal{Q}. \tag{44}$$

We start with the upper bound. Lemma 11 and Eq. 44 give

$$\pi_t \le \bar{c}_{(2)} + c_{\max}\beta(m_{i,t}) \le \bar{c}_{(2)} + c_{\max}\beta(g(t)).$$

We next turn to the lower bound. Since $\hat{\mathcal{Q}}_t = \mathcal{Q}$, the selected provider satisfies $i \in \mathcal{Q}$, and the payment rule gives

$$\pi_t = \min\left\{c_{\max}\beta(m_{i,t}) + \min_{j \in \mathcal{Q}\setminus\{i\}}\left(b_{j,t} - c_{\max}\beta(m_{j,t})\right),\ c_{\max}\right\}.$$

For every $j \in \mathcal{Q} \setminus \{i\}$, Lemma 5 and Eq. 44 give

$$b_{j,t} - c_{\max}\beta(m_{j,t}) \ge \bar{c}_j - 2c_{\max}\beta(g(t)). \tag{45}$$

If $i = i^*$, then $\bar{c}_j \geq \bar{c}_{(2)}$ for every $j \in \mathcal{Q} \setminus \{i^*\}$, and Eq. 45 yields

$$c_{\max}\beta(m_{i^*,t}) + \min_{j \in \mathcal{Q}\setminus\{i^*\}} \big(b_{j,t} - c_{\max}\beta(m_{j,t})\big) \geq \bar{c}_{(2)} - 2c_{\max}\beta(g(t)) \geq \bar{c}_{(2)} - 3c_{\max}\beta(g(t)).$$

If instead $i \neq i^*$, then $i^* \in \mathcal{Q} \setminus \{i\}$ and $\bar{c}_j \geq \bar{c}_{i^*}$ for every $j \in \mathcal{Q} \setminus \{i\}$, so Eq. 45 yields

$$\min_{j \in \mathcal{Q}\setminus\{i\}} \big(b_{j,t} - c_{\max}\beta(m_{j,t})\big) \geq \bar{c}_{i^*} - 2c_{\max}\beta(g(t)).$$

Further, Lemma 6 gives $\bar{c}_i - \bar{c}_{i^*} \leq 2c_{\max}\beta(m_{i,t})$, and since $\bar{c}_i \geq \bar{c}_{(2)}$,

$$\bar{c}_{i^*} \geq \bar{c}_{(2)} - 2c_{\max}\beta(m_{i,t}).$$

Therefore,

$$\begin{aligned} c_{\max}\beta(m_{i,t}) + \min_{j \in \mathcal{Q}\setminus\{i\}} \big(b_{j,t} - c_{\max}\beta(m_{j,t})\big) &\geq \bar{c}_{(2)} - c_{\max}\beta(m_{i,t}) - 2c_{\max}\beta(g(t)) \\ &\geq \bar{c}_{(2)} - 3c_{\max}\beta(g(t)), \end{aligned}$$

where the last inequality uses Eq. 44. In both cases, the first argument of the minimum defining π_t is at least $\bar{c}_{(2)} - 3c_{\max}\beta(g(t))$, and, since $\bar{c}_{(2)} \leq c_{\max}$, so is the second. Hence,

$$\pi_t \geq \bar{c}_{(2)} - 3c_{\max}\beta(g(t)),$$

which completes the proof. □

Theorem 11. *When all providers choose their equilibrium strategy $\boldsymbol{\sigma}^*$, with probability at least $1 - \delta$, for every time step $t > T_0$ that is not an exploration step,* i.e., *$\mathcal{S}_t = \emptyset$, the payment satisfies*

$$\bar{c}_{(2)} - 3\,c_{\max}\,\beta(g(t)) \;\leq\; \pi_t \;\leq\; \bar{c}_{(2)} + c_{\max}\,\beta(g(t)). \tag{46}$$

Hence, $\pi_t \to \bar{c}_{(2)}$ as $t \to \infty$ along these time steps. Moreover, for every $T > T_0$,

$$\begin{aligned} \frac{1}{T}\sum_{t=1}^{T} \big|\pi_t - \bar{c}_{(2)}\big| \leq \frac{T_0}{T} c_{\max} + k\left(\frac{N}{T}\right)^{1-\alpha} \big(c_{\max} - \bar{c}_{(2)}\big) \\ + \frac{6\,c_{\max}}{2-\alpha}\sqrt{\frac{1}{2k}\left(\frac{N}{T}\right)^{\alpha} \ln\left(\frac{2\pi^2 N k^2}{3\delta}\left(\frac{T}{N}\right)^{2\alpha}\right)}. \end{aligned} \tag{47}$$

Proof. We work under the event $\mathcal{E}$, which by Lemma 1 occurs with probability at least $1 - \delta$.

The first claim, Eq.46, is Lemma 16. Since $g(t) \to \infty$ and $\beta(x) \to 0$ as $x \to \infty$, we have $\beta(g(t)) \to 0$, and hence $\pi_t \to \bar{c}_{(2)}$ in non-exploration steps.

For the second claim, fix $T > T_0$. In $\mathcal{E}$, every post-initialization step is either an exploration step or a non-exploration step, so we partition the steps $t \leq T$ into the three sets

$$\mathcal{R}_1 = \{t \leq T_0\}, \qquad \mathcal{R}_2 = \{T_0 < t \leq T : \mathcal{S}_t \neq \emptyset\}, \qquad \mathcal{R}_3 = \{T_0 < t \leq T : \mathcal{S}_t = \emptyset\}.$$

The set $\mathcal{R}_1$ contains at most T_0 steps, each of which contributes at most $c_{\max}$ to the sum in the left-hand side of Eq. 47 by Lemma 15. Hence,

$$\sum_{t \in \mathcal{R}_1} \big|\pi_t - \bar{c}_{(2)}\big| \leq T_0\, c_{\max}.$$

On every step in $\mathcal{R}_2$, the selected provider is paid $c_{\max}$, so the step contributes $c_{\max} - \bar{c}_{(2)}$. By Lemma 4, there are at most $N(\lceil g(T) \rceil - 1) \leq N g(T) = kN(T/N)^{\alpha}$ exploration steps. Hence,

$$\sum_{t \in \mathcal{R}_2} \big|\pi_t - \bar{c}_{(2)}\big| \leq kN\left(\frac{T}{N}\right)^{\alpha} \big(c_{\max} - \bar{c}_{(2)}\big).$$

Finally, for every $t \in \mathcal{R}_3$, Lemma 16 gives

$$\big|\pi_t - \bar{c}_{(2)}\big| \leq 3c_{\max}\beta(g(t)).$$

Since $1 \le \bar{M} < g(t) \le g(T)$ and $1/g(t) = (N/t)^\alpha/k$,

$$\beta(g(t)) = \sqrt{\frac{1}{2g(t)} \ln\!\left(\frac{2\pi^2 N g(t)^2}{3\delta}\right)} \le \sqrt{\frac{1}{2k} \ln\!\left(\frac{2\pi^2 N k^2}{3\delta} \left(\frac{T}{N}\right)^{2\alpha}\right)} \left(\frac{N}{t}\right)^{\alpha/2}.$$

Moreover,

$$\sum_{t=1}^{T} t^{-\alpha/2} \le \int_0^T x^{-\alpha/2}\, dx = \frac{2}{2-\alpha}\, T^{1-\alpha/2}.$$

Combining the two preceding inequalities gives

$$\sum_{t \in \mathcal{R}_3} \left|\pi_t - \bar{c}_{(2)}\right| \le \frac{6 c_{\max}}{2-\alpha} \sqrt{\frac{N^\alpha}{2k} \ln\!\left(\frac{2\pi^2 N k^2}{3\delta} \left(\frac{T}{N}\right)^{2\alpha}\right)}\, T^{1-\alpha/2}.$$

Summing the contributions of $\mathcal{R}_1$, $\mathcal{R}_2$, and $\mathcal{R}_3$ and dividing by T, together with

$$\frac{kN(T/N)^\alpha}{T} = k\left(\frac{N}{T}\right)^{1-\alpha}, \qquad \frac{N^{\alpha/2} T^{1-\alpha/2}}{T} = \left(\frac{N}{T}\right)^{\alpha/2},$$

proves equation 47. □

Proof of Theorem 4. Theorem 4 is the second claim of Theorem 11.

Corollary 12. *Under the assumptions of Theorem 11, for N, k, δ fixed and $T \to \infty$, the right-hand side of Eq. 47 is*

$$O\Big(T^{-(1-\alpha)}\Big) + O\Big(T^{-\alpha/2}\sqrt{\ln T}\Big) = O\Big(T^{-\min\{1-\alpha,\, \alpha/2\}}\sqrt{\ln T}\Big).$$

The exponent $\min\{1-\alpha, \alpha/2\}$ is maximized over $\alpha \in (0,1)$ at $\alpha = 2/3$, for which the bound is $O\Big(T^{-1/3}\sqrt{\ln T}\Big)$.

Proof. The first term of Eq. 47 is $O(1/T)$ and the second is $O(T^{-(1-\alpha)})$; since $\alpha > 0$, the first is dominated by the second. In the third term, $\ln\!\big(\frac{2\pi^2 N k^2}{3\delta}(T/N)^{2\alpha}\big) = O(\ln T)$, so the term is $O(T^{-\alpha/2}\sqrt{\ln T})$. Finally, $1-\alpha$ is decreasing and $\alpha/2$ is increasing in α, so $\min\{1-\alpha, \alpha/2\}$ is maximized where they coincide, $1-\alpha = \alpha/2$, i.e., at $\alpha = 2/3$, where both equal $1/3$. □

B.10 Robust optimality of the truthful strategy

Throughout, we fix a provider $i \in [N]$, the strategies $\boldsymbol{\sigma}_{-i}$ of the other providers, and a history H_t of the platform before step $t > N$, containing the observed history $H_{i,t}$ of provider i. For every step $s \ge t$, we write

$$G_{i,s} = \big(\pi_s - \hat{c}_{i,s}\big)\, \mathbb{1}\{I_s = i\}$$

for the net payment of provider i at step s, so that $U_{i|H_t}(\sigma_i, \boldsymbol{\sigma}_{-i}; \boldsymbol{\theta}_{-i}, T) = \mathbb{E}[\sum_{s=t}^{T} G_{i,s} \,|\, H_t]$ by Eq. 7.

The proof rests on a single observation: once the history before a step is fixed, the net payment of provider i at that step depends on its own bid only through whether it is assigned the query, and bidding the cost estimate is the best choice.

Lemma 17. *Fix a step $s > N$ and a history H_s of the platform before step s, containing the observed history $H_{i,s}$. For $b \in [0, c_{\max}]$, let $G_{i,s}(b)$ denote the net payment of provider i at step s if it bids $b_{i,s} = b$. Then, for every $b \in [0, c_{\max}]$,*

a) $G_{i,s}(\hat{c}_{i,s}) \ge 0$;

b) $G_{i,s}(\hat{c}_{i,s}) \ge G_{i,s}(b)$;

c) if $\mathcal{S}_s = \emptyset$, $i \in \hat{\mathcal{Q}}_s$, and $(\hat{c}_{i,s} - P_{i,s})(b - P_{i,s}) < 0$, then $G_{i,s}(b) < G_{i,s}(\hat{c}_{i,s})$.

Proof. If $\hat{\mathcal{Q}}_s = \emptyset$ or $\mathcal{S}_s \neq \emptyset$, the allocation does not depend on the bids and the selected provider is paid $c_{\max}$. Hence $G_{i,s}(b) = (c_{\max} - \hat{c}_{i,s})\ \mathbb{1}\{I_s = i\} \geq 0$ for every b, which proves a) and b), while c) does not apply. If $\mathcal{S}_s = \emptyset$ and $i \notin \hat{\mathcal{Q}}$, provider i is not selected whatever it bids, so $G_{i,s}(b) = 0$ for every b and all three claims hold.

It remains to consider $\mathcal{S}_s = \emptyset$ and $i \in \hat{\mathcal{Q}}_s$. By Proposition 1, provider i is selected if $b < P_{i,s}$. Whenever it is selected, the payment rule (Eq. 5) pays it $\pi_s = \min\{P_{i,s}, c_{\max}\}$, which does not depend on b. Therefore,

$$G_{i,s}(b) = \big(\min\{P_{i,s}, c_{\max}\} - \hat{c}_{i,s}\big)\ \mathbb{1}\{I_s = i\} \in \big\{0,\ \min\{P_{i,s}, c_{\max}\} - \hat{c}_{i,s}\big\}. \tag{48}$$

We first compute $G_{i,s}(\hat{c}_{i,s})$. If $\hat{c}_{i,s} < P_{i,s}$, provider i is selected and, since $\hat{c}_{i,s} \leq c_{\max}$, $\min\{P_{i,s}, c_{\max}\} \geq \min\{\hat{c}_{i,s}, c_{\max}\} = \hat{c}_{i,s}$, so $G_{i,s}(\hat{c}_{i,s}) = \min\{P_{i,s}, c_{\max}\} - \hat{c}_{i,s} \geq 0$. If $\hat{c}_{i,s} > P_{i,s}$, provider i is not selected and $G_{i,s}(\hat{c}_{i,s}) = 0$, while $\min\{P_{i,s}, c_{\max}\} - \hat{c}_{i,s} \leq P_{i,s} - \hat{c}_{i,s} < 0$. If $\hat{c}_{i,s} = P_{i,s}$, the factor $\min\{P_{i,s}, c_{\max}\} - \hat{c}_{i,s}$ in Eq. 48 equals $\min\{\hat{c}_{i,s}, c_{\max}\} - \hat{c}_{i,s} = 0$, so $G_{i,s}(\hat{c}_{i,s}) = 0$ whether or not i is selected. In all three cases,

$$G_{i,s}(\hat{c}_{i,s}) = \big(\min\{P_{i,s}, c_{\max}\} - \hat{c}_{i,s}\big)^+ \geq 0 \tag{49}$$

which proves a). Claim b) follows since, by Eq. 48, $G_{i,s}(b)$ is either 0 or $\min\{P_{i,s}, c_{\max}\} - \hat{c}_{i,s}$, both of which are at most the right-hand side of Eq. 49.

For c), suppose first that $\hat{c}_{i,s} < P_{i,s} < b$. Then bidding $\hat{c}_{i,s}$ gets the query while bidding b does not, so $G_{i,s}(b) = 0$. Moreover, $\hat{c}_{i,s} < b \leq c_{\max}$ and $\hat{c}_{i,s} < P_{i,s}$ give $\min\{P_{i,s}, c_{\max}\} > \hat{c}_{i,s}$, so $G_{i,s}(\hat{c}_{i,s}) > 0 = G_{i,s}(b)$ by Eq. 49. Suppose instead that $b < P_{i,s} < \hat{c}_{i,s}$. Then bidding b gets the query while bidding $\hat{c}_{i,s}$ does not, so $G_{i,s}(\hat{c}_{i,s}) = 0$, whereas $G_{i,s}(b) = \min\{P_{i,s}, c_{\max}\} - \hat{c}_{i,s} \leq P_{i,s} - \hat{c}_{i,s} < 0$ by Eq. 48. □

Proof of Theorem 2.

Proof of Theorem 2. To lighten the notation, we omit the fixed arguments $\boldsymbol{\sigma}_{-i}$ and $\boldsymbol{\theta}_{-i}$, and write $U_{i|H_t}(\sigma_i; T)$. By Eq. 8, $\sigma_i^*(H_{i,s}) = \hat{c}_{i,s}$ at every step s.

Claim i). Under σ_i^*, provider i bids $\hat{c}_{i,s}$ at every step $s \geq t$, so Lemma 17a), applied at step s to the realized history before it, gives $G_{i,s} \geq 0$. Summing over $s = t, \ldots, T$ and taking the expectation conditional on H_t yields $U_{i|H_t}(\sigma_i^*; T) \geq 0$ for every $T \geq t$ and every $\boldsymbol{\theta}_{-i}$.

Claim ii). By Lemma 17a), every term $G_{i,s}$ is non-negative under σ_i^*, so $U_{i|H_t}(\sigma_i^*; T)$ is non-decreasing in T, and its infimum over $T \geq t$ is attained at $T = t$. Conditional on H_t, all standing bids, selection counts, and estimates entering the allocation and payment rules at step t are fixed, so the net payment at step t does not depend further on $\boldsymbol{\theta}_{-i}$, and $U_{i|H_t}(\sigma_i^*; t) = \mathbb{E}[G_{i,t}(\hat{c}_{i,t}) \,|\, H_t]$ for every $\boldsymbol{\theta}_{-i}$. Hence,

$$\inf_{T \geq t,\, \boldsymbol{\theta}_{-i}} U_{i|H_t}(\sigma_i^*; T) = \mathbb{E}\big[G_{i,t}(\hat{c}_{i,t}) \,|\, H_t\big]. \tag{50}$$

Now fix any strategy σ_i and let $b = \sigma_i(H_{i,t})$ be its bid at step t. Taking $T = t$ in the infimum and then applying Lemma 17b) for every realization of the platform's draw at step t,

$$\inf_{T \geq t,\, \boldsymbol{\theta}_{-i}} U_{i|H_t}(\sigma_i; T) \leq U_{i|H_t}(\sigma_i; t) = \mathbb{E}\big[G_{i,t}(b) \,|\, H_t\big] \leq \mathbb{E}\big[G_{i,t}(\hat{c}_{i,t}) \,|\, H_t\big]. \tag{51}$$

Combining Eqs. 50 and 51 proves claim ii).

Claim iii). Under the assumptions of claim iii), step t is not an exploration step, so $G_{i,t}(b)$ and $G_{i,t}(\hat{c}_{i,t})$ are determined by H_t, and Lemma 17c) gives $G_{i,t}(b) < G_{i,t}(\hat{c}_{i,t})$. The last inequality in Eq. 51 is therefore strict, and so is the inequality in claim ii). □

C BAYESIAN COST ESTIMATION

The analysis in Appendix B considers providers that bid the empirical mean $\hat{c}_{i,t}$ of their observed generation costs. In this section, we show that the robust optimality of bidding one's current cost estimate (Theorem 2) holds verbatim for any estimator, and that the user-facing guarantees of Section 4 extend to Bayesian estimators whose prior is eventually dominated by the data, with the cost-estimation radius widened by a factor $1+\gamma$ and unchanged asymptotic rates. No proof of Appendix B is repeated; instead, we identify the properties of the bids that those proofs use and verify them for Bayesian estimates.

Estimation model. Each provider $i \in [N]$ forms its cost estimate through a map e_i from histories to $[0, c_{\max}]$, and we write $e_{i,t} = e_i(H_{i,t})$ for its estimate before step t. Since $H_{i,t}$ consists of the first $m_{i,t}$ observations of provider i, we also write $e_{i,m}$ for the estimate after m selections, so that $e_{i,t} = e_{i,m_{i,t}}$, in the same way that $\hat{q}_{i,t} = \hat{q}_{i,m_{i,t}}$ and $\hat{c}_{i,t} = \hat{c}_{i,m_{i,t}}$. For a Bayesian provider, $e_{i,m}$ is the posterior mean of the cost of serving its next query given a prior and its m observed costs; it lies in $[0, c_{\max}]$ because the cost does. We require that the prior does not pull the estimate away from the empirical mean by more than a fixed multiple of the confidence radius.

Assumption 1. *There exists a constant $\gamma \geq 0$ such that, for every provider $i \in [N]$ and every $m \geq 1$,*

$$\left|e_{i,m} - \hat{c}_{i,m}\right| \leq \gamma\, c_{\max}\, \beta(m). \tag{52}$$

The constant γ is a property of the estimation rule, and $\gamma = 0$ recovers the empirical mean.

Lemma 18. *Under Assumption 1, on the event $\mathcal{E}$ of Lemma 1, for every provider $i \in [N]$ and every $m \geq 1$,*

$$\left|e_{i,m} - \bar{c}_i\right| \leq (1+\gamma)\, c_{\max}\, \beta(m). \tag{53}$$

Proof. On $\mathcal{E}$, $|\hat{c}_{i,m} - \bar{c}_i| \leq c_{\max}\beta(m)$ by Eq. 13. By the triangle inequality and Eq. 52,

$$\left|e_{i,m} - \bar{c}_i\right| \leq \left|e_{i,m} - \hat{c}_{i,m}\right| + \left|\hat{c}_{i,m} - \bar{c}_i\right| \leq \gamma c_{\max}\beta(m) + c_{\max}\beta(m). \qquad \square$$

Bayesian variant of the mechanism. The platform uses a known upper bound γ satisfying Assumption 1—any larger valid bound may be used, since Eq. 52 is monotone in γ—and replaces the cost-estimation radius $c_{\max}\beta(m_{i,t})$ by $(1+\gamma)c_{\max}\beta(m_{i,t})$ in the adjusted bids of the allocation rule (Eq. 3) and in the critical payment (Eq. 4), *i.e.*,

$$\begin{aligned} I_t &= \arg\min_{i \in \hat{\mathcal{Q}}_t} \big(b_{i,t} - (1+\gamma)c_{\max}\beta(m_{i,t})\big), && \text{if } \mathcal{S}_t = \emptyset, \\ P_{i,t} &= (1+\gamma)c_{\max}\beta(m_{i,t}) + \min_{j \in \hat{\mathcal{Q}}_t \setminus \{i\}} \big(b_{j,t} - (1+\gamma)c_{\max}\beta(m_{j,t})\big). \end{aligned} \tag{54}$$

The quality estimates $\hat{q}_{i,t}$, the estimated-qualified set $\hat{\mathcal{Q}}_t$, the exploration set $\mathcal{S}_t$, the schedule g, the initialization and exploration payments $c_{\max}$, and the payment cap in Eq. 5 are unchanged, and $\gamma = 0$ recovers the mechanism of Section 2. Providers bid their current estimate, $b_{i,t} = e_{i,t}$, and the utility in Eq. 7 is evaluated with $e_{i,s}$ in place of $\hat{c}_{i,s}$.

The widening of the radius is necessary: if the platform kept $c_{\max}\beta$, the adjusted bid of i^* would only satisfy $b_{i^*,t} - c_{\max}\beta(m_{i^*,t}) \leq \bar{c}_{i^*} + \gamma c_{\max}\beta(m_{i^*,t})$, and the optimism property $b_{i,t} - c_{\max}\beta(m_{i,t}) \leq \bar{c}_i$ on which Lemma 6 rests would fail.

