BETTER MODELS CAN REDUCE ACCURACY UNDER STRATEGIC PRICING

Anonymous authors
Paper under double-blind review

ABSTRACT

Recent progress in language models is assessed against fixed computation budgets, but provider prices may change when they compete for users. We show that a model can become more accurate at every computation level and can generate higher equilibrium revenue, yet users on a fixed budget may incur lower accuracy after providers change prices. We model the market of combining verifiable outputs from models priced independently by providers, and show that every pure-strategy equilibrium is characterized by the maximizers of $xf'(x)$, where f is the negative log failure probability. This characterization allows us to identify an explicit family of improvements that reduces equilibrium accuracy on an open parameter region, even when each task is eventually solvable, and holds under sufficiently small strictly positive serving costs. We perform experiments with 610,560 Qwen and Llama generations on LiveCodeBench and MATH-500, comparing fixed vs. equilibrium prices. Non-thinking replacements have gains of 3.8 and 2.2 percentage points from checkpoints at fixed prices, but become losses of 1.7 and 1.6 points after repricing. Thinking replacements have gains that remain positive, but are reduced by 1.4 and 0.5 points respectively. Moreover, we bound profitable price deviations and discuss the implications of these results for the shape of test-time scaling curves.

1 INTRODUCTION

More inference computation can lead to better performance of language models via repeated sampling, verification, and search (Wang et al., 2023; Yao et al., 2023; Lightman et al., 2024; Snell et al., 2025). Users can also aggregate calls to models (Chen et al., 2021; Brown et al., 2024; Olmedo et al., 2026). At fixed prices, an upgrade that improves every available allocation cannot reduce the accuracy attainable at a fixed spending budget.

However, the upgrade also modifies incentives for price-setting by providers. Providers maximize revenue, while users maximize accuracy subject to a fixed budget. A provider can change its price even if its model is unchanged, and therefore an upgrade might change what users can afford to compute. Routing and cascading strategies optimize model usage well at fixed prices (Chen et al., 2023; Ong et al., 2025; Dekoninck et al., 2025). Nevertheless, the performance of models at fixed prices is not necessarily predictive of performance after re-pricing by providers. We explore the relationship between model quality, inference prices, and purchased computation for problems with verifiable solutions.

Our main result is a reversal: Improving a model's accuracy-compute curve can reduce the user's equilibrium success probability. We establish our result for two providers with independent task-dependent success probabilities, and a perfect verifier, for models in a family that can solve any task with positive probability. Our new model is more accurate at every positive computation level and captures a higher share of budget but has a higher equilibrium price, reducing the quantity of computation purchased from it enough to lower the overall accuracy. The reversal takes place on an open parameter interval and persists for sufficiently small positive serving costs.

$f_j(x)$ denotes provider j's negative log failure probability and $H_j(x) = xf_j'(x)$. Equilibrium quantities correspond precisely to global maximizers of H_j. Finite equilibria exist if and only if every H_j attains a finite supremum at a finite positive quantity. Peak heights determine revenue shares. In the latent-rate family, improving the failure curve's tail constant can reduce equilibrium accuracy,

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.

and increasing its tail exponent raises equilibrium accuracy. For pointwise improvements with unique equilibrium quantities, equilibrium accuracy cannot decrease when the ratio of new to old marginal capability is nondecreasing with computation. Near-maximizers may be used to find approximate equilibria with explicit bounds on the revenue attainable via any unilateral price deviation.

Our paired experiments replace Qwen checkpoints and keep a Llama competitor fixed. On LiveCodeBench and MATH-500, we compare the upgrade at the old prices with the outcome after providers reprice. We retain cross-model dependence with joint success estimates aligned by task in purchasing decisions. Repricing reduces all four gains measured at fixed prices and reverses both non-thinking comparisons. We vary serving costs, call limits, the value of unspent funds, and stopping after verified success. Continuous-price best responses quantify incentives to deviate, with additional bounds that account for uncertainty in the buyer's allocation.

We make three contributions. First, we characterize all pure strategy equilibria of the inference-pricing game, and derive a constructive approximation bound. Second, we show that pointwise model improvement can decrease accuracy at equilibrium, even if neither model has permanently unsolvable tasks, and that this effect holds under sufficiently small positive serving costs. Third, we investigate checkpoint replacement under endogenous prices using task-aligned generations, continuous-price best responses, and bounds on profitable deviations under sampling uncertainty.

2 Problem Formulation

Failure Curves. For each model, fix a generation procedure and take its repetitions as units of computation. Decisions about output length, prompting, or search can affect the failure curve itself. On a randomly drawn task, model j has latent success rate $\Lambda_j \geq 0$. Conditional on this rate, computation x has no verifiably correct candidate with probability $e^{-\Lambda_j x}$. The population failure curve and its negative logarithm are

$$F_j(x) = \mathbb{E}[e^{-\Lambda_j x}], \qquad f_j(x) = -\log F_j(x). \tag{1}$$

Independent candidate generation has this form: if each candidate succeeds with probability q, all k candidates fail with probability $(1-q)^k$. Taking $\Lambda = -\log(1-q)$ gives the continuous extension.

A failure also provides information about the task. Write $\nu_{j,x}$ for the distribution of Λ_j conditional on continued failure, with $d\nu_{j,x}(\lambda) = e^{-\lambda x} d\nu_j(\lambda)/F_j(x)$. Exponential-family moment identities (Wainwright and Jordan, 2008) then give

$$f_j'(x) = \mathbb{E}_{\nu_{j,x}}[\Lambda_j], \qquad f_j''(x) = -\operatorname{Var}_{\nu_{j,x}}(\Lambda_j). \tag{2}$$

Marginal capability is the mean latent success rate among the tasks still unsolved. Continued failure gives greater relative weight to lower-rate tasks. Variation in task difficulty therefore accounts for diminishing marginal capability in (2).

Buyer and Providers. The buyer purchases computation from $m \geq 2$ providers. Provider j quotes a linear price $p_j > 0$ per unit. For independent latent success rates and generation randomness across providers, a perfect verifier gives combined success probability

$$U(x) = 1 - \prod_j F_j(x_j) = 1 - \exp\left(-\sum_j f_j(x_j)\right). \tag{3}$$

Maximizing success with an inference budget $B > 0$ is thus equivalent to

$$x(p) = \arg\max_{x \geq 0} \left\{ \sum_j f_j(x_j) : \sum_j p_j x_j \leq B \right\}. \tag{4}$$

The theoretical analysis covers capability functions with $f_j(0) = 0$, continuity on $[0, \infty)$, twice continuous differentiability on $(0, \infty)$, $f_j' > 0$, $f_j'' < 0$, and $f_j'(x) \to 0$ as $x \to \infty$; they need not be bounded. Strict concavity gives a unique allocation characterized by constrained optimality conditions (Boyd and Vandenberghe, 2004). Strict monotonicity makes the budget constraint bind.

Providers initially maximize revenue $R_j(p) = p_j x_j(p)$; Section 3.5 introduces serving costs, with profits $\pi_j = (p_j - c_j)x_j$.

Each price must be evaluated at the allocation it induces. A provider makes more per unit when it increases its price, but the buyer can reallocate computation. A price vector p is an equilibrium if $R_j(p) \geq R_j(q, p_{-j})$ for every provider j and every $q > 0$, with the buyer reoptimizing after each deviation. The benchmark considers the budget as funds allocated to inference procurement: keeping it has no direct utility. We also explore integer calls, a money-valued outside option, and stopping after verified success.

Cross-Model Dependence. Even when drawn independently, candidates may fail on the same challenging tasks (Kim et al., 2025). Under correlated latent rates, the negative log failure function $f_{\text{joint}}(x) = -\log \mathbb{E}[e^{-\sum_j \Lambda_j x_j}]$ is concave but in general nonseparable. The Hessian is the negative tilted covariance matrix, $-\operatorname{Cov}_x(\Lambda)$. Our equilibrium characterization involves the separable buyer problem (4). The experiments use joint success surfaces aligned with the tasks to preserve the cross-model dependencies observed.

3 Equilibrium Analysis

3.1 Equilibrium Characterization

Define provider j's log-compute marginal capability, its peak value, and its maximizers by

$$H_j(x) = x f_j'(x), \qquad M_j = \sup_{x>0} H_j(x), \qquad Z_j = \arg\max_{x>0} H_j(x). \tag{5}$$

Since $H_j = df_j/d\log x$, its peak is the greatest marginal capability gain from a proportional increase in computation.

Theorem 1 (Equilibrium characterization)**.** *A positive price pure strategy equilibrium exists under the assumptions of Section 2 if and only if each H_j attains a finite supremum at a finite positive quantity. All equilibria are obtained by choosing $z_j \in Z_j$ and defining*

$$x_j^* = z_j, \qquad p_j^* = \frac{B f_j'(z_j)}{\sum_h M_h}, \qquad R_j^* = B\frac{M_j}{\sum_h M_h}. \tag{6}$$

The equilibrium is unique if each H_j has a unique maximizer.

Proof. Let $\mu^* = \sum_h M_h/B$. At the proposed prices, $f_j'(z_j) = \mu^* p_j^*$ and $\sum_j p_j^* z_j = B$, so the buyer's optimality conditions are satisfied. A profitable unilateral price change would leave less of the budget for competing providers. At their fixed prices, demands are nonincreasing in the common budget multiplier; the new multiplier must therefore satisfy $\mu' > \mu^*$. The deviating provider would then have

$$H_j(x_j') = \mu' R_j' > \mu^* R_j^* = M_j,$$

contradicting the definition of the global maximum.

Conversely, every provider must have positive equilibrium demand: one with no demand could enter at a sufficiently small positive price. Fix competitors' prices and let $D_{-j}(\mu)$ be their total expenditure at multiplier μ. On $D_{-j}(\mu) \leq B$, the function $G(\mu) = \mu[B - D_{-j}(\mu)]$ is continuous and strictly increasing from zero to infinity. To implement any quantity $z > 0$, a firm solves $G(\mu) = H_j(z)$ and sets $q = f_j'(z)/\mu$. A larger value of H_j raises the multiplier and strictly reduces positive expenditure on competitors. Thus every equilibrium quantity maximizes H_j. Summing the buyer's first-order conditions yields (6). Appendix A proves the demand and boundary claims. □

Theorem 1 covers price deviations of any size, including those that reorder the buyer's allocation. At a smooth interior maximum,

$$-\frac{x f_j''(x)}{f_j'(x)} = 1. \tag{7}$$

This condition identifies a candidate peak whose global optimality must still be checked. Bounded capability is sufficient for existence, because concavity then guarantees that $H_j(x) \to 0$ at both endpoints. Unbounded capability can also have a finite peak, as the next construction shows.

3.2 Solvable Tasks and Finite Equilibria

We consider the family of failure functions

$$F_{\alpha,\kappa,\tau}(x) = \left(1 + \frac{x}{\tau}\right)^{-\kappa} \exp\left(-\alpha \frac{x}{\tau + x}\right), \qquad \alpha, \kappa, \tau > 0. \tag{8}$$

This family is induced by a latent rate distribution. That is, let $G \sim \mathrm{Gamma}(\kappa, 1)$, $N \sim \mathrm{Poisson}(\alpha)$, and $E_r \sim \mathrm{Exp}(1)$ be independent random variables. Let $\Lambda = (G + \sum_{r=1}^{N} E_r)/\tau$. The Laplace transforms of these distributions result in (8) (Schilling et al., 2012). Since $G > 0$ a.s., all tasks have positive rates and $F(x) \to 0$. Furthermore, when $\kappa > 1$, $\mathbb{E}[1/\Lambda] \le \tau/(\kappa - 1)$, so the proposed distribution can have finite expected solution time.

The associated capability is $f(x) = \alpha x/(\tau + x) + \kappa \log(1 + x/\tau)$. For $t = x/\tau$,

$$H(\tau t) = \frac{\alpha t}{(1+t)^2} + \frac{\kappa t}{1+t}, \qquad \frac{dH(\tau t)}{dt} = \frac{\alpha + \kappa + (\kappa - \alpha)t}{(1+t)^3}. \tag{9}$$

When $\alpha > \kappa$, the unique maximum and the capability purchased at equilibrium are

$$z = \tau \frac{\alpha + \kappa}{\alpha - \kappa}, \qquad M = \frac{(\alpha + \kappa)^2}{4\alpha}, \tag{10}$$

$$V^* := f(z) = \frac{\alpha + \kappa}{2} + \kappa \log \frac{2\alpha}{\alpha - \kappa}. \tag{11}$$

When $\alpha \le \kappa$, H increases toward κ without attaining it. This market has a finite equilibrium if and only if every $\alpha_j > \kappa_j$. In that region, its success probability is $1 - e^{-\sum_j V_j^*} < 1$, though the technical limit is one. Scaling the market's spending budget scales prices and revenues in (6) without altering purchased quantities.

3.3 Model Improvement and Equilibrium Accuracy

Theorem 2 (Price-mediated reversal)**.** *Assume that all providers follow* (8) *and* $\alpha_j > \kappa_j$*. Increase one provider's parameter from* α_j^{old} *to* α_j^{new} *(*$\kappa_j < \alpha_j^{\mathrm{old}} < \alpha_j^{\mathrm{new}} < 2\kappa_j$*; other parameters fixed). The provider's success probability increases for any fixed* $x > 0$*. At the new equilibrium, its price and revenue both increase. Its purchased computation amount decreases, and the buyer's overall success probability decreases.*

Proof. Pointwise improvement follows from $\partial_\alpha f(x) = x/(\tau + x) > 0$. By (10)–(11),

$$\frac{\partial z}{\partial \alpha} = -\frac{2\kappa\tau}{(\alpha - \kappa)^2} < 0, \quad \frac{\partial V^*}{\partial \alpha} = \frac{1}{2} - \frac{\kappa^2}{\alpha(\alpha - \kappa)} < 0, \quad \frac{\partial M}{\partial \alpha} = \frac{1}{4}\left(1 - \frac{\kappa^2}{\alpha^2}\right) > 0. \tag{12}$$

The middle inequality holds in $(\kappa, 2\kappa)$. None of the competitors' equilibrium quantities or capabilities vary. The sum of the capabilities falls, and the improved provider's budget share rises. Thus its price R_j^*/z_j increases, as its numerator rises and denominator falls. Integrating the derivatives over the stated interval establishes the finite comparisons. □

Increasing α adds an independent nonnegative compound-Poisson term to each task's rate. Coupling the two processes in this way ensures that no task becomes harder. At unchanged prices, the previous allocation is still affordable, and improves strictly. In Figure 1, $\kappa = 0.2$, $\tau = 1$, and increasing α from 0.24 to 0.36 leads to a drop in equilibrium quantity from 11 to 3.5. In the right panel, $\kappa = \tau = 1$, while α varies. Repricing moves the purchase to a lower point on the improved curve.

