THE PRICE OF A BUDGET IN TIERED PROCUREMENT

Anonymous authors
Paper under double-blind review

ABSTRACT

A platform buys one contract per period from providers that each offer several service tiers. A tier's quality and its cost multiplier are public, the provider's base cost is private, and the platform pays at most a budget B. Without the budget the answer is classical: the welfare-optimal allocation is truthful under threshold payments. With it, the truthful payment for a high tier carries the option value of the tiers below the one awarded, so it can exceed B even when the awarded tier's own threshold does not, and fixing the tier before seeing the report gives up the choice of tier that welfare optimality needs. Clipping the payment at the budget keeps that choice, and what the budget then costs is the quality it cannot certify. Three results follow. First, with two tiers the best worst-case ratio $\rho^\star(q, \lambda, B)$ of any deterministic truthful budget-feasible mechanism is exact, a budget-clipped VCG menu attains it, and the ratio is 1 exactly when the budget stops binding on the two tiers, falling to zero as it approaches that threshold from below. Second, when each tier's cost is a separate private number, no truthful budget-feasible mechanism has a positive worst-case ratio as soon as B separates two tier qualities, while the same clipped menu still guarantees welfare at least the offline optimum minus the part of the best tier's quality the budget cannot cover; the public multiplier is what turns an additive guarantee into a multiplicative one. Third, when qualities must be learned, running the clipped menu on lower confidence bounds keeps every incentive property at every round and across rounds, never buys below true value, and has regret $\tilde{O}(\sqrt{T})$ against the offline clipped menu. Exact enumeration on 2955 discrete instances, 300 continuous instances and 600 online runs confirms the bounds, shows that the tie-broken threshold rule that looks optimal on a coarse grid drops to a median worst-case ratio of 0.12 on continuous instances, and that the learning layer's regret stays sublinear where the rule it replaces is linear. A budget costs a platform exactly the quality it cannot certify.

1 INTRODUCTION

A platform that buys capacity from external providers buys it in tiers. A language-model provider sells the same request at several service levels (Ong et al., 2025; Chen et al., 2024), a cloud provider sells the same job on several instance classes, and a logistics provider quotes several delivery speeds. In each case a tier's quality is a public number the platform can measure, and the tier's cost to the provider is a public multiple of a base cost the provider knows and the platform does not. The platform has a budget for the one contract it awards per period, with per-request routing and capacity constraints outside the model. The question is what a mechanism can guarantee when it must be truthful, must never pay a provider less than the cost of the tier it buys, and must never pay more than the budget.

Without a budget the answer is classical: the provider's type is one number, the welfare-optimal allocation is monotone in it, and the threshold payment of (Myerson, 1981) makes truthful reporting dominant (Nisan et al., 2007). Tiers change two things. The truthful payment for a high tier carries the value of the option to fall back to a lower tier, so it can exceed the budget even when the awarded tier's threshold is within it. The natural repair, fixing the tier before looking at the report and paying the multiplier times the threshold, gives up the choice of tier that welfare optimality needs. The gap between what a budget-feasible truthful mechanism can deliver and the offline optimum is the object of this paper.

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.

The mechanism we study keeps the menu and clips the prices. A menu is what a platform can run, and offering every tier at $\min\{q_t - W_{-i},\, B\}$, where W_{-i} is the best welfare the other providers can deliver, leaves each menu independent of the provider's own report, so truthfulness, individual rationality and the budget hold by construction. The clip can move the provider's choice to a cheaper tier than welfare wants, and Section 5 shows this mismatch is exactly what the budget costs.

Exact enumeration on 2955 discrete instances, 300 continuous instances and 600 online runs confirms the bounds. On the discrete family that motivated this work the clipped menu is never below the predicted floor and meets it exactly in 2474 instances. A threshold rule with quality tie-breaks attains ratio 1 on all $20{,}568$ profiles of the coarse grid it was tuned on, yet drops to a median worst-case ratio of 0.12 on continuous instances, where the clipped menu's median is 0.95. Online, the certified-price menu has zero rounds of negative welfare in 600 runs and sublinear regret, while the threshold rule it replaces has linear regret.

Contributions. With n providers, m tiers, public qualities $q_{i,t}$ and multipliers $\lambda_{i,t}$, private base costs θ_i and a budget B on the payment (Section 3):

The price of a budget is exact and has a cliff. With two tiers we give the best worst-case welfare ratio $\rho^\star(q, \lambda, B)$ of any deterministic truthful, individually rational, budget-feasible mechanism in closed form (Theorem 2); the clipped menu attains it, the ratio equals 1 exactly when the budget stops binding, and below that threshold it approaches 0. With n providers and two tiers the same menu is worst-case optimal with ratio $\min_i \rho^\star(q_i, \lambda_i, B)$ (Theorem 3); with three or more tiers we quantify the gap (Section 8).

The public multiplier separates a multiplicative guarantee from an additive one. If each tier's cost is a separate private number, no truthful budget-feasible mechanism has a positive worst-case ratio once the budget falls strictly between two tier qualities (Theorem 4). The clipped menu still guarantees welfare at least the offline optimum minus the part of the best contract's quality the budget cannot cover, and only the single-parameter structure converts this additive loss into the floor $\rho^\star$.

Learned qualities keep the guarantees if prices are certified. Running the clipped menu on lower confidence bounds is truthful, individually rational and budget-feasible at every round, when the bounds hold never awards a tier whose true quality is below its cost, and has regret $\tilde{O}(\sqrt{T})$ (Theorem 5). From awarded contracts the provider controls the sample: a plain lower bound invites withholding (Proposition 1), a running-maximum one makes accepting every affordable exploration contract optimal for a fixed type, and the regret is $T^{2/3}$ (Theorem 6).

Claims come in three strengths: Theorems 2, 3 and 4 hold deterministically on every profile; Theorem 5(b,c) and Theorem 6(iii) hold with probability at least $1 - 1/T$, when the confidence intervals cover the truth; the experiments are empirical, exact enumeration on the discrete families and sampled profiles on the continuous ones.

2 Related work

Budget-feasible and service-level mechanism design. Budget-feasible mechanism design buys a *set* under a hard budget on payments, with constant-factor approximations for additive and subadditive valuations (Singer, 2010; Liu et al., 2024; Anagnostides et al., 2026b); a monopolist blocks every such approximation in the multidimensional case (Neogi et al., 2025). The levels-of-service model, closest to us, has a budgeted buyer take a prefix of an agent's ladder (Amanatidis et al., 2025a; Chan & Chen, 2014). We let the budget bind, award one tier, and give an exact instance-dependent ratio with a matching lower bound; Theorem 4 is that obstruction under our net-welfare benchmark, and the public multiplier is what removes it.

Payment shape and multidimensional screening. Multi-outcome single-parameter payments carry an integral term (Myerson, 1981; Archer & Tardos, 2001), read off a step function in (1): our Theorem 1 is that characterization in the tiered notation, and the sum in (1) is where the option value of lower tiers enters. Multidimensional screening characterizes the menus a principal facing a vector-typed agent should offer (Rochet & Choné, 1998; Armstrong, 1996; Rochet & Stole, 2003), but maximizes profit with no budget. Capping the *agents* (Che & Gale, 1998; Dobzinski et al., 2012) concerns a private budget and Pareto optimality; for one good, the unique Pareto optimal, incentive

Table 1: Notation.

symbol	meaning
$n,\ m$	providers $i \in [n]$, tiers $t \in [m]$
$q_{i,t},\ \lambda_{i,t}$	public quality and multiplier of tier t
θ_i	private base cost of provider i
B	budget; no payment exceeds it
$(x)^+$	positive part, $\max\{x, 0\}$
W_{-i}	best welfare of the other providers
$\mathrm{OPT}(\theta)$	welfare optimum under the budget
$\rho_M(I)$	worst-case welfare ratio of M
$\rho^\star$	$\sup \rho_M$ over deterministic DSIC, IR, BF M

compatible and individually rational mechanism truncates valuations before VCG (Le, 2018). We fix the budget, with the lower bound of Theorem 2(a) chosen as in the prior-free template of (Hartline & Roughgarden, 2008); the clipped menu keeps the option value and clips the payment, parting from that truncation once a second tier is available.

Online procurement and bandit mechanisms. Deterministic truthful bandit mechanisms with click feedback need exploration that reports cannot influence (Babaioff et al., 2009); randomization and implicit payments relax that requirement (Babaioff et al., 2010), and the $T^{2/3}$ rate is the price of truthfulness there too (Devanur & Kakade, 2009). Budgeted online procurement uses posted prices (Badanidiyuru et al., 2012; Charalampopoulos et al., 2025; Amanatidis et al., 2025b). Our learning layer runs the clipped menu on lower confidence bounds built from probes no report can reach, and the widths are the Hoeffding half-width (Hoeffding, 1963) under the low-noise condition of (Audibert & Tsybakov, 2007). What separates our learning results is where the samples come from: with probes the platform controls the sample, and from awarded contracts the provider does. We therefore freeze a running-maximum lower bound, since a plain one invites withholding. Appendix D and Table 7 there continue the comparison.

3 PRELIMINARIES

The procurement environment. A platform awards one contract among n providers. Provider $i \in [n]$ can deliver any tier $t \in [m]$, and tier t has public quality $q_{i,t} \geq 0$ and public multiplier $\lambda_{i,t} > 0$; we index tiers so that $\lambda_{i,1} < \cdots < \lambda_{i,m}$. Provider i has a private base cost $\theta_i \geq 0$, and delivering tier t costs it $\lambda_{i,t}\theta_i$. The platform pays at most a budget $B > 0$. An instance is $I = (q, \lambda, B)$.

Mechanisms and incentive properties. A deterministic direct mechanism maps reports $r \in \mathbb{R}^n_{\geq 0}$ to an outcome $o(r) \in \{\varnothing\} \cup ([n] \times [m])$ and a payment $p(r) \geq 0$ to the selected provider. Provider i with type θ_i awarded tier t at price p gets utility $p - \lambda_{i,t}\theta_i$, and 0 otherwise. The mechanism is *dominant-strategy incentive compatible* (DSIC) if truthful reporting is optimal for every i, θ_i and r_{-i}; *ex-post individually rational* (IR) if a provider awarded tier t under truthful reports is paid at least $\lambda_{i,t}\theta_i$; and *budget-feasible* (BF) if $p(r) \leq B$ for every r.

Welfare benchmark. The welfare of an outcome (i, t) at profile θ is $q_{i,t} - \lambda_{i,t}\theta_i$, and of $\varnothing$ it is 0. The benchmark is the offline optimum: what a buyer who knew θ would choose, over every contract whose cost fits the budget,

$$\mathrm{OPT}(\theta) = \max\Big\{0,\ \max_{(i,t):\, \lambda_{i,t}\theta_i \leq B}\ q_{i,t} - \lambda_{i,t}\theta_i\Big\}.$$

Write $W_M(\theta)$ for the welfare of mechanism M under truthful reports. The *worst-case ratio* of M on I is $\rho_M(I) = \sup\{\rho \leq 1 :\ W_M(\theta) \geq \rho\,\mathrm{OPT}(\theta) \text{ for all } \theta\}$; a mechanism that buys at a loss on some profile where $\mathrm{OPT} = 0$ has $\rho_M(I) = -\infty$. The constraint in OPT falls on the cost of the tier, the constraint on the mechanism on the payment. Table 1 collects the notation.

4 THE SHAPE OF TRUTHFUL PAYMENTS

Truthful payments carry an option value. Fix i and r_{-i} and write $x_i(\theta_i) \in \{0\} \cup \{\lambda_{i,t}\}_t$ for the multiplier of the tier awarded to i when it reports θ_i, with $x_i = 0$ when i is not selected. The provider's cost is $\theta_i x_i$, so the setting is single-parameter and the characterization of (Myerson, 1981; Archer & Tardos, 2001) applies (Nisan et al., 2007, Ch. 9).

Theorem 1 (shape of truthful payments)**.** *A mechanism is* DSIC *and* IR *if and only if, for every i and r_{-i}, $x_i(\cdot)$ is non-increasing and a selected provider is paid $p_i(\theta_i) = \theta_i x_i(\theta_i) + \int_{\theta_i}^{\infty} x_i(s)\, ds$. Writing $\mu_1 > \cdots > \mu_K$ for the multipliers x_i takes on intervals and $b_1 < \cdots < b_K$ for the thresholds at which it steps down ($x_i = \mu_k$ on (b_{k-1}, b_k), $b_0 = 0$), the payment for tier with multiplier μ_k is*

$$p_i = \mu_k b_k + \sum_{j>k} \mu_j \,(b_j - b_{j-1}). \tag{1}$$

The first term is the awarded tier's multiplier times its own threshold. The sum is the value of the option to fall back to lower tiers as the report rises, and it is what separates a tiered contract from a single-tier one. Three consequences organize the rest of the paper.

Corollary 1. *Fix i and r_{-i}. Here $\varnothing$ has welfare 0, and ties are broken toward the larger multiplier and against $\varnothing$.* (a) *If the awarded tier does not depend on θ_i ($K = 1$), the truthful payment is $\lambda_{i,t}\, b_1$: the multiplier times the threshold in base units.* (b) *Paying the base-unit threshold b_1 itself violates* IR *on $\theta_i \in (b_1/\lambda_{i,t}, b_1]$ whenever $\lambda_{i,t} > 1$, and every such provider has a profitable misreport (any $r_i > b_1$).* (c) *The welfare-optimal feasible allocation, which awards $\arg\max_{o \in \{\varnothing\} \cup \{t:\lambda_{i,t}\theta_i \le B\}} \mathrm{welfare}(o)$ to the provider with the largest such value, has non-increasing x_i and is therefore implementable; its payment* (1) *can exceed B even when $\mu_k b_k \le B$. Example 1 instantiates this.*

Part (a) is the payment rule of the threshold mechanisms we compare against: fix the tier from public data, compute the base-unit threshold τ from the other reports, pay $\min\{\lambda_t \tau, B\}$ and award only if the report is at most that price over λ_t. Part (b) is why paying τ itself fails, and part (c) is the tension the next section resolves: the allocation welfare wants is truthful but not budget-feasible, and the allocation that is trivially budget-feasible is not the one welfare wants.