Proposition 13. *Consider the Bayesian variant of the mechanism in Eq. 54 with all providers bidding $b_{i,t} = e_{i,t}$, and let $L_i^{(\gamma)} = M\big(\Delta_i/((1+\gamma)c_{\max})\big)$, where $M(\cdot)$ is defined in Theorem 3.*

- *i) Without Assumption 1: Lemmas 1, 2, 3, 4, 8, 10, 12, 13, 14, and 15, Theorem 7, and the quantities M_i, $\bar{M}$, and T_0 hold exactly as stated, under arbitrary bids; and Proposition 1, Lemma 17, Theorem 2, and the robust-equilibrium property that follows it hold with $e_{i,t}$ in place of $\hat{c}_{i,t}$, for any estimator e_i with values in $[0, c_{\max}]$.*
- *ii) Under Assumption 1: Lemmas 5, 6, 7, 9, 11, and 16, Theorems 3, 8, 9, 11, and 4, and Corollaries 5, 6, 10, and 12 hold after replacing each cost-estimation radius $c_{\max}\beta(\cdot)$ by $(1+\gamma)c_{\max}\beta(\cdot)$ in their statements, with the corresponding replacement $L_i \to L_i^{(\gamma)}$.*

Proof. Claim i). The results in the first group concern the quality estimates, the estimated-qualified set, the exploration rule, or the confidence radius alone; their proofs do not involve the bids or the cost estimates, all of which the Bayesian variant leaves unchanged, so they hold as stated under arbitrary bids. For the second group, the proofs of Proposition 1 and Lemma 17 use only that the allocation rule selects a minimizer of $b_{j,t} - \rho_{j,t}$ over $\hat{\mathcal{Q}}_t$ and pays $\min\{P_{i,t}, c_{\max}\}$ with $P_{i,t} = \rho_{i,t} + \min_{j \in \hat{\mathcal{Q}}_t \setminus \{i\}}(b_{j,t} - \rho_{j,t})$, for arbitrary radii $\rho_{j,t} \geq 0$ determined by the history, together with the fact that the bid compared to $P_{i,t}$ and the estimate entering the net payment $G_{i,s}$ coincide and lie in $[0, c_{\max}]$. Both hold for $\rho_{j,t} = (1+\gamma)c_{\max}\beta(m_{j,t})$ and $e_{i,t} \in [0, c_{\max}]$. The proof of Theorem 2 uses only Lemma 17 and that the truthful strategy bids the estimate entering Eq. 7 at every step. None of these proofs uses the event $\mathcal{E}$ or the form of the estimator.

Claim ii). The only property of the bids used in the listed proofs is the adjusted-bid bound of Lemma 5, which in the Bayesian variant reads, on $\mathcal{E}$ and by Lemma 18,

$$\bar{c}_i - 2(1+\gamma)c_{\max}\beta(m_{i,t}) \leq b_{i,t} - (1+\gamma)c_{\max}\beta(m_{i,t}) \leq \bar{c}_i, \tag{55}$$

with the same one-line proof. Downstream of this bound, the radius enters only through the choice $a = \Delta_i / c_{\max}$ in the application of Lemma 2 (proof of Lemma 7), which becomes $a = \Delta_i / ((1+\gamma)c_{\max})$ and yields $L_i^{(\gamma)}$, and the choices $\kappa = c_{\max}$ and $\kappa = 2c_{\max}$ in the applications of Lemma 10 (proofs of Theorems 9 and 8), which become $\kappa = (1+\gamma)c_{\max}$ and $\kappa = 2(1+\gamma)c_{\max}$; both lemmas hold for arbitrary $a > 0$ and $\kappa \geq 0$. Every other occurrence of $c_{\max}$ in these proofs—the bid range, the initialization and exploration payments, the payment cap, the bound $|\pi_t - \bar{c}_{(2)}| \leq c_{\max}$, and the factor $c_{\max} - \bar{c}_{(2)}$—refers to the range of costs and payments, which is unchanged. Hence each listed proof goes through verbatim under the stated replacement. □

Resulting guarantees. Two consequences of Proposition 13ii) are worth recording. First, Theorem 3 holds with $M(\Delta_i / c_{\max})$ replaced by $L_i^{(\gamma)} = M\big(\Delta_i / ((1+\gamma)c_{\max})\big)$: the ordinary-step selection threshold of a qualified suboptimal provider grows with γ, while the terms for unqualified providers, the exploration term, and the quality guarantee of Theorem 7 are unchanged. Second, Theorem 4 holds with the coefficient $\frac{6\,c_{\max}}{2-\alpha}$ of its last term replaced by $\frac{6(1+\gamma)\,c_{\max}}{2-\alpha}$, and its first two terms unchanged; likewise, Lemma 16 gives $\bar{c}_{(2)} - 3(1+\gamma)c_{\max}\beta(g(t)) \leq \pi_t \leq \bar{c}_{(2)} + (1+\gamma)c_{\max}\beta(g(t))$ on every non-exploration step $t > T_0$. Hence, for fixed γ, the average absolute deviation of the payments from $\bar{c}_{(2)}$ is $O\big(T^{-(1-\alpha)}\big) + O\big((1+\gamma)\,T^{-\alpha/2}\sqrt{\ln T}\big)$, the exponent is still maximized at $\alpha = 2/3$ (Corollary 12), and $\frac{1}{T}\sum_{t=1}^{T} \pi_t \to \bar{c}_{(2)}$ as $T \to \infty$. The average excess-payment bound of Corollary 10 receives the same radius substitution in its last term only.

Provider incentives. By Proposition 13i), bidding one's current cost estimate, $\sigma_i^*(H_{i,t}) = e_{i,t}$, is worst-case optimal in the sense of Theorem 2 for any estimator with values in $[0, c_{\max}]$, in the Bayesian variant of the mechanism as well as in the original one. The robust-optimality result is thus estimator-agnostic: once a provider's utility is evaluated relative to its current estimate, the threshold structure of the allocation rule makes bidding that estimate optimal. Assumption 1 plays no role in this result; it is needed only to relate the providers' estimates to their true mean costs, and thereby obtain the user-facing guarantees above.

D ROBUSTNESS TO DEVIATIONS FROM COST-ESTIMATE BIDDING

Theorem 2 shows that, for a provider, bidding its current cost estimate is worst-case optimal, but a provider may nevertheless deviate from it. In this section, we quantify how the user-facing guarantees of Section 4 change when bids may fall below or above the cost estimate by bounded amounts.

Assumption 2. *There exist constants $d_-, d_+ \geq 0$ such that every provider $i \in [N]$ follows a strategy σ_i with*

$$-d_- \;\leq\; \sigma_i(H_{i,t}) - \hat{c}_{i,t} \;\leq\; d_+ \qquad \textit{for every history } H_{i,t},\ t > N. \tag{56}$$

The constant d_- bounds underbidding and d_+ overbidding relative to the empirical mean cost; the equilibrium strategy $\boldsymbol{\sigma}^*$ of Eq. 8 is the case $d_- = d_+ = 0$.

The good-event, quality-estimation, and forced-exploration results of Appendix B do not use the providers' bids. In particular, Lemmas 1, 2, 3, 4, 8, 10, 12, 13, 14, and 15, Theorem 7, and the quantities M_i, $\bar{M}$, and T_0 hold under Assumption 2 exactly as stated; unqualified providers are therefore eliminated, and the quality guarantee holds, unchanged. The remaining results depend on the bids only through Lemma 5, whose counterpart under deviations is the following:

Lemma 19. *Under Assumption 2, on the event $\mathcal{E}$, for every provider $i \in [N]$ and step $t > N$,*

$$\bar{c}_i - 2c_{\max}\beta(m_{i,t}) - d_- \;\leq\; b_{i,t} - c_{\max}\beta(m_{i,t}) \;\leq\; \bar{c}_i + d_+. \tag{57}$$

Proof. On $\mathcal{E}$, $|\hat{c}_{i,t} - \bar{c}_i| \leq c_{\max}\beta(m_{i,t})$ by Eq. 13; combine with $\hat{c}_{i,t} - d_- \leq b_{i,t} \leq \hat{c}_{i,t} + d_+$. □

The adjusted bid is no longer optimistic, since its upper bound exceeds $\bar{c}_i$ by d_+, and its lower bound is lowered by d_-. Tracing these two shifts through the proofs of Appendix B gives the results below.

Selection and identification. Underbidding by a competitor and overbidding by i^* both help the competitor win a non-exploration round, so the two deviations enter the selection analysis.

Proposition 14. *Under Assumption 2, on the event $\mathcal{E}$:*

i) On every non-exploration step $t > N$, the selected provider satisfies

$$\bar{c}_{I_t} - \bar{c}_{i^*} \;\leq\; 2c_{\max}\beta(m_{I_t,t}) + d_- + d_+. \tag{58}$$

ii) Every qualified provider $i \in \mathcal{Q} \setminus \{i^*\}$ with $\Delta_i > d_- + d_+$ is selected after initialization at most $\max\{L_i^{(d)}, \lceil g(T) \rceil - 1\}$ times, where

$$L_i^{(d)} \;\coloneqq\; M\!\left(\frac{\Delta_i - d_- - d_+}{c_{\max}}\right) \tag{59}$$

and $M(\cdot)$ is defined in Theorem 3.

iii) If $\Delta_i > d_- + d_+$ for every $i \in \mathcal{Q} \setminus \{i^*\}$, then Theorem 3 and Corollaries 5 and 6 hold with $M(\Delta_i/c_{\max})$ replaced by $L_i^{(d)}$.

Proof. Claim i). As in the proof of Lemma 6: on $\mathcal{E}$, $i^* \in \hat{\mathcal{Q}}_t$ by Lemma 3, so the selected provider $I_t = i$ satisfies $b_{i,t} - c_{\max}\beta(m_{i,t}) \leq b_{i^*,t} - c_{\max}\beta(m_{i^*,t})$. Bounding the left side from below and the right side from above with Eq. 57 gives $\bar{c}_i - 2c_{\max}\beta(m_{i,t}) - d_- \leq \bar{c}_{i^*} + d_+$. This holds for every $i \in \hat{\mathcal{Q}}_t$, qualified or not, since Eq. 57 holds for every provider.

Claim ii). At every non-exploration selection of $i \in \mathcal{Q} \setminus \{i^*\}$, claim i) gives $(\Delta_i - d_- - d_+)/c_{\max} \leq 2\beta(m_{i,t})$, and Lemma 2 with $a = (\Delta_i - d_- - d_+)/c_{\max} > 0$ gives $m_{i,t} \leq L_i^{(d)}$. The rest is the proof of Lemma 9 with L_i replaced by $L_i^{(d)}$.

Claim iii). The proofs of Theorem 3 and its corollaries use the bids only through Lemma 9, which claim ii) replaces. □

The condition $\Delta_i > d_- + d_+$ cannot be dropped: if $\Delta_i \le d_- + d_+$, the bids $b_{i,t} = \hat{c}_{i,t} - d_-$ and $b_{i^*,t} = \hat{c}_{i^*,t} + d_+$ satisfy Assumption 2 and, once the two providers have comparable selection counts, give i an adjusted bid at most that of i^* up to the confidence radii, so i may win every ordinary round. When some qualified competitor is within $d_- + d_+$ of i^*, the platform therefore need not identify i^*; claim i) shows what it guarantees instead: on every ordinary round it selects a provider whose mean cost exceeds $\bar{c}_{i^*}$ by at most $d_- + d_+$ plus a vanishing term, and that provider is qualified on all but the rounds bounded by Lemma 8.

Theorem 15. *Under Assumption 2, on the event $\mathcal{E}$, the generation-cost regret of Theorem 8 satisfies*

$$R_T^{\mathrm{gen}} \;\le\; (N-1)\lceil g(T)\rceil c_{\max} + 2c_{\max}\sqrt{2N(T-N)\ln\!\left(\frac{2\pi^2 NT^2}{3\delta}\right)} + (d_- + d_+)(T-N). \quad (60)$$

If moreover $\Delta_i > d_- + d_+$ for every $i \in \mathcal{Q} \setminus \{i^*\}$, the second bound of Theorem 8 holds with L_i replaced by $L_i^{(d)}$.

Proof. In the proof of the first bound of Theorem 8, the inequality $\bar{c}_i - 2c_{\max}\beta(m_{i,t}) \le \bar{c}_{i^*}$ on non-exploration steps is replaced by Eq. 58. Splitting $a_t \le 2c_{\max}\beta(m_{I_t,t}) + (d_- + d_+)$, Lemma 10 with $\kappa = 2c_{\max}$ bounds the sum of the first terms as before, and the second terms contribute at most $(d_- + d_+)(T - N)$. The second bound uses the bids only through Lemma 9, which Proposition 14ii) replaces. □

Dividing by T, the average generation-cost regret is at most $d_- + d_+$ up to a term of order $T^{-(1-\alpha)} + \sqrt{\ln T/T}$, whatever the cost gaps.

Payments. Overbidding and underbidding affect the two sides of the payment bounds separately: the payment is set by the adjusted bid of the cheapest competing provider, which overbidding raises by at most d_+ and underbidding lowers by at most d_-.

Lemma 20. *Under Assumption 2, on the event $\mathcal{E}$, on every non-exploration step $t > N$ with $I_t = i$,*

$$\pi_t \;\le\; \min\bigl\{c_{\max},\; \bar{c}_{(2)} + c_{\max}\beta(m_{i,t}) + d_+\bigr\}, \quad (61)$$

and, if moreover $t > T_0$,

$$\pi_t \;\ge\; \bar{c}_{(2)} - 3c_{\max}\beta(g(t)) - d_-. \quad (62)$$

Proof. Upper bound. As in the proof of Lemma 11, with $j = i^*$ if $i \ne i^*$ and j the second-cheapest qualified provider if $i = i^*$: by Eq. 57, $P_{i,t} \le c_{\max}\beta(m_{i,t}) + b_{j,t} - c_{\max}\beta(m_{j,t}) \le c_{\max}\beta(m_{i,t}) + \bar{c}_j + d_+ \le c_{\max}\beta(m_{i,t}) + \bar{c}_{(2)} + d_+$.

Lower bound. For $t > T_0$, Lemma 14 gives $\hat{\mathcal{Q}}_t = \mathcal{Q}$ and $m_{j,t} \ge g(t)$ for every $j \in \mathcal{Q}$, so $\beta(m_{j,t}) \le \beta(g(t))$ by Lemma 12. Write $A_{j,t} = b_{j,t} - c_{\max}\beta(m_{j,t})$, so that $P_{i,t} = c_{\max}\beta(m_{i,t}) + \min_{j \in \mathcal{Q}\setminus\{i\}} A_{j,t}$. For every $j \in \mathcal{Q} \setminus \{i, i^*\}$, Eq. 57 and $\bar{c}_j \ge \bar{c}_{(2)}$ give $A_{j,t} \ge \bar{c}_{(2)} - 2c_{\max}\beta(g(t)) - d_-$. If $i = i^*$, this covers every $j \in \mathcal{Q} \setminus \{i\}$ and $P_{i,t} \ge \bar{c}_{(2)} - 2c_{\max}\beta(g(t)) - d_-$. If $i \ne i^*$, the remaining competitor is i^*, and since i was selected, $A_{i^*,t} \ge A_{i,t}$, so that

$$\begin{aligned} c_{\max}\beta(m_{i,t}) + A_{i^*,t} &\ge c_{\max}\beta(m_{i,t}) + A_{i,t} = b_{i,t} \ge \hat{c}_{i,t} - d_- \\ &\ge \bar{c}_i - c_{\max}\beta(g(t)) - d_- \ge \bar{c}_{(2)} - c_{\max}\beta(g(t)) - d_- \end{aligned}$$

using $i \in \mathcal{Q} \setminus \{i^*\}$ in the last step. Hence $P_{i,t} \ge \bar{c}_{(2)} - 2c_{\max}\beta(g(t)) - d_-$ in both cases, and $\pi_t = \min\{P_{i,t}, c_{\max}\} \ge \bar{c}_{(2)} - 2c_{\max}\beta(g(t)) - d_- \ge \bar{c}_{(2)} - 3c_{\max}\beta(g(t)) - d_-$, since $\bar{c}_{(2)} \le c_{\max}$. □

At $d_- = d_+ = 0$, Eqs. 61 and 62 are Lemmas 11 and 16.

Theorem 16. *Under Assumption 2, with probability at least $1 - \delta$, for every $T \ge T_0$,*

$$\begin{aligned} \frac{1}{T}\sum_{t=1}^{T}\bigl|\pi_t - \bar{c}_{(2)}\bigr| \le \frac{T_0}{T}\,c_{\max} + k\left(\frac{N}{T}\right)^{1-\alpha}\bigl(c_{\max} - \bar{c}_{(2)}\bigr) + \max\{d_-, d_+\} \\ + \frac{6\,c_{\max}}{2-\alpha}\sqrt{\frac{1}{2k}\left(\frac{N}{T}\right)^{\alpha}\ln\!\left(\frac{2\pi^2 Nk^2}{3\delta}\left(\frac{T}{N}\right)^{2\alpha}\right)} \end{aligned} \quad (63)$$

and consequently

$$\bar{c}_{(2)} - d_- \;\le\; \liminf_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} \pi_t \;\le\; \limsup_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} \pi_t \;\le\; \bar{c}_{(2)} + d_+. \tag{64}$$

Moreover, the average excess payment of Corollary 10 is bounded by its right-hand side plus d_+.

Proof. In the proof of Theorem 11, the only change is on the non-exploration steps $t > T_0$, where Lemma 20 gives $|\pi_t - \bar{c}_{(2)}| \le 3c_{\max}\beta(g(t)) + \max\{d_-, d_+\}$ in place of $3c_{\max}\beta(g(t))$; summing the additional term over at most T steps and dividing by T gives Eq. 63. For Eq. 64, the same decomposition applied to $\pi_t - \bar{c}_{(2)}$ without absolute values gives, for the non-exploration steps $t > T_0$, $-3c_{\max}\beta(g(t)) - d_- \le \pi_t - \bar{c}_{(2)} \le c_{\max}\beta(g(t)) + d_+$, while the first T_0 steps and the exploration steps contribute $O(T_0/T + T^{-(1-\alpha)})$ to the average; the vanishing terms are those of Eq. 63. The excess-payment claim follows from Eq. 61 in the proof of Theorem 9, whose $\sqrt{\cdot}$ term is unchanged and which acquires the additional term $d_+(T - N)$. □

Summary. Under bounded deviations, the platform still serves qualified providers on all but a bounded number of rounds; it identifies i^* whenever every qualified competitor is more than $d_- + d_+$ costlier, and otherwise selects a qualified provider within $d_- + d_+$ of $\bar{c}_{i^*}$ in mean cost; and the average payment converges to within $[\bar{c}_{(2)} - d_-,\ \bar{c}_{(2)} + d_+]$ at the rate of Theorem 4. Underbidding thus affects the lower side of the payment guarantee by at most d_-, whereas overbidding affects the upper side by at most d_+; neither affects the quality guarantee, and both slow identification through the reduced gaps $\Delta_i - d_- - d_+$. Setting $d_- = d_+ = 0$ recovers Theorems 3, 8, and 4 verbatim. Finally, for the Bayesian variant of Appendix C, with Assumption 2 stated relative to $e_{i,t}$ in place of $\hat{c}_{i,t}$, the same statements hold with $c_{\max}\beta(\cdot)$ replaced by $(1+\gamma)c_{\max}\beta(\cdot)$ throughout and $L_i^{(d)}$ replaced by $M\big((\Delta_i - d_- - d_+)/((1+\gamma)c_{\max})\big)$, since the proofs above use the empirical mean only through Eq. 13, which Lemma 18 replaces.

E EXPERIMENTAL DETAILS

In this section, we describe the experimental setup we use to evaluate our platform (Algorithm 1). Here, our goal is twofold: first, to instantiate the model of Section 2 using recorded generations from real LLMs, so that provider qualities and generation costs are grounded in observed data; and second, to compare the platform against practical routing policies that a platform could deploy today. In what follows, we describe the data, the construction of providers, the experimental settings, the routing policies we compare against, and the metrics we report.

Recorded generation data. We evaluate the platform using a publicly available dataset of recorded LLM generations.[11] The data cover `GSM8K` (Cobbe et al., 2021) and `GPQA` (Rein et al., 2024), spanning a wide range of difficulty. For every model–question pair, the data contain 128 recorded generations, each annotated with a binary correctness label, an output-token count, and, when needed for Best-of-n_i service, a score from the reward model `ArmoRM-Llama3-8B-v0.1` (Wang et al., 2024a). Our simulator retains only these quantities and a question identifier; i.e., neither the prompts nor the generated text are used. The data contain 1,319 `GSM8K` questions and 546 `GPQA` questions. Since Best-of-n_i service requires the correctness, token, and reward records to be aligned for every included model, we remove one `GPQA` question that fails this alignment check jointly from all models whenever Best-of-n_i variants are in use, leaving 545 `GPQA` questions in those experiments; no questions are removed from `GSM8K`. As a result, every retained model has the same question set and 128 recorded generations per question.

Base models and listed prices. We use the nine base models listed in Table 2, together with their publicly listed prices. To determine the per-output-token price of each model, we follow the procedure of Velasco et al. (2026a): we refer to the Hugging Face list of inference providers[12] and, for each model, compute the average output-token price across the listed providers that offer access to that model. More formally, we denote by λ_m the resulting listed price of model m in USD per 10^6 output tokens. Throughout, all monetary quantities are reported in units of 10^{-6} USD, so that the numerical value of λ_m coincides with the price per output token in these units.

Model	λ_m	Model	λ_m
`Llama-3-8B`	0.1455	`Qwen2-0.5B`	0.1000
`Llama-3.1-8B`	0.1245	`Qwen2-1.5B`	0.1000
`Llama-3.2-1B`	0.1000	`Qwen2-7B`	0.2000
`Llama-3.2-3B`	0.0800	`Qwen2.5-3B`	0.0650
		`Qwen2.5-7B`	0.1465

Table 2: Listed output-token prices used in the experiments, in USD per 10^6 output tokens.

Provider construction. To instantiate providers with a cost–quality trade-off, we build each provider from a base model m and a service multiplicity n_i, and write $i = (m, n_i)$ to identify the provider by this pair. A base provider, with $n_i = 1$, serves a query with a single generation of model m, whereas a Best-of-n_i provider generates n_i candidate responses, incurs the generation cost of all n_i candidates, and returns the candidate with the highest recorded reward score. More concretely, for each question, the simulator samples n_i distinct generations uniformly without replacement from the 128 recorded generations of model m and, if $n_i > 1$, returns the sampled generation with the highest reward score, resolving ties uniformly through the random ordering of the sampled generations. The correctness of the response, which serves as the score $r(x_t, y_t)$ of Section 2, is the binary correctness of the returned generation, while the token count used for pricing and cost includes all n_i candidates. We consider three rosters. The *full* roster contains all nine base models with $n_i = 1$. The *ladder* roster contains

$$m \in \{\texttt{Llama-3-8B}, \texttt{Llama-3.2-1B}, \texttt{Qwen2-1.5B}, \texttt{Qwen2.5-7B}\}$$

[11] `https://huggingface.co/datasets/Human-Centric-Machine-Learning/strategic-ttc-data`

[12] `https://huggingface.co/docs/inference-providers/index`, consulted on December 30, 2025.