3.4 Scaling-Curve Shape and Equilibrium Accuracy

The equilibrium effect depends on how the scaling curve improves. The horizontal acceleration $\widetilde{f}(x) = f(sx)$ changes the equilibrium to $\widetilde{z} = z/s$ and $\widetilde{p} = sp$ with no change in capability or revenue. The reparameterization $y = sx$, $q = p/s$ keeps expenditure the same, $px = qy$. By contrast, increasing the tail exponent κ with $\kappa < \alpha$ gives

$$\frac{\partial V^*}{\partial \kappa} = \frac{1}{2} + \log \frac{2\alpha}{\alpha - \kappa} + \frac{\kappa}{\alpha - \kappa} > 0. \tag{13}$$

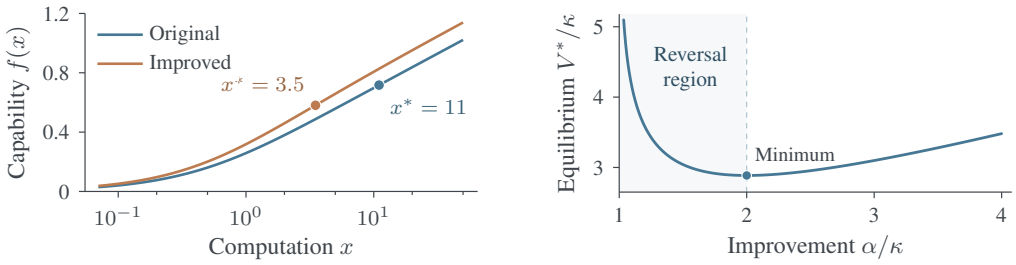

Figure 1: **Pointwise improvement can reduce purchased accuracy.** Left: Dots represent equilibrium purchases at the old and new curves. Right: Equilibrium capability decreases for $1 < \alpha/\kappa < 2$.

The original failure curve has tail $F(x) \sim e^{-\alpha}(x/\tau)^{-\kappa}$. Both of these changes improve accuracy for fixed computation, but improving the tail constant can decrease equilibrium accuracy whereas increasing the tail exponent increases it.

We establish a sufficient condition on relative marginal gains that rules out reversal. Consider the zero-cost independent provider model of Section 2 with functions f and g satisfying the buyer assumptions and unique finite positive maximizers z_f and z_g of $xf'(x)$ and $xg'(x)$. If $g(x) \geq f(x)$ for all $x \geq 0$ and $g'(x)/f'(x)$ is nondecreasing on $(0, \infty)$, then $z_g \geq z_f$ and $g(z_g) \geq f(z_f)$. Thus, when relative marginal gains do not decline with computation, the maximizing quantity cannot move to the left. Replacing a provider's capability function under these conditions cannot reduce equilibrium success. The proof is in Proposition 5.

3.5 Serving Costs

Allowing for a positive cost of serving computations preserves the accuracy reversal near a strict zero-cost equilibrium. The following conditions describe local price adjustments and global deviations.

Theorem 3 (Cost robustness)**.** *Let f_j be C^3 near its equilibrium quantity and assume all H_j have unique peaks z_j with $H_j''(z_j) < 0$, and for each j, its peak exceeds both endpoint limit superiors. For sufficiently small nonnegative cost vectors, there exists a locally unique pure-strategy equilibrium near the zero-cost equilibrium, which varies continuously with costs. Strict comparisons of buyer success and provider profit between two such markets continue for sufficiently small costs.*

Fix competitors' prices and denote the revenue obtained by choosing quantity x as $r_j(H_j(x))$ with $r_j' > 0$ near equilibrium. The first-order condition for profit is $r_j'(H_j(x_j))H_j'(x_j) - c_j = 0$. With zero cost, the Jacobian of these simultaneous conditions is diagonal with negative entries $r_j'(M_j)H_j''(z_j)$. The implicit function theorem implies the existence of a nearby solution. Strict peak and endpoint gaps rule out profitable distant deviations, which proves equilibrium; Appendix C contains the full argument. The family in (8) meets these conditions when $\alpha > \kappa$.

4 Equilibrium Computation

Quantities near the maxima of H_j lead to approximate equilibria with explicit revenue bounds.

Proposition 4 (Approximate equilibrium)**.** *Assume $M_j = \sup_x H_j(x) < \infty$. Given $z_j > 0$, let $S = \sum_j H_j(z_j)$ and $p_j = Bf_j'(z_j)/S$. The buyer obtains z and any unilateral price deviation satisfies*

$$\frac{R_j(q, p_{-j})}{R_j(p)} \leq \frac{M_j}{H_j(z_j)}. \tag{14}$$

The profitable deviation increases the multiplier strictly above S/B and thus the deviating provider earns at most BM_j/S. Dividing by its current revenue we obtain the bound. Taking $H_j(z_j) \geq M_j/(1+\varepsilon)$ yields a $(1+\varepsilon)$-approximate equilibrium even if the finite supremum is not attained. Any valid upper bound on M_j leads to a valid upper bound on revenue.

Revenue guarantees and quantity predictions must be checked separately. Near a strict quadratic maximum, a uniform perturbation of H_j by δ moves the maximizer by at most $O(\sqrt{\delta})$ and changes the achieved value by at most 2δ (Bonnans and Shapiro, 2000). Appendix D.6 extends these bounds on profitable price deviations to numerical optimization error. Nearly flat peaks can lead to similar revenues at vastly different quantities. An approximate-equilibrium guarantee alone does not imply accurate quantity predictions; we also check quantity predictions and curve fit.

Empirical Demand. In the continuous model, demand is separable. In empirical price comparisons, the buyer optimizes a *task-aligned joint success surface* over a finite allocation space $\mathcal{K} \subset \mathbb{Z}_+^m$ that includes zero. At prices p, it maximizes $\widehat{U}(k)$ subject to $p \cdot k \leq B$, breaking ties first by fewer total calls, then lexicographically. The incentive to deviate is measured by the budget-normalized largest profitable deviation (Bosshard et al., 2020):

$$\mathrm{Gap}(p) = \frac{1}{B} \max_j \left\{ \sup_{q>0} (q - c_j)\widehat{k}_j(q, p_{-j}) - (p_j - c_j)\widehat{k}_j(p) \right\}. \tag{15}$$

The supremum over continuous prices is computed with a finite scan. With competitors' prices fixed, an allocation with $k_j > 0$ is affordable up to

$$q_k = \frac{B - \sum_{h \neq j} p_h k_h}{k_j}. \tag{16}$$

Between consecutive positive thresholds, demand is constant and profit increases linearly with price whenever sales are positive. The allocation remains affordable at the right endpoint. So a best response occurs at a threshold or at a price above B with no sales. The ranked event sweep in Algorithm 2 computes each provider's best response in $O(|\mathcal{K}| \log |\mathcal{K}|)$ time.

We initialize prices from finite near-maximizers of fitted H_j curves and perform damped, multi-start best-response search on joint empirical demand (Algorithm 1 in Appendix F.3). A continuous-price deviation test is run on every visited profile. The selected profile is the one minimizing $\mathrm{Gap}(p)$, breaking ties first by success in calibration and then by lexicographic price order. Numerical acceptance is determined by this test, not iteration convergence; equilibrium comparison requires accepted old and new profiles.

The empirical gap $\mathrm{Gap}(p)$ measures deviation incentives in the estimated game. Uncertainty in estimated success can also affect the buyer's allocation. We bound profitable deviations for every buyer choice consistent with simultaneous success intervals (Appendix D; see Balcan et al. (2019) for sample-based incentive estimation). When allocations are close to a tie, even a small success estimation error can affect demand. Controlling success error alone is therefore insufficient to control deviation incentives.

Finite observations may not be sufficient to identify the tail behavior that determines the existence of equilibrium. Appendix E provides a two-point testing argument (Tsybakov, 2009): consider two markets, one with a finite equilibrium and one without, whose observation laws have a total variation distance of at most $N\eta T$ for N tasks and horizon T. The difference between the two is a constant increase of η in one provider's latent success rate.

5 Experiments

5.1 Experimental Setup

We compare the two Qwen replacements against the same Llama-3.1-8B-Instruct competitor (Yang et al., 2025; Grattafiori et al., 2024): Qwen3-4B without thinking to Qwen3-4B-Instruct-2507 (N), and Qwen3-4B with thinking to Qwen3-4B-Thinking-2507 (T). The generation limits and sampling parameters are matched for each pair. The replacement direction is determined by the checkpoint chronology.

The experiments use 454 executable problems from the LiveCodeBench release_v6, dated August 1, 2024 to April 30, 2025 (Jain et al., 2025), and 500 MATH-500 problems (Hendrycks et al., 2021). The problems in each domain are divided equally into calibration and test sets, stratified by task difficulty and mathematical subject. For the five service configurations, 128 independent candidates per task

and configuration yield 610,560 generations. Temperature is 0.7, top-p is 0.95, and token limits are 8,192 without thinking and 16,384 with thinking. Executable tests check code correctness, and equivalence to the reference answer checks mathematical correctness. Models, splits, and generation settings are fixed before test evaluation. Appendix F describes prompting, validation, resource accounting, and selection under public tests.

For task i and model j, let c_{ij} be the number of correct outputs among the $n = 128$ candidates. We generalize the standard pass@k estimator (Chen et al., 2021) to estimate joint success as

$$\widehat{U}(k) = 1 - \frac{1}{N}\sum_{i=1}^{N}\prod_{j}\frac{\binom{n-c_{ij}}{k_j}}{\binom{n}{k_j}}. \tag{17}$$

We multiply failure probabilities across models within each task before averaging over tasks. This order retains the dependence induced by shared task difficulty. Our main allocation set is $\{0,\dots,64\}^2$, with 32 and 128 calls as alternatives. Smooth fitted curves initialize the price search; all reported purchase decisions use (17).

Prices and Budgets. A purchase is one candidate under fixed decoding settings. The main comparison uses $B = 1$ with zero marginal cost. For positive cost, we set $c_j = \gamma a_j / a_{\text{ref}}$ where a_j is the measured generation time per candidate and a_{ref} is the corresponding generation time of old non-thinking Qwen. We test $\gamma \in \{0, 10^{-4}, 10^{-3}, 10^{-2}\}$. Matched-cost comparisons use the larger of the two measured costs for both checkpoints to focus on differences in model quality. Separate comparisons use the cost of each checkpoint. For each fixed cost, we compare budgets $B_0/4$ to $4B_0$.

Evaluation Conditions. For each replacement, U_A uses the old services at their chosen prices; U_B replaces the checkpoint at those prices and reoptimizes the buyer's allocation; U_C also allows repricing. We decompose the change as

$$\Delta_{\text{tech}} = U_B - U_A, \qquad \Delta_{\text{price}} = U_C - U_B, \qquad \Delta_{\text{net}} = U_C - U_A. \tag{18}$$

Calibration chooses prices and allocations, which remain fixed during testing. Profile selection minimizes (15), then maximizes calibration success, then uses lexicographic price order. Acceptance requires a normalized gap at most 0.01. Baselines are fixed quotes, cost-plus quotes, equal spending, pass@1 routing, and greedy marginal success per price. Paired 95% confidence intervals use 2,000 bootstrap replicates of complete tasks, preserving every provider's outputs (Efron, 1979).

5.2 Model Upgrades and Repricing

All four replacements result in performance improvements at fixed prices, and all four gains are reduced after repricing (Table 1). The non-thinking gains of 3.8 and 2.2 points on LiveCodeBench and MATH-500 are reduced to losses of 1.7 and 1.6 points against the old market. The thinking gains of 2.2 and 0.8 points at fixed prices are reduced to gains of 0.8 and 0.3 points after repricing. Figure 2 splits the upgrade and repricing effects.

Table 1: Success before and after model replacement. Rates are percentages and changes are percentage points. N denotes non-thinking and T denotes thinking.

Domain	Mode	U_A	U_B	U_C	Δ_{tech}	Δ_{price}	Δ_{net}
LiveCodeBench	N	72.6	76.4	70.9	3.8	-5.5	-1.7
	T	84.1	86.3	84.9	2.2	-1.4	$+0.8$
MATH-500	N	94.8	97.0	93.2	2.2	-3.8	-1.6
	T	98.0	98.8	98.3	0.8	-0.5	$+0.3$

The mean net change across replacements is -0.5 points on coding with interval $[-1.6, 0.7]$ and -0.7 points on mathematics with interval $[-1.3, -0.1]$. Coding replacements have opposite net effects, with the interval for their average spanning zero. Across the four comparisons, the upgraded provider's geometric-mean price ratio is 1.32, its call ratio is 0.83, and its mean revenue share increases by 3.7 points. The changes in price and quantity are consistent with the mechanism identified by Theorem 2.

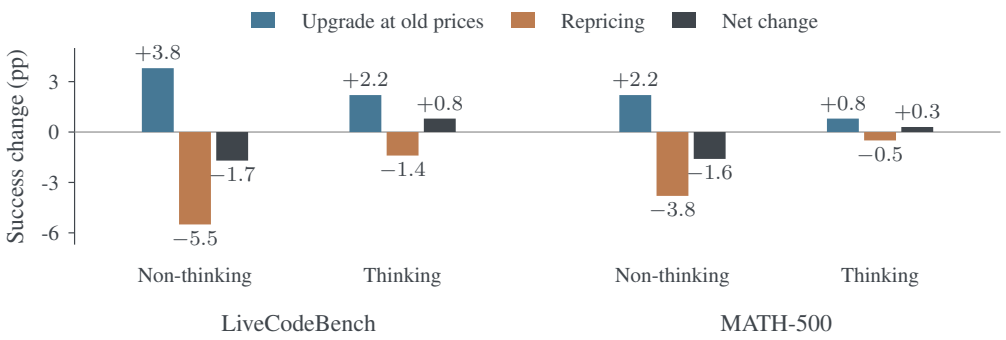

Figure 2: **Repricing reduces all four fixed-price gains.** The bars indicate the paired effects in (18). Non-thinking replacements are less successful than the old market.