5 THE PRICE OF A BUDGET

By the taxation principle (Nisan et al., 2007, Ch. 9), a deterministic DSIC mechanism presents each provider i, given the other reports, with a menu of prices $(p_{i,t})_{t \in [m]}$. The provider takes a utility-maximizing tier or nothing, so IR is automatic and BF says $p_{i,t} \le B$. Every mechanism in this paper is stated as a menu.

Definition 1 (budget-clipped VCG menu, BC-VCG)**.** Given reports r, let $W_{-i}(r) = \max\{0, \max_{j \ne i,\, t:\, \lambda_{j,t} r_j \le B} (q_{j,t} - \lambda_{j,t} r_j)\}$ be the best feasible welfare the other providers offer. Provider i faces the menu

$$p_{i,t}(r_{-i}) = \min\{\, q_{i,t} - W_{-i}(r),\ B \,\}, \qquad t \in [m],$$

takes $t_i \in \arg\max_t (p_{i,t} - \lambda_{i,t} r_i)$, ties broken toward the larger $q_{i,t} - \lambda_{i,t} r_i$, and has demand $u_i = \max_t (p_{i,t} - \lambda_{i,t} r_i)$. The contract goes to the smallest-index provider with $u_i \ge 0$ at price p_{i,t_i}; if there is none, nothing is bought.

The prices have the VCG externality form $q_{i,t} - W_{-i}$ (Vickrey, 1961; Clarke, 1971; Groves, 1973); without the clip, and with W_{-i} taken over all tiers, the mechanism is VCG, and the provider's choice coincides with the welfare-optimal tier because all prices are shifted by the same constant. The clip is what the budget forces, and it is the only place where the provider's choice and the welfare-optimal choice can part. Clipping at the budget also restores feasibility in the budget-constrained auction literature (Le, 2018), where the clip falls on valuations before VCG runs; the two coincide for one provider and part under competition (E1, Section 8).

5.1 TWO TIERS: THE EXACT RATIO

For one provider with tiers $(q_1, \lambda_1), (q_2, \lambda_2)$, $q_1 < q_2$, define $a^\star = (q_2 - q_1)/(\lambda_2 - \lambda_1)$, the type at which the two tiers' welfares cross (for $B < q_2$ the budget-feasible optimum leaves tier 2 at

$\min\{a^\star, B/\lambda_2\}$), and $a = (B - q_1)^+/(\lambda_2 - \lambda_1)$, where the menu $(\min\{q_1, B\}, B)$ switches. Let $r(\theta) = (q_1 - \lambda_1\theta)/(q_2 - \lambda_2\theta)$ be the welfare ratio of buying tier 1 where tier 2 is optimal. Define

$$\rho^\star(q, \lambda, B) = \begin{cases} 1 & \text{if } B \geq \min\{q_2,\ q_1\lambda_2/\lambda_1\}, \\ r(a) = \dfrac{q_1 - \lambda_1 a}{q_2 - \lambda_2 a} & \text{if } B < \min\{q_2, q_1\lambda_2/\lambda_1\} \text{ and } q_1/\lambda_1 \geq q_2/\lambda_2, \\ r(B/\lambda_2) = \dfrac{q_1 - \lambda_1 B/\lambda_2}{q_2 - B} & \text{otherwise.} \end{cases} \tag{2}$$

When $q_2 \leq q_1$ the second tier is dominated and we set $\rho^\star(q, \lambda, B) := 1$.

Theorem 2 (price of a budget, two tiers)**.** *For one provider and two tiers with $q_1 < q_2$:* (a) *every deterministic* DSIC, IR, BF *mechanism has worst-case ratio at most $\rho^\star(q, \lambda, B)$;* (b) BC-VCG *has worst-case ratio exactly $\rho^\star(q, \lambda, B)$;* (c) $\rho^\star = 1$ *if and only if $q_2 \leq q_1$ or $B \geq \min\{q_2,\ q_1\lambda_2/\lambda_1\}$.*

Proof sketch; Appendix A. Assume $B < \min\{q_2, q_1\lambda_2/\lambda_1\}$; the complementary case is checked directly in the appendix. A menu (p_1, p_2) with $p_t \leq B$ serves tier 2 only for $\theta \leq (p_2 - p_1)/(\lambda_2 - \lambda_1)$ and tier 1 up to p_1/λ_1. Any menu with positive ratio must set $p_1 = \min\{q_1, B\}$: a lower p_1 leaves types just below $\min\{q_1, B\}/\lambda_1$ unserved, where tier 2 is unaffordable at any price up to B and $\mathrm{OPT} > 0$, and a higher p_1 buys at a loss where $\mathrm{OPT} = 0$. With p_1 pinned and $p_2 \leq B$ the menu's switch point is at most a, while the optimum switches at $\min\{a^\star, B/\lambda_2\} > a$ exactly when $B < \min\{q_2, q_1\lambda_2/\lambda_1\}$; on the mismatch interval the ratio $r(\theta)$ is monotone with the sign of $\lambda_2 q_1 - \lambda_1 q_2$, so its infimum is $r(a)$ or $r(B/\lambda_2)$. BC-VCG has $p_1 = \min\{q_1, B\}$ and $p_2 = B$, so it switches exactly at a and agrees with the optimum outside the mismatch interval. □

Two features of (2) matter. The budget is harmless once it covers either the top tier's quality or the low tier's quality scaled up to the top tier's multiplier. The second condition is the cheaper one exactly when $q_1/\lambda_1 < q_2/\lambda_2$, and then the guarantee is discontinuous at the threshold: the clipped price makes the provider switch to tier 1 slightly early, on types where tier 1 is nearly worthless, and by part (a) no mechanism avoids this. Figure 1 (Appendix A) plots (2).

5.2 Several providers

Competition enters the menu only through W_{-i}, which shifts every quality of provider i down by the same amount before the clip. The two-tier analysis therefore transfers to the shifted instance.

Theorem 3 (n providers)**.** *On every instance,* BC-VCG *is* DSIC, IR *and* BF*; at most one provider has positive demand; and at every profile θ with $\mathrm{OPT}(\theta) > 0$ a provider is awarded and, with $(w, t^\star)$ the welfare-optimal feasible contract and $W_{-w} = W_{-w}(\theta)$,*

$$W_{\text{BC-VCG}}(\theta) \ \geq\ \mathrm{OPT}(\theta) - (q_{w,t^\star} - B - W_{-w})^+ \qquad \textit{and}$$
$$W_{\text{BC-VCG}}(\theta) \ \geq\ \rho_{\text{BC-VCG}}\big(q_w - W_{-w}, \lambda_w, B\big) \cdot \mathrm{OPT}(\theta),$$

where $\rho_{\text{BC-VCG}}(\tilde{q}, \lambda, B)$ is the single-provider worst-case ratio of BC-VCG *on the instance with qualities $\tilde{q}$. When $m = 2$, moreover, $W_{\text{BC-VCG}}(\theta) \geq \rho^\star(q_w, \lambda_w, B) \cdot \mathrm{OPT}(\theta)$ with the unshifted $\rho^\star$ of* (2)*, so that $\rho_{\text{BC-VCG}}(I) \geq \min_i \rho^\star(q_i, \lambda_i, B)$, while every deterministic* DSIC, IR, BF *mechanism M has $\rho_M(I) \leq \min_i \rho^\star(q_i, \lambda_i, B)$. Hence for two tiers* BC-VCG *is worst-case optimal on every instance with any number of providers, with ratio exactly $\min_i \rho^\star(q_i, \lambda_i, B)$.*

For real-valued $\tilde{q}$, $\rho_{\text{BC-VCG}}(\tilde{q}, \lambda, B)$ is the worst-case ratio of the single-provider menu $\min\{\tilde{q}_t, B\}$ against $\max\{0, \max_{t:\lambda_t\theta \leq B}(\tilde{q}_t - \lambda_t\theta)\}$; tiers with $\tilde{q}_t < 0$ carry a negative price and are never taken, so an accepted tier has welfare at least $p_t - \lambda_t\theta \geq 0$ and the ratio lies in $[0, 1]$. If no type gives the shifted instance a positive optimum, set $\rho_{\text{BC-VCG}}(\tilde{q}, \lambda, B) := 1$; at such a profile $u_w \geq 0$ gives shifted welfare at least 0, hence welfare at least $W_{-w} = \mathrm{OPT}$.

The additive bound says that whatever the number of tiers, the budget costs at most the part of the optimal contract's quality it cannot cover after competition. The multiplicative bound with the unshifted $\rho^\star$ is the one the experiments test profile by profile: competition does not erode the guarantee, since W_{-w} can only shrink the mismatch interval while r on it is unchanged. The matching upper bound holds because the type space is unbounded: on the profiles where every provider but i has cost above B, any mechanism reduces to a single-provider menu for i and Theorem 2(a) applies.

6 PRIVATE COST PER TIER: WHAT THE MULTIPLIER BUYS

Drop the public multiplier and let provider i's type be a vector $c_i = (c_{i,1}, \ldots, c_{i,m})$ of private costs, one per tier. This is the model for a provider that prices each tier directly, with no base cost for the public multiplier to scale. The offline optimum is unchanged with $c_{i,t}$ in place of $\lambda_{i,t}\theta_i$, and BC-VCG is defined verbatim: the menu $p_{i,t} = \min\{q_{i,t} - W_{-i}, B\}$ does not involve i's own report, so it remains a menu mechanism in the vector domain.

Theorem 4 (vector costs)**.** *In the vector-cost domain:* (a) BC-VCG *is* DSIC, IR, BF*, awards a provider whenever* $\mathrm{OPT}(\theta) > 0$*, and satisfies* $W_{\text{BC-VCG}}(\theta) \geq \mathrm{OPT}(\theta) - (q_{w,t^\star} - B - W_{-w})^+$ *on those profiles, with* $(w, t^\star)$ *the welfare-optimal feasible contract.* (b) *If some provider has two tiers with* $0 < q_{i,s} \leq B < q_{i,t}$*, then every deterministic* DSIC, IR, BF *mechanism has worst-case ratio* $\rho_M(I) \leq 0$ *on that instance.*

Part (b), proved in Appendix A, is where the two domains part: the adversary sets the expensive tier's cost at the budget and the cheap tier's cost just below its price, so the cheap tier is barely worth buying and the expensive one is worth a lot. A public multiplier forbids that profile: $c_t = B$ pins $\theta = B/\lambda_t$, hence $c_s = \lambda_s B/\lambda_t$, and the ratio of the two welfares is $r(B/\lambda_t)$. It is positive whenever $B < q_s\lambda_t/\lambda_s$, and it is the third branch of (2) when $q_1/\lambda_1 < q_2/\lambda_2$. When $B \geq q_s\lambda_t/\lambda_s$, Theorem 2 gives $\rho^\star = 1$ and the profile is no longer a mismatch profile. The single-parameter structure is thus exactly what converts the additive guarantee of part (a), which both domains share, into a multiplicative one. The obstruction is the monopoly a provider gets by offering both an affordable and an unaffordable tier, which the multidimensional budget-feasible literature names as why the standard benchmark admits no constant-factor approximation (Neogi et al., 2025); we keep OPT and shrink the domain instead.

Part (a) also corrects a natural but wrong design: a rule that ranks every (i, t) pair by $q_{i,t}/c_{i,t}$ and pays the winner its quality divided by the best *other* score, counting the winner's own losing tiers among the competition. With $q = \left[\begin{smallmatrix} 2 & 2 \\ 2 & 2 \end{smallmatrix}\right]$, $B = 2$ and costs $(1, 1)$ and $(2, 2)$ for the two providers, the first wins with score 2 against a best other score of 2 (its own second tier) and is paid 1; reporting 2 for that unawarded tier lowers the best other score to 1 and raises the payment to 2. The rule itself is what fails: on the same 32 instances the corrected menu has no profitable misreport (Section 8).

7 LEARNED QUALITIES AND CERTIFIED PRICES

Qualities in a real market are estimates. We keep the instance fixed over T rounds, let the platform learn q from probes it controls, and ask what survives.

Protocol. Each round τ has four steps. (1) *Probe:* the platform measures one pair (i, t), the next on a fixed round-robin schedule, observing an independent $y \sim \mathrm{Bernoulli}(q_{i,t}/Q)$ for a known scale $Q \geq \max q_{i,t}$; the probe is a measurement, not a contract. (2) *Estimate:* from all probes so far it forms, for every pair, the scaled empirical mean $\hat{q}_{i,t} = Q \cdot \frac{1}{N_{i,t}} \sum y$ and the Hoeffding half-width $h_{i,t} = Q\sqrt{\log(2nmT/\delta)/(2N_{i,t})}$ over $N_{i,t}$ probes (Hoeffding, 1963; Lattimore & Szepesvári, 2019), and sets $\underline{q}_{i,t} = \max\{0, \hat{q}_{i,t} - h_{i,t}\}$, $\overline{q}_{i,t} = \min\{Q, \hat{q}_{i,t} + h_{i,t}\}$ (unprobed pairs get 0 and Q). (3) *Report:* providers report costs r^τ for this round's contract. (4) *Award:* the platform runs BC-VCG on the instance $(\underline{q}^\tau, \lambda, B)$. We call this LCB-BC-VCG; running step (4) on $\overline{q}^\tau$ instead gives UCB-BC-VCG.

