HIGH-PRECISION FIRST-PRICE THROTTLING EQUILIBRIA VIA POSITIVE-TREE CONTINUATION

Anonymous authors
Paper under double-blind review

ABSTRACT

First-price throttling lets budget-constrained advertisers randomize participation without changing bids, but ill-conditioned payment Jacobians complicate high-precision equilibrium computation. To overcome this obstacle, we represent a normalized continuation direction by positive tree polynomials. This yields a deterministic algorithm that computes a rational δ-approximate equilibrium in $\mathrm{poly}(N, \log(1/\delta))$ time for every finite rational first-price market of encoding length N. We then characterize extensions through payment evaluation and derive computable revenue and coordinate-error certificates. Finally, learning results distinguish residual accuracy from participation accuracy, with matching sample bounds for a cyclic boundary family.

1 INTRODUCTION

Budget-constrained advertisers use *throttling* to control participation: advertiser i enters each auction with probability θ_i and keeps her bid unchanged. In first-price auctions, increasing her participation can raise her expected payment, while competition can lower it. An equilibrium exhausts the budget of every advertiser who participates with probability below one. This produces a nonlinear complementarity system with a unique solution, but uniqueness alone supplies neither an efficient high-precision algorithm nor a useful error bound.

For finite first-price markets, Chen et al. (2021) characterized the equilibrium as the coordinatewise largest feasible vector and gave decentralized multiplicative dynamics with running time polynomial in $1/\delta$, where δ is residual accuracy. Chen & Li (2026) explicitly asked whether the dependence could be improved to $\mathrm{poly}(n, m, \log(1/\delta))$. We answer this question affirmatively. This distinction matters even for a unique equilibrium: the solution can be irrational, and the payment Jacobian can become arbitrarily ill conditioned near full participation.

The Optimization Obstacle. Root-finder convergence alone does not give a global running-time bound. We instead *normalize* the continuation direction, expressing it as the stationary vector of an augmented graph and hence a ratio of positive directed-tree polynomials. Positivity prevents cancellation in a complex neighborhood controlled solely by the degree in each pair $(\theta_i, 1-\theta_i)$. This yields certified Taylor steps, including after coordinates freeze. Restarting at rational iterates and bounding payment-ratio drift controls global error without tracking one exact trajectory. No bid-separation or condition-number bound is needed.

Main Theorem. Approximation preserves budget feasibility exactly and relaxes only the budget-exhaustion condition; Section 2 gives the formal definition.

Theorem 1.1 (High-Precision Computation)**.** *Given a rational first-price market of encoding length N and a precision parameter $p \geq 1$ in unary, a deterministic algorithm outputs a rational 2^{-p}-approximate throttling equilibrium in $\mathrm{poly}(N, p)$ time. Equivalently, the complexity for residual accuracy δ is $\mathrm{poly}(N, \log(1/\delta))$, apart from reading the encoding of δ.*

The theorem permits ties and imposes no lower bound on the Jacobian's smallest singular value. It guarantees residual accuracy, not $O(\delta)$ coordinate error. The earlier multiplicative update can require $\Omega(1/\delta)$ rounds even with one bidder and one good (Appendix D); this is a lower bound for that update rule, not for all equilibrium algorithms. Our algorithm is centralized.

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.

Further Results. Two entry conditions extend its positive-tree structure beyond single-slot auctions; an equivalence to pointed reliability then identifies the payment-evaluation obstacle for succinct matroids (Theorem 4.2). We next ask what the computed residual certifies. Theorem 5.1 gives a computable bound on revenue and coordinate error, exact on a cyclic family where residual δ coexists with coordinate error $\delta^{1/n}$. Finally, for unknown markets, Theorems 5.3 and 5.4 distinguish learning the equilibrium conditions from identifying participation probabilities at a boundary.

1.1 Related Work

Auction Equilibrium Computation. First-price *pacing* scales bids and admits a convex-program characterization (Conitzer et al., 2022). Throttling instead randomizes participation; the finite-market dynamics of Chen et al. (2021) are the direct computational baseline for Theorem 1.1. Chen et al. (2026) establish linear convergence for diffuse autobidding under distributional regularity and a positive lower bound on equilibrium multipliers. Our guarantee covers every finite rational first-price throttling market, including ties and boundary degeneracy, without a condition-number assumption. Second-price pacing and throttling face PPAD-hardness even at constant approximation (Chen et al., 2024a; Chen & Li, 2026; Liu, 2026), although exact second-price pacing becomes tractable with a constant number of buyers or goods (Yan et al., 2026; Huang et al., 2026).

Equilibrium Learning and Budget Control. Market-equilibrium algorithms address large instances and unknown item populations (Gao & Kroer, 2020; Gao et al., 2021). Statistical work studies Fisher-market inference (Liao et al., 2023), pacing-market A/B tests (Liao & Kroer, 2023), bootstrapping (Liao & Kroer, 2024), and interference (Liao et al., 2025). Online autobidding instead studies regret and welfare over repeated auctions (Wang et al., 2023; Balseiro et al., 2024; Lucier et al., 2024); Chen et al. (2024b) analyze dynamic throttling for a single advertiser. Our statistical results estimate a market-wide throttling equilibrium from i.i.d. contexts with fully observed bids and priorities. They separate learning the equilibrium conditions from identifying participation probabilities, with sharp sample bounds at a nonregular boundary. This observation model does not address censored feedback or pathwise budget control.

Analytic and Combinatorial Methods. The matrix-tree theorem (Chaiken, 1982) and high-order Taylor methods (Bournez et al., 2012) underlie our algorithm. The new ingredients are an analytic bound controlled by polynomial degree for the normalized tree field and a polynomial-time rational implementation. For matroid auctions, we connect high-precision equilibrium computation to exact pointed reliability by reductions in both directions, building on classical matroid and reliability theory (Oxley, 2011; Provan, 1986).

2 Preliminaries

There are n bidders and m goods. Bidder i has budget $B_i > 0$ and bids $b_{ij} \geq 0$, all rational. Decreasing bids and fixed tie-breaking specify a strict priority order for each good; write $k \succ_j i$ when k precedes i. Participation is independent across bidders, with vector $\theta \in [0,1]^n$. Expected first-price payments are

$$p_{ij}(\theta) = \theta_i b_{ij} \prod_{k \succ_j i} (1 - \theta_k), \qquad P_i(\theta) = \sum_j p_{ij}(\theta). \tag{1}$$

The input length N includes the binary encodings of bids, budgets, and priorities. Exact equilibrium is the complementarity system

$$0 \leq \theta \leq \mathbf{1}, \qquad B_i - P_i(\theta) \geq 0, \qquad (1 - \theta_i)(B_i - P_i(\theta)) = 0 \quad \text{for all } i. \tag{2}$$

Definition 2.1 (Approximate Equilibrium (Chen et al., 2021))**.** For $0 < \delta < 1$, a δ-approximate throttling equilibrium satisfies

$$P_i(\theta) \leq B_i, \qquad \theta_i < 1 - \delta \implies P_i(\theta) \geq (1 - \delta) B_i \quad \text{for all } i. \tag{3}$$

Budget feasibility is exact; only complementarity is relaxed. We call δ *residual accuracy*, as distinct from coordinate accuracy $\|\theta - \theta^\star\|_\infty \leq \eta$. At a feasible vector, the residual $\max_i \min\{1 - \theta_i, 1 - P_i(\theta)/B_i\}$ is at most δ exactly when (3) holds.

3 POSITIVE-TREE CONTINUATION

The algorithm raises participation while keeping active payments approximately proportional to budgets. We construct this path, implement rational steps, and bound their cumulative error.

As preprocessing, set $\theta_i = 1$ for bidders with $M_i := \sum_j b_{ij} = 0$ and remove them; their presence cannot change any positive payment. Return $\mathbf{1}$ if none remain, and otherwise relabel the retained bidders as $[n]$. Their coordinates remain in $(0, 1)$ throughout the computation.

3.1 NORMALIZED CONTINUATION PATH

Let A be the active bidders; all other probabilities are frozen. For every retained bidder, put $u_i = \log(\theta_i/(1-\theta_i))$ and let $H_A = [\partial P_i/\partial u_k]_{i,k\in A}$. Differentiation gives

$$H_{ii} = (1-\theta_i)P_i, \qquad H_{ik} = -\theta_k \sum_{j:k\succ_j i} p_{ij} \le 0 \quad (i \ne k). \tag{4}$$

Entry by an active bidder cannot lower the active group's total revenue: she either displaces a weakly lower active bid or adds revenue to the group. The event that all other bidders are absent makes the increase strictly positive. Thus H_A is strictly column diagonally dominant, so it is a nonsingular M-matrix and $H_A^{-1} \ge 0$ (Berman & Plemmons, 1994).

Introduce a target budget fraction $0 < t < 1$. We initialize P_i/t close to B_i and preserve this ratio along the exact path. With $u_0 = \log(t/(1-t))$, differentiating the desired relation $P_i \propto t$ gives $H_A\, \mathrm{d}u_A/\,\mathrm{d}u_0 = (1-t)P_A$. This direction may be large, so we reparameterize time by requiring the sum of all moving logit speeds to be one. In the new time variable s, the equations are

$$H_A y = (1-t)aP_A, \qquad a + \mathbf{1}^\top y = 1, \qquad \dot{u}_0 = a, \quad \dot{u}_i = y_i. \tag{5}$$

The inverse-positivity just established gives $a > 0$, $y_i > 0$, and all logit speeds at most one. The chain rule yields

$$\dot{P}_A = (1-t)aP_A, \qquad \dot{t} = t(1-t)a, \qquad \frac{\mathrm{d}}{\mathrm{d}s}\log\frac{P_i}{t} = 0 \quad (i \in A). \tag{6}$$

Active payments therefore increase proportionally to t. Frozen bidders' payments can only decrease as competitors' participation increases. Normalizing the sum of logit speeds also makes path length measurable: $\Phi = u_0 + \sum_i u_i$ increases at unit speed, with frozen logits retained as constants.

When t approaches one, active bidders nearly exhaust their budgets. A bidder approaching full participation earlier satisfies the other branch of approximate complementarity and can be frozen without compromising feasibility. Budgets thus determine the starting ratios, while the payment map determines the continuation direction.

3.2 ANALYTIC STEP BOUNDS FROM POSITIVE TREES

Bounded real speeds alone do not bound the high derivatives needed for high-precision integration. We obtain such bounds from a positive representation of the entire normalized field. For distinct active bidders k, i, define the arc weights

$$w_{k\to i} = -H_{ik}, \qquad \kappa_k = H_{kk} - \sum_{i\in A\setminus\{k\}} w_{k\to i} > 0.$$

Use arcs $k \to i$ of weight $w_{k\to i}$. Augment this graph by a vertex 0, with arcs $k \to 0$ of weight κ_k and $0 \to i$ of weight $(1-t)P_i$. At vertex i, outgoing flow is $H_{ii}y_i$ and incoming flow is $\sum_{k\ne i} w_{k\to i}y_k + (1-t)P_i a$. Thus (5) is precisely stationary flow balance with total mass one. If T_v sums the weights of directed spanning trees oriented toward v, the directed matrix-tree theorem gives

$$a = T_0/Z, \qquad y_i = T_i/Z, \qquad Z = \sum_{v\in\{0\}\cup A} T_v. \tag{7}$$

Every arc weight has nonnegative coefficients in the logistic factors $x_k, 1-x_k$ ($x_0 = t$, $x_i = \theta_i$); for κ_k, expand the nonnegative expected revenue increment used above. Each arc has degree at most two

in its tail's pair and at most one in every other pair. A tree has at most n arcs and one per nonroot tail, giving degree at most $D = n + 1$ in each pair.

The analytic argument applies to any positive polynomial ratio with this degree bound; the tree representation supplies the required structure here.

Lemma 3.1 (Analytic Bound for Positive Polynomial Ratios)**.** *Let T_v be a finite family of polynomials represented with nonnegative coefficients in factors $x_k, 1 - x_k$, with at least one nonzero member. Suppose every monomial has degree at most an integer $D \geq 1$ in each pair. For $x_k(u) = e^{u_k}/(1 + e^{u_k})$, the ratios $F_v(u) = T_v(x(u))/\sum_w T_w(x(u))$ extend holomorphically around every real u to $\|z\|_1 < \rho := 2/D$, with*

$$\sum_v |F_v(u + z)| \leq 2 \quad (\|z\|_1 \leq \rho). \tag{8}$$

The radius and bound are independent of the coefficient magnitudes.

Proof. Multiply each monomial by powers of $x_k + (1 - x_k) = 1$ to give every pair degree D. Substituting $x_k = e^{u_k}/(1 + e^{u_k})$ cancels common denominators. Centering exponents at $D/2$ gives

$$F_v(u) = \frac{S_v(u)}{\sum_w S_w(u)}, \qquad S_v(u) = \sum_\alpha c_{v,\alpha} e^{(\alpha - D\mathbf{1}/2)\cdot u}, \quad c_{v,\alpha} \geq 0, \quad 0 \leq \alpha_k \leq D.$$

At $u + z$, every nonzero term has argument of magnitude at most $(D/2)\|\operatorname{Im} z\|_1 \leq 1$, hence real part greater than half its modulus. The denominator cannot vanish, and summing these inequalities gives (8). □

For the tree ratios (7), $D = n + 1$, so $F = (a, y)$ has radius $\rho = 2/(n + 1)$ independently of bids, budgets, and $\det H_A$. Frozen factors are positive constants; the same bound therefore holds after every active-set change. This conclusion uses the positive representation and its degree bound, beyond inverse-positivity of H_A.

Cauchy's estimate in the same ℓ_1 norm bounds $\|DF\|$ by $4/\rho$ on the half-radius ball. Picard iteration then gives an analytic solution on the complex time disc $|s| \leq 2h$, where $h = 1/[8(n + 1)]$. For each moving probability x, let $\omega = x/(1 - x)$ and denote its order-q Taylor polynomial by $\omega_{[q]}$. Then

$$|\omega(\sigma) - \omega_{[q]}(\sigma)| \leq 2\omega(0)2^{-q} \quad (0 \leq \sigma \leq h). \tag{9}$$

Thus local logit error ξ is achievable with $q = O(\log(1/\xi))$. Lemmas A.5 and A.7 give the constants and account for rounding.

Appendix A.3 illustrates this distinction: in a two-bidder market, the inverse Jacobian diverges while the normalized direction has a finite limit. The same analytic bound also covers the rational restarts used by the algorithm.