All old/new profile pairs satisfy the numerical acceptance threshold in both domains. Median held-out empirical deviation gaps are 0.006 for coding and 0.004 for mathematics. Allowing every buyer choice compatible with the success intervals gives upper bounds of 0.041 and 0.028. These bounds account for demand uncertainty beyond the empirical deviation test; full diagnostics are in Appendix G.

5.3 Sensitivity Analysis

Task Dependence and Curve Prediction. The median absolute log error of call predictions from the fitted marginal curves is 0.27 for coding and 0.18 for mathematics. 71.4% and 64.3% of fits have marginal capability peaks in the observed range. The lower bounds of joint improvement over the finite evaluation range are -0.3 and -0.1 points. Both meet the 0.5-point noninferiority criterion for joint success.

Multiplying marginal failure rates without task alignment causes joint-success errors of 3.6 and 1.9 points, and repricing estimates are 1.1 and 0.6 points more negative (Table 2). All main comparisons keep task alignment. At fixed quotes, equal spending, pass@1 routing, and greedy allocation are 2.7, 5.8, and 0.9 success points worse than the optimized allocation. These numbers quantify the value of portfolio optimization at fixed prices.

Table 2: Sensitivity of the net effect. Values are shown in percentage points relative to the primary variant. For acceptance, the normalized best-response gap must be no greater than 0.01.

Modification	Coding	Math	Accepted (%)
Independent marginal failures	-1.1	-0.6	100.0
32-call maximum	$+0.7$	$+0.2$	100.0
128-call maximum	-0.3	-0.1	87.5
Matched cost, $\gamma = 10^{-3}$	$+0.1$	$+0.1$	100.0
Money-valued outside option	$+0.8$	$+0.5$	87.5
Success-triggered stopping	$+0.4$	$+0.2$	87.5
Public-test selection	-0.6	n/a	75.0

Costs and Call Limits. With matched cost $\gamma = 10^{-3}$, the mean net effects are -0.4 and -0.6 points; individual costs increase them by 0.2 and 0.1 points. This is consistent with the local theorem for positive serving costs. At 32 calls, the net effects are $+0.2$ points on coding and -0.5 points on mathematics; at 128 calls, the net effects are -0.8 points for both. The share of selected allocations on the boundary drops from 12.5% at 64 calls to 3.1% at 128 calls. Call limits affect both the magnitude and sign of the net effect.

Spending and Stopping Decisions. The net effects are between -0.9 and $+0.6$ points for buyers maximizing $vU(k) - p \cdot k$ with $v/B \in \{0.5, 1, 2, 4, 8\}$. At low task values, keeping unspent funds

weakens the hard-budget effect. Stopping after verified success increases the net effects by 0.4 and 0.2 points compared to the primary buyer model. Both alternatives change demand and provider receipts, so prices are recomputed for each buyer model.

Keeping the full old portfolio available at its original prices preserves attainable accuracy (Appendices E.2 and G.4).

6 Related Work

Test-Time Scaling. Additional inference computation can enhance reasoning using chain-of-thought prompting, repeated sampling, and verification (Wei et al., 2022; Wang et al., 2023; Cobbe et al., 2021; Lightman et al., 2024). Search and empirical scaling laws explore how to distribute inference computation across tasks (Yao et al., 2023; Snell et al., 2025; Wu et al., 2025). Routing and cascading select services at quoted prices (Chen et al., 2023; Ong et al., 2025; Dekoninck et al., 2025). Computational arbitrage explores combinations of services (Olmedo et al., 2026). Providers select inference effort (Artola Velasco et al., 2026). Here, providers select unit prices, users buy computation, and equilibrium quantities depend on the full scaling curve.

Competition and Model Quality. Individually more accurate algorithms can negatively affect collective performance through monoculture (Kleinberg and Raghavan, 2021). Improved representations and model expressivity can reduce equilibrium predictive welfare (Jagadeesan et al., 2023; Handina and Mazumdar, 2024). Here, a pointwise improvement of an accuracy-compute curve can reduce equilibrium accuracy even with fixed cross-provider dependence and perfect verification. The loss is due to endogenous prices reducing the computation purchased from the improved model. Our budget allocation setting builds on price competition with decreasing consumption value (Caro and Martínez-de-Albéniz, 2012).

Information Pricing. Query and model prices are limited by repeated and combined purchases (Deep and Koutris, 2017; Chen et al., 2019; Wu et al., 2026). Data market equilibria depend on buyer externalities (Hossain and Chen, 2024) and deviations available to sellers (Chaudhury et al., 2026). Our demand depends on verification success and equilibrium existence on the log-compute marginal capability.

7 Discussion and Conclusion

Pointwise improvements in model accuracy can increase provider revenue and reduce the accuracy users attain with a fixed budget. At zero cost, equilibrium quantities maximize $xf'(x)$, with peak heights determining revenue shares and capability at those quantities determining user accuracy. An upgrade can raise this peak and shift it toward smaller quantities, reducing the accuracy of the purchased portfolio despite improving the underlying curve. Reversal occurs even when every task is eventually solvable and persists under sufficiently small positive serving costs. In our paired experiments, repricing reduces all four fixed-price gains and reverses both non-thinking improvements.

How a scaling curve improves determines the direction of this response. Horizontal acceleration induces a proportional price increase in the zero-cost model and leaves equilibrium accuracy unchanged. Within the latent-rate family, improving the tail constant can reduce equilibrium accuracy, whereas increasing the tail exponent raises it. For a pointwise upgrade from f to g with unique equilibrium quantities, a nondecreasing ratio g'/f' ensures that neither purchased computation nor equilibrium accuracy declines (Proposition 5). Comparisons of test-time scaling should report changes in $xf'(x)$ alongside accuracy at fixed computation.

Task-aligned empirical demand accommodates integer calls and cross-model dependence. Continuous-price best responses quantify providers' incentives to deviate from the reported prices. Statistical bounds also account for changes in the buyer's allocation under uncertainty in estimated success probabilities. Allowing buyers to retain unspent funds or stop after verified success changes both demand and provider receipts. Extending this analysis to heterogeneous buyers, capacity constraints, and subscriptions would clarify how alternative purchasing and pricing arrangements change the gains from model improvements.

AI USE STATEMENT

We used large language models (LLMs) to improve grammar and correct spelling errors in the manuscript. The authors take full responsibility for the final content of this work.

REFERENCES

Ander Artola Velasco, Dimitrios Rontogiannis, Stratis Tsirtsis, and Manuel Gomez-Rodriguez. Test-time compute games. *arXiv preprint arXiv:2601.21839*, 2026.

Maria-Florina Balcan, Travis Dick, and Ellen Vitercik. Dispersion for data-driven algorithm design, online learning, and private optimization. In *IEEE Symposium on Foundations of Computer Science*, 2018.

Maria-Florina Balcan, Tuomas Sandholm, and Ellen Vitercik. Estimating approximate incentive compatibility. In *ACM Conference on Economics and Computation*, 2019.

David Blackwell. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24(2):265–272, 1953.

J. Frédéric Bonnans and Alexander Shapiro. *Perturbation Analysis of Optimization Problems*. Springer, 2000.

Vitor Bosshard, Benedikt Bünz, Benjamin Lubin, and Sven Seuken. Computing Bayes-Nash equilibria in combinatorial auctions with verification. *Journal of Artificial Intelligence Research*, 69:531–570, 2020.

Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

Bradley Brown, Jordan Juravsky, Ryan Ehrlich, Ronald Clark, Quoc V. Le, Christopher Ré, and Azalia Mirhoseini. Large Language Monkeys: Scaling inference compute with repeated sampling. *arXiv preprint arXiv:2407.21787*, 2024.

Felipe Caro and Victor Martínez-de-Albéniz. Product and price competition with satiation effects. *Management Science*, 58(7):1357–1373, 2012.

Bhaskar Ray Chaudhury, Jugal Garg, Eklavya Sharma, and Jiaxin Song. Equilibrium pricing in oligopolistic data markets. *arXiv preprint arXiv:2608.14018*, 2026.

Lingjiao Chen, Paraschos Koutris, and Arun Kumar. Towards model-based pricing for machine learning in a data marketplace. In *ACM SIGMOD International Conference on Management of Data*, 2019.

Lingjiao Chen, Matei Zaharia, and James Zou. FrugalGPT: How to use large language models while reducing cost and improving performance. *arXiv preprint arXiv:2305.05176*, 2023.

Mark Chen, Jerry Tworek, Heewoo Jun, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.

Yiling Chen and Bo Waggoner. Informational substitutes. In *IEEE Symposium on Foundations of Computer Science*, 2016.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.

Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley, second edition, 2006.

Shaleen Deep and Paraschos Koutris. The design of arbitrage-free data pricing schemes. In *International Conference on Database Theory*, 2017.

Jasper Dekoninck, Maximilian Baader, and Martin Vechev. A unified approach to routing and cascading for LLMs. In *International Conference on Machine Learning*, 2025.

Asen L. Dontchev and R. Tyrrell Rockafellar. *Implicit Functions and Solution Mappings: A View from Variational Analysis*. Springer, second edition, 2014.

Bradley Efron. Bootstrap methods: Another look at the jackknife. *The Annals of Statistics*, 7(1):1–26, 1979.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, et al. The Llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.

Tinashe Handina and Eric Mazumdar. Understanding model selection for learning in strategic environments. In *Advances in Neural Information Processing Systems*, volume 37, 2024.

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the MATH dataset. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, volume 1, 2021.

Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.

Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. In *International Conference on Learning Representations*, 2020.

Safwan Hossain and Yiling Chen. Equilibrium of data markets with externality. In *International Conference on Machine Learning*, 2024.

Meena Jagadeesan, Michael I. Jordan, Jacob Steinhardt, and Nika Haghtalab. Improved Bayes risk can yield reduced social welfare under competition. In *Advances in Neural Information Processing Systems*, volume 36, 2023.

Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. LiveCodeBench: Holistic and contamination free evaluation of large language models for code. In *International Conference on Learning Representations*, 2025.

Stephen Joe and Frances Y. Kuo. Constructing Sobol sequences with better two-dimensional projections. *SIAM Journal on Scientific Computing*, 30(5):2635–2654, 2008.

Elliot Myunghoon Kim, Avi Garg, Kenny Peng, and Nikhil Garg. Correlated errors in large language models. In *International Conference on Machine Learning*, 2025.

Jon Kleinberg and Manish Raghavan. Algorithmic monoculture and social welfare. *Proceedings of the National Academy of Sciences*, 118(22):e2018340118, 2021.

Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. In *International Conference on Learning Representations*, 2024.

Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. Is your code generated by ChatGPT really correct? Rigorous evaluation of large language models for code generation. In *Advances in Neural Information Processing Systems*, volume 36, 2023.

Paul Milgrom and Chris Shannon. Monotone comparative statics. *Econometrica*, 62(1):157–180, 1994.

Paul Milgrom and Ilya Segal. Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601, 2002.

Ricardo Olmedo, Bernhard Schölkopf, and Moritz Hardt. Computational arbitrage in AI model markets. *arXiv preprint arXiv:2603.22404*, 2026.

Isaac Ong, Amjad Almahairi, Vincent Wu, Wei-Lin Chiang, Tianhao Wu, Joseph E. Gonzalez, M. Waleed Kadous, and Ion Stoica. RouteLLM: Learning to route LLMs from preference data. In *International Conference on Learning Representations*, 2025.

Qwen Team. Qwen3-4B. *Model card*, 2025a.

Qwen Team. Qwen3-4B-Instruct-2507. *Model card*, 2025b.

Qwen Team. Qwen3-4B-Thinking-2507. *Model card*, 2025c.

René L. Schilling, Renming Song, and Zoran Vondraček. *Bernstein Functions: Theory and Applications*. De Gruyter, second edition, 2012.

Robert J. Serfling. *Approximation Theorems of Mathematical Statistics*. Wiley, 1980.

Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling LLM test-time compute optimally can be more effective than scaling parameters for reasoning. In *International Conference on Learning Representations*, 2025.

Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer, 2009.

Martin J. Wainwright and Michael I. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(1–2):1–305, 2008.

Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V. Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. In *International Conference on Learning Representations*, 2023.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc V. Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems*, volume 35, 2022.

Yangzhen Wu, Zhiqing Sun, Shanda Li, Sean Welleck, and Yiming Yang. Inference scaling laws: An empirical analysis of compute-optimal inference for LLM problem-solving. In *International Conference on Learning Representations*, 2025.

Yihang Wu, Zhengyu Jin, Yicheng Fu, Jinfei Liu, and Kui Ren. Arbitrage-free data pricing. *arXiv preprint arXiv:2606.10451*, 2026.

An Yang, Anfeng Li, Baosong Yang, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.

Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of Thoughts: Deliberate problem solving with large language models. In *Advances in Neural Information Processing Systems*, volume 36, 2023.

A Proofs of the Equilibrium Characterization

A.1 Buyer Demand

For $t > 0$, we define

$$d_j(t) = \begin{cases} (f_j')^{-1}(t), & t < f_j'(0+), \\ 0, & t \geq f_j'(0+). \end{cases} \tag{19}$$

The case $t \geq f_j'(0+)$ does not occur when $f_j'(0+) = \infty$. By strict concavity, f_j' is strictly decreasing. Under the assumptions in Section 2, d_j is continuous and nonincreasing, strictly decreasing wherever it is positive, and $d_j(t) \to \infty$ as $t \downarrow 0$ and $d_j(t) \to 0$ as $t \to \infty$.

For any price vector p with all entries positive, the budget multiplier is the unique solution of

$$\sum_j p_j d_j(\mu p_j) = B, \tag{20}$$

where the left-hand side is continuous, ranges between infinity and zero, and is strictly decreasing everywhere positive. Its associated allocation satisfies the necessary and sufficient optimality conditions for the strictly concave buyer problem

$$f_j'(x_j) = \mu p_j \quad \text{if } x_j > 0, \qquad f_j'(0+) \leq \mu p_j \quad \text{if } x_j = 0. \tag{21}$$

These conditions also apply to the case of deviations that exclude some providers from the purchased portfolio.

A.2 Positive Equilibrium Revenue

Assume provider j has no demand in a candidate equilibrium. Let μ be the multiplier on the buyer's allocation across its competitors. Pick $q > 0$ such that $f_j'(0+)/q > \mu$. Then if j has no demand in (q, p_{-j}), its competitors have the same optimization problem and multiplier. The boundary condition in (21) would then imply $f_j'(0+) \leq \mu q$, a contradiction. The deviation leads to positive demand and positive revenue. All providers have positive demand at every positive-price equilibrium.