Theorem 5 (certified prices)**.** (a) *For every realization of the probes, the estimates* $(\underline{q}^\tau, \overline{q}^\tau)$ *are computed without any report, and provider* i*'s round-*τ *menu depends on* r^τ_{-i} *only. Hence, at every round, after every history and for every* r^τ_{-i}*, truthful reporting is a dominant strategy in the round-*τ *game, and the mechanism is* IR *and* BF *at every round. If in addition the other providers' reports do not depend on* i*'s past reports (opponents playing history-independent strategies, truthful reporting among them), then truthful reporting at every round maximizes the total utility of a provider with a fixed type.* (b) *On the event* $\mathcal{G} = \{\underline{q}^\tau \leq q \leq \overline{q}^\tau \text{ for all } \tau\}$*, which has probability at least* $1 - \delta$*,* LCB-BC-VCG *under truthful reporting never awards a contract whose welfare* $q_{i,t} - \lambda_{i,t}\theta_i$ *at the true qualities is negative.* (c) *Suppose the per-round profiles* θ^τ *are i.i.d. across rounds with coordinates independent across providers, and that each* θ_i *has a density bounded by* $\bar{f}_i$*. Let* $s := \min_i \min\{\lambda_{i,1},\ \min_{t'<t}(\lambda_{i,t} - \lambda_{i,t'})\}$ *and* $L := 2\max_i \bar{f}_i/s$*. Then, with* $\delta = 1/T$ *and*

$K := n\big(m + \binom{m}{2}\big) \le nm^2$*, the expected regret of* LCB-BC-VCG *against the offline menu on the true qualities is*

$$\mathbb{E}\sum_{\tau=1}^{T}\Big[W_{\text{BC-VCG}(q)}(\theta^\tau) - W_{\text{LCB-BC-VCG}}(\theta^\tau)\Big] \le 16KLQ^2\sqrt{nmT\log(2nmT^2)} + 4KLQ^2nm + 2Q.$$

Proof sketch; Appendix A. (a) The probe schedule and hence $\underline{q}^\tau$ do not depend on reports, so each round is BC-VCG on a fixed instance and Theorem 3 applies round by round; provider i's report reaches no later menu. (b) A provider that accepts tier t at round τ is paid at most $\underline{q}^\tau_{i,t} \le q_{i,t}$ and accepts only if the price covers its cost. (c) On $\mathcal{G}$ every entry of $\underline{q}^\tau$ is within $2\tilde{h}^\tau$ of q, every price in the menu moves by at most $4\tilde{h}^\tau$, and the outcome differs from the outcome on q only if some comparison margin under q is within $8\tilde{h}^\tau$ of a tie. Where the outcomes agree the welfares at the true qualities agree; where they differ and $\mathcal{G}$ holds the loss is at most Q by (b); off $\mathcal{G}$ it is at most $2Q$. Summing $\tilde{h}^\tau$ over the round-robin schedule gives the bound. □

Part (b) is the reason to price from the lower bound: a price built from $\overline{q}$ promises a quality the platform has not certified, and when the promise exceeds the true quality the platform pays more than the contract is worth. Incentive compatibility constrains the provider's surplus, and the loss here falls on the platform. UCB-BC-VCG has regret of the same order: on $\mathcal{G}$ an accepted price is at most $\overline{q} \le q + 2\tilde{h}^\tau$, so a round on which the two menus disagree loses at most $Q + 2\tilde{h}^\tau \le 2Q$ and the proof of (c) goes through with twice the constant. Empirically its regret is smaller (Section 8), and it buys at a loss in a positive fraction of rounds there; it has no analogue of part (b). Part (a) separates probes from awards: a provider's report never reaches the estimator, so the exploration-separation that truthful bandit mechanisms require (Babaioff et al., 2009) holds by construction. The obstacle to an intertemporal dominant-strategy statement is the other route: opposing reports may depend on i's own past reports, which is why part (a) carries its additional condition. Part (c) is a statement about profiles that vary: its margin condition is the low-noise condition of (Audibert & Tsybakov, 2007), and every margin is affine in one θ_i with slope at least s in absolute value, so it holds for independent θ_i uniform on $[0, \Theta_i]$ with $L = 2/(s\min_i \Theta_i)$. The trivial bound is $2TQ$, so part (c) improves on it as long as $KLQ\sqrt{nm\log(2nmT^2)}$ stays below $\sqrt{T}$. On a fixed profile at a comparison boundary the two menus can disagree at every round, and no learning rate helps.

7.1 Learning from awarded contracts

Suppose instead that quality becomes known only when a contract is awarded. Observations then reach the estimator through the award, the award depends on reports, and the separation of part (a) is gone; Babaioff et al. (2009) show this is not a technicality, since truthful learning there requires the exploration decisions themselves to be report-independent. Let $Q \ge \max_{i,t} q_{i,t}$ be a known scale, let the observation after awarding tier t to provider i be an independent $y \sim \text{Bernoulli}(q_{i,t}/Q)$, and assume

$$\pi := \min_{i,t} \Pr[\lambda_{i,t}\theta_i \le B] > 0, \tag{3}$$

so every pair is affordable with positive probability. ETC-LCB has two phases. *Exploration:* for $\min\{nmN, T\}$ rounds it works through the pairs in a fixed round-robin order, posting the price B on the assigned tier and $-\infty$ everywhere else. If the provider accepts, which under truthful play it does exactly when $\lambda_{i,t}\theta_i \le B$, the platform records y and updates a *running-maximum* lower bound: after the k-th accepted observation of the pair it forms $\underline{q}^{(k)}_{i,t} = (\hat{q}^{(k)}_{i,t} - h^{(k)}_{i,t})^+$ from the first k observations and keeps $\underline{q}^{\max}_{i,t} = \max_{k' \le k} \underline{q}^{(k')}_{i,t}$; unobserved pairs keep 0. *Commitment:* it freezes $\underline{q}^{\max}$ and runs LCB-BC-VCG on it for the remaining rounds. The running maximum is what a plain lower bound lacks: one more observation can lower $\underline{q}^{(k)}$ but never $\underline{q}^{\max}$ (Proposition 1).

Theorem 6 (learning from awarded contracts)**.** (i) *Under* ETC-LCB*, every round's menu depends on* r^τ_{-i} *only, and in exploration rounds on no report at all, so the mechanism is* DSIC, IR *and* BF *at every round.* (ii) *Fix a provider* i *with a fixed type* θ_i *and suppose the other providers' actions (reports and acceptance decisions) do not depend on* i*'s past actions. Then, for every realization of all observation streams, reporting truthfully and accepting exactly the affordable exploration contracts gives* i *a total*

utility at least that of any other strategy. (iii) *Under the assumptions of Theorem 5(c) and* (3)*, with* $N = \lceil (T^2\ell/(\pi nm))^{1/3} \rceil$, $\ell := \log(2nmT^2)$ *and* $\delta = 1/T$*, the expected regret against the offline menu on the true qualities is at most*

$$\Big[(nm)^{2/3}(Q+B) + 8\,(nm)^{1/6}KLQ^2\Big]\Big(\frac{T^2\ell}{\pi}\Big)^{1/3} + (nm+2)(Q+B)$$

whenever $T \geq 512\,nm\log^2(nmT)/\pi^2$*, and at most* $T(Q+B)$ *otherwise; the rate is* $T^{2/3}$*.*

Proposition 1 (withholding under a plain lower bound)**.** *Replace* $\underline{q}^{\max}$ *in* ETC-LCB *by the plain frozen bound* $(\hat{q}_{i,t}^{(N_{i,t})} - h_{i,t}^{(N_{i,t})})^+$*. Then there are instances, horizons and histories of positive probability at which a provider with an affordable exploration contract strictly prefers to decline it: the declined observation could lower its own commitment price, and over enough commitment rounds that loss exceeds the surplus* $B - \lambda_{i,t}\theta_i$ *of accepting.*

Exploration costs at most $Q + B$ per round, since the benchmark welfare is at most Q and the mechanism's is at least $-B$; balancing that against the frozen-estimate error at the stated N gives the bound, and N uses π but not L. ETC-LCB pays a $T^{2/3}$ rate where Theorem 5 pays $\sqrt{T}$, the same rate deterministic truthful bandit mechanisms pay in Babaioff et al. (2009); whether it is necessary here is open. Part (ii) is the one place where learning and incentives interact: the platform learns provider i's quality only from contracts i accepts, so i controls its own sample. Proposition 1 shows that with a plain lower bound this control is worth exercising, and the running maximum removes the incentive because an accepted contract can only raise i's future prices and lower its competitors' demand. Updating $\underline{q}$ after every award would reopen the channel: the estimate would depend on which contracts were awarded and the award on the reports.

8 Experiments

All mechanisms are implemented in one package; discrete families are checked by exact rational enumeration over all type profiles and all unilateral misreports, so every violation count below is exact, while continuous instances are evaluated in floating point on sampled profiles. Appendix B lists generators, seeds and further details.

E1: the discrete family, and what a coarse grid hides. On the family that motivated this work (2955 instances, two or three providers, two tiers with $\lambda = (1, 2)$, integer-grid qualities, $\theta_i \in \{1, 2\}$, $B \in \{2, 3, 4\}$, 20,568 profiles), Appendix Table 2 puts seven payment rules side by side. Paying the base-unit threshold breaks IR 1052 times and DSIC 3552 times, the latter always a losing winner reporting its way out of the deal, as Corollary 1(b) predicts. The threshold rules that fix the tier from public data and pay $\lambda_t\tau$ have no violations; with index tie-breaks they reach ratio $1/2$ at worst, with quality tie-breaks ratio 1 on every profile. Unclipped VCG exceeds the budget on 2110 profiles in 706 instances, its prices being blind to the budget. BC-VCG has no violations, average ratio 0.9992, worst-case ratio $1/2$ on 15 instances, and in every one of the 2955 instances its measured worst ratio is at or above the instance floor $\min_\theta \rho_{\text{BC-VCG}}(q_w - W_{-w}, \lambda_w, B)$ of Theorem 3, with equality in 2474 (2460 of them at floor 1). The feasible baselines the budget literature suggests do worse: truncating qualities at the budget before unclipped VCG (Le, 2018) averages 0.9448, and posting $\min\{q_{i,t}, B\}$ while ignoring competition (Badanidiyuru et al., 2012) averages 0.7122.

The quality-tie rule therefore looks better than BC-VCG on this grid, and Theorem 2(a) says that cannot be true of any mechanism on the underlying continuous instances. Appendix B resolves the 15 instances: the cost grid hides the types the tie-break fails on.

E2: continuous instances. We draw 300 instances with $n \in \{1, 2, 3, 5\}$, $m \in \{2, 3, 4\}$, multiplier ladders with ratio 2 to 5 and per-provider jitter, qualities concave in the multiplier, and B between 0.15 and 1.6 times the top quality; each is evaluated on 400 profiles with $\theta_i \sim U(0, q_{i,1}/\lambda_{i,1})$. Appendix Table 3 reports the worst and mean ratio per instance, and Figure 2 (Appendix B) the distribution behind it. Ties no longer occur, so the two threshold rules coincide: their median instance-worst ratio is 0.12 and 95% of instances fall below $1/2$, against a median instance-worst ratio of 0.95 and a profile-mean ratio of 0.976 for BC-VCG, with no instance below 0.18. On the 101 two-tier instances its measured worst ratio is never below the predicted floor, while the threshold rule falls below it on

82. On those 60 non-trivial two-tier instances the worst ratio over 400 sampled profiles is *strictly* above the closed form in every case, while a witness profile at $\theta = a + 10^{-15}$ reproduces it: the infimum is a limit, as Theorem 2(b) states. Le-truncated VCG, the closest feasible baseline, has median 0.84 and falls below $1/2$ on 31% of instances against 24% for BC-VCG, and unclipped VCG overruns the budget on 26,222 of the 120,000 profiles (119,521 with OPT > 0 are used for ratios).

E6: five-tier price ladders. On 200 instances of synthetic five-tier ladders ($n \in \{2, 3, 4\}$, $m = 5$ tiers, concave qualities, B between 0.2 and 1.5 times the top quality), the same ordering holds empirically outside the scope of the two-tier theory. Appendix Table 3 (lower block): the threshold rule has median instance-worst ratio 0.19 and falls below $1/2$ on 89% of instances, BC-VCG has median 1.00, profile-mean ratio 0.994 and 6% below $1/2$, and unclipped VCG overruns the budget on 5535 of the 80,000 profiles. The extra tiers widen the gap between the two rules and leave BC-VCG ahead on every statistic reported.

E3 and E4: beyond the two-tier, single-parameter case. A grid search over all menus on 120 instances with $m \in \{3, 4\}$ finds a menu beating BC-VCG by more than 10^{-3} on 3 of them and never by more than 0.029, all in exact arithmetic: BC-VCG is not worst-case optimal beyond two tiers. In the vector-cost domain, the ranking rule of Section 6 admits 704 profitable misreports on the 32 instances and 736,899 on 60 random ones, with 638 underpayments, while BC-VCG admits none, at worst profile ratio $1/3$ and profile-mean ratio 0.997.

E5: learning. We run $T = 5{,}000/10{,}000/20{,}000$ rounds with five seeds on 60 discrete and 60 continuous instances, profiles redrawn every round (Appendix B), and compare LCB-BC-VCG, UCB-BC-VCG and the estimated threshold rule that preceded this work. Figure 3 (Appendix B) shows cumulative regret against the offline BC-VCG on the true qualities. LCB-BC-VCG and UCB-BC-VCG are sublinear, with finite-horizon fitted log–log slopes of 0.32 and 0.29 and bootstrap intervals $[0.18, 0.41]$ and $[0.20, 0.39]$ on the continuous instances, while the threshold rule is linear at 1.02 with interval $[1.01, 1.02]$, its regret growing fourfold when the horizon does. LCB-BC-VCG has no round with negative welfare at the true qualities in any of the 600 runs at $T = 20{,}000$; UCB-BC-VCG loses 11,333 over 8267 such rounds on the continuous instances and the threshold rule 45,608 over 16,525. The two menus trade regret for loss-making rounds: at $T = 20{,}000$ UCB-BC-VCG has regret 9263 against 12,795 for LCB-BC-VCG, and 11,333 of cumulative negative welfare against 0.

E8: exploration. ETC-LCB learns from awarded contracts alone. On single-provider instances where $\rho^\star = 1$ the regret is consistent with the $T^{2/3}$ rate of Theorem 6, with finite-horizon fitted slope 0.6982 (bootstrap interval $[0.6931, 0.7021]$) and 0.6885 (interval $[0.6740, 0.6988]$) on continuous instances, realizing 0.909 and 0.564 of OPT: exploration is expensive exactly where affordability is rare, taking 86% of the $T = 10{,}000$ horizon at $\pi = 0.052$. Every negative-welfare round is an exploration round, and the committed phase has none over 1,606,480 rounds.

9 Conclusion

A budget on payments and a choice among tiers pull a truthful mechanism in opposite directions. The tier welfare wants cannot be priced inside the budget, since its truthful payment carries the option value of the lower tiers, and a rule that fixes the tier before the report stays inside the budget while giving up that option. For two tiers we gave the exact best worst-case ratio any deterministic truthful, individually rational, budget-feasible mechanism can guarantee; a budget-clipped VCG menu attains it, and the budget costs nothing above a threshold we give in closed form. The menu extends to any number of providers, and to private cost vectors, where the public multiplier turns "the budget costs at most the quality it cannot cover" into a positive ratio. Learned qualities keep the per-round guarantees: certified prices lose $\tilde{O}(\sqrt{T})$ welfare when the bounds hold, and from awarded contracts a running-maximum bound stops a fixed-type provider withholding, at regret $T^{2/3}$. Beyond two tiers the clipped menu is not always worst-case optimal, and the exact ratio is open; on the instances we search the gap to the best menu stays below 3% (Section 8), so the loss from not knowing it appears small.