Benchmark	Roster	N	K	$q_{\min}$	$\lvert\mathcal{Q}\rvert$	i^*
GSM8K	full	9	1319	0.7640	3	Qwen2.5-3B
GPQA	full	9	546	0.0790	8	Qwen2-1.5B
GSM8K	ladder	12	1319	0.8050	3	Qwen2.5-7B
GPQA	ladder	12	545	0.2396	4	Llama-3-8B
GPQA	strong-competition	4	546	0.1451	2	Llama-3.1-8B

Table 3: Experimental settings. The benchmark size K equals the number of aligned benchmark questions available to the corresponding roster.

at each multiplicity $n_i \in \{1, 2, 4\}$, for a total of 12 providers. The ladder roster therefore introduces a direct cost–quality trade-off within each base model, in addition to the variation across model families. The *strong-competition* roster is used on GPQA only and contains the newest large model and the smallest model of each family, Llama-3.1-8B, Llama-3.2-1B, Qwen2.5-7B and Qwen2-0.5B, all with $n_i = 1$.

Empirical quality, price, and cost. We now specify how the quantities of Section 2 are instantiated from the recorded data. To this end, it is important to distinguish the listed price, the query price, and the generation cost. The *listed price* λ_m is the public per-token rate in Table 2; it is known before routing and does not depend on the realized output length. We define the realized *query price* $V_{i,t}$ of provider $i = (m, n_i)$ in step t as the listed price λ_m times the total number of output tokens across the n_i recorded generations sampled in that step, i.e., the amount a platform using the model at its listed rate would be invoiced once the output lengths are known. Thus, a Best-of-n_i query is billed for every generated candidate, not only for the returned response. Since the data do not contain the providers' true generation costs, we construct the private realized *generation cost* as

$$c_{i,t} \coloneqq \frac{V_{i,t}}{1+\mu}, \qquad \mu = 0.25, \tag{65}$$

where μ is a (fixed) margin. It is worth noting that the margin μ is used only to construct the latent cost observations, and it is neither supplied to nor used by the platform. Building on the above, the empirical quality q_i of provider i is their expected correctness under a uniformly sampled benchmark question and the generation-sampling procedure above. For a base provider, it is the mean correctness over all recorded question–generation pairs, while for a Best-of-n_i provider, it is computed exactly from the empirical joint distribution of reward scores and correctness, including uniform tie-breaking. Analogously, the cost distribution $\mathcal{P}_i^c$ of provider i is the distribution of the realized generation cost $c_{i,t}$ in Eq. equation 65 under the same sampling procedure as q_i, and $\bar{c}_i$ is its mean. Since query price and generation cost are proportional by construction, the mean query price of provider i is $\bar{v}_i = (1+\mu)\bar{c}_i$. These quantities are used only for evaluation and for constructing perfect-information reference quantities; they are not given to the platform. Finally, in our primary experiments, providers bid their empirical-mean cost estimate $\hat{c}_{i,t}$, *i.e.*, they follow the equilibrium strategy σ_i^* of Eq. equation 8, and payments follow the rule of Eq. equation 5.

Settings and platform parameters. For each roster and benchmark, we place the quality threshold $q_{\min}$ at the midpoint between two adjacent empirical provider qualities. Table 3 reports, for each of the resulting five settings, the exact threshold, the benchmark size K, the number of empirically qualified providers, and the lowest-cost qualified provider i^* of Eq. equation 10.

We set the confidence level to $\delta = 0.05$. The output-token cap entering the cost bound is fixed at $L_{\max} = 512$, and the cost bound $c_{\max}$ of Section 2, which also caps every payment, is set to

$$c_{\max} = L_{\max} \max_{i=(m,n_i)} n_i \lambda_m, \tag{66}$$

which gives $c_{\max} \approx 102$ for the full roster, $c_{\max} \approx 300$ for the ladder roster and $c_{\max} \approx 75$ for the strong-competition roster. Since private costs are obtained by deflating query prices by $1+\mu$, this ceiling also bounds every realized provider cost, as required by our model. Interestingly, the two full-roster settings provide qualitatively different tests of routing by listed price. The lowest listed rate belongs to Qwen2.5-3B, which is i^* on GSM8K and is qualified but not i^* on GPQA. Listed-price routing is therefore, respectively, aligned with and only partially informative about the platform's target decision.

Episode construction and randomization. At the beginning of every episode, the benchmark questions are placed in a seeded random permutation. For every step and every provider, the simulator preassigns a potential outcome using the recorded-generation sampling rule above; only the selected provider's outcome is then revealed, while the potential outcomes of nonselected providers remain unavailable to the platform. This preserves the platform's information design of Section 2 while allowing different policies to be evaluated on matched potential outcomes. The first N steps initialize the platform as prescribed in Section 2, selecting every provider exactly once in a seeded random order, and all subsequent steps follow the platform's selection and payment rules using the pre-step state.

Exploration steps. In every experiment of this appendix the platform runs with the exploration rule of Eq. equation 3. At the beginning of step $t > N$, if some estimated-qualified provider $i \in \hat{\mathcal{Q}}_t$ (Eq. equation 2), which we also refer to as an *eligible* provider, has been selected fewer than

$$g(t) = k\,(t/N)^{\alpha}, \qquad k = 2,\ \alpha = 3/4,$$

times, *i.e.*, if $\mathcal{S}_t \neq \emptyset$, the step is an exploration step: one of the under-sampled eligible providers in $\mathcal{S}_t$ is drawn uniformly at random, serves the query and is paid $c_{\max}$; otherwise it is a non-exploration step, selected and paid according to Eqs. equation 3 and equation 5. Exploration steps count toward every selection statistic, regret, generation cost, query price and amount paid; the winning bid and the payment on non-exploration steps are reported over those steps only. The routing baselines and the no-quality-filter ablation do not explore.

Routing policies. Our primary comparison uses three routing policies, none of which observes private generation costs: (i) *uniform routing among eligible providers*, (ii) *lowest-listed-price routing among eligible providers*, and (iii) *lowest-listed-price routing without quality screening*. Policy (i) constructs the same online estimated-qualified set $\hat{\mathcal{Q}}_t$ as the platform, i.e., it may use the observed scores, the selection counts, the threshold $q_{\min}$, and the confidence radius $\beta(m_{i,t})$ of Eq. equation 2, and selects uniformly at random from $\hat{\mathcal{Q}}_t$; it uses neither bids nor listed prices, and thus isolates the contribution of cost-sensitive selection after online quality screening. Policy (ii) constructs the same $\hat{\mathcal{Q}}_t$ and then minimizes the public effective listed rate

$$\bar{\lambda}_i := n_i \lambda_m \qquad \text{for } i = (m, n_i),$$

breaking ties uniformly at random; it observes listed rates but neither provider bids nor generation costs, and the realized response length is not known at selection time. Policy (iii) is static: it minimizes $\bar{\lambda}_i$ over the entire roster, irrespective of observed quality, and requires only the public listed rates. It therefore provides a direct measure of the reliability lost when routing ignores the quality requirement. In addition, we consider a *no-quality-filter* ablation of the platform, which retains the optimistic bidding index $b_{i,t} - c_{\max}\beta(m_{i,t})$ of Eq. equation 3, the confidence radius $\beta(m_{i,t})$, and the critical payment of Eq. equation 4, but minimizes the index over the full roster rather than over $\hat{\mathcal{Q}}_t$; this ablation isolates the contribution of online quality screening. The platform pays π_t as in Eq. equation 5, the routing policies pay the query price of the selected provider, and both are reported as the amount paid.

Horizons and long-run evaluation. To examine longer-run behavior, we additionally run every setting for $T = 70{,}000$ queries and 30 independent runs, paired across rules, corresponding to 54 passes over the `GSM8K` benchmark and 129 passes over `GPQA`, with the final pass only partially used. At the beginning of each additional pass, the simulator draws a fresh seeded permutation of all benchmark questions and fresh recorded generations for every provider–question pair, so that the same generation may reappear in different passes; extending the horizon leaves the complete first pass unchanged. This resampling procedure approximates repeated draws from each provider's empirical outcome distribution. It does not create new questions or new model generations, and the resulting long-horizon evidence remains conditional on the observed benchmark and generation pools.

Metrics. Our main metrics are the quality regret $R_T^{\mathrm{qual}} = \sum_{t=1}^{T}(q_{\min} - q_{I_t})^+$ of Theorem 7 and the generation-cost regret $R_T^{\mathrm{gen}} = \sum_{t=1}^{T}(\bar{c}_{I_t} - \bar{c}_{i^*})^+$ of Theorem 8, reported both cumulatively and per query. Note that both regrets are computed using the empirical provider parameters q_i and $\bar{c}_i$

rather than the realized correctness labels and costs, which are subject to sampling noise. We report realized accuracy separately as

$$\widehat{\mathrm{Acc}}_T = \frac{1}{T}\sum_{t=1}^{T} r(x_t, y_t),$$

together with the share of unqualified selections,

$$\frac{1}{T}\sum_{t=1}^{T} \mathbb{1}\{I_t \notin \mathcal{Q}\},$$

the selection share of i^*, provider-level selection counts, and whether i^* is the unique most selected provider. The economic outcomes comprise realized generation cost per query, query price per query, and the amount paid per query under the policy-specific convention above. For the platform and its no-quality-filter ablation, we additionally report the winning standing bid, the total payment on non-exploration steps, the cumulative excess payment $E_T^{\mathrm{pay}} = \sum_{t=1}^{T}(\pi_t - \bar{c}_{(2)})^+$ bounded in Theorem 9, and the provider payoff $\sum_{t:I_t=i}(\pi_t - c_{i,t})$. Where useful, these quantities are reported both over the full episode and after initialization, since initialization pays $c_{\max}$ by construction. Payment diagnostics compare the realized payments on non-exploration steps with the upper bound $\pi_t \le \min\{c_{\max},\, \bar{c}_{(2)} + c_{\max}\,\beta(m_{I_t,t})\}$ of Lemma 11 and with the lower bound $\pi_t \ge \max\{0,\, \bar{c}_{I_t} - c_{\max}\,\beta(m_{I_t,t})\}$, which holds on $\mathcal{E}$ because, by Proposition 1, the selected provider bids $\hat{c}_{I_t,t} \le P_{I_t,t}$. We further report the realized frequency of the good event $\mathcal{E}$ of Lemma 1 and compare the post-initialization selections of providers $i \neq i^*$ with the corresponding bounds M_i of Lemma 8 for unqualified providers and $\max\{L_i,\, \lceil g(T)\rceil - 1\}$ of Lemma 9 for qualified ones, where L_i is defined in Lemma 7. Accordingly, when reporting $B_{\mathrm{id}}(T)$ and T_{id}, we replace the exploration term $k(T/N)^\alpha$ of Theorem 3 by the cap $\lceil g(T)\rceil - 1 < k(T/N)^\alpha$ of Lemma 9, which yields a slightly tighter bound that is also valid by the proof of Theorem 3. For the long-horizon experiment, we additionally report the size of the estimated-qualified set $\hat{\mathcal{Q}}_t$, the probability that i^* is the unique most selected provider, and moving-average selection and payment statistics with window size K.

Aggregation and uncertainty. Episode-level tables report the mean and standard deviation across runs. For one-pass figures, points and selection shares are shown with normal-approximation 95% intervals, $\overline{x} \pm 1.96\, s_x/\sqrt{R}$, where R is the number of runs and s_x is their sample standard deviation. Long-horizon trajectory figures use the pointwise median across runs as the central trajectory and the empirical 10th–90th percentile range as a band. Long-horizon checkpoint tables report means across runs, while identification probabilities are the corresponding empirical fractions of runs.

F GSM8K WEAK-COMPETITION EXPERIMENTS

This section reports the complete GSM8K results for the full and ladder rosters; Appendix E provides the experimental details. The platform runs with the exploration rule described there, $k = 2$ and $\alpha = 3/4$. One-pass results are averaged over 30 independent runs, paired across rules and reported with 95% confidence intervals, while long-horizon trajectories report the median and 10–90% range over 30 independent runs. Monetary quantities are per query unless stated otherwise. A moving average has window size $K = 1{,}319$, one pass over the benchmark, so that it reflects a rule's current behaviour rather than its whole history. For the platform and its no-quality-filter ablation the amount paid is the payment π_t; for the routing baselines it is the query price of the selected provider. On the full roster $\bar{c}_{(2)} = 34.4$ lies above the mean query price of i^*, $\bar{v}_{i^*} = 20.1$, so this is the weak-competition case in which a payment that approaches $\bar{c}_{(2)}$ exceeds the listed price of i^*.

F.1 PROVIDERS AND ENVIRONMENT

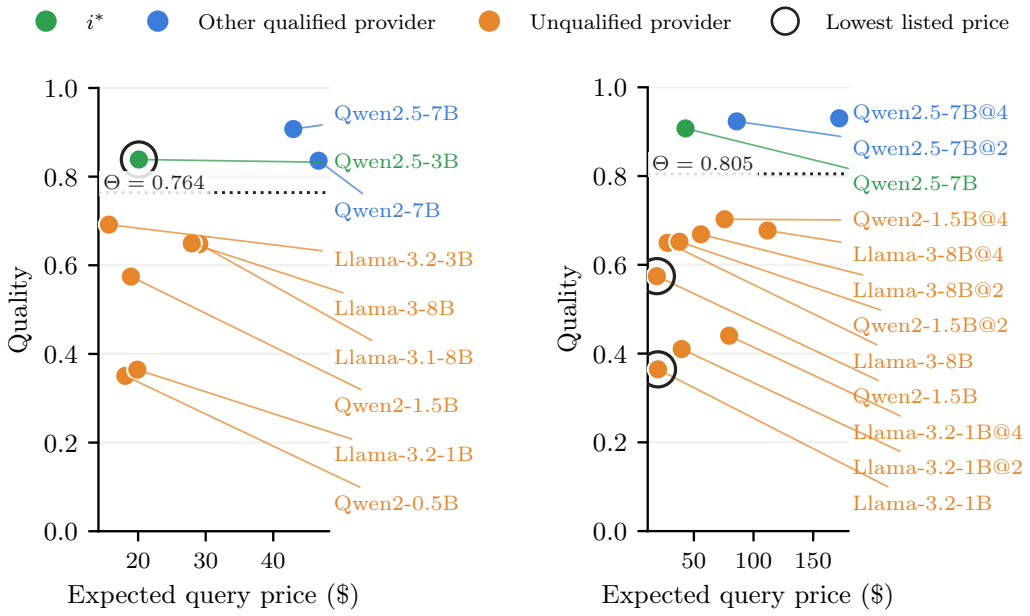

(a) Quality against mean query price, full roster

(b) Quality against mean query price, ladder roster

Figure 2: GSM8K providers. Each point is one provider at its quality q_i and mean query price; the dotted line is the quality threshold $q_{\min}$, and the ring marks the provider with the lowest listed price. Monetary quantities are in units of 10^{-6} USD. On the full roster the lowest listed price belongs to i^*; on the ladder roster it is shared by two unqualified providers.

Table 4: GSM8K setting constants for the two rosters. Costs and prices are expected values per query.

Quantity	Full roster	Ladder roster
Providers N	9	12
Benchmark size K (queries in one pass)	1,319	1,319
Quality threshold $q_{\min}$	0.764	0.805
Qualified providers $\lvert\mathcal{Q}\rvert$	3	3
i^*	Qwen2.5-3B	Qwen2.5-7B
$\bar{c}_{i^*}$	16.1	34.4
$\bar{v}_{i^*}$	20.1	43.0
$\bar{c}_{(2)}$	34.4	68.7
$\bar{c}_{(2)} - \bar{c}_{i^*}$	18.3	34.4
Lowest listed price	Qwen2.5-3B	Llama-3.2-1B, Qwen2-1.5B
qualified	yes	no
equals i^*	yes	no
$c_{\max}$	102.4	300.0
$\sum_{i \notin \mathcal{Q}} M_i$	28,296	23,132
$\sum_{i \in \mathcal{Q} \setminus \{i^*\}} L_i$	3,464	6,012
$\lceil g(T) \rceil - 1$	1,656	1,334
$B_{\mathrm{id}}(T)$	31,977	30,010
T_{id}	63,740	59,735

Table 5: GSM8K providers, sorted by mean generation cost. Listed price is in USD per 10^6 output tokens; tokens are the mean number of generated tokens per query (summed over the n_i generations of a Best-of-n_i provider); query price and $\bar{c}_i$ are expected values per query. The gap is $\varepsilon_i = q_{\min} - q_i$ for an unqualified provider and $\Delta_i = \bar{c}_i - \bar{c}_{i^*}$ for a qualified one, with the corresponding selection bound M_i from equation 25 or L_i from equation 24; where the exploration-step cap $\lceil g(T) \rceil - 1$ at $T = 70{,}000$ exceeds L_i, the applicable cap $\max\{L_i, \lceil g(T) \rceil - 1\}$ of Lemma 9 follows it.

Provider	Listed price	Tokens	Query price	$\bar{c}_i$	q_i	Qualified	Gap	Bound
Full roster								
Llama-3.2-3B	0.08	196	15.7	12.5	0.692	no	$\varepsilon_i = 0.072$	$M = 15{,}141$
Qwen2-0.5B	0.1	181	18.1	14.5	0.351	no	$\varepsilon_i = 0.413$	$M = 299$
Qwen2-1.5B	0.1	189	18.9	15.1	0.574	no	$\varepsilon_i = 0.190$	$M = 1{,}770$
Llama-3.2-1B	0.1	199	19.9	15.9	0.364	no	$\varepsilon_i = 0.400$	$M = 324$
Qwen2.5-3B (i^*)	0.065	309	20.1	16.1	0.839	yes	–	–
Llama-3-8B	0.1455	192	27.9	22.3	0.650	no	$\varepsilon_i = 0.114$	$M = 5{,}466$
Llama-3.1-8B	0.1245	233	29.1	23.3	0.648	no	$\varepsilon_i = 0.116$	$M = 5{,}296$
Qwen2.5-7B	0.1465	293	43.0	34.4	0.907	yes	$\Delta_i = 18.3$	$L = 2{,}025$
Qwen2-7B	0.2	234	46.7	37.4	0.836	yes	$\Delta_i = 21.3$	$L = 1{,}439$; cap 1,656
Ladder roster								
Qwen2-1.5B	0.1	189	18.9	15.1	0.574	no	$\varepsilon_i = 0.231$	$M = 1{,}159$
Llama-3.2-1B	0.1	199	19.9	15.9	0.364	no	$\varepsilon_i = 0.441$	$M = 264$
Llama-3-8B	0.1455	192	27.9	22.3	0.650	no	$\varepsilon_i = 0.155$	$M = 2{,}809$
Qwen2-1.5B@2	0.1	379	37.9	30.3	0.652	no	$\varepsilon_i = 0.153$	$M = 2{,}910$
Llama-3.2-1B@2	0.1	398	39.8	31.8	0.410	no	$\varepsilon_i = 0.395$	$M = 340$
Qwen2.5-7B (i^*)	0.1465	293	43.0	34.4	0.907	yes	–	–
Llama-3-8B@2	0.1455	384	55.9	44.7	0.669	no	$\varepsilon_i = 0.136$	$M = 3{,}768$
Qwen2-1.5B@4	0.1	757	75.7	60.6	0.703	no	$\varepsilon_i = 0.102$	$M = 7{,}118$
Llama-3.2-1B@4	0.1	795	79.5	63.6	0.440	no	$\varepsilon_i = 0.365$	$M = 408$
Qwen2.5-7B@2	0.1465	587	85.9	68.7	0.923	yes	$\Delta_i = 34.4$	$L = 5{,}544$
Llama-3-8B@4	0.1455	768	111.7	89.4	0.677	no	$\varepsilon_i = 0.128$	$M = 4{,}356$
Qwen2.5-7B@4	0.1465	1,173	171.9	137.5	0.930	yes	$\Delta_i = 103.1$	$L = 468$; cap 1,334

F.2 ONE-PASS BEHAVIOR

Table 6: GSM8K one-pass outcomes, mean ± 95% half-width over 30 independent runs, paired across rules. Generation cost and amount paid are per query; the amount paid is the payment π_t for the platform and its ablation, and the query price for the other rules.

