INCENTIVE COMPATIBILITY MECHANISM DESIGN FOR THE AUTO-BIDDING META-GAME

Anonymous authors
Paper under double-blind review

ABSTRACT

In auto-bidding, advertisers no longer bid per impression but set a high-level target (e.g. tCPA), competing in an *Auto-bidding Meta-Game* (AMG) over reported targets. Running a truthful per-impression auction underneath does *not* make this meta-game truthful—an advertiser can still gain by misreporting its target. Within the platform-supported uniform-bidding interface, we formalize the AMG as a Myerson single-parameter environment and design a *deployable* Dominant-Strategy Incentive Compatible (DSIC) mechanism for it: with exact allocation information under first-price, truthfulness is restored by an end-of-day rebate that returns the advertiser's information rent and leaves the auction untouched. The truthful profile is also welfare-optimal under first-price. For other pricing rules the same payment construction is DSIC whenever the induced allocation is monotone; second-price numerical audits expose allocation nonmonotonicity and integration-sensitive deviation gains, so they do not validate that condition, though welfare optimality does not carry over. Deployment estimates the allocation curve G offline; its signed settlement implements an approximately DSIC payment, not the exact rebate in general. We prove that uniform allocation error yields approximate DSIC. Under an additional uniform quadratic-growth condition, damped best responses track a neighborhood of truth; convergence to truth requires vanishing estimation error, not merely repeated estimation. We evaluate these theoretical predictions on a synthetic simulator across market sizes and value/conversion-rate distributions. An online A/B test on a deployed tCPA product gives an encouraging behavioral signal: a significant market-wide gain in cost, and significant gains in reported targets for the cohort it was designed to move.

1 INTRODUCTION

Auto-bidding has become a critical tool for advertisers to optimize bids across digital impressions, maximizing conversions subject to high-level constraints. Under Target-CPA (tCPA), for example, the advertiser specifies a target cost per acquisition and the platform algorithm sets per-impression bids to maximize conversions while keeping the average cost per conversion below that target. The conversions an advertiser obtains and the cost it pays are thus determined jointly by the auto-bidding algorithm and the underlying per-impression auction. An advertiser's only strategic lever is the high-level target it reports; advertisers therefore effectively compete in a *meta-game* over reported targets, which we call the *Auto-bidding Meta-Game* (AMG).

Classical per-impression auctions such as VCG-style truthful keyword auctions Aggarwal et al. (2006) and GSP Edelman et al. (2007); Varian (2007) presuppose direct per-impression bidding—an assumption auto-bidding breaks, since advertisers control only high-level parameters. Running a classical DSIC auction underneath does *not* make the AMG itself DSIC: prior work shows that advertisers can generically profit by misreporting their target in this report-to-agent setting Kolumbus & Nisan (2022b); Li & Tang (2022). A related line treats auto-bidders as *value maximizers* under ROI/budget constraints Aggarwal et al. (2019); Balseiro et al. (2021b); Deng et al. (2021), but, like the classical auctions above, analyzes at the per-impression level. We instead treat the reported target itself as the primitive strategic object, and restoring truthfulness of the meta-game over reports is the central problem we study.

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.

In this work we take the AMG as the object of design directly: abstracting away the per-impression auction, we study the relationship among the reported target, the allocation volume, and the payment, and seek a mechanism under which truthful reporting is dominant—the truthful report then coinciding with the advertiser's private per-conversion value, a single scalar type. We take *first-price* as the underlying pricing rule, the semantics a binding tCPA target induces (the advertiser pays its own declared price per conversion), under which the truthful mechanism is realized by an allocation-preserving rebate. The Myerson construction also applies to a second-price underlying auction conditional on monotonicity of its induced allocation. Our single-deviation audits expose numerical nonmonotonicity and integration sensitivity in the second-price implementation rather than validating its truthfulness; welfare optimality remains specific to first-price. A full discussion of related work is deferred to Appendix J.

1.1 Our Results

Our novelty is not that a truthful report equals value—for a single quasi-linear advertiser that is classical Myerson—but *where* and under *what information* truthfulness is secured: unlike prior work that *controls* the allocation curve, assuming the value v known, we treat the allocation G as exogenous and *unknown* and intervene only through a single end-of-day transfer, shifting the difficulty to *estimating* G from logs. We make three contributions.

- **A welfare-optimal DSIC mechanism for the AMG.** Within the platform-supported uniform-bidding strategy class, we formalize the AMG and identify it as a *Myerson single-parameter environment*, importing the full payment identity rather than a monotonicity condition alone. Instantiated on the efficient allocation, it is DSIC and, under first-price, realized by an end-of-day rebate that leaves the underlying auction untouched. Truthful reporting is welfare-optimal under first-price; for any other pricing rule, including second-price, the DSIC conclusion is conditional on monotonicity of the rule's induced allocation.
- **An estimator with conditional robustness.** Uniform error ϵ yields $O(\epsilon)$ approximate DSIC. Report displacement and the resulting $O(\sqrt{\epsilon})$ first-price welfare bound additionally require uniform quadratic growth. Damped best responses track an error-dependent neighborhood and converge to truth when the error vanishes. We separate iid replay error from grid interpolation error (Section 4.1).
- **Validation on synthetic and production evidence.** On a synthetic simulator with known ground truth we test the theoretical predictions across market sizes and distributions, including a single-deviation audit that identifies unresolved numerical limitations of the second-price implementation. An online A/B test on a deployed tCPA product provides a behavioral signal consistent with the predicted response to a subsidy, delivered through a discount schedule derived from our payment formula rather than a direct instantiation of the rebate (Section 5.1).

2 Preliminaries

2.1 Auto-bidding and the AMG

We consider N advertisers on a Target-CPA (tCPA) auto-bidding product; advertiser i holds a private per-conversion value v_i and reports a target τ_i. It does not bid per impression: it delegates to a platform algorithm and observes only aggregate conversions and cost. Under tCPA the constraint bounds the *average* cost per conversion—over won impressions yielding G_i conversions at total payment P_i,

$$\frac{P_i}{G_i} \;\le\; \tau_i, \tag{1}$$

and the algorithm bids to maximize conversions subject to it. We isolate this constraint and assume budgets are non-binding—a budget merely adds a spending ceiling, whereas the CPA constraint is what couples the report to the bids—and model advertisers as utility maximizers (Appendix I).

Uniform bidding as the platform interface. We restrict the platform algorithm to the *uniform-bidding* policy class used by the tCPA interface: it bids on every impression j in proportion to the

predicted conversion probability, scaled by one global coefficient,

$$b_{i,j} \;=\; \alpha_i \cdot \tau_i \cdot \mathrm{pCVR}_{i,j}, \tag{2}$$

where α_i is the impression-invariant pacing coefficient tuned to meet (1). Ranking multiplies the bid by the click-through rate $\mathrm{pCTR}_{i,j}$, so we collect both prediction factors into a single *per-impression conversion rate*

$$q_{i,j} \;:=\; \mathrm{pCTR}_{i,j} \cdot \mathrm{pCVR}_{i,j}. \tag{3}$$

The ranking score is then $\alpha_i \tau_i \, q_{i,j}$, and the algorithm bids in proportion to $\tau_i \, q_{i,j}$, which equals the per-impression value $v_i \, q_{i,j}$ under truthful reporting when $\alpha_i = 1$. We do not claim that uniform bidding is an advertiser's unrestricted optimum over arbitrary per-impression bid vectors. Rather, it is the platform-provided policy interface: the advertiser chooses only the scalar target τ_i, while the platform sets the impression-invariant coefficient α_i. All optimality and incentive statements below are relative to this restricted strategy class. The value of α_i is pinned down by the underlying pricing rule, which we turn to next.

Pricing rule: first-price as the primary setting. The AMG layers on a per-impression auction whose *allocation* is highest-bid-wins, $\arg\max_i \alpha_i \tau_i \, q_{i,j}$, common to all standard rules; they differ only in price. We take *first-price* as the primary setting, the rule the tCPA product is closest to: when the constraint binds the advertiser pays its declared τ_i per conversion, so the natural payment $P_i^{\mathrm{nat}} = \tau_i G_i$ (Definition 2) follows from the constraint, and the effective bid coefficient equals the report, with $\alpha_i \equiv 1$ in (2). A second-price rule can instead require $\alpha_i > 1$ when the CPA target binds; Section 3.2 treats that case.

Assumption 1 (Regular allocation)**.** Fix competitors' reports $\vec{\tau}_{-i}$. The bounded allocation $G_i(\tau_i; \vec{\tau}_{-i})$ is non-negative and non-decreasing on $[0, \bar{v}]$. Under first-price this follows from highest-score-wins with $\alpha_i \equiv 1$ and report-independent tie-breaking. Neither continuity nor strict increase follows: finite impression populations yield step functions. We adopt a no-allocation convention at zero report, so $G_i(0; \cdot) = 0$. For other pricing rules monotonicity is a separate condition on the complete pacing-induced allocation; Appendix C gives finite-instance evidence only. Quantitative growth assumptions are imposed separately where needed.

Aggregate model. Fixing rivals' reports $\vec{\tau}_{-i}$, the mechanism acts on advertiser i's expected total conversions $G_i(\tau_i)$ and payment $P_i(\tau_i)$; the per-impression value is $v_{i,j} = v_i \, q_{i,j}$ (3) and truthful reporting means $\tau_i = v_i$. We abbreviate $G_i(\tau_i), P_i(\tau_i)$ when $\vec{\tau}_{-i}$ is fixed.

Definition 1 (Auto-bidding Meta-Game)**.** *The **Auto-bidding Meta-Game (AMG)** is the N-player game in which:*

- ***Type and action.*** *Advertiser i has a private per-conversion value v_i (its type) and chooses a single action, the reported tCPA $\tau_i \in [0, \bar{v}]$, where $\bar{v}$ is a commonly known upper bound on the value support (so $v_i \le \bar{v}$ for every i).*
- ***Mechanism.*** *Given the profile $\vec{\tau}$, the platform's per-impression auction maps the uniform bids $b_{i,j} = \alpha_i \tau_i \, pCVR_{i,j}$ (2) to an allocation $G_i(\tau_i; \vec{\tau}_{-i})$ (set by its ranking rule) and a payment $P_i(\tau_i; \vec{\tau}_{-i})$ (set by its pricing rule).*
- ***Payoff.*** *Advertiser i's payoff $U_i(\tau_i, v_i)$ depends on the value obtained $v_i G_i$ and the payment P_i; we specify it as the quasi-linear utility of a utility maximizer in Section 2.2.*

Definition 2 (Natural AMG)**.** *The **natural AMG** is the instance in which the platform does not intervene: when the tCPA constraint* (1) *binds, the average cost per conversion equals τ_i, so*

$$P_i^{\mathrm{nat}}(\tau_i; \vec{\tau}_{-i}) \;=\; \tau_i \, G_i(\tau_i; \vec{\tau}_{-i})$$

(approximately, since pacing meets the constraint only up to slack). This is the status-quo baseline our mechanism modifies; Section 2.2 shows it is not DSIC.

Remark 1 (The settlement is our only design lever). The platform could instead redesign the ranking (reserves, boosts) to change G; we deliberately do not, intervening only through an allocation-preserving *transfer*. The construction can be layered on an underlying auction whenever its induced allocation is monotone in the advertiser's report; our unconditional auction-specific results focus on first-price. Appendix I records why the action space runs to $\bar{v}$ rather than v_i.

2.2 ADVERTISER MODEL

We model each advertiser as a **utility maximizer** (UM): it maximizes net surplus, the total value of the conversions it wins minus what it pays,

$$U_i(\tau_i, v_i) \;=\; v_i\, G_i(\tau_i) - P_i(\tau_i); \tag{4}$$

Appendix I discusses this modeling choice. Each advertiser selects a reported target $\tau_i \in [0, \bar{v}]$ to maximize (4); a best response, when it exists, satisfies

$$\tau_i^* \;\in\; \arg\max_{\tau_i \in [0,\bar{v}]} U_i(\tau_i, v_i). \tag{5}$$

Why the natural AMG is generally not DSIC. Under the natural payment $P_i^{\mathrm{nat}}(\tau_i) = \tau_i G_i(\tau_i)$ (Definition 2), the utility (4) is $U_i(\tau_i, v_i) = G_i(\tau_i)(v_i - \tau_i)$. Truthful $\tau_i = v_i$ yields zero utility, whereas any lower report with $G_i(\tau_i) > 0$ yields strictly positive utility. Thus shading is strictly profitable whenever such a report exists; without it, a strict incentive is not implied. Our goal is a payment rule P—replacing P^{nat}—under which truthful reporting is a dominant strategy.

Definition 3. *A mechanism (G, P) is **Dominant-Strategy Incentive Compatible (DSIC)** if, for every advertiser i with true valuation v_i, and for any $\vec{\tau}_{-i}$:*

$$v_i \cdot G_i(v_i;\, \vec{\tau}_{-i}) - P_i(v_i) \;\geq\; v_i \cdot G_i(\tau_i';\, \vec{\tau}_{-i}) - P_i(\tau_i') \qquad \forall\, \tau_i' \neq v_i.$$

3 MECHANISM DESIGN

3.1 THE DSIC FRAMEWORK: SINGLE-PARAMETER STRUCTURE AND MYERSON PAYMENTS

Fixing rivals' reports $\vec{\tau}_{-i}$, advertiser i chooses a scalar report τ_i, and its utility $v_i G_i(\tau_i; \vec{\tau}_{-i}) - P_i(\tau_i; \vec{\tau}_{-i})$ is linear in its private type v_i, with allocation and payment depending only on reports. This gives the AMG its single-parameter linear structure; Assumption 1 separately supplies allocation monotonicity, allowing us to apply Myerson's characterization Myerson (1981); Hartline (2017) and realize its payment through an allocation-preserving rebate.

Theorem 1 states this characterization at the meta-game level. Appendix A gives a derivative-free proof covering monotone step allocations.

Theorem 1 (Myerson at the meta-game level)**.** *Among AMG mechanisms with $P_i(0; \vec{\tau}_{-i}) = 0$ for every rival profile, a mechanism $M = (G, P)$ is DSIC (Definition 3) if and only if:*

- *the allocation rule $G_i(\tau_i)$ is monotonically non-decreasing in the reported tCPA τ_i, and*
- *the payment rule follows the integral form, normalized so that $P_i(0) = 0$,*

$$P_i(\tau_i) \;=\; \tau_i \cdot G_i(\tau_i) - \int_0^{\tau_i} G_i(z)\, dz.$$

First-price ranking implies non-decrease at fixed rivals; other pricing rules require monotonicity of the complete pacing-induced allocation. The payment, however, does not: under the natural CPA-binding payment one has $P_i(\tau_i) \approx \tau_i G_i(\tau_i)$, so to enforce the Myerson form we must estimate G_i in order to compute the integral $\int_0^{\tau_i} G_i(z)\, dz$.

We first assume that G_i can be estimated accurately and give the implementation of the corrected payment. The mechanism does not touch the running auction; it acts only as an end-of-day transfer:

1. **Run the auction unchanged.** Record auction outcomes and gross spend. In the expected-quantity model, $C_i = \tau_i G_i(\tau_i)$ under first-price; realized conversions and spend fluctuate and are not identified with their expectations.
2. **Compute the target payment.** Set $\widehat{P}_i(\tau_i) = \tau_i \widehat{G}_i(\tau_i) - \int_0^{\tau_i} \widehat{G}_i(z)\, dz$.
3. **Apply a signed settlement.** The correction from gross expected spend is

$$S_i := C_i - \widehat{P}_i(\tau_i) = \tau_i\big(G_i(\tau_i) - \widehat{G}_i(\tau_i)\big) + \int_0^{\tau_i} \widehat{G}_i(z)\, dz. \tag{6}$$

Exact estimation recovers the non-negative rebate $S_i = \int_0^{\tau_i} G_i$; the stronger one-sided condition $0 \le \widehat{G}_i \le G_i$ also ensures $S_i \ge 0$. These statements concern expected settlement; random realized billing requires its own accounting convention.