AI USE STATEMENT

An AI coding assistant wrote the enumeration code, the online simulator, the table and figure generators and the number audit under the authors' specification, and drafted text that the authors revised. Every count and ratio reported here is recomputed from the released code. We have reviewed all AI-assisted work and take responsibility for the final content of this work, including text, claims and artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

The mechanisms, verifiers, instance generators and the online simulator are released in the anonymous supplementary material. The core modules ship compiled, so the results reproduce without exposing the implementation, while the experiment drivers, which only build the instances and tabulate what the compiled primitives return, are plain Python; the package depends on nothing beyond the standard library and NumPy. The discrete families are enumerated in exact rational arithmetic over all type profiles and all unilateral misreports, so the reported violation counts are exact.

REFERENCES

Georgios Amanatidis, Sophie Klumper, Evangelos Markakis, Guido Schäfer, and Artem Tsikiridis. Partial allocations in budget-feasible mechanism design: bridging multiple levels of service and divisible agents. *ACM Transactions on Economics and Computation*, 13(2):1–28, 2025a.

Georgios Amanatidis, Evangelos Markakis, Christodoulos Santorinaios, Guido Schäfer, Panagiotis Tsamopoulos, and Artem Tsikiridis. Online budget-feasible mechanism design with predictions. In *International Symposium on Algorithmic Game Theory*, pp. 402–421. Springer, 2025b.

Ioannis Anagnostides, Alkis Kalavasis, and Tuomas Sandholm. Barriers to welfare maximization with no-regret learning, 2024. URL https://arxiv.org/abs/2411.01720.

Ioannis Anagnostides, Kshipra Bhawalkar, Christopher Liaw, Aranyak Mehta, Renato Paes Leme, Yifeng Teng, Grigoris Velegkas, and Weiqiang Zheng. Compensation design, 2026a. URL https://arxiv.org/abs/2607.14438.

Ioannis Anagnostides, Kshipra Bhawalkar, Christopher Liaw, Aranyak Mehta, Grigoris Velegkas, and Weiqiang Zheng. From compensation design to budget-feasible mechanisms: A constant approximation for subadditive valuations, 2026b. URL https://arxiv.org/abs/2608.04337.

Aaron Archer and Éva Tardos. Truthful mechanisms for one-parameter agents. In *Proceedings 42nd IEEE Symposium on Foundations of Computer Science*, pp. 482–491. IEEE, 2001.

Mark Armstrong. Multiproduct nonlinear pricing. *Econometrica: Journal of the Econometric Society*, pp. 51–75, 1996.

Jean-Yves Audibert and Alexandre B Tsybakov. Fast learning rates for plug-in classifiers. 2007.

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

Moshe Babaioff, Robert D. Kleinberg, and Aleksandrs Slivkins. Truthful mechanisms with implicit payment computation. In *Proceedings of the 11th ACM Conference on Electronic Commerce*, EC '10, pp. 43–52, New York, NY, USA, 2010. Association for Computing Machinery. ISBN 9781605588223. doi: 10.1145/1807342.1807349. URL https://doi.org/10.1145/1807342.1807349.

Ashwinkumar Badanidiyuru, Robert Kleinberg, and Yaron Singer. Learning on a budget: posted price mechanisms for online procurement. In *Proceedings of the 13th ACM conference on electronic commerce*, pp. 128–145, 2012.

Martin Bichler, Maximilian Fichtl, and Matthias Oberlechner. Computing bayes nash equilibrium strategies in auction games via simultaneous online dual averaging. In *Proceedings of the 24th ACM Conference on Economics and Computation*, pp. 294–294, 2023.

Hau Chan and Jing Chen. Truthful multi-unit procurements with budgets. In *International Conference on Web and Internet Economics*, pp. 89–105. Springer, 2014.

Andreas Charalampopoulos, Dimitris Fotakis, Panagiotis Patsilinakos, and Thanos Tolias. A competitive posted-price mechanism for online budget-feasible auctions. In *Proceedings of the 26th ACM Conference on Economics and Computation*, pp. 1046–1075, 2025.

Yeon-Koo Che and Ian Gale. Standard auctions with financially constrained bidders. *The Review of Economic Studies*, 65(1):1–21, 1998.

Haolong Chen, Zhengyuan Xin, Liang Zhang, Lei Xue, and Guangxu Zhu. Error-aware reverse auction mechanism for large language model routing, 2026. URL https://arxiv.org/abs/2608.12719.

Lingjiao Chen, Matei Zaharia, and James Zou. FrugalGPT: How to use large language models while reducing cost and improving performance. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. URL https://openreview.net/forum?id=cSimKw5p6R. Featured Certification.

Edward H Clarke. Multipart pricing of public goods. *Public choice*, 11:17–33, 1971.

Vincent Conitzer and Tuomas Sandholm. Complexity of mechanism design. In *Proceedings of the Eighteenth conference on Uncertainty in artificial intelligence*, pp. 103–110, 2002.

Nikhil R. Devanur and Sham M. Kakade. The price of truthfulness for pay-per-click auctions. In *Proceedings of the 10th ACM Conference on Electronic Commerce*, EC '09, pp. 99–106, New York, NY, USA, 2009. Association for Computing Machinery. ISBN 9781605584584. doi: 10.1145/1566374.1566388. URL https://doi.org/10.1145/1566374.1566388.

Shahar Dobzinski, Ron Lavi, and Noam Nisan. Multi-unit auctions with budget limits. *Games and Economic Behavior*, 74(2):486–503, 2012.

Theodore Groves. Incentives in teams. *Econometrica: Journal of the Econometric Society*, pp. 617–631, 1973.

Jason D Hartline and Tim Roughgarden. Optimal mechanism design and money burning. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pp. 75–84, 2008.

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

Yunlong Hou, Zixin Zhong, and Vincent YF Tan. On the benefits of free exploration for regret minimization in multi-armed bandits. *arXiv preprint arXiv:2605.25789*, 2026.

Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics*, 49(2):1055–1080, 2021.

Nathanael Hyafil and Craig Boutilier. Regret minimizing equilibria and mechanisms for games with strict type uncertainty. In *UAI*, pp. 268–277. AUAI Press, 2004.

Levente Kocsis and Csaba Szepesvári. Bandit based monte-carlo planning. In *European conference on machine learning*, pp. 282–293. Springer, 2006.

Tor Lattimore and Csaba Szepesvári. Bandit algorithms. *Cambridge University*, 355, 2019.

Phuong Le. Pareto optimal budgeted combinatorial auctions. *Theoretical Economics*, 13(2):831–868, 2018.

Xiang Liu, Hau Chan, Minming Li, and Weiwei Wu. Budget feasible mechanisms: A survey. International Joint Conferences on Artificial Intelligence Organization, 2024.

Andreas Maurer and Massimiliano Pontil. Empirical bernstein bounds and sample-variance penalization. In *COLT 2009 - The 22nd Conference on Learning Theory, Montreal, Quebec, Canada, June 18-21, 2009*, 2009. URL http://www.cs.mcgill.ca/%7Ecolt2009/papers/012.pdf#page=1.

Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

Robert Müller and Clemens Müller. Cattle trade: A multi-agent benchmark for llm bluffing, bidding, and bargaining, 2026. URL https://arxiv.org/abs/2605.14537.

Rian Neogi, Kanstantsin Pashkovich, and Chaitanya Swamy. Multidimensional budget-feasible mechanism design. *arXiv preprint arXiv:2508.09367*, 2025.

Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani. *Algorithmic game theory*. Cambridge university press, 2007.

Isaac Ong, Amjad Almahairi, Vincent Wu, Wei-Lin Chiang, Tianhao Wu, Joseph E Gonzalez, Mohammed Kadous, and Ion Stoica. Routellm: Learning to route llms from preference data. In *International Conference on Learning Representations*, volume 2025, pp. 34433–34448, 2025.

Fabian R Pieroth, Nils Kohring, and Martin Bichler. Deep reinforcement learning for equilibrium computation in multistage auctions and contests. *Management Science*, 2025.

Riccardo Poiani, Martino Bernasconi, and Andrea Celli. Regret minimization in single-dimensional contract-design with binary actions. *arXiv preprint arXiv:2606.06125*, 2026.

Jean-Charles Rochet and Philippe Choné. Ironing, sweeping, and multidimensional screening. *Econometrica*, pp. 783–826, 1998.

Jean-Charles Rochet and Lars A Stole. The economics of multidimensional screening. *Econometric Society Monographs*, 35:150–197, 2003.

Aviad Rubinstein and Junyao Zhao. Strategizing against no-regret learners in first-price auctions. In *Proceedings of the 25th ACM Conference on Economics and Computation*, pp. 894–921, 2024.

Shai Shalev-Shwartz. Online learning and online convex optimization. *Machine Learning*, 4(2): 107–194, 2011.

Yaron Singer. Budget feasible mechanisms. In *2010 IEEE 51st Annual Symposium on foundations of computer science*, pp. 765–774. IEEE, 2010.

Emanuel Tewolde, Xiao Zhang, David Guzman Piedrahita, Vincent Conitzer, and Zhijing Jin. Coopeval: Benchmarking cooperation-sustaining mechanisms and llm agents in social dilemmas. *arXiv preprint arXiv:2604.15267*, 2026.

William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37, 1961.

Ian Waudby-Smith and Aaditya Ramdas. Estimating means of bounded random variables by betting. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 86(1), 2024.

Brian Zhang, Gabriele Farina, Ioannis Anagnostides, Federico Cacciamani, Stephen McAleer, Andreas Haupt, Andrea Celli, Nicola Gatti, Vincent Conitzer, and Tuomas Sandholm. Computing optimal equilibria and mechanisms via learning in zero-sum extensive-form games. *Advances in Neural Information Processing Systems*, 36:2659–2699, 2023.

Itai Zilberstein, Pranav Rajbhandari, Steve Chien, and Tuomas Sandholm. Distributed constraint optimization via online learning and iterative pricing with application to large-scale satellite scheduling, 2026. URL https://arxiv.org/abs/2607.25835.

A PROOFS

Throughout, fix i and r_{-i}. In a deterministic DSIC mechanism two reports of i that receive the same tier t receive the same payment: if reports r, r' both obtain t at payments $p < p'$, then every type reporting r gains by reporting r'. Hence every tier i can obtain at all is obtained at a unique price $p_{i,t}(r_{-i})$; set $p_{i,t} = -\infty$ for tiers i cannot obtain. DSIC and IR together say that the option i receives at type θ_i maximizes $p_{i,t} - \lambda_{i,t}\theta_i$ over the obtainable tiers and the outside option (utility 0); ties are resolved by the mechanism. IR is what puts the outside option into the comparison, and BF is $p_{i,t} \le B$ for every obtainable tier. This is the taxation principle (Nisan et al., 2007, Ch. 9). Conversely, a profile of such menus together with a selection rule that awards at most one provider and gives every provider a utility-maximizing option is DSIC; for BC-VCG the selection rule is the smallest-index rule (a tie in both utility and welfare between two tiers of one provider, which leaves the outcome's welfare unchanged, is broken toward the smaller tier index) and Lemma 4 shows it never has to deny a provider with positive demand.

A.1 THEOREM 1 AND COROLLARY 1

Fix i and r_{-i}. Provider i's cost for the outcome with multiplier x is $\theta_i x$, so i is a single-parameter agent with allocation $x_i(\theta_i) \in \{0\} \cup \{\lambda_{i,t}\}_t$, and the characterization of DSIC allocations as monotone (here non-increasing, since higher cost must not buy a larger multiplier) with the payment $p_i(\theta_i) = \theta_i x_i(\theta_i) + \int_{\theta_i}^{\infty} x_i(s)\, ds$ is the standard one (Myerson, 1981)(Nisan et al., 2007, Thm. 9.36); the constant of integration is 0 because the model pays only the selected provider (Section 3). The integral is finite without using BF: by DSIC at $\theta_i = 0$, type 0 reporting θ_i would receive $p_i(\theta_i)$ at zero cost, so $p_i(\theta_i) \le p_i(0) < \infty$ for every selected θ_i; by IR, $\lambda_{i,1}\theta_i \le \lambda_{i,t}\theta_i \le p_i(\theta_i)$; hence no type above $p_i(0)/\lambda_{i,1}$ is selected and x_i vanishes beyond a finite b_K. If i is selected on no set of positive length, $K = 0$, $x_i = 0$ almost everywhere, and any selected type (only $\theta_i = 0$, by monotonicity) is paid 0 by the envelope formula. Otherwise list the multipliers $\mu_1 > \cdots > \mu_K$ that x_i takes on sets of positive length and the thresholds $0 = b_0 < b_1 < \cdots < b_K$ between them; since x_i takes finitely many values and is non-increasing, $x_i = \mu_k$ on the open interval (b_{k-1}, b_k) and $x_i = 0$ beyond b_K. At a breakpoint b_k the two adjacent tiers give the provider the same utility (their payments differ by $b_k(\mu_k - \mu_{k+1})$), and the mechanism may award either of them, or a tier taken at that single type only; the payments of the listed tiers do not depend on this choice, so the theorem does not fix it (Corollary 1 fixes it for the allocation it studies). A tier of multiplier μ awarded only at the single type b (a breakpoint, or $b = 0$) is paid $b\mu + \int_b^{\infty} x_i(s)\, ds$, the envelope formula with its own threshold b; it is not among the μ_k. For $\theta_i \in (b_{k-1}, b_k)$,

$$p_i = \theta_i\mu_k + \mu_k(b_k - \theta_i) + \sum_{j>k} \mu_j(b_j - b_{j-1}) = \mu_k b_k + \sum_{j>k} \mu_j(b_j - b_{j-1}),$$

which is (1); at $\theta_i = b_k$ the same computation gives $\mu_k b_k + \sum_{j>k} \mu_j(b_j - b_{j-1})$ if tier μ_k is awarded and $\mu_{k+1}b_{k+1} + \sum_{j>k+1} \mu_j(b_j - b_{j-1})$ if tier μ_{k+1} is, each the payment (1) of the tier awarded.