Policy	Unqualified share	R_T^{qual}/T	R_T^{gen}/T	Accuracy	Share of i^*	Generation cost	Amount paid
Full roster							
Platform	0.692 ± 0.003	0.099 ± 0.002	3.81 ± 0.02	0.692 ± 0.004	0.179 ± 0.003	18.7 ± 0.0	37.0 ± 0.3
Uniform among eligible	0.598 ± 0.007	0.094 ± 0.002	7.05 ± 0.12	0.707 ± 0.003	0.136 ± 0.004	22.5 ± 0.1	28.1 ± 0.2
Lowest listed price among eligible	0.005 ± 0.000	0.001 ± 0.000	0.04 ± 0.00	0.837 ± 0.002	0.994 ± 0.000	16.1 ± 0.0	20.1 ± 0.0
Lowest listed price, no quality filter	0.005 ± 0.000	0.001 ± 0.000	0.04 ± 0.00	0.837 ± 0.002	0.994 ± 0.000	16.1 ± 0.0	20.1 ± 0.0
No quality filter (ablation)	0.799 ± 0.003	0.182 ± 0.001	2.30 ± 0.03	0.600 ± 0.004	0.135 ± 0.002	17.2 ± 0.0	17.9 ± 0.1
Ladder roster							
Platform	0.771 ± 0.005	0.148 ± 0.003	13.69 ± 0.15	0.684 ± 0.003	0.125 ± 0.005	42.1 ± 0.4	118.1 ± 1.5
Uniform among eligible	0.690 ± 0.005	0.126 ± 0.003	24.81 ± 0.37	0.711 ± 0.003	0.103 ± 0.003	55.3 ± 0.5	69.2 ± 0.7
Lowest listed price among eligible	0.446 ± 0.031	0.087 ± 0.005	0.20 ± 0.00	0.775 ± 0.007	0.552 ± 0.031	28.2 ± 0.4	35.2 ± 0.5
Lowest listed price, no quality filter	0.998 ± 0.000	0.333 ± 0.001	0.20 ± 0.00	0.477 ± 0.004	0.001 ± 0.000	15.8 ± 0.1	19.8 ± 0.1
No quality filter (ablation)	0.855 ± 0.001	0.220 ± 0.000	7.57 ± 0.09	0.600 ± 0.003	0.091 ± 0.001	33.5 ± 0.1	36.0 ± 0.1

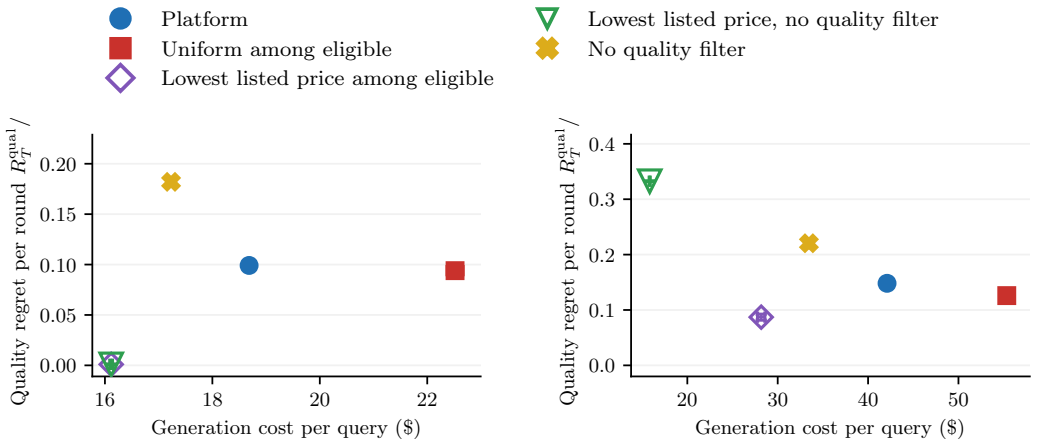

(a) Quality regret against generation cost, full roster

(b) Quality regret against generation cost, ladder roster

Figure 3: GSM8K one pass: quality regret per step against generation cost per query for every routing rule, mean with 95% confidence intervals across runs. Monetary quantities are in units of 10^{-6} USD. The two listed-price rules coincide on the full roster, where the lowest listed price is i^*.

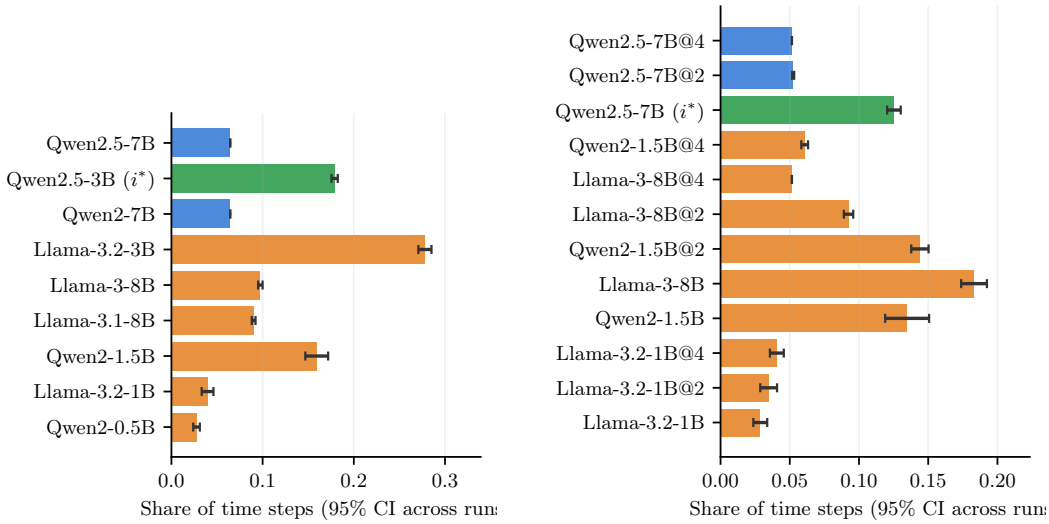

(a) Platform selection shares, full roster

(b) Platform selection shares, ladder roster

Figure 4: `GSM8K` one pass: share of steps the platform assigned to each provider, exploration steps included, mean with 95% confidence intervals across runs. Green marks i^*, blue the other qualified providers and orange the unqualified providers; the platform never observes these labels.

F.3 LONG-HORIZON LEARNING AND IDENTIFICATION

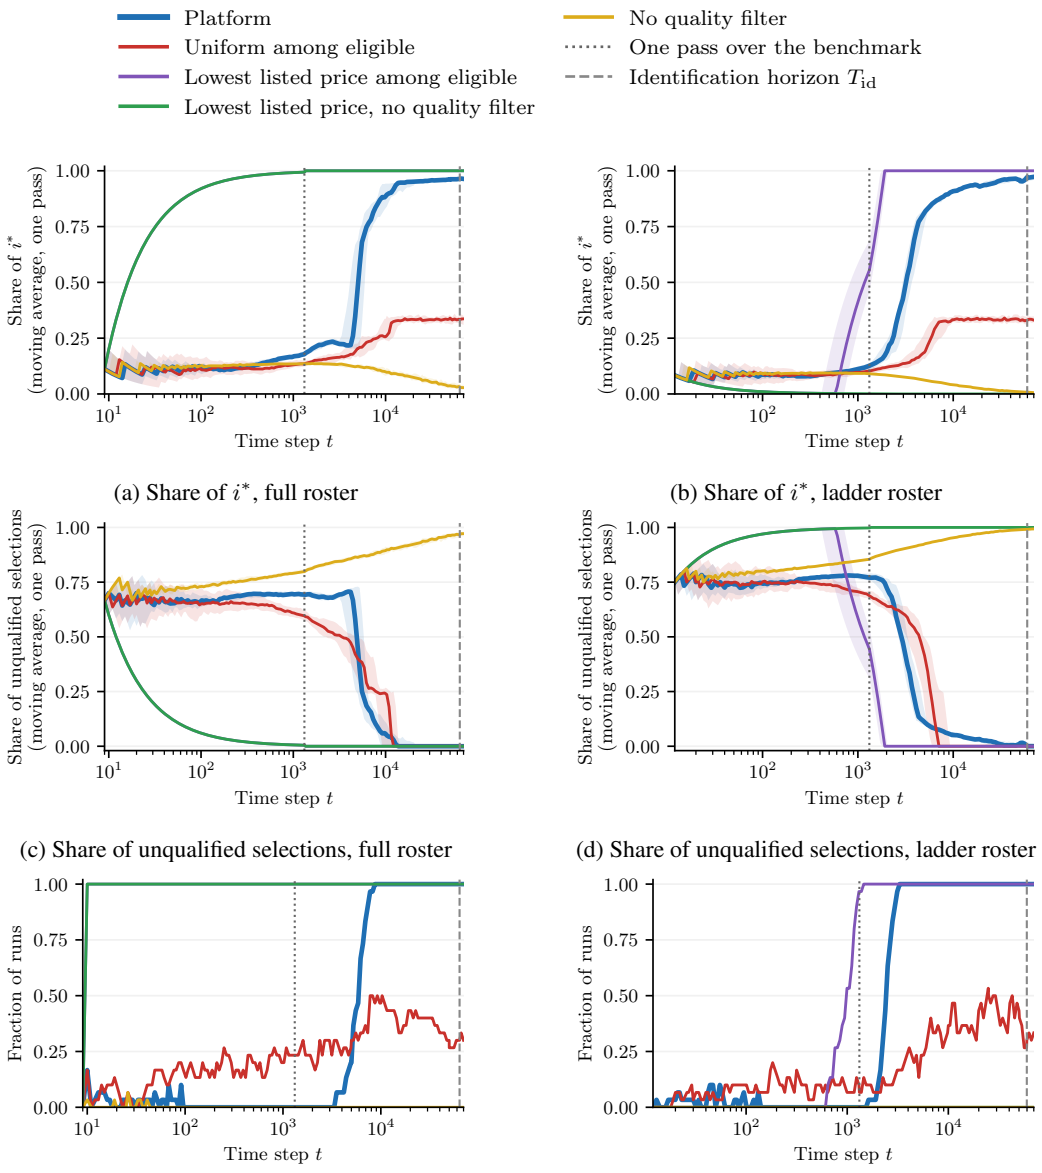

(a) Share of i^*, full roster

(b) Share of i^*, ladder roster

(c) Share of unqualified selections, full roster

(d) Share of unqualified selections, ladder roster

(e) i^* is the unique most selected provider, full roster

(f) i^* is the unique most selected provider, ladder roster

Figure 5: `GSM8K` long horizon: identification of i^* under every routing rule. Shares are moving averages with window size K, exploration steps included, median over 30 independent runs with a 10–90% band; the bottom row is the fraction of runs in which i^* has strictly more selections than every other provider. Vertical lines mark one pass and the identification horizon T_{id}. On the full roster the two listed-price rules coincide.

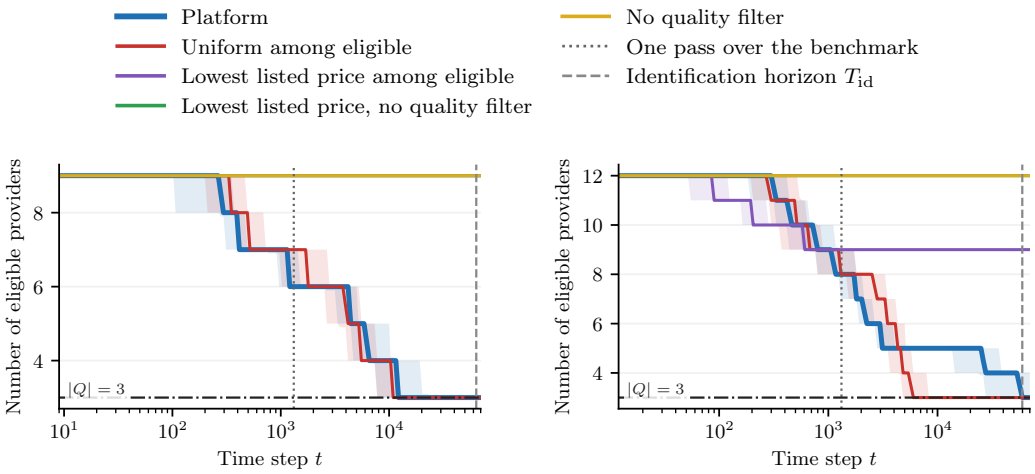

(a) Number of eligible providers, full roster

(b) Number of eligible providers, ladder roster

Figure 6: GSM8K long horizon: the number of providers not yet ruled out by the quality filter under every rule, median over 30 independent runs with a 10–90% band. The two rules without a filter stay at N; the listed-price rule inside a quality set learned with the same filter rules out only the providers it samples before settling on i^*; the dash-dotted line marks $|\mathcal{Q}|$.

Table 7: GSM8K long horizon: outcomes at $T = 70{,}000$, means over 30 independent runs. Shares of i^* and of unqualified selections are cumulative unless marked as moving averages and include the exploration steps; "i^* most selected" is the fraction of runs in which i^* is the unique most selected provider; $|\hat{\mathcal{Q}}_t|$ is the number of eligible providers at T. Generation cost and amount paid are cumulative per query; for the platform the amount paid includes the exploration steps.

Policy	Share of i^*	Moving avg.	Unqualified share	R_t^{qual}/t	R_t^{gen}/t	Generation cost	Amount paid	i^* most selected	$\lvert\hat{\mathcal{Q}}_t\rvert$
Full roster									
Platform	0.899	0.964	0.054	0.0056	1.077	17.1	28.2	1.00	3.0
Uniform among eligible	0.314	0.333	0.057	0.0059	12.580	28.6	35.7	0.30	3.0
Lowest listed price among eligible	1.000	1.000	0.000	0.0000	0.001	16.1	20.1	1.00	9.0
Lowest listed price, no quality filter	1.000	1.000	0.000	0.0000	0.001	16.1	20.1	1.00	9.0
No quality filter (ablation)	0.061	0.029	0.935	0.1466	0.219	13.8	13.8	0.00	9.0
Ladder roster									
Platform	0.914	0.970	0.048	0.0070	3.643	37.9	50.5	1.00	3.2
Uniform among eligible	0.319	0.334	0.044	0.0064	44.670	78.9	98.6	0.33	3.0
Lowest listed price among eligible	0.992	1.000	0.008	0.0016	0.004	34.3	42.8	1.00	9.0
Lowest listed price, no quality filter	0.000	0.000	1.000	0.3355	0.004	15.5	19.4	0.00	12.0
No quality filter (ablation)	0.022	0.007	0.974	0.2934	0.609	18.5	18.5	0.00	12.0

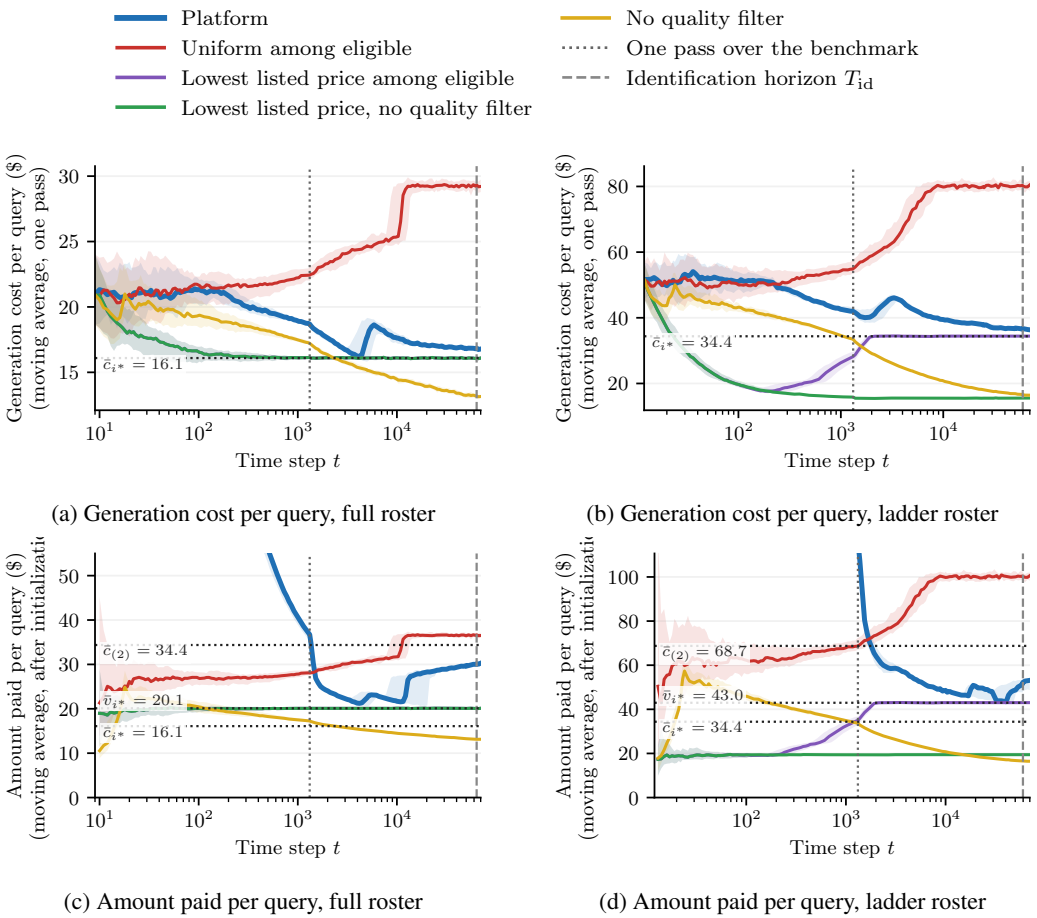

(a) Generation cost per query, full roster

(b) Generation cost per query, ladder roster

(c) Amount paid per query, full roster

(d) Amount paid per query, ladder roster

Figure 7: GSM8K long horizon: generation cost per query as a moving average with window size K (top) and amount paid per query as a moving average with window size K over the steps after initialization (bottom), under every rule, median over 30 independent runs with a 10–90% band. Monetary quantities are in units of 10^{-6} USD. The reference lines mark $\bar{c}_{i^*}$, the mean query price $\bar{v}_{i^*}$ of i^* and $\bar{c}_{(2)}$. The platform's payment π_t includes its exploration steps at $c_{\max}$, which are most of the steps early in the run; the vertical axis of the bottom panels is cut, so the platform's curve enters from above.

Table 8: GSM8K exploration steps of the platform, means over 30 independent runs, over one pass and over the long run of $T = 70{,}000$ queries. Shares are of the steps after initialization; payments and costs are per step of the corresponding horizon, the long-run values as moving averages with window size K at T. The last two rows give the platform's moving-average payment π_t on non-exploration steps and over all steps at T, per query.

Quantity	Full roster		Ladder roster	
	one pass	long run	one pass	long run
Exploration steps (mean)	314	3,458	422	3,348
Share of exploration steps after initialization	0.240	0.049	0.323	0.048
Share of exploration steps at T (moving avg.)	–	0.036	–	0.030
Exploration-step payment per step	24.40	3.73	95.92	9.14
Exploration-step generation cost per step	6.58	1.27	22.26	3.01
Share of exploration steps given to i^*	0.057	0.005	0.047	0.006
Share of exploration steps given to unqualified providers	0.408	0.037	0.641	0.307
Payment π_t at T, non-exploration steps (moving avg.)	–	27.52	–	44.27
Payment π_t at T, all steps (moving avg.)	–	30.25	–	52.09

F.4 Regret and theoretical bounds

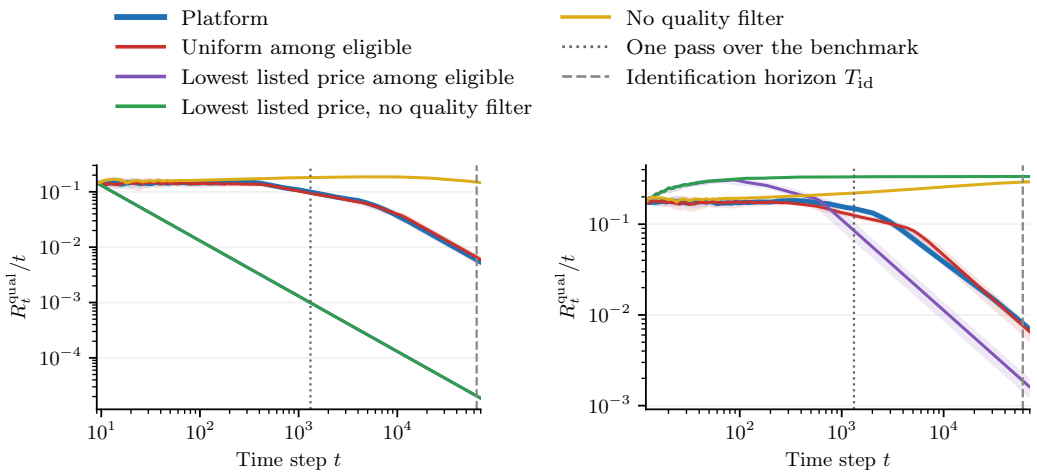

(a) Average quality regret, full roster

(b) Average quality regret, ladder roster

(c) Average generation regret, full roster

(d) Average generation regret, ladder roster

Figure 8: `GSM8K` long horizon: average quality regret R_t^{qual}/t and average generation regret R_t^{gen}/t of every routing rule, median over 30 independent runs with a 10–90% band, on logarithmic axes. Monetary quantities are in units of 10^{-6} USD. The platform's regrets include its exploration steps.

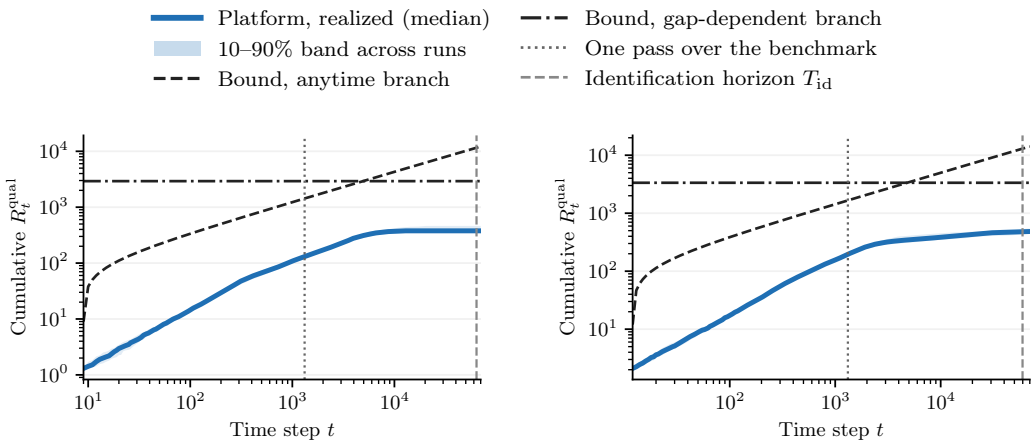

(a) Quality regret, full roster

(b) Quality regret, ladder roster

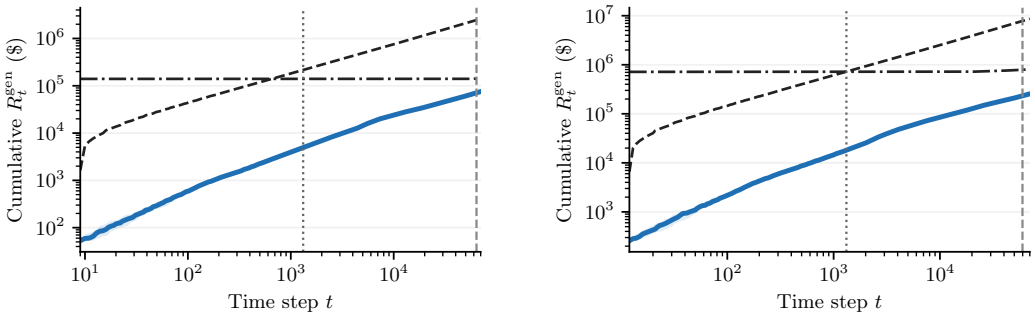

(c) Generation regret, full roster

(d) Generation regret, ladder roster

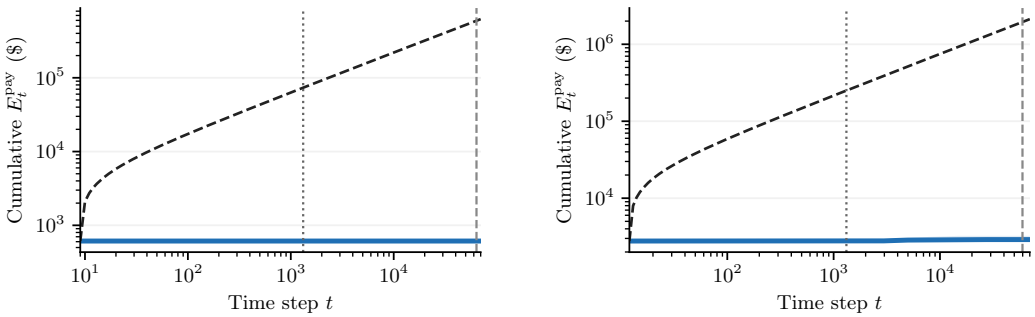

(e) Excess payment, full roster

(f) Excess payment, ladder roster

Figure 9: `GSM8K` long horizon: the platform's cumulative quality regret, generation regret and excess payment against their bounds, evaluated at every time step. Monetary quantities are in units of 10^{-6} USD. The bounds are those for the platform with exploration steps, each with an anytime branch and a gap-dependent branch; the excess-payment bound equation 37 without its exploration term has a single branch. The regrets include the exploration steps; the excess payment excludes them.

