Constructive Lean 4 Formalization and Exact-Rational Certificates for Query-Sample Tradeoffs in Revenue Maximization

Lluis Eriksson*

November 25, 2026

Abstract

Yaonan Jin, Wei Tang, Yifan Wang, and Mengxiao Zhang (`arXiv:2611.31982`, Nov. 2026) established tight query-sample tradeoffs for $(1-\varepsilon)$-approximate monopoly revenue maximization across all four standard valuation families: MHR (L-shaped corner at $S = \Theta(\varepsilon^{-3/2}\log(1/\delta))$, $Q = \Theta(\varepsilon^{-2}\log(1/\delta))$), quasi-MHR (continuous slope--1 frontier $S \cdot Q = \Theta(\varepsilon^{-4}\log^2(1/\delta))$ for $c \in [0,1]$), regular (L-shaped corner at $\Theta(\varepsilon^{-3}\log(1/\delta))$), and quasi-regular (logarithmic boundary separating $Q = \Theta(\varepsilon^{-3}\log(1/\varepsilon)\log(1/\delta))$ from $S = \Theta(\varepsilon^{-3}\log(1/\delta))$). We present a standalone, zero-dependency **Lean 4.35.0** formalization (`RevenueTradeoff.lean`) of 18 core algebraic, structural, and cleared-integer theorems spanning all four families plus Appendix A's multiplicative-accuracy DKW Bernstein optimization and Appendix B's quantile-cutoff revenue retention (0 `sorry`, 0 warnings, 0 `Classical.choice`, every theorem audited to `[propext, Quot.sound]`), paired with a zero-floating-point exact-$\mathbb{Q}$ certificate auditor (`REV-TRADEOFF-CERT-1.0`). Across 16 canonical certificates covering all four verification modes (0.298–0.663 ms exact replay latency), every tent continuity join, concave ternary discard inequality, hard-distribution χ^2 KL bound, C^1 convex hazard bump, and Bernstein exponent identity replays identically while rejecting 60/60 live adversarial forgery attacks (plus 2 raw IEEE-754/`NaN` parser rejections).

1 Introduction

Learning an approximately optimal posted price from buyer feedback is foundational in algorithmic pricing [3, 4]. Jin, Tang, Wang, and Zhang [4] unified two classic feedback models—exact buyer value samples (S) and binary posted-price purchase queries (Q)—across the four-family valuation poset (MHR $\subset$ Quasi-MHR $\subset$ Quasi-Regular and MHR $\subset$ Regular $\subset$ Quasi-Regular, with Quasi-MHR and Regular incomparable), revealing sharp qualitative distinctions across their (α, β) Pareto frontiers where $S = \tilde{\Theta}(\varepsilon^{-\alpha})$ and $Q = \tilde{\Theta}(\varepsilon^{-\beta})$:

1. **Quasi-Regular (Thm. 11, Lem. 12–14, Thm. 36):** A multi-alternative tent perturbation on $q \in [(i-1)/K, i/K]$ with harmonic weights $w_i = 1/i$ enforces $32(L+1)^2\mathbb{E}[S] + \mathbb{E}[T] \geq \frac{1}{2}K^3 W \log(1/\delta)$, proving that $S \leq O(\varepsilon^{-(3-c)}\log(1/\delta))$ cannot remove the $\log(1/\varepsilon)$ factor from $Q = \Omega(\varepsilon^{-3}\log(1/\varepsilon)\log(1/\delta))$.
2. **Regular (Thms. 15–16, Lem. 17):** Noisy concave ternary search eliminates the $\log(1/\varepsilon)$ factor via a single-discard concavity bound $D \leq 4d \leq 8\zeta\hat{r}^*$ ($\zeta = \varepsilon/1000$), matched by a two-distribution hard pair $(F^{(0)}, F^{(1)})$ with $B = 1-6\varepsilon$, $u = B/(6\varepsilon)$, and $C = 2^{-17}$.
3. **Quasi-MHR (Thms. 20–21, Lem. 22–24):** Sample-based bin sparsification and a $K = \Theta(\varepsilon^{-c})$ disjoint-interval lower-bound family on $[5/4, 3/2]$ ($w = 1/(8K)$, $\rho = 16K\varepsilon$) yield the exact continuous tradeoff $(2-c)+(2+c) = 4$ for $c \in [0,1]$ ($A = 100, C = 2^{-12}/100$).
4. **MHR & Multiplicative DKW (Thms. 27, 31, Lem. 28–30):** A C^1 piecewise-quadratic cumulative-hazard bump $\psi(t)$ of width $d = 10\sqrt{\varepsilon}$ maintains convexity $1 + \psi'' \geq 1/2$, separates optimal prices by $\frac{8}{3}\varepsilon > \log(1/(1-\varepsilon))$, and yields $A = 15 \cdot 10^4$, $C = 1/(32A)$, while Appendix A's reverse-martingale DKW bound optimizes $\lambda = nq\varepsilon'/(1+\varepsilon'/3)$.