(a) With $K = 1$ the sum is empty. (b) If the selected provider is paid $b_1 < \lambda_{i,t}\theta_i$, which happens on $\theta_i \in (b_1/\lambda_{i,t}, b_1]$ when $\lambda_{i,t} > 1$, its utility is negative; any report above b_1 yields utility 0. (c) The map $\theta_i \mapsto \arg\max_{o \in \{\varnothing\} \cup \{t : \lambda_{i,t}\theta_i \le B\}} \mathrm{welfare}(o)$ selects a non-increasing multiplier (with $x_i = 0$ at $\varnothing$) because $q_{i,t} - \lambda_{i,t}\theta_i$ has strictly decreasing differences in $(\lambda_{i,t}, \theta_i)$ and the feasible set shrinks with θ_i; whether i is selected at all is monotone for the same reason.

Example 1. One provider, $q = (1, 2)$, $\lambda = (1, 2)$, $B = 3/2$. The welfare-optimal feasible allocation awards tier 2 for $\theta \le 3/4$ (the largest θ with $2\theta \le B$) and tier 1 for $\theta \in (3/4, 1]$. By (1) the payment for tier 2 is $2 \cdot \frac{3}{4} + 1 \cdot (1 - \frac{3}{4}) = \frac{7}{4} > B$, although the tier's own threshold term $2 \cdot \frac{3}{4} = \frac{3}{2}$ fits the budget exactly.

A.2 THEOREM 2

One provider, tiers $(q_1, \lambda_1), (q_2, \lambda_2)$ with $\lambda_1 < \lambda_2$ and $q_1 < q_2$. Write $\bar{p}_1 = \min\{q_1, B\}$, $a^\star = (q_2 - q_1)/(\lambda_2 - \lambda_1)$, $a = (B - q_1)^+/(\lambda_2 - \lambda_1)$, $\bar{\theta} = \min\{a^\star, B/\lambda_2\}$ and $r(\theta) = (q_1 - \lambda_1\theta)/(q_2 - \lambda_2\theta)$. In the regime $B < q_2$ of the two lemmas below, $B/\lambda_2 < q_2/\lambda_2$, so the offline optimum buys tier 2 for $\theta \le \bar{\theta}$ (affordable, non-negative welfare, at least as good as tier 1), tier 1 for $\theta \in (\bar{\theta}, \bar{p}_1/\lambda_1]$ when

this interval is non-empty, and nothing beyond. Note $r'(\theta)$ has the sign of $\lambda_2 q_1 - \lambda_1 q_2$, and $r(\theta) < 1$ for $\theta < a^\star$.

Lemma 1. *$a < \bar{\theta}$ if and only if $B < \min\{q_2,\, q_1\lambda_2/\lambda_1\}$.*

Proof. If $B \le q_1$ then $a = 0 < \bar{\theta}$ and $B < q_2$, $B \le q_1 < q_1\lambda_2/\lambda_1$. If $B > q_1$: $a < a^\star$ iff $B - q_1 < q_2 - q_1$ iff $B < q_2$; and $a < B/\lambda_2$ iff $(B - q_1)\lambda_2 < B(\lambda_2 - \lambda_1)$ iff $B\lambda_1 < q_1\lambda_2$. □

Lemma 2 (the cheap tier's price is pinned)**.** *Assume $B < \min\{q_2, q_1\lambda_2/\lambda_1\}$. Any menu (p_1, p_2) with $p_t \le B$ and worst-case ratio $\rho > 0$ has $p_1 = \bar{p}_1$.*

Proof. The assumption gives $B/\lambda_2 < q_1/\lambda_1$ and $B/\lambda_2 < B/\lambda_1$, hence $B/\lambda_2 < \bar{p}_1/\lambda_1$. Suppose $p_1 < \bar{p}_1$ and pick $\theta \in (\max\{p_1/\lambda_1, B/\lambda_2\}, \bar{p}_1/\lambda_1)$. Tier 1 is not accepted ($p_1 < \lambda_1\theta$), tier 2 is not accepted ($p_2 \le B < \lambda_2\theta$), and $\mathrm{OPT}(\theta) \ge q_1 - \lambda_1\theta > q_1 - \bar{p}_1 \ge 0$, so the ratio at θ is 0. Suppose $p_1 > \bar{p}_1$. If $\bar{p}_1 = B$ this violates BF. If $\bar{p}_1 = q_1 < B$, pick $\theta \in (q_1/\lambda_1, p_1/\lambda_1)$. Since $\theta > q_1/\lambda_1 > B/\lambda_2$, tier 2 is not acceptable ($p_2 \le B < \lambda_2\theta$), while tier 1 is ($p_1 > \lambda_1\theta$), so the provider takes tier 1 with welfare $q_1 - \lambda_1\theta < 0$. If $\mathrm{OPT}(\theta) > 0$ the ratio at θ is negative; if $\mathrm{OPT}(\theta) = 0$ the mechanism buys at a loss where the optimum buys nothing, so its ratio is $-\infty$ by the convention of Section 3. In every case $\rho \le 0$. □

Lemma 3 (the mismatch interval)**.** *Assume $B < \min\{q_2, q_1\lambda_2/\lambda_1\}$ and let (p_1, p_2) be a menu with $p_1 = \bar{p}_1$, $p_2 \le B$. For every $\theta \in (a, \bar{\theta})$ the provider does not take tier 2, so its welfare ratio at θ is at most $r(\theta)$; hence its worst-case ratio is at most $\inf_{\theta\in(a,\bar{\theta})} r(\theta) = \rho^\star(q, \lambda, B)$.*

Proof. Tier 2 is preferred to tier 1 only if $p_2 - \lambda_2\theta \ge p_1 - \lambda_1\theta$, i.e. $\theta \le (p_2 - \bar{p}_1)/(\lambda_2 - \lambda_1) \le (B - \bar{p}_1)^+/(\lambda_2 - \lambda_1) = a$, where we used $p_2 \le B$ and, when $\bar{p}_1 = B$, that $p_2 - \bar{p}_1 \le 0$. On $(a, \bar{\theta})$ tier 2 is optimal with $\mathrm{OPT} = q_2 - \lambda_2\theta > 0$ and the provider takes tier 1 or nothing, giving ratio at most $\max\{0, r(\theta)\} = r(\theta)$; on the interval the numerator of r is positive because $\theta < \bar{\theta} \le B/\lambda_2 < q_1/\lambda_1$, and the denominator is positive because $\theta < \bar{\theta} \le B/\lambda_2$ gives $q_2 - \lambda_2\theta > q_2 - B > 0$. Since r is monotone, its infimum over the open interval is $r(a)$ when $\lambda_2 q_1 \ge \lambda_1 q_2$ and $r(\bar{\theta})$ otherwise; in the latter case $a^\star > q_1/\lambda_1 > B/\lambda_2$, so $\bar{\theta} = B/\lambda_2$. This is (2). □

Proof of Theorem 2. (a) If $B \ge \min\{q_2, q_1\lambda_2/\lambda_1\}$ the bound $\rho^\star = 1$ is trivial. Otherwise a mechanism with ratio ≤ 0 satisfies the bound, and one with positive ratio has $p_1 = \bar{p}_1$ by Lemma 2 and ratio at most $\rho^\star$ by Lemma 3.

(b) BC-VCG offers $p_1 = \bar{p}_1$ and $p_2 = \min\{q_2, B\}$. If $B \ge q_2$ the menu is (q_1, q_2), price differences equal quality differences, the provider's choice is the welfare-optimal tier whenever it buys, it buys whenever some tier has non-negative welfare, and affordability is automatic since a price of at most $q_2 \le B$ covers the cost; the ratio is 1. If $q_1\lambda_2/\lambda_1 \le B < q_2$, then $a \ge B/\lambda_2$ by the computation in Lemma 1, and $B/\lambda_2 < a^\star$ because $q_2 > B \ge q_1\lambda_2/\lambda_1$ gives $\lambda_2(q_2 - q_1) > (\lambda_2 - \lambda_1)q_2 > (\lambda_2 - \lambda_1)B$; so for $\theta \le B/\lambda_2$ the provider takes tier 2 (its utility $B - \lambda_2\theta \ge q_1 - \lambda_1\theta$ there, with equality resolved toward tier 2 by the tie rule of Definition 1) and so does the optimum, and for $\theta > B/\lambda_2$ tier 2 is neither affordable to the optimum nor accepted, while tier 1 is offered at q_1 and taken exactly when its welfare is non-negative; the ratio is 1. If $B < \min\{q_2, q_1\lambda_2/\lambda_1\}$, the menu is $(\bar{p}_1, B)$ with switch point exactly a. On $[0, a]$ the provider takes tier 2 and so does the optimum ($a < \bar{\theta}$). On $(a, \bar{\theta})$ the provider takes tier 1, which it accepts since $\bar{p}_1 - \lambda_1\theta > \bar{p}_1 - \lambda_1\bar{\theta} \ge 0$ (Lemma 2 gives $B/\lambda_2 < \bar{p}_1/\lambda_1$), while tier 2 is optimal, for ratio exactly $r(\theta)$. At $\theta = \bar{\theta}$ two cases arise: if $\bar{\theta} = a^\star$ both tiers have equal welfare and the ratio is 1; if $\bar{\theta} = B/\lambda_2 < a^\star$ the optimum still buys tier 2 (cost exactly B) while the provider, whose tier-2 utility is $B - \lambda_2\bar{\theta} = 0 < \bar{p}_1 - \lambda_1\bar{\theta}$, takes tier 1, for ratio $r(B/\lambda_2)$. On $(\bar{\theta}, \bar{p}_1/\lambda_1]$ both take tier 1, and beyond it both buy nothing. The worst-case ratio is therefore $\inf_{(a,\bar{\theta}]} r$, which equals $r(a)$ when r is non-decreasing ($\lambda_2 q_1 \ge \lambda_1 q_2$): a limit not attained when r is strictly increasing, and attained on the whole interval when $\lambda_2 q_1 = \lambda_1 q_2$, where r is constant; and $r(B/\lambda_2)$ when r is decreasing (attained at $\bar{\theta} = B/\lambda_2$). Both agree with (2).

(c) is Lemma 1 together with (b): the mismatch interval is non-empty exactly when $B < \min\{q_2, q_1\lambda_2/\lambda_1\}$, and on it $r < 1$.

If $q_2 \le q_1$, tier 2 has lower quality at higher cost and is never optimal; in the BC-VCG menu $p_2 \le p_1$ and $\lambda_2 > \lambda_1$, so the provider never prefers tier 2 and the ratio is 1. □

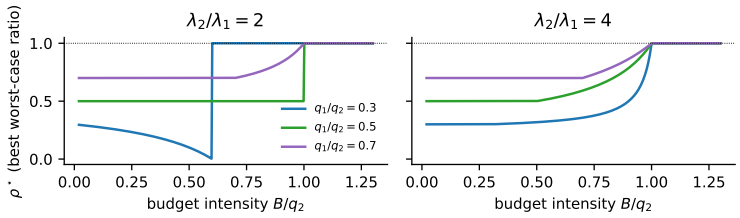

Figure 1: The closed form (2) against budget intensity B/q_2, for two multiplier ratios and three quality ratios q_1/q_2. Each curve is 1 on the plateau $B \geq \min\{q_2, q_1\lambda_2/\lambda_1\}$ and drops discontinuously at its edge.

A.3 Theorem 3

Lemma 4. *Under* BC-VCG*, if $u_i > 0$ then $u_j < 0$ for every $j \neq i$. If $u_i \geq 0$ then provider i has a feasible tier with welfare at least W_{-i}.*

Proof. $u_i \geq 0$ means some t has $\min\{q_{i,t} - W_{-i}, B\} \geq \lambda_{i,t}\theta_i$, hence $\lambda_{i,t}\theta_i \leq B$ and $q_{i,t} - \lambda_{i,t}\theta_i \geq W_{-i}$, which is the second claim. If moreover $u_i > 0$ then $q_{i,t} - \lambda_{i,t}\theta_i > W_{-i} \geq$ the best feasible welfare of any $j \neq i$, while $u_j \geq 0$ would give j a feasible tier with welfare $\geq W_{-j} \geq q_{i,t} - \lambda_{i,t}\theta_i$, a contradiction. □

Proof of Theorem 3. Incentives. Provider i's menu depends on r_{-i} only. If i is awarded under some report, it receives a tier from its menu at the menu price, so its utility is at most $u_i(\theta_i)$, the demand at its true type. If $u_i(\theta_i) > 0$, truthful reporting wins: by Lemma 4 no other provider has non-negative demand, so i is the smallest-index provider with $u \geq 0$. If $u_i(\theta_i) = 0$, truthful reporting yields 0, and no report yields more. If $u_i(\theta_i) < 0$, any award yields negative utility and truthful reporting yields 0. Hence DSIC; IR and BF hold because the awarded price is a menu price with $u_i \geq 0$ and $p \leq B$.

Award when $\mathrm{OPT} > 0$. Let $(w, t^\star)$ be welfare-optimal and feasible with $\mathrm{OPT} = q_{w,t^\star} - \lambda_{w,t^\star}\theta_w \geq W_{-w}$. Then $\min\{q_{w,t^\star} - W_{-w}, B\} - \lambda_{w,t^\star}\theta_w \geq \min\{\mathrm{OPT} - W_{-w}, B - \lambda_{w,t^\star}\theta_w\} \geq 0$, so $u_w \geq 0$ and some provider is awarded. If a provider $i \neq w$ is awarded, then $u_i \geq 0$ and by Lemma 4 $u_w \leq 0$, so $u_i = u_w = 0$; but then the tier i takes has $u_i \geq 0$, hence is feasible with welfare at least $W_{-i} \geq \mathrm{OPT}$, so i is also welfare-optimal and the analysis below applies to i in place of w.

Additive bound. Let t' be the tier w takes. Then $\min\{q_{w,t'} - W_{-w}, B\} - \lambda_{w,t'}\theta_w \geq \min\{q_{w,t^\star} - W_{-w}, B\} - \lambda_{w,t^\star}\theta_w$, and $q_{w,t'} - \lambda_{w,t'}\theta_w \geq \min\{q_{w,t'} - W_{-w}, B\} - \lambda_{w,t'}\theta_w + W_{-w} \geq \min\{\mathrm{OPT},\ B - \lambda_{w,t^\star}\theta_w + W_{-w}\} = \mathrm{OPT} - (q_{w,t^\star} - B - W_{-w})^+$.