Table 9: GSM8K theoretical diagnostics of the platform over the long run ($T = 70{,}000$, 30 independent runs). Realized quantities are means over runs; each regret bound is the smaller of its two branches at T, for the platform with exploration steps. Selections and regrets include the exploration steps and are compared with $B_{\mathrm{id}}(T)$ and the bounds that account for them; the excess payment is given without the exploration steps, against equation 37 without its exploration term, and with them, against equation 37. The good event is checked at every selection count of every provider; the payment bounds are those of Appendix E, checked on non-exploration steps.

Quantity	Full roster	Ladder roster
$B_{\mathrm{id}}(T)$	31,977	30,010
T_{id}	63,740	59,735
Horizon T of the long run	70,000	70,000
Selections of providers other than i^* after initialization (mean)	7,094	6,034
ratio to $B_{\mathrm{id}}(T)$	0.222	0.201
Runs in which i^* is the most selected provider at $t = 5{,}000$	0.20	1.00
at $t = 10{,}000$	1.00	1.00
at $t = T_{\mathrm{id}}$	1.00	1.00
at $t = T$	1.00	1.00
R_T^{qual}, realized (mean)	392	487
bound equation 31	2,924	3,349
ratio	0.134	0.145
R_T^{gen}, realized (mean)	75,384	254,988
bound equation 34	144,428	805,387
ratio	0.522	0.317
E_T^{pay} without the exploration steps, realized (mean)	612	2,956
bound equation 37 without its exploration term	623,706	2,121,117
ratio	0.001	0.001
E_T^{pay} incl. exploration steps, realized (mean)	235,875	777,199
bound equation 37	1,637,592	5,823,569
ratio	0.144	0.133
Good event held, one pass (runs)	30 of 30	30 of 30
Good event held, long run, platform	30 of 30	30 of 30
Good event held, long run, no quality filter	30 of 30	30 of 30
Payment-bound violations, one pass (steps)	0	0

Table 10: `GSM8K` long horizon: selections of every provider other than i^* by the platform after initialization, over 30 independent runs of $T = 70{,}000$ queries, against its cap, M_i for an unqualified provider and $\max\{L_i, \lceil g(T) \rceil - 1\}$ for a qualified one (Lemma 9), shown as $\lceil g(T) \rceil - 1$ where the exploration-step cap exceeds L_i; the last row of each roster compares the total with $B_{\mathrm{id}}(T)$. Providers are sorted by mean generation cost, as in Table 5. Selections are the mean and the maximum over runs of the number of post-initialization selections, exploration steps included; the exploration column is the mean number of exploration steps among them; the last selection is the mean over runs of the step of the provider's last selection; the ratio is the mean number of selections divided by the cap.

Provider	Gap	Selections mean	Selections max	Exploration (mean)	Last selection (mean step)	Bound	Ratio
Full roster							
Llama-3.2-3B	$\varepsilon_i = 0.072$	2,008	3,240	16	4,182	$M = 15{,}141$	0.133
Qwen2-0.5B	$\varepsilon_i = 0.413$	35	59	16	280	$M = 299$	0.118
Qwen2-1.5B	$\varepsilon_i = 0.190$	221	361	17	1,184	$M = 1{,}770$	0.125
Llama-3.2-1B	$\varepsilon_i = 0.400$	51	142	17	391	$M = 324$	0.158
Llama-3-8B	$\varepsilon_i = 0.114$	745	1,271	30	8,845	$M = 5{,}466$	0.136
Llama-3.1-8B	$\varepsilon_i = 0.116$	721	1,156	32	11,280	$M = 5{,}296$	0.136
Qwen2.5-7B	$\Delta_i = 18.3$	1,656	1,656	1,656	69,977	$L = 2{,}025$	0.818
Qwen2-7B	$\Delta_i = 21.3$	1,656	1,656	1,656	69,978	$\lceil g(T) \rceil - 1 = 1{,}656$	1.000
All providers other than i^*		7,094		3,440		$B_{\mathrm{id}}(T) = 31{,}977$	0.222
Ladder roster							
Qwen2-1.5B	$\varepsilon_i = 0.231$	181	340	16	963	$M = 1{,}159$	0.157
Llama-3.2-1B	$\varepsilon_i = 0.441$	37	76	15	337	$M = 264$	0.140
Llama-3-8B	$\varepsilon_i = 0.155$	401	820	17	1,746	$M = 2{,}809$	0.143
Qwen2-1.5B@2	$\varepsilon_i = 0.153$	426	599	18	2,144	$M = 2{,}910$	0.146
Llama-3.2-1B@2	$\varepsilon_i = 0.395$	45	103	18	482	$M = 340$	0.132
Llama-3-8B@2	$\varepsilon_i = 0.136$	490	1,015	25	3,231	$M = 3{,}768$	0.130
Qwen2-1.5B@4	$\varepsilon_i = 0.102$	1,059	1,334	215	49,140	$M = 7{,}118$	0.149
Llama-3.2-1B@4	$\varepsilon_i = 0.365$	58	106	46	1,013	$M = 408$	0.141
Qwen2.5-7B@2	$\Delta_i = 34.4$	1,334	1,334	954	69,935	$L = 5{,}544$	0.241
Llama-3-8B@4	$\varepsilon_i = 0.128$	669	1,227	669	28,411	$M = 4{,}356$	0.154
Qwen2.5-7B@4	$\Delta_i = 103.1$	1,334	1,334	1,334	69,934	$\lceil g(T) \rceil - 1 = 1{,}334$	1.000
All providers other than i^*		6,034		3,328		$B_{\mathrm{id}}(T) = 30{,}010$	0.201

F.5 Payments and provider outcomes

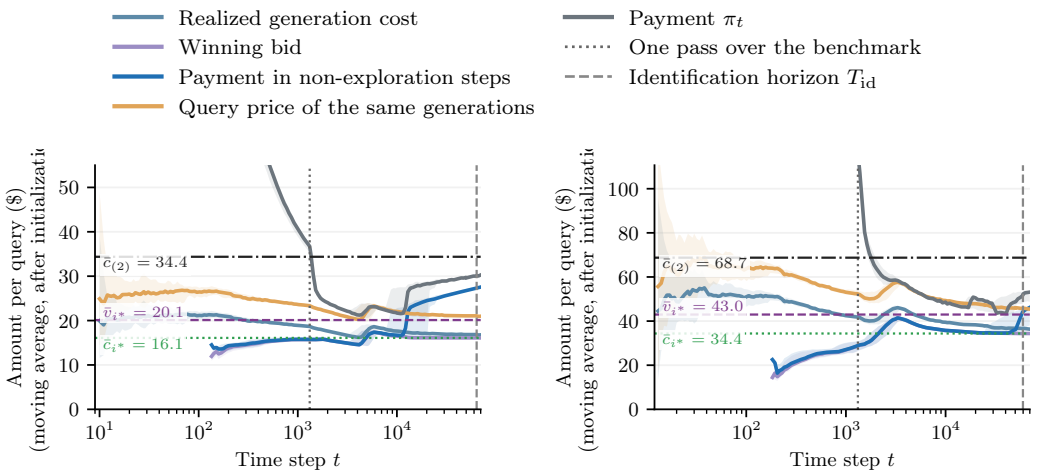

(a) Cost, bid, payment and query price, full roster

(b) Cost, bid, payment and query price, ladder roster

Figure 10: `GSM8K` long horizon: the platform's realized generation cost, the query price of the same generations and the payment π_t, as a moving average with window size K over the steps after initialization, and its winning bid and payment on non-exploration steps in the same window, median over 30 independent runs with a 10–90% band. Monetary quantities are in units of 10^{-6} USD. The dotted, dash-dotted and dashed lines mark $\bar{c}_{i^*}$, $\bar{c}_{(2)}$ and the mean query price $\bar{v}_{i^*}$ of i^*. The vertical axis is cut; π_t enters from above, since early in the run most steps are exploration steps and paid $c_{\max}$.

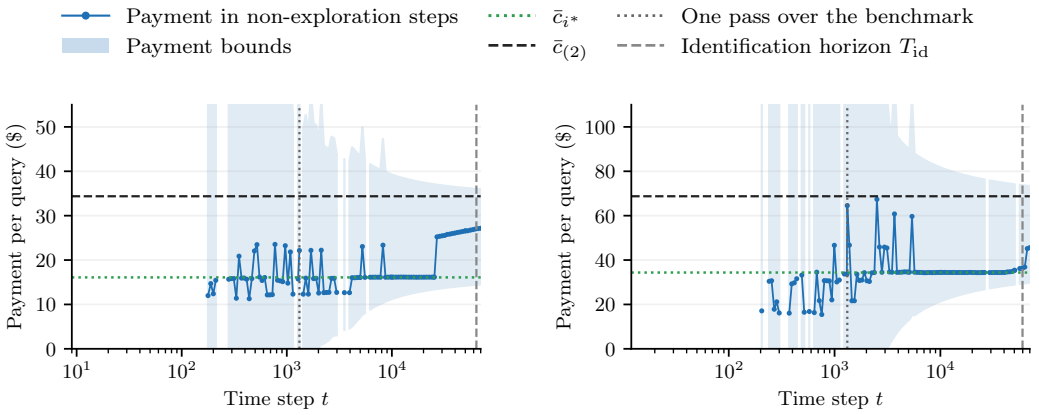

(a) Payment on non-exploration steps and its bounds, full roster

(b) Payment on non-exploration steps and its bounds, ladder roster

Figure 11: `GSM8K` long horizon: the payment π_t of the first run on non-exploration steps, sampled at logarithmically spaced time steps, inside the payment bounds of Appendix E for the provider selected in that step; sampled steps that were exploration steps are left blank. Monetary quantities are in units of 10^{-6} USD. The bounds depend on the selected provider's cost and selection count, and close in as the counts grow.

Table 11: GSM8K one pass: the platform's economic quantities, mean $\pm$ 95% half-width over 30 independent runs. The payment π_t is over all steps; the winning bid, the payment on non-exploration steps and the second excess-payment row are over non-exploration steps only. A provider's payoff is its total payment minus its total generation cost over the pass.

Quantity	Full roster	Ladder roster
Generation cost per query, whole pass	18.68 ± 0.05	42.10 ± 0.36
Query price per query, whole pass	23.35 ± 0.06	52.62 ± 0.44
Payment π_t per query, whole pass	37.03 ± 0.29	118.10 ± 1.49
Generation cost per query, after initialization	18.67 ± 0.05	42.01 ± 0.36
Query price per query, after initialization	23.34 ± 0.06	52.51 ± 0.45
Payment π_t per query, after initialization	36.58 ± 0.29	116.43 ± 1.51
Winning bid, non-exploration steps	15.75 ± 0.07	28.77 ± 0.75
Payment π_t per query, non-exploration steps	15.80 ± 0.07	28.93 ± 0.75
Exploration-step payment per query, after initialization	24.40 ± 0.36	95.92 ± 1.98
Share of exploration steps after initialization	0.240 ± 0.004	0.323 ± 0.007
Excess payment per step E_T^{pay}/T, whole pass	16.676 ± 0.238	76.051 ± 1.524
Excess payment per query, non-exploration steps	0.000 ± 0.000	0.002 ± 0.004
Mean provider payoff over the pass	2688.6 ± 43.6	8353.3 ± 170.9
Lowest provider payoff over providers and runs	903.8	2814.4

Table 12: GSM8K one pass: each provider's selections and payoff under the platform over 30 independent runs, exploration steps included, with providers sorted by mean generation cost as in Table 5. The last column counts the runs in which the provider's payoff over the pass was negative.

Provider	Selections (mean)	Payoff (mean)	Payoff (min)	Payoff (max)	Negative runs
Full roster					
Llama-3.2-3B	366.6	1506.2	1264.8	1675.0	0
Qwen2-0.5B	36.3	1516.9	903.8	1872.4	0
Qwen2-1.5B	210.2	1576.2	1257.8	1877.1	0
Llama-3.2-1B	52.3	1584.3	1166.0	1943.0	0
Qwen2.5-3B (i^*)	236.0	1618.9	1243.1	1848.1	0
Llama-3-8B	128.6	2464.5	1764.6	3205.6	0
Llama-3.1-8B	118.9	2613.9	1708.6	3915.3	0
Qwen2.5-7B	85.0	5795.4	5587.7	5932.6	0
Qwen2-7B	85.0	5520.5	5261.4	5733.9	0
Ladder roster					
Qwen2-1.5B	177.8	4814.9	4303.1	4955.2	0
Llama-3.2-1B	37.8	4671.6	2814.4	4881.2	0
Llama-3-8B	241.5	4908.6	4503.1	5163.9	0
Qwen2-1.5B@2	189.9	5156.8	4516.7	5882.3	0
Llama-3.2-1B@2	45.8	5225.4	4016.8	5863.3	0
Qwen2.5-7B (i^*)	165.2	5567.6	4594.4	6821.7	0
Llama-3-8B@2	121.9	6622.6	5220.5	7875.4	0
Qwen2-1.5B@4	80.2	11988.5	6944.9	16058.1	0
Llama-3.2-1B@4	53.6	10912.0	3421.6	16011.6	0
Qwen2.5-7B@2	69.2	15030.6	10862.7	15919.4	0
Llama-3-8B@4	68.0	14278.0	13922.0	14696.6	0
Qwen2.5-7B@4	68.0	11062.7	10464.0	11551.1	0

G GPQA WEAK-COMPETITION EXPERIMENTS

This section reports the complete GPQA results for the full and ladder rosters, in the same form as the GSM8K section; Appendix E provides the experimental details, and the platform runs with the exploration rule described there. One-pass results are averaged over 30 independent runs, paired across rules and reported with 95% confidence intervals, while long-horizon trajectories report the median and 10–90% range over 30 independent runs. Monetary quantities are per query unless stated otherwise. A moving average has window size K, one pass over the benchmark, *i.e.*, 546 steps on the full roster and 545 on the ladder roster, where the reward alignment drops one question. The amount paid follows the convention of Appendix F. On both rosters the identification horizon T_{id} lies beyond the long run of $T = 70{,}000$ queries, so the long-horizon results describe learning before that horizon, and no vertical line marks it in the figures. On both rosters $\bar{c}_{(2)}$ lies above the mean query price $\bar{v}_{i^*}$ of i^*, as on GSM8K, but here unqualified providers remain in $\hat{\mathcal{Q}}_t$ throughout the run, so the payment on non-exploration steps does not approach $\bar{c}_{(2)}$; the strong-competition roster of Appendix H is the GPQA setting in which it does.

G.1 PROVIDERS AND ENVIRONMENT

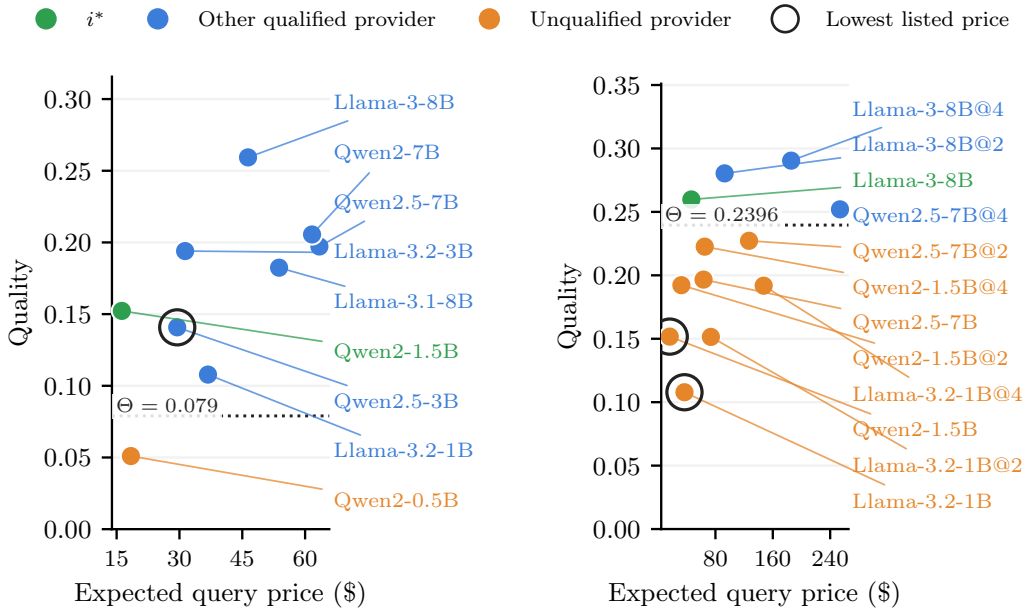

(a) Quality against mean query price, full roster

(b) Quality against mean query price, ladder roster

Figure 12: GPQA providers. Each point is one provider at its quality q_i and mean query price; the dotted line is the quality threshold $q_{\min}$, and the ring marks the provider with the lowest listed price. Monetary quantities are in units of 10^{-6} USD. On the full roster the lowest listed price belongs to a qualified provider that is not i^*; on the ladder roster it is shared by two unqualified providers.

Table 13: GPQA setting constants for the two rosters. Costs and prices are mean values per query.

Quantity	Full roster	Ladder roster
Providers N	9	12
Benchmark size K (queries in one pass)	546	545
Quality threshold $q_{\min}$	0.079	0.2396
Qualified providers $\lvert\mathcal{Q}\rvert$	8	4
i^*	Qwen2-1.5B	Llama-3-8B
$\bar{c}_{i^*}$	13.0	37.1
$\bar{v}_{i^*}$	16.2	46.4
$\bar{c}_{(2)}$	23.6	74.2
$\bar{c}_{(2)} - \bar{c}_{i^*}$	10.6	37.1
Lowest listed price	Qwen2.5-3B	Llama-3.2-1B, Qwen2-1.5B
qualified	yes	no
equals i^*	no	no
$c_{\max}$	102.4	300.0
$\sum_{i \notin \mathcal{Q}} M_i$	120,055	1,211,234
$\sum_{i \in \mathcal{Q} \setminus \{i^*\}} L_i$	17,090	5,222
$\lceil g(T) \rceil - 1$	1,656	1,334
$B_{\mathrm{id}}(T)$	141,210	1,218,576
T_{id}	315,294	2,540,927

Table 14: GPQA providers, sorted by mean generation cost. Listed price is in USD per 10^6 output tokens; tokens are the mean number of generated tokens per query (summed over the n_i generations of a Best-of-n_i provider); query price and $\bar{c}_i$ are mean values per query. The gap is $\varepsilon_i = q_{\min} - q_i$ for an unqualified provider and $\Delta_i = \bar{c}_i - \bar{c}_{i^*}$ for a qualified one, with the corresponding selection bound M_i from equation 25 or L_i from equation 24; where the exploration-step cap $\lceil g(T) \rceil - 1$ at $T = 70{,}000$ exceeds L_i, the applicable cap $\max\{L_i, \lceil g(T) \rceil - 1\}$ of Lemma 9 follows it.

Provider	Listed price	Tokens	Query price	$\bar{c}_i$	q_i	Qualified	Gap	Bound
Full roster								
Qwen2-1.5B (i^*)	0.1	162	16.2	13.0	0.152	yes	–	–
Qwen2-0.5B	0.1	184	18.4	14.7	0.051	no	$\varepsilon_i = 0.028$	$M = 120{,}055$
Qwen2.5-3B	0.065	453	29.5	23.6	0.141	yes	$\Delta_i = 10.6$	$L = 6{,}861$
Llama-3.2-3B	0.08	392	31.3	25.1	0.194	yes	$\Delta_i = 12.1$	$L = 5{,}106$
Llama-3.2-1B	0.1	368	36.8	29.4	0.108	yes	$\Delta_i = 16.5$	$L = 2{,}564$
Llama-3-8B	0.1455	319	46.4	37.1	0.259	yes	$\Delta_i = 24.1$	$L = 1{,}085$; cap 1,656
Llama-3.1-8B	0.1245	432	53.8	43.0	0.182	yes	$\Delta_i = 30.0$	$L = 658$; cap 1,656
Qwen2-7B	0.2	308	61.7	49.3	0.206	yes	$\Delta_i = 36.3$	$L = 426$; cap 1,656
Qwen2.5-7B	0.1465	433	63.4	50.7	0.197	yes	$\Delta_i = 37.7$	$L = 390$; cap 1,656
Ladder roster								
Qwen2-1.5B	0.1	163	16.3	13.0	0.152	no	$\varepsilon_i = 0.088$	$M = 9{,}964$
Qwen2-1.5B@2	0.1	325	32.5	26.0	0.192	no	$\varepsilon_i = 0.047$	$M = 38{,}751$
Llama-3.2-1B	0.1	368	36.8	29.5	0.108	no	$\varepsilon_i = 0.132$	$M = 4{,}063$
Llama-3-8B (i^*)	0.1455	319	46.4	37.1	0.260	yes	–	–
Qwen2.5-7B	0.1465	433	63.4	50.7	0.197	no	$\varepsilon_i = 0.043$	$M = 47{,}908$
Qwen2-1.5B@4	0.1	651	65.1	52.1	0.223	no	$\varepsilon_i = 0.017$	$M = 357{,}226$
Llama-3.2-1B@2	0.1	736	73.6	58.9	0.151	no	$\varepsilon_i = 0.088$	$M = 9{,}906$
Llama-3-8B@2	0.1455	638	92.8	74.2	0.280	yes	$\Delta_i = 37.1$	$L = 4{,}674$
Qwen2.5-7B@2	0.1465	866	126.9	101.5	0.227	no	$\varepsilon_i = 0.012$	$M = 705{,}391$
Llama-3.2-1B@4	0.1	1,473	147.3	117.8	0.192	no	$\varepsilon_i = 0.048$	$M = 38{,}025$
Llama-3-8B@4	0.1455	1,276	185.6	148.5	0.290	yes	$\Delta_i = 111.4$	$L = 392$; cap 1,334
Qwen2.5-7B@4	0.1465	1,732	253.7	203.0	0.252	yes	$\Delta_i = 165.8$	$L = 156$; cap 1,334

G.2 ONE-PASS BEHAVIOR

Table 15: GPQA one-pass outcomes, mean ± 95% half-width over 30 independent runs, paired across rules. Generation cost and amount paid are per query; the amount paid is the payment π_t for the platform and its ablation, and the query price for the other rules. Within one pass the quality filter rules out no provider on either roster, so the two listed-price rules coincide.