A.3 Implementing a Target Quantity

Given p_{-j}, let

$$D_{-j}(\mu) = \sum_{h \neq j} p_h d_h(\mu p_h). \tag{22}$$

There is a unique $\mu_B > 0$ such that $D_{-j}(\mu_B) = B$. Let

$$G_j(\mu) = \mu\big[B - D_{-j}(\mu)\big]. \tag{23}$$

This function is continuous and strictly increasing on $[\mu_B, \infty)$. To see this, let $\mu_2 > \mu_1 \geq \mu_B$, then

$$G_j(\mu_2) - G_j(\mu_1) = (\mu_2 - \mu_1)\big[B - D_{-j}(\mu_2)\big] + \mu_1\big[D_{-j}(\mu_1) - D_{-j}(\mu_2)\big]. \tag{24}$$

Both terms on the right hand side are non-negative. If the first term is zero, then $D_{-j}(\mu_2) = B > 0$, so some competitor demand is strictly reduced and therefore the second term is positive. Furthermore, $G_j(\mu_B) = 0$ and $G_j(\mu) \to \infty$ as $\mu \to \infty$.

To implement any $z > 0$, find a unique $\mu(z)$ such that

$$G_j(\mu(z)) = H_j(z), \qquad q(z) = \frac{f_j'(z)}{\mu(z)}. \tag{25}$$

Then

$$q(z)z = \frac{H_j(z)}{\mu(z)} = B - D_{-j}(\mu(z)).$$

The target quantity, competitors' demands and the implemented price satisfy (21) and budget equality. Provider j's revenue is

$$r_j(H_j(z)) = B - D_{-j}\big(G_j^{-1}(H_j(z))\big). \tag{26}$$

It is nondecreasing in $H_j(z)$ and strictly increasing at any configuration at which the competitors receive positive expenditure.

A.4 Proof of Theorem 1

In equilibrium all providers are active. Suppose that x_j does not maximize H_j. Then there exists a $z > 0$ such that $H_j(z) > H_j(x_j)$. By (25), setting z increases the multiplier and strictly decreases the competitors' aggregate expenditure. It increases provider j's revenue. This is a contradiction to the definition of an equilibrium. It follows that H_j attains a finite global maximum at the equilibrium quantity.

In the other direction, let $z_j \in Z_j$ be arbitrary and specify prices via (6). Since $\mu^* = \sum_h M_h/B$, the buyer's optimality conditions are satisfied. If a deviation increased provider j's revenue, the aggregate receipts for the other providers must have decreased, so $\mu' > \mu^*$. The deviating provider has positive demand and satisfies

$$H_j(x_j') = \mu' R_j' > \mu^* R_j^* = M_j,$$

a contradiction. This establishes sufficiency for arbitrary positive-price deviations.

Finally, at any equilibrium, multiplying $f_j'(x_j) = \mu p_j$ by x_j and summing gives $\mu B = \sum_j M_j$. Substitution recovers all prices and revenues in (6). Unique maximizers yield a unique price vector.

A.5 Endpoint Behavior and Curvature

Continuity at zero and concavity yield

$$0 \le x f'(x) \le f(x) - f(0) \longrightarrow 0 \quad (x \downarrow 0). \tag{27}$$

For bounded f, concavity also leads to

$$0 \le x f'(x) \le 2\big[f(x) - f(x/2)\big] \longrightarrow 0 \quad (x \to \infty). \tag{28}$$

Thus, a continuous positive H has to attain a finite maximum. For any smooth interior maximum, one has $H'(x) = f'(x) + x f''(x) = 0$ and thus (7). When several stationary points exist, their values must be compared.

A.6 Latent-Rate Representation

The exponential family moment identities (Wainwright and Jordan, 2008) imply that

$$f'(x) = \frac{\mathbb{E}[\Lambda e^{-\Lambda x}]}{\mathbb{E}[e^{-\Lambda x}]}, \tag{29}$$

$$f''(x) = -\frac{\mathbb{E}[\Lambda^2 e^{-\Lambda x}]}{\mathbb{E}[e^{-\Lambda x}]} + \left(\frac{\mathbb{E}[\Lambda e^{-\Lambda x}]}{\mathbb{E}[e^{-\Lambda x}]}\right)^2. \tag{30}$$

These quantities are respectively the mean and negative variance of Λ in the failure-tilted distribution. For correlated rates, the gradient and Hessian of $-\log \mathbb{E}[e^{-\Lambda \cdot x}]$ are $\mathbb{E}_x[\Lambda]$ and $-\operatorname{Cov}_x(\Lambda)$, respectively. The joint objective is concave and may not be additive across providers.

Let $t = x/\tau$ in the family (8). The Laplace transform of the Gamma component is $(1+t)^{-\kappa}$. We compute

$$\mathbb{E}\left[e^{-t\sum_{r=1}^{N} E_r} \mid N\right] = (1+t)^{-N}, \tag{31}$$

$$\mathbb{E}[(1+t)^{-N}] = \exp\{\alpha[(1+t)^{-1} - 1]\} = \exp\left(-\frac{\alpha t}{1+t}\right). \tag{32}$$

This proves (8). The rate is a.s. positive. If $\kappa > 1$ then $\Lambda \ge G/\tau$, so that $\mathbb{E}[1/\Lambda] \le \tau \mathbb{E}[1/G] = \tau/(\kappa - 1)$.

When $\alpha > \kappa$, the derivative in (9) changes sign once at $t^* = (\alpha + \kappa)/(\alpha - \kappa)$. Upon substituting this into (9) we obtain (10) and (11). When $\alpha \le \kappa$ the sign is positive, and $H(\tau t) \uparrow \kappa$, so the supremum is not attained.

A.7 Finite Improvements and Price Responses

All derivatives in (12) have consistent signs for $\kappa < \alpha < 2\kappa$. Integrating over any closed subinterval thus proves Theorem 2 for a finite upgrade. Defining $M_{-j} = \sum_{h \neq j} M_h$, the upgraded provider's revenue satisfies

$$\frac{dR_j^*}{d\alpha_j} = B \frac{M_j' M_{-j}}{(M_j + M_{-j})^2} > 0. \tag{33}$$

Its unit price satisfies

$$\frac{d \log p_j^*}{d\alpha_j} = M_j' \left(\frac{1}{M_j} - \frac{1}{\sum_h M_h} \right) - \frac{z_j'}{z_j} > 0. \tag{34}$$

The competitors' quantities remain fixed while their prices and revenues fall by the common factor arising from an increase in $\sum_h M_h$. The old allocation at unchanged prices remains feasible and improves strictly. A reduction at the new equilibrium cannot be explained by a reduction in fixed-price achievable accuracy.

For a horizontal acceleration, $\widetilde{f}_j(x) = f_j(s_j x)$ implies $\widetilde{H}_j(x) = H_j(s_j x)$, so $\widetilde{M}_j = M_j$, $\widetilde{z}_j = z_j / s_j$, and $\widetilde{p}_j = s_j p_j$. Revenue and delivered capability are unchanged. More generally, effective computation $y_j = s_j x_j$ and price $q_j = p_j / s_j$ preserve expenditures. Given physical cost $c_j x_j$, the effective cost is c_j / s_j per unit of y_j. An unchanged positive physical cost does not preserve the original price game.

B Further Properties of Pricing Equilibria

B.1 Revenue Shares and Price Responses

Let $M_\Sigma = \sum_h M_h$ and $s_j = M_j / M_\Sigma$. In a unique-maximizer equilibrium (as in Theorem 1), we have

$$R_j^* = B s_j, \qquad p_j^* = \frac{B s_j}{z_j}. \tag{35}$$

We can perform general comparative statics for changes in an arbitrary smooth model. Suppose a parameter θ only changes provider j's capability curve, and its maximum remains unique, interior, and nondegenerate. Dots denote derivatives with respect to θ. By the envelope theorem (Milgrom and Segal, 2002) and first-order condition,

$$\dot{M}_j = \partial_\theta H_j(z_j, \theta), \qquad \dot{z}_j = -\frac{\partial_{x\theta} H_j(z_j, \theta)}{\partial_{xx} H_j(z_j, \theta)}. \tag{36}$$

Other variables z_h, M_h remain unchanged for competitors $h \neq j$. Differentiating (35), we obtain

$$\frac{\dot{R}_j^*}{R_j^*} = (1 - s_j) \frac{\dot{M}_j}{M_j}, \qquad \frac{\dot{p}_j^*}{p_j^*} = (1 - s_j) \frac{\dot{M}_j}{M_j} - \frac{\dot{z}_j}{z_j}, \tag{37}$$

$$\frac{\dot{R}_h^*}{R_h^*} = -s_j \frac{\dot{M}_j}{M_j}, \qquad \frac{\dot{p}_h^*}{p_h^*} = -s_j \frac{\dot{M}_j}{M_j}. \tag{38}$$

The equilibrium capability $V_j^* = f_j(z_j, \theta)$ satisfies

$$\dot{V}_j^* = \partial_\theta f_j(z_j, \theta) + f_j'(z_j, \theta) \dot{z}_j. \tag{39}$$

The first term is nonnegative for a pointwise improvement. The second term measures the change of equilibrium computation. The sign of $\dot{M}_j$ determines whether the provider's revenue increases or decreases. Similarly, the sign of (39) determines the response of the user's success rate. In the independent providers case,

$$\frac{dU^*}{d\theta} = \exp\left(-\sum_h V_h^* \right) \dot{V}_j^*. \tag{40}$$

A change in a competitor's generation curve is not necessary for its price or revenue to respond.

For finite changes, the competitors' prices and revenues have the same ratio,

$$\frac{p_h^{\text{new}}}{p_h^{\text{old}}} = \frac{R_h^{\text{new}}}{R_h^{\text{old}}} = \frac{M_j^{\text{old}} + M_{-j}}{M_j^{\text{new}} + M_{-j}}, \qquad h \neq j. \tag{41}$$

The firms' purchased quantities are not changed. This equation covers the response to competition in Theorem 2 without fixing the competitor's price between the two equilibria.

B.2 Improvements that Preserve Equilibrium Capability

Pointwise improvement alone permits the reversal in Theorem 2. The following condition orders equilibrium outcomes through relative marginal gains. The approach is comparative statics, see Milgrom and Shannon (1994).

Proposition 5 (Monotone relative marginal improvement)**.** *We assume that f and g satisfy buyer assumptions, and they have unique finite maximizers z_f and z_g of $xf'(x)$ and $xg'(x)$. We assume that $g(x) \geq f(x)$ for every $x \geq 0$ and $r(x) = g'(x)/f'(x)$ is nondecreasing. Then $z_g \geq z_f$ and $g(z_g) \geq f(z_f)$. If we replace one or more of the independent providers with functions that satisfy these conditions, then equilibrium success will not decrease.*

Proof. For any $x < z_f$, uniqueness of the old maximum implies $xf'(x) < z_f f'(z_f)$. The positivity and monotonicity of r then give

$$xg'(x) = r(x)xf'(x) < r(z_f)z_f f'(z_f) = z_f g'(z_f).$$

No point below z_f can be the new maximizer, so $z_g \geq z_f$. Monotonicity of g and pointwise dominance now give

$$g(z_g) \geq g(z_f) \geq f(z_f).$$

Other providers keep their quantities by Theorem 1. The same argument can be applied to each changed curve to prove the portfolio claim. □

This condition differs from previous comparisons across arbitrary decision problems (Blackwell, 1953; Chen and Waggoner, 2016) as it focuses on marginal gains across computation levels for a specific inference problem. In the Gaussian example in Appendix H, reducing variance from a to $a' < a$ at sampling variance b leads to the following ratio of marginal gains:

$$\frac{g'(x)}{f'(x)} = \left(\frac{ax+b}{a'x+b}\right)^2, \tag{42}$$

which is increasing. This variance reduction increases both the equilibrium computation and private precision. A pure horizontal acceleration does not need to satisfy Proposition 5. Its equilibrium effect is derived in Appendix B.5 via price reparameterization.

B.3 Behavior across the Latent-Rate Family

Let $\rho = \alpha/\kappa > 1$. Divide (10)–(11) by the respective scale parameters:

$$\frac{z}{\tau} = \frac{\rho+1}{\rho-1}, \qquad \frac{M}{\kappa} = \frac{(\rho+1)^2}{4\rho}, \qquad \frac{V^*}{\kappa} = \frac{\rho+1}{2} + \log\frac{2\rho}{\rho-1}. \tag{43}$$

Differentiating the last expression with respect to ρ:

$$\frac{d(V^*/\kappa)}{d\rho} = \frac{1}{2} - \frac{1}{\rho(\rho-1)}. \tag{44}$$

Equilibrium capability decreases on $(1, 2)$ and increases on $(2, \infty)$. The unique minimum occurs at $\rho = 2$:

$$z = 3\tau, \qquad M = \frac{9\kappa}{8}, \qquad V^* = \kappa\left(\frac{3}{2} + \log 4\right). \tag{45}$$

These expressions describe the change of sign without a specific budget or number of providers.

Near the boundary of existence, let $\rho = 1 + u$ with $u \downarrow 0$. Equation (43) gives

$$\frac{z}{\tau} = \frac{2}{u} + 1, \tag{46}$$

$$\frac{M}{\kappa} = 1 + \frac{u^2}{4(1+u)}, \tag{47}$$

$$\frac{V^*}{\kappa} = 1 + \frac{u}{2} + \log \frac{2(1+u)}{u}. \tag{48}$$

The purchased computation diverges and the finite maximum tends to the boundary value κ. The difference $M - \kappa$ is quadratic in u. A curve might have a finite equilibrium maximum with a small vertical separation from its tail, rendering peak estimation difficult in this region.

For large ρ, $z/\tau \to 1$ and

$$\frac{M}{\kappa} = \frac{\rho}{4} + \frac{1}{2} + \frac{1}{4\rho}, \qquad \frac{V^*}{\kappa} = \frac{\rho}{2} + \frac{1}{2} + \log 2 + O(\rho^{-1}). \tag{49}$$

In the same family, stronger improvements eventually lead to more capability being bought at equilibrium. The reversal region is a proper subset of the equilibrium region.

B.4 Budget Changes and Provider Entry

At zero serving cost, when B is multiplied by $b > 0$, all equilibrium prices and revenues are multiplied by b while z remains invariant. Indeed, this follows directly from (6). It states that in a market where prices can be adjusted to the new budget, all prices and revenues are scaled by the same factor. At given prices, an increased budget widens the choice of consumers.