Shifted multiplicative bound. If $\widetilde{\mathrm{OPT}} = 0$ the claim is the remark after the theorem. Otherwise the choice of w is that of a single provider facing the menu $\min\{\tilde{q}_t, B\}$ with $\tilde{q} = q_w - W_{-w}$, whose offline optimum is $\widetilde{\mathrm{OPT}} = \mathrm{OPT} - W_{-w} \geq 0$ (feasibility is unaffected by the shift). By definition of $\rho_{\text{BC-VCG}}$, $\tilde{q}_{t'} - \lambda_{w,t'}\theta_w \geq \rho_{\text{BC-VCG}}(\tilde{q}, \lambda_w, B)\,\widetilde{\mathrm{OPT}}$, and adding W_{-w} to both sides gives welfare at least $\rho_{\text{BC-VCG}}\,\widetilde{\mathrm{OPT}} + W_{-w} \geq \rho_{\text{BC-VCG}}\,\mathrm{OPT}$, because $\widetilde{\mathrm{OPT}} = \mathrm{OPT} - W_{-w}$, $W_{-w} \geq 0$ and $\rho_{\text{BC-VCG}} \leq 1$.

Unshifted bound for $m = 2$. Fix θ with $\mathrm{OPT}(\theta) > 0$. If the awarded provider is some $i \neq w$, then by Lemma 4 its tier has welfare at least $W_{-i} \geq \mathrm{OPT}$, so the ratio is 1. Otherwise w is awarded with $u_w \geq 0$; write $W = W_{-w}$, which does not depend on θ_w, and $q = q_w$, $\lambda = \lambda_w$. If $q_2 \leq q_1$ (where $\rho^\star := 1$), then $p_1 \geq p_2$ and $(p_1 - \lambda_1\theta_w) - (p_2 - \lambda_2\theta_w) = (p_1 - p_2) + (\lambda_2 - \lambda_1)\theta_w \geq 0$, so w takes tier 1, or tier 2 only when both utilities and both welfares coincide; the tier taken is feasible and its welfare $q_1 - \lambda_1\theta_w$ is the largest feasible welfare of w, hence equals OPT, for ratio 1. So let $q_1 < q_2$. If $q_2 - W \leq B$ neither price is clipped, utility differences equal welfare differences, and with ties resolved toward the larger welfare (Definition 1) the tier taken is welfare-optimal, for ratio 1. So assume $q_2 - W > B$, where $p_2 - p_1 = B - \min\{q_1 - W, B\} = (B - q_1 + W)^+$ is at most $q_2 - q_1$, so tier 2 is never taken where tier 1 has strictly larger welfare. The only remaining case is that tier 2 is optimal, i.e. $\theta_w \leq \bar{\theta} = \min\{a^\star, B/\lambda_2\}$, and w takes tier 1, so tier 1 is taken only if

$(\lambda_2 - \lambda_1)\theta_w > p_2 - p_1$, i.e. $\theta_w > \tilde{a} := (B - q_1 + W)^+/(\lambda_2 - \lambda_1) \geq a$ (at equality the tie rule picks tier 2, whose welfare is larger for $\theta_w < a^\star$); tier 1 is accepted at a price $\leq q_1 - W \leq q_1$ so its welfare $q_1 - \lambda_1\theta_w$ is non-negative, and the ratio is $r(\theta_w)$ with the *unshifted* q. Since $\theta_w \in (\tilde{a}, \bar{\theta}] \subseteq (a, \bar{\theta}]$ and r is monotone, $r(\theta_w) \geq \inf_{(a,\bar{\theta}]} r = \rho^\star(q_w, \lambda_w, B)$ by the computation in Theorem 2(b). (When $B \geq q_1\lambda_2/\lambda_1$ the interval $(a, \bar{\theta}]$ is empty by Lemma 1, so this case does not arise and the ratio is $1 = \rho^\star$.)

Upper bound for every mechanism. Fix i and fix every other provider's type at a value $\theta_j > B/\lambda_{j,1}$, so that no tier of j is feasible; these providers contribute nothing to OPT at any θ_i. Restricted to the profiles that vary only in θ_i, mechanism M is a deterministic DSIC mechanism for the single provider i, hence a menu with prices at most B, and Theorem 2(a) gives $\rho_M(I) \leq \rho^\star(q_i, \lambda_i, B)$ when $q_{i,1} < q_{i,2}$, while $\rho_M(I) \leq 1 = \rho^\star(q_i, \lambda_i, B)$ holds by definition when $q_{i,2} \leq q_{i,1}$; take the minimum over i. Applied to BC-VCG itself (whose menu at these profiles is $\min\{q_{i,t}, B\}$) this shows its ratio is exactly $\min_i \rho^\star(q_i, \lambda_i, B)$. □

A.4 THEOREM 4

(a) The incentive argument, Lemma 4, the award-when-OPT > 0 argument and the additive bound in the proof of Theorem 3 never use the product structure $c_{i,t} = \lambda_{i,t}\theta_i$; replacing $\lambda_{i,t}\theta_i$ by $c_{i,t}$ throughout gives (a).

Proof of (b). Take one provider with tiers s, t and menu prices $p_s, p_t \leq B$ (by the taxation principle every deterministic DSIC mechanism is such a menu, a tier it never offers having price $-\infty$, which the case $p_s < q_s$ covers); other providers and tiers are made irrelevant by large costs, as in the lower bound above, and costs range over all of $\mathbb{R}^m_{\geq 0}$, so the profiles below are admissible. If $p_s < q_s$, choose $c_s \in (p_s, q_s)$ and $c_t > B$: both tiers are declined, so nothing is awarded, yet $\mathrm{OPT} \geq q_s - c_s > 0$. If $p_s \geq q_s$, choose $c_t = B$ and $c_s = p_s - \varepsilon$: tier s yields utility $\varepsilon > 0 \geq p_t - B$, so it is taken with welfare $q_s - p_s + \varepsilon \leq \varepsilon$, while $\mathrm{OPT} \geq q_t - B > 0$. Letting $\varepsilon \to 0$ gives an infimum at most 0 (negative if $p_s > q_s$). □

Extra tiers of the same provider are removed by giving them cost above B at the profiles used.

A.5 THEOREM 5

(a) The probe at round τ is the τ-th element of a fixed schedule and its outcome is a draw from the environment (the scale Q is positive, so the Bernoulli parameters are defined), so the matrices $\underline{q}^\tau, \overline{q}^\tau$ are measurable with respect to the probe outcomes alone. Round τ is BC-VCG on the instance $(\underline{q}^\tau, \lambda, B)$, which Theorem 3 (whose proof does not use $q \geq 0$ beyond $W_{-i} \geq 0$, guaranteed by the $\max\{0, \cdot\}$) shows is DSIC, IR and BF. Under the additional condition of (a) the sequence $(r^\tau_{-i})_\tau$ is the same along every strategy of i, so the total utility is a sum of per-round utilities each maximized by truth.

(b) By Hoeffding's inequality and a union bound over the nm pairs and the T rounds, $\Pr[\mathcal{G}] \geq 1 - \delta$. On $\mathcal{G}$ an awarded provider i with tier t is paid $p = \min\{\underline{q}^\tau_{i,t} - W, B\} \leq \underline{q}^\tau_{i,t} \leq q_{i,t}$ (with $W \geq 0$) and accepts only if $p \geq \lambda_{i,t}\theta_i$, so $q_{i,t} - \lambda_{i,t}\theta_i \geq 0$.

(c) Define $\tilde{h}^\tau = \max_{i,t} \min\{Q/2, h^\tau_{i,t}\}$, with $h^\tau_{i,t} = +\infty$ for an unprobed pair; $\tilde{h}^\tau$ is deterministic because the schedule is fixed. On $\mathcal{G}$, $0 \leq q_{i,t} - \underline{q}^\tau_{i,t} \leq \min\{Q, 2h^\tau_{i,t}\} \leq 2\tilde{h}^\tau$ for every pair, hence $|W_{-i}(\underline{q}^\tau) - W_{-i}(q)| \leq 2\tilde{h}^\tau$, every menu price moves by at most $4\tilde{h}^\tau$, every utility $p_{i,t} - \lambda_{i,t}\theta_i$ by at most $4\tilde{h}^\tau$, and every difference of two such utilities by at most $8\tilde{h}^\tau$. The outcome of BC-VCG is a function of the signs of the $K = n(m + \binom{m}{2})$ comparisons "utility of tier t for provider i against 0" and "tier t against tier t' for provider i" (the tier a provider takes and the choice of the winning index among providers with $u \geq 0$ are determined by these signs, and under the density assumption a tie has probability 0), so the outcomes under $\underline{q}^\tau$ and q coincide unless some comparison under q has margin at most $8\tilde{h}^\tau$; by a union bound over the K comparisons this event has probability at most $8KL\tilde{h}^\tau$ over the draw of θ^τ, because each margin is affine in the single coordinate θ_i it involves (the slope is $-\lambda_{i,t}$ against 0 and $\lambda_{i,t'} - \lambda_{i,t}$ between two tiers, at least s in absolute value), and by

independence of the coordinates the conditional probability of landing within η of a tie is at most $2\eta\bar{f}_i/s \le L\eta$. When the outcomes coincide the the welfares at the true qualities coincide; when they differ on $\mathcal{G}$ the loss is at most Q by (b), and off $\mathcal{G}$ it is at most $2Q$, since the mechanism either buys nothing ($\mathrm{OPT} \le Q$) or pays a price at most $\underline{q}^{\tau}_{i,t} \le Q$ for a tier whose cost is at most that price, so its welfare is at least $-Q$. Let E_τ be the event that some comparison under q has margin at most $8\tilde{h}^\tau$. The round-τ regret is at most $Q\,\mathbf{1}_{E_\tau} + 2Q\,\mathbf{1}_{\mathcal{G}^c}$, so its expectation is at most $8KLQ\tilde{h}^\tau + 2Q\delta$. Under round-robin, every pair has $N^\tau \ge \lfloor \tau/(nm) \rfloor$ probes at round τ, so $\tilde{h}^\tau \le Q/2$ for every τ, and for $\tau \ge nm$, $\lfloor \tau/(nm) \rfloor \ge \tau/(2nm)$ gives $\tilde{h}^\tau \le Q\sqrt{nm\log(2nmT/\delta)/\tau}$. Summing, $\sum_\tau \tilde{h}^\tau \le nmQ/2 + Q\sqrt{nm\log(2nmT/\delta)}\cdot 2\sqrt{T}$, and with $\delta = 1/T$ the total expected regret is at most $8KLQ\big(nmQ/2 + 2Q\sqrt{nmT\log(2nmT^2)}\big) + 2\delta TQ = 16KLQ^2\sqrt{nmT\log(2nmT^2)} + 4KLQ^2nm + 2Q$.

A.6 Theorem 6 and Proposition 1

(i) In an exploration round assigned to (i,t) the menu of provider i is the price B on tier t and $-\infty$ elsewhere, and every other provider has the empty menu, so no report enters; in a commitment round the menu is the LCB-BC-VCG menu on the frozen $\underline{q}^{\max}$, which depends on r^τ_{-i} only. Every price is at most B and a contract is awarded only to a provider with $u_i \ge 0$, so IR and BF hold, and a report reaches only the current round's menus of the other providers, so truthful reporting is optimal round by round.

(ii) Coupling. The observation model of Section 7.1 is that the observations of a pair are independent $\mathrm{Bernoulli}(q_{i,t}/Q)$ draws, independent across pairs and of the cost profiles; the coupling below only names them in advance. Attach to every pair (i,t) an i.i.d. stream $y^{(1)}_{i,t}, y^{(2)}_{i,t}, \ldots$ of $\mathrm{Bernoulli}(q_{i,t}/Q)$ draws; the k-th accepted contract of the pair reveals $y^{(k)}_{i,t}$, whenever it occurs. Fix the streams and the other providers' actions, which by assumption are the same along every strategy of i. Let $\sigma^\star$ be "report truthfully, accept an exploration contract iff $\lambda_{i,t}\theta_i \le B$" and σ any strategy. Write $N^\sigma_{i,t}$ for the number of observations of (i,t) at the end of exploration under σ. For an *affordable* tier ($\lambda_{i,t}\theta_i \le B$), $\sigma^\star$ accepts every offer, so $N^\sigma_{i,t} \le N^{\sigma^\star}_{i,t}$ and the observations under σ are a prefix of those under $\sigma^\star$; since $\underline{q}^{\max}$ is the running maximum along the prefix, $\underline{q}^{\max,\sigma}_{i,t} \le \underline{q}^{\max,\sigma^\star}_{i,t}$. The other providers' estimates coincide under σ and $\sigma^\star$, hence so does W_{-i} in every commitment round.

Exploration rounds. Under $\sigma^\star$ provider i earns $B - \lambda_{i,t}\theta_i \ge 0$ on affordable offers and 0 otherwise; under σ it earns at most that, since accepting an unaffordable offer yields $B - \lambda_{i,t}\theta_i < 0$.

Commitment rounds. Fix a round and r_{-i}. Under truthful reporting i's realized utility is $\max\{0, u_i^{\sigma^\star}(\theta_i)\}$: if $u_i > 0$ it is the unique provider with non-negative demand (Lemma 4) and is awarded; if $u_i \le 0$ it earns 0. Under σ its utility is at most $\max\{0, u_i^\sigma(\theta_i)\}$, where $u_i^\sigma(\theta_i) = \max_t\big(\min\{\underline{q}^{\max,\sigma}_{i,t} - W_{-i}, B\} - \lambda_{i,t}\theta_i\big)$ is the demand at its true type on the σ-menu (any award is a menu item at its menu price). For an affordable tier the σ-price is at most the $\sigma^\star$-price by the prefix inequality; for an unaffordable tier every price is at most $B < \lambda_{i,t}\theta_i$, so its term is negative under both strategies and never attains the maximum in $\max\{0, u_i\}$. Hence $\max\{0, u_i^\sigma(\theta_i)\} \le \max\{0, u_i^{\sigma^\star}(\theta_i)\}$, and the round utility under σ is at most that under $\sigma^\star$. Summing over rounds gives (ii). (Observations of i's unaffordable tiers, which σ may collect at a loss, enter no price i can accept and enter W_{-j} only through tiers feasible at i's report, which under truthful reporting they are not.)