Policy	Unqualified share	R_T^{qual}/T	R_T^{gen}/T	Accuracy	Share of i^*	Generation cost	Amount paid
Full roster							
Platform	0.185 ± 0.006	0.005 ± 0.000	14.22 ± 0.04	0.154 ± 0.004	0.214 ± 0.006	27.4 ± 0.2	64.4 ± 0.2
Uniform among eligible	0.112 ± 0.005	0.003 ± 0.000	18.84 ± 0.22	0.161 ± 0.005	0.109 ± 0.005	31.9 ± 0.3	39.9 ± 0.4
Lowest listed price among eligible	0.002 ± 0.000	0.000 ± 0.000	10.72 ± 0.00	0.143 ± 0.004	0.002 ± 0.000	23.7 ± 0.0	29.6 ± 0.1
Lowest listed price, no quality filter	0.002 ± 0.000	0.000 ± 0.000	10.72 ± 0.00	0.143 ± 0.004	0.002 ± 0.000	23.7 ± 0.0	29.6 ± 0.1
No quality filter (ablation)	0.233 ± 0.007	0.007 ± 0.000	9.73 ± 0.14	0.139 ± 0.004	0.279 ± 0.008	22.6 ± 0.2	24.0 ± 0.3
Ladder roster							
Platform	0.715 ± 0.001	0.047 ± 0.000	32.96 ± 0.04	0.201 ± 0.006	0.092 ± 0.001	64.4 ± 0.2	196.5 ± 0.8
Uniform among eligible	0.664 ± 0.008	0.039 ± 0.001	42.11 ± 0.85	0.211 ± 0.005	0.087 ± 0.005	75.9 ± 1.0	94.9 ± 1.2
Lowest listed price among eligible	0.993 ± 0.000	0.108 ± 0.000	0.94 ± 0.00	0.130 ± 0.003	0.002 ± 0.000	22.3 ± 0.2	27.9 ± 0.3
Lowest listed price, no quality filter	0.993 ± 0.000	0.108 ± 0.000	0.94 ± 0.00	0.130 ± 0.003	0.002 ± 0.000	22.3 ± 0.2	27.9 ± 0.3
No quality filter (ablation)	0.794 ± 0.002	0.056 ± 0.000	15.94 ± 0.20	0.193 ± 0.005	0.113 ± 0.001	45.4 ± 0.3	51.0 ± 0.5

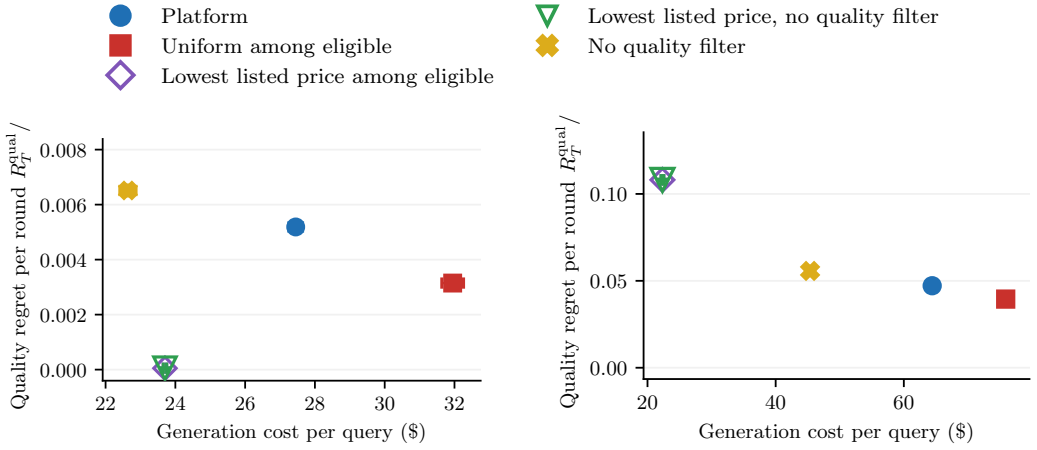

(a) Quality regret against generation cost, full roster

(b) Quality regret against generation cost, ladder roster

Figure 13: GPQA one pass: quality regret per step against generation cost per query for every routing rule, mean with 95% confidence intervals across runs. Monetary quantities are in units of 10^{-6} USD. On both rosters the two listed-price rules coincide, since the quality filter rules out no provider within one pass.

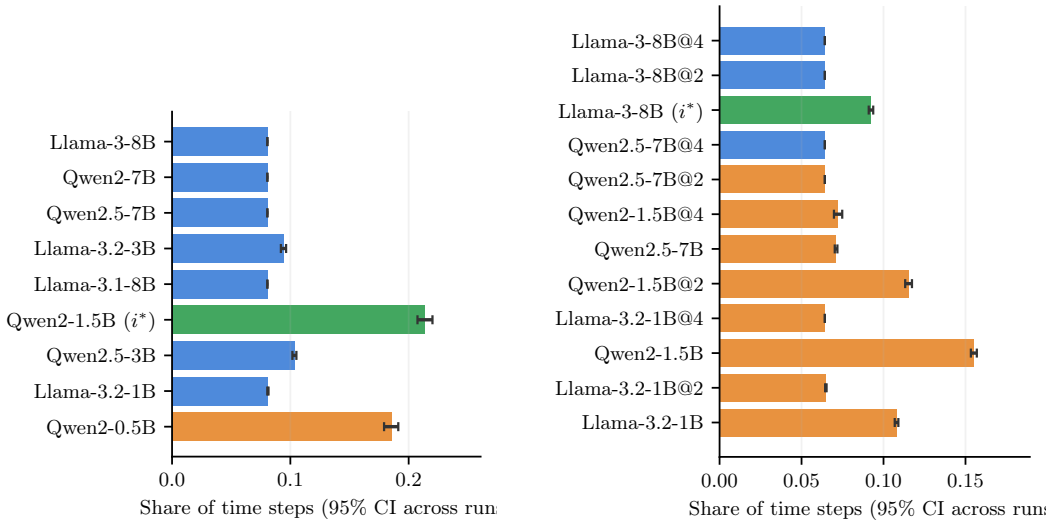

(a) Platform selection shares, full roster

(b) Platform selection shares, ladder roster

Figure 14: GPQA one pass: share of steps the platform assigned to each provider, exploration steps included, mean with 95% confidence intervals across runs. Green marks i^*, blue the other qualified providers and orange the unqualified providers; the platform never observes these labels.

G.3 Long-horizon learning and identification

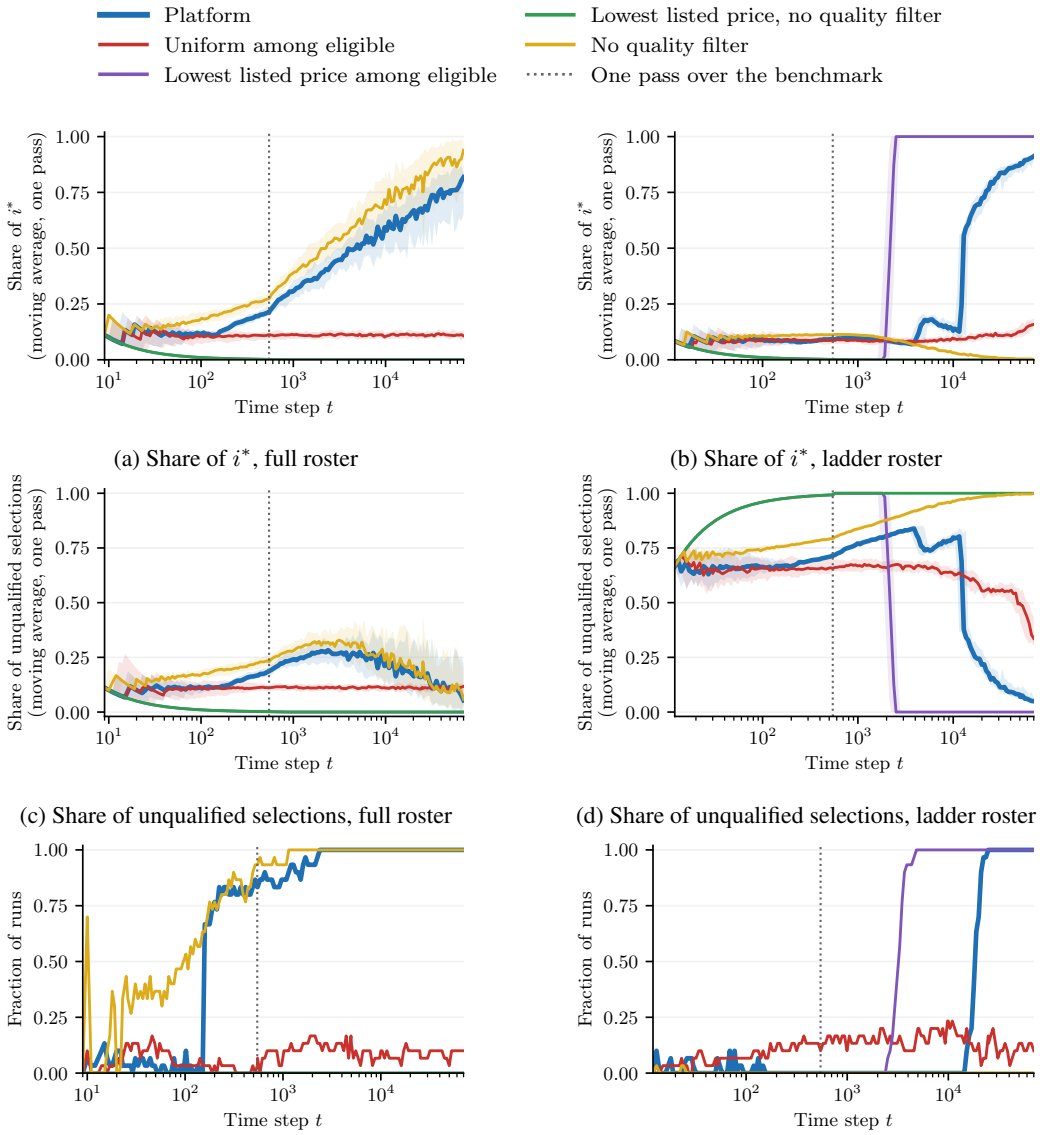

(a) Share of i^*, full roster

(b) Share of i^*, ladder roster

(c) Share of unqualified selections, full roster

(d) Share of unqualified selections, ladder roster

(e) i^* is the unique most selected provider, full roster

(f) i^* is the unique most selected provider, ladder roster

Figure 15: GPQA long horizon: identification of i^* under every routing rule. Shares are moving averages with window size K, exploration steps included, median over 30 independent runs with a 10–90% band; the bottom row is the fraction of runs in which i^* has strictly more selections than every other provider. The vertical line marks one pass over the benchmark; the identification horizon lies beyond T. On the full roster the two listed-price rules coincide.

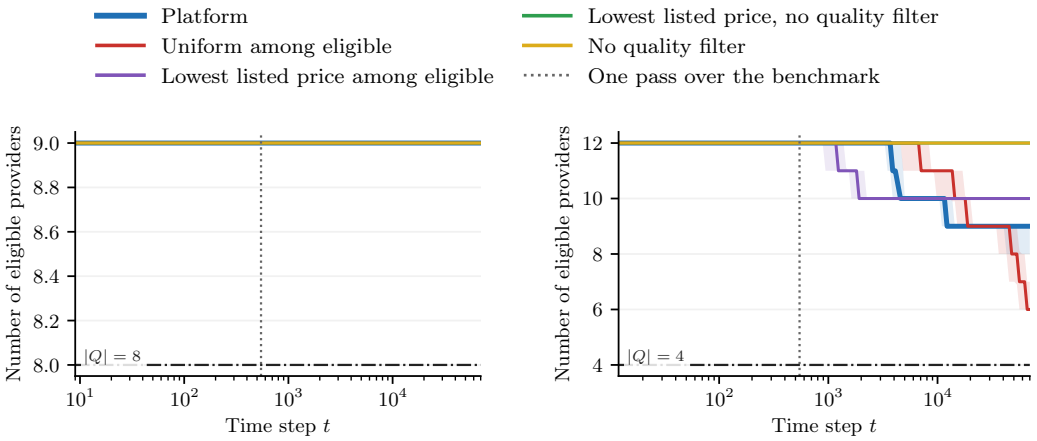

(a) Number of eligible providers, full roster

(b) Number of eligible providers, ladder roster

Figure 16: GPQA long horizon: the number of providers not yet ruled out by the quality filter under every rule, median over 30 independent runs with a 10–90% band. On the full roster the single unqualified provider is never ruled out, so every curve stays at N; on the ladder roster the two rules without a filter stay at N, and the dash-dotted line marks $|\mathcal{Q}|$.

Table 16: GPQA long horizon: outcomes at $T = 70{,}000$, means over 30 independent runs. Shares of i^* and of unqualified selections are cumulative unless marked as moving averages and include the exploration steps; "i^* most selected" is the fraction of runs in which i^* is the unique most selected provider; $|\hat{\mathcal{Q}}_t|$ is the number of eligible providers at T. Generation cost and amount paid are cumulative per query; for the platform the amount paid includes the exploration steps.

Policy	Share of i^*	Moving avg.	Unqualified share	R_t^{qual}/t	R_t^{gen}/t	Generation cost	Amount paid	i^* most selected	$\lvert\hat{\mathcal{Q}}_t\rvert$
Full roster									
Platform	0.676	0.779	0.158	0.0044	4.234	17.2	27.4	1.00	9.0
Uniform among eligible	0.111	0.111	0.111	0.0031	18.785	31.8	39.7	0.10	9.0
Lowest listed price among eligible	0.000	0.000	0.000	0.0000	10.582	23.6	29.5	0.00	9.0
Lowest listed price, no quality filter	0.000	0.000	0.000	0.0000	10.582	23.6	29.5	0.00	9.0
No quality filter (ablation)	0.801	0.913	0.172	0.0048	0.709	13.7	13.7	1.00	9.0
Ladder roster									
Platform	0.701	0.909	0.242	0.0117	10.288	46.0	63.2	1.00	8.5
Uniform among eligible	0.117	0.161	0.533	0.0196	54.062	89.7	112.2	0.10	6.3
Lowest listed price among eligible	0.972	1.000	0.027	0.0028	0.007	36.6	45.7	1.00	10.0
Lowest listed price, no quality filter	0.000	0.000	1.000	0.1098	0.007	21.2	26.6	0.00	12.0
No quality filter (ablation)	0.017	0.003	0.980	0.0847	0.635	15.8	15.9	0.00	12.0

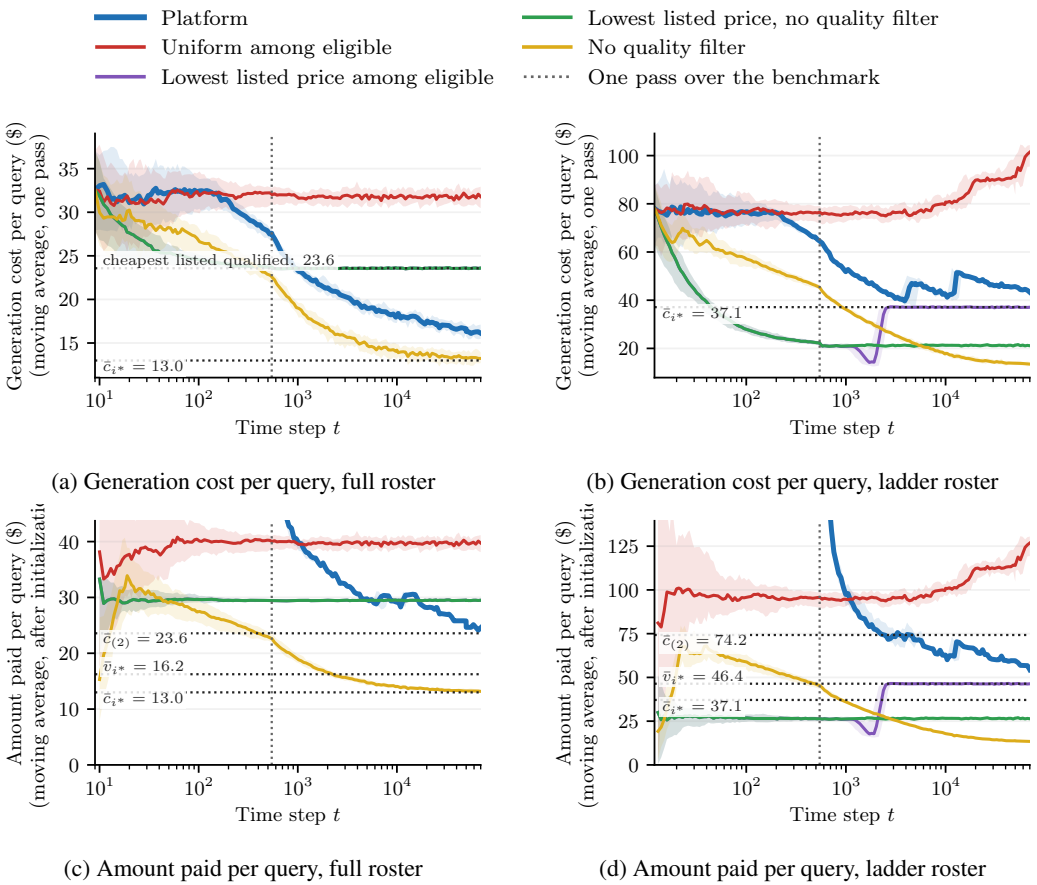

(a) Generation cost per query, full roster

(b) Generation cost per query, ladder roster

(c) Amount paid per query, full roster

(d) Amount paid per query, ladder roster

Figure 17: GPQA long horizon: generation cost per query as a moving average with window size K (top) and amount paid per query as a moving average with window size K over the steps after initialization (bottom), under every rule, median over 30 independent runs with a 10–90% band. Monetary quantities are in units of 10^{-6} USD. The reference lines mark $\bar{c}_{i^*}$, the mean query price $\bar{v}_{i^*}$ of i^*, $\bar{c}_{(2)}$ and, on the full roster, the cost of the qualified provider with the lowest listed price. The platform's payment π_t includes its exploration steps at $c_{\max}$; the vertical axis of the bottom panels is cut, so the platform's curve enters from above.

Table 17: GPQA exploration steps of the platform, means over 30 independent runs, over one pass and over the long run of $T = 70{,}000$ queries. Shares are of the steps after initialization; payments and costs are per step of the corresponding horizon, the long-run values as moving averages with window size K at T. The last two rows give the platform's moving-average payment π_t on non-exploration steps and over all steps at T, per query.

Quantity	Full roster		Ladder roster	
	one pass	long run	one pass	long run
Exploration steps (mean)	295	10,967	326	6,916
Share of exploration steps after initialization	0.550	0.157	0.612	0.099
Share of exploration steps at T (moving avg.)	–	0.128	–	0.066
Exploration-step payment per step	55.39	13.13	179.69	19.78
Exploration-step generation cost per step	19.62	4.68	52.58	8.38
Share of exploration steps given to i^*	0.055	0.001	0.059	0.003
Share of exploration steps given to unqualified providers	0.057	0.002	0.629	0.419
Payment π_t at T, non-exploration steps (moving avg.)	–	13.19	–	37.49
Payment π_t at T, all steps (moving avg.)	–	24.63	–	54.80

G.4 Regret and theoretical bounds

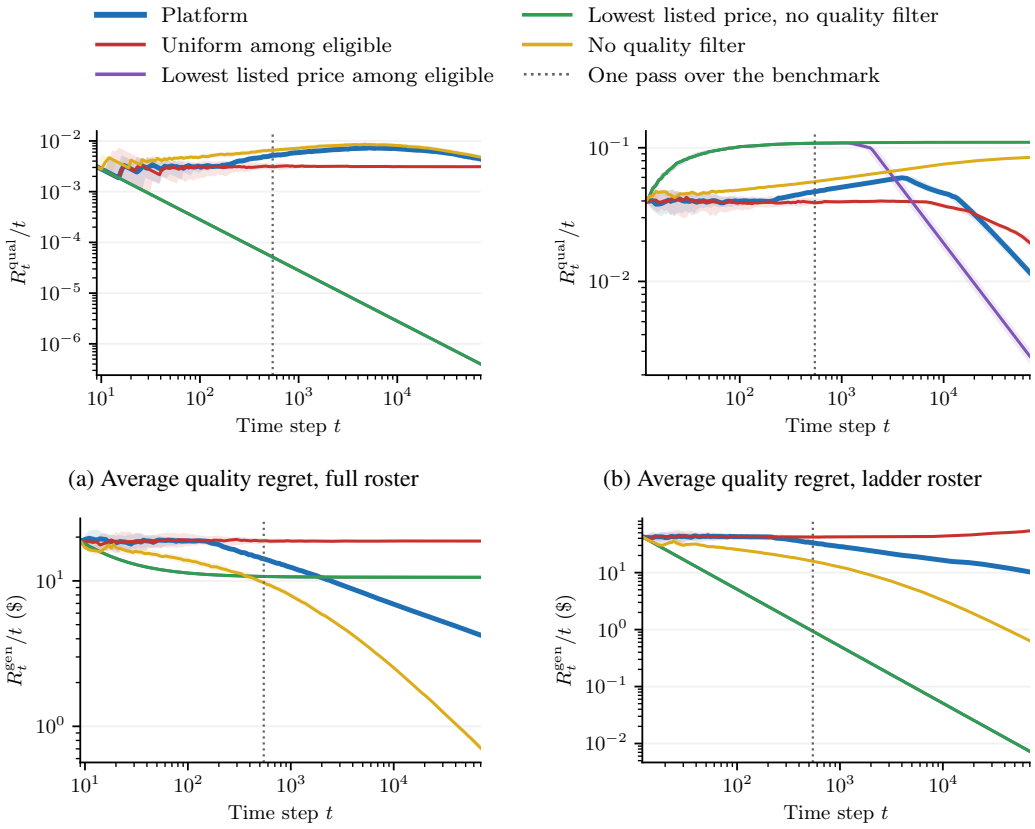

(a) Average quality regret, full roster

(b) Average quality regret, ladder roster

(c) Average generation regret, full roster

(d) Average generation regret, ladder roster

Figure 18: GPQA long horizon: average quality regret R_t^{qual}/t and average generation regret R_t^{gen}/t of every routing rule, median over 30 independent runs with a 10–90% band, on logarithmic axes. Monetary quantities are in units of 10^{-6} USD. The platform's regrets include its exploration steps.