3.3 A Rational Algorithm

Set $\varepsilon = 2^{-p}$. Choose the largest dyadic power t_0 satisfying

$$t_0 \leq \min\left\{\frac{1}{4}, \frac{\varepsilon}{1000n} \min_i \frac{M_i}{B_i}\right\}, \qquad \theta_i^0 = t_0 \frac{B_i}{M_i}. \tag{10}$$

Then $\theta_i^0 \leq \varepsilon/(1000n)$, so each conditional payment differs from its singleton value M_i by at most an $\varepsilon/1000$ fraction. Consequently, $(1 - \varepsilon/1000)B_i \leq P_i(\theta^0)/t_0 \leq B_i$. This supplies the budget-proportional starting point needed by the path.

Algorithm 1 uses a single subroutine, $\mathsf{Step}_h(t, \theta; A)$: advance the active flow for $h = 1/[8(n + 1)]$ by rational Taylor expansion and odds rounding, returning a rational endpoint with coordinatewise logit error at most ξ. The tolerance and rounding grid are fixed from the input before the loop; Appendix A.5 gives their exact values. Frozen coordinates are unchanged by the subroutine.

Algorithm 1 Positive-tree continuation

Require: Rational market after zero-bidder preprocessing and accuracy $\varepsilon = 2^{-p}$.
1: Initialize $(t, \theta) = (t_0, \theta^0)$ by (10); set $A = [n]$.
2: **while** $A \neq \varnothing$ and $t < 1 - \varepsilon/2$ **do**
3: $\quad(\widetilde{t}, \widetilde{\theta}) \leftarrow \mathsf{Step}_h(t, \theta; A)$.
4: $\quad(t, \theta) \leftarrow \max\{(t, \theta), (\widetilde{t}, \widetilde{\theta})\}$ coordinatewise.
5: $\quad A \leftarrow \{i \in A : \theta_i < 1 - \varepsilon/2\}$.
6: **end while**
7: **return** θ, restoring removed zero bidders at one.

Implementing the Certified Step. The tree formula is used for analysis; the subroutine computes the direction by solving the rational linear system

$$\mathcal{A}(t, \theta) \begin{pmatrix} y \\ a \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \qquad \mathcal{A} = \begin{pmatrix} H_A & -(1-t)P_A \\ \mathbf{1}^\top & 1 \end{pmatrix}, \qquad \det \mathcal{A} = Z > 0. \tag{11}$$

Write local series as $z(s) = \sum_{r \geq 0} z^{[r]} s^r$, and put $v = (y, a)$ and $b = (0, \ldots, 0, 1)^\top$. Coefficient extraction gives

$$\mathcal{A}^{[0]} v^{[r]} = \mathbf{1}_{r=0} b - \sum_{j=1}^{r} \mathcal{A}^{[j]} v^{[r-j]}. \tag{12}$$

At order r, state coefficients through order r determine $\mathcal{A}^{[0]}, \ldots, \mathcal{A}^{[r]}$; all direction coefficients on the right are already known. Solve the same matrix for $v^{[r]}$, then use $\dot{t} = t(1-t)a$ and $\dot{\theta}_i = \theta_i(1-\theta_i)y_i$ to obtain the state coefficients of order $r+1$. Lemma A.7 certifies the truncation and rounding errors. No spanning trees or expanded multivariate polynomials are enumerated.

Use order $q = \lceil \log_2(8/\xi) \rceil$, evaluate the odds series at h, and round each result to the fixed dyadic grid, with ties rounded down. Convert back by $x = \omega/(1+\omega)$ exactly. Rounding odds controls relative error even near the probability boundary; the log of the ratio between rounded and exact odds is the logit error. The coordinatewise maximum in the algorithm then enforces monotonicity. Because the exact endpoint is no smaller than the old state, this operation cannot enlarge the local error. The loop uses only rational comparisons and requires no numerical logarithms or event-time search.

3.4 Global Error and Running Time

We first impose $\xi \leq h/[4(n+1)]$ and bound how many steps can occur; we then choose a smaller tolerance to control their cumulative error. Every moving coordinate starts below $1 - \varepsilon/2$ and gains at most $h + \xi \leq 1/8$ in logit. It therefore remains below $1 - \varepsilon/4$, including on the step when it freezes; frozen coordinates retain this bound. Monotonicity bounds all coordinates away from zero by their initial values. If L bounds each input numerator and denominator length, initialization gives a computable $Q = O(L + p + \log(nm))$ such that every stored logit lies in $[-Q, p+2]$.

Each exact step increases Φ by h, and rounding and clamping leave an increase of at least $h - (n+1)\xi \geq 3h/4$. The available range of Φ is at most $(n+1)(Q+p+2)$. Hence

$$K_* := \left\lceil \frac{2(n+1)(Q+p+2)}{h} \right\rceil = O(n^2(Q+p))$$

is an a priori step bound valid for every tolerance satisfying the initial constraint. We can therefore fix

$$\xi = \min\left\{ \frac{h}{4(n+1)}, \frac{\varepsilon}{1000(n+1)(K_*+1)} \right\} \tag{13}$$

before the loop, without a circular dependence on the number of steps.

Residual Drift and Budget Feasibility. A generic trajectory-error bound can amplify local errors through the Jacobian of the flow. Here correctness is controlled by payment ratios, which are invariant on every exact segment, including one restarted from a rounded state. For a nonnegative multilinear payment, $|\partial \log P_i / \partial u_k| \leq 1$ (Lemma A.9); also $|\partial \log t / \partial u_0| \leq 1$. Thus a rounded step changes

$\log(P_i/t)$ by at most $(n+1)\xi$ while i is active. Summing these changes, rather than comparing entire trajectories, bounds the total drift by $\varepsilon/1000$. Initialization therefore gives

$$(1-\varepsilon/1000)e^{-\varepsilon/1000}tB_i \le P_i \le e^{\varepsilon/1000}tB_i$$

until bidder i freezes. Since $t \le 1-\varepsilon/4$, the upper bound is strictly below B_i; frozen payments never increase. If any bidders remain active at termination, then $t \ge 1-\varepsilon/2$, and the lower bound exceeds $(1-\varepsilon)B_i$. Frozen bidders already have $\theta_i \ge 1-\varepsilon/2$. The returned vector therefore satisfies (3).

Running Time. The coefficient recurrence uses $O(mn^3(q+1)^3)$ arithmetic operations per step, with $q = O(p + \log(n(K_*+1)))$. Together with the bound K_*, this gives a polynomial number of operations in (N, p). All intermediate rational values have polynomial encoding length; the details are in Appendix A.4 and the complexity analysis in Appendix A. This proves Theorem 1.1. The polylogarithmic dependence on $1/\delta$ comes from geometric Taylor accuracy along a path of polynomial logit length, without a condition-number bound on the payment Jacobian.

Accuracy and Verification. The stopping rule allows either near budget exhaustion or near full participation. Bidders can therefore freeze without first identifying which bidders participate fully at the exact equilibrium. The returned vector can be checked by evaluating payments and testing (3); no exact equilibrium is needed to verify the guarantee. A small residual alone does not identify participation accurately. Section 5 provides revenue and coordinate certificates that clarify what the computed approximation certifies. The polynomial bound establishes tractability; practical speed relative to simpler updates remains to be evaluated.

4 Payment Conditions and the Evaluation Boundary

The proof above uses both a sign structure and an efficient way to evaluate payments. We now separate these requirements to determine which other auction rules admit the same algorithm. For participating set S, let $\pi_i^j(S) \ge 0$ be bidder i's payment in auction j, with $\pi_i^j(S) = 0$ for $i \notin S$. Averaging over independent participation gives expected payments that are affine in each probability separately, hence multilinear. Assume, for every $k \notin S$,

$$\begin{aligned} &\pi_i^j(S \cup \{k\}) \le \pi_i^j(S) \quad (i \in S), \\ &\pi_k^j(S \cup \{k\}) \ge \sum_{i \in S} \big[\pi_i^j(S) - \pi_i^j(S \cup \{k\})\big]. \end{aligned} \tag{14}$$

Entry weakly lowers incumbent payments and does not lower total revenue. These conditions are preserved by nonnegative sums and expectations.

Theorem 4.1 (Structural and Computational Extension)**.** *With positive budgets, every market satisfying* (14) *has a unique exact equilibrium, the coordinatewise largest feasible vector. Equilibrium participation and revenue are nondecreasing in budgets. If a single polynomial-time algorithm evaluates payments exactly at rational probabilities, a rational δ-approximate equilibrium is computable in* $\mathrm{poly}(N, \log(1/\delta))$ *time.*

To see the structural implication, consider the revenue of a group A when $k \in A$ enters. The entrant covers the losses of all incumbents, hence also the losses of those in A. The group's revenue increment is therefore nonnegative. If k has positive singleton payment and all other bidders are absent, the increment is strictly positive. Expanding this expectation gives exactly the positive column surplus needed by the augmented graph. The same degree bound then yields the same analytic radius.

Competitor monotonicity also makes the feasible set closed under coordinatewise maximum: raising competitors' participation cannot violate a bidder's budget if her own probability is held fixed. The largest feasible vector satisfies complementarity, and strict increase of a differing group's revenue proves uniqueness. For the computational claim, polynomial-time point evaluation recovers truncated series by univariate interpolation, so the extension does not require an explicit monomial expansion of the payments (Appendix B.1).

In first-price matroid auctions, greedy allocation accepts bidders in bid order subject to independence. Inserting one bidder removes at most one lower-priority winner, proving (14). For a distinguished

matroid element e, sample all other elements independently with probabilities x and define pointed reliability by

$$Q_{M,e}(x) = \Pr_{S\sim x}[e \notin \mathrm{cl}_M(S)]. \tag{15}$$

Here cl_M denotes matroid closure: the event means that adding e increases the rank of S. If H_i is bidder i's higher-priority set, her payment in one auction is $P_i(\theta) = \theta_i b_i Q_{M|(H_i\cup\{i\}),i}(\theta_{H_i})$, where $M|E$ restricts the matroid to elements in E.

Theorem 4.2 (Equivalence to Reliability Evaluation)**.** *Fix a uniform succinct matroid representation with an explicitly indexed ground set, polynomial-time restriction, deletion, contraction and loop testing, polynomial-size transformed representations, and singleton matroids. Exact rational evaluation of* (15) *and high-precision equilibrium computation are polynomial-time Turing equivalent. Laminar-matroid and position-auction markets admit* $\mathrm{poly}(N, \log(1/\delta))$ *algorithms; graphic-matroid markets are* #P*-hard at this precision.*

Reliability queries supply payment evaluations. Conversely, private auctions pin ordinary participation probabilities sufficiently accurately that one high-precision equilibrium query recovers an exact reliability value by rational rounding. Laminar reliability admits a rank-distribution dynamic program; graphic reliability is two-terminal network unreliability (Provan, 1986). Position payments are nonnegative sums of uniform-matroid payments. Appendix B.2 proves these statements. The hardness concerns succinct input and high precision, not fixed constant δ or an explicitly listed exponential payment table.

5 Residual Certificates and Learning

5.1 Revenue and Coordinate Certificates

The evaluation results explain when a small residual is computable. We now examine what that residual says about the equilibrium itself. Budget feasibility provides an order relation to the exact equilibrium; revenue then turns this relation into a computable error bound. Let $\theta^\star$ be the exact equilibrium, $V = \sum_i P_i$, and $V^\star = V(\theta^\star)$. Write $\theta^{i\leftarrow a}$ for replacement of coordinate i by a, and define

$$\begin{aligned} c_i(\theta) &= P_i(\theta^{i\leftarrow 1}), \qquad g_k(\theta) = V(\theta^{k\leftarrow 1}) - V(\theta^{k\leftarrow 0}) = \partial_k V(\theta), \\ \mathcal{E}(\theta) &= \sum_i \big(\min\{B_i, c_i(\theta)\} - P_i(\theta)\big). \end{aligned} \tag{16}$$

Here c_i is bidder i's payment conditional on entry with competitors fixed, and g_k is the marginal revenue from increasing θ_k. The quantity $\mathcal{E}$ sums the additional payments that budgets and these conditional payments could individually permit. Competitors' future entry can only reduce each conditional payment, which makes this sum an upper bound on the remaining equilibrium revenue gap.

Theorem 5.1 (Computable Revenue and Coordinate Certificate)**.** *In any market satisfying* (14)*, every feasible* θ *satisfies* $\theta \le \theta^\star$ *and*

$$0 \le g_k(\theta)(\theta_k^\star - \theta_k) \le V^\star - V(\theta) \le \mathcal{E}(\theta) \quad \textit{for every } k. \tag{17}$$

For $g_k(\theta) > 0$*, the coordinate error is at most* $\min\{1-\theta_k, \mathcal{E}(\theta)/g_k(\theta)\}$*. If* θ *is a* δ*-approximate equilibrium, then also*

$$V^\star - V(\theta) \le \delta V^\star, \qquad \mathcal{E}(\theta) \le \delta \min\left\{\sum_i B_i, \frac{V(\theta)}{1-\delta}\right\}.$$

Proof. The largest-feasible characterization gives $\theta \le \theta^\star$. Put $q_i = \min\{B_i, c_i(\theta)\}$. Competitor monotonicity gives $P_i(\theta^\star) \le q_i$. Since V is nondecreasing and affine in each coordinate,

$$g_k(\theta)(\theta_k^\star - \theta_k) = V(\theta^{k\leftarrow\theta_k^\star}) - V(\theta) \le V^\star - V(\theta) \le \mathcal{E}(\theta).$$

For an approximate equilibrium, either $\theta_i \ge 1-\delta$, giving $P_i = \theta_i c_i \ge (1-\delta)q_i$, or complementarity gives $P_i \ge (1-\delta)B_i \ge (1-\delta)q_i$. Thus $V \ge (1-\delta)\sum_i q_i \ge (1-\delta)V^\star$ and $\mathcal{E} = \sum_i (q_i - P_i) \le \delta \sum_i q_i$. Using $\sum_i q_i \le \min\{\sum_i B_i, V/(1-\delta)\}$ proves both bounds. □

At rational θ, the certificate uses $O(n)$ exact payment-vector evaluations and does not require $\theta^\star$. It certifies a coordinate target whenever all computed error bounds meet that target. If $g_k = 0$, retain the bound $1 - \theta_k$; zero marginal revenue need not imply zero coordinate error. Small residual alone does not ensure small coordinate bounds, because the revenue marginals may also be small. The following family makes this obstruction exact.