The settlement must not perturb the current allocation: it is applied after the auction and excluded from its pacing feedback (Appendix A).

3.2 Welfare optimality: first-price versus second-price

We compare equilibria: under the natural payment advertisers shade to a Nash equilibrium $\vec{\tau}^{\mathrm{nat}}$, whereas under the exact DSIC mechanism truthful reporting is a weakly dominant-strategy equilibrium, $\vec{\tau}^* = \vec{v}$. Writing $G_i^\bullet$ for the equilibrium allocations, define

$$W^\bullet := \sum_i v_i G_i^\bullet, \qquad \mathrm{Rev}^\bullet := \sum_i P_i^\bullet, \qquad U^\bullet := W^\bullet - \mathrm{Rev}^\bullet, \qquad \bullet \in \{\mathrm{nat}, *\}.$$

Throughout, $W^{\mathrm{opt}} := \sum_j \max_i v_i\, q_{i,j}$ denotes the first-best (value-optimal) welfare, which upper-bounds the welfare of any allocation.

Theorem 2 (Welfare optimality and individual rationality)**.** *Under the exact DSIC mechanism truthful reporting is a weakly dominant-strategy equilibrium, $\vec{\tau}^* = \vec{v}$, and:*

1. (Optimal social welfare, under first-price.) *$W^* = W^{\mathrm{opt}} \ge W^{\mathrm{nat}}$, with strict inequality if shading awards impressions to strictly lower-value winners with positive probability. (Under second-price this fails; see below.)*
2. (Individual rationality, any monotone induced allocation.) *Each advertiser obtains non-negative utility, $U_i^* = \int_0^{v_i} G_i(z; \vec{v}_{-i})\, dz \ge 0$, directly from the Myerson payment. Thus this part extends to any pricing rule satisfying Assumption 1.*

The proof (weak dominance from Myerson's identity, welfare optimality at the truthful profile from the first-price value ranking, and IR from the payment identity) is in Appendix A. Uniqueness requires additional assumptions and is not asserted here. The revenue comparison is more delicate—not even sign-definite in general—and is taken up separately in Section 3.3.

Welfare optimality is not robust to second-price. Welfare optimality (1) relies on the first-price identity $\alpha_i \equiv 1$. Under a second-price rule the winner pays less than its bid, so pacing to the same CPA target inflates the bid by a multiplier α_i Conitzer et al. (2022b),

$$b_{i,j} \;=\; \alpha_i\, \tau_i\, \mathrm{pCVR}_{i,j}, \qquad \alpha_i = \alpha_i(\tau_i; \vec{\tau}_{-i}) \text{ solves } \tfrac{C_i}{G_i} = \tau_i, \tag{7}$$

and, when a binding solution exists, the multiplier can exceed one and differ across advertisers facing different competition. Heterogeneous multipliers can distort the truthful ranking $\arg\max_i \alpha_i v_i\, q_{i,j}$, but do not necessarily change a winner (Appendix C measures the effect). In finite markets the CPA equality need not be attainable; a pacing selection rule must then be specified. This is a positive argument for the first-price primary setting, not a defect: the rebate is built on the first-price identity $\alpha_i \equiv 1$, under which report equals coefficient and $C_i = \tau_i G_i(\tau_i) = P_i^{\mathrm{nat}}(\tau_i)$ exactly. For a different pricing rule, Theorem 1 gives DSIC through the Myerson payment on its induced allocation G_i^r *provided that* G_i^r is monotone in i's own report. We do not claim a general proof of this property for every second-price pacing implementation; Appendix C reports a numerical audit, including nonmonotonicity and unresolved integration sensitivity, for the implementation used in our simulator.

3.3 Revenue: natural versus DSIC

Unlike welfare, the effect of switching from the natural payment to the DSIC mechanism on *platform revenue* has no clean sign: the efficient allocation raises the pool of value extracted, but the Myerson payment returns each advertiser's information rent $\int_0^{v_i} G_i$, and whether the net change is positive depends on the whole value profile, with no clean closed form. We therefore characterize the change *empirically* in Section 5.1. The relevant quantities are computable from historical logs—$W^{\mathrm{nat}}, \mathrm{Rev}^{\mathrm{nat}}$ by replaying observed bids and W^*, Rev^* by replaying at the truthful profile—so the welfare and revenue change is measurable on the simulator against ground truth. Throughout, the reported $\Delta\mathrm{Rev} = \mathrm{Rev}^* - \mathrm{Rev}^{\mathrm{nat}}$ is in simulator value units summed over the market's M impressions, and advertiser *cost* likewise denotes gross spend in those units.

4 ESTIMATING G FROM LOGS

The mechanism of Section 3 needs only one object it does not already possess: the allocation curve $G_i(\cdot)$, from which the Myerson payment and its signed end-of-day settlement are computed. Everything else—the highest-bid allocation, the reported target τ_i, the first-price identity $\alpha_i \equiv 1$—is inherited from the live auction. This section is therefore about one deployment question: *how to estimate G from data the platform already logs*, and how the mechanism degrades when that estimate is imperfect.

We keep the notation of the model (Section 2.1). The allocation $G_i(\tau_i; \vec{\tau}_{-i}) = \mathbb{E}[\sum_j q_{i,j}\, \mathbb{I}\{i \text{ wins } j\}]$ ($\mathbb{I}\{\cdot\}$ the indicator function, equal to 1 if the event holds and 0 otherwise) is measured in *expected conversions*, where $q_{i,j} = \mathrm{pCTR}_{i,j}\, \mathrm{pCVR}_{i,j}$ (3) is the per-impression conversion rate. In a production stack the auction ranks by $\mathrm{ECPM} = \mathrm{bid} \cdot \mathrm{pCTR}$, so both prediction factors—the click-through rate that scales the bid into ECPM and the per-click conversion rate inside the bid—enter the ranking; the quantity credited to G_i when i wins is the same product $q_{i,j}$. As in the primary setting, we take the underlying pricing to be first-price, so the reported target is the effective bid coefficient ($\alpha_i \equiv 1$).

4.1 THE OFFLINE-SIMULATOR ESTIMATOR

We estimate G from historical auction data by offline replay against fixed rivals, then bound how the mechanism's guarantees degrade when that estimate is imperfect. The historical data records, for every exposure, the full candidate queue of advertising units (winners and losers alike). Concretely, the candidate queue for an exposure has the form

$$Q = \{(\mathrm{bid}_1, \tau_1, \alpha_1, \mathrm{pCTR}_1, \mathrm{pCVR}_1, \mathrm{rank}_1), \ldots, (\mathrm{bid}_N, \tau_N, \alpha_N, \mathrm{pCTR}_N, \mathrm{pCVR}_N, \mathrm{rank}_N)\}.$$

The advertising system sorts candidates in descending order of ECPM, and rank_i records i's position in that order.

For report r, replay holds rivals fixed and compares $rq_{i,j}$ with the highest competing score; after that score is precomputed, each exposure requires one comparison. Aggregate replay values are fitted by weighted least-squares isotonic regression with positive weights, then linearly interpolated between grid nodes. This uses allocation values, not derivatives. A full-domain guarantee requires a deterministic grid $0 = r_0 < \cdots < r_K = \bar{v}$, not merely multiplicative reports near the current target. The following results separate node estimation from interpolation; proofs are in Appendix A.

Proposition 1 (Weighted-threshold replay bound)**.** *Fix rivals and draw n iid impressions. Let $q \in [0, Q]$ and X be the highest competing score divided by q; set the contribution to zero when $q = 0$. For $f_r = q\mathbf{1}\{X \le r\}$, let $g(r) = \mathbb{E}f_r$, $g_n(r) = n^{-1}\sum_{j=1}^{n} f_r(j)$, and $G = Mg$ for a market of M impressions. With probability at least $1 - \delta$,*

$$\sup_r |g_n(r) - g(r)| \le a_n(\delta) := \frac{4Q}{\sqrt{n}} + Q\sqrt{\frac{\log(1/\delta)}{2n}}.$$

Fitting replay nodes $y_k = Mg_n(r_k)$ as in Remark 2 yields $\epsilon \le Ma_n(\delta) + Lh$ if the total G is L-Lipschitz. The same bound holds with strict threshold comparisons, provided the convention is consistent in truth and replay. A no-allocation rule at $r = 0$ is implemented by setting both functions there to zero.

For fixed M, Q, L and a mesh $h = O(n^{-1/2})$, the full-domain error is $O(\sqrt{\log(1/\delta)/n})$ for $\delta \le 1/2$; a fixed mesh leaves an interpolation term. Without Lipschitz regularity retain $\omega_G(h)$, not this rate. Ties may be resolved by fixed priority: the threshold comparison is then closed or open according to that priority (see the proof). These iid population guarantees do not automatically apply to subsampling a reused finite log or to drifting traffic.

4.2 ROBUSTNESS TO ESTIMATION ERROR

The estimator of Section 4.1 yields $\widehat{G}_i \neq G_i$ in general; this subsection bounds how far that imperfection can move the mechanism's guarantees. When $\widehat{G}_i$ has uniform error $\sup_{\tau \in [0,\bar{v}]} |G_i(\tau) - \widehat{G}_i(\tau)| \le \epsilon$, payment and incentive errors are bounded without slope assumptions; report displacement

and the resulting welfare control require additional quadratic growth. The estimated payment is the Myerson form on $\widehat{G}_i$, $\widehat{P}_i(\tau) = \tau\widehat{G}_i(\tau) - \int_0^\tau \widehat{G}_i$.

Theorem 3 (Robustness of the estimated mechanism)**.** *Fix any $\vec{\tau}_{-i}$ and let $G(\tau) \equiv G_i(\tau; \vec{\tau}_{-i})$, $\sup_\tau |G - \widehat{G}| \le \epsilon$. Then:*

1. (Payment/utility deviation.) *$|P(\tau) - \widehat{P}(\tau)| \le 2\epsilon\tau$ and $|U(v,v) - \widehat{U}(v,v)| \le 2\epsilon v$ for all $\tau, v \in [0, \bar{v}]$.*

2. (Approximate IC.) *The gain from any single deviation, $\Delta_{IC}(v) := \sup_\tau\{vG(\tau) - \widehat{P}(\tau)\} - \{vG(v) - \widehat{P}(v)\}$, is at most $4\epsilon\bar{v}$; the mechanism is $4\epsilon\bar{v}$-approximately DSIC. This holds for* every *fixed $\vec{\tau}_{-i}$, so it applies verbatim day by day as rivals adapt.*

3. (Report displacement, under quadratic growth.) *Write $U_v(r) = vG(r) - P(r)$ and $D_v(r) = U_v(v) - U_v(r)$. Suppose, separately, that $D_v(r) \ge \mu(r - v)^2/2$ for every candidate report $r \in [0, \bar{v}]$, with $\mu > 0$. Any maximizer r^* of $\widehat{U}_v(r) = vG(r) - \widehat{P}(r)$, if it exists, satisfies $|r^* - v| \le \sqrt{8\epsilon\bar{v}/\mu}$. A sufficient growth condition is absolute continuity of G on every segment joining v to r, with $G' \ge \mu$ almost everywhere on that entire segment. If additionally the total G is L-Lipschitz, the advertiser's own value change at fixed rivals is at most $\bar{v}L|r^* - v|$. This is not an aggregate welfare bound.*

Proofs are in Appendix A (triangle inequality; DSIC of the true mechanism at fixed $\vec{\tau}_{-i}$; a separate quadratic-growth condition, not global strong concavity).

Part (3) concerns one advertiser at fixed rivals. Aggregate welfare instead follows directly from any report vector's distance to truth, without requiring a limiting profile or equilibrium.

Corollary 4 (Aggregate welfare deviation)**.** *Under first-price with $q_{i,j} \in [0, 1]$, every report vector $\vec{\tau}$ satisfies*

$$W^{\mathrm{opt}} - W(\vec{\tau}) \le 2M\|\vec{\tau} - \vec{v}\|_\infty.$$

A vector within Δ of truth therefore loses at most $2M\Delta$. Choosing $\Delta = \sqrt{8\epsilon\bar{v}/\mu}$ gives $O(\sqrt{\epsilon})$ for fixed $M, \bar{v}, \mu$, whenever the displacement conditions hold.

Proof in Appendix A.

Sample complexity and dynamics. Proposition 1 gives a weighted empirical-process bound under iid sampling, fixed rivals and a refining mesh, not a direct application of ordinary DKW. For adaptive rivals over T days, condition on history and draw fresh iid logs from each day's fixed-rival distribution. Use confidence $\delta/(NT)$ and union bound over N advertisers and T days. Reused logs and drift need separate analysis; finite-horizon probability bounds do not establish infinite-horizon convergence.

Proposition 2 (Tracking and conditional convergence)**.** *At each day t and the actual rival profile, assume full-domain error $\|\widehat{G}_i^{(t)} - G_i^{(t)}\|_\infty \le \epsilon_t$ and quadratic growth $D_{i,v_i}^{(t)}(r) \ge \mu(r - v_i)^2/2$ for all advertisers and reports, with the same $\mu > 0$. Suppose maximizers $\widehat{r}_i^{(t)}$ of $v_i G_i^{(t)}(r) - \widehat{P}_i^{(t)}(r)$ exist and $\vec{\tau}^{(t+1)} = (1-\lambda)\vec{\tau}^{(t)} + \lambda\widehat{\vec{r}}^{(t)}$, with fixed $\lambda \in (0, 1]$. Define $d_t = \|\vec{\tau}^{(t)} - \vec{v}\|_\infty$ and $b_t = \sqrt{8\epsilon_t\bar{v}/\mu}$. Then*

$$d_{t+1} \le (1-\lambda)d_t + \lambda b_t, \qquad d_t \le (1-\lambda)^t d_0 + \lambda\sum_{s=0}^{t-1}(1-\lambda)^{t-1-s}b_s.$$

For $\epsilon_t \le \epsilon$, $\limsup_t d_t \le \sqrt{8\epsilon\bar{v}/\mu}$; a limit need not exist. If $\epsilon_t \to 0$, then $d_t \to 0$ from every initial profile in the report domain.

The proof is in Appendix A. Stabilizing or repeatedly recomputing an estimate does not imply truth convergence. Monotonicity alone does not make the exact best-response map single-valued; quadratic growth is the extra condition. E3 is a finite-grid proxy-utility diagnostic, not a verification of these assumptions. Separately, $0 \le \widehat{G} \le G$ preserves truthful individual rationality since $\widehat{P}(v) \le v\widehat{G}(v) \le vG(v)$; otherwise truthful utility is at least $-2\epsilon v$.

Table 1: E1, first-price: bid-shade ratio $s_i = \tau_i/v_i$ among active advertisers when they best-respond to the natural payment, across value/conversion-rate (q) distributions ($N = 10$, $M = 2000$, 60 markets each, active-only). Shading is pervasive in every configuration—the median active advertiser reports 75–83% of value, with a left tail down to p10 ≈ 0.53.

v dist	q dist	$\mathrm{med}(s_i)$	p10/p90
uniform	uniform	0.825	0.73 / 0.90
lognormal	uniform	0.825	0.68 / 0.90
exponential	Beta(2,5)	0.750	0.53 / 0.87
uniform	Beta(2,5)	0.750	0.67 / 0.87

5 Experiments

5.1 Evaluation questions

We organize the evaluation around three questions a deployment decision would hinge on, and address them on the synthetic simulator; Appendix H complements this with an initial signal from a production A/B test, though a full production replication of Q1–Q3 remains for future work:

Q1. Is bid-shading real and severe under the current natural payment $P^{\mathrm{nat}}(\tau) = \tau G(\tau)$—how far below value do advertisers report? (E1)

Q2. How do measured deviation gains change with replay sample size? (E2)

Q3. How do finite-run trajectories behave across tested starts and scales? (E3)

Appendix B compares self-selected finite-grid search profiles with the truthful benchmark. Exact-G runs are near truthful; estimated runs improve over the tested sample sizes, without establishing an asymptotic limit. Appendix C exposes numerical nonmonotonicity and integration-sensitive deviation gains in the second-price implementation, without establishing its incentive compatibility. Appendix D varies scale ($N = 5$ to 100) and distributions: exact-mechanism welfare at truth remains optimal, while revenue gains vary. Appendix F studies a shared noisy pCVR signal, not incentive compatibility against true-conversion utility when it differs from that signal.