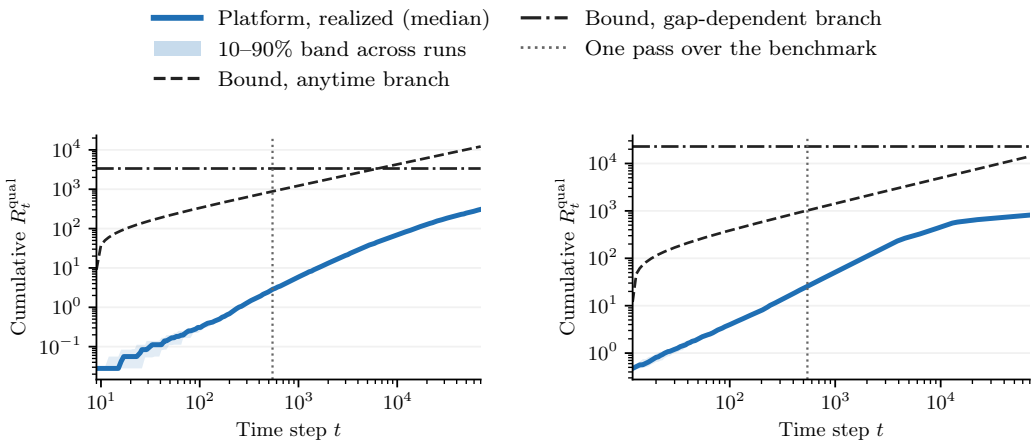

(a) Quality regret, full roster

(b) Quality regret, ladder roster

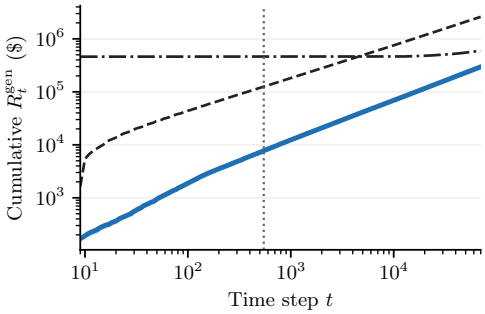

(c) Generation regret, full roster

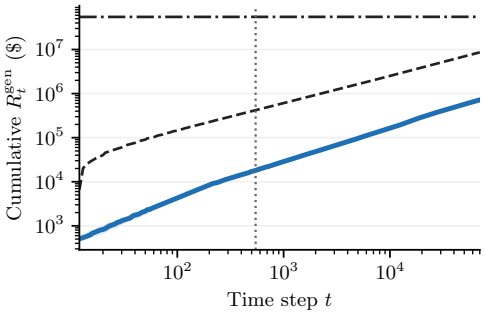

(d) Generation regret, ladder roster

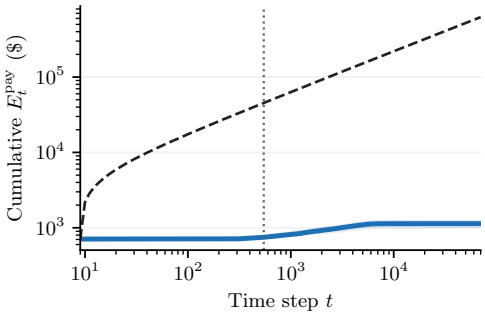

(e) Excess payment, full roster

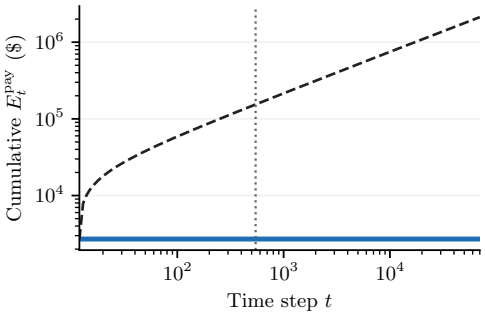

(f) Excess payment, ladder roster

Figure 19: GPQA long horizon: the platform's cumulative quality regret, generation regret and excess payment against their bounds, evaluated at every time step. Monetary quantities are in units of 10^{-6} USD. The bounds are those for the platform with exploration steps, Theorems 7 and 8, each with an anytime branch and a gap-dependent branch; the excess-payment bound equation 37 without its exploration term has a single branch. The regrets include the exploration steps; the excess payment excludes them.

Table 18: GPQA theoretical diagnostics of the platform over the long run ($T = 70{,}000$, 30 independent runs). Realized quantities are means over runs; each regret bound is the smaller of its two branches at T, for the platform with exploration steps. Selections and regrets include the exploration steps and are compared with $B_{\mathrm{id}}(T)$ and the bounds that account for them; the excess payment is given without the exploration steps, against equation 37 without its exploration term, and with them, against equation 37. The identification horizon T_{id} lies beyond T on both rosters, so the corresponding row is empty. The good event is checked at every selection count of every provider; the payment bounds are those of Appendix E, checked on non-exploration steps.

Quantity	Full roster	Ladder roster
$B_{\mathrm{id}}(T)$	141,210	1,218,576
T_{id}	315,294	2,540,927
Horizon T of the long run	70,000	70,000
Selections of providers other than i^* after initialization (mean)	22,641	20,907
ratio to $B_{\mathrm{id}}(T)$	0.160	0.017
Runs in which i^* is the most selected provider at $t = 5{,}000$	1.00	0.00
at $t = 10{,}000$	1.00	0.00
at $t = T_{\mathrm{id}}$	–	–
at $t = T$	1.00	1.00
R_T^{qual}, realized (mean)	309	818
bound equation 31	3,362	14,133
ratio	0.092	0.058
R_T^{gen}, realized (mean)	296,412	720,159
bound equation 34	596,889	8,642,653
ratio	0.497	0.083
E_T^{pay} without the exploration steps, realized (mean)	1,149	2,710
bound equation 37 without its exploration term	623,803	2,121,051
ratio	0.002	0.001
E_T^{pay} incl. exploration steps, realized (mean)	865,643	1,564,346
bound equation 37	1,798,620	5,735,545
ratio	0.481	0.273
Good event held, one pass (runs)	30 of 30	30 of 30
Good event held, long run, platform	30 of 30	30 of 30
Good event held, long run, no quality filter	30 of 30	30 of 30
Payment-bound violations, one pass (steps)	0	0

Table 19: GPQA long horizon: selections of every provider other than i^* by the platform after initialization, over 30 independent runs of $T = 70{,}000$ queries, against its cap, M_i for an unqualified provider and $\max\{L_i, \lceil g(T) \rceil - 1\}$ for a qualified one (Lemma 9), shown as $\lceil g(T) \rceil - 1$ where the exploration-step cap exceeds L_i; the last row of each roster compares the total with $B_{\mathrm{id}}(T)$. Providers are sorted by mean generation cost, as in Table 14. Selections are the mean and the maximum over runs of the number of post-initialization selections, exploration steps included; the exploration column is the mean number of exploration steps among them; the last selection is the mean over runs of the step of the provider's last selection; the ratio is the mean number of selections divided by the cap.

Provider	Gap	Selections mean	Selections max	Exploration (mean)	Last selection (mean step)	Bound	Ratio
Full roster							
Qwen2-0.5B	$\varepsilon_i = 0.028$	11,049	12,859	17	69,838	$M = 120{,}055$	0.092
Qwen2.5-3B	$\Delta_i = 10.6$	1,656	1,656	1,255	69,981	$L = 6{,}861$	0.241
Llama-3.2-3B	$\Delta_i = 12.1$	1,656	1,656	1,406	69,980	$L = 5{,}106$	0.324
Llama-3.2-1B	$\Delta_i = 16.5$	1,656	1,656	1,649	69,980	$L = 2{,}564$	0.646
Llama-3-8B	$\Delta_i = 24.1$	1,656	1,656	1,656	69,979	$\lceil g(T) \rceil - 1 = 1{,}656$	1.000
Llama-3.1-8B	$\Delta_i = 30.0$	1,656	1,656	1,656	69,980	$\lceil g(T) \rceil - 1 = 1{,}656$	1.000
Qwen2-7B	$\Delta_i = 36.3$	1,656	1,656	1,656	69,980	$\lceil g(T) \rceil - 1 = 1{,}656$	1.000
Qwen2.5-7B	$\Delta_i = 37.7$	1,656	1,656	1,656	69,980	$\lceil g(T) \rceil - 1 = 1{,}656$	1.000
All providers other than i^*		22,641		10,951		$B_{\mathrm{id}}(T) = 141{,}210$	0.160
Ladder roster							
Qwen2-1.5B	$\varepsilon_i = 0.088$	1,379	1,979	16	4,064	$M = 9{,}964$	0.138
Qwen2-1.5B@2	$\varepsilon_i = 0.047$	5,486	7,649	17	12,345	$M = 38{,}751$	0.142
Llama-3.2-1B	$\varepsilon_i = 0.132$	545	716	17	4,314	$M = 4{,}063$	0.134
Qwen2.5-7B	$\varepsilon_i = 0.043$	2,983	3,157	27	69,918	$M = 47{,}908$	0.062
Qwen2-1.5B@4	$\varepsilon_i = 0.017$	2,563	3,051	37	69,716	$M = 357{,}226$	0.007
Llama-3.2-1B@2	$\varepsilon_i = 0.088$	1,281	1,456	113	59,060	$M = 9{,}906$	0.129
Llama-3-8B@2	$\Delta_i = 37.1$	1,334	1,334	1,334	69,936	$L = 4{,}674$	0.285
Qwen2.5-7B@2	$\varepsilon_i = 0.012$	1,334	1,334	1,334	69,936	$M = 705{,}391$	0.002
Llama-3.2-1B@4	$\varepsilon_i = 0.048$	1,334	1,334	1,334	69,936	$M = 38{,}025$	0.035
Llama-3-8B@4	$\Delta_i = 111.4$	1,334	1,334	1,334	69,936	$\lceil g(T) \rceil - 1 = 1{,}334$	1.000
Qwen2.5-7B@4	$\Delta_i = 165.8$	1,334	1,334	1,334	69,936	$\lceil g(T) \rceil - 1 = 1{,}334$	1.000
All providers other than i^*		20,907		6,897		$B_{\mathrm{id}}(T) = 1{,}218{,}576$	0.017

G.5 PAYMENTS AND PROVIDER OUTCOMES

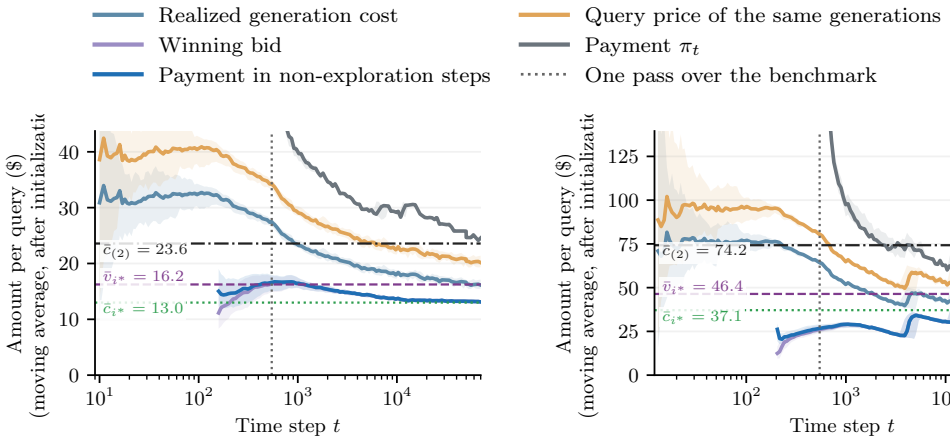

(a) Cost, bid, payment and query price, full roster

(b) Cost, bid, payment and query price, ladder roster

Figure 20: GPQA long horizon: the platform's realized generation cost, the query price of the same generations and the payment π_t, as a moving average with window size K over the steps after initialization, and its winning bid and payment on non-exploration steps in the same window, median over 30 independent runs with a 10–90% band. Monetary quantities are in units of 10^{-6} USD. The dotted, dash-dotted and dashed lines mark $\bar{c}_{i^*}$, $\bar{c}_{(2)}$ and the mean query price $\bar{v}_{i^*}$ of i^*. The vertical axis is cut; π_t enters from above, since early in the run most steps are exploration steps and paid $c_{\max}$.

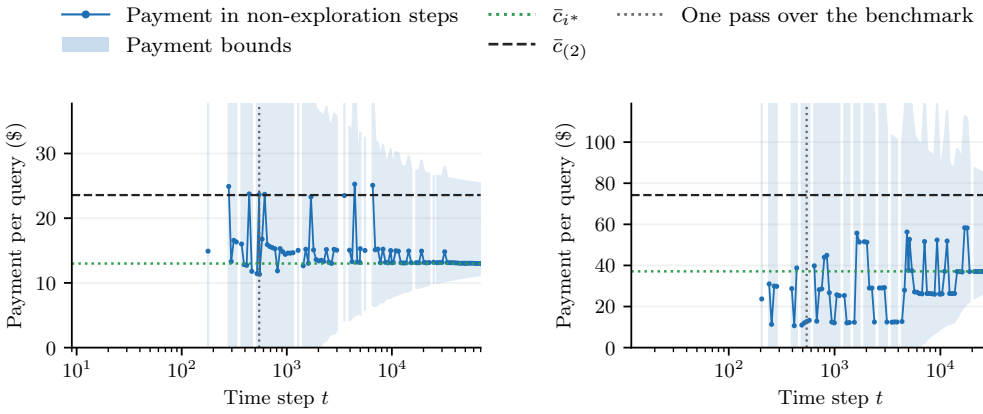

(a) Payment on non-exploration steps and its bounds, full roster

(b) Payment on non-exploration steps and its bounds, ladder roster

Figure 21: GPQA long horizon: the payment π_t of the first run on non-exploration steps, sampled at logarithmically spaced time steps, inside the payment bounds of Appendix E for the provider selected in that step; sampled steps that were exploration steps are left blank. Monetary quantities are in units of 10^{-6} USD. The bounds depend on the selected provider's cost and selection count, and close in as the counts grow; the vertical axis is cut at $1.6\,\bar{c}_{(2)}$, above which the upper bound sits at $c_{\max}$ in the early time steps.

Table 20: GPQA one pass: the platform's economic quantities, mean $\pm$ 95% half-width over 30 independent runs. The payment π_t is over all steps; the winning bid, the payment on non-exploration steps and the second excess-payment row are over non-exploration steps only. A provider's payoff is its total payment minus its total generation cost over the pass.

Quantity	Full roster	Ladder roster
Generation cost per query, whole pass	27.45 ± 0.20	64.42 ± 0.21
Query price per query, whole pass	34.31 ± 0.24	80.53 ± 0.26
Payment π_t per query, whole pass	64.43 ± 0.25	196.53 ± 0.84
Generation cost per query, after initialization	27.37 ± 0.20	64.14 ± 0.21
Query price per query, after initialization	34.21 ± 0.25	80.18 ± 0.26
Payment π_t per query, after initialization	63.80 ± 0.25	194.20 ± 0.86
Winning bid, non-exploration steps	16.35 ± 0.45	26.43 ± 0.59
Payment π_t per query, non-exploration steps	16.59 ± 0.45	26.96 ± 0.59
Exploration-step payment per query, after initialization	55.39 ± 0.47	179.69 ± 1.08
Share of exploration steps after initialization	0.550 ± 0.005	0.612 ± 0.004
Excess payment per step E_T^{pay}/T, whole pass	44.018 ± 0.352	140.199 ± 0.810
Excess payment per query, non-exploration steps	0.182 ± 0.042	0.000 ± 0.000
Mean provider payoff over the pass	2243.8 ± 19.5	5999.7 ± 40.6
Lowest provider payoff over providers and runs	1343.3	2870.6

Table 21: GPQA one pass: each provider's selections and payoff under the platform over 30 independent runs, exploration steps included, with providers sorted by mean generation cost as in Table 14. The last column counts the runs in which the provider's payoff over the pass was negative.

Provider	Selections (mean)	Payoff (mean)	Payoff (min)	Payoff (max)	Negative runs
Full roster					
Qwen2-1.5B (i^*)	116.7	1555.1	1343.3	1986.5	0
Qwen2-0.5B	101.1	1594.5	1435.4	1888.2	0
Qwen2.5-3B	56.5	1883.0	1560.1	2205.5	0
Llama-3.2-3B	51.4	2100.5	1593.1	2870.4	0
Llama-3.2-1B	44.2	3063.6	1996.9	3278.2	0
Llama-3-8B	44.0	2861.2	2719.4	2994.6	0
Llama-3.1-8B	44.0	2609.1	2493.5	2728.7	0
Qwen2-7B	44.0	2247.9	1844.5	2560.2	0
Qwen2.5-7B	44.0	2278.8	2122.5	2431.9	0
Ladder roster					
Qwen2-1.5B	84.5	4935.4	4805.1	5068.8	0
Qwen2-1.5B@2	62.8	4886.1	4441.9	5491.0	0
Llama-3.2-1B	58.8	4942.7	4600.5	5403.9	0
Llama-3-8B (i^*)	50.3	5308.7	4857.2	5847.9	0
Qwen2.5-7B	38.7	6965.4	5994.6	7906.3	0
Qwen2-1.5B@4	39.4	6897.7	4887.2	8577.4	0
Llama-3.2-1B@2	35.3	8251.0	7355.8	8529.1	0
Llama-3-8B@2	35.0	7929.8	7581.1	8159.3	0
Qwen2.5-7B@2	35.0	6918.9	6730.4	7165.6	0
Llama-3.2-1B@4	35.0	6327.6	6072.9	6828.6	0
Llama-3-8B@4	35.0	5278.7	4828.8	5951.7	0
Qwen2.5-7B@4	35.0	3354.2	2870.6	3942.7	0

H GPQA STRONG-COMPETITION EXPERIMENT

This section reports the results for the strong-competition roster on GPQA, in the same form as the two preceding sections; Appendix E provides the experimental details, and the platform runs with the exploration rule described there. The roster is the newest large model and the smallest model of each family, and the threshold is placed so that both large models qualify and both small ones do not. The setting is chosen for its payment ordering: the second-lowest qualified cost $\bar{c}_{(2)} = 50.7$ lies below the mean query price $\bar{v}_{i^*}$ of i^*, 53.8, and above $\bar{c}_{i^*} = 43.0$, so a payment that approaches $\bar{c}_{(2)}$ pays the providers more than the generation cost of i^* and the platform less than the listed price of i^*. Conventions are those of Appendix G; a moving average has window size $K = 546$. The exploration-step cap at $T = 70{,}000$, $\lceil g(T) \rceil - 1 = 3{,}043$, lies below $L_i = 6{,}659$ of the single qualified competitor, so the cap is L_i itself, and $B_{\mathrm{id}}(T) = 76{,}846$; the identification horizon $T_{\mathrm{id}} = 153{,}697$ lies beyond the long run, so no vertical line marks it.

H.1 PROVIDERS AND ENVIRONMENT

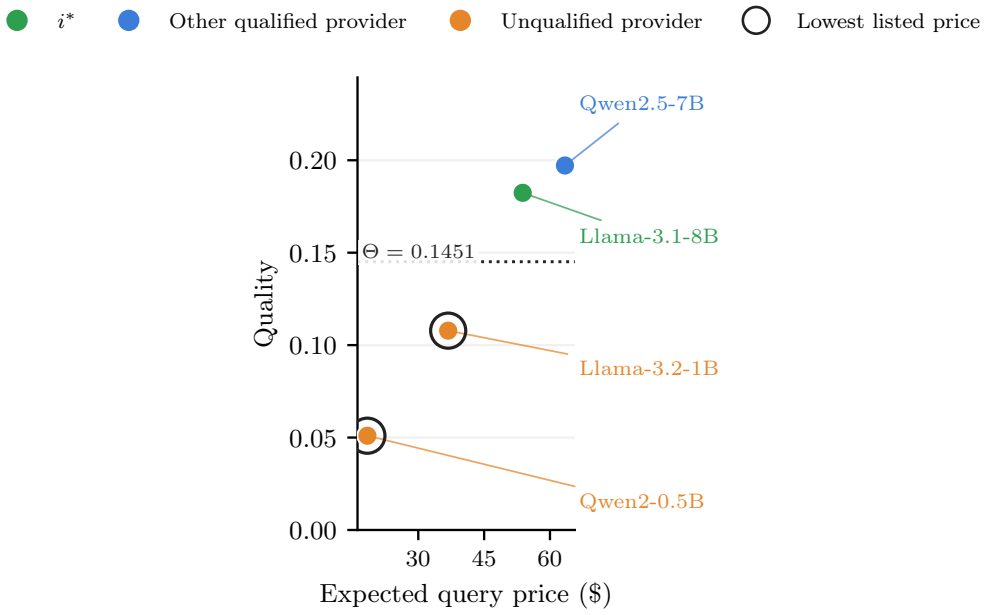

(a) Quality against mean query price

Figure 22: GPQA strong-competition roster. Each point is one provider at its quality q_i and mean query price; the dotted line is the quality threshold $q_{\min}$, and the ring marks the providers with the lowest listed price, both unqualified. Monetary quantities are in units of 10^{-6} USD.

Table 22: GPQA strong-competition setting constants. Costs and prices are mean values per query.

Quantity	Strong-competition roster
Providers N	4
Benchmark size K (queries in one pass)	546
Quality threshold $q_{\min}$	0.1451107
Qualified providers $\lvert\mathcal{Q}\rvert$	2
i^*	Llama-3.1-8B
$\bar{c}_{i^*}$	43.0
$\bar{v}_{i^*}$	53.8
$\bar{c}_{(2)}$	50.7
$\bar{c}_{(2)} - \bar{c}_{i^*}$	7.7
Lowest listed price	Llama-3.2-1B, Qwen2-0.5B
qualified	no
equals i^*	no
$c_{\max}$	75.0
$\sum_{i \notin \mathcal{Q}} M_i$	70,187
$\sum_{i \in \mathcal{Q} \setminus \{i^*\}} L_i$	6,659
$\lceil g(T) \rceil - 1$	3,043
$B_{\mathrm{id}}(T)$	76,846
T_{id}	153,697

Table 23: GPQA strong-competition providers, sorted by mean generation cost; columns as in Table 14.