Let us model an additional independent provider who enters with a finite maximizing quantity z_0, peak $M_0 > 0$ and capability $V_0 = f_0(z_0)$. Let M_Σ and V_Σ represent the sums over incumbent providers. In the enlarged market equilibrium, every incumbent retains its quantity, and

$$\frac{p_j^{\text{entry}}}{p_j^{\text{old}}} = \frac{R_j^{\text{entry}}}{R_j^{\text{old}}} = \frac{M_\Sigma}{M_\Sigma + M_0}. \tag{50}$$

The increase in success is

$$U^{\text{entry}} - U^{\text{old}} = e^{-V_\Sigma}(1 - e^{-V_0}) > 0. \tag{51}$$

The analysis treats the entrant as an additional independent source of solutions. Two vendors offering the same model and sharing the task failures are substitute suppliers that do not constitute two independent capability contributions.

B.5 Speed Changes and Heterogeneous Buyers

A horizontal scaling has a broader invariance than its effect on a representative buyer. Let us index buyers by i. Let the outcome payoffs be given in terms of effective quantities $y_{ij} = s_j x_{ij}$. Buyers may have different budgets and different preferences, which need not be additive. Let $q_j = p_j/s_j$. Then

$$\sum_j p_j x_{ij} = \sum_j q_j y_{ij} \tag{52}$$

for each buyer and each allocation. When the feasible sets of quantities transform as $y_{ij} = s_j x_{ij}$, feasible choices and optimal outcomes are related by this change of variables. A consistent transformation of tie-breaking rules preserves the selected demand.

For zero serving cost, provider revenue is

$$p_j \sum_i x_{ij} = q_j \sum_i y_{ij}. \tag{53}$$

Each unilateral positive-price deviation in one market leads to one in the other, so the equilibrium sets correspond. At cost c_j per physical computation unit,

$$(p_j - c_j) \sum_i x_{ij} = (q_j - c_j/s_j) \sum_i y_{ij}. \tag{54}$$

Pure acceleration changes the effective serving cost to c_j/s_j. This argument relies on quantity constraints converting with the effective units. Fixed latency constraints, integer call restrictions, and physical capacity constraints can alter the feasible sets.

B.6 Bounded Curves Can Have Arbitrarily Low Utilization

The one-half precision ratio in Appendix H is specific to its rational capability function. Consider

$$f_r(x) = 1 - (1+x)^{-r}, \qquad r > 0. \tag{55}$$

All curves are increasing, strictly concave and with $f_r(x) \le 1$. The marginal capability derivative is

$$H_r'(x) = \frac{r(1-rx)}{(1+x)^{r+2}}. \tag{56}$$

The unique equilibrium quantity is $z_r = 1/r$ where

$$f_r(z_r) = 1 - \left(\frac{r}{r+1}\right)^r \longrightarrow 0 \qquad (r \downarrow 0). \tag{57}$$

The limit follows by taking logarithms of the power term. For m independent identical providers, equilibrium success approaches zero while unlimited-computation success remains at $1 - e^{-m}$.

Here is a latent rate example. Let $N \sim \mathrm{Poisson}(1)$ and let its jumps be $\mathrm{Gamma}(r, 1)$. The Laplace transform of the compound Poisson rate process is

$$\exp\big((1+x)^{-r} - 1\big) = e^{-f_r(x)}. \tag{58}$$

This example has a nonzero probability of failure even at time infinity. The unbounded family used in Theorem 2 does not. It shows that boundedness and concavity are not sufficient to have any uniform positive utilization fraction.

C Positive Serving Costs

C.1 Profit from Implemented Computation

Fix prices of the competitors close to the zero-cost equilibrium. All competitors are active, close to the equilibrium allocation and the functions in (23) are smooth, so that their derivatives are

$$D'_{-j}(\mu) < 0, \qquad G_j'(\mu) = B - D_{-j}(\mu) - \mu D'_{-j}(\mu) > 0.$$

Differentiating (26), we obtain

$$r_j'(h) = \frac{-D'_{-j}(\mu)}{B - D_{-j}(\mu) - \mu D'_{-j}(\mu)} > 0, \qquad \mu = G_j^{-1}(h). \tag{59}$$

The profit of a deviation implementing x is

$$\pi_j(x) = r_j(H_j(x)) - c_j x. \tag{60}$$

The previous expression also accounts for the dependence of the buyer's response to the seller's price.

C.2 Local Equilibrium Conditions

Let x be a positive allocation. It is implemented by

$$\mu(x) = \frac{\sum_h H_h(x_h)}{B}, \qquad p_j(x) = \frac{f_j'(x_j)}{\mu(x)}. \tag{61}$$

Let $A_j(x) > 0$ be (59) against $p_{-j}(x)$. The simultaneous first order conditions are

$$E_j(x, c) = A_j(x) H_j'(x_j) - c_j = 0. \tag{62}$$

At $(z, 0)$, $H_j'(z_j) = 0$ so

$$\left.\frac{\partial E_j}{\partial x_k}\right|_{(z,0)} = \mathbf{1}\{j = k\} A_j(z) H_j''(z_j). \tag{63}$$

The invertibility of (63) gives a locally unique continuously differentiable $x(c)$ and its prices by the implicit function theorem (Dontchev and Rockafellar, 2014). In particular,

$$\left.\frac{\partial x_j}{\partial c_k}\right|_{c=0} = \frac{\mathbf{1}\{j=k\}}{A_j(z)H_j''(z_j)}. \tag{64}$$

Own computation decreases to first order with own cost. This is local to the zero cost equilibrium.

C.3 Global Optimality of Nearby Responses

Let I_j be a sufficiently small interval around z_j. By the uniqueness of the global maximum and the strict endpoint gaps, we have some $\delta_j > 0$ such that

$$H_j(x) \le M_j - \delta_j \qquad (x \notin I_j). \tag{65}$$

At the zero-cost equilibrium, r_j is strictly increasing on the relevant interval up to M_j, for competitors receive positive expenditure there. Quantities outside I_j incur a strictly positive revenue loss. Continuity implies uniform loss for competitors' prices in a sufficiently small compact neighborhood. Endpoint gaps cover arbitrarily small and large quantities and compactness covers the region in between.

The cost term $-c_j x$ cannot increase the profit of a distant deviation. The change of profit of the nearby candidate is continuous and all best responses are in I_j for sufficiently small costs. The second derivative at the zero-cost optimum of I_j is

$$r_j'(M_j)H_j''(z_j) < 0.$$

The above expression remains strictly negative if intervals are shrunk. The local solution of (62) is the unique best response in I_j and a global best response. This proves Theorem 3. Buyer success and provider profit are continuous on this branch and strict comparisons between two markets can be preserved for sufficiently small costs.

For (8), $H(0) = 0$, $H(\infty) = \kappa < M$ and the finite maximum has negative second derivative when $\alpha > \kappa$. The theorem applies to each finite upgrade strictly within the reversal interval. The equilibrium curve is locally unique in a neighborhood of the zero-cost solution.

C.4 First-Order Effects of Serving Costs

The local equilibrium system determines the initial price reaction to costs. Let $s_j = M_j / \sum_h M_h$ and take derivatives at the zero cost equilibrium. Setting $H_h'(z_h) = 0$ gives

$$f_h''(z_h) = -\frac{f_h'(z_h)}{z_h} = -\frac{\mu p_h}{z_h}, \tag{66}$$

while

$$D_{-j}'(\mu) = \sum_{h \ne j} \frac{p_h^2}{f_h''(z_h)} = -\frac{B - R_j}{\mu} \tag{67}$$

by differentiation of the competitor demand. Then (59) simplifies to

$$A_j(z) = r_j'(M_j) = \frac{1 - s_j}{\mu}. \tag{68}$$

Therefore (63) gives

$$\left.\frac{\partial x_j}{\partial c_j}\right|_{c=0} = \frac{\mu}{(1-s_j)H_j''(z_j)} < 0, \qquad \left.\frac{\partial x_j}{\partial c_k}\right|_{c=0} = 0 \quad (k \ne j). \tag{69}$$

Finally, the common multiplier satisfies $\mu(x) = \sum_h H_h(x_h)/B$, and its derivative with respect to any cost is zero at $c = 0$, since each $H_h'(z_h)$ is zero. Differentiating $p_j = f_j'(x_j)/\mu$ gives

$$\left.\frac{\partial p_j}{\partial c_j}\right|_{c=0} = \frac{f_j''(z_j)}{(1-s_j)H_j''(z_j)} > 0, \qquad \left.\frac{\partial p_j}{\partial c_k}\right|_{c=0} = 0 \quad (k \ne j). \tag{70}$$

Therefore, own prices rise and own quantities decrease to first order. Revenue is locally stationary in all cost directions since $R_j = H_j(x_j)/\mu$, and both the numerator and denominator of this ratio have a zero derivative. Profit satisfies

$$\left.\frac{\partial \pi_j}{\partial c_k}\right|_{c=0} = -\mathbf{1}\{j=k\} z_j. \tag{71}$$

These derivatives characterize the local equilibrium branch of Theorem 3 and allow for heterogeneous capability functions.

C.5 Curvature at an Interior Positive-Cost Equilibrium

Consider any differentiable equilibrium where all firms have positive demand. Let

$$e_j = -\frac{x_j f_j''(x_j)}{f_j'(x_j)}, \qquad T = \sum_h \frac{R_h}{e_h}, \qquad \sigma_j = \frac{R_j/e_j}{T}. \tag{72}$$

Differentiating the buyer condition (4) for a unilateral change in p_j,

$$\frac{dx_h}{dp_j} = \frac{p_h \mu' + \mu \mathbf{1}\{h=j\}}{f_h''(x_h)}. \tag{73}$$

Substituting $f_h''(x_h) = -e_h \mu p_h / x_h$ and differentiating the budget equality (4) gives

$$0 = x_j - \frac{x_j}{e_j} - \frac{\mu'}{\mu} T, \qquad \frac{\mu'}{\mu} = \frac{x_j(1-1/e_j)}{T}. \tag{74}$$

The resulting revenue elasticity is

$$\frac{d\log R_j}{d\log p_j} = \left(1 - \frac{1}{e_j}\right)(1-\sigma_j). \tag{75}$$

When $p_j > c_j$, $\pi_j = (1 - c_j/p_j) R_j$ and differentiating at a stationary price gives

$$\left(\frac{1}{e_j} - 1\right)(1-\sigma_j) = \frac{c_j}{p_j - c_j}. \tag{76}$$

At zero cost, $e_j = 1$ and we recover (7). At positive cost, $1 - \sigma_j > 0$ and (76) requires $e_j < 1$. For the Gaussian curve $f_j(x) = x/(a_j x + b_j)$,

$$e_j = \frac{2 a_j x_j}{a_j x_j + b_j}, \tag{77}$$

so every such positive-cost equilibrium has $x_j < b_j/a_j$. These are necessary conditions for an interior stationary response. Global optimality of the small-cost branch is proved in Theorem 3. Appendix H considers the symmetric Gaussian market.

D Approximation and Statistical Guarantees

D.1 Proof of Proposition 4

Substituting $S = \sum_j H_j(z_j)$ and $p_j = B f_j'(z_j)/S$, we have $\mu = S/B$ and the buyer receives z. The current revenue is $B H_j(z_j)/S$. A profitable deviation fulfils $\mu' > S/B$ and

$$R_j' = \frac{H_j(x_j')}{\mu'} \le \frac{M_j}{\mu'} < \frac{B M_j}{S}.$$

Dividing the above by the current revenue establishes (14). The inequality trivially holds for a nonprofitable deviation.

Arbitrarily good approximations of equilibria exist, even with a finite but non-attained supremum. Specifically, we have $f(x) = \kappa \log(1 + x/\tau)$ with $M = \kappa$ and

$$\frac{M}{H(x)} = 1 + \frac{\tau}{x}.$$

Note that $x \ge \tau/\varepsilon$ yields a $(1+\varepsilon)$-approximate equilibrium by Proposition 4.

D.2 Continuous-Price Best Responses

Fix p_{-j}. An allocation with $k_j > 0$ is affordable if and only if

$$q \le q_k = \frac{B - \sum_{h \neq j} p_h k_h}{k_j}. \tag{78}$$

For allocations with $k_j = 0$, affordability is independent of this price. Between consecutive positive thresholds, each interval $(a, b]$ has a fixed feasible set and hence a fixed highest-ranked allocation. If the provider sells a positive quantity, its profit increases with price and is maximized at b within that interval. If its quantity is zero, its profit is zero. Prices above B yield zero sales. A global best response can thus be found by checking the finite set of positive thresholds together with one price above B.

Demand exhibits finitely many switching points of affordability, a structure that has also arisen in discontinuous data-driven optimization (Balcan et al., 2018). A buyer preference ranked heap removes an allocation after its threshold has been evaluated. The time complexity of sorting and processing is $O(|\mathcal{K}| \log |\mathcal{K}|)$. The empirical game is solved at its threshold values with numerical profit tolerance $10^{-8}B$.

For a buyer maximizing $v\widehat{U}(k) - p \cdot k$, there can be additional changes in rankings at

$$q = \frac{v[\widehat{U}(k) - \widehat{U}(\ell)] - \sum_{h \neq j} p_h (k_h - \ell_h)}{k_j - \ell_j}, \qquad k_j \neq \ell_j. \tag{79}$$

These events can also be added to the affordability events, partitioning the price axis again into regions of constant choice. This is reflected in the money-valued outside-option comparison.

D.3 Task-Aligned Success Estimation

For task i, provider j has per-candidate success probability θ_{ij}. For n independent candidates, let c_{ij} be correct. A uniformly random subset of k_j of these candidates has no successes with conditional probability

$$\frac{\binom{n-c_{ij}}{k_j}}{\binom{n}{k_j}}. \tag{80}$$

As a symmetric subset average (Serfling, 1980), its expectation over the generated bank is $(1-\theta_{ij})^{k_j}$. Candidate banks are independent across providers, conditional on the task. Taking the product within each task has expectation $\prod_j (1-\theta_{ij})^{k_j}$. Averaging over tasks yields the joint failure probability without assuming provider task difficulty independence.

Let

$$Y_i(k) = 1 - \prod_j \frac{\binom{n-c_{ij}}{k_j}}{\binom{n}{k_j}} \in [0, 1]. \tag{81}$$

With Q allocations and N independent task blocks, Hoeffding (1963) and a union bound yield

$$\mathbb{P}\left\{ \max_k |\widehat{U}(k) - U(k)| > \rho_N \right\} \le \delta, \qquad \rho_N = \sqrt{\frac{\log(2Q/\delta)}{2N}}. \tag{82}$$

Conditional on a fixed benchmark, this construction generalizes to independent candidate banks, with U interpreted as expected benchmark-average success. Confidence intervals at the task level also retain the pairing of services and allocations.