(iii) Write $M = nmN$. If $M \ge T$ the regret is at most $T(Q+B) \le M(Q+B)$. If $M < T$ but $T < 512\,nm\log^2(nmT)/\pi^2$ the same trivial bound $T(Q+B)$ is the one stated. Otherwise the exploration phase costs at most $M(Q+B)$, as each round's benchmark welfare is at most Q and the mechanism's at least $-B$. Under $\sigma^\star$ the acceptances of pair (i,t) are i.i.d. with probability $p_{i,t} \ge \pi$, so $N_{i,t} \sim \mathrm{Binomial}(N, p_{i,t})$ and the multiplicative Chernoff bound (Shalev-Shwartz, 2011) gives $\Pr[N_{i,t} < \pi N/2] \le e^{-\pi N/8}$; a union bound over the nm pairs gives the count event $\mathcal{C} = \{N_{i,t} \ge \pi N/2\ \forall (i,t)\}$ probability at least $1 - nm\,e^{-\pi N/8} \ge 1 - 1/T$ as soon as $N \ge (8/\pi)\log(nmT)$, which the stated N satisfies when $T^2\ell \ge 512\,nm\log^3(nmT)/\pi^2$, in particular when $T \ge 512\,nm\log^2(nmT)/\pi^2$. Let $\mathcal{G}' = \{\,|Q\bar{y}^{(k)}_{i,t} - q_{i,t}| \le h^{(k)}_{i,t}\ \forall (i,t),\ \forall k \le T\}$, where

$\bar{y}_{i,t}^{(k)}$ is the mean of the first k observations of the pair and $h_{i,t}^{(k)} = Q\sqrt{\ell/(2k)}$. By the two-sided Hoeffding inequality each of the nmT events fails with probability at most $2e^{-\ell} = 1/(nmT^2)$, so $\Pr[\mathcal{G}'] \geq 1 - 1/T = 1 - \delta$. On $\mathcal{G}'$, $\underline{q}_{i,t}^{(k)} = (Q\bar{y}_{i,t}^{(k)} - h_{i,t}^{(k)})^+ \leq q_{i,t}$ for every k, hence $\underline{q}_{i,t}^{\max} \leq q_{i,t}$; and $q_{i,t} - \underline{q}_{i,t}^{\max} \leq q_{i,t} - \underline{q}_{i,t}^{(N_{i,t})} \leq 2h_{i,t}^{(N_{i,t})}$, because either $Q\bar{y}_{i,t}^{(N_{i,t})} - h_{i,t}^{(N_{i,t})} \geq 0$ and the difference is at most $2h_{i,t}^{(N_{i,t})}$ directly, or it is negative, $\underline{q}_{i,t}^{(N_{i,t})} = 0$ and $q_{i,t} \leq Q\bar{y}_{i,t}^{(N_{i,t})} + h_{i,t}^{(N_{i,t})} < 2h_{i,t}^{(N_{i,t})}$. On $\mathcal{E} = \mathcal{C} \cap \mathcal{G}'$ every frozen half-width is at most $h = Q\sqrt{\ell/(\pi N)}$. Conditionally on the exploration history, each commitment round is BC-VCG on a fixed instance within $2h$ of q against a fresh profile θ^τ, so the argument of Theorem 5(c) applies with $\tilde{h}^\tau$ replaced by h: the outcomes differ from those on q with probability at most $8KLh$ over θ^τ, and when they differ the loss is at most Q, since on $\mathcal{G}'$ a price is at most the true quality. The commitment phase therefore costs at most $8KLQ^2T\sqrt{\ell/(\pi N)}$ in expectation. Off $\mathcal{E}$, which has probability at most $2/T$, the loss is at most $T(Q+B)$, contributing at most $2(Q+B)$. With $(T^2\ell/(\pi nm))^{1/3} \leq N \leq (T^2\ell/(\pi nm))^{1/3} + 1$,

$$M(Q+B) \leq (nm)^{2/3}(Q+B)\Big(\tfrac{T^2\ell}{\pi}\Big)^{1/3} + nm(Q+B),$$
$$8KLQ^2T\sqrt{\tfrac{\ell}{\pi N}} \leq 8\,(nm)^{1/6}KLQ^2\Big(\tfrac{T^2\ell}{\pi}\Big)^{1/3},$$

and adding the three contributions gives the bound.

Proof of Proposition 1. One provider, one tier, $\lambda = 1$, $Q = B = 1$, $q \in (0,1)$, $N \geq \ell$ exploration offers, $h_k = \sqrt{\ell/(2k)}$. Condition on the first $N-1$ accepted observations all being 1, an event of probability $q^{N-1} > 0$, and consider the last offer. Declining freezes the price $p_- = 1 - h_{N-1}$; accepting freezes $p_+ = 1 - h_N$ if $y = 1$ and $p_0 = (N-1)/N - h_N$ if $y = 0$. Since $h_{N-1} - h_N \leq \sqrt{\ell/2}/(2(N-1)\sqrt{N}) < 1/N$ for $N \geq \ell$, we have $p_0 < p_-$; pick a type $c \in (p_0, p_-)$ with $p_- - c \geq \frac{q}{1-q}(h_{N-1} - h_N) + \varepsilon$ for some $\varepsilon > 0$, which is possible because $\frac{1}{1-q}(h_{N-1} - h_N) < 1/N - (h_{N-1} - h_N) = p_- - p_0$ for N large. The contract is affordable ($c < p_- \leq B$). Over the $T - N$ commitment rounds the provider earns $(p-c)^+$ per round at the frozen price p. Accepting minus declining, in expectation, is

$$\begin{aligned}&(1-c) + (T-N)\big[q(p_+ - c) - (p_- - c)\big]\\ &\quad = (1-c) + (T-N)\big[q(h_{N-1} - h_N) - (1-q)(p_- - c)\big] \leq (1-c) - (T-N)(1-q)\varepsilon,\end{aligned}$$

which is negative once $T - N > (1-c)/((1-q)\varepsilon)$. Under the running maximum the frozen price is $\max\{p_-, p\} \geq p_-$ whatever y is, and accepting is weakly better. □

Why (ii) fixes the type. With types redrawn every round the acceptance strategy $\sigma^\star$ is no longer optimal. One provider, one tier, $\lambda = 1$, $Q = B = 1$, $\theta^\tau \in \{\frac{1}{2}, \frac{6}{5}\}$ with equal probability. An exploration offer at type $\frac{6}{5}$ is unaffordable and accepting it loses $\frac{1}{5}$, but if the earlier observations were all 1 and this one is too (probability q), it raises the running-maximum price from $1 - h_{k-1}$ to $1 - h_k$; over the commitment rounds with type $\frac{1}{2}$, about $(T-N)/2$ of them, the provider then earns $h_{k-1} - h_k$ more per round, which exceeds $\frac{1}{5}$ once T is large. The deviation buys information at a loss. The commitment-phase analysis in (iii) is unchanged along it, since it only adds observations, all covered by $\mathcal{G}'$ (in such an exploration round the mechanism's welfare is $q_{i,t} - \lambda_{i,t}\theta_i$, which can fall below the $-B$ that the exploration cost in (iii) uses); but a dynamic dominant-strategy statement for redrawn types is false, which is why (ii) is stated for a fixed type.

B Experimental details

Discrete family (E1). Two providers with qualities $q_{i,t} \in \{1,2,3,4\}$ or three providers with $q_{i,t} \in \{1,2,3\}$, two tiers, $\lambda_i = (1,2)$ for every provider, $B \in \{2,3,4\}$: $256 \cdot 3 + 729 \cdot 3 = 2955$ instances. Costs $\theta_i \in \{1,2\}$, so an instance has 2^n profiles and the family has 20,568. DSIC is checked in its strongest discrete form: for every true profile, every report profile and every provider whose report differs from its type, truthful reporting must weakly dominate. IR and BF are checked at the truthful profile. Ratios use the guarded welfare of the original study (a profile with a violation

Table 2: E1: the discrete family, 2955 instances and 20,568 profiles, exact enumeration. Left block: violation counts, over (profile, misreport) pairs for DSIC and over profiles for IR and budget, with the number of instances carrying a violation in parentheses. Right block: welfare ratios against the offline optimum, over all profiles, on guarded welfare (a profile with a violation scores 0); the average offline optimum is 1.4600.

	violations			welfare ratio		
rule	DSIC	IR	budget	worst	avg.	avg. welfare
base-unit threshold	3552 (561)	1052 (571)	0	0	0.9459	1.3939
threshold, index ties	0	0	0	1/2	0.9928	1.4456
threshold, quality ties	0	0	0	1	1.0000	1.4600
Le-truncated VCG	0	0	0	0	0.9448	1.3794
posted price	0	0	0	0	0.7122	0.9975
VCG (unclipped)	0	0	2110 (706)	0	0.8805	1.2431
BC-VCG	0	0	0	1/2	0.9992	1.4583

Table 3: E2 (300 continuous instances, uniform costs) and E6 (200 five-tier ladder instances, log-normal costs), 400 sampled profiles each. Left block: statistics of the per-instance worst ratio over the sampled profiles — its minimum, median and mean over instances, and the fraction of instances below 1/2 and below 0.1. Right block: mean ratio and budget overruns over all profiles.

		worst ratio, per instance					mean ratio, per profile	
family	rule	min	median	mean	$< 1/2$	< 0.1	mean ratio	overruns
E2	threshold rule (index or quality ties)	0.00	0.12	0.171	0.95	0.43	0.573	0
	Le-truncated VCG	0.04	0.84	0.718	0.31	0.01	0.944	0
	posted price	0.00	0.00	0.203	0.81	0.72	0.660	0
	VCG (unclipped; overruns score 0)	0.00	0.00	0.470	0.53	0.53	0.782	26222
	BC-VCG	0.18	0.95	0.762	0.24	0.00	0.976	0
E6	threshold rule (index or quality ties)	0.00	0.19	0.248	0.89	0.18	0.693	0
	Le-truncated VCG	0.13	1.00	0.845	0.11	0.00	0.986	0
	posted price	0.00	0.01	0.022	0.99	0.96	0.680	0
	VCG (unclipped; overruns score 0)	0.00	1.00	0.530	0.47	0.47	0.931	5535
	BC-VCG	0.35	1.00	0.875	0.06	0.00	0.994	0

scores 0); for menu mechanisms this coincides with the realized welfare since they have no violations. "avg. welfare" is the per-instance average over profiles, then averaged over instances. The five rules are the base-unit threshold rule (winner by $\max_t q_{i,t}/\lambda_{i,t}$ per unit of report, tier by the same public score, price $\min\{\tau, B\}$), the capped threshold rule (same winner and tier, price $\min\{\lambda_t\tau, B, q_t\}$, award iff the report is at most price$/\lambda_t$), its variant with both ties broken toward higher quality, unclipped VCG (Definition 1 without the clip and with W_{-i} over all tiers), and BC-VCG. The predicted floor of Theorem 3 is $\min_\theta \rho_{\text{BC-VCG}}(q_w - W_{-w}, \lambda_w, B)$, the minimum over the profiles with $\mathrm{OPT}(\theta) > 0$, with w the welfare-optimal provider at that profile and $\rho_{\text{BC-VCG}}$ set to 0 when a shifted quality is negative.

The 15 instances on which BC-VCG has ratio 1/2 are also the only ones where the quality-tie rule beats it, and one of them shows how the enumeration misses the failure. At $q = \left[\begin{smallmatrix} 1 & 1 \\ 2 & 4 \end{smallmatrix}\right]$, $\lambda_t = (1, 2)$, $B = 2$ and $\theta = (2, \theta_1)$ the two tiers tie on q_t/λ_t, and the quality tie-break fixes tier 2 at price 2. Types $\theta_1 \le 1$ accept with surplus $2 - 2\theta_1 \ge 0$, while types $\theta_1 \in (1, 2)$, where $\mathrm{OPT} = 2 - \theta_1 > 0$, are offered nothing they can accept. The family's cost grid $\{1, 2\}$ contains no such type, so the enumeration never sees the failure. BC-VCG offers both tiers at price 2 and the type $\theta_1 = 1$ takes the cheap one, for ratio $1/2 = \rho^\star$.

Continuous instances (E2, E5). For each instance: $n \in \{1, 2, 3, 5\}$ and $m \in \{2, 3, 4\}$ uniformly; a ladder ratio $\kappa \in \{2, 3, 4, 5\}$; for each provider $\lambda_{i,1} = 1$ and $\lambda_{i,t} = \lambda_{i,t-1}\,\kappa\,U(0.8, 1.25)$; a base quality $U(1, 10)$ and exponent $\alpha \sim U(0.3, 0.9)$ with $q_{i,t} = \text{base} \cdot \lambda_{i,t}^{\alpha} \cdot U(0.9, 1.1)$, sorted;

Table 4: E2 stratified, on the same per-instance worst ratios as Table 3. The median is the upper median over the instances of the stratum; the $B/\max q$ blocks cut at the terciles of β, 0.690 and 1.171.

		threshold rule		BC-VCG	
stratum	inst.	median	below $1/2$	median	below $1/2$
$n = 1$	75	0.248	88.0%	0.991	25.3%
$n = 2$	84	0.100	92.9%	1.000	23.8%
$n = 3$	66	0.116	100.0%	0.955	21.2%
$n = 5$	75	0.097	100.0%	0.770	26.7%
$m = 2$	101	0.216	85.1%	0.919	20.8%
$m = 3$	90	0.135	100.0%	1.000	22.2%
$m = 4$	109	0.066	100.0%	0.788	29.4%
$B/\max q$ low	101	0.084	98.0%	0.419	64.4%
$B/\max q$ mid	100	0.165	92.0%	0.954	8.0%
$B/\max q$ high	99	0.132	94.9%	1.000	0.0%

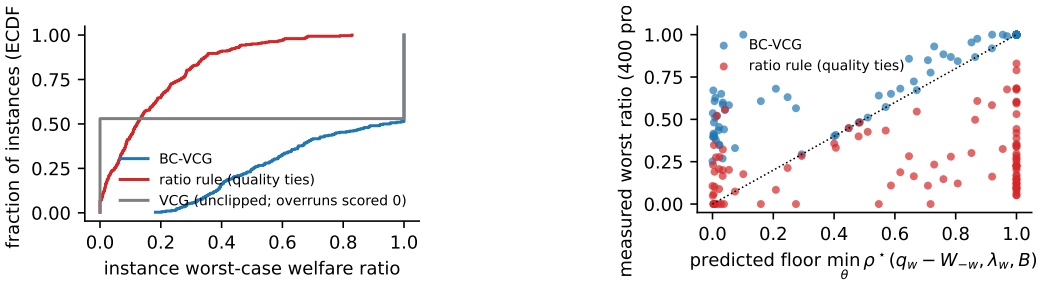

Figure 2: E2, 300 continuous instances. Left: distribution of the per-instance worst-case welfare ratio. Right: the two-tier instances, predicted floor $\min_\theta \rho^\star$ of Theorem 3 against the measured worst ratio; BC-VCG is on or above the diagonal on every instance and the tie-broken threshold rule is not.