Provider	Listed price	Tokens	Query price	$\bar{c}_i$	q_i	Qualified	Gap	Bound
Strong-competition roster								
Qwen2-0.5B	0.1	184	18.4	14.7	0.051	no	$\varepsilon_i = 0.094$	$M = 8{,}074$
Llama-3.2-1B	0.1	368	36.8	29.4	0.108	no	$\varepsilon_i = 0.037$	$M = 62{,}113$
Llama-3.1-8B (i^*)	0.1245	432	53.8	43.0	0.182	yes	–	–
Qwen2.5-7B	0.1465	433	63.4	50.7	0.197	yes	$\Delta_i = 7.7$	$L = 6{,}659$

H.2 ONE-PASS BEHAVIOR

Table 24: GPQA strong-competition one-pass outcomes, mean ± 95% half-width over 30 independent runs, paired across rules; columns and the amount-paid convention as in Table 15. Within one pass the quality filter rules out no provider, so the two listed-price rules coincide.

Policy	Unqualified share	R_T^{qual}/T	R_T^{gen}/T	Accuracy	Share of i^*	Generation cost	Amount paid
Strong-competition roster							
Platform	0.707 ± 0.000	0.058 ± 0.000	1.13 ± 0.00	0.100 ± 0.004	0.147 ± 0.000	26.4 ± 0.1	46.0 ± 0.3
Uniform among eligible	0.503 ± 0.008	0.033 ± 0.001	1.93 ± 0.04	0.136 ± 0.005	0.246 ± 0.007	34.4 ± 0.2	42.9 ± 0.3
Lowest listed price among eligible	0.996 ± 0.000	0.066 ± 0.000	0.01 ± 0.00	0.085 ± 0.004	0.002 ± 0.000	22.1 ± 0.2	27.6 ± 0.2
Lowest listed price, no quality filter	0.996 ± 0.000	0.066 ± 0.000	0.01 ± 0.00	0.085 ± 0.004	0.002 ± 0.000	22.1 ± 0.2	27.6 ± 0.2
No quality filter (ablation)	0.930 ± 0.002	0.081 ± 0.000	0.22 ± 0.01	0.071 ± 0.004	0.042 ± 0.002	18.6 ± 0.2	19.2 ± 0.2

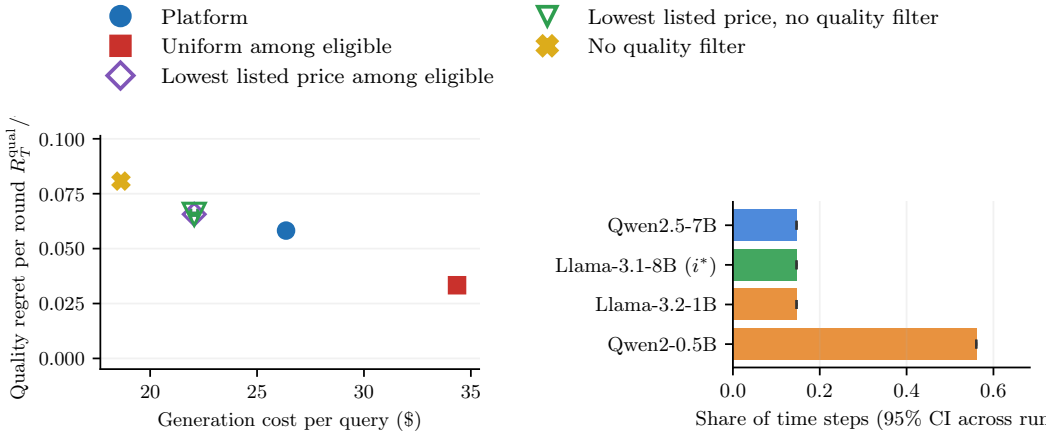

(a) Quality regret against generation cost

(b) Platform selection shares

Figure 23: GPQA strong-competition one pass: quality regret per step against generation cost per query for every routing rule (left), and the share of steps the platform assigned to each provider, exploration steps included (right), means with 95% confidence intervals across runs. Monetary quantities are in units of 10^{-6} USD. In the right panel green marks i^*, blue the other qualified provider and orange the unqualified providers.

H.3 LONG-HORIZON LEARNING AND IDENTIFICATION

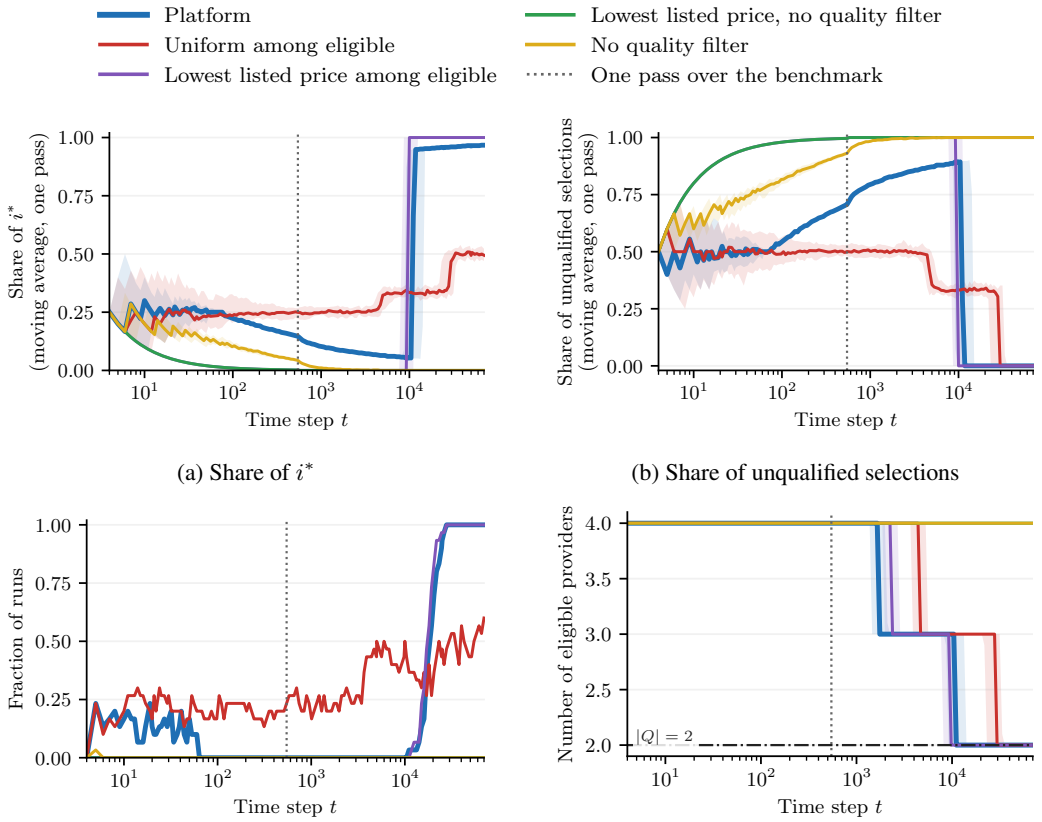

(a) Share of i^*

(b) Share of unqualified selections

(c) i^* is the unique most selected provider

(d) Number of eligible providers

Figure 24: GPQA strong-competition long horizon: identification of i^* under every routing rule. Shares are moving averages with window size K, exploration steps included, median over 30 independent runs with a 10–90% band; the leader panel is the fraction of runs in which i^* has strictly more selections than every other provider; the last panel is the number of providers not yet ruled out, with the dash-dotted line at $|\mathcal{Q}|$. The vertical line marks one pass over the benchmark. The two listed-price rules coincide until the filter rules out the two cheap providers.

Table 25: GPQA strong-competition long horizon: outcomes at $T = 70{,}000$, means over 30 independent runs; columns as in Table 16.

Policy	Share of i^*	Moving avg.	Unqualified share	R_t^{qual}/t	R_t^{gen}/t	Generation cost	Amount paid	i^* most selected	$\lvert\hat{\mathcal{Q}}_t\rvert$
Strong-competition roster									
Platform	0.821	0.967	0.135	0.0060	0.335	41.3	47.0	1.00	2.0
Uniform among eligible	0.428	0.498	0.144	0.0063	3.291	44.1	55.1	0.60	2.0
Lowest listed price among eligible	0.865	1.000	0.135	0.0060	0.000	41.0	51.2	1.00	2.0
Lowest listed price, no quality filter	0.000	0.000	1.000	0.0657	0.000	22.1	27.6	0.00	4.0
No quality filter (ablation)	0.001	0.000	0.999	0.0939	0.003	14.8	14.8	0.00	4.0

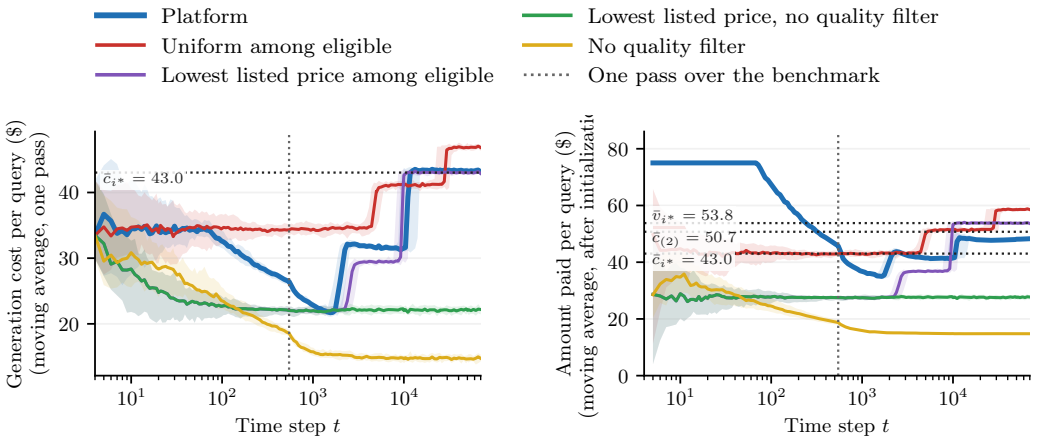

(a) Generation cost per query

(b) Amount paid per query

Figure 25: GPQA strong-competition long horizon: generation cost per query as a moving average with window size K (left) and amount paid per query as a moving average with window size K over the steps after initialization (right), under every rule, median over 30 independent runs with a 10–90% band. Monetary quantities are in units of 10^{-6} USD. The reference lines mark $\bar{c}_{i^*}$, $\bar{c}_{(2)}$ and the mean query price $\bar{v}_{i^*}$ of i^*. The platform's payment π_t includes its exploration steps at $c_{\max}$; the vertical axis of the right panel is cut.

Table 26: GPQA strong-competition exploration steps of the platform, means over 30 independent runs, over one pass and over the long run; rows as in Table 17.

Quantity	Strong-competition roster one pass	long run
Exploration steps (mean)	253	4,005
Share of exploration steps after initialization	0.467	0.057
Share of exploration steps at T (moving avg.)	–	0.033
Exploration-step payment per step	34.76	2.47
Exploration-step generation cost per step	18.30	1.70
Share of exploration steps given to i^*	0.312	0.188
Share of exploration steps given to unqualified providers	0.375	0.051
Payment π_t at T, non-exploration steps (moving avg.)	–	47.37
Payment π_t at T, all steps (moving avg.)	–	48.28

H.4 Regret and theoretical bounds

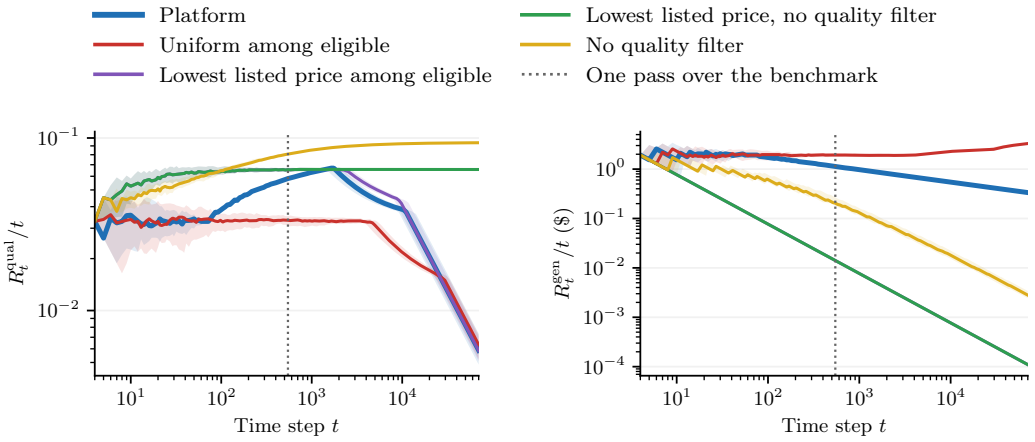

(a) Average quality regret

(b) Average generation regret

Figure 26: GPQA strong-competition long horizon: average quality regret R_t^{qual}/t and average generation regret R_t^{gen}/t of every routing rule, median over 30 independent runs with a 10–90% band, on logarithmic axes. Monetary quantities are in units of 10^{-6} USD. The platform's regrets include its exploration steps.

Table 27: GPQA strong-competition theoretical diagnostics of the platform over the long run ($T = 70{,}000$, 30 independent runs); rows as in Table 18.

Quantity	Strong-competition roster
$B_{\text{id}}(T)$	76,846
T_{id}	153,697
Horizon T of the long run	70,000
Selections of providers other than i^* after initialization (mean)	12,508
ratio to $B_{\text{id}}(T)$	0.163
Runs in which i^* is the most selected provider at $t = 5{,}000$	0.00
at $t = 10{,}000$	0.00
at $t = T_{\text{id}}$	–
at $t = T$	1.00
R_T^{qual}, realized (mean)	420
bound equation 31	3,076
ratio	0.136
R_T^{gen}, realized (mean)	23,433
bound equation 34	51,268
ratio	0.457
E_T^{pay} without the exploration steps, realized (mean)	111
bound equation 37 without its exploration term	300,158
ratio	0.000
E_T^{pay} incl. exploration steps, realized (mean)	97,343
bound equation 37	595,681
ratio	0.163
Good event held, one pass (runs)	30 of 30
Good event held, long run, platform	30 of 30
Good event held, long run, no quality filter	30 of 30
Payment-bound violations, one pass (steps)	0

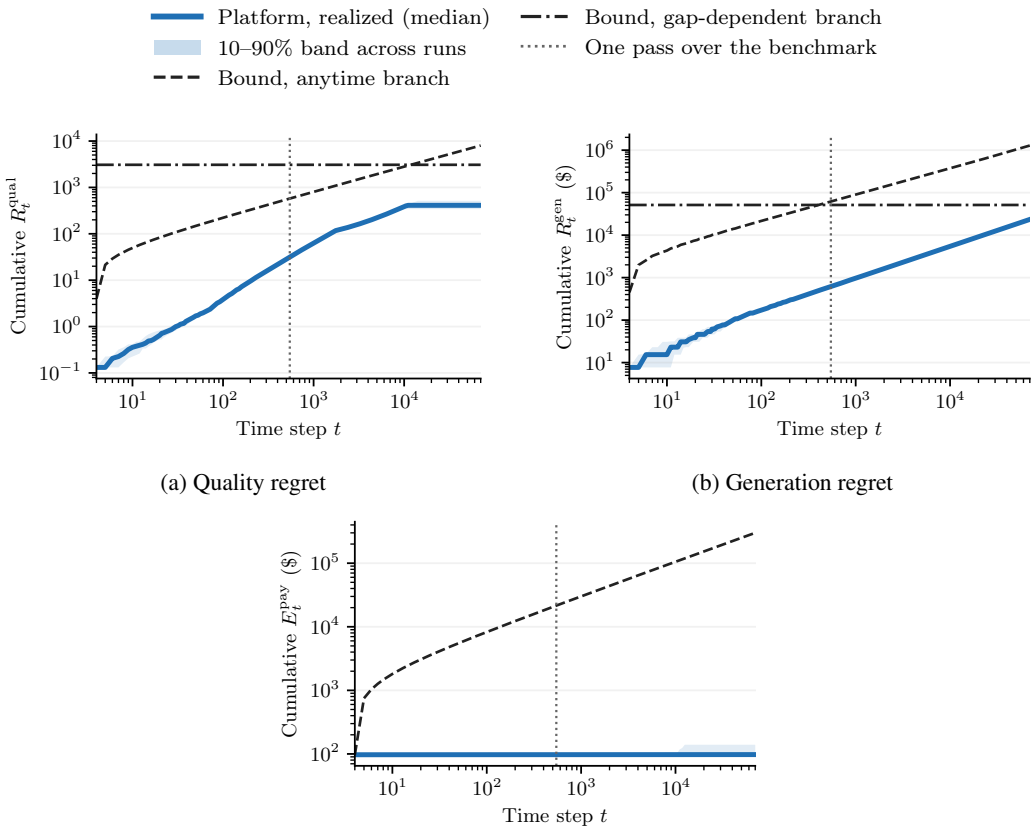

(a) Quality regret

(b) Generation regret

(c) Excess payment

Figure 27: GPQA strong-competition long horizon: the platform's cumulative quality regret, generation regret and excess payment against their bounds, evaluated at every time step. Monetary quantities are in units of 10^{-6} USD. The bounds are those of Theorems 7 and 8 and of equation 37 without its exploration term, as in Figure 19. The regrets include the exploration steps; the excess payment excludes them.

Table 28: GPQA strong-competition long horizon: selections of every provider other than i^* by the platform after initialization, over 30 independent runs of $T = 70{,}000$ queries, against its cap; columns as in Table 19.

Provider	Gap	Selections mean	Selections max	Exploration (mean)	Last selection (mean step)	Bound	Ratio
Strong-competition roster							
Qwen2-0.5B	$\varepsilon_i = 0.094$	1,172	1,585	16	1,746	$M = 8{,}074$	0.145
Llama-3.2-1B	$\varepsilon_i = 0.037$	8,293	12,542	190	10,980	$M = 62{,}113$	0.134
Qwen2.5-7B	$\Delta_i = 7.7$	3,043	3,043	3,043	69,999	$L = 6{,}659$	0.457
All providers other than i^*		12,508		3,249		$B_{\text{id}}(T) = 76{,}846$	0.163

H.5 PAYMENTS AND PROVIDER OUTCOMES

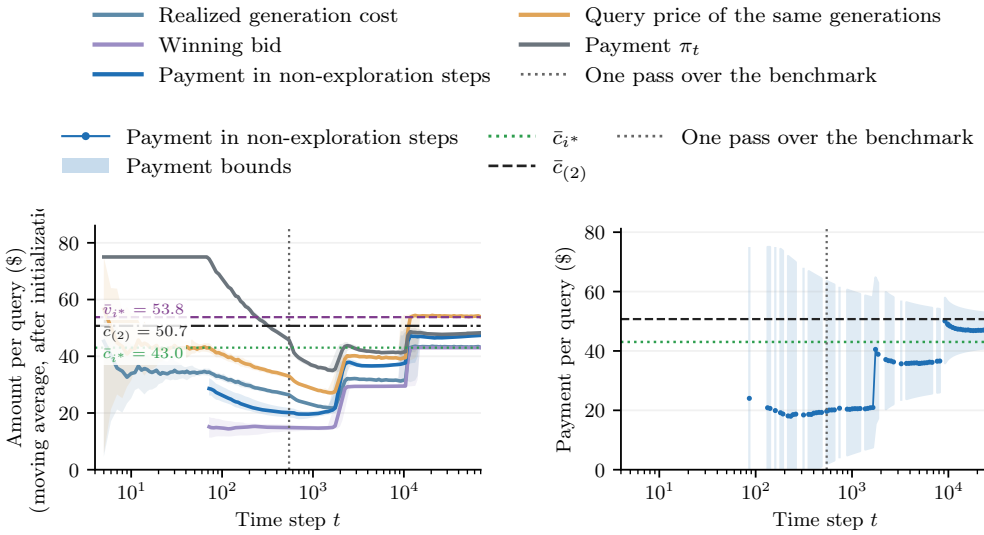

(a) Cost, bid, payment and query price

(b) Payment on non-exploration steps and its bounds

Figure 28: GPQA strong-competition long horizon. Left: the platform's realized generation cost, the query price of the same generations and the payment π_t, as a moving average with window size K over the steps after initialization, and its winning bid and payment on non-exploration steps in the same window, median over 30 independent runs with a 10–90% band; the dotted, dash-dotted and dashed lines mark $\bar{c}_{i^*}$, $\bar{c}_{(2)}$ and the mean query price $\bar{v}_{i^*}$ of i^*, and the vertical axis is cut. Right: the payment π_t of the first run on non-exploration steps, sampled at logarithmically spaced time steps, inside the payment bounds of Appendix E for the provider selected in that step; sampled steps that were exploration steps are left blank. Monetary quantities are in units of 10^{-6} USD.

Table 29: GPQA strong-competition one pass: the platform's economic quantities, mean $\pm$ 95% half-width over 30 independent runs; rows as in Table 20.

Quantity	Strong-competition roster
Generation cost per query, whole pass	26.36 ± 0.13
Query price per query, whole pass	32.94 ± 0.16
Payment π_t per query, whole pass	45.96 ± 0.27
Generation cost per query, after initialization	26.29 ± 0.13
Query price per query, after initialization	32.87 ± 0.16
Payment π_t per query, after initialization	45.75 ± 0.27
Winning bid, non-exploration steps	14.66 ± 0.32
Payment π_t per query, non-exploration steps	20.13 ± 0.51
Exploration-step payment per query, after initialization	34.76 ± 0.00
Share of exploration steps after initialization	0.467 ± 0.000
Excess payment per step E_T^{pay}/T, whole pass	11.428 ± 0.000
Excess payment per query, non-exploration steps	0.000 ± 0.000
Mean provider payoff over the pass	2675.9 ± 35.2
Lowest provider payoff over providers and runs	1726.3

Table 30: GPQA strong-competition one pass: each provider's selections and payoff under the platform over 30 independent runs, exploration steps included, with providers sorted by mean generation cost. The last column counts the runs in which the provider's payoff over the pass was negative.

Provider	Selections (mean)	Payoff (mean)	Payoff (min)	Payoff (max)	Negative runs
Strong-competition roster					
Qwen2-0.5B	306.0	2581.9	1726.3	3657.6	0
Llama-3.2-1B	80.0	3637.5	3389.7	3815.3	0
Llama-3.1-8B (i^*)	80.0	2564.5	2386.3	2730.8	0
Qwen2.5-7B	80.0	1919.8	1749.9	2130.2	0