Proposition 5.2 (A Sharp Residual–Coordinate Separation)**.** *For every $n \geq 2$, a rational market with n bidders, n goods, and unit bids and budgets has exact equilibrium $\mathbf{1}$, while $(1-r)\mathbf{1}$ is an r^n-approximate equilibrium with coordinate error r, for every $0 < r < 1$. The certificate in Theorem 5.1 is exact on this family.*

Proof. Use the n cyclic shifts of a priority order. Each bidder appears once at each rank, so $P_i(t\mathbf{1}) = t\sum_{j=0}^{n-1}(1-t)^j = 1-(1-t)^n$. Full participation is the exact equilibrium. At $t = 1-r$, payments equal $1-r^n$, which proves the residual claim. Moreover, $V(\theta) = n[1-\prod_i(1-\theta_i)]$ and, at $(1-r)\mathbf{1}$,

$$\mathcal{E}(\theta) = nr^n, \qquad g_k(\theta) = nr^{n-1}, \qquad \mathcal{E}(\theta)/g_k(\theta) = r.$$

□

For $n = 20$ and $r = 1/2$, residual 2^{-20} coexists with coordinate error $1/2$. This is a limitation of residual-to-coordinate conversion, not a computational lower bound: the exact equilibrium in this example is known, and $\log(1/\eta^n) = n\log(1/\eta)$ still gives a polynomial dependence on n and $\log(1/\eta)$. With an unknown market, however, the same boundary geometry has a statistical consequence.

5.2 Residual Learning under Full Observation

When payments are estimated from data, the residual–coordinate distinction also determines what can be learned. We give a feasible residual guarantee, then use the cyclic geometry to quantify the cost of coordinate identification.

Let each i.i.d. context Z specify all bids, priorities, and slot qualities of a rational single-slot or position auction. Full observation means that these data are seen for every bidder, allowing evaluation of $f_i(\theta; Z)$, the payment averaged over participation coins at fixed Z. Realized payments alone do not supply this feedback. Define

$$P_i(\theta) = \mathbb{E} f_i(\theta; Z), \qquad \widehat{P}_i(\theta) = T^{-1}\sum_{t=1}^{T} f_i(\theta; Z_t).$$

Budgets are known and positive, and a known $R \geq 1$ satisfies

$$0 \leq f_i(\theta; Z) \leq RB_i \quad \text{for all } i, \theta, Z. \tag{18}$$

If a bound $R < 1$ is available, full participation is already exact.

Theorem 5.3 (Learning Equilibrium Conditions)**.** *For $0 < \delta \leq 1/2$ and $0 < \alpha < 1$, a $\delta/8$-approximate empirical equilibrium with budgets $(1-\delta/4)B_i$ is a population δ-approximate equilibrium with probability at least $1-\alpha$, using*

$$T = O\left(\frac{R}{\delta^2}\left[n + \log\frac{n}{\alpha}\right]\right) \tag{19}$$

contexts. It is computable in time polynomial in the sampled input length and $\log(1/\delta)$.

Proof idea. Multilinearity expresses each payment as a convex mixture of Boolean-vertex payments. Relative Bernstein concentration and a union bound over $n2^n$ bidder–vertex pairs yield, simultaneously for every i, θ, $|\widehat{P}_i(\theta) - P_i(\theta)| \leq (\delta/16)(B_i + P_i(\theta))$. At the empirical solution, the budget slack absorbs this estimation error. Also, $\theta_i < 1-\delta$ implies $\theta_i < 1-\delta/8$, so empirical complementarity applies whenever population complementarity is required:

$$\frac{P_i}{B_i} \leq \frac{1-3\delta/16}{1-\delta/16} < 1, \qquad \theta_i < 1-\delta \implies \frac{P_i}{B_i} \geq \frac{(1-\delta/8)(1-\delta/4)-\delta/16}{1+\delta/16} \geq 1-\delta/2.$$

Thus population budgets remain feasible and complementarity is satisfied. The vertices are used only for the concentration bound; the algorithm does not enumerate them. Weight each sampled context by $1/T$ and apply Theorem 1.1 or its position-auction extension to the resulting rational market. Appendix C proves the concentration step. □

A one-bidder lower bound of $\Omega((R/\delta^2)\log(1/\alpha))$ matches R, δ, α for fixed n when $R \geq 2$, $\delta \leq 1/8$, and $\alpha \leq 1/4$ (Theorem C.2); matching dimension dependence remains open. The earlier $\mathrm{poly}(1/\delta)$ solver also gives polynomial-time residual learning because the sample size already grows as δ^{-2}. Theorem 1.1 improves the empirical solver's precision dependence.

For coordinates, a known bound $\Gamma^\star := \min_k \partial_k V(\theta^\star) \geq \underline{\Gamma} > 0$ allows residual tolerance $\delta_\eta = \min\{1/2, \underline{\Gamma}\eta/\sum_i B_i\}$ to ensure error η (Corollary C.5). With other parameters fixed, this gives $O(\eta^{-2}\log(2/\alpha))$ samples. The positive-marginal bound is an additional assumption; uniqueness alone does not supply it. It can hold at full participation, as the private-auction example in Appendix C.3 shows. The cyclic family below instead has vanishing revenue marginals and a slower coordinate-identification rate.

5.3 Sharp Coordinate Identification at a Boundary

To obtain a statistical version of Proposition 5.2, draw a priority order J uniformly from the same n cyclic shifts, and independently draw a common bid $X \in \{0, 2\}$ with mean $\mu \in [0, 4/3]$. Each budget is $1/n$; thus $R = 2n$, independent of the target coordinate accuracy. Cyclic symmetry and uniqueness give $\theta^\star = t^\star(\mu)\mathbf{1}$, where

$$t^\star(\mu) = \begin{cases} 1, & 0 \leq \mu \leq 1, \\ 1 - (1 - 1/\mu)^{1/n}, & 1 < \mu \leq 4/3. \end{cases} \tag{20}$$

This follows from $P_i^\mu(t\mathbf{1}) = \mu[1 - (1-t)^n]/n$. At $\mu = 1$, the equilibrium map has a $1/n$-order boundary modulus.

Theorem 5.4 (Matching Sample Bounds for the Cyclic Family)**.** *For $n \geq 2$, $0 < \eta \leq 1/8$, and $0 < \alpha \leq 1/4$, an estimator attaining coordinate error at most η with probability at least $1 - \alpha$ throughout this family requires $T \geq \frac{1}{4}(2\eta)^{-2n}\log(1/\alpha)$ samples. Conversely, $T \geq 8(2\eta)^{-2n}\log(2/\alpha)$ samples suffice. Both constants are independent of n.*

Proof idea. Compare $\mu_0 = 1$ with $\mu_1 = 1/(1 - d^n)$ for $2\eta < d < 3\eta$. Their equilibria are distance d apart, but their one-sample KL divergence is $O(d^{2n})$. Binary testing and $d \downarrow 2\eta$ give the lower bound. For the upper bound, estimate μ with a confidence interval of width $(2\eta)^n$, map its endpoints through $t^\star$, and return their midpoint. The $1/n$-Hölder bound on (20) gives an interval of width at most 2η. Appendix C.4 specifies the estimator and proves the constants. □

Here $\Gamma^\star = \mu(1-t^\star)^{n-1}$ vanishes at $\mu = 1$. No fixed positive lower bound on marginal revenue holds throughout the family. Estimating the mean bid to its usual $T^{-1/2}$ accuracy is therefore insufficient to identify participation at that rate: the equilibrium map transforms mean perturbations of order η^n into coordinate changes of order η. The matching bounds quantify the sample cost of this flat response. The upper-bound estimator guarantees coordinates, not budget feasibility; the lower bound applies even without imposing feasibility. This is a statistical identification result for the specified family, not a general coordinate-computation lower bound.

6 Conclusion

Positive-tree continuation enables efficient high-precision equilibrium computation even with ill-conditioned payment Jacobians. However, extensions still depend on efficient payment evaluation, and residual accuracy need not imply accurate participation estimates. The revenue certificates and boundary learning bounds quantify this gap. Our guarantees concern expected budgets and, for learning, fully observed contexts. Future work should address practical performance and bandit feedback, as well as optimal dimension dependence in residual learning and coordinate rates beyond the cyclic family.

Reproducibility Statement

All results are theoretical. Algorithm 1 summarizes the continuation loop, and Lemma A.7 implements its certified step. Appendix A gives the rational coefficient recurrence, parameter choices, rounding rules, stopping tests, and bit bounds. Appendix B proves the payment-rule extensions and reliability reductions. Appendix C contains the statistical proofs and specifies the estimator for the cyclic family. The worked examples in Section 5 and Appendix A.3 are analytic calculations.

AI Use Statement

Generative AI tools were used to search and compare related literature; propose and refine auxiliary mathematical claims and proof steps; develop, check, and simplify proof arguments; revise the organization and English prose; and prepare the LaTeX source and finite-instance consistency checks. The authors are responsible for the final content, including all AI-assisted claims, arguments, citations, and artifacts.

References

Santiago R. Balseiro, Kshipra Bhawalkar, Zhe Feng, Haihao Lu, Vahab Mirrokni, Balasubramanian Sivan, and Di Wang. A field guide for pacing budget and ROS constraints. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 2607–2638. PMLR, 2024. URL https://proceedings.mlr.press/v235/balseiro24a.html.

Abraham Berman and Robert J. Plemmons. *Nonnegative Matrices in the Mathematical Sciences*. SIAM, 1994. doi: 10.1137/1.9781611971262. URL https://doi.org/10.1137/1.9781611971262.

Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013. doi: 10.1093/acprof:oso/9780199535255.001.0001. URL https://academic.oup.com/book/26549.

Olivier Bournez, Daniel S. Graça, and Amaury Pouly. On the complexity of solving initial value problems. arXiv:1202.4407, 2012. URL https://arxiv.org/abs/1202.4407. Conference version at ISSAC 2012.

Seth Chaiken. A combinatorial proof of the all minors matrix tree theorem. *SIAM Journal on Algebraic and Discrete Methods*, 3(3):319–329, 1982. doi: 10.1137/0603033. URL https://epubs.siam.org/doi/10.1137/0603033.

Li Chen, Jamie Morgenstern, and Yuanyuan Yang. Beyond the PPAD hardness of auto-bidding auctions. arXiv:2608.01889, 2026. URL https://arxiv.org/abs/2608.01889.

Xi Chen and Yuhao Li. Constant inapproximability of pacing equilibria in second-price auctions. In Vittorio Bilò, Yang Cai, and Zhiyi Huang (eds.), *Web and Internet Economics*, volume 16266 of *Lecture Notes in Computer Science*, pp. 294–306. Springer, 2026. doi: 10.1007/978-3-032-18660-7_15. URL https://arxiv.org/abs/2501.15295. WINE 2025; extended version: arXiv:2501.15295v2.

Xi Chen, Christian Kroer, and Rachitesh Kumar. Throttling equilibria in auction markets. arXiv:2107.10923, 2021. URL https://arxiv.org/abs/2107.10923. Revised 2023; conference version at WINE 2021.

Xi Chen, Christian Kroer, and Rachitesh Kumar. The complexity of pacing for second-price auctions. *Mathematics of Operations Research*, 49(4):2109–2135, 2024a. doi: 10.1287/moor.2022.0009. URL https://doi.org/10.1287/moor.2022.0009.

Zhaohua Chen, Chang Wang, Qian Wang, Yuqi Pan, Zhuming Shi, Zheng Cai, Yukun Ren, Zhihua Zhu, and Xiaotie Deng. Dynamic budget throttling in repeated second-price auctions. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(9):9598–9606, 2024b. doi: 10.1609/aaai.v38i9.28816. URL https://ojs.aaai.org/index.php/AAAI/article/view/28816.

Vincent Conitzer, Christian Kroer, Debmalya Panigrahi, Okke Schrijvers, Nicolas E. Stier-Moses, Eric Sodomka, and Christopher A. Wilkens. Pacing equilibrium in first price auction markets. *Management Science*, 68(12):8515–8535, 2022. doi: 10.1287/mnsc.2022.4310. URL https://doi.org/10.1287/mnsc.2022.4310.

Shaddin Dughmi, Tim Roughgarden, and Mukund Sundararajan. Revenue submodularity. *Theory of Computing*, 8(5):95–119, 2012. doi: 10.4086/toc.2012.v008a005. URL https://theoryofcomputing.org/articles/v008a005/.

Tara Fife and James Oxley. Laminar matroids. *European Journal of Combinatorics*, 62:206–216, 2017. doi: 10.1016/j.ejc.2017.01.002. URL https://arxiv.org/abs/1606.08354.

Yuan Gao and Christian Kroer. First-order methods for large-scale market equilibrium computation. In *Advances in Neural Information Processing Systems*, volume 33, pp. 21738–21750, 2020. URL https://proceedings.neurips.cc/paper/2020/hash/f75526659f31040afeb61cb7133e4e6d-Abstract.html.

Yuan Gao, Alex Peysakhovich, and Christian Kroer. Online market equilibrium with application to fair division. In *Advances in Neural Information Processing Systems*, volume 34, pp. 27305–27318, 2021. URL https://proceedings.neurips.cc/paper/2021/hash/e562cd9c0768d5464b64cf61da7fc6bb-Abstract.html.

Yiyang Huang, Yonglei Yan, Zihe Wang, and Zhengyang Liu. Pacing equilibria in second-price auctions with few goods. arXiv:2605.09332, 2026. URL https://arxiv.org/abs/2605.09332.

Luofeng Liao and Christian Kroer. Statistical inference and A/B testing for first-price pacing equilibria. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 20868–20905. PMLR, 2023. URL https://proceedings.mlr.press/v202/liao23a.html.

Luofeng Liao and Christian Kroer. Bootstrapping Fisher market equilibrium and first-price pacing equilibrium. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 29748–29781. PMLR, 2024. URL https://proceedings.mlr.press/v235/liao24b.html.

Luofeng Liao, Yuan Gao, and Christian Kroer. Statistical inference for Fisher market equilibrium. In *International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=KemSBwOYJC.

Luofeng Liao, Christian Kroer, Sergei Leonenkov, Okke Schrijvers, Liang Shi, Nicolas Stier-Moses, and Congshan Zhang. Interference among first-price pacing equilibria: A bias and variance analysis. In *International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=6bDJ3CIm5w.

Zhengyang Liu. Tight inapproximability of pacing and throttling equilibria in second-price auctions. arXiv:2608.16682, 2026. URL https://arxiv.org/abs/2608.16682.

Brendan Lucier, Sarath Pattathil, Aleksandrs Slivkins, and Mengxiao Zhang. Autobidders with budget and ROI constraints: Efficiency, regret, and pacing dynamics. In *Proceedings of the 37th Annual Conference on Learning Theory*, volume 247 of *Proceedings of Machine Learning Research*, pp. 3642–3643. PMLR, 2024. URL https://proceedings.mlr.press/v247/lucier24a.html.

James Oxley. *Matroid Theory*. Oxford University Press, 2 edition, 2011. doi: 10.1093/acprof:oso/9780198566946.001.0001. URL https://academic.oup.com/book/34846.

J. Scott Provan. The complexity of reliability computations in planar and acyclic graphs. *SIAM Journal on Computing*, 15(3):694–702, 1986. doi: 10.1137/0215050. URL https://epubs.siam.org/doi/10.1137/0215050.