D.4 Revenue Bounds under Sampling Uncertainty

Assume $L(k) \le U(k) \le U^+(k)$ for all k. Let

$$\mathcal{P}(p) = \left\{ k : p \cdot k \le B,\ U^+(k) \ge \max_{\ell : p \cdot \ell \le B} L(\ell) \right\}. \tag{83}$$

Every true optimal allocation is in $\mathcal{P}(p)$: its upper success bound is at least the lower success bound of every feasible alternative. Minimizing provider profit over this set gives a lower bound on its current profit. Maximizing over the corresponding set after a deviation gives an upper bound on its deviating profit. Hence,

$$\mathrm{Gap}_{\mathrm{true}}(p) \le \frac{1}{B} \max_j \Big[\sup_{q>0} \max_{k \in \mathcal{P}(q, p_{-j})} (q - c_j) k_j - \min_{k \in \mathcal{P}(p)} (p_j - c_j) k_j \Big]. \tag{84}$$

The same affordability partition is used to evaluate this bound. Equation (82) gives a finite-sample construction of simultaneous intervals. Uniform control over choices is needed to construct sample-based bounds on profitable deviations (Balcan et al., 2019). Uncertainty in success estimates can change the buyer's preferred allocation; a small utility error alone does not imply a small bound on revenue deviations.

D.5 Derivative-Free Bounds and Curve Fitting

Concavity implies the following secant bounds for $r > 1$,

$$\frac{f(rx) - f(x)}{r - 1} \le H(x) \le \frac{f(x) - f(x/r)}{1 - 1/r}. \tag{85}$$

If all the failure probabilities are at least $q_{\min} > 0$, and have an absolute error of at most $\epsilon \le q_{\min}/2$, the negative log error is at most $2\epsilon/q_{\min}$. The errors in the lower and upper secant bounds are at most $4\epsilon/[q_{\min}(r-1)]$ and $4\epsilon/[q_{\min}(1-1/r)]$, respectively.

For a fitted mixture

$$f(x) = \sum_{r=-4}^{10} \left[a_r \frac{x}{2^r + x} + b_r \log(1 + x/2^r) \right], \qquad a_r, b_r \ge 0, \tag{86}$$

its peak satisfies

$$\sup_x H(x) \le \frac{1}{4} \sum_r a_r + \sum_r b_r, \qquad \lim_{x \to \infty} H(x) = \sum_r b_r. \tag{87}$$

The change of variable $x = t/(1-t)$ maps the entire domain to $t \in [0, 1]$, covering the limit at $t = 1$. Subdivision of intervals bounds the fitted function's global supremum. This bound applies to the fitted family. A finite range of observations is insufficient to determine the true tail of a scaling function.

D.6 Stability of Estimated Equilibrium Quantities

The local stability of an optimizer depends on its curvature and on its separation from other candidates (Bonnans and Shapiro, 2000). The next lemma addresses the stability of the peaks from Proposition 4, allowing for error in the estimated function and in its numerical maximization.

Lemma 6 (Perturbation of a strict maximum)**.** *Let H have a unique maximizer z and maximum M on $\mathcal{X} \subset (0, \infty)$. Suppose that $I \subset \mathcal{X}$ is an interval, $z \in \mathrm{int}(I)$, H is twice continuously differentiable on an open neighbourhood of I, and*

$$H''(x) \le -\lambda < 0 \quad (x \in I), \qquad \sup_{x \in \mathcal{X} \setminus I} H(x) \le M - \eta, \tag{88}$$

for $\lambda, \eta > 0$. Suppose that $\sup_{x \in \mathcal{X}} |\widehat{H}(x) - H(x)| \le \delta$, and let $\widehat{z} \in \mathcal{X}$ with

$$\widehat{H}(\widehat{z}) \ge \sup_{x \in \mathcal{X}} \widehat{H}(x) - \xi. \tag{89}$$

If $2\delta + \xi < \eta$, then $\widehat{z} \in I$ and

$$M - H(\widehat{z}) \le 2\delta + \xi, \qquad |\widehat{z} - z| \le \sqrt{\frac{2(2\delta + \xi)}{\lambda}}. \tag{90}$$

Proof. The approximation guarantees

$$H(z) \le \widehat{H}(z) + \delta \tag{91}$$
$$\le \widehat{H}(\widehat{z}) + \xi + \delta \tag{92}$$
$$\le H(\widehat{z}) + \xi + 2\delta. \tag{93}$$

An arbitrary point outside I loses at least η, and this contradicts $2\delta + \xi < \eta$. If the point is in I, then using Taylor's theorem and $H'(z) = 0$,

$$H(z) - H(\widehat{z}) \ge \frac{\lambda}{2}(\widehat{z} - z)^2. \tag{94}$$

Thus, the two inequalities give the result. □

If each provider satisfies the lemma on the full positive computation domain and $2\delta_j + \xi_j < M_j$, then we can construct prices from the true derivatives at the estimated quantities to get

$$\frac{R_j(q, p_{-j})}{R_j(p)} \le \frac{M_j}{H_j(\widehat{z}_j)} \le \frac{M_j}{M_j - 2\delta_j - \xi_j}. \tag{95}$$

This analysis works for quantities taken from an estimated curve and implemented under the underlying demand. Prices taken from estimated derivatives also alter implementation. The empirical best-response calculation in Section 4 evaluates prices directly against the joint success table.

The quantity sensitivity is governed by λ. The peak-to-tail gap in (43) can be small when close to the boundary $\alpha/\kappa = 1$, so the assumption $2\delta + \xi < \eta$ may require accurate tail estimates. Limiting the error in predicted failure probability is not sufficient to control the error in quantity in (90).

D.7 Integer Computation at Fixed Prices

Fix prices p and a computation increment $h > 0$. Let x^* solve the continuous buyer problem and define $\underline{x}_j = h\lfloor x_j^*/h \rfloor$. This is feasible because $\underline{x}_j \le x_j^*$. For increasing concave f_j with $f_j(0) = 0$,

$$f_j(u + t) - f_j(u) \le f_j(t) \le f_j(h) \qquad (u \ge 0,\ 0 \le t \le h). \tag{96}$$

The first inequality follows from the decreasing increments of a concave function. So the best allocation $x^{(h)}$ on the grid satisfies

$$0 \le \sum_j f_j(x_j^*) - \sum_j f_j(x_j^{(h)}) \le \Delta_h := \sum_j f_j(h). \tag{97}$$

Its success probability loss is upper bounded by

$$0 \le U(x^*) - U(x^{(h)}) \le 1 - e^{-\Delta_h} \le \Delta_h. \tag{98}$$

The middle inequality holds because capability loss is written as $d \le \Delta_h$, and $e^{-V_h}(1 - e^{-d}) \le 1 - e^{-\Delta_h}$. Continuity at zero implies $\Delta_h \to 0$ as $h \downarrow 0$, a purchasing guarantee at fixed prices. The integer pricing game is evaluated using its own best responses, as its demand switches at affordability boundaries.

D.8 Global Bounds for the Fitted Mixture

For positive scales and nonnegative coefficients, the fitted family has the form

$$f(x) = \sum_s \left[a_s \frac{x}{s + x} + b_s \log(1 + x/s) \right]. \tag{99}$$

Its marginal-capability curve is

$$H(x) = \sum_s \left[a_s \frac{sx}{(s + x)^2} + b_s \frac{x}{s + x} \right]. \tag{100}$$

In the marginal-capability curve, the contribution from each rational term of the fitted capability reaches $a_s/4$ at $x = s$; the contribution from each logarithmic term increases to b_s. Hence

$$\sup_x H(x) \le \frac{1}{4}\sum_s a_s + \sum_s b_s, \qquad \lim_{x\to\infty} H(x) = \sum_s b_s. \tag{101}$$

A sharper upper bound is obtained by subdividing the transformed domain $t = x/(1+x) \in [0,1]$. With $D_s(t) = s + (1-s)t$,

$$H\left(\frac{t}{1-t}\right) = \sum_s \left[\frac{a_s s\, t(1-t)}{D_s(t)^2} + \frac{b_s t}{D_s(t)}\right]. \tag{102}$$

The expression extends continuously to both endpoints, taking values zero and $\sum_s b_s$.

For any interval $I = [u, v] \subset [0,1]$ each rational term can be bounded by the largest of its endpoints or $a_s/4$ if $s/(1+s) \in I$; the logarithmic term can be bounded by its value at v. These bounds are summed to give an upper bound on the interval. Evaluating the interval endpoints and midpoint gives lower bounds on the global supremum. The intervals with the largest upper bounds are repeatedly subdivided to refine the bounds. Convergence is guaranteed by continuity on the compact transformed domain when all intervals whose upper bound remains competitive are subdivided.

When the fitted family is maximized at $t < 1$, then there is a finite quantity. When the supremum is approached only at $t = 1$, then Proposition 4 still implies finite approximate quantities. This establishes bounds for the fitted family. In order to extend these bounds to the underlying model, we need to establish estimation bounds for the same domain.

D.9 Cross-Model Dependence in Joint Success

Let $W_j(k_j)$ be task-conditional failure probability after k_j candidates. Joint failure is $\mathbb{E}[W_1(k_1)W_2(k_2)]$. Multiplying marginal failure probabilities gives $\mathbb{E}[W_1(k_1)]\mathbb{E}[W_2(k_2)]$. Their induced success probabilities relate as:

$$U_{\text{product}}(k) - U_{\text{joint}}(k) = \text{Cov}(W_1(k_1), W_2(k_2)). \tag{103}$$

If $W_1(k_1)$ and $W_2(k_2)$ are positively correlated in task difficulty, then the marginal-product success estimate is optimistic, consistent with the shared failures observed by Kim et al. (2025). Negative correlation makes the marginal-product success estimate pessimistic. The identity is allocation-specific, so it does not determine the sign of a change in an optimized market outcome. Demand can move to another allocation when the success surface changes.

Equation (17) is calculated by multiplying inside each task and then averaging. It is natural to multiply the two subset failure probabilities as the candidates can be generated independently conditional on the task. Task difficulties do not need to be independent between the two models. Similar reasoning applies if there are more than two providers: higher-order joint moments are obtained using the task-wise product.

The marginal success curves of checkpoints can improve without improving every joint allocation with a competitor. Checkpoints can solve different subsets of tasks. Joint comparisons examine

$$U_{\text{new}}(k_1, k_2) - U_{\text{old}}(k_1, k_2) \tag{104}$$

on the common allocation grid. This is also the important quantity for the fixed-price feasibility argument for correlated outputs across providers.

D.10 Confidence Bounds and Allocation Changes

Assume that all intervals $[L(k), U^+(k)]$ cover all possible population success values simultaneously. If a feasible allocation k^* satisfies

$$L(k^*) > \max_{k \neq k^*,\, p\cdot k \le B} U^+(k), \tag{105}$$

it is the unique population best response at p. If intervals overlap, multiple allocations may remain consistent with optimality. The plausible set in (83) considers all such allocations.

For a domain of N independent task blocks and Q allocation vectors, the Hoeffding sample-size construction is sufficient to achieve uniform half-width of at most ϵ with

$$N \geq \frac{\log(2Q/\delta)}{2\epsilon^2}. \tag{106}$$

For m providers and at most K calls per provider, $Q = (K+1)^m$. Candidate repetitions improve within-task success probability estimates. Tasks still represent the independent blocks for across-task inference.

The best-response profit bound uses the minimum current profit and the maximum deviating profit over their respective sets of plausible allocations. This constraint ensures an upper bound. Replacing either set by the empirically preferred allocation can remove a population-optimal choice. The upper bound may be loose when nearly-tied allocations have different payments made to a provider, but nearly equal success probabilities.

At a fixed competitor price vector, the feasible and plausible sets remain constant between affordability boundaries. The lower and upper provider profit envelopes are piecewise affine. The boundary sweep method used to determine the empirical gap also considers these envelopes. The intervals of all values must be computed simultaneously on the allocation grid, using the same coverage event for each price.

E Identifying Equilibrium Existence

E.1 Finite-Observation Lower Bound

Consider the market (8) with all $\alpha_j > \kappa_j$, which has a unique finite equilibrium. Now add a constant $\eta > 0$ to the latent rate of one of the providers. The failure and capability are then

$$\widetilde{F}_j(x) = e^{-\eta x} F_j(x), \qquad \widetilde{f}_j(x) = f_j(x) + \eta x. \tag{107}$$

This modified market is a nonnegative latent rate model, with limiting marginal capability $\eta > 0$.

The market cannot have an equilibrium with a finite positive price. In any equilibrium candidate, all providers must have positive revenue, by the entry argument in Appendix A. A perturbed provider can play a small but positive price q. This has positive demand, and the optimality condition implies $\mu \geq \eta/q$. If q is small enough, this is strictly larger than the finite value $f_h'(0+)/p_h$ of each other provider. All other providers have zero demand, so the deviating provider gets the full budget and hence improves over the candidate revenue.

Now observe candidate-success events or first-success times on N tasks, at most T computation per task. Couple the two experiments by adding to the perturbed provider an independent rate-η success clock, without changing the original process. The observations differ only when the added clock fires during the observed computation. Their total variation distance satisfies

$$\mathrm{TV}(P_0, P_\eta) \leq N(1 - e^{-\eta T}) \leq N\eta T. \tag{108}$$

Their two-point testing inequality (Tsybakov, 2009) implies a lower bound on the worst-case error,

$$\frac{1 - \mathrm{TV}(P_0, P_\eta)}{2} \geq \frac{1 - N\eta T}{2}. \tag{109}$$

Fixing N, T, one can choose η so close to zero that this lower bound is arbitrarily close to $1/2$. The observation model encompasses success events and first-success times. The lower bound concerns equilibrium existence over unbounded computation. For finite sets of allocations, a direct best-response test is given in Appendix D.

E.2 Price-Commitment Guarantee

At any private-provider price vector, spending B on a service with price $c_0 > 0$ and failure curve F_0 is feasible. The buyer's best success satisfies

$$U_{\text{achieved}}(B) \geq 1 - F_0(B/c_0). \tag{110}$$

Substituting B/c_0 into the bound $F_0(x) \le C(1 + x/\tau_0)^{-\kappa_0}$ with $\kappa_0 > 0$ limits the buyer's failure probability to $C(1 + B/(c_0\tau_0))^{-\kappa_0}$. The bound holds for arbitrary private-provider prices and arbitrary dependence between the backup model and the private model.