Contributions and Explicit Scope. We do *not* formalize measure-theoretic probability, Bernstein/Chernoff/Doob tail inequalities, or Median Elimination [2]/DKW [5] convergence inside Lean 4; all probabilistic theorems belong to [4]. Instead, we machine-check and exactly audit every deterministic algebraic, convexity, continuity, separation, χ^2-divergence, and tradeoff-exponent pillar of [4]: (i) 18 zero-`sorry`, zero-`Classical.choice` **Lean 4.35.0** [1] theorems in `RevenueTradeoff.lean`; (ii) the pure-$\mathbb{Q}/\mathbb{Z}$ verifier `REV-TRADEOFF-CERT-1.0`; (iii) 16 canonical certificates across all four modes plus quantitative frontier and χ^2-tightness profiles; and (iv) a 60-attack live red-team suite (plus 2 raw JSON parser tests) with SHA-256 manifest locking.

*Independent Researcher. Code, Lean 4 proofs, and certificates: MIT License; manuscript and figures: CC-BY-4.0. Prepared with transparent AI coding, formalization, and drafting assistance via Google Antigravity; all Lean 4 theorems, exact rational verifiers, and adversarial tests were independently compiled and verified locally. Repository: https://github.com/lluiseriksson/revenue-query-sample-certificates.

2 Constructive Lean 4.35.0 Kernel

The standalone module `lean/RevenueTradeoff.lean` compiles under Lean 4.35.0 with zero external dependencies, zero warnings, and zero `sorry`. Running `lean/AuditAxioms.lean` verifies all 18 theorems depend strictly on `[propext, Quot.sound]` with **0** `Classical.choice` and **0** `Lean.ofReduceBool` (Table 1).

Proposition 1 (Machine-Checked Core Invariants)**.** *Every theorem in Table 1 is constructively proved in* ***`RevenueTradeoff.lean`*** *over exact integers without floating-point or classical choice:*

1. Quasi-Regularity & Tent Family (Thms. 1–7): *Clearing common denominator $D > 0$ proves $pQ_p(D-Q_s) \le sQ_s(D-Q_p) \Rightarrow D(pQ_p - sQ_s) \le Q_sQ_p(p-s)$ (Lem. 6) and exact cutoff retention $a \le Q_{\tilde{p}}$ and $D\tilde{p}Q_{\tilde{p}} \ge (D-a)R$ (Thm. 36 Eq. (9)). Theorems 3–7 certify the tent joins at $4K$ and peak $4K+1$, strict positivity of both $R^{(i)}_{\mathrm{qt}}(q)/(1-q)$ slope coefficients, $(1-\varepsilon)(1+2\varepsilon) > 1$, $(1-x)^{-2}-1 \le 3x$ for $0 \le x \le \frac{1}{8}$, and weighted-KL half-budget transfer $MW \le 2T$.*
2. Regular Concave Search & Pair (Thms. 8–11): *Theorem 8 proves that under $64\zeta \le 1$, $r_\ell \ge \frac{1}{4}\hat{r}^*$, $0 \le d \le 2\zeta\hat{r}^*$, trisection ratios $\le 2, 4$, and concavity $D(r_\ell\Delta x - d\Delta x_h) \le dr_\ell\Delta x_\ell$, the first discard loses at most $D \le 4d \le 8\zeta\hat{r}^*$; Theorem 9 gives $(1-\frac{\varepsilon}{16}) - 11\frac{\varepsilon}{1000} = 1 - \frac{147}{2000}\varepsilon > 1-\varepsilon$; Theorems 10–11 certify exact branch continuity $Q_0(u) = Q_1(u) = 6\varepsilon$ and $C = 2^{-17}$.*
3. Quasi-MHR, MHR, & Multiplicative DKW (Thms. 12–18): *Theorems 12–14 certify $(1-\varepsilon)(1