Jacob T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *Journal of the ACM*, 27(4):701–717, 1980. doi: 10.1145/322217.322225. URL https://pages.cs.wisc.edu/~cs787-1/Schwartz1980.pdf.

Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer, 2009. doi: 10.1007/b13794. URL https://doi.org/10.1007/b13794.

Qian Wang, Zongjun Yang, Xiaotie Deng, and Yuqing Kong. Learning to bid in repeated first-price auctions with budgets. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 36494–36513. PMLR, 2023. URL https://proceedings.mlr.press/v202/wang23ao.html.

Yonglei Yan, Zihe Wang, and Zhengyang Liu. Pacing equilibria in second-price auctions with few buyers. *Proceedings of the AAAI Conference on Artificial Intelligence*, 40(20):17302–17309, 2026. doi: 10.1609/aaai.v40i20.38782. URL https://doi.org/10.1609/aaai.v40i20.38782.

A Proof of the High-Precision Algorithm

Roadmap. This appendix has four proof modules. The present section develops the normalized continuation path, its directed-tree representation, the rational local solver, and the global error analysis. Appendix B extends the structural argument to general payment rules and proves the matroid-reliability reductions. Appendix C contains the statistical proofs, and Appendix D gives the lower bound for the previous multiplicative update.

Bidders with Zero Bids. Let $M_i = \sum_j b_{ij}$. If $M_i = 0$, then $P_i \equiv 0$; set $\theta_i = 1$, delete the bidder during the computation, and restore her at the end. Hence every bidder considered in the continuation proof has $M_i > 0$. If no bidder remains, return the all-one vector immediately.

A.1 The Normalized Payment Flow

Fix a set A of *active* bidders. Bidders outside A are frozen, so only the coordinates in A move. For every retained bidder, use the log-odds coordinate

$$u_i = \log \frac{\theta_i}{1-\theta_i}, \qquad \theta_i = \frac{\mathrm{e}^{u_i}}{1+\mathrm{e}^{u_i}}. \tag{21}$$

Let

$$H_A(\theta) = \left[\frac{\partial P_i}{\partial u_k}\right]_{i,k\in A} \tag{22}$$

be the active payment Jacobian. Frozen logits stay constant and remain in the potential used to count steps.

A.1.1 The Jacobian Is an M-Matrix

Differentiating (1) gives, for $i, k \in A$,

$$H_{ii} = (1-\theta_i)P_i, \tag{23}$$

$$H_{ik} = -\theta_k \sum_{j:\, k \succ_j i} p_{ij} \quad (i \neq k). \tag{24}$$

Thus $H_{ii} > 0$ and $H_{ik} \leq 0$ for $i \neq k$.

Lemma A.1. *Assume $0 < \theta_i < 1$ for every nonzero bidder. For every active bidder k,*

$$H_{kk} > \sum_{i\in A\setminus\{k\}} |H_{ik}|. \tag{25}$$

Consequently, H_A is a nonsingular M-matrix and $H_A^{-1} \geq 0$ entrywise.

Proof. Let $v_A^j(S)$ be the winning bid on item j if the winner belongs to A, and zero otherwise. Adding $k \in A$ cannot lower v_A^j: if k wins, she displaces a weakly lower bid. Thus, for $S \sim \theta_{-k}$,

$$\begin{aligned}\sum_{i\in A} H_{ik} &= \theta_k(1-\theta_k)\sum_j \mathbb{E}_S[v_A^j(S\cup\{k\}) - v_A^j(S)] \\ &\geq \theta_k(1-\theta_k)M_k \prod_{\ell\neq k}(1-\theta_\ell) > 0.\end{aligned} \tag{26}$$

The lower bound retains only $S = \varnothing$. Since off-diagonal entries are nonpositive, this is strict column dominance. Writing $H_A^\top = D(I-Q)$ with $D > 0$, $Q \geq 0$, and row sums of Q below one gives $(H_A^\top)^{-1} = \sum_{r\geq 0} Q^r D^{-1} \geq 0$. □

A.1.2 NORMALIZE ALL LOGIT SPEEDS

Introduce a scale $0 < t < 1$ and its logit $u_0 = \log(t/(1-t))$. At a state with active set A, define

$$H_A y = (1-t)aP_A, \qquad a + \mathbf{1}^\top y = 1, \qquad \dot{u}_0 = a, \quad \dot{u}_i = y_i \ (i \in A). \tag{27}$$

Equivalently,

$$a = \frac{1}{1 + (1-t)\mathbf{1}^\top H_A^{-1} P_A}, \qquad y = (1-t)aH_A^{-1}P_A. \tag{28}$$

Every $P_i > 0$, and H_A^{-1} is nonnegative with positive diagonal. Therefore

$$a > 0, \qquad y_i > 0, \qquad a + \sum_{i \in A} y_i = 1. \tag{29}$$

Frozen coordinates are held fixed. All moving logit speeds lie in $(0, 1]$.

Lemma A.2. *Along* (27)*, t and every active θ_i increase, and*

$$\dot{P}_A = (1-t)aP_A, \qquad \frac{\mathrm{d}}{\mathrm{d}s} \log \frac{P_i}{t} = 0 \quad (i \in A). \tag{30}$$

Frozen bidders' payments are nonincreasing.

Proof. The chain rule gives $\dot{P}_A = H_A y = (1-t)aP_A$ and $\dot{t} = t(1-t)a$. Hence $\dot{P}_i/P_i = \dot{t}/t$. Positivity in (29) gives monotonicity. Equation (1) makes each P_i nonincreasing in competitors' probabilities, so frozen bidders' payments cannot rise. □

A.2 A POSITIVE DIRECTED-TREE FORMULA

The main numerical issue is that H_A may be badly conditioned. We now show that the *normalized* direction (a, y) nevertheless has a uniform analytic neighborhood depending only polynomially on n.

For $k, i \in A$, $k \neq i$, define

$$w_{k \to i} := -H_{ik} \geq 0, \qquad \kappa_k := H_{kk} - \sum_{i \in A \setminus \{k\}} w_{k \to i} > 0.$$

Equations (24) and (26) show more than positivity: each $w_{k\to i}$ and κ_k is a polynomial with nonnegative coefficients in factors of the form θ_ℓ and $1 - \theta_\ell$, after frozen coordinates are treated as constants. Each such polynomial has total degree at most $n + 1$. The weight $(1-t)P_i$ also has degree at most $n + 1$ when t and $1 - t$ are included among the logistic factors.

Construct a directed weighted graph on the vertex set $\{0\} \cup A$ with arcs

$$\begin{aligned} &k \to i \text{ of weight } w_{k\to i} && (k, i \in A,\ k \neq i),\\ &k \to 0 \text{ of weight } \kappa_k && (k \in A),\\ &0 \to i \text{ of weight } (1-t)P_i && (i \in A). \end{aligned}$$

Its column Laplacian is

$$\mathcal{L}_A = \begin{pmatrix} (1-t)\sum_{i \in A} P_i & -\kappa^\top \\ -(1-t)P_A & H_A \end{pmatrix}. \tag{31}$$

The columns sum to zero. Furthermore, (27) is exactly

$$\mathcal{L}_A \begin{pmatrix} a \\ y \end{pmatrix} = 0, \qquad a + \mathbf{1}^\top y = 1. \tag{32}$$

Lemma A.3 (Directed-Tree Formula)**.** *For each vertex $v \in \{0\} \cup A$, let T_v be the sum of the weights of all directed spanning trees in which every vertex different from v has one outgoing arc and every directed path ends at v. Then*

$$a = \frac{T_0}{Z}, \qquad y_i = \frac{T_i}{Z}, \qquad Z := T_0 + \sum_{i \in A} T_i. \tag{33}$$

Each T_v and Z is a nonzero polynomial with nonnegative coefficients in the logistic factors $t, 1 - t, \theta_k, 1 - \theta_k$, of degree at most

$$d := n(n+1). \tag{34}$$

The degree in each individual pair $x_k, 1 - x_k$ is at most $n + 1$, where $x_0 = t$ and $x_i = \theta_i$.

Proof. The directed matrix-tree theorem (Chaiken, 1982) identifies the principal cofactor at v with T_v. The graph is strongly connected, since the arcs $k \to 0$ and $0 \to k$ have positive weights for every k. Thus all $T_v > 0$, and the cofactor vector spans the right nullspace of the column Laplacian. Normalizing it gives (33). Each tree has $|A| \le n$ arcs, each of degree at most $n+1$ in the logistic factors. Products and sums preserve nonnegative coefficients, giving the degree bound. An arc with tail k uses at most two factors from the pair $x_k, 1-x_k$ and at most one from any other pair. Since a tree has at most one arc from each tail, its degree in any individual pair is at most $|A|+1 \le n+1$. □

A.2.1 A Zero-Free Complex Neighborhood

Corollary A.4 (Uniformity Over Active Sets)**.** *For every active set A, every assignment of the frozen participation probabilities in $(0,1)$, and every real state $u = (u_0, (u_i)_{i\in A})$, the vector field*

$$F(u) := (a(u), (y_i(u))_{i\in A})$$

extends holomorphically to $\|z\|_1 < \rho := 2/(n+1)$ around u. For $\|z\|_1 \le \rho$,

$$|a(u+z)| + \sum_{i\in A} |y_i(u+z)| \le 2. \tag{35}$$

The same radius and bound hold at every restart of the algorithm; they are independent of A, the frozen state, $\det H_A$, all bid magnitudes, all budgets, and ε.

Proof. This applies Lemma 3.1 with $D = n+1$; we spell out why the bound is uniform over frozen states. Frozen logistic factors are positive real constants, so they only change the nonnegative coefficients below and never the degree bound. By Lemma A.3, each monomial has degree at most D in every moving pair. Multiply each monomial by enough copies of $x_k + (1-x_k) = 1$ to obtain

$$T_v(x) = \sum_{\alpha} c_{v,\alpha} \prod_k x_k^{\alpha_k} (1-x_k)^{D-\alpha_k}, \qquad c_{v,\alpha} \ge 0, \quad 0 \le \alpha_k \le D.$$

Substitution of $x_k = e^{u_k}/(1+e^{u_k})$ gives a common denominator. Cancel it and the common factor $e^{(D/2)\sum_k u_k}$ to get

$$F_v(u) = \frac{S_v(u)}{\sum_w S_w(u)}, \qquad S_v(u) = \sum_{\alpha} c_{v,\alpha} e^{(\alpha - D\mathbf{1}/2)\cdot u}.$$

For real u and $\|z\|_1 \le 2/D$, every term at $u+z$ has argument of magnitude at most $(D/2)\|\operatorname{Im} z\|_1 \le 1$. Since $\cos(1) > 1/2$, for each v,

$$|S_v(u+z)| \le \sum_{\alpha} c_{v,\alpha} \left| e^{(\alpha - D\mathbf{1}/2)\cdot(u+z)} \right| \le \sec(1) \operatorname{Re} S_v(u+z) < 2 \operatorname{Re} S_v(u+z).$$

At least one tree term is present, so $\operatorname{Re} \sum_v S_v(u+z) > 0$. Therefore the denominator is nonzero and

$$\sum_v |S_v(u+z)| < 2 \operatorname{Re} \sum_v S_v(u+z) \le 2 \left| \sum_v S_v(u+z) \right|,$$

which proves (35). An active-set change is handled only after the current step ends; restarting with the new set therefore uses the same radius without requiring analyticity across the change. □

A.2.2 Uniform Taylor Continuation

Lemma A.5 (Uniform Local Analytic Step)**.**

$$h := \frac{1}{8(n+1)} \tag{36}$$

is a valid continuation step for every active set and every real state with $0 < t < 1$ and $0 < \theta_i < 1$. More precisely, the solution of (27) is holomorphic in complex time on $|s| \le 2h$, and the odds $\omega = x/(1-x)$ have order-q Taylor approximation $\omega_{[q]}$ satisfying

$$|\omega(\sigma) - \omega_{[q]}(\sigma)| \le 2\omega(0)\, 2^{-q} \quad (0 \le \sigma \le h) \tag{37}$$

for every moving coordinate $x \in \{t\} \cup \{\theta_i : i \in A\}$.

Proof. Put $\rho = 2/(n+1)$. On the complex ℓ_1 ball of radius ρ, the field $F = (a, y)$ satisfies $\|F\|_1 \le 2$. For $\|v\|_1 \le 1$ and u in the half-radius ball, the vector-valued Cauchy formula applied to $\zeta \mapsto F(u + \zeta v)$ gives

$$\|DF(u)v\|_1 \le \frac{\sup_{|\zeta|<\rho/2} \|F(u+\zeta v)\|_1}{\rho/2} \le \frac{4}{\rho}.$$

On the complex time disc $|s| \le 2h = \rho/8$, the Picard map sends paths in the half-radius ball into the radius-$\rho/4$ ball and has contraction factor at most $1/2$. Thus the solution exists holomorphically on this disc. Since $\omega(s) = \omega(0)e^{u(s)-u(0)}$, each moving coordinate satisfies

$$|\omega(s)| \le e^{\rho/4}\omega(0) \le 2\omega(0).$$

Cauchy's coefficient estimate followed by a geometric sum proves (37). □

A.3 Example: A Diverging Inverse, a Bounded Direction

Take two bidders and two unit-bid goods with opposite priorities and unit budgets. At $\theta = (1-r)\mathbf{1}$, $0 < r < 1$, each payment is $1 - r^2$, and

$$H_A = (1-r)r \begin{pmatrix} 1+r & -(1-r) \\ -(1-r) & 1+r \end{pmatrix}.$$

The eigenvalues are $2(1-r)r^2$ and $2(1-r)r$, so its spectral condition number is $1/r$. At any $t \in (0,1)$, the normalized direction is

$$a = \frac{r^2}{r^2 + (1-t)(1+r)}, \qquad y_1 = y_2 = \frac{1-a}{2}. \tag{38}$$

Along the ideal budget-proportional curve $P_i = tB_i$, we have $t = 1 - r^2$, so $a = 1/(2+r)$ and $y_1 = y_2 = (1+r)/[2(2+r)]$. As $r \downarrow 0$, this direction tends to $(1/2, 1/4, 1/4)$ while the inverse Jacobian diverges. Initialization (10) instead gives $\theta_i^0 = t_0/2$ and $P_i(\theta^0)/t_0 = C := 1 - t_0/4$; the unrounded path before freezing obeys $t = (1-r^2)/C$. Lemma 3.1 applies to every real state, covering both curves and every rational restart. It supplies the complex bound needed to high-order steps.