A complete retained portfolio will retain its original feasible allocation too. Let a_0 denote the allocation (if the utilities are estimated). The buyer retains a_0 unless there is some new allocation a with $L(a) \ge U^+(a_0)$. On the simultaneous coverage event we have

$$U(a) \ge L(a) \ge U^+(a_0) \ge U(a_0).$$

If there is no such new allocation, the buyer retains a_0. In general the retention of one model from a portfolio does not imply that the old complete allocation is feasible.

F Experimental Details

F.1 Model Configurations and Data Splits

The five service configurations are the original Qwen3-4B with thinking disabled, Qwen3-4B-Instruct-2507, the original Qwen3-4B with thinking enabled, Qwen3-4B-Thinking-2507, and Llama-3.1-8B-Instruct. The first two form the non-thinking replacement, and the next two form the thinking replacement. Llama is the competitor in both markets. The original and new services used their native chat templates, and the original Qwen checkpoint received the corresponding thinking flag for the service in which it was used (Qwen Team, 2025a;b;c; Yang et al., 2025; Grattafiori et al., 2024). No model was fine-tuned on the evaluation tasks.

LiveCodeBench uses `release_v6`, which is filtered based on release dates from August 1, 2024 to April 30, 2025. Filtering is based on executable reference tests before candidate generation. We retain the full code-generation test suite. MATH-500 retains all original problems, answers, subjects, and difficulty levels. Calibration and test splits are stratified and equally split based on coding difficulty or mathematical subject and difficulty (seed 2026).

Table 3: Data and generation counts. The original Qwen model has thinking and non-thinking modes, and each mode is a service configuration.

Dataset	Problems	Calibration	Test
LiveCodeBench release_v6	454	227	227
MATH-500	500	250	250
Total	954	477	477

There are 128 candidates for each task and service, with a total of 610,560 generations. There were 1,238,617,984 tokens generated. All candidates that failed, were truncated, or had formatting issues are included in these numbers.

F.2 Prompts, Decoding, and Verification

The problem statement, examples, and interface to function or input/output are included in the prompt, along with the request "Return a correct Python 3 solution. Place the final solution in a single Python code block". The response is evaluated based on the last Python code block (if present), or the full response if no code block is present. The candidates run in isolated containers with network access disabled. Syntax errors, incorrect input or output interface, timeouts, and any failed tests are scored as failures.

Mathematical prompts include the original problem and ask for a final answer in `\boxed{}`. Answers are normalized and compared using string equality or symbolic equivalence, with a five-second symbolic-check timeout. Unparseable outputs are counted as failures. Reference solutions and correctness labels are never given to generation prompts.

We use 0.7 temperature, nucleus sampling with top-p 0.95 (Holtzman et al., 2020), no top-k filtering, and repetition penalty 1.0. Non-thinking services have an output cap of 8,192 tokens and thinking

services have a 16,384 token output cap. For matched checkpoints, we use the same output cap. We use BF16 inference, up to 32 samples in a batch, and do not cache prefixes between requests. We generate candidate seeds deterministically by hashing task id, service id, and sample index. We track candidate generation time, token counts, and verifier time.

The primary metric reports whether the purchased candidates have a reference-verifiable solution. Code ranking can be affected by test coverage (Liu et al., 2023), so we separately evaluate candidate availability and public-test selection. In the coding selection experiment, the public test examples choose the answer and the hidden tests score it. We visit the candidates in a fixed random order. We return the first candidate that passes all the public test examples, or we fail if none does. Cross-provider candidates are interleaved round-robin, with the initial ordering of the providers specified by the task ID. We recompute the success surface and prices using the new selector.

F.3 Curve Fitting and Price Search

For choosing the regularization we use five-fold cross validation of the calibration set in (86). The coefficients are determined such that the squared error of the negative log-failure is minimized, including an ℓ_2 penalty with penalty factor set to one of $\{10^{-6}, 10^{-4}, 10^{-2}, 1\}$. The three-parameter curve from (8) is fitted to the same calibration data. The adjustment of $(N\widehat{F} + 1/2)/(N + 1)$ is done only to get finite logarithmic values when fitting. The success estimates are left as they are.

In price search, we apply Algorithm 1 and continuous price best responses (Algorithm 2). As our primary allocation set, we use $\{0, \dots, 64\}^2$; in our call-limiting test we use 32 or 128 instead of 64. For the starting price profiles, we use the fitted curve solution and seven equal price profiles of one to 64 affordable calls, as well as 16 deterministic Sobol points in log-price space (Joe and Kuo, 2008). Each starting profile is explored for at most 500 price sweeps with geometric damping of one half. For each price profile visited, we check for profitable deviations at continuous prices. We choose the price profile with the lowest normalized gap. If there are ties, we break ties by prioritizing higher calibration success and then by lexicographic price order. For all selection decisions, we use calibration data. We never use test outcomes to select a price profile.

Numerical acceptance requires $\mathrm{Gap}(p) \le 0.01$ in the empirical game with best-response profit computed with tolerance $10^{-8}B$. To compare the equilibria of old and new price profiles we require acceptance of both price profiles. The minimum gap for all instances including those that have failed acceptance is reported for coverage. Variation between different accepted initializations measures the sensitivity of the price selection computation.

Algorithm 1 Price construction and best-response search

Require: $\widehat{f}_j, \widehat{U}, \mathcal{K}, B, c$, initial prices $\mathcal{P}_0$, sweeps T, $\eta \in (0, 1]$, $\varepsilon_0 > 0$

$\widehat{H}_j(x) \leftarrow x\widehat{f}_j'(x)$, choose $z_j > 0$ with $\widehat{H}_j(z_j) \ge \sup_x \widehat{H}_j(x)/(1 + \varepsilon_0)$

$p_j^{\mathrm{curve}} \leftarrow B\widehat{f}_j'(z_j)/\sum_h \widehat{H}_h(z_h)$, $\mathcal{P}_0 \leftarrow \mathcal{P}_0 \cup \{p^{\mathrm{curve}}\}$

$\mathcal{V} \leftarrow \mathcal{P}_0$

for $p^{(0)} \in \mathcal{P}_0$ **do**

  $p \leftarrow p^{(0)}$

  for $(t, j) \in \{1, \dots, T\} \times \{1, \dots, m\}$ in lexicographic order **do**

    $(v_j, q_j) \leftarrow \mathrm{BR}_j(p_{-j}, \widehat{U}, \mathcal{K}, B, c_j)$ ▷ Algorithm 2

    $p_j \leftarrow p_j^{1-\eta} q_j^{\eta}$, $\mathcal{V} \leftarrow \mathcal{V} \cup \{p\}$

  end for

end for

$p^* \leftarrow \arg\min_{p \in \mathcal{V}} \mathrm{Gap}(p)$ under the fixed tie rule

$g_j \leftarrow [\mathrm{BR}_j^{\mathrm{value}}(p^*_{-j}) - (p_j^* - c_j)\widehat{k}_j(p^*)]/B$

return $p^*, \widehat{k}(p^*), \widehat{U}(\widehat{k}(p^*)), (g_j)_{j=1}^m$

F.4 Costs, Budgets, and Comparisons

The baseline budget is $B_0 = 1$. With a_j denoting measured mean generation time per candidate, and a_{ref} the measured value of old non-thinking Qwen, the serving costs are

$$c_j(\gamma) = \gamma \frac{a_j}{a_{\text{ref}}}, \qquad \gamma \in \{0, 10^{-4}, 10^{-3}, 10^{-2}\}. \tag{111}$$

Replacement costs are matched for old and new, picking the larger old/new value. The individual costs of each checkpoint are also compared. Budget multipliers are $B/B_0 \in \{1/4, 1/2, 1, 2, 4\}$ with costs held fixed in B_0 units. The verification cost comparison adds normalized CPU cost $\eta v_j / v_{\text{ref}}$, $\eta \in \{0, 10^{-4}, 10^{-3}\}$.

The three conditions are the old services at old prices, the upgraded service at old prices, and the upgraded service at recomputed prices. The buyer reoptimizes in each condition. A fixed-allocation comparison replaces the service but retains the old call allocation. Fixed-price references include equal unit quotes, marginal-cost quotes for positive costs, and twice-cost quotes. Equal spending, pass@1 routing, and marginal-success-per-price greedy allocation are compared with the optimal empirical allocation at the same quotes.

The equal-spending rule spends the budget equally between providers, buying the (affordable) integer quantity from each provider, and spending the remainder of the budget on calls with their marginal success per unit price. The pass@1 router spends on the provider with the highest pass@1 calibration probability. The greedy choice repeatedly purchases the (feasible) call with the highest marginal success-to-price ratio. None of these choices uses test-task outcomes in their decisions.

F.5 Unspent Funds and Early Stopping

The money-valued buyer maximizes

$$v\widehat{U}(k) - p \cdot k \quad \text{subject to } p \cdot k \le B, \qquad v/B \in \{0.5, 1, 2, 4, 8\}. \tag{112}$$

Its price-deviation test includes the ranking events in (79). The buyer may have unspent funds or no calls selected.

Our early stopping benchmark follows model cascade benchmarks Chen et al. (2023); Dekoninck et al. (2025). We compare two-stage schedules $1^{k_1}2^{k_2}$ and their reverse orders, as well as alternating schedules with the same call limits. All schedules terminate after successful call verification, and are feasible only if their worst-case payments are at most B. For a schedule $s_1, \ldots, s_L$,

$$\mathbb{E}[N_j] = \sum_{t=1}^{L} \mathbf{1}\{s_t = j\} \mathbb{P}(\text{all earlier calls fail}). \tag{113}$$

We estimate this expectation using failures of task-aligned prefixes. A provider's profit is $(p_j - c_j)\mathbb{E}[N_j]$. Within an interval where the feasible policy set does not change, the policy chosen by the system and its corresponding expected call counts are fixed. Hence, profit is affine in the provider's price on each interval. We recompute prices given their corresponding expected receipts. Our comparison is restricted to this finite set of policies.

F.6 Statistical Analysis

Primary confidence intervals were based on 2,000 task-level paired bootstrap replicates (Efron, 1979). They preserved all the providers and candidate counts in each task. Prices and the allocation chosen were held fixed from the calibration stage. A secondary analysis was also performed in which calibration fits and price searches were rerun. Point estimates, paired intervals, empirical deviation gaps, and bounds for profitable deviations were reported.

The measured complete reversal has a positive lower confidence bound for the fixed price gain, a negative upper confidence bound for the net change, and accepted old and new price profiles. The joint capability over the measured allocation set is compared with a 0.5 percentage point noninferiority margin. The condition is an empirical one for the finite range success probability.

The full legacy intervention keeps the old services and prices. The new allocation is selected if its simultaneous lower bound on the success probability is better than the upper bound of the legacy

Table 4: Services used in the paired comparisons. N and T denote the two replacements. The two pairs have the same competitor. The output count limit includes all tokens output, including reasoning tokens.

Checkpoint	Mode	Role	Token limit
Qwen3-4B	Non-thinking	Old N	8,192
Qwen3-4B-Instruct-2507	Non-thinking	New N	8,192
Qwen3-4B	Thinking	Old T	16,384
Qwen3-4B-Thinking-2507	Thinking	New T	16,384
Llama-3.1-8B-Instruct	Instruction	Competitor	8,192

allocation, otherwise the legacy allocation is kept. In the partial legacy intervention, only the old Qwen is kept. The backup price multipliers are 0.5, 1 and 2. Retention is the ability to reach the old success level within the 0.5-point noninferiority margin.

F.7 Service Configuration and Candidate Reuse

The five configurations used in both domains are listed in Table 4. The thinking and non-thinking settings have identical competitor banks. For each task and configuration, we create a bank of 128 candidates under fixed decoding conditions. This same bank can be used at different prices and call allocations. Changing the purchasing decision does not require new generations.

The candidate bank contains 128 independently sampled outputs for each task and service. The subset estimator of Eq. (17) is used for purchased quantities up to 128. These subsets may overlap in the candidate bank, resulting in dependent estimates used to evaluate different allocations. Task-level resampling preserves this dependency. The same bank is reused across different monetary budgets. The monetary budget changes which allocations are available, whereas the estimate of the capability surface is fixed.

Unsuccessful and truncated generations are included in the physical resource accounting. Generated length variation is accounted for in the measured mean cost of each purchased candidate, generated under the configuration in Table 4. In the matched-cost comparison, a common cost is selected for the old and new checkpoint, taking the larger of the two measured costs. In the checkpoint-specific comparison, each measured mean is used. In either case, a candidate does not equate to the same physical computation in all five configurations.

F.8 Paired Outcomes on Held-Out Tasks

Let $\widehat{U}_o^{\mathrm{cal}}$ and $\widehat{U}_n^{\mathrm{cal}}$ denote the old and new joint success tables on the calibration tasks. The price search returns prices p_o and p_n. Applying the fixed tie rule, we have

$$k_A = \operatorname*{arg\,max}_{k \in \mathcal{K},\ p_o \cdot k \le B} \widehat{U}_o^{\mathrm{cal}}(k), \tag{114}$$

$$k_B = \operatorname*{arg\,max}_{k \in \mathcal{K},\ p_o \cdot k \le B} \widehat{U}_n^{\mathrm{cal}}(k), \tag{115}$$

$$k_C = \operatorname*{arg\,max}_{k \in \mathcal{K},\ p_n \cdot k \le B} \widehat{U}_n^{\mathrm{cal}}(k). \tag{116}$$

For a held-out task i, let $Y_i^o(k)$ and $Y_i^n(k)$ denote the subset-based success estimates for the corresponding checkpoint. Let

$$A_i = Y_i^o(k_A), \qquad B_i = Y_i^n(k_B), \qquad C_i = Y_i^n(k_C). \tag{117}$$

The reported rates are averages of A_i, B_i, C_i. The three effects are averages of $B_i - A_i$, $C_i - B_i$, and $C_i - A_i$. The decomposition holds prior to averaging, and is preserved under every paired bootstrap resample.