$B = \beta \max q$ with $\beta \sim U(0.15, 1.6)$. Profiles are $\theta_i \sim U(0, q_{i,1}/\lambda_{i,1})$ independently, 400 per instance, and profiles with OPT $= 0$ are skipped (119,521 remain). Seed 0 for E2; E5 uses the first 60 instances of seed 1. For the two threshold rules with $n = 1$ (no competitor, $\tau = \infty$) the price is $\min\{B, q_t\}$. The posted-price baseline collapses to a median instance-worst ratio of 0.00 on this family.

Stratifying those instance-level ratios by provider count, tier count and budget ratio keeps the ordering inside every stratum (Table 4). The budget blocks carry the interpretation: where $B/\max q \le 0.690$, BC-VCG's median instance-worst ratio is 0.419 and 64.4% of instances fall below $1/2$, while above 1.171 the median is 1.000 and no instance falls below $1/2$, the regime in which every unclipped price already fits the budget.

Menu search (E3). One provider, $m = 3$ (80 instances) or $m = 4$ (40), λ a sorted sample of distinct integers in $[1, 8]$, q sorted uniform in $[0.5, 8]$, $B \sim U(0.3, 1.1) \max q$. Prices $p_t \in [0, \min\{q_t, B\}]$ on 26 (resp. 13) equally spaced points per tier; ratios evaluated on 2501 equally spaced types. The three instances with a gap above 10^{-3} have gaps 0.0015, 0.0279 and 0.0285. The largest is certified exactly: for $q = (1405/502,\ 3507/988,\ 2763/716,\ 4444/613)$, $\lambda = (1, 2, 5, 8)$ and $B = 3042/865$, the best menu the grid finds, $p = (1405/502,\ 2666/827,\ 3042/865,\ 0)$, has ratio 0.5186 against BC-VCG's 0.4896, a gap of 0.0290. Both ratios are exact rational infima: between adjacent breakpoints the menu's tier and the optimal tier are fixed, so the ratio is a fractional linear function of the type whose infimum is a one-sided limit at an endpoint. The smallest BC-VCG ratio on this family is 0.019.

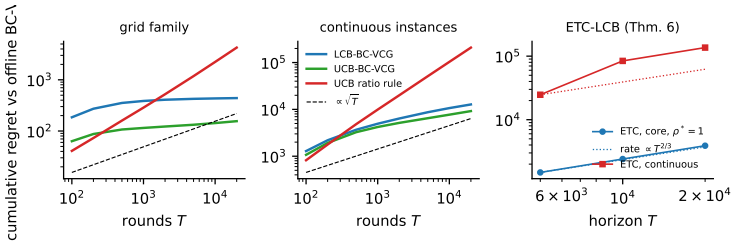

Figure 3: E5 and E8, cumulative regret against offline BC-VCG, log–log. Left and centre: 300 runs each. Right: ETC-LCB regret at each horizon; dotted lines are the $T^{2/3}$ rate of Theorem 6.

Vector costs (E4). The original family has $n = m = 2$, $q_{i,t} \in \{2, 3\}$, $B \in \{2, 3\}$, cost vectors in $\{1, 2\}^2$: 32 instances. The random family has $n, m \in \{2, 3\}$, $q_{i,t} \in \{1, \ldots, 8\}$, $B \in \{1, \ldots, 8\}$, cost vectors in $\{1, 2, 3\}^2$ or $\{1, 2\}^3$; 60 instances, seed 0. Table 5 reports the counts.

Table 5: E4: vector costs, exact enumeration. The original family (32 instances) and the random family (60 instances). "own-tier competition" is the rule that prices the winner against the best remaining score including its own tiers.

family	rule	DSIC (inst.)	IR	budget	worst ratio	avg. ratio
original, 32	own-tier competition	704 (32)	0	0	1	1.000
original, 32	BC-VCG	0 (0)	0	0	1	1.000
random, 60	own-tier competition	736899 (49)	638	0	0	0.937
random, 60	BC-VCG	0 (0)	0	0	1/3	0.997

Online (E5). $T = 20{,}000$, five seeds per instance, one probe per round on the round-robin schedule, $\delta = 0.05$, profiles redrawn each round as in E2. The estimated threshold rule uses $\overline{q}$ for the winner score and the price cap and the point estimate for the tier, as the original study did. Regret is measured against BC-VCG run on the true q at the same profile. Over the horizon LCB-BC-VCG realizes 0.907 of OPT on the continuous instances against 0.950 for the offline menu it learns, and 0.987 against 0.998 on the discrete family (each ratio is aggregate welfare over the horizon divided by aggregate OPT). Table 6 gives the horizon totals.

The threshold rule is linear even on the true qualities: its slope is 1.00 with interval $[1.00, 1.00]$ and its regret 210,241 at $T = 20{,}000$, against 211,662 for the estimated version. The linearity comes from the mechanism itself. Pricing from the point estimate instead of the certified bound inverts the comparison: BC-VCG on $\hat{q}$ has regret 1603, eight times below the certified menu's 12,795, at the price of 9634 loss-making rounds worth 18,825 of welfare, both summed over the 300 runs of the continuous family.

Price ladders (E6). 200 instances, seed 1: $n \in \{2, 3, 4\}$ providers, $m = 5$ tiers with multipliers $(1, 3, 5, 15, 25)$ times per-provider jitter, qualities concave in the multiplier, B between 0.2 and 1.5 times the top quality, and log-normal costs with median half the value of the cheapest tier; 400 sampled profiles each. 79,769 profiles in total.

Infimum (E7). The 101 two-tier instances of E2. The predicted floor is $\min_i \rho^\star(q_i, \lambda_i, B)$, the value Theorem 3 gives once every provider but the worst is made infeasible. The witness profile places the worst provider at $a + 10^{-15}$ in the increasing branch or at B/λ_2 in the decreasing branch, and every other provider at a cost above B, all in exact rational arithmetic; on it the measured ratio agrees with the closed form to within 10^{-15}. Median gap between the measured worst ratio and the closed form over the 60 non-trivial instances: 0.016, at most 0.688. All 60 lie in the branch where r increases, so BC-VCG approaches the closed form only as $\theta \to a$: the infimum is a limit, not attained by any profile.

Table 6: E5 (online families) and ETC-LCB of Theorem 6. Per family, 300 runs (60 instances $\times$ 5 seeds); ETC groups are described in the text. "welfare/OPT" is total realized welfare over total offline optimum at $T = 20{,}000$; "neg. / loss" counts rounds with negative realized welfare and their total magnitude; regret is cumulative against offline BC-VCG at $T = 5{,}000/10{,}000/20{,}000$ with a 95% bootstrap interval (by instance) at $T = 20{,}000$; the last column is a log–log slope of that regret over the three horizons. "threshold rule (true q)" is that rule with no estimation error and "BC-VCG on point estimate" is BC-VCG on $\hat{q}$ without confidence bounds, so the pair separates the mechanism gap from the estimation error.

family	mechanism	welfare/OPT	neg. / loss	regret vs BC-VCG at T			slope
				5,000	10,000	20,000	
discrete	offline BC-VCG (true q)	0.998	–	−−	−−	76	−−
	LCB-BC-VCG	0.987	0 / 0	385	411	441	0.09
	UCB-BC-VCG	0.994	51 / 15	141	157	179	0.17
	threshold rule (est.)	0.893	573 / 220	1283	2395	4362	0.93
	threshold rule (true q)	0.879	0 / 0	1217	2426	4854	1.00
	BC-VCG on point estimate	0.997	219 / 82	61	66	77	0.16
continuous	offline BC-VCG (true q)	0.950	–	−−	−−	14733	−−
	LCB-BC-VCG	0.907	0 / 0	8129	10284	12795	0.32
	UCB-BC-VCG	0.919	8267 / 11333	6127	7522	9263	0.29
	threshold rule (est.)	0.246	16525 / 45608	52229	105354	211662	1.02
	threshold rule (true q)	0.245	0 / 0	52836	105545	210241	1.00
	BC-VCG on point estimate	0.945	9634 / 18825	1141	1375	1603	0.23
ETC, core, $\rho^\star = 1$	ETC-LCB	0.909	0 / 0	1480	2406	3897	0.70
ETC, core, $\rho^\star < 1$	ETC-LCB	0.800	0 / 0	−1656	−3150	−5999	−−
ETC, continuous	ETC-LCB	0.564	109233 / 1238489	24695	84179	136311	0.69

Exploration (E8). ETC-LCB at $T = 5{,}000/10{,}000/20{,}000$, five seeds. Group 1: single-provider instances $q = (1, 3)$, $\lambda = (1, 2)$, $B \in \{1, 1.5, 2, 3\}$ plus one two-provider instance, where $\rho^\star$ is 1 for $B \geq 2$. Group 2: 20 continuous instances drawn from E2; four of them cannot fit the exploration phase nmN inside $T = 5{,}000$, at M/T up to 1.16, so that horizon rests on the other 16. On group 2 the exploration rounds number 109,233, losing 1,238,489 in welfare, from posting B for a tier whose quality may fall below its cost. The exploration length is $N = \lceil (T^2 \log(2nmT^2)/(\pi nm))^{1/3} \rceil$ with $\pi = \min_{i,t} \min\{1,\ B\lambda_{i,1}/(\lambda_{i,t} q_{i,1})\}$ the analytic value under $\theta_i \sim U(0, q_{i,1}/\lambda_{i,1})$, and $Q = \max_{i,t} q_{i,t}$. Slopes are fitted on the horizons where the benchmark is attainable: for group 2 the 16 instances present at all three horizons, 80 trajectories in all, with the bootstrap resampling single trajectories. Where $\rho^\star < 1$ the regret against offline BC-VCG is negative, from -1656 to -5999.

C Reproducibility

The anonymous supplementary material ships the implementation (model, mechanisms, exact verifier, closed form $\rho^\star$, menu search, online simulator) as a compiled extension module, together with one driver per experiment; the drivers write JSON records. The theorem checks run on the exported verifier and closed form: on random instances they compare the closed form (2) against the exact worst-case ratio of the BC-VCG menu (400 instances), test the "if and only if" of Theorem 2(c) (400), that no grid menu beats BC-VCG for two tiers (60 instances, 21^2 menus each), that random menus never exceed $\rho^\star$ (200 instances, 30 menus each), the incentive properties of BC-VCG and the budget overruns of VCG by exhaustive enumeration in both domains, and the multiplicative and additive bounds of Theorems 3 and 4 profile by profile, including the impossibility of Theorem 4(b) on a grid of menus.

D Further related work

Table 7 places the closest lines of work against ours. Every cell records what the cited paper states; "n/a" means the paper does not take a position on that axis.

Table 7: The closest lines of work on the axes this paper turns on.

work	private information	where the budget binds	feedback	objective / benchmark	guarantee
Singer (2010), Liu et al. (2024)	provider costs	total payment for a set	n/a	buyer value vs cost-feasible optimum	constant-factor approximation
Amanatidis et al. (2025a), Chan & Chen (2014)	provider costs	one payment budget; assumed to afford any single agent	n/a	concave value vs cost-feasible optimum	constant-factor approximation to a concave value
Neogi et al. (2025)	vector costs	total payment	n/a	buyer value vs cost-feasible optimum	no constant-factor approximation
Le (2018)	valuations	on the valuations, truncated before VCG	n/a	Pareto optimality, not an approximation	unique IR, IC and Pareto optimal rule for one good
Babaioff et al. (2009)	per-arm values	n/a	clicks	regret vs the best arm	exploration must be report-independent
Badanidiyuru et al. (2012), Amanatidis et al. (2025b)	valuations	total payment over the horizon	arrivals	regret vs the offline benchmark	posted prices are competitive
Theorems 2 and 3 (ours)	one base cost per provider	a single posted contract price	n/a	net welfare vs budget-feasible OPT	exact instance-dependent ratio with matching lower bound (two tiers); additive bound for any m
Theorem 5 (ours)	base costs; qualities unknown to the platform, learned	price ceiling each round	probes	regret vs the clipped menu on the true qualities	$\tilde{O}(\sqrt{T})$ regret under a margin condition, no loss-making award on $\mathcal{G}$
Theorem 6 (ours)	base costs; qualities unknown to the platform, learned	price ceiling each round	awarded contracts	regret vs the clipped menu on the true qualities	$T^{2/3}$ regret; fixed-type acceptance dominance

Budget-feasible mechanism design has since moved to subadditive valuations through a compensation-design reduction (Anagnostides et al., 2026a;b); welfare maximization meets barriers under no-regret learning even without a budget (Anagnostides et al., 2024). Theorem 5(c) uses the low-noise condition of (Audibert & Tsybakov, 2007) with the Hoeffding half-width (Hoeffding, 1963) and standard confidence intervals (Kocsis & Szepesvári, 2006; Shalev-Shwartz, 2011; Lattimore & Szepesvári, 2019); variance-adaptive (Maurer & Pontil, 2009) and time-uniform (Howard et al., 2021; Waudby-Smith & Ramdas, 2024) bounds would sharpen the widths in the probe protocol at the same rate. Language-model service markets appear as reverse auctions over routing (Chen et al., 2026), negotiation benchmarks (Müller & Müller, 2026), and classical cooperation mechanisms played by language-model agents (Tewolde et al., 2026); the tiered prices such providers publish are one observable proxy for the multipliers, though they do not show that costs scale this way. A computational and learning-side backdrop comes from learning in auction games (Bichler et al., 2023; Pieroth et al., 2025), equilibrium computation (Conitzer & Sandholm, 2002; Zhang et al., 2023), online pricing for constrained allocation (Zilberstein et al., 2026), regret minimization for single-dimensional contract design (Poiani et al., 2026), mechanism design under type uncertainty (Hyafil & Boutilier, 2004), exploration granted before regret accrues (Hou et al., 2026), and strategic principals facing learning agents (Rubinstein & Zhao, 2024). Our computations are exact rational enumerations over small discrete instances, which is also how we show what such an enumeration can miss.