A.4 Rational Evaluation of One Step

Although the analysis uses logits, the algorithm stores only rational (t, θ). In these coordinates the ODE is

$$\dot{t} = t(1-t)a, \qquad \dot{\theta}_i = \theta_i(1-\theta_i)y_i \quad (i \in A). \tag{39}$$

At a fixed rational state, (y, a) is the solution of

$$\underbrace{\begin{pmatrix} H_A(\theta) & -(1-t)P_A(\theta) \\ \mathbf{1}^\top & 1 \end{pmatrix}}_{\mathcal{A}(t,\theta)} \begin{pmatrix} y \\ a \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}. \tag{40}$$

The matrix-tree theorem gives $T_0 = \det H_A$. The Schur complement and (28) therefore give

$$\det \mathcal{A} = \det H_A\big(1 + (1-t)\mathbf{1}^\top H_A^{-1} P_A\big) = \frac{T_0}{a} = Z > 0. \tag{41}$$

Analyticity controls the truncation error. The next lemma separately bounds the bits needed to represent the retained Taylor coefficients.

Lemma A.6. *Consider a rational ODE and a rational function with the same denominator,*

$$x' = A(x)/B(x), \qquad g(x) = U(x)/B(x), \qquad B(x_0) \neq 0,$$

at a rational initial point x_0. Suppose the number of variables, degree bounds, the bit lengths of a common coefficient denominator and of every coefficient numerator, and the encoding length of x_0 are $\mathrm{poly}(L)$. The polynomials need not be sparse. Then the Taylor coefficients of $g(x(s))$ through order q have $\mathrm{poly}(L, q)$ bits. This includes state coordinates, by taking $U = x_j B$.

Proof. Clear one common coefficient denominator so A, B, U have integer coefficients, and let $D \geq 1$ bound their degrees. The differential operator $\mathcal{D} = (A/B) \cdot \nabla$ satisfies

$$\mathcal{D}^r g = \frac{U_r}{B^{2r+1}}, \qquad U_0 = U, \quad U_{r+1} = BA \cdot \nabla U_r - (2r+1) U_r A \cdot \nabla B.$$

Thus U_r has integer coefficients and degree at most $(2r+1)D$. For a polynomial V, write $\|V\|_1$ for the sum of its absolute coefficients. The recurrence also gives

$$\|U_{r+1}\|_1 \leq C(2r+1)\|U_r\|_1, \qquad C = \max\left\{1, 2D\|B\|_1 \sum_j \|A_j\|_1\right\}.$$

Hence $\|U_r\|_1 \leq \|U\|_1 [C(2r+1)]^r$. A degree-$D$ polynomial in m variables has at most $(D+1)^m$ potential monomials. Since m, D, and every coefficient length are polynomially bounded, $\log C = \mathrm{poly}(L)$ even when the expanded support is exponential. Let c be the product of the coordinate denominators of x_0 and put $b = c^D B(x_0) \in \mathbb{Z} \setminus \{0\}$. Then

$$[s^r] g(x(s)) = \frac{c^{(2r+1)D} U_r(x_0)}{b^{2r+1} r!}.$$

The numerator is an integer of magnitude at most $\|U_r\|_1 \max\{c, \|c x_0\|_\infty\}^{(2r+1)D}$. Both numerator and denominator therefore have $\mathrm{poly}(L, r)$ bits. □

Lemma A.7. *Suppose the current state is rational, lies in the open participation cube with $0 < t < 1$, and has encoding length at most G. Given rational $\sigma \in [0, h]$ and $\xi \in (0, 1]$, one can compute a rational endpoint satisfying*

$$\max_{k \in \{0\} \cup A} \left| \log \frac{\widehat{x}_k/(1-\widehat{x}_k)}{x_k(\sigma)/(1-x_k(\sigma))} \right| \leq \xi,$$

using Taylor order $q = \lceil \log_2(8/\xi) \rceil$ and $\mathrm{poly}(N, G, \mathrm{bits}(\xi), \mathrm{bits}(\sigma))$ bit operations, where $x_0 = t$ and $x_i = \theta_i$ for $i \in A$.

Proof. Write $z(s) = \sum_{r \geq 0} z^{[r]} s^r$ for $z = t, \theta, y, a$, and write $\mathcal{A}(s) = \mathcal{A}(t(s), \theta(s)) = \sum_{r \geq 0} \mathcal{A}^{[r]} s^r$. Let $v^{[r]} = (y^{[r]}, a^{[r]})$ and $b = (0, \ldots, 0, 1)^\top$. The product formulas for P and H give their truncated series by ordinary convolution, without expanding their monomial supports. At order r, the state coefficients through order r are already known. Extracting the coefficient of (40) gives a linear system

$$\mathcal{A}^{[0]} v^{[r]} = \mathbf{1}_{\{r=0\}} b - \sum_{j=1}^{r} \mathcal{A}^{[j]} v^{[r-j]}.$$

The right-hand side uses state coefficients through order r and direction coefficients of orders below r. Thus the same rational matrix is solved once per order. Equation (39) then determines $t^{[r+1]}$ and $\theta^{[r+1]}$ by one convolution and division by $r+1$. Compute the odds series $\omega_k = x_k/(1-x_k)$ by formal division, using $1 - x_k(0) > 0$. A conservative rational-operation count is $O(mn^3(q+1)^3)$. Indeed, at a given order r, recomputing each truncated product in the explicit formulas for P and H by ordinary convolution takes at most $O(mn^3(r+1)^2)$ operations. Sum this bound over $r \leq q$. Forming the recurrence right-hand sides, solving the $(|A|+1)$-dimensional systems by elimination, and computing the state and odds series take at most $O(n^3(q+1) + n^2(q+1)^2)$ further operations. These are arithmetic counts; the following argument bounds their bit cost. No tree enumeration or expanded multivariate polynomial representation is used.

By Cramer's rule and (41), the state derivatives and (y, a) are rational functions with common denominator Z and polynomial degrees and coefficient heights. To include the odds as rational functions, use the common denominator $B = Z \prod_k (1 - x_k)$ over the state coordinates. It is nonzero at the restart point and still has polynomial degree and height. Lemma A.6 gives $\mathrm{poly}(N, G, q)$-bit coefficients for the states, directions, and odds in one application.

The coefficient recurrence computes these same coefficients. Each right-hand side and the fixed matrix $\mathcal{A}(t(0), \theta(0))$ have polynomial bit length. Reducing rational arithmetic after each operation and using fraction-free elimination therefore gives polynomial bit cost per order. The temporary products

and sums have polynomial length because they combine only polynomially many polynomial-length coefficients. No condition-number bound is used.

For each coordinate, Lemma A.5 gives odds truncation error at most $\xi\omega(0)/4$. Round to the nearest multiple of a dyadic mesh $\Delta \le \xi\omega(0)/2$, with ties rounded down; the rounding error is at most $\xi\omega(0)/4$. For this local lemma, take the coarsest such mesh with $\Delta \le 1$. In the full algorithm, use the fixed mesh specified below instead. Since $\omega(\sigma) \ge \omega(0)$,

$$\left|\frac{\widehat{\omega}}{\omega(\sigma)} - 1\right| \le \frac{\xi}{2}, \qquad \left|\log \frac{\widehat{\omega}}{\omega(\sigma)}\right| \le \xi.$$

The second bound uses $|\log(1+z)| \le 2|z|$ for real $|z| \le 1/2$. Return $\widehat{x} = \widehat{\omega}/(1+\widehat{\omega}) \in (0,1)$ exactly. The dyadic odds and the returned probability have $O(G + \log(1/\xi))$ bits. □

A.5 The Rational Algorithm

We now give the complete finite-precision algorithm. Constants such as 1000 are only used to keep the inequalities visibly separated. For the first formulation of Theorem 1.1, take $\delta = 2^{-p}$; for a general requested δ, define the internal accuracy below.

Set

$$p := \max\left\{1, \left\lceil \log_2 \frac{1}{\delta} \right\rceil\right\}, \qquad \varepsilon := 2^{-p}. \tag{42}$$

Then $0 < \varepsilon \le \min\{\delta, 1/2\}$, and an ε-approximate equilibrium is also a δ-approximate equilibrium.

For every remaining bidder let $M_i = \sum_j b_{ij}$. Define

$$s_0 := \min\left\{\frac{1}{4}, \frac{\varepsilon}{1000n} \min_i \frac{M_i}{B_i}\right\}, \tag{43}$$

and set

$$t_0 := 2^{-\lceil \log_2(1/s_0)\rceil}, \qquad \theta_i^0 = t_0 \frac{B_i}{M_i}. \tag{44}$$

Thus t_0 is the largest number of the form 2^{-k} not exceeding s_0, $s_0/2 < t_0 \le s_0$, and $\theta_i^0 \le \varepsilon/(1000n)$.

Use the rational step h from Lemma A.5. Define a computable polynomial upper bound K_* on the number of steps. Let

$$\lambda := \min\left\{t_0, \min_i \theta_i^0, \frac{\varepsilon}{4}\right\}, \qquad Q := \left\lceil \log_2 \frac{1}{\lambda} \right\rceil + 2. \tag{45}$$

The potential argument in Section A.6.3 shows that one may take

$$K_* := \left\lceil \frac{2(n+1)(Q+p+2)}{h} \right\rceil. \tag{46}$$

Finally set the local logit tolerance

$$\xi := \min\left\{\frac{h}{4(n+1)}, \frac{\varepsilon}{1000(n+1)(K_*+1)}\right\}. \tag{47}$$

All these quantities have $\mathrm{poly}(N,p)$-bit encodings. Fix once and for all

$$b_{\mathrm{grid}} := Q + \lceil \log_2(2/\xi) \rceil, \qquad \Delta := 2^{-b_{\mathrm{grid}}} \le \xi 2^{-Q-1}. \tag{48}$$

Every odds endpoint is rounded to the nearest multiple of Δ, with ties rounded down. The bound $\omega(0) \ge 2^{-Q}$ makes this mesh valid at each restart. Its validity is inductive: initialization satisfies the bound; a valid local step and the clamp preserve it by (58). Thus there is no assumption about the encoding length of future iterates. All dyadic exponents and threshold tests are computed by integer comparisons; no logarithm or transcendental function is evaluated numerically.

Iteration. Start from (t_0, θ^0) with all nonzero bidders active. Repeat the following steps:

1. If $A = \varnothing$ or $t \geq 1 - \varepsilon/2$, output the current θ.
2. Use Lemma A.7 to advance the active flow by time h, using the fixed grid (48) for odds rounding and converting to probabilities exactly. Replace any returned coordinate below its old value by that old value.
3. Freeze every bidder satisfying
$$\theta_i \geq 1 - \frac{\varepsilon}{2}. \tag{49}$$

Only stored rational numbers are compared with the threshold. The clamp in step 2 preserves monotonicity without increasing local error, because the exact endpoint is at least the old value. As proved below, every stored coordinate remains at most $1 - \varepsilon/4$. Consequently, if the scale triggers termination, then
$$1 - \frac{\varepsilon}{2} \leq t \leq 1 - \frac{\varepsilon}{4}. \tag{50}$$
Every step has the same length; no crossing time needs to be located.

A.6 Correctness and Complexity

A.6.1 Initialization

Lemma A.8. *For every nonzero bidder,*
$$\left(1 - \frac{\varepsilon}{1000}\right) B_i \leq \frac{P_i(\theta^0)}{t_0} \leq B_i. \tag{51}$$

Proof. Every initial participation probability is at most $\varepsilon/(1000n)$. Hence, for every i, j,
$$\prod_{k \succ_j i} (1 - \theta_k^0) \geq 1 - \sum_{k \succ_j i} \theta_k^0 \geq 1 - \frac{\varepsilon}{1000}.$$
Using (44),
$$\frac{P_i(\theta^0)}{t_0} = \frac{B_i}{M_i} \sum_j b_{ij} \prod_{k \succ_j i} (1 - \theta_k^0),$$
which gives both bounds. □

A.6.2 Rounding Errors Do Not Need Global Stability

Lemma A.9. *Let f be multilinear and nonnegative on $[0,1]^n$. At every interior point with $f(\theta) > 0$,*
$$-\theta_k \leq \frac{\partial \log f}{\partial u_k} \leq 1 - \theta_k, \qquad \left|\frac{\partial \log f}{\partial u_k}\right| \leq 1.$$

Proof. Write $f = (1 - \theta_k) f_0 + \theta_k f_1$, where $f_0, f_1 \geq 0$ are the values with coordinate k fixed to zero and one. Then
$$\frac{\partial \log f}{\partial u_k} = \frac{\theta_k f_1}{f} - \theta_k \in [-\theta_k, 1 - \theta_k].$$
□

For an active bidder define
$$R_i(t, \theta) := \log \frac{P_i(\theta)}{t}. \tag{52}$$

Lemma A.10. *A numerical step changes R_i by at most $(n+1)\xi$ for every bidder that is active during that step.*

Proof. The exact segment preserves R_i by Lemma A.2. Lemma A.9 gives $|\partial R_i/\partial u_k| \le 1$, while $|\partial R_i/\partial u_0| = 1 - t \le 1$. The mean-value theorem in u therefore shows that a coordinatewise error at most ξ changes R_i by at most $(n+1)\xi$. Monotone clamping does not enlarge the local error: if an approximate new coordinate falls below the old one, the exact new coordinate lies above the old one, so the old value is still within the same error interval. □

The step bound in Section A.6.3 uses only coordinate monotonicity and the upper constraint on ξ, independently of payment drift. By (47), over all at most $K_* + 1$ rounded segments, every bidder, until the moment she freezes or the algorithm terminates, satisfies

$$\left|R_i - R_i(t_0, \theta^0)\right| \le \frac{\varepsilon}{1000}. \tag{53}$$

This direct restart argument is why no global condition number or global Gronwall estimate is needed.

A.6.3 Polynomially Many Steps

Let

$$\Phi(t, \theta) = \log \frac{t}{1-t} + \sum_i \log \frac{\theta_i}{1-\theta_i}, \tag{54}$$

where frozen coordinates remain in the sum but no longer move. During an exact full step,

$$\Delta\Phi = \int_0^h \left(a + \sum_{i \in A} y_i\right) \mathrm{d}s = h. \tag{55}$$

After rounding and monotone clamping,

$$\Delta\Phi \ge h - (n+1)\xi \ge \frac{3h}{4}. \tag{56}$$

At the start of a step, t and every active θ_i are below $1 - \varepsilon/2$. Their logit speeds are at most one, and rounding changes each logit by at most ξ. Since $h + \xi \le 1/8$, every such coordinate x has, throughout the exact step and after rounding,

$$\frac{x}{1-x} \le \frac{2e^{1/8}}{\varepsilon}, \qquad 1 - x \ge \frac{\varepsilon}{\varepsilon + 2e^{1/8}} \ge \frac{\varepsilon}{4}. \tag{57}$$

Frozen coordinates retain this bound. Monotonicity gives $t \ge t_0$ and $\theta_i \ge \theta_i^0$. Hence, with Q from (45), throughout each exact step and at every stored state,

$$t, 1-t, \theta_i, 1-\theta_i \ge 2^{-Q}. \tag{58}$$

Initially each logit is at least $-Q$. At every stored state each logit is at most $\log(4/\varepsilon) \le p + 2$. The total increase of Φ is therefore at most $(n+1)(Q+p+2)$. Together with (56), this bounds the number of steps by

$$\frac{4(n+1)(Q+p+2)}{3h} < K_* = O(n^2(Q+p)).$$

Thus the algorithm terminates within the stated bound, and (50) follows from (57).