The fixed allocation comparison $Y_i^n(k_A)$ corresponds to replacement at the originally purchased allocation. The move from k_A to k_B corresponds to the change in allocation due to buyer reoptimization at unaltered prices. The move from k_B to k_C corresponds to the allocation that was chosen after

Algorithm 2 Best response over affordability thresholds

Require: $j, p_{-j}, B, c_j, \widehat{U}$, finite $\mathcal{K} \subset \mathbb{Z}_+^m$ with $0 \in \mathcal{K}$
$\mathcal{K}_0 \leftarrow \{k \in \mathcal{K} : k_j = 0,\ \sum_{h \neq j} p_h k_h \leq B\}$
$t_k \leftarrow (B - \sum_{h \neq j} p_h k_h)/k_j$ for each $k \in \mathcal{K}$ with $k_j > 0$
$\{b_1 < \cdots < b_L\} \leftarrow \mathrm{sort}\big(\{B\} \cup \{t_k : 0 < t_k \leq B\}\big)$
$\mathcal{H} \leftarrow \mathcal{K}_0 \cup \{k \in \mathcal{K} : k_j > 0,\ t_k > 0\}$, ranked by buyer preference
$(v, q) \leftarrow (0, 2B)$ ▷ Exit yields zero sales
for $r = 1, \ldots, L$ **do**
  $k \leftarrow \mathrm{top}(\mathcal{H})$, $w \leftarrow (b_r - c_j)k_j$
  if $w > v$ **then**
    $(v, q) \leftarrow (w, b_r)$
  end if
  $\mathcal{H} \leftarrow \mathcal{H} \setminus \{k \in \mathcal{K} : k_j > 0,\ t_k = b_r\}$
end for
return $v = \max_{u>0}(u - c_j)\widehat{k}_j(u, p_{-j})$ and a maximizing price q

the change in prices. All allocations were chosen before the test and are not influenced by the test outcome.

Bootstrap resamples are obtained by resampling complete task indices with replacement, computing the three means and their differences, and using the 2.5th and 97.5th percentiles to calculate paired 95 percent confidence intervals. Domain averages are computed by assigning equal weight to the two replacement pairs. Retaining both pairs in each task resample preserves their dependence because they share a competitor bank.

F.9 Event-Sweep Implementation

Algorithm 2 is the implementation of the best-response oracle used in Algorithm 1, and in the deviation tests. For the hard-budget buyer, allocations can be ordered first by empirical success, then by fewer total calls, then lexicographically, as this ordering does not depend on the deviating price.

Allocations with $k_j = 0$ are still feasible if the cost of the competitors is in the budget. The all-zero allocation is always present. A priority queue is used to maintain the preference ordering of allocations when they are removed. There is at most one affordability event created per allocation for the deviating provider, and the complexity in Section 4 follows.

At each price threshold, the allocations leaving the algorithm remain affordable and remain part of the buyer's choice set until the threshold is considered. Removing them before checking the price threshold may lose a best response. Algorithm 2 removes allocations after evaluating the price threshold, so the heap contains the preferred allocation throughout the next interval. The price-independent ranking is fixed so the algorithm computes an attained maximum without one-sided approximations.

For a money-valued buyer the allocation utilities are affine in the deviating price inside their affordability intervals. Pairwise crossings in (79) result in additional ranking events where the selected allocation can differ from its adjacent limiting choices. Their evaluation and both one-sided limits determine the supremum, and an interior price whose profit is within $10^{-8}B$ of the limiting profit gives a numerical witness. The potentially quadratic ranking events above can be outside the hard-budget complexity bound. Early-stopping policies have fixed success and expected usage between feasibility events, so their profit maximization relies on the same affordability argument.

F.10 Expected Payments with Early Stopping

For a fixed ordered schedule of providers $s_1, \ldots, s_L$, let $k_h(t-1)$ be the number of calls to provider h made before position t. With task-specific candidate success probabilities, the probability of seeing the t-th call is

$$Q_i(t) = \prod_h (1 - \theta_{ih})^{k_h(t-1)}. \tag{118}$$

Task-aligned subset estimates replace each of these power terms when constructing the empirical table. The expected provider usage and total payment are

$$\mathbb{E}[N_j] = \frac{1}{N} \sum_{i=1}^{N} \sum_{t=1}^{L} \mathbf{1}\{s_t = j\} Q_i(t), \tag{119}$$

$$\mathbb{E}[\text{payment}] = \sum_{j} p_j \mathbb{E}[N_j]. \tag{120}$$

The probability of success is one minus the probability that every scheduled call fails. Using conditional independent sampling, the probability of success is the same regardless of the order of the same multiset of providers. However, the expected usage depends on the order of calls. For instance, one can try a more promising provider earlier in the schedule and avoid subsequent purchases on tasks it succeeds at.

The feasibility constraint $\sum_{t=1}^{L} p_{s_t} \leq B$ enforces that every realized payment is within budget. The provider's expected profit is $(p_j - c_j)\mathbb{E}[N_j]$. The providers reprice using these expected receipts and corresponding expected serving costs. The policy class evaluated includes two provider block orders and alternating schedules at the same call limits.

F.11 Reporting Price-Profile Coverage

A quoted profile is accepted if its largest normalized empirical deviation gain is at most 0.01. A replacement pair is accepted if both its old and new profiles are. The sensitivity analysis also reports the fraction of individual profiles that pass the threshold, as a given setting may only satisfy it on one side of a replacement.

Profiles are selected by minimum gap first, second by greater calibration success, and finally by lexicographic price. Test success is only measured on the selected profiles. The range across the accepted initializations of these strategies is a measure of how much the results can vary depending on the price profiles this method creates.

A complete measured reversal requires a positive lower confidence bound on fixed-price gain and a negative upper confidence bound on net change, with both prices accepted. A negative repricing component alone is attenuation. Joint noninferiority over the finite allocation range is a separate test with a 0.5 point margin and task-paired intervals.

G Additional Empirical Results

G.1 Curve Prediction and Deviation Bounds

Table 5 summarizes fit quality, finite-range peak locations, and price-deviation diagnostics. The flexible mixture fits the failure curves with lower error in both domains. The statistical deviation bounds take into account uncertainty in the buyer's allocation, and are therefore larger than the empirical gaps. All four primary old/new price-profile pairs pass the numerical deviation test. Uncertainty increases the bounds by a median of 0.030 normalized budget units.

G.2 Domain Averages and Allocation Baselines

Table 6 reports equally-weighted averages over the checkpoint replacements. The non-thinking net effect is negative and the thinking net effect is positive on LiveCodeBench. The mean's confidence interval includes zero. The mean interval is below zero on MATH-500. The repricing effect is negative in all four primary comparisons. Across the various settings, the upgraded provider's geometric mean price ratio is 1.32, its purchased call ratio is 0.83, and its mean revenue-share increase is 3.7 percentage points.

The net effect range between accepted price profile selections is $[-2.2, 1.3]$ percentage points. This summarizes the net effects of calculated profile selections. The main selection rule is fixed on calibration data. Table 7 shows the loss due to alternative buying rules at constant prices.

Table 5: Curve and price diagnostics. Percentage points are used for failure curve errors and joint surface errors. Log quantity error and normalized deviation gaps are dimensionless.

Quantity	LiveCodeBench	MATH-500
Three-parameter failure-curve error	2.9	1.4
Mixture failure-curve error	1.7	0.8
Peaks inside observed range (%)	71.4	64.3
Median absolute log error in calls	0.27	0.18
Independence-projection joint error	3.6	1.9
Minimum joint-improvement lower bound	−0.3	−0.1
Accepted old/new profile pairs	2/2	2/2
Median held-out empirical deviation gap	0.006	0.004
Median statistical deviation upper bound	0.041	0.028

Table 6: Mean effects across the two checkpoint replacements, in percentage points. The intervals are task paired 95 % intervals for the net effect.

Domain	Technical gain	Repricing	Net change	Net interval
LiveCodeBench	+3.0	−3.5	−0.5	$[-1.6, 0.7]$
MATH-500	+1.5	−2.2	−0.7	$[-1.3, -0.1]$

G.3 Call Limits, Costs, and Unspent Funds

The net effects of the 32-call limit are 0.2 and -0.5 percentage points on coding and math. Increasing the limit to 128 calls yields effects of -0.8 in both cases. The fraction of selected allocations hitting the cap is 12.5% at 64 calls, but drops to 3.1% at 128 calls. The matched cost results for $\gamma = 10^{-3}$ are -0.4 and -0.6 points. Using measured costs for specific checkpoints elevates them by 0.2 and 0.1 points.

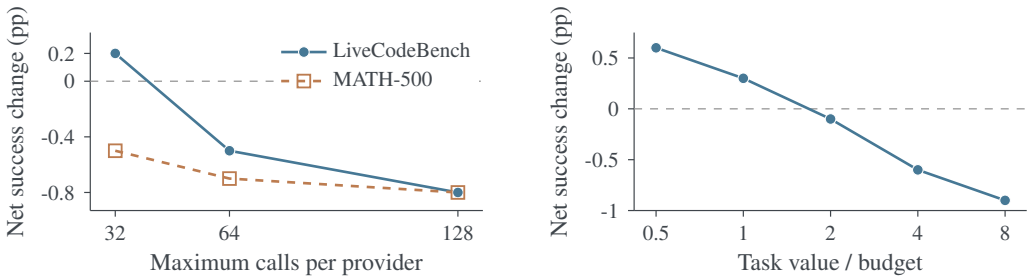

Figure 3: **Purchasing constraints change the effect of repricing.** Left: Net changes at different integer call limits. Right: Net changes for a buyer that values retained funds. A higher v/B implies that the buyer places more weight on task success relative to money.

Success-triggered stopping changes the net effect by $+0.4$ and $+0.2$ points in comparison to the primary setting. Public-test selection changes the coding effect by -0.6 points. Numerical acceptance rates for these settings are 87.5% and 75.0%, respectively. The corresponding rates for the other changes to the buyer model are given in Table 2.

G.4 Price Commitment

We see a difference of $+0.1$ points with interval $[-0.2, 0.4]$ by retaining the complete old service portfolio compared to that portfolio. Retaining only the old Qwen shows a difference of -0.8 points. In table 9 we see how often a fixed backup retains the old outcome within the noninferiority margin. Affordability affects the strength of the retained option.

Table 7: Purchasing rules compared at the same quoted prices. Entries show the change in success rates compared with optimal allocation (percentage points).

Purchasing rule	Success change
Equal spending	-2.7
Pass@1 routing	-5.8
Marginal-success-per-price greedy	-0.9

Table 8: The money-valued outside option. Net changes are percentage points.

Task value / budget, v/B	0.5	1	2	4	8
Net success change	$+0.6$	$+0.3$	-0.1	-0.6	-0.9

H Gaussian Prediction

For the information allocation problem, the pricing characterization is also valid. Let

$$Y_{jr} = \theta + G + Z_j + \varepsilon_{jr}, \tag{121}$$

with zero-mean Gaussian components having variances g, $a_j > 0$, $b_j > 0$. We assume the component G is common across all providers, the component Z_j is common to all calls to provider j, and the other components are independent of each other. The average of x_j calls has private variance $a_j + b_j/x_j$. For integer x_j, this is sample averaging, extended to a continuous function of its precision.

The Gaussian precision aggregation method (Cover and Thomas, 2006) is an unbiased linear combination with weights summing to 1 and mean squared error

$$g + \sum_j w_j^2 (a_j + b_j/x_j).$$

The minimum mean squared error is obtained for weights proportional to private precision, and the minimum error is given by

$$\mathcal{E}(x) = g + \frac{1}{\sum_j f_j(x_j)}, \qquad f_j(x) = \frac{x}{a_j x + b_j}. \tag{122}$$

Optimizing the sum of private precision aligns with the buyer's problem. Here

$$H_j(x) = \frac{b_j x}{(a_j x + b_j)^2}, \qquad z_j = \frac{b_j}{a_j}, \qquad M_j = \frac{1}{4a_j}. \tag{123}$$

Then, the equilibrium prices, revenues, and prediction error are given by

$$p_j^* = \frac{B}{b_j S}, \qquad R_j^* = \frac{B}{a_j S}, \qquad \mathcal{E}^* = g + \frac{2}{S}, \tag{124}$$

where $S = \sum_j 1/a_j$. In the case of unlimited computation, the error is $g + 1/S$. Thus, the buyer gets half of the private precision. This ratio is related to the private precision, while the common component of the error, g, remains unchanged.

For identical providers of parameters a, b and serving cost $c > 0$, define

$$d = \frac{aB}{bm}, \qquad K = \frac{m+1}{m-1}. \tag{125}$$

The unique symmetric equilibrium price satisfies

$$P = \frac{c + d + \sqrt{(c+d)^2 + 4Kcd}}{2}. \tag{126}$$

Consider differentiating the buyer's demands at a symmetric price. Each provider's expenditure derivative is

$$\left.\frac{\partial(p_j x_j)}{\partial p_j}\right|_{p_j = P} = \frac{m-1}{2am}\left(\frac{aB}{mP} - b\right). \tag{127}$$

Table 9: Retention of the old success level with a fixed-price backup.

Backup price multiplier	0.5	1	2
Configurations retaining old success (%)	93.8	78.1	46.9

The first-order condition for $(1 - c/p_j)p_j x_j$ is $P^2 - (c + d)P - Kcd = 0$, so the price is (126).

This is a global best response. Prices below cost cannot improve the deviator's profit in the positive profit equilibrium. For $p \geq c$, the requirement to exclude all competitors is

$$\sqrt{pP} - p \geq aB/b.$$

However, the root equation gives $P - c = d(1 + Kc/P) < 2aB/[b(m - 1)]$, and

$$\sqrt{pP} - p \leq (P - c)/2 < aB/[b(m - 1)] \leq aB/b,$$

a contradiction. On the interval of positive demand and profit, set $t = \sqrt{p}$, $L = (m - 1)\sqrt{P}$, $A = aB + b(m - 1)P$, and $C = bL$. The deviator's profit is

$$\pi(t) = \frac{(t^2 - c)(A - Ct)}{at(t + L)}. \tag{128}$$

The logarithm of its profit has second derivative

$$-\frac{2(t^2 + c)}{(t^2 - c)^2} - \frac{C^2}{(A - Ct)^2} + \frac{1}{t^2} + \frac{1}{(t + L)^2} < 0. \tag{129}$$

The stationary point is the unique maximum on this interval. The zero-profit boundary regions cannot improve it. This proves existence and uniqueness within symmetric equilibria.

The error in question satisfies

$$\mathcal{E}^*(B) = g + \frac{a}{m} + \frac{bP}{B} \longrightarrow g + \frac{2a}{m} \qquad (B \to \infty). \tag{130}$$

At marginal cost prices it is $g + a/m + bc/B$ and thus tends to $g + a/m$. Positive serving costs do not eliminate the additional equilibrium error in this specialization.