5.2 Setup

We use a synthetic AMG simulator with known values v_i. Each market draws advertiser values and impression conversion rates $q_{i,j}$; independent markets use fixed seeds. Value laws are uniform, log-normal, exponential, or Pareto, and q laws are uniform or $\mathrm{Beta}(2, 5)$ (default: log-normal values and uniform q). Metrics include shade $s_i = \tau_i/v_i$, welfare and revenue relative to the natural baseline, restricted-grid deviation gain, and active-advertiser mean squared report error. E3 fixes the active set at truthful allocation. These finite-log diagnostics are not full-domain error or IC certificates. Appendix G details the metrics.

5.3 Experiments

E1. Shading in the natural game (Q1). *Answer (Q1): yes—shading is real and severe.* When advertisers best-respond to the natural payment $P^{\mathrm{nat}} = \tau G$ ($N = 10$, $M = 2000$, 60 markets/config), the median active shade ratio $s_i = \tau_i/v_i$ is well below one across all value/q distributions (Table 1). Under-reporting is the norm in the current system—the incentive our mechanism removes.

E2. DSIC under estimated G (Q2). *Answer (Q2): the measured restricted-grid deviation gain decreases over the tested sample sizes.* We subsample n impressions without replacement from a fixed synthetic log ($N = 6$, $M = 2 \times 10^4$). Replay and deviations use 50 reports in $[0.02v_i, v_i]$, with $(0, 0)$ added for isotonic fitting; the fit is linearly interpolated and clipped outside its fitted range. Error is evaluated only at the 50 nodes, and payments use 400-point trapezoidal integration. These diagnostics neither cover over-reports up to $\bar{v}$ nor certify the full-domain error or IC supremum of Theorem 3. Table 2 shows declining measured error and deviation gain, but does not establish an asymptotic rate or validate the theorem's bound. At $n = 10^4$, mean deviation gain is under 2% of

Table 2: E2: nodewise error and restricted-grid deviation gain versus replay sample size ($N = 6$, $M = 2 \times 10^4$, 20 markets, log-normal values). Error uses 50 evaluation nodes, not a full-domain supremum. Each advertiser's gain is normalized by gross truthful value $v_i G_i(v_i)$; values below 10^{-9} are excluded. Advertiser ratios are averaged within each market before computing the mean and p95 across markets. Declining diagnostics do not certify an asymptotic rate or the theoretical IC bound.

sample size n	100	300	1000	3000	10000
nodewise $\max \lvert\widehat{G} - G\rvert/M$	0.029	0.018	0.0095	0.0052	0.0020
grid gain (market mean)	0.323	0.156	0.064	0.032	0.017
grid gain (p95 market mean)	0.885	0.534	0.226	0.087	0.051

gross truthful value $v_i G_i(v_i)$, not net surplus; the p95 of market means is about 5%. Numerical integration error is also present.

E3. Dynamic stability under co-adaptation (Q3). *Answer (Q3): reports approach a near-truthful region in the tested finite runs.* E3 uses a 50-point grid to maximize the proxy utility $(v_i - r)\widehat{G}_i(r) + \int_0^r \widehat{G}_i$, not the true-utility objective $v_i G_i(r) - \widehat{P}_i(r)$ of Proposition 2. Each day re-subsamples the same fixed log, rather than drawing fresh iid traffic.

Tested starts and scales (Table 8): $N \in \{4, 10, 50\}$ and initial multipliers $\{0.3, 0.5, 0.9, 1.5\}$ give near-unit median active shade after 60 days. The simulation uses $\bar{v} = \max_i v_i$ and clips initial reports to that bound, unlike a fixed public support bound. Residuals are mean squared report errors over advertisers active at truth. Trajectories approach a noisy floor; the exact-G comparison has $N = 50$ residual 0.011. These observations do not isolate estimation as the sole cause of residual error, establish convergence from every start, or verify uniform quadratic growth and vanishing full-domain error.

From simulator to production: an online A/B test. The questions above concern the synthetic simulator, where values are known; as an initial real-world check, we ran an online A/B test on a deployed tCPA product, applying a discount from our payment formula through the platform's existing billing channel—a test of advertiser responsiveness to that discount, not an end-to-end realization of the mechanism. Because advertisers' private values are unobserved, the test cannot identify truthful reports or verify DSIC. Cost rose $+4.38\%$ market-wide ($p < 0.01$); on the pre-specified low-budget-utilization cohort the discount targeted, the reported target rose $+2.54\%$ ($p = 0.02$) and budget $+4.02\%$ ($p < 0.01$). Appendix H gives the methodology, scope, and caveats.

6 Conclusion

We design an allocation-preserving, prior-free payment correction for the Auto-bidding Meta-Game within a uniform-bidding interface. With exact allocation information it is DSIC whenever the induced allocation is monotone; under first-price it is a nonnegative rebate and the truthful equilibrium is welfare-optimal. Estimated payments generally require signed transfers. Uniform allocation error controls payment error and IC slack; report displacement and welfare control require additional quadratic growth. Damped true-utility best responses track a neighborhood under bounded error and converge to truth only when that error vanishes. Our finite-grid synthetic diagnostics show small deviation gains and near-truthful reports in the tested settings, not full-domain statistical or global dynamic validation. Second-price audits expose numerical nonmonotonicity and unresolved integration sensitivity, not empirical validation of DSIC. The production A/B test offers a behavioral signal, not a DSIC test: values are unobserved and the billing channel is not the exact mechanism. End-to-end estimation, allocation-preserving settlement, and stability under drift remain open.

AI use statement

In this work, we used generative AI tools to generate the code that produces the synthetic datasets used in our simulated experiments (Section 5.1), to help develop theoretical framing and definitions, formulate and refine mathematical claims, provide critical ingredients for several proofs and assist in writing them up, and to give feedback on the experimental methodology. We have used generative AI

tools to propose or refine the experimental hypotheses (Q1–Q3, Section 5.1) or to interpret the resulting numbers; assisting with translation, cleaning/reformatting datasets, and qualitative/thematic data analysis are not applicable to this work. Additionally, we used generative AI tools to implement and debug the mechanism, estimator, and simulation code, to edit the manuscript's prose for readability, to format bibliography entries, and to summarize and identify relevant related work (Appendix J). We have reviewed all AI-assisted work: every proof step was checked by hand against the stated theorem, all AI-generated synthetic-data code was checked to confirm it draws correctly from the stated distributions, and all AI-assisted code was tested against the paper's own tables and figures for numerical agreement before inclusion. We take responsibility for the final content of this work, including text, claims, and artifacts produced with the aid of generative AI.

Ethics statement

The online A/B test reported in Appendix H was run through the platform's standard, pre-existing experimentation pipeline used for all bidding-product tests, and it manipulates only advertiser-side billing parameters (a per-conversion discount) already at the platform's disposal; it does not involve human subjects, collect or expose any end-user personal data, or change what data any party can observe. All reported metrics (cost, reported target, budget) are aggregated at the cohort level. We are not aware of any ethical concerns raised by this work.

Reproducibility statement

All simulated results in Section 5.1 are computed from a fixed random seed and are exactly reproducible; every table and figure caption states the market distributions, sample sizes, and estimator needed to regenerate it. Appendix A gives complete proofs of every theorem and proposition stated in the main text. Appendix H documents the online A/B test's estimation methodology, delivery mechanism, and scope in full; the underlying production logs cannot be released publicly, and this limitation is stated explicitly in that appendix.

References

Gagan Aggarwal, Ashish Goel, and Rajeev Motwani. Truthful auctions for pricing search keywords. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, EC '06, pp. 1–7, New York, NY, USA, 2006. Association for Computing Machinery. ISBN 1595932364. doi: 10.1145/1134707.1134708. URL https://doi.org/10.1145/1134707.1134708.

Gagan Aggarwal, Ashwinkumar Badanidiyuru, and Aranyak Mehta. Autobidding with constraints. In *Proceedings of the 15th International Conference on Web and Internet Economics (WINE)*, pp. 17–30, 2019.

Gagan Aggarwal, Ashwinkumar Badanidiyuru, Santiago R. Balseiro, Kshipra Bhawalkar, Yuan Deng, Zhe Feng, Gagan Goel, Christopher Liaw, Haihao Lu, Mohammad Mahdian, Jieming Mao, Aranyak Mehta, Vahab Mirrokni, Renato Paes Leme, Andrés Perlroth, Georgios Piliouras, Jon Schneider, Ariel Schvartzman, Balasubramanian Sivan, Kelly Spendlove, Yifeng Teng, Di Wang, Hanrui Zhang, Mingfei Zhao, Wennan Zhu, and Song Zuo. Auto-bidding and auctions in online advertising: A survey. *ACM SIGecom Exchanges*, 22(1):159–183, 2024. doi: 10.1145/3699824.3699838.

Yeganeh Alimohammadi, Aranyak Mehta, and Andrés Perlroth. Incentive compatibility in the auto-bidding world. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC '23)*, pp. 63. ACM, 2023. doi: 10.1145/3580507.3597725.

Moshe Babaioff, Richard Cole, Jason Hartline, Nicole Immorlica, and Brendan Lucier. Non-quasi-linear agents in quasi-linear mechanisms. In *12th Innovations in Theoretical Computer Science Conference (ITCS 2021)*, 2021. doi: 10.4230/LIPIcs.ITCS.2021.84.

Santiago Balseiro, Yuan Deng, Jieming Mao, Vahab Mirrokni, and Song Zuo. Robust auction design in the auto-bidding world. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021a.

Santiago R. Balseiro and Yonatan Gur. Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Science*, 65:3952–3968, 2019.

Santiago R. Balseiro, Yuan Deng, Jieming Mao, Vahab Mirrokni, and Song Zuo. The landscape of auto-bidding auctions: Value versus utility maximization. *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, pp. 132–133, 2021b.

Yeon-Koo Che and Ian Gale. The optimal mechanism for selling to a budget-constrained buyer. *Journal of Economic Theory*, 92(2):198–233, 2000. doi: 10.1006/jeth.1999.2639.

Xi Chen, Christian Kroer, and Rachitesh Kumar. The complexity of pacing for second-price auctions. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC '21)*, 2021. doi: 10.1145/3465456.3467578.

Vincent Conitzer, Christian Kroer, Debmalya Panigrahi, Okke Schrijvers, Eric Sodomka, Nicolás E. Stier-Moses, and Christopher A. Wilkens. Pacing equilibrium in first price auction markets. *Management Science*, 68(12), 2022a. doi: 10.1287/mnsc.2022.4310.

Vincent Conitzer, Christian Kroer, Eric Sodomka, and Nicolas E. Stier-Moses. Multiplicative pacing equilibria in auction markets. *Operations Research*, 70(2):963–989, 2022b. doi: 10.1287/opre.2021.2167.

Yuan Deng, Sébastien Lahaie, Vahab Mirrokni, and Song Zuo. A data-driven metric of incentive compatibility. In *Proceedings of The Web Conference 2020 (WWW '20)*, 2020. doi: 10.1145/3366423.3380249.

Yuan Deng, Jieming Mao, Vahab Mirrokni, and Song Zuo. Towards efficient auctions in an auto-bidding world. In *Proceedings of the Web Conference (WWW)*, pp. 3965–3973, 2021.

Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, March 2007. doi: 10.1257/aer.97.1.242. URL `https://www.aeaweb.org/articles?id=10.1257/aer.97.1.242`.

Yiding Feng, Brendan Lucier, and Aleksandrs Slivkins. Strategic budget selection in a competitive autobidding world. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC '24)*, 2024. doi: 10.1145/3618260.3649688.

Amos Fiat, Kira Goldner, Anna R. Karlin, and Elias Koutsoupias. The FedEx problem. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC '16)*, 2016. doi: 10.1145/2940716.2940752.

Jason Gaitonde, Yingkai Li, Bar Light, Brendan Lucier, and Aleksandrs Slivkins. Budget pacing in repeated auctions: Regret and efficiency without convergence. In *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*, 2023. doi: 10.4230/LIPIcs.ITCS.2023.52.

Negin Golrezaei, Ilan Lobel, and Renato Paes Leme. Auction design for ROI-constrained buyers. In *Proceedings of the Web Conference 2021 (WWW '21)*, pp. 3941–3952, 2021. doi: 10.1145/3442381.3449841.

Jason D. Hartline. Mechanism design and approximation. Manuscript, book in preparation, `http://jasonhartline.com/MDnA/`, 2017. Accessed 2026.

Benjamin Heymann. Cost per action constrained auctions. In *Proceedings of the 14th Workshop on the Economics of Networks, Systems and Computation (NetEcon)*, 2019. doi: 10.1145/3338506.3340269.

Yoav Kolumbus and Noam Nisan. Auctions between regret-minimizing agents. In *Proceedings of the ACM Web Conference 2022 (WWW '22)*, 2022a. doi: 10.1145/3485447.3512055.

Yoav Kolumbus and Noam Nisan. How and why to manipulate your own agent: On the incentives of users of learning agents. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022b.

Bin Li, Xiao Yang, Daren Sun, Zhi Ji, Zhen Jiang, Cong Han, and Dong Hao. Incentive mechanism design for ROI-constrained auto-bidding, 2020.

Juncheng Li and Pingzhong Tang. Vulnerabilities of single-round incentive compatibility in auto-bidding: Theory and evidence from roi-constrained online advertising markets, 2022.

Christopher Liaw and Wennan Zhu. Risk-averse and optimistic advertiser incentive compatibility in auto-bidding, 2025.

Christopher Liaw, Aranyak Mehta, and Andrés Perlroth. Efficiency of non-truthful auctions in auto-bidding: The power of randomization. In *Proceedings of the ACM Web Conference 2023 (WWW '23)*, 2023. doi: 10.1145/3543507.3583492.

Hongtao Lv, Xiaohui Bei, Zhenzhe Zheng, and Fan Wu. Auction design for bidders with ex post ROI constraints. In *Web and Internet Economics (WINE 2023)*, pp. 492–508, 2023. doi: 10.1007/978-3-031-48974-7_28.

Linghui Meng, Chun Gan, Shengsheng Niu, Chengcheng Zhang, Chenchen Li, Chuan Yang, Yi Mao, Xin Zhu, Jie He, Zhangang Lin, and Ching Law. JD-BP: A joint-decision generative framework for auto-bidding and pricing, 2026.

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

Andrés Perlroth and Aranyak Mehta. Auctions without commitment in the auto-bidding world. In *Proceedings of the ACM Web Conference 2023 (WWW '23)*, 2023. doi: 10.1145/3543507.3583416.

Hal R. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, 2007. ISSN 0167-7187. doi: https://doi.org/10.1016/j.ijindorg.2006.10.002. URL `https://www.sciencedirect.com/science/article/pii/S0167718706001317`.

Christopher A. Wilkens, Ruggiero Cavallo, Rad Niazadeh, and Samuel Taggart. Mechanism design for value maximizers. *CoRR*, abs/1607.04362, 2016.

A Deferred proofs

This appendix collects the proofs deferred from Section 3 (the welfare and individual-rationality guarantees of Theorem 2) and Section 4 (the estimation-robustness and convergence results). Theorem 1 instantiates Myerson's classical characterization Myerson (1981); Hartline (2017) and is proved below without differentiability assumptions.