A.6.4 Budget Feasibility and Approximate Complementarity

Lemma A.11 (Correctness of the Output)**.** *The vector returned by the algorithm is an ε-approximate throttling equilibrium.*

Proof. Initialization and (53) give, until bidder i freezes,

$$(1 - \varepsilon/1000)e^{-\varepsilon/1000} t B_i \le P_i \le e^{\varepsilon/1000} t B_i.$$

At every stored state, (57) gives $t \le 1 - \varepsilon/4$, so $P_i \le (1 - \varepsilon/4)e^{\varepsilon/1000} B_i < B_i$ while i is active. The strict inequality follows from $\log(1 - \varepsilon/4) \le -\varepsilon/4$. After freezing, her payment can only decrease by Lemma A.2. Thus all budgets remain feasible.

Every frozen bidder has $\theta_i \geq 1 - \varepsilon/2 > 1 - \varepsilon$. If any bidder remains active at termination, then $t \geq 1 - \varepsilon/2$ and

$$P_i \geq (1 - \varepsilon/2)(1 - \varepsilon/1000)e^{-\varepsilon/1000}B_i \geq (1 - \varepsilon/2 - 2\varepsilon/1000)B_i > (1 - \varepsilon)B_i,$$

using $e^{-x} \geq 1 - x$ and $\prod_j (1 - a_j) \geq 1 - \sum_j a_j$ for $a_j \in [0, 1]$. Restored zero bidders have participation one and payment zero. These are precisely the approximate-equilibrium conditions. □

A.6.5 Bit Complexity

Assume every nonzero input numerator and denominator has at most L bits. Then

$$2^{-L} \leq B_i \leq 2^L, \qquad b_{ij} > 0 \implies 2^{-L} \leq b_{ij} \leq 2^L, \qquad 2^{-L} \leq M_i \leq m2^L.$$

Equations (43)–(44) imply

$$\log \frac{1}{t_0} + \max_i \log \frac{1}{\theta_i^0} = O(L + p + \log(nm)). \tag{59}$$

Complexity Ledger. First, (59) gives $Q = \mathrm{poly}(N, p)$, so $K_* = O(n^2(Q + p)) = \mathrm{poly}(N, p)$ and $\log(1/\xi) = \mathrm{poly}(N, p)$. In particular, b_{grid} and $q = \lceil \log_2(8/\xi) \rceil$ are polynomially bounded.

Next bound the stored state *before* using any coefficient-size lemma. By (58), restart odds lie in $[2^{-Q}, 2^Q]$. Their exact endpoint grows by at most e^h; truncation and rounding change it by at most $\xi\omega(0)/2$. Thus rounded odds have $O(Q)$ integer bits and b_{grid} fractional bits. Exact conversion $\widehat{x} = \widehat{\omega}/(1 + \widehat{\omega})$ preserves $O(Q + b_{\mathrm{grid}})$ bits. Clamping or freezing retains an earlier rational value. Initial probabilities need not lie on the odds grid, and their encodings can be longer than $O(Q + b_{\mathrm{grid}})$. Their lengths are $O(N + p)$, however, so a uniform per-coordinate bound is $O(N + Q + b_{\mathrm{grid}}) = O(N + Q + q)$. Induction therefore gives total restart-state length $G = O(n(N + Q + q)) = \mathrm{poly}(N, p)$, independently of the Taylor coefficient bounds. This also accounts for a coordinate that is repeatedly clamped to its initial value.

Finally apply Lemma A.6 with this bound on G. Every Taylor coefficient has $\mathrm{poly}(N, p)$ bits, and the convolution and fraction-free solves in Lemma A.7 take polynomial bit time. At most K_* such steps are performed. The total running time and output length are $\mathrm{poly}(N, p)$; since $p = O(1 + \log(1/\delta))$, Lemma A.11 proves Theorem 1.1.

Source of the Polylogarithmic Precision Dependence. The earlier multiplicative tâtonnement changes an underspending bidder by a factor $1/(1 - \delta) = 1 + \Theta(\delta)$ and has an $O(1/\delta)$ dependence on precision in its round bound (Chen et al., 2021). Here the normalized path has polynomial length in logit coordinates, while one additional Taylor coefficient reduces local error geometrically. Thus the requested precision enters through its bit length.

B General Payment Rules and Matroid Reliability

B.1 Payment Conditions, Uniqueness, and Computation

The high-precision proof uses properties that extend beyond the single-slot payment formula. We isolate them below, prove existence and uniqueness, and specify the evaluation access needed for polynomial-bit computation.

B.1.1 Payments Under Entry

For each auction j and participating set $S \subseteq [n]$, let $\pi_i^j(S) \in \mathbb{R}_{\geq 0}$ be bidder i's payment, with $\pi_i^j(S) = 0$ whenever $i \notin S$. Independent participation with probabilities θ gives

$$P_i(\theta) = \sum_j \sum_{S \subseteq [n]} \pi_i^j(S) \prod_{k \in S} \theta_k \prod_{k \notin S} (1 - \theta_k). \tag{60}$$

The structural statements allow real payments. The computational statements assume a rational instance with the evaluation access below.

Definition B.1 (Payment Conditions)**.** We impose the following conditions for every auction j, every $S \subseteq [n]$, and every $k \notin S$,

$$\pi_i^j(S \cup \{k\}) \le \pi_i^j(S) \qquad \text{for every } i \in S, \tag{61}$$

$$\pi_k^j(S \cup \{k\}) \ge \sum_{i \in S} \big(\pi_i^j(S) - \pi_i^j(S \cup \{k\})\big). \tag{62}$$

Thus entry weakly lowers every incumbent payment, and the entrant's payment covers the total payment displaced from incumbents.

Let

$$M_i := \sum_j \pi_i^j(\{i\}).$$

If $M_i = 0$, nonnegativity and repeated use of (61) imply $P_i \equiv 0$. Moreover, (62) implies that adding such a bidder cannot lower any positive incumbent payment. We may therefore delete the bidder and restore her with participation one. Henceforth all retained bidders have $M_i > 0$.

Lemma B.2. *In a market satisfying* (61)–(62)*, P_i is nonincreasing in every competitor coordinate θ_k, $k \neq i$. More generally, for every $C \subseteq [n]$,*

$$R_C(\theta) := \sum_{i \in C} P_i(\theta)$$

is nondecreasing in each coordinate θ_k with $k \in C$.

If $k \in C$, $P_k(\theta) > 0$, and $\theta_\ell < 1$ for every $\ell \in C \setminus \{k\}$, then

$$\frac{\partial R_C}{\partial \theta_k}(\theta) \ge \frac{P_k(\theta)}{\theta_k} \prod_{\ell \in C \setminus \{k\}} (1 - \theta_\ell) > 0. \tag{63}$$

Proof. Fix k and expose the random participation set $S \subseteq [n] \setminus \{k\}$ of all other bidders. For $i \neq k$,

$$\frac{\partial P_i}{\partial \theta_k} = \sum_j \mathbb{E}_S \left[\pi_i^j(S \cup \{k\}) - \pi_i^j(S)\right] \le 0$$

by (61). For $k \in C$,

$$\frac{\partial R_C}{\partial \theta_k} = \sum_j \mathbb{E}_S \Bigg[\pi_k^j(S \cup \{k\}) - \sum_{i \in C \cap S} \big(\pi_i^j(S) - \pi_i^j(S \cup \{k\})\big)\Bigg]. \tag{64}$$

Every loss in the inner sum is nonnegative, and the sum is taken over only a subset of the incumbents in S. Hence (62) makes the bracket nonnegative.

When every member of $C \setminus \{k\}$ is absent, the bracket is just k's payment upon entry. By (61), removing these competitors can only increase that payment, so its conditional expectation is at least $P_k(\theta)/\theta_k$. The event has probability $\prod_{\ell \in C \setminus \{k\}}(1 - \theta_\ell)$, proving (63). □

B.1.2 The Exact Equilibrium Is the Largest Feasible Vector

Let

$$\mathcal{F}(B) := \{\theta \in [0,1]^n : P_i(\theta) \le B_i \ \forall i\}$$

be the budget-feasible set.

Lemma B.3. *If $\theta, \widehat{\theta} \in \mathcal{F}(B)$, then their coordinatewise maximum $\theta \vee \widehat{\theta}$ also belongs to $\mathcal{F}(B)$. Consequently, $\mathcal{F}(B)$ has a unique coordinatewise largest element.*

Proof. Put $\psi = \theta \vee \widehat{\theta}$. For each bidder i, choose one of the two feasible vectors, say θ, with $\psi_i = \theta_i$. Since $\psi_{-i} \geq \theta_{-i}$ coordinatewise and competitor participation can only lower i's payment by Lemma B.2,

$$P_i(\psi) \leq P_i(\theta) \leq B_i.$$

Thus ψ is feasible.

Choose $\theta^\star$ maximizing $\sum_i \theta_i$ on the nonempty compact feasible set. For any feasible θ, max closure makes $\theta \vee \theta^\star$ feasible, so maximality of the sum forces $\theta \vee \theta^\star = \theta^\star$. Thus $\theta^\star$ is its largest element. □

Theorem B.4 (Existence and Uniqueness Under the Payment Conditions)**.** *Every market satisfying* (61)–(62) *with positive budgets has a unique exact throttling equilibrium $\theta^\star$. It is the coordinatewise largest budget-feasible vector.*

Proof. Let $\theta^\star$ be the largest feasible vector from Lemma B.3. Suppose $\theta_i^\star < 1$ and $P_i(\theta^\star) < B_i$. Holding all competitors fixed, P_i is linear in θ_i. Hence $\theta_i^\star$ can be increased slightly without violating i's budget; all other payments weakly decrease by Lemma B.2. This contradicts maximality. Therefore $\theta_i^\star < 1$ implies $P_i(\theta^\star) = B_i$, so $\theta^\star$ is an exact equilibrium.

Let θ be another exact equilibrium. Feasibility gives $\theta \leq \theta^\star$. If $C := \{i : \theta_i < \theta_i^\star\}$ is nonempty, then $\theta_i < 1$ and $P_i(\theta) = B_i > 0$ for every $i \in C$. Choose $k \in C$. Group monotonicity and affinity in θ_k give

$$\sum_{i \in C} B_i \geq R_C(\theta^\star) \geq R_C(\theta) + (\theta_k^\star - \theta_k)\partial_k R_C(\theta) > R_C(\theta) = \sum_{i \in C} B_i.$$

The strict inequality is (63). This contradiction proves uniqueness. □

Corollary B.5. *If $B' \geq B$ coordinatewise, then $\theta^\star(B') \geq \theta^\star(B)$. Moreover total expected revenue $\sum_i P_i(\theta^\star(B))$ is nondecreasing in the budget vector.*

Proof. The inclusion $\mathcal{F}(B) \subseteq \mathcal{F}(B')$ and the largest-feasible characterization give the participation comparison. Taking $C = [n]$ in Lemma B.2 shows that total revenue is coordinatewise nondecreasing in θ. □

B.1.3 The Jacobian and Exact Evaluation

Lemma B.6. *Consider a market satisfying* (61)–(62) *and a state with $0 < \theta_i < 1$ for every nonzero bidder. In log-odds coordinates, every active principal Jacobian $H_A = [\partial P_i / \partial u_k]_{i,k \in A}$ is a nonsingular M-matrix. Moreover, each off-diagonal weight $-H_{ik}$ and each column surplus*

$$\kappa_k^A := H_{kk} - \sum_{i \in A \setminus \{k\}} (-H_{ik})$$

is a polynomial with nonnegative coefficients in the factors θ_ℓ and $1 - \theta_\ell$; the surplus is strictly positive.

Proof. The chain rule and Lemma B.2 give

$$H_{ik} = \theta_k(1 - \theta_k)\partial_k P_i \leq 0 \quad (i \neq k), \qquad \kappa_k^A = \theta_k(1 - \theta_k)\partial_k R_A.$$

For $-H_{ik}$, the conditional payment differences in (61) are nonnegative. For κ_k^A, use the nonnegative bracket in (64) with $C = A$. Expanding these expectations gives the claimed nonnegative polynomials, each of degree at most $n + 1$. The empty-set event also gives

$$\kappa_k^A \geq \theta_k(1 - \theta_k) M_k \prod_{\ell \neq k} (1 - \theta_\ell) > 0.$$

Thus H_A is strictly column diagonally dominant; the Neumann-series argument in Lemma A.1 proves $H_A^{-1} \geq 0$. □

Definition B.7 (Exact Payment Evaluation)**.** The compact instance explicitly indexes its n bidders, so n is at most the instance length. A *uniform exact evaluator* is a deterministic binary Turing algorithm that, on an instance and rational $\theta \in [0, 1]^n$, outputs every $P_i(\theta)$ as a reduced binary rational in time polynomial in the combined instance and point encoding lengths. One polynomial running-time bound applies to all instances and query points.

Singleton payments are obtained at Boolean points. The following observation also supplies all series needed by the algorithm.

Lemma B.8. *Under Definition B.7, rational order-q series $x(s)$ with interior constant terms permit computation of $P_i(x(s))$ and $H_{ik}(x(s))$ modulo s^{q+1} in polynomial bit complexity.*

Proof. For any multilinear polynomial f on n variables, $f(x(s))$ has degree at most nq. Evaluate it at $nq+1$ equally spaced rational nodes $s_r = r\lambda$ and interpolate. Choose the dyadic λ small enough that every $x_k(s_r)$ stays in $[0,1]$: if coefficient magnitudes are at most 2^G and constant terms have boundary distance at least 2^{-Q_0}, one may take

$$\lambda \le \frac{2^{-G-Q_0-2}}{(nq+1)(q+1)}.$$

At these nodes, the change in any coordinate is at most $q2^G|s_r| \le 2^{-Q_0}/4$. The nodes, evaluations, and interpolation system all have polynomial encoding length, so exact rational elimination has polynomial bit cost. Apply this first to $f = P_i$. Multilinearity gives

$$H_{ik}(x(s)) = x_k(s)(1-x_k(s))\big[P_i(x(s)^{k\leftarrow 1}) - P_i(x(s)^{k\leftarrow 0})\big].$$

The two restricted polynomials are recovered by the same interpolation, then multiplied by the displayed quadratic factor and truncated. □

Interpolation supplies the required series, but their coefficient lengths must also be controlled. The next lemma extracts a common-denominator bound from the output lengths of point evaluations.

Lemma B.9. *Let $f \in \mathbb{Q}[x_1,\dots,x_n]$ be multilinear, $n \ge 1$, and put $K = 2n+1$. Suppose every value of f on both*

$$\{0,1\}^n \quad \text{and} \quad \{0,1/K,\dots,(K-1)/K\}^n$$

has reduced denominator at most 2^H. If D is the least positive integer such that Df has integer coefficients, then

$$\log_2 D \le (K+2)H.$$

If in addition $|f(v)| \le 2^H$ at every Boolean vertex, all monomial coefficients have numerators and a common denominator of $O(nH+n)$ bits.

Proof. Write $g = Df \in \mathbb{Z}[x_1,\dots,x_n]$. For every prime power $p^a \parallel D$, at least one coefficient of g is nonzero modulo p. If $p \le K$, a nonzero multilinear polynomial cannot vanish on the entire Boolean cube, because Boolean values determine its coefficients by inclusion–exclusion over every field. At a vertex with $g(v) \not\equiv 0 \pmod p$, the reduced denominator of $f(v)$ contains p^a. Thus the contribution of each such prime to $\log_2 D$ is at most H, and all primes $p \le K$ contribute at most KH.

For $p > K$, define $g_K(r) = K^n g(r_1/K,\dots,r_n/K)$. Modulo p this is a nonzero polynomial of total degree at most n. On the grid $\{0,\dots,K-1\}^n$, the Schwartz–Zippel bound (Schwartz, 1980) gives

$$\Pr[g_K(r) \not\equiv 0 \pmod p] \ge 1 - n/K > 1/2.$$

Since $p \nmid K$, on this event the reduced denominator of $f(r/K)$ retains the full factor p^a. Taking the expectation of the logarithm of that denominator yields

$$H \ge \frac{1}{2} \sum_{\substack{p^a \parallel D \\ p>K}} a \log_2 p.$$

Combining the two prime ranges proves the denominator bound. Finally, inclusion–exclusion expresses each monomial coefficient as a signed sum of at most 2^n Boolean values; the claimed numerator bound follows: its absolute value after multiplication by D is at most $D2^{H+n}$. This argument bounds a common denominator, not just the denominators of individual coefficients. □

Corollary B.10. *For a set-payment representation with the access in Definition B.7 on n explicitly indexed bidders, every payment polynomial has degree and coefficient height polynomial in the compact instance length, even when its expanded support is exponential.*

Proof. Use rational point evaluation at the points in Lemma B.9. Every such query has polynomial encoding length. Because a Turing machine must write its reduced rational output, the uniform polynomial running-time bound gives one polynomial bound H on all output numerator and denominator lengths. Apply the lemma to each of the polynomially many payment functions. This is only a size argument: neither the grid nor the expanded payment polynomials are enumerated by the algorithm. □

Theorem B.11 (Effective Computation Under the Payment Conditions)**.** *Every market satisfying* (61)–(62) *and Definition B.7 admits a deterministic algorithm computing a rational δ-approximate throttling equilibrium in*

$$\mathrm{poly}\left(N, \log \frac{1}{\delta}\right)$$

bit operations, where N is the compact input length.

Proof. Lemma B.6 supplies the nonnegative off-diagonal weights and positive column surpluses. Together with the root weights $(1-t)P_i$, all arc weights have degree at most $n+1$ in logistic factors. In an arc with tail k, only the pair $\theta_k, 1-\theta_k$ can occur twice; every other pair occurs at most once. Thus the tree polynomials have total degree at most $n(n+1)$ and degree at most $n+1$ in each pair. Frozen probabilities enter these expansions only as nonnegative real constants. Corollary A.4 and the uniform analytic step therefore apply with the same radius for every active set, without enumerating the expansions.

Lemma B.8 replaces explicit expansion in Lemma A.7: the coefficient recurrence still solves the same rational augmented matrix at every order, while the oracle supplies the truncated series of P and H. Corollary B.10 gives a uniform polynomial height bound for the hidden multilinear payment polynomials. At every rational restart point, Cramer's rule again turns the normalized flow into a rational ODE with polynomial degree and height. Lemma A.6 also covers the odds as in Lemma A.7, so every coefficient supplied to subsequent oracle calls has polynomial bit length. This closes the finite-precision recurrence without expanding the payment functions.

For initialization, conditional payment given participation is at most M_i by (61) and equals M_i when no competitor participates. Hence the proof of Lemma A.8 applies. Frozen payments never rise by Lemma B.2, and the restart-rounding estimate follows from Lemma A.9. All remaining path-length, coordinate-size, and correctness arguments are identical to those proving Theorem 1.1. □

Corollary B.12. *If all rational values $\pi_i^j(S)$ are listed explicitly, a δ-approximate equilibrium is computable in* $\mathrm{poly}(N_{\mathrm{exp}}, \log(1/\delta))$ *bit operations.*

Proof. Evaluate (60) by summing the explicitly listed Bernoulli monomials in polynomial time. Theorem B.11 applies. □

B.2 Matroid First-Price Auctions and Pointed Reliability

In auction j, let M_j be a matroid on an explicitly specified subset of eligible bidders; bidders outside this subset have zero payment in that auction. Given a participating set, process bidders in strict bid-priority order and select a bidder whenever adding her preserves independence. Every selected bidder pays her bid. We call this a first-price matroid auction.

B.2.1 Entry in a Greedy Matroid Allocation

Lemma B.13. *Let $G(S)$ be the greedy basis of the restriction $M|S$ under a fixed strict priority order, and let $k \notin S$. Then $G(S \cup \{k\})$ is one of*

$$G(S), \qquad G(S) \cup \{k\}, \qquad G(S) \cup \{k\} \setminus \{e\},$$

where in the last case e has lower priority than k.

Proof. Greedy accepts an element i exactly when $i \notin \mathrm{cl}_M(S_{\succ i})$: the accepted prefix is a basis of the processed prefix (Oxley, 2011). Adding k only enlarges the higher-priority set of an incumbent. Monotonicity of closure therefore gives

$$G(S \cup \{k\}) \setminus \{k\} \subseteq G(S).$$

Both greedy sets are bases, and $\mathrm{rank}(S \cup \{k\}) - \mathrm{rank}(S) \in \{0, 1\}$. If k is rejected, cardinalities force equality of the two greedy sets. If k is accepted, the displayed inclusion and the rank bound show that at most one incumbent is removed. All choices before k are unchanged, so any removed incumbent has lower priority. □

Theorem B.14 (Entry in Matroid Auctions)**.** *Every first-price matroid auction satisfies* (61)–(62)*. Consequently every finite market of such auctions has a unique exact throttling equilibrium.*

Proof. By Lemma B.13, inserting bidder k either changes nothing, adds k, or replaces one selected bidder e by k. Existing selected bidders other than e keep the same payment, while e's payment falls to zero. In the exchange case k has at least as high a bid as e, so k's first-price payment covers the only displaced payment. This is precisely (61)–(62). Apply Theorem B.4. □

B.2.2 The Pointed-Reliability Representation

For a matroid M, an element e, and independent inclusion probabilities $x = (x_f)_{f \neq e}$, define

$$Q_{M,e}(x) := \Pr_{S \sim x}\big[e \notin \mathrm{cl}_M(S)\big]. \tag{65}$$

This is a multilinear polynomial in x.

Lemma B.15. *In a first-price matroid auction, let H_i be the bidders with higher priority than i. Then*

$$P_i(\theta) = \theta_i b_i Q_{M|(H_i \cup \{i\}),i}(\theta_{H_i}). \tag{66}$$

For a market of several auctions, the corresponding terms add.

Proof. Conditional on bidder i participating, greedy accepts her exactly when she is not spanned by the active higher-priority elements. Lower-priority bidders are processed later and cannot affect this event. A selected bidder pays b_i. □

We now make the representation boundary precise. Fix a uniform succinct matroid class $\mathfrak{C}$. An input explicitly indexes a ground set $E = [r]$, with r at most the representation length. Deterministic polynomial-time algorithms implement restriction, deletion, contraction, and loop testing, and every transformed representation has size polynomial in the original size; singleton matroids are available as well. Define Pointed-Reliability$_{\mathfrak{C}}$ to output the exact reduced rational value of (65) at a binary-encoded rational product distribution.

Theorem B.16 (Reliability Completeness)**.** *For every uniform representation class $\mathfrak{C}$ satisfying the access and closure assumptions above, the following two functional problems are polynomial-time Turing equivalent:*

(i) exact evaluation of the pointed-reliability polynomial for $\mathfrak{C}$ at rational product probabilities;

(ii) given a first-price $\mathfrak{C}$-matroid market and a precision parameter p in unary, return a rational 2^{-p}-approximate throttling equilibrium of polynomial encoding length.

At most one query to (ii) suffices to solve (i), using one common auction with equal bids and private unit-bid singleton auctions.

Proof of (i) ⇒ (ii). Lemma B.15 evaluates every payment with a polynomial number of reliability queries. Every oracle answer lies in $[0, 1]$ and has denominator dividing the product of the query-probability denominators, hence polynomial output length. The coefficient-height argument in Corollary B.10 therefore applies also relative to this oracle: it uses output length, not the cost of obtaining that output. Matroid auctions satisfy the payment conditions by Theorem B.14. Execute the algorithm of Theorem B.11, replacing each payment evaluation by these queries. The interpolation and rational-series recurrences make only polynomially many queries of polynomial encoding length, with polynomial additional bit work. This is a polynomial-time Turing reduction, not an assumption that reliability itself is efficiently computable. □

Proof of (ii) ⇒ (i). Let (M, e, x) be a pointed-reliability instance. Delete every element with $x_f = 0$ and contract every element with $x_f = 1$. If e is a loop, return zero. If no ordinary element remains, return one. Otherwise let $m \geq 1$ be their number, write $x_f = a_f/d_f$ in lowest terms, and put $D = \prod_f d_f$. Its binary length is $\log D = \sum_f \log d_f$, hence polynomial in the reliability input. Since e is not a loop and $0 < a_f < d_f$, the empty-set event gives

$$DQ_{M,e}(x) \in \mathbb{Z}, \qquad Q_{M,e}(x) \geq \prod_f (1 - x_f) \geq \frac{1}{D}, \qquad x_f \leq 1 - \frac{1}{D}.$$

Choose the dyadic accuracy

$$p = \lceil \log_2(4mD) \rceil, \qquad \varepsilon = 2^{-p} \leq \frac{1}{4mD} \leq \frac{1}{8}. \tag{67}$$

In one common matroid auction on M, every bidder bids ε, and e is last in the tie-breaking order. Each ordinary bidder f also has a private singleton auction with bid one. Set

$$B_f = x_f, \qquad B_e = \frac{\varepsilon}{2D}.$$

The unary length of p is polynomial in the input length, and all bids and budgets have polynomial binary encodings. Query an ε-approximate equilibrium θ.

An ordinary bidder's conditional payment lies in $[1, 1 + \varepsilon]$. Feasibility gives $\theta_f \leq x_f < 1 - \varepsilon$, so approximate complementarity gives $\theta_f(1 + \varepsilon) \geq (1 - \varepsilon)x_f$. Thus

$$\frac{(1 - \varepsilon)x_f}{1 + \varepsilon} \leq \theta_f \leq x_f, \qquad 0 \leq x_f - \theta_f \leq 2\varepsilon. \tag{68}$$

Pointed reliability is nonincreasing in each probability, so $Q_{M,e}(\theta_{-e}) \geq Q_{M,e}(x) \geq 1/D$. Since $P_e = \varepsilon\theta_e Q_{M,e}(\theta_{-e})$, feasibility forces $\theta_e \leq 1/2 < 1 - \varepsilon$. Approximate complementarity implies $\theta_e > 0$ and, with $\widehat{Q} = 1/(2D\theta_e)$, gives

$$Q_{M,e}(\theta_{-e}) \leq \widehat{Q} \leq \frac{Q_{M,e}(\theta_{-e})}{1 - \varepsilon} \leq \frac{1}{1 - \varepsilon}. \tag{69}$$

Coupling the product measures with common uniforms bounds the change of the event probability by the sum of coordinate changes. Therefore

$$0 \leq \widehat{Q} - Q_{M,e}(x) \leq \frac{\varepsilon}{1 - \varepsilon} + \sum_f (x_f - \theta_f) < 4m\varepsilon \leq \frac{1}{D}. \tag{70}$$

Here $\varepsilon/(1 - \varepsilon) < 2\varepsilon$ and $m \geq 1$. Since $DQ_{M,e}(x)$ is an integer, return $\lfloor D\widehat{Q} \rfloor / D$. This uses one equilibrium query and polynomial-time rational arithmetic. □

B.2.3 Exact Evaluation for Laminar Matroids

A laminar matroid is specified by a laminar family (Fife & Oxley, 2017). We may add the full ground set as a root of capacity n and replace each capacity c_A by $\min\{c_A, |A|\}$. A set I is independent iff $|I \cap A| \leq c_A$ for every node A.

Lemma B.17. *Exact pointed reliability for laminar matroids is computable in polynomial bit complexity.*

Proof. Adding one element changes rank by zero or one, so

$$Q_{M,e}(x) = \mathbb{E}\big[\operatorname{rank}(S \cup \{e\}) - \operatorname{rank}(S)\big]. \tag{71}$$

It therefore suffices to compute expected rank twice, with e always present and always absent.

Represent the laminar family as a rooted tree. At a node A, let X_A count participating elements belonging to no child, and let R_A be the rank within the subtree. Then

$$R_A = \min\left\{ c_A,\ X_A + \sum_{C \text{ child of } A} R_C \right\}. \tag{72}$$

Indeed, independent sets from distinct children can be united, with only the capacity of A left to impose. The summands use disjoint participation coins. Convolve their distributions and cap the result at c_A. Each distribution has at most $n+1$ entries. Running this recursion with the participation probability of e fixed to one and to zero gives the two expectations in (71). All probabilities have a common denominator dividing the product of the input denominators, so both the arithmetic count and bit lengths are polynomial. □

Corollary B.18 (Laminar-Matroid Equilibria)**.** *First-price laminar-matroid markets have a unique exact equilibrium, and a rational δ-approximate equilibrium is computable in* $\mathrm{poly}(N, \log(1/\delta))$ *bit operations. This includes uniform and partition matroids.*

Proof. Combine Theorems B.14 and B.16 with Lemma B.17. □

B.2.4 Position Auctions as Sums of Uniform-Matroid Auctions

Auction j has rational slot qualities $\alpha_{j1} \geq \cdots \geq \alpha_{js_j} \geq 0$, with $\alpha_{j,s_j+1} = 0$. The participating bidder at rank $r \leq s_j$ pays $\alpha_{jr} b_{ij}$; bidders below the last slot pay zero.