Restatement. Among AMG mechanisms with $P_i(0;\vec{\tau}_{-i}) = 0$ for every rival profile, a mechanism $M = (G, P)$ is DSIC if and only if the allocation $G_i(\tau_i)$ is monotonically non-decreasing in the reported tCPA τ_i and the payment takes the normalized integral form $P_i(\tau_i) = \tau_i G_i(\tau_i) - \int_0^{\tau_i} G_i(z)\,dz$ with $P_i(0) = 0$.

Proof of Theorem 1. Fix an advertiser i and an arbitrary rival-report profile $\vec{\tau}_{-i}$. Throughout this proof write

$$G(r) := G_i(r;\vec{\tau}_{-i}), \qquad P(r) := P_i(r;\vec{\tau}_{-i}), \qquad U_v(r) := vG(r) - P(r),$$

where $v \in [0, \bar{v}]$ is the advertiser's true type and $r \in [0, \bar{v}]$ is its report. Thus $U_v(r)$ compares reports while holding both the true type and rivals fixed.

Sufficiency. Suppose G is non-decreasing and the payment is

$$P(r) = rG(r) - \int_0^r G(s)\,ds.$$

The bounded monotone allocation is Riemann integrable, even if it has jumps. Substituting the payment into utility gives

$$\begin{aligned} U_v(r) &= vG(r) - \left(rG(r) - \int_0^r G(s)\,ds\right) \\ &= (v-r)G(r) + \int_0^r G(s)\,ds, \\ U_v(v) &= \int_0^v G(s)\,ds. \end{aligned}$$

Consequently, the gain from reporting truthfully rather than reporting r is

$$\begin{aligned} U_v(v) - U_v(r) &= \int_0^v G(s)\,ds - \int_0^r G(s)\,ds - (v-r)G(r) \\ &= \int_r^v G(s)\,ds - (v-r)G(r). \end{aligned}$$

If $r \le v$, the integration runs from the lower report to the true type, so

$$U_v(v) - U_v(r) = \int_r^v \big[G(s) - G(r)\big]\,ds \ge 0.$$

Indeed, $s \in [r, v]$ implies $G(s) \ge G(r)$ by monotonicity, and the interval has non-negative length.

If $r \ge v$, reverse the integration limits and use $-(v-r) = r - v$ to obtain

$$\begin{aligned} U_v(v) - U_v(r) &= -\int_v^r G(s)\,ds + (r-v)G(r) \\ &= \int_v^r \big[G(r) - G(s)\big]\,ds \ge 0. \end{aligned}$$

Here $s \in [v, r]$ implies $G(s) \le G(r)$, again making the integrand non-negative over an interval of non-negative length. At $r = v$ the difference is zero. Thus truth-telling is a weakly optimal report for every type and every rival profile, establishing DSIC. This argument requires neither differentiability nor strict monotonicity and allows ties between optimal reports.

Necessity: monotonicity. Conversely, suppose the mechanism is DSIC. For any two types $0 \le a < b \le \bar{v}$, apply DSIC first to type a deviating to report b, and then to type b deviating to report a:

$$\begin{aligned} aG(a) - P(a) &\ge aG(b) - P(b), \\ bG(b) - P(b) &\ge bG(a) - P(a). \end{aligned}$$

Adding these inequalities cancels the payments and yields

$$(b-a)\big[G(b) - G(a)\big] \ge 0.$$

Since $b - a > 0$, we have $G(b) \ge G(a)$; hence G is non-decreasing.

Necessity: the payment identity. Define truthful utility by $V(v) := vG(v) - P(v)$. Rewriting the same two DSIC inequalities in terms of V gives

$$\begin{aligned} V(a) &\ge V(b) - (b-a)G(b), \\ V(b) &\ge V(a) + (b-a)G(a). \end{aligned}$$

Rearranging therefore bounds each truthful-utility increment:

$$(b-a)G(a) \le V(b) - V(a) \le (b-a)G(b).$$

For $v > 0$, take any partition $\pi : 0 = t_0 < t_1 < \cdots < t_m = v$ and write $\Delta t_k = t_k - t_{k-1}$. Apply this bound on each subinterval and sum. The utility increments telescope, giving

$$\sum_{k=1}^{m} G(t_{k-1})\Delta t_k \le V(v) - V(0) \le \sum_{k=1}^{m} G(t_k)\Delta t_k.$$

These are the lower and upper Riemann sums of G, since G is non-decreasing. Its finite endpoint values bound it on $[0, v]$, and the gap between these sums satisfies

$$\begin{aligned} 0 &\le \sum_{k=1}^{m} \big[G(t_k) - G(t_{k-1})\big]\Delta t_k \\ &\le \|\pi\|\big[G(v) - G(0)\big] \longrightarrow 0, \qquad \|\pi\| := \max_k \Delta t_k \longrightarrow 0. \end{aligned}$$

Thus bounded monotonicity ensures Riemann integrability without a continuity assumption, and squeezing between the sums yields

$$V(v) - V(0) = \int_0^v G(s)\,ds.$$

This identity also holds at $v = 0$. Since $V(0) = -P(0)$, solving for the payment gives

$$P(v) = vG(v) - \int_0^v G(s)\,ds + P(0).$$

The normalization $P(0) = 0$ fixes the additive constant and recovers the claimed formula. Without normalization, that constant may depend on rivals' reports but not on the advertiser's own report, so it cancels in all report comparisons. Since the advertiser and rival profile were arbitrary, both directions hold for the full mechanism. □

On the regularity assumption. Highest-score-wins under first-price proves only non-decrease: increasing a report cannot remove a win, given fixed rivals and report-independent tie-breaking. Continuity and positive slope require further distributional assumptions, and finite logs generally give steps. Flat regions may yield multiple best responses for a fixed rival profile, but this alone proves neither multiple equilibria of the full game nor welfare irrelevance under joint report changes. Strict increase over the entire report domain for every rival profile would make truthful reporting the unique best response and the truthful profile the unique equilibrium. We do not assume this. Under other pricing rules monotonicity must be established for the complete pacing-induced allocation.

Why settlement is an end-of-day transfer. The signed settlement is applied after allocation rather than during the auction because intervening during the auction—e.g. lowering the per-click price—would slacken the cost constraint and let pacing acquire extra traffic, inflating G_i; settling afterwards preserves allocation only if the transfer is excluded from its pacing feedback.

Restatement. Under the exact DSIC mechanism truthful reporting is a weakly dominant-strategy equilibrium, $\vec{\tau}^* = \vec{v}$; under first-price, $W^* = W^{\mathrm{opt}} \geq W^{\mathrm{nat}}$ (strict if shading selects strictly lower-value winners with positive probability); and under any pricing rule with a monotone induced allocation, $U_i^* = \int_0^{v_i} G_i(z; \vec{v}_{-i})\,dz \geq 0$.

Proof of Theorem 2 (welfare optimality and individual rationality). Fix a type profile $\vec{v} \in [0, \bar{v}]^N$. The settlement leaves the allocation unchanged and replaces the natural payment by the exact normalized Myerson payment. For each advertiser i, Theorem 1 applies at every fixed rival profile, and therefore gives

$$v_i G_i(v_i; \vec{\tau}_{-i}) - P_i(v_i; \vec{\tau}_{-i}) \geq v_i G_i(r; \vec{\tau}_{-i}) - P_i(r; \vec{\tau}_{-i}) \quad \text{for every } r \in [0, \bar{v}].$$

In particular, this holds at $\vec{\tau}_{-i} = \vec{v}_{-i}$ for every i, so no advertiser can profitably deviate unilaterally from $\vec{v}$. Truth is thus a weakly dominant-strategy equilibrium. The inequalities need not be strict, so this argument neither establishes uniqueness nor selects among other possible equilibria.

(1) Welfare under first-price. First condition on a realization of the M impressions and their conversion rates. Let $x_{i,j}(\vec{\tau})$ be the allocation indicator, or the conditional allocation probability if ties are randomized, for advertiser i on impression j. Feasibility means

$$x_{i,j}(\vec{\tau}) \geq 0, \qquad \sum_i x_{i,j}(\vec{\tau}) \leq 1.$$

Write $m_j := \max_i v_i q_{i,j}$. Non-negative values and feasibility imply, for any allocation,

$$\sum_i v_i q_{i,j} x_{i,j}(\vec{\tau}) \leq m_j \sum_i x_{i,j}(\vec{\tau}) \leq m_j.$$

Under first-price, $\alpha_i = 1$, so at truth the ranking score is exactly $v_i q_{i,j}$. If $m_j > 0$, the winning advertiser attains m_j; splitting a tie among maximizers yields the same value. If $m_j = 0$, any winner or no allocation has value zero. Thus the truthful allocation attains the upper bound impression by impression:

$$\sum_i v_i q_{i,j} x_{i,j}(\vec{v}) = m_j, \qquad W^* = \sum_{j=1}^M m_j = W^{\mathrm{opt}}.$$

The natural-mechanism allocation is also feasible. Consequently, at any natural equilibrium under comparison, if such an equilibrium exists,

$$W^{\mathrm{opt}} - W^{\mathrm{nat}} = \sum_{j=1}^M \left[m_j - \sum_i v_i q_{i,j} x_{i,j}(\vec{\tau}^{\mathrm{nat}}) \right] \geq 0.$$

The comparison in fact holds for every natural report profile, without proving existence of a natural equilibrium. For random impressions, take expectations of these realization-wise identities and inequalities; the same notation then denotes expected total welfare, consistently with expected conversions G_i.

For strictness, each realized loss is non-negative. If shading assigns an impression to a strictly lower-value winner with positive probability, the total loss Y is positive with positive probability. Since $\{Y > 0\} = \bigcup_{m\geq 1}\{Y \geq 1/m\}$, some m has $\Pr(Y \geq 1/m) > 0$, and hence

$$\mathbb{E}Y \geq m^{-1}\Pr(Y \geq 1/m) > 0.$$

This proves the stated strict comparison in the random setting; a positive loss on any impression suffices for a fixed finite market. Merely changing between tied value maximizers does not make the loss positive. Under other pricing rules the score can instead be $\alpha_i v_i q_{i,j}$, so the impression-wise equality with m_j is not justified.

(2) Individual rationality for any monotone induced allocation. Now fix advertiser i and the truthful rival profile. Substitute the exact payment, retaining the same allocation on both sides:

$$\begin{aligned} U_i^* &= v_i G_i(v_i; \vec{v}_{-i}) - P_i(v_i; \vec{v}_{-i}) \\ &= v_i G_i(v_i; \vec{v}_{-i}) - \left[v_i G_i(v_i; \vec{v}_{-i}) - \int_0^{v_i} G_i(z; \vec{v}_{-i})\, dz \right] \\ &= \int_0^{v_i} G_i(z; \vec{v}_{-i})\, dz \geq 0. \end{aligned}$$

The last inequality uses $G_i \geq 0$ and $v_i \geq 0$; at $v_i = 0$ the integral is zero. Bounded monotonicity makes the integral well-defined even for step allocations. This calculation works at any fixed rival profile and uses no first-price identity. Monotonicity is needed to implement the DSIC payment on the induced allocation, not to turn this non-negative integral into a welfare-optimal ranking. □

Remark 2 (Grid projection and interpolation). Let $0 = r_0 < \cdots < r_K = \bar{v}$ be the full-domain grid, let $h := \max_{k<K}(r_{k+1} - r_k)$, and suppose G is non-decreasing and

$$c_k := G(r_k), \qquad |y_k - c_k| \leq \eta \quad (0 \leq k \leq K).$$

For fixed positive weights $w_0, \ldots, w_K$, weighted least-squares isotonic regression minimizes $\sum_{j=0}^{K} w_j (z_j - y_j)^2$ over $z_0 \leq \cdots \leq z_K$. Its unique fitted vector has the min–max representation

$$T(y)_k = \widetilde{y}_k = \max_{0\leq a\leq k} \min_{k\leq b\leq K} \frac{\sum_{j=a}^{b} w_j y_j}{\sum_{j=a}^{b} w_j}.$$

All denominators are positive. Thus increasing any input coordinate cannot decrease any interval average, its minimum over b, or the subsequent maximum over a. In particular, $u \leq y$ coordinatewise implies $T(u) \leq T(y)$, with the same weights used in both fits.

To check that T fixes any non-decreasing vector c, choose $a = k$: every average over $[k, b]$ is at least c_k, and the choice $b = k$ attains c_k, so $T(c)_k \geq c_k$. Conversely, for every $a \leq k$, choosing $b = k$ gives an average at most c_k, so the inner minimum is at most c_k and $T(c)_k \leq c_k$. Therefore $T(c) = c$. Both $c - \eta\mathbf{1}$ and $c + \eta\mathbf{1}$ are also non-decreasing. Applying order preservation to the nodewise error inequalities yields

$$c - \eta\mathbf{1} = T(c - \eta\mathbf{1}) \leq T(y) \leq T(c + \eta\mathbf{1}) = c + \eta\mathbf{1}, \qquad \max_k |\widetilde{y}_k - c_k| \leq \eta.$$

This is a deterministic statement for any realized positive weights; no probabilistic independence of those weights is needed.

On $[r_k, r_{k+1}]$, the linear interpolant is

$$\widehat{G}(r) = (1 - \theta)\widetilde{y}_k + \theta\widetilde{y}_{k+1}, \qquad \theta = \frac{r - r_k}{r_{k+1} - r_k} \in [0, 1].$$

Because the fitted nodes are ordered, this lies between them. Monotonicity of G similarly puts $G(r)$ between c_k and c_{k+1}. The node bound therefore gives both one-sided error inequalities:

$$\begin{aligned} \widehat{G}(r) - G(r) &\leq (c_{k+1} + \eta) - c_k, \\ G(r) - \widehat{G}(r) &\leq c_{k+1} - (c_k - \eta). \end{aligned}$$

Since $r_{k+1} - r_k \le h$, taking the supremum over every grid interval proves

$$\|\widehat{G} - G\|_\infty \le \eta + \omega_G(h), \qquad \omega_G(h) := \sup_{\substack{r,s\in[0,\bar{v}]\\ |r-s|\le h}} |G(r) - G(s)|.$$

The same reasoning applies to any interpolant staying between adjacent fitted nodes, not to an interpolant that overshoots them. If the total allocation G is L-Lipschitz, then $\omega_G(h) \le Lh$, giving the upper bound $\eta + Lh$, not an equality in general. The full-domain grid is essential: the argument does not control extrapolation beyond its endpoints. Monotonicity alone does not make $\omega_G(h)$ tend to zero; for an interior jump of size $J > 0$, points on opposite sides of that jump give $\omega_G(h) \ge J$ for every $h > 0$.

Restatement. Fix rivals and draw n iid impressions; for $f_r = q\mathbf{1}\{X \le r\}$ with $g(r) = \mathbb{E}f_r$ and empirical mean g_n, with probability at least $1 - \delta$ one has $\sup_r |g_n(r) - g(r)| \le a_n(\delta) = \frac{4Q}{\sqrt{n}} + Q\sqrt{\log(1/\delta)/(2n)}$, and fitting the replay nodes $y_k = Mg_n(r_k)$ as in Remark 2 yields $\epsilon \le Ma_n(\delta) + Lh$ when the total G is L-Lipschitz.

Proof of Proposition 1. Fix the advertiser and rivals before drawing iid observations $Z_j = (q_j, X_j)$, with $0 \le q_j \le Q$; no independence between q_j and X_j is needed. Start with $f_r(Z_j) = q_j\mathbf{1}\{X_j \le r\}$ and define

$$Z := \sup_{r\in[0,\bar{v}]} |g_n(r) - g(r)|.$$

Symmetrization and threshold ordering. Let Z_j' be an independent ghost sample, g_n' its empirical mean, and σ_j iid Rademacher signs independent of both samples. Jensen's inequality, exchangeability of each pair (Z_j, Z_j'), and the triangle inequality give

$$\begin{aligned}
\mathbb{E}Z &\le \mathbb{E}\sup_r |g_n(r) - g_n'(r)| \\
&= \mathbb{E}\sup_r \left| \frac{1}{n}\sum_j \sigma_j\big(f_r(Z_j) - f_r(Z_j')\big) \right| \\
&\le 2\mathbb{E}\sup_r \left| \frac{1}{n}\sum_j \sigma_j f_r(Z_j) \right|.
\end{aligned}$$

Conditional on the observations, let π order thresholds increasingly, independently of the signs. Every selected set is a prefix (tied thresholds enter together), so

$$S_k := \sum_{\ell=1}^{k} \sigma_{\pi(\ell)} q_{\pi(\ell)}, \quad S_0 = 0, \qquad \sup_r \left| \frac{1}{n}\sum_j \sigma_j f_r(Z_j) \right| \le \frac{1}{n}\max_{0\le k\le n} |S_k|.$$

The partial sums form a square-integrable martingale in the ordered signs, with $\mathbb{E}_\sigma S_n^2 = \sum_j q_j^2$. Cauchy–Schwarz and Doob's L^2 inequality yield

$$\mathbb{E}_\sigma \max_k |S_k| \le \big(\mathbb{E}_\sigma \max_k |S_k|^2\big)^{1/2} \le 2\Big(\sum_j q_j^2\Big)^{1/2} \le 2Q\sqrt{n}, \qquad \mathbb{E}Z \le \frac{4Q}{\sqrt{n}}.$$

Concentration and total-allocation error. Replacing one observation changes every empirical mean by at most Q/n, hence $|Z - \widetilde{Z}| \le \|g_n - \widetilde{g}_n\|_\infty \le Q/n$. McDiarmid's inequality gives, for $Q > 0$ and $u > 0$,

$$\Pr\{Z > \mathbb{E}Z + u\} \le e^{-2nu^2/Q^2}.$$

Taking $u = Q\sqrt{\log(1/\delta)/(2n)}$ for $0 < \delta < 1$ proves $Z \le a_n(\delta)$ with probability at least $1 - \delta$; $Q = 0$ is trivial. This event is uniform over the full report domain, so simultaneously

$$\max_k |y_k - G(r_k)| \le Ma_n(\delta), \qquad \epsilon = \|\widehat{G} - G\|_\infty \le Ma_n(\delta) + \omega_G(h) \le Ma_n(\delta) + Lh.$$

Here Remark 2 applies to the full-domain grid without a node union bound. The last inequality requires the *total* curve $G = Mg$ to be L-Lipschitz; otherwise retain $\omega_G(h)$. The sample size n need not equal the market size M.

Conventions and measurability. Strict comparisons also select prefixes. For fixed-priority ties, write

$$f_r(Z_j) = q_j\big(\mathbf{1}\{X_j < r\} + e_j\mathbf{1}\{X_j = r\}\big), \qquad e_j \in \{0, 1\}.$$

Here e_j records whether the advertiser wins equality. Within each tied group, put observations with $e_j = 1$ first. This sign-independent ordering again makes every selected set a prefix, so all bounds remain valid. Truth and replay must use the same convention. When $q_j = 0$, set its contribution to zero and assign any threshold; a no-allocation rule at $r = 0$ sets both curves there to zero and adds only the empty prefix.

To justify measurability, the finite measure $\nu(B) = \mathbb{E}[q\mathbf{1}\{X \in B\}]$ has a countable deterministic atom set A. At any interior report outside A, the population curve is continuous, while the empirical value lies between its one-sided limits. Its discrepancy is therefore bounded by the larger limiting discrepancy, approached along rational reports. Thus Z equals the supremum over rational reports, A, and the domain endpoints (restricted to the domain). Empirical and signed-sample suprema require only finitely many thresholds and intervening intervals. This also covers mixed tie priorities and the zero-report convention.

These are iid fixed-rival population guarantees, not automatic guarantees for reused logs, sampling without replacement, or traffic drift. With fresh conditional-iid logs, one may condition on the history fixing rivals and union-bound over a finite adaptive horizon; an infinite horizon needs additional probability control. □

Restatement. Uniform allocation error bounds payment error and IC slack without a slope assumption. Report displacement additionally requires the quadratic-growth condition of Theorem 3; own-value displacement additionally uses Lipschitz continuity.

Proof of Theorem 3 (robustness of the estimated mechanism). Fix advertiser i, an arbitrary rival profile $\vec{\tau}_{-i}$, and a type $v \in [0, \bar{v}]$. Suppress the advertiser and rival indices throughout. The estimate affects only the payment, not the current allocation. In particular, the utilities being compared are

$$U_v(r) = vG(r) - P(r), \qquad \widehat{U}_v(r) = vG(r) - \widehat{P}(r),$$

not a proxy utility in which G is replaced by $\widehat{G}$. The integral in the estimated payment is well-defined for the fitted curve of Section 4.1; more generally the payment definition presupposes an integrable estimate. Define the allocation and payment errors by

$$e(r) := \widehat{G}(r) - G(r), \qquad E(r) := \widehat{P}(r) - P(r), \qquad |e(r)| \le \epsilon.$$

(1) Payment and utility deviation. Subtract the two normalized payment identities to obtain

$$E(r) = re(r) - \int_0^r e(z)\,dz.$$

Both the endpoint error and the integrated error contribute to the bound. For every $r \in [0, \bar{v}]$, the triangle inequality and the uniform error assumption give

$$\begin{aligned} |E(r)| &\le r|e(r)| + \left|\int_0^r e(z)\,dz\right| \\ &\le r\epsilon + \int_0^r |e(z)|\,dz \le r\epsilon + \int_0^r \epsilon\,dz = 2\epsilon r. \end{aligned}$$

Since the allocation term is identical in the two utilities,

$$\widehat{U}_v(r) = U_v(r) - E(r), \qquad |\widehat{U}_v(r) - U_v(r)| = |E(r)| \le 2\epsilon r.$$

Taking $r = v$ proves the truthful-utility claim, including the zero-report endpoint.

(2) Approximate IC. The true normalized mechanism is DSIC by Theorem 1, so $U_v(r) - U_v(v) \le 0$ for every report r. Insert the payment error at both reports:

$$\begin{aligned}\widehat{U}_v(r) - \widehat{U}_v(v) &= U_v(r) - U_v(v) - E(r) + E(v) \\ &\le |E(r)| + |E(v)| \\ &\le 2\epsilon(r+v) \le 4\epsilon\bar{v}.\end{aligned}$$

The last step uses the common report and type domain. Taking the supremum over all r yields $\Delta_{IC}(v) \le 4\epsilon\bar{v}$; because reporting v is feasible, this supremum gain is also non-negative. No maximizer, derivative, or positive slope is needed for this part.

The rival profile was arbitrary: this is a conditional guarantee at each profile where the stated uniform estimation bound holds. To obtain a common approximately-DSIC bound over all rival profiles, the same error bound must hold at every such profile. For adaptive rivals, the algebra applies separately at each day's actual profile on its full-domain error event; it does not turn an estimate accurate at one profile into an estimate accurate at another.

(3) From quadratic growth to report displacement. Assume now the separate condition $D_v(r) := U_v(v) - U_v(r) \ge \mu(r-v)^2/2$ for every candidate report, with $\mu > 0$. Let r^* be any maximizer of $\widehat{U}_v$ whose existence is assumed. Since the truthful report is feasible,

$$\widehat{U}_v(r^*) \ge \widehat{U}_v(v).$$

Using $U_v = \widehat{U}_v + E$ and then the sign of this utility difference gives

$$\begin{aligned}\frac{\mu}{2}(r^*-v)^2 &\le D_v(r^*) \\ &= \widehat{U}_v(v) - \widehat{U}_v(r^*) + E(v) - E(r^*) \\ &\le |E(v)| + |E(r^*)| \\ &\le 2\epsilon(v + r^*) \le 4\epsilon\bar{v}.\end{aligned}$$

Divide by $\mu/2 > 0$ and take non-negative square roots to obtain

$$|r^* - v| \le \sqrt{\frac{8\epsilon\bar{v}}{\mu}}.$$

The constant 8 is the product of the utility-comparison constant 4 and the factor 2 in the growth normalization. This conclusion holds for every estimated maximizer; it does not require choosing a unique one. Without quadratic growth, approximate IC bounds utility loss but need not bound report distance, for example on a flat allocation interval.

Verification of the sufficient slope condition. The derivative-free Myerson identity proved above gives, for any r,

$$D_v(r) = \begin{cases} \displaystyle\int_r^v [G(s) - G(r)]\, ds, & r \le v, \\ \displaystyle\int_v^r [G(r) - G(s)]\, ds, & r \ge v. \end{cases}$$

Suppose G is absolutely continuous on the entire segment joining v and r, with $G'(t) \ge \mu$ almost everywhere there. For $r \le v$, the fundamental theorem for absolutely continuous functions implies $G(s) - G(r) = \int_r^s G'(t)\, dt$. Substitution and integration over the triangular region $r \le t \le s \le v$ yield

$$\begin{aligned}D_v(r) &= \int_r^v \int_r^s G'(t)\, dt\, ds \\ &= \int_r^v \int_t^v G'(t)\, ds\, dt = \int_r^v (v-t) G'(t)\, dt \\ &\ge \mu \int_r^v (v-t)\, dt = \frac{\mu}{2}(v-r)^2.\end{aligned}$$

For $r \geq v$, use $G(r) - G(s) = \int_s^r G'(t)\,dt$ and the region $v \leq s \leq t \leq r$ instead:

$$\begin{aligned} D_v(r) &= \int_v^r \int_s^r G'(t)\,dt\,ds \\ &= \int_v^r \int_v^t G'(t)\,ds\,dt = \int_v^r (t-v)G'(t)\,dt \\ &\geq \mu \int_v^r (t-v)\,dt = \frac{\mu}{2}(r-v)^2. \end{aligned}$$

The order changes are justified by the non-negative integrands (or by absolute integrability on a bounded segment). At $r = v$ both sides vanish. Establishing this sufficient condition for every candidate report requires control on every entire connecting segment, not just a positive derivative near truth. It is a sufficient condition for the assumed growth inequality, not a consequence of bounded monotonicity; finite step allocations generally do not satisfy such a slope condition.

Own-value change, with an additional Lipschitz assumption. If the total allocation G is L-Lipschitz on the report domain, then

$$\begin{aligned} |vG(r^*) - vG(v)| &= v|G(r^*) - G(v)| \\ &\leq vL|r^* - v| \leq \bar{v}L|r^* - v| \leq \bar{v}L\sqrt{\frac{8\epsilon\bar{v}}{\mu}}. \end{aligned}$$

Here rivals remain fixed in both allocations, and L is measured in total conversions per unit of report. The expression controls only this advertiser's value, not its utility or the changes in other advertisers' values. It cannot be summed as an aggregate-welfare comparison when all rival reports change at once; Corollary 4 treats that separate question. □

Restatement. Under first-price with $q_{i,j} \in [0,1]$, every report vector $\vec{\tau}$ satisfies $W^{\mathrm{opt}} - W(\vec{\tau}) \leq 2M\|\vec{\tau} - \vec{v}\|_\infty$, so a profile within Δ of truth loses at most $2M\Delta$ in welfare.

Proof of Corollary 4 (aggregate welfare deviation). Fix any report profile $\vec{\tau}$ and set $\Delta_\tau := \|\vec{\tau} - \vec{v}\|_\infty$. The argument compares two complete allocations at this same report profile and the truthful benchmark, rather than adding unilateral own-value changes.

(1) Compare the two winners on one impression. Fix a realized impression j. Include a dummy outside option with value, report, and conversion rate all equal to zero. Choose any value-optimal winner $a \in \arg\max_i v_i q_{i,j}$, and let b be the winner selected by the first-price allocation at reports $\vec{\tau}$. Highest-bid-wins under first-price means

$$\tau_b q_{b,j} \geq \tau_a q_{a,j}.$$

If the mechanism makes no allocation when all report scores are zero, take b to be the dummy; the inequality still holds. Ties among positive scores may be resolved by the prescribed rule, since only this weak inequality is needed. Likewise, the choice among value-optimal ties does not change the optimal value. If a or b is the dummy, its report error is zero and hence still at most Δ_τ.

(2) Bound the loss using two report errors. Optimality of a gives a non-negative welfare loss. Add and subtract the two report scores and apply the winner inequality:

$$\begin{aligned} 0 \leq v_a q_{a,j} - v_b q_{b,j} &= (v_a - \tau_a)q_{a,j} + (\tau_a q_{a,j} - \tau_b q_{b,j}) + (\tau_b - v_b)q_{b,j} \\ &\leq (v_a - \tau_a)q_{a,j} + (\tau_b - v_b)q_{b,j} \\ &\leq |v_a - \tau_a|q_{a,j} + |\tau_b - v_b|q_{b,j} \\ &\leq \Delta_\tau(q_{a,j} + q_{b,j}) \leq 2\Delta_\tau. \end{aligned}$$

The absolute-value step uses $q_{i,j} \geq 0$, and the last step uses $q_{i,j} \leq 1$. Thus the bound covers the value lost by the displaced optimal winner and the value delivered by its replacement. It also covers $a = b$, when the actual loss is zero, and the zero-allocation case.

(3) Sum the impression-wise losses. Because the welfare benchmark selects a value-maximizing winner separately on each impression,

$$W^{\mathrm{opt}} - W(\vec{\tau}) = \sum_{j=1}^{M} \big(v_{a(j)} q_{a(j),j} - v_{b(j)} q_{b(j),j}\big) \\ \le \sum_{j=1}^{M} 2\Delta_\tau = 2M\|\vec{\tau} - \vec{v}\|_\infty.$$

Here $a(j), b(j)$ denote the two winners on impression j. This is a realization-wise bound. For random traffic or randomized tie-breaking, take expectations on both sides at a fixed report vector; no independence among impressions is required for this step. If the report vector is itself random, the realization-wise inequality instead gives an expected loss bounded by $2M\mathbb{E}\Delta_\tau$, or by $2M\Delta$ whenever $\Delta_\tau \le \Delta$ almost surely.

(4) Apply a simultaneous displacement guarantee. If every coordinate of one report vector satisfies $|\tau_i - v_i| \le \Delta$, then its maximum coordinate error is at most Δ, yielding the stated $2M\Delta$ bound. In particular, whenever the hypotheses of Theorem 3(3) provide this simultaneous guarantee with common error ϵ and growth constant μ,

$$\Delta = \sqrt{\frac{8\epsilon\bar{v}}{\mu}} \quad \Longrightarrow \quad W^{\mathrm{opt}} - W(\vec{\tau}) \le 2M\sqrt{\frac{8\epsilon\bar{v}}{\mu}}.$$

For fixed $M, \bar{v}, \mu$, this is $O(\sqrt{\epsilon})$. Neither existence of an equilibrium nor convergence to a limiting profile is needed. A displacement bound for only one advertiser, or an own-value comparison at different rival profiles, does not by itself supply the simultaneous coordinate bounds used here. □

Restatement. Assume at each day t and the actual rival profile a full-domain error $\|\widehat{G}_i^{(t)} - G_i^{(t)}\|_\infty \le \epsilon_t$ and quadratic growth $D_{i,v_i}^{(t)}(r) \ge \mu(r - v_i)^2/2$ with a common $\mu > 0$, that maximizers $\widehat{r}_i^{(t)}$ exist, and the damped update $\vec{\tau}^{(t+1)} = (1-\lambda)\vec{\tau}^{(t)} + \lambda\widehat{\vec{r}}^{(t)}$ with fixed $\lambda \in (0, 1]$. With $d_t = \|\vec{\tau}^{(t)} - \vec{v}\|_\infty$ and $b_t = \sqrt{8\epsilon_t\bar{v}/\mu}$, then $d_{t+1} \le (1-\lambda)d_t + \lambda b_t$ and the induced convolution bound hold; if $\epsilon_t \le \epsilon$ then $\limsup_t d_t \le \sqrt{8\epsilon\bar{v}/\mu}$, and if $\epsilon_t \to 0$ then $d_t \to 0$ from every initial profile.

Proof of Proposition 2. Work on any sequence of estimates and reports satisfying the proposition's assumptions. At day t, freeze the actual rivals $\vec{\tau}_{-i}^{(t)}$ for advertiser i. All maximizers below are of actual utility, with $G_i^{(t)}$ in the value term and $\widehat{P}_i^{(t)}$ in the payment term.

(1) Bound every selected best response at the current rivals. The full-domain error bound and common quadratic-growth constant permit application of Theorem 3(3) to every advertiser at this profile. Since $v_i \le \bar{v}$, every chosen maximizer satisfies

$$|\widehat{r}_i^{(t)} - v_i| \le \sqrt{\frac{8\epsilon_t\bar{v}}{\mu}} = b_t, \qquad \|\widehat{\vec{r}}^{(t)} - \vec{v}\|_\infty \le b_t.$$

This does not require a unique estimated maximizer. All coordinates use the same day-t profile before the synchronous update. Because the report domain is an interval and the update is a convex combination, reports stay in the domain from every admissible initial profile.

(2) Obtain the one-step tracking inequality. Subtract $v_i = (1-\lambda)v_i + \lambda v_i$ from the update. For each i, the triangle inequality gives

$$|\tau_i^{(t+1)} - v_i| = |(1-\lambda)(\tau_i^{(t)} - v_i) + \lambda(\widehat{r}_i^{(t)} - v_i)| \\ \le (1-\lambda)|\tau_i^{(t)} - v_i| + \lambda|\widehat{r}_i^{(t)} - v_i| \\ \le (1-\lambda)d_t + \lambda b_t.$$

Taking the maximum over advertisers yields

$$d_{t+1} \le (1-\lambda)d_t + \lambda b_t.$$

Thus the proof controls distance to truth directly; it does not assert that the best-response map is a contraction between arbitrary pairs of profiles.

(3) Unroll the recursion. First let $0 < \lambda < 1$ and abbreviate $\rho = 1 - \lambda \in (0, 1)$. For $t = 0$ the convolution sum is empty and the proposed bound is equality. If it holds at time t, then

$$\begin{aligned} d_{t+1} &\le \rho d_t + \lambda b_t \\ &\le \rho^{t+1} d_0 + \lambda \sum_{s=0}^{t-1} \rho^{t-s} b_s + \lambda b_t \\ &= \rho^{t+1} d_0 + \lambda \sum_{s=0}^{t} \rho^{t-s} b_s. \end{aligned}$$

Induction therefore proves, for all $t \ge 0$,

$$d_t \le \rho^t d_0 + \lambda \sum_{s=0}^{t-1} \rho^{t-1-s} b_s.$$

For $\lambda = 1$, the one-step inequality is simply $d_{t+1} \le b_t$. The displayed convolution formula has the same meaning for $t \ge 1$: only its last term survives, interpreting the zeroth power as 1; the $t = 0$ bound is handled separately as $d_0 \le d_0$.

(4) Bounded error gives a neighborhood, not necessarily a limit. If $\epsilon_t \le \epsilon$, set $b = \sqrt{8\epsilon\bar{v}/\mu}$, so $b_t \le b$. For $0 < \lambda < 1$, the geometric sum satisfies

$$\lambda \sum_{s=0}^{t-1} \rho^{t-1-s} = \lambda \sum_{k=0}^{t-1} \rho^k = \lambda \frac{1 - \rho^t}{1 - \rho} = 1 - \rho^t.$$

Consequently

$$d_t \le \rho^t d_0 + b(1 - \rho^t), \qquad \limsup_{t \to \infty} d_t \le b.$$

For $\lambda = 1$, the same limsup follows from $d_t \le b$ for $t \ge 1$. These are upper bounds on distance; they do not imply that the reports, or even their distances, have a limit under persistent error.

(5) Vanishing error gives convergence to truth. If $\epsilon_t \to 0$, then $b_t \to 0$. For $0 < \lambda < 1$, fix any $a > 0$ and choose S such that $b_s \le a$ for every $s \ge S$. For $t > S$, split the convolution at S:

$$\begin{aligned} d_t &\le \rho^t d_0 + \lambda \sum_{s=0}^{S-1} \rho^{t-1-s} b_s + \lambda \sum_{s=S}^{t-1} \rho^{t-1-s} b_s \\ &\le \rho^t d_0 + \rho^{t-S} \lambda \sum_{s=0}^{S-1} \rho^{S-1-s} b_s + a(1 - \rho^{t-S}). \end{aligned}$$

The finite prefix in the middle is a fixed finite number times ρ^{t-S}, so it tends to zero; the initial-profile term also tends to zero. Hence $\limsup_t d_t \le a$. Since $a > 0$ was arbitrary and $d_t \ge 0$, it follows that $d_t \to 0$. For $\lambda = 1$, this follows directly from $d_{t+1} \le b_t \to 0$.

These are deterministic implications along a path satisfying the full-domain error and growth assumptions at every actual rival profile. No independence across days is needed for the algebra. A statistical interpretation requires an event on which the corresponding pathwise assumptions hold: a finite-horizon union bound alone cannot establish the infinite-horizon premise, and reusing logs or merely stabilizing an estimate does not imply vanishing error. □

A bounded error need not produce a limiting report. Take one advertiser, $\bar{v} = 1$, $v = 1/2$, and $G(r) = r/2$. Then $P(r) = r^2/4$ and $D_v(r) = (r - v)^2/4$, so quadratic growth holds with $\mu = 1/2$. Alternate estimates $\widehat{G}^{(t)}(r) = (1 + a_t)r/2$, where a_t alternates between 0.1 and -0.1. The uniform error is 0.05, and maximizing the actual utility $vG(r) - \widehat{P}^{(t)}(r)$ gives $r_t^* = v/(1 + a_t)$, alternating between $5/11$ and $5/9$. With $\lambda = 1$, reports do not converge, although the neighborhood guarantee holds. Even a fixed nonzero a_t gives a stable biased report rather than truth.

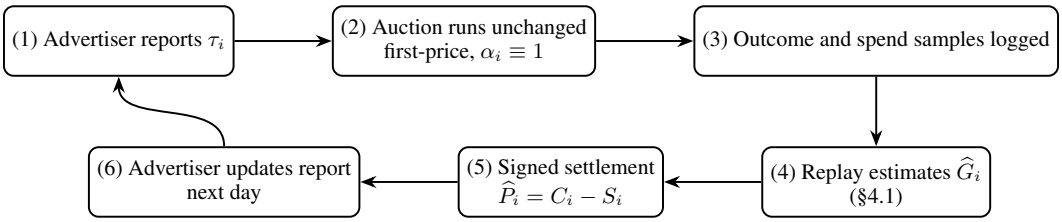

Figure 1: Conceptual allocation-preserving settlement loop. C_i and $\widehat{P}_i$ denote expected charges; logged outcomes are samples, not the entire expected allocation curve. Proposition 2 assumes true-utility best responses and uniform growth/error control, whereas the experiments use finite-grid proxy-utility search. Daily re-estimation alone supplies neither condition.

Table 3: Finite-run search profiles ($N = 5$, 30 markets): exact-G search versus proxy-utility search on an n-impression estimate. The squared-error column averages over all advertisers; median shade uses advertisers active at truth. The assumed-truthful row is a reference; near-truthful reports and improving revenue in these runs are not asymptotic guarantees.

environment	$N^{-1}\|\vec{\tau} - \vec{v}\|^2$	med. shade	W/W^{opt}	$\overline{\Delta\mathrm{Rev}}$
assumed $\vec{\tau} = \vec{v}$	—	—	1.0000	$+117.4$
ideal (exact G)	0.0014	1.000	1.0000	$+117.8$
realistic, n=100	0.0220	0.982	0.9994	$+80.5$
realistic, n=300	0.0113	0.997	0.9999	$+104.6$
realistic, n=1000	0.0026	1.000	1.0000	$+115.3$
realistic, n=3000	0.0014	1.000	1.0000	$+117.8$

B Emergence of truthful reporting under self-selected search

E3 examines finite-run co-adaptation. This appendix compares welfare and revenue at self-selected finite-grid search profiles with the truthful benchmark. Estimated search maximizes proxy utility using $\widehat{G}$ in both the value and payment terms; it is not the true-utility best response assumed by Proposition 2.

Finite-run search (Table 3): exact-G search yields a near-truthful profile (residual 0.0014), with welfare ratio 1.0000 at reported precision. Estimated search increases revenue gain from $+80$ at $n = 100$ to $+118$ at $n = 3000$. These finite-grid proxy-utility results do not prove exact truthfulness, generalize to other pricing rules, or isolate estimation as the sole source of the remaining gap.

C Conditional extension beyond first-price

The welfare optimality of Theorem 2 is specific to first-price. For any other pricing rule r, define the induced allocation

$$G_i^r(\tau_i; \vec{\tau}_{-i}) \;:=\; \mathbb{E}\Big[\textstyle\sum_j q_{i,j}\, \mathbb{I}\{\alpha_i \tau_i\, q_{i,j} \text{ is the top score on } j\}\Big],$$

where the pacing multiplier α_i is set by the rule (Eq. 7); the ideal first-price model has $\alpha_i \equiv 1$, whereas α_i can exceed one under second-price. If $G_i^r(\cdot; \vec{\tau}_{-i})$ is monotone in i's report, Theorem 1 gives DSIC via the Myerson payment $P_i^{\mathrm{IC},r}(\tau_i) = \tau_i G_i^r(\tau_i) - \int_0^{\tau_i} G_i^r(z)\, dz$; only the allocation curve on which the payment is evaluated changes. This is a conditional application of Myerson's theorem, not a general proof that every second-price pacing equilibrium induces a monotone G_i^r. Truth is a weakly dominant report under these conditions, not necessarily the unique optimum.

Actual-kernel unilateral audit. We rerun 60 paired markets with seeds $20260710 + k$, $k = 0, \ldots, 59$, $N = 5$, $M = 100$, iid $v \sim U[0.2, 1]$ and $q \sim U[0.05, 1]$. For each of all 300 advertisers per rule, including inactive ones, rivals' final truthful-run reports *and multipliers* are frozen. Only the focal multiplier is recomputed with the unmodified 50-step CPA-feasibility bisection capped at 20, under both FP and GSP; FP is not replaced by an $\alpha = 1$ surrogate. The focal kernel wins

Table 4: Pricing-rule comparison on 60 paired markets ($N = 5$, $M = 100$, uniform values and q). The raw numerical-payment gain $\max_i \widehat{\delta}_i$ is in simulator utility units, using 16-panel quadrature and the refined report grid; it is not rounded to zero at the 10^{-10} counting tolerance. IC and shade aggregate all 300 advertisers per rule; multiplier dispersion uses truthfully active advertisers. Welfare ratios and multiplier dispersion are market means. Revenue change is mean±s.e.m. over markets, retaining the original final profiles and 60-midpoint payment approximation, not the focal-recomputed audit payments. GSP gains remain integration-sensitive and do not certify exact-payment IC violations or DSIC.

Pricing	$\max_i \widehat{\delta}_i$	$\mathrm{med}(s_i)$	$\Delta W \geq 0$	$\Delta\mathrm{Rev} > 0$	$\overline{\Delta\mathrm{Rev}}$	W/W^{opt}	$\frac{\max_i \alpha_i}{\min_i \alpha_i}$
First-price	1.78×10^{-14}	0.700	**60/60**	56/60	3.2±0.3	1.00	1.00
Second-price	2.17×10^{-2}	0.725	1/60	43/60	0.9±0.4	0.97	1.82

only by strictly exceeding the rival score; the joint auction instead resolves exact ties by index. We record both the stored truthful allocation and its focal recomputation, rather than assuming final-iterate convergence. The report domain is $[0, 1]$, not $[0, v_i]$. For each report r we explicitly compute $\widehat{P}_i(r) = rG_i(r) - \widehat{I}_i(r)$ and $\widehat{U}_i(r) = v_iG_i(r) - \widehat{P}_i(r)$, and report the unrounded, unclipped maximum $\widehat{\delta}_i = \max_{r\in S_i} \widehat{U}_i(r) - \widehat{U}_i(v_i)$. The scan contains truth, 256- and 512-interval uniform grids, candidate allocation breakpoints, their adjacent floating-point reports and 10^{-10} flanks, and interval midpoints. The absolute tolerance 10^{-10} is used only to count exceedances, not to set residuals to zero. We integrate the *actual* kernel by composite midpoint quadrature with 16 and 32 panels per candidate interval. Candidate thresholds and an independent real-arithmetic staircase integral are diagnostics only; no allocation is monotonized. These are finite-search, numerical-payment residuals, distinct from the simulator's natural-payment deviation gain and from its revenue identity error.

The payment is an ex-post Myerson evaluation on this fixed competitor landscape: it is not a per-impression critical-price implementation and does not alter the controller or budget feedback. Re-solving rivals after deviations would define a different audit. Neither a finite scan nor agreement of two quadrature resolutions proves full-domain DSIC against arbitrary opponents; nonmonotonicity and unresolved quadrature differences specifically prevent a second-price IC validation here.

Observed residuals and precision. For FP, the raw maximum/mean/p95 across 300 advertisers are 1.77636×10^{-14}, 1.66162×10^{-15}, and 7.10543×10^{-15}; no advertiser or market exceeds 10^{-10}. Doubling the report grid changes the maximum gain per advertiser by at most 3.55×10^{-15}; increasing integration from 16 to 32 panels changes it by at most 2.84×10^{-14}, with a refined maximum of 3.55×10^{-14}. The largest integral change is 2.20×10^{-13} and discrepancy from the threshold reference is 9.59×10^{-14}. No sampled allocation decrease or truthful-recomputation discrepancy is observed. These are numerical checks, not a certified quadrature error bound. For GSP, the raw maximum/mean/p95 are 2.16850×10^{-2}, 7.68713×10^{-4}, and 5.26165×10^{-3}; 75/300 advertisers in 49/60 markets exceed 10^{-10}. Grid doubling changes a gain by up to 2.16850×10^{-2}; 16-to-32-panel refinement changes a gain by up to 3.81064×10^{-3} (77 advertisers exceed the tolerance), although the refined overall maximum remains 2.16850×10^{-2}. The largest integral change is 4.44075×10^{-3} and threshold-reference discrepancy is 1.65455×10^{-2}. We observe 180 allocation decreases across 71 advertisers, including a maximum decrease of 1.55864 on the grid/flank scan excluding adjacent-float probes; within-interval allocation disagreement reaches 1.48293. Stored versus focal-recomputed truthful allocation and multiplier differences reach 43.1755 and 17.8930, respectively, so the stored final profile cannot be assumed to be a converged pacing equilibrium. These diagnostics preclude interpreting this table as either zero IC slack or a certified violation under exact Myerson payments.

D PERSISTENCE ACROSS MARKET SIZE AND DISTRIBUTION

This appendix checks whether the static first-price guarantees—welfare optimality, shading, revenue—persist as the market scales and the distributions vary (active advertisers). *Scale*: sweeping N from 5 to 100 (with M scaled up proportionally), welfare optimality stays exact ($\min W^*/W^{\mathrm{opt}} = 1.000$ at every N) and revenue rises in essentially every market ($38/40$, $28/30$, $19/20$, $15/15$), with the

Table 5: Persistence across value × conversion-rate distributions ($N = 10$, $M = 2000$, 60 markets/cell, first-price, active-only). Welfare optimality ($\min W^*/W^{\text{opt}} = 1.000$, $\Delta W \geq 0$ in $60/60$) is distribution-free; shading is pervasive across the tested cells, and the revenue gain is strong except under the heavy-tailed pareto value law. (Scale sweep N=5–100 in text above; all $\min W^*/W^{\text{opt}} = 1.000$.) ΔRev is in simulator value units summed over the M impressions (Section 3.3), as in Tables 3, 6, reported as mean±s.e.m. across the 60 markets. Shade medians here are from an independent set of market draws from those in Table 1, so a shared configuration (e.g. exponential value, Beta$(2, 5)$ conversion rate) can differ by Monte-Carlo noise between the two tables.

value dist	q dist	med. shade	$\min W^*/W^{\text{opt}}$	$\Delta W \geq 0$	% ΔRev>0	mean ΔRev
uniform	U[.05, 1]	0.825	1.000	60/60	100%	+45.0±2.2
uniform	U[0, 1]	0.806	1.000	60/60	100%	+46.8±1.8
uniform	Beta(2, 5)	0.750	1.000	60/60	100%	+38.6±1.5
lognormal	U[.05, 1]	0.825	1.000	60/60	93%	+92.5±8.0
lognormal	U[0, 1]	0.825	1.000	60/60	95%	+70.5±6.7
lognormal	Beta(2, 5)	0.775	1.000	60/60	98%	+65.8±3.5
exponential	U[.05, 1]	0.769	1.000	60/60	97%	+121.8±9.6
exponential	U[0, 1]	0.750	1.000	60/60	97%	+121.5±8.5
exponential	Beta(2, 5)	0.713	1.000	60/60	97%	+89.8±7.0
pareto	U[.05, 1]	0.812	1.000	60/60	77%	−11.0±15.9
pareto	U[0, 1]	0.806	1.000	60/60	68%	+17.6±13.3
pareto	Beta(2, 5)	0.775	1.000	60/60	78%	+10.3±6.1

Table 6: Robustness under a product conversion-rate law $q = \text{U}[0, 1] \cdot \text{U}[0, 1]$, first-price, active-only ($N = 10$, $M = 2000$, 60 markets/cell). The main conclusions of Table 5 persist: welfare optimality is exact and distribution-free, shading pervasive, revenue rising in the large majority of markets. ΔRev is reported as mean±s.e.m. across the 60 markets.

value dist	med. shade	$\min W^*/W^{\text{opt}}$	$\Delta W \geq 0$	% ΔRev>0	mean ΔRev
uniform	0.675	1.000	60/60	100%	+50.0±1.9
lognormal	0.688	1.000	60/60	100%	+91.1±4.0
exponential	0.650	1.000	60/60	100%	+96.6±5.7
pareto	0.700	1.000	60/60	87%	+31.1±6.8

active shade if anything rising mildly ($0.73 \to 0.93$)—sharper competition thins the winners' shading without removing it. ***Distribution*** (Table 5): crossing four value laws with three q laws, welfare optimality is theoretically ***distribution-free*** and exact in every tested cell ($\min W^*/W^{\text{opt}} = 1.000$, $\Delta W \geq 0$ in $60/60$); shading is empirically pervasive across these cells (median 0.71–0.83). The revenue gain is robust (93–100% of markets) under light/medium tails but degrades under pareto (68–78%, mean occasionally slightly negative); we report both fraction and mean so the heavy-tail attenuation is visible.

E Robustness under a product conversion-rate law

The main text samples the per-impression conversion rate q directly, since only the product $q = \text{pCTR} \cdot \text{pCVR}$ enters the model. To confirm this reparameterisation hides nothing, we re-run the first-price main checks (the distribution sweep of Appendix D) with q drawn as a genuine product $q_{i,j} = \text{pCTR}_{i,j} \cdot \text{pCVR}_{i,j}$ of two independent U$[0, 1]$ factors—right-skewed and concentrated near zero, as a real click-through × conversion product would be—across the four value laws ($N = 10$, $M = 2000$, 60 markets/cell, active-only). Table 6 shows the conclusions are unchanged: welfare optimality is exact in every cell ($\min W^*/W^{\text{opt}} = 1.000$, $\Delta W \geq 0$ in $60/60$), shading is pervasive (median active shade 0.65–0.70), and revenue rises in the large majority of markets (100% under light/medium tails, 87% even under pareto—if anything better than the uniform-q sweep of Table 5, since the product law concentrates q and softens dispersion).

Table 7: Shared pCVR prediction noise ($N = 5$, 30 markets). Squared error averages over all advertisers; median shade uses advertisers active at truth and the shared-signal utility, not true-conversion utility. Welfare uses true pCVR and declines with noise. Predictions are floored at 10^{-4} but uncapped; about 9% exceed 1 at $\delta = 0.3$. No general incentive or sampling-error immunity is claimed.

pCVR noise δ	med. shade	$W_{\text{true}}/W^{\text{opt}}$	$N^{-1}\|\vec{\tau} - \vec{v}\|^2$
0.00	1.000	1.0000	0.0014
0.05	1.000	0.9982	0.0014
0.10	1.000	0.9928	$< 10^{-3}$
0.20	1.000	0.9739	$< 10^{-3}$
0.30	1.000	0.9479	$< 10^{-3}$

F Robustness to pCVR prediction error

This experiment varies a shared noisy prediction rather than allocation-curve sampling error. Advertisers evaluate value using the same noisy signal used for ranking and payment, while welfare is scored on true $\text{pCVR}^\star$. The median shade 1.000 at every tested noise level thus reflects a common-signal objective, not robustness of DSIC against true-conversion utility. True welfare decreases from 1.0000 at $\delta = 0$ to 0.9479 (about 0.95) at $\delta = 0.3$. This does not imply immunity to sampling error either, despite welfare 0.999 at $n = 100$ in Appendix B. The signal is $\max(\text{pCVR}^\star(1 + \mathcal{N}(0, \delta)), 10^{-4})$, floored but not capped: at $\delta = 0.3$, roughly 9% of entries exceed 1 (maximum about 2.3). Results at high noise therefore concern this uncapped score model, not valid conversion probabilities throughout; the $q \in [0, 1]$ welfare bound cannot be applied to these scores unchanged.

G Per-experiment summary and risks

Full experimental setup. This expands the setup summarized in Section 5.1.

- **Markets.** A *market* is one random instance of the AMG: a draw of the N advertisers' values and of the M impressions' conversion rates $q_{i,j}$. Every table averages over many independent markets (count stated per table) from a fixed seed and is exactly reproducible; production replay additionally requires counterfactual coverage, stable rivals, and compatible logging and settlement (Appendix H).
- **Realism of G.** Known exactly (E1); estimated once from a fixed log by offline replay (Section 4.1, E2); or re-estimated daily as advertisers and the platform co-adapt (Section 4, E3).
- **Payments and pricing.** Under first-price (the primary setting, ideally $\alpha_i \equiv 1$), compare the incumbent natural payment $P^{\text{nat}} = \tau G$ with the Myerson payment P^{IC} on the induced allocation. Appendix C also runs second-price on paired markets: its natural revenue is actual second-price cost, and its post-hoc Myerson audit does not establish DSIC.
- **Distributions.** Values uniform / log-normal / exponential / (heavy-tailed) pareto; q (3) uniform or $\text{Beta}(2, 5)$ (default: log-normal v, uniform q). We sample q directly since only the product $q = \text{pCTR} \cdot \text{pCVR}$ enters the model (Appendix E confirms this hides nothing).
- **Metrics.**
  - *Bid-shade ratio* $s_i = \tau_i / v_i$: median and p10/p90 over *active* advertisers—those winning positive mass, $G_i > 0$. A shut-out advertiser's shade ratio is uninformative; captions abbreviate this *active-only*. Appendix C is an explicit exception: its shade median and unilateral IC audit include all five advertisers per market; only its multiplier dispersion is active-only.
  - *Welfare and revenue changes* $\Delta W, \Delta\text{Rev}$ vs. the natural baseline.
  - *Restricted-grid deviation gain*: maximize $v_i G_i(r) - \widehat{P}_i(r)$ over the 50 E2 nodes in $[0.02 v_i, v_i]$, subtract its truthful value, and normalize by $v_i G_i(v_i)$ (excluding denominators below 10^{-9}). Payments use numerical quadrature. Average advertisers within

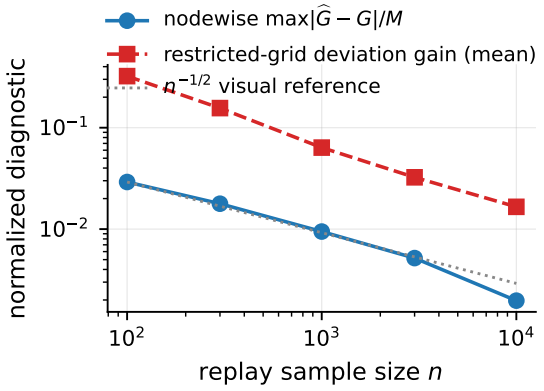

Figure 2: E2: nodewise error and restricted-grid deviation gain decrease over tested sample sizes. The dotted $n^{-1/2}$ line is a visual reference, not a fitted or validated full-domain rate. The nodewise error is divided by M; deviation gains use each advertiser's gross truthful value before averaging. Same data as Table 2.

Table 8: E3: finite-run proxy-utility search ($M = 4000$, 60 days, 15 markets/cell). Entries are final mean squared errors over advertisers active at truth, with median active shade in parentheses. Tested initial reports are $c\vec{v}$, clipped to $\bar{v} = \max_i v_i$, including $c = 1.5$. Near-unit medians at these starts and scales do not imply global convergence.

Tested starts: final active mean squared error (median shade)					
N	frac. active	$\vec{\tau}^{(0)}$=$0.3\vec{v}$	$0.5\vec{v}$	$0.9\vec{v}$	$1.5\vec{v}$
4	1.00	0.0001 (1.00)	0.0000 (1.00)	0.0001 (1.00)	0.0001 (1.00)
10	0.89	0.0032 (1.00)	0.0028 (1.00)	0.0034 (1.00)	0.0053 (1.00)
50	0.33	0.0401 (1.00)	0.0391 (1.00)	0.0347 (1.00)	0.0357 (1.00)

markets before reporting the mean and p95 of market means. Nodewise error is not the full-domain ϵ in the theorem.

– *E3 report residual*: $|A|^{-1}\sum_{i\in A}(\tau_i^{(t)} - v_i)^2$, where A is the fixed set with truthful conversions exceeding 10^{-9}. This is a mean squared error, not an unnormalized squared norm; the table also reports the active fraction.

Per-experiment summary. E1 documents shading under natural payment. E2 measures decreasing restricted-grid gains, with mean below 2% of gross truthful value and p95 of market means about 5% at $n = 10^4$; it does not validate a full-domain rate or IC bound. Second-price audits over 60 paired markets expose numerical allocation nonmonotonicity and integration-sensitive deviation gains; they do not validate DSIC. Exact first-price welfare at truth is optimal; revenue improvements remain empirical. Shared noisy-pCVR tests do not evaluate incentives under true-conversion utility. E3 uses $\lambda = 0.5$, daily $n = 3000$ subsamples of the same log, and 50 reports in $[0.02\bar{v}, \bar{v}]$, optimizing proxy utility. Its near-truthful finite trajectories do not establish global or infinite-horizon convergence. Appendix B provides analogous finite-search welfare and revenue comparisons.

The main caveat is scope: every number above is from the synthetic simulator, where ground-truth v_i makes IC slack and the welfare/revenue change measurable. Appendix H adds deployment evidence—an online A/B test showing significant gains in advertiser cost and reported targets—but does not close the truthfulness gap because advertisers' private values are unobserved. It also applies the offline-replay estimator of Section 4.1 to only a small seed set of units and extends it to the rest by an unvalidated clustering approximation, and delivers the resulting discount through the platform's existing billing-correction channel rather than the rebate that instantiates Theorem 1 (Eq. 6) directly;

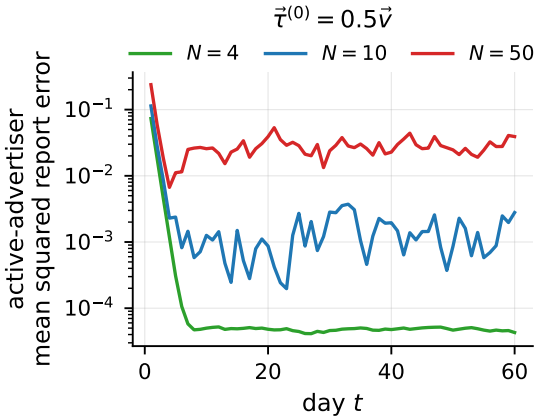

Figure 3: E3 finite trajectories from $0.5\vec{v}$, $N \in \{4, 10, 50\}$. The plotted residual is active-advertiser mean squared report error. Initial decline and a noisy floor do not isolate a cause or establish a limiting report. Same setting as Table 8.

running the offline-replay estimator itself, end-to-end, on production logs remains the next step. A secondary risk is closed-loop stability under real drift: a live market with faster non-stationarity could stress the daily re-estimation, which must satisfy uniform quadratic growth and full-domain error control for the tracking guarantee of Proposition 2; the test's own discount sequence (Appendix H) is computed once rather than re-estimated daily, a weaker, static analogue of that loop, and we flag closed-loop stability under real drift as the first thing a live deployment should monitor. We also did not add head-to-head simulator baselines for the payment-correction mechanisms of Li et al., Lv et al., and Deng et al. discussed in Appendix J: each is not a like-for-like alternative under our own assumptions rather than merely an unbuilt one, as detailed there.

H An online A/B test on a deployed tCPA product

We ran an online A/B test (2026-08-19 to 2026-09-15) on a deployed tCPA auto-bidding product matching this paper's problem setting, applying a discount to advertisers' effective billed CPA using a schedule computed from this paper's own payment formula. The market-wide treatment group (150,000 units) and control group (100,000 units) differ in cost by $+4.38\%$ ($p < 0.01$)—most of the treatment group received no subsidy, so this number is already diluted by that unsubsidized remainder; restricted to the subsidized cohort specifically—the low-budget-utilization units identified in the Findings below—the reported tCPA target rose $+2.54\%$ ($p = 0.02$) and budget rose $+4.02\%$ ($p < 0.01$) (Table 9). Because advertisers' private values are unobserved, these outcomes cannot reveal whether reports equal values and therefore cannot verify DSIC, truthful reporting, or welfare optimality. This appendix instead evaluates the predicted direction of advertiser response and gives the estimation methodology, how the discount is delivered in production and its relation to Theorem 1, and what remains open.

Estimating G and P^{IC}. For a trial multiplier $r_i \in [0, 2]$ applied to advertiser i's self-reported tCPA, we estimate $\widehat{G}_i(r_i \cdot \mathrm{tCPA}_i)$ and the Myerson payment $\widehat{P}_i(r_i \cdot \mathrm{tCPA}_i)$ by the same offline replay as Section 4.1, and store the ratio $\widehat{P}_i(r_i \cdot \mathrm{tCPA}_i)/\widehat{G}_i(r_i \cdot \mathrm{tCPA}_i)$—the corrected price per conversion at that bid level—as a lookup sequence indexed by r_i where the estimated conversion volume is positive. This is the quantity Eq. (6)'s signed settlement is built from, expressed as an average price rather than a lump transfer, the natural unit for a billing-correction channel that adjusts the price charged per conversion. Direct replay covered roughly 8,000 units, against 400,000+ actively-bidding units on the platform; every other unit's sequence is approximated by nearest-cluster lookup on two features—the advertiser's self-reported tCPA and the unit's realized spend—with all units in a cluster

sharing one sequence. This clustering approximation's own effectiveness has not been separately validated and is the least-tested link in the pipeline.

Online use. Each day, the platform locates the advertiser's current bid by its shift from its initial bid within the stored sequence and interpolates to set that day's billing target. The sequence itself, however, is computed once, up front, rather than re-estimated from fresh logs each day—different from a refreshed estimator, but neither refresh alone nor a fixed lookup guarantees convergence: Proposition 2 requires uniform growth and vanishing full-domain error for truth convergence. The test approximates a fixed-G Myerson ratio with a static lookup table rather than closing the estimation loop against a drifting market.

How the discount is applied. The computed sequence is not delivered as a standalone rebate; it is executed through the platform's existing billing-correction channel, described in Meng et al. (2026). That mechanism solves a different problem: it corrects the *ex-post* gap between an advertiser's realized and target CPA that arises from prediction error and feedback latency, via a learned joint bidding-and-pricing policy under second-price pacing, using a uniform per-unit refund shown to clear the historical deficit while leaving the forward bidding problem equivalent to the undeficited one (their Theorem 3.1). This differs from what Theorem 1 solves—correcting *ex-ante* strategic misreporting of the target itself, in closed form, under first-price—so routing our discount sequence through their channel does not inherit the exact allocation-preservation the rebate has by construction (the warning after Eq. (6)): the channel was not built to hold G_i exactly fixed in our sense, and some interaction with allocation is possible in principle. We see this as an ordinary consequence of two correction mechanisms, built at different times for different problems, sharing one settlement channel in production—not a defect of either—and a routine cost of moving a mechanism from the whiteboard into an existing billing system. We read the results below as evidence that the effect our mechanism predicts—subsidizing payment raises reported targets—holds even when delivered through infrastructure built for another purpose.

Findings. Not every unit's behavior matches the utility-maximizer model of Section 2.2: the intended response to this discount schedule is a higher reported target; Eq. (5) alone does not imply monotone comparative statics for arbitrary discounts, but many units instead lowered their bid, and adjusted their budget about as often as they adjusted τ, not only at the extremes as the model would suggest. The subsidy was targeted, in advance, at units with low budget utilization—cost at most $0.25\times$ budget, i.e. far from a binding budget constraint—a cohort fixed from prior rounds of experimentation as the regime where the predicted response concentrates, rather than selected post hoc here; on this pre-specified subset the predicted direction holds cleanly, and it is this subset that produces the aggregate result in Table 9. The cohort is fixed in advance, so it is not a post-hoc search over subgroups; but membership still conditions on budget utilization $\mathrm{cost}/\mathrm{budget} \leq 0.25$, a quantity realized during the test, so the subgroup tCPA and budget deltas should be read as descriptive of that pre-specified low-utilization regime rather than as an unconditional average treatment effect. We read this as budget non-bindingness standing in for the assumption implicit in the utility-maximizer objective $v_i G_i(\tau_i) - P_i(\tau_i)$ (Eq. (4)), which carries no budget term: the model should fit best when the budget constraint is slack, and fit worse as it binds and behavior is increasingly governed by pacing against the cap rather than by the value–payment trade-off—the regime the value-maximizer alternative deferred to Appendix I is built to capture.

Table 9: Online A/B test on a deployed tCPA product (2026-08-19 to 2026-09-15). Cost is the market-wide treatment (150k units) vs. control (100k units) delta; tCPA and budget deltas are restricted to the subsidized cohort (units with budget utilization $\mathrm{cost}/\mathrm{budget} \leq 0.25$, Findings above) within the treatment group. All figures are relative differences in advertiser-side quantities—cost is advertiser gross spend, not platform revenue—between the groups, not simulator value units.

Metric	Scope	Δ	p-value
Cost	market-wide (150k vs. 100k units)	$+4.38\%$	< 0.01
Reported tCPA	subsidized cohort	$+2.54\%$	0.02
Budget	subsidized cohort	$+4.02\%$	< 0.01

Conclusion. Despite the delivery mechanism being an approximation executed through infrastructure built for a different purpose, and despite the utility-maximizer model fitting only the low-budget-utilization subset of units, the test provides a production behavioral signal consistent with the mechanism's predicted response direction: subsidizing payment in the direction and shape our formula prescribes measurably raised the reported target and budget of the cohort it was designed to move, with a market-wide cost gain significant at $p < 0.01$. It is not an empirical validation of DSIC or Theorem 1, because true advertiser values are unobserved and the implemented channel is not the theorem's exact rebate. The gaps to close are the ones flagged above: validate the clustering approximation used beyond the $\sim$8k directly-replayed units, validate full-domain error and growth conditions under repeated estimation and drift, rather than treating daily refresh as sufficient for Proposition 2, and extend the advertiser model itself—utility maximization alone does not cover units whose budget constraint binds, and richer behavioral models for that regime, alongside the value-maximizer extension of Appendix I, are the natural next step.

I On the utility-maximizer modeling choice

The main text models each advertiser as a utility maximizer (UM), maximizing $U_i = v_i G_i(\tau_i) - P_i(\tau_i)$. This appendix records why that choice is the right primitive for isolating the meta-game, and why we defer the value-maximizer case to future work.

Why the action space runs to $\bar{v}$. The reported tCPA τ_i ranges over $[0, \bar{v}]$, not $[0, v_i]$ (Definition 1), with $\bar{v}$ a commonly known upper bound on the value support, so that over-reporting is a feasible deviation—otherwise incentive compatibility would be tested only against under-reporting—for every advertiser with $v_i < \bar{v}$. Only in the degenerate case $v_i = \bar{v}$ does $\tau_i = v_i$ already sit at the top of the action space, and there the over-report case of the DSIC argument (Theorem 1) is vacuous rather than violated. Fixing $\bar{v}$ as a public bound rather than the realized maximum $\max_k v_k$ keeps each advertiser's action set independent of rivals' private values, at no cost to the argument, which needs only that the set contains v_i and permits over-reporting.

Why utility maximization. Treating the reported tCPA as a single-parameter type is a reduction the auto-bidding literature has used before Li et al. (2020); what the UM primitive adds on top of it is a clean single-parameter *Myerson* environment in which allocation, payment, welfare, and revenue can all be answered exactly. Under UM the reported target enters a quasi-linear utility that is linear in the private per-conversion value, giving the single-parameter structure described in Section 3.1; together with the separate allocation-monotonicity assumption, this allows us to apply Myerson's payment characterization (Theorem 1). Under first-price, truthful reporting induces highest-value-wins allocation, so the welfare-optimality of Theorem 2 follows directly, and the revenue change is measured empirically.

Value maximization as future work. The value-maximizer (VM) model prevalent in the auto-bidding literature Aggarwal et al. (2019); Balseiro et al. (2021b); Li et al. (2020) is often the more realistic behavioral assumption, but it couples the incentive problem to budget/return constraints and forgoes the single-parameter structure that makes the meta-game analysis exact. Layering a value-maximizing objective onto the meta-game is a substantial extension in its own right—requiring the constraint that caps a VM's report to be modeled explicitly—and we leave it to future work. Our stance is that the UM analysis isolates the phenomena intrinsic to the meta-game formulation, on top of which the value-maximizing case can later be studied.

What a binding budget would change. The main-text analysis assumes budgets are non-binding. If advertisers strategically choose both a target and a spend cap, the action space is no longer the scalar report studied here, and incentive constraints must cover both controls. A fixed exogenous budget does not by itself destroy the single-parameter structure: one could condition on that budget and apply Theorem 1 if the resulting allocation remains monotone and the corrected payment is feasible. However, budget pacing can change the allocation and its dependence on the target, so our first-price ranking, welfare, and allocation-preserving settlement arguments do not automatically extend. Private or strategically reported budgets require a separate model. The low-utilization production cohort is closer to our non-binding-budget setting, but does not verify these theoretical conditions.

J Related Work

This section gives the full related-work discussion referenced from the Introduction (Section 1), with the individual results and contrasts stated in detail.

Reporting to one's own agent (the meta-game). The problem closest to ours is the incentive of an advertiser *reporting to its own auto-bidder*. Kolumbus and Nisan formalize this as a *meta-game* between users of regret-minimizing agents and show that misreporting one's own parameters is generically profitable Kolumbus & Nisan (2022b). In a companion paper restricted to a single auction instance, the same authors show the direction of this effect depends on the pricing rule: the second-price auction *loses* incentive compatibility while the first-price auction *gains* it Kolumbus & Nisan (2022a). Alimohammadi et al. (2023) sharpen the negative side: first- and second-price auctions—and in fact any scale-invariant, symmetric truthful auction—are *not* incentive-compatible in this report-to-agent game. The only two routes to compatibility they identify are both restrictive:

- pinning agents to a suboptimal uniform-bidding policy, or
- requiring advertisers' values to be well-aligned across queries.

Li & Tang (2022) likewise expose vulnerabilities of single-round IC under auto-bidding and show that even finding the induced equilibrium is computationally hard, while Feng et al. (2024) bound the liquid-welfare price of anarchy of the budget-selection metagame. We likewise adopt a platform-supported uniform-bidding interface, but do not claim that this restriction alone creates incentive compatibility: the natural AMG remains manipulable within that interface. Our positive contribution is to modify the payment with an aggregate rebate that restores exact DSIC within the stated strategy class, without requiring values to be aligned across queries. This departs from Liaw & Zhu (2025) in the opposite direction: they keep the second-price auction unchanged and instead relax the solution concept (to an ordinal, equilibrium-selection-based notion of IC), whereas we keep the solution concept (exact DSIC) and instead change the mechanism.

Closest in language, Li et al. (2020) give the aggregate single-parameter reduction we adopt: a reported target maps to episode-level conversions and cost, with a per-impression $\Leftrightarrow$ aggregate IC equivalence. But the problem they solve is structurally the reverse of ours. In their setting:

- the platform *controls* delivery, reshaping each advertiser's bidding environment through PID controllers to *force* the realized cost and conversions onto a platform-chosen curve;
- that conversion curve is treated as a known, designed object; and
- for profit-maximizing advertisers, the mechanism must know the private value v to escape a triviality barrier—with v unknown, only trivial mechanisms exist.

We take the opposite stance on every point:

- the allocation G is *exogenous and unknown*, fixed by the underlying auction and rivals and left untouched by our rebate;
- v is kept private throughout; and
- our central difficulty is accordingly not *controlling* delivery but *estimating* G from logs—the dynamics question Li et al. leave as future work.

With exact allocation information, truthful reporting is weakly dominant regardless of rivals. This static incentive property does not establish behavioral convergence; estimated payments and co-adaptation require the separate conditions above. The mechanical link between the two papers is minor by comparison: the rent their forced payment $P = \tau G$ omits, $\int_0^\tau G$, is exactly what our Myerson rebate returns. Because their only non-trivial construction presupposes exactly the private-value knowledge our DSIC result is designed to avoid needing, and their prescribed payment $cpa_a(t_a) = t_a$ coincides with our own already-tabulated natural baseline $P^{\mathrm{nat}} = \tau G$ (Table 1), there is no distinct, fairly-constructible payment rule of theirs left to add as a separate row in our experiments.

Value- versus utility-maximization. A growing line models auto-bidders as value maximizers under ROI/CPA constraints rather than classical utility maximizers Aggarwal et al. (2019); Balseiro et al. (2021b). This line studies equilibria, price-of-anarchy bounds, and welfare loss Deng et al. (2021), as well as learning dynamics of budget-constrained bidders Balseiro & Gur (2019). The value-maximizer primitive itself was founded by Wilkens et al. (2016), who show it reshapes feasible and optimal mechanisms and renders GSP truthful. More generally, non-quasi-linear agents optimally reduce to a single value-scaling (ROI) report in quasi-linear mechanisms Babaioff et al. (2021). Whether this behavioral shift should change the design primitive is an active question Aggarwal et al. (2024). We adopt the utility-maximizer primitive throughout, which isolates a clean single-parameter Myerson environment; extending the meta-game analysis to value-maximizing agents is left to future work.

Optimal and truthful auction design for auto-bidding. A recent line designs revenue-optimal or truthful auctions directly for auto-bidders, typically in a Bayesian single-auction model where each advertiser has a value v_i, an RoS/CPA target τ_i, and possibly a budget Aggarwal et al. (2024). Golrezaei et al. (2021) characterize the revenue-optimal auction for utility-maximizing bidders under RoS/ROI constraints via *modified virtual values*. Heymann (2019) shows a CPA constraint already breaks second-price truthfulness and derives the constrained equilibrium in a single auction. Lv et al. (2023) give a single-parameter Myerson-style DSIC characterization for per-impression auctions with an *ex-post* RoS constraint and—closest to us in language—implement the DSIC payment as *the Myerson payment with a per-bidder rebate*, though their revenue-optimal auction is worked out only for a single, randomized bidder; the rebate there is an in-auction payment rule serving one purpose-built auction, not an aggregate transfer layered on an arbitrary auction. Balseiro et al. (2021b) map the revenue-optimal landscape across value/utility maximization and private-value/private-target structures, where for utility maximizers the optimum requires ironing or a *subsidy*. Two contrasts recur across this line:

- *Direction.* Their subsidies exist to restore feasibility of an RoS constraint, whereas our rebate returns the advertiser's information rent to secure DSIC.
- *Form.* Their optima are solved per instance from the dual and depend on the value prior, whereas our rebate is closed-form, deterministic, prior-free, and allocation-preserving. Its DSIC guarantee extends to another underlying pricing rule whenever that rule's induced allocation is monotone; Appendix C gives a second-price numerical audit with unresolved monotonicity and integration issues, not a validation of this condition.

Setting Lv et al.'s ROI multiplier $M_i = 1$—removing the ex-post constraint their rebate exists to satisfy—collapses their payment $p_i(v) = M_i\tilde{p}_i(v) - \max_{0\le z\le v}\{M_i\tilde{p}_i(z) - zx_i(z)\}$ to the plain Myerson payment exactly, up to relabeling ($\tau_i \leftrightarrow v_i$, $G_i \leftrightarrow x_i$)—so absent that constraint, which our setting does not have, their construction is not a distinct mechanism to benchmark against. Robust designs in the same world instead tune reserves and boosts to bound welfare/revenue loss across bidder types Balseiro et al. (2021a), and randomization can push efficiency past VCG Liaw et al. (2023); both stay at the per-impression layer rather than the meta-game.

Estimating incentives from data, and dynamics. Because our allocation G is unknown and estimated from logs, we connect to several lines of work that ground incentives empirically or study the dynamics of auto-bidding markets.

- *A data-driven IC metric.* Deng et al. (2020) define a metric computable directly from black-box auction logs: an envelope-condition test equal to the gap between the payment that truthfully implements the allocation and the one actually charged. This gap is exactly the information rent our rebate returns. This is a diagnostic metric rather than a mechanism, and the payment-like object inside its derivation is the identical Myerson identity our rebate already instantiates—so it offers a way to *measure* the rent we return, not an alternative rule for returning it.
- *Dynamics of pacing equilibria.* Budget-pacing equilibria are essentially unique and convex-programmable under first price Conitzer et al. (2022a) but PPAD-complete under second price Chen et al. (2021), and efficiency/regret guarantees can be obtained even *without* convergence Gaitonde et al. (2023). Our exact payment instead gives a static weak-dominance

result at the meta-game level. It does not remove the need to analyze estimated-payment learning dynamics or guarantee selection among indifferent reports.

- *Optimal auctions and single-parameter types.* Our study of revenue under a value-dispersed market connects to the classical theory of optimal auctions with financially constrained buyers Che & Gale (2000). Our single-scalar-type modeling also heeds the caution of the FedEx problem Fiat et al. (2016), where a seemingly low-dimensional type already breaks single-parameter Myerson; whether tCPA is a genuine scalar type is a question our reduction supports treating affirmatively in the utility-maximizer setting.
- *Commitment.* Perlroth & Mehta (2023), complementarily, explain tCPA adoption through auctioneer commitment—a lens orthogonal to our rebate, which presupposes commitment.