Corollary B.19 (Position-Auction Equilibria)**.** *Multi-slot first-price position auctions satisfy* (61)–(62) *and have a unique exact throttling equilibrium. A rational δ-approximate equilibrium is computable in* $\mathrm{poly}(N_{\mathrm{pos}}, \log(1/\delta))$ *bit operations.*

Proof. Let $\pi_i^{j,r}(S)$ be i's first-price payment in the uniform-matroid auction that accepts the first r bidders from S, with the same bids and priorities. For every participating set,

$$\pi_i^j(S) = \sum_{r=1}^{s_j} (\alpha_{jr} - \alpha_{j,r+1}) \pi_i^{j,r}(S). \tag{73}$$

Indeed, a bidder at rank q is accepted exactly in the terms with $r \geq q$, whose coefficients sum to α_{jq} (or zero if $q > s_j$). Absorb each positive coefficient into the bids of its uniform-matroid auction and discard zero coefficients. This gives a polynomial-size laminar-matroid market with exactly the same payments for every θ. Apply Theorem B.14 and Corollary B.18. □

B.2.5 Graphic Matroids and Network Reliability

For the graphic matroid of an undirected graph $G = (V, E)$ and an additional edge $e = (s, t)$,

$$e \in \mathrm{cl}(S) \quad \Longleftrightarrow \quad s \text{ and } t \text{ are connected in } (V, S).$$

Thus pointed reliability is exactly two-terminal network unreliability. Exact evaluation is #P-hard (Provan, 1986).

Corollary B.20 (Graphic-Matroid Hardness)**.** *Computing high-precision approximate throttling equilibria in succinct first-price graphic-matroid markets is #P-hard under polynomial-time Turing reductions. More precisely, a polynomial-time algorithm in the binary market input and* $\log(1/\delta)$ *would place exact two-terminal network unreliability—a #P-hard evaluation problem—in* FP*. The hardness holds with one common graphic-matroid auction having equal bids and private unit-bid singleton auctions.*

Proof. The graphic-matroid representation supports restriction, deletion, contraction, and loop testing in polynomial time. Two-terminal unreliability is an instance of pointed reliability, so the reverse reduction in Theorem B.16 recovers it exactly from a high-precision equilibrium oracle. □

C Proofs for the Statistical Results

Lemma C.1. *Let $0 < \gamma \leq 1$ and $0 < \alpha < 1$. If*

$$T \geq \frac{4R}{\gamma^2} \left[n \log 2 + \log \frac{2n}{\alpha} \right], \tag{74}$$

then with probability at least $1 - \alpha$, simultaneously for every i and $\theta \in [0,1]^n$,

$$|\widehat{P}_i(\theta) - P_i(\theta)| \leq \gamma \big(B_i + P_i(\theta) \big). \tag{75}$$

Proof. Fix a bidder i and a vertex v. For $g = f_i(v; Z)/B_i \in [0, R]$ and $\mu = \mathbb{E}g$, we have $\mathrm{Var}(g) \le R\mu$. Bernstein's inequality (Boucheron et al., 2013) with $u = \gamma(1+\mu)$ gives

$$\Pr(|\widehat{\mu} - \mu| > u) \le 2\exp\left(-\frac{Tu^2}{2(R\mu + Ru/3)}\right) \le 2\exp\left(-\frac{T\gamma^2}{4R}\right).$$

The last step uses $\mu + \gamma(1+\mu)/3 \le 4(1+\mu)/3$ and $\mu \ge 0$. A union bound over $n2^n$ pairs gives (75) at all vertices. To extend it to every θ, use multilinearity. For $v \in \{0,1\}^n$,

$$w_v(\theta) = \prod_{k:v_k=1} \theta_k \prod_{k:v_k=0} (1-\theta_k), \qquad P_i(\theta) = \sum_v w_v(\theta) P_i(v). \tag{76}$$

The same identity holds for $\widehat{P}_i$. Since $w_v \ge 0$ and $\sum_v w_v = 1$, relative bounds at the vertices imply

$$|\widehat{P}_i(\theta) - P_i(\theta)| \le \gamma(B_i + P_i(\theta)) \quad \text{for every } i, \theta. \tag{77}$$

The vertex weights are used only in the proof; the learner does not enumerate them. □

C.1 Empirical-to-Population Transfer

Proof of Theorem 5.3. Apply Lemma C.1 with $\gamma = \delta/16$ and the sample size in (19). Work on the resulting uniform concentration event. If $\theta_i < 1 - \delta$, then also $\theta_i < 1 - \delta/8$, so the empirical complementarity condition applies. Empirical feasibility and complementarity therefore imply, respectively,

$$\frac{P_i}{B_i} \le \frac{1 - 3\delta/16}{1 - \delta/16} < 1, \qquad \theta_i < 1 - \delta \implies \frac{P_i}{B_i} \ge \frac{(1-\delta/8)(1-\delta/4) - \delta/16}{1 + \delta/16} \ge 1 - \delta/2.$$

These are the population equilibrium conditions. Assigning weight $1/T$ to each sampled context gives a rational market of polynomial encoding length; Theorem 1.1 or Corollary B.19 computes the required approximate equilibrium. □

C.2 A One-Bidder Sample Lower Bound

Theorem C.2 (One-Bidder Sample Lower Bound)**.** *For $R \ge 2$, $0 < \delta \le 1/8$, and $0 < \alpha \le 1/4$, any learner successful on every admissible context distribution needs $\Omega((R/\delta^2)\log(1/\alpha))$ samples in the worst case, even with one bidder and one good.*

Proof of Theorem C.2. Set $B = 1$ and $b = \lfloor R \rfloor \ge R/2$. A context contains a bid of b with probability q and a zero bid otherwise. Compare

$$q_- = (1 - 2\delta)/b, \qquad q_+ = (1 + 2\delta)/b.$$

Then $P_\pm(\theta) = (1 \pm 2\delta)\theta$. Under the minus distribution every δ-approximate equilibrium has $\theta \ge 1 - \delta$. Under the plus distribution feasibility forces $\theta \le 1/(1+2\delta) < 1 - \delta$. A successful learner therefore distinguishes the two distributions with both error probabilities at most α.

For Bernoulli parameters $u, v \in (0,1)$,

$$\mathrm{KL}(\mathrm{Ber}(u) \,\|\, \mathrm{Ber}(v)) \le \frac{(u-v)^2}{v(1-v)}.$$

Here $q_+ \ge 1/b$, $1 - q_+ \ge 3/8$, and $q_+ - q_- = 4\delta/b$, so the one-sample divergence is at most $128\delta^2/(3b)$. Data processing through the test and additivity over samples (Tsybakov, 2009) give

$$\frac{128T\delta^2}{3b} \ge \mathrm{kl}(1-\alpha, \alpha) \ge \frac{1}{4}\log\frac{1}{\alpha}.$$

Since $b \ge R/2$, this proves the lower bound, also for randomized learners. □

C.3 Revenue Marginals and Weighted Errors

Lemma C.3. *In single-slot, position, and matroid first-price markets, including their nonnegative sums and population averages, total expected revenue $V = \sum_i P_i$ satisfies*

$$0 \le x \le y \le \mathbf{1} \quad \Longrightarrow \quad \partial_k V(x) \ge \partial_k V(y) \quad \textit{for every } k.$$

Proof. In one matroid auction, order bids as $b_1 \ge \cdots \ge b_n \ge b_{n+1} = 0$ and put $H_r = \{1, \dots, r\}$. For a participating set S, greedy allocation accepts $\operatorname{rank}_M(S \cap H_r)$ elements among the first r. Thus revenue is

$$v(S) = \sum_{r=1}^{n} (b_r - b_{r+1}) \operatorname{rank}_M(S \cap H_r).$$

Matroid rank is submodular (Oxley, 2011), so v has diminishing set marginals. This is the usual weighted-rank argument (Dughmi et al., 2012). Since

$$\partial_k V(x) = \mathbb{E}_{S \sim x_{-k}}[v(S \cup \{k\}) - v(S)],$$

couple the sets at $x \le y$ by the same uniform random variables. The smaller set has the larger marginal, proving the claim. Single-slot auctions are uniform-matroid auctions; position auctions are nonnegative sums of these by (73). Sums and expectations preserve the inequality. □

Corollary C.4 (Weighted Coordinate Errors)**.** *In any market satisfying* (14)*, define*

$$M_k = \sum_j \pi_k^j(\{k\}), \qquad \Lambda_k^\star = M_k \prod_{\ell \ne k} (1 - \theta_\ell^\star), \qquad \Gamma_k^\star = \partial_k V(\theta^\star). \tag{78}$$

For population payments, M_k is the expected singleton payment. Every δ-approximate equilibrium, $0 < \delta < 1$, satisfies

$$\sum_k w_k(\theta_k^\star - \theta_k) \le \delta V^\star \quad \textit{with } w_k = \Lambda_k^\star. \tag{79}$$

In the auction classes of Lemma C.3, the same bound holds with $w_k = \Gamma_k^\star \ge \Lambda_k^\star$. For either choice, if $w_{\min} := \min_k w_k > 0$, then $\|\theta - \theta^\star\|_\infty \le \delta V^\star / w_{\min}$.

Proof. For $x \le \theta^\star$, nonnegative revenue increments and the event that all other bidders are absent give $\partial_k V(x) \ge M_k \prod_{\ell \ne k}(1 - x_\ell) \ge \Lambda_k^\star$. Lemma C.3 gives $\partial_k V(x) \ge \Gamma_k^\star$ in its stated classes. Integrate the applicable bound along $[\theta, \theta^\star]$ and apply Theorem 5.1:

$$\sum_k w_k(\theta_k^\star - \theta_k) \le V^\star - V(\theta) \le \delta V^\star.$$

□

Corollary C.5 (Coordinates under Positive Revenue Marginals)**.** *Under Theorem 5.3's assumptions, suppose a known bound $\Gamma^\star := \min_k \partial_k V(\theta^\star) \ge \underline{\Gamma} > 0$ holds. Run the residual learner with $\delta_\eta = \min\{1/2, \underline{\Gamma}\eta / \sum_i B_i\}$. It achieves coordinate error η with probability $1 - \alpha$ and sample bound* (19) *with $\delta = \delta_\eta$.*

Proof. Revenue has decreasing coordinate marginals for these auction classes (Lemma C.3), so feasible θ has $g_k(\theta) \ge \Gamma^\star$. Theorem 5.1 gives

$$\|\theta - \theta^\star\|_\infty \le \delta_\eta V^\star / \Gamma^\star \le \eta,$$

since $V^\star \le \sum_i B_i$. Apply Theorem 5.3 with $\delta = \delta_\eta$. For fixed $\underline{\Gamma}$, budgets, n, and R, the resulting sample rate is $O(\eta^{-2} \log(2/\alpha))$. □

The positive-marginal bound is an additional assumption, not a consequence of uniqueness, and it does not require an interior equilibrium.

For private unit-bid auctions with budgets $1/2$, the equilibrium is $\frac{1}{2}\mathbf{1}$ and $\Gamma_k^\star = 1$, whereas $\Lambda_k^\star = 2^{-(n-1)}$. With budgets two, the equilibrium is $\mathbf{1}$ and $\Gamma_k^\star$ is still one.

C.4 Sharp Coordinate Estimation on the Cyclic Family

Proof of Theorem 5.4. For the lower bound, choose $d \in (2\eta, 3\eta)$, put $s = d^n \le 1/4$, and compare $\mu_0 = 1$ with $\mu_1 = 1/(1-s) \le 4/3$. Their equilibria are $\mathbf{1}$ and $(1-d)\mathbf{1}$, so their radius-η balls are disjoint. A successful estimator distinguishes these laws with both errors at most α. Only $X/2$ changes, from Bernoulli mean $p_0 = 1/2$ to $p_1 = 1/(2(1-s))$; J has the same independent law. Data processing and additivity of Bernoulli KL divergence give

$$\mathrm{kl}(1-\alpha, \alpha) \le T\,\mathrm{kl}(p_0, p_1) = \frac{T}{2}\log\left(1 + \frac{s^2}{1-2s}\right) \le Ts^2.$$

The last inequality uses $\log(1+x) \le x$ and $s \le 1/4$. Since $\mathrm{kl}(1-\alpha, \alpha) \ge \frac{1}{4}\log(1/\alpha)$, $T \ge \frac{1}{4}d^{-2n}\log(1/\alpha)$ for every such d. Letting $d \downarrow 2\eta$ gives the stated bound, also for randomized estimators.

For the upper bound, clip the sample mean of X to $[0, 4/3]$, obtaining $\widehat{\mu}$. Set $r = (2\eta)^n/2$ and

$$L = \max\{0, \widehat{\mu} - r\}, \qquad U = \min\{4/3, \widehat{\mu} + r\}, \qquad \widehat{\theta} = \tfrac{1}{2}\big[t^\star(L) + t^\star(U)\big]\mathbf{1}.$$

The map $h(a) = \max\{0, 1 - 1/a\}$, extended by zero at $a = 0$, is 1-Lipschitz on $[0, 4/3]$. Since $|u^{1/n} - v^{1/n}| \le |u - v|^{1/n}$ for $u, v \ge 0$,

$$t^\star(L) - t^\star(U) \le (U - L)^{1/n} \le 2\eta.$$

If $|\widehat{\mu} - \mu| \le r$, monotonicity places the true equilibrium between the two endpoints, within η of their midpoint. Clipping cannot increase error, so Hoeffding's inequality gives $\Pr(|\widehat{\mu} - \mu| > r) \le 2e^{-Tr^2/2}$. The stated sample size makes this at most α, including at $\mu = 1$. □

D Lower Bound for the Multiplicative Update

Proposition D.1 (Tight Precision Dependence of the Multiplicative Update)**.** *For every $0 < \delta \le 1/4$, the multiplicative tâtonnement of Chen, Kroer, and Kumar can require $\Omega(1/\delta)$ update rounds.*

Proof. With one bidder, bid $b = 1$, and budget $B = 1/2$, payment is $P(\theta) = \theta$. The prescribed initialization is $\theta^0 = B/(2b) = 1/4$, and every update multiplies θ by $1/(1-\delta)$ until $\theta \ge (1-\delta)/2$. Therefore the number T of updates satisfies

$$T \ge \frac{\log(2(1-\delta))}{\log(1/(1-\delta))} \ge \frac{\log(3/2)}{2\delta} = \Omega(1/\delta),$$

since $\delta \le 1/4$ implies $2(1-\delta) \ge 3/2$ and $\log(1/(1-\delta)) \le 2\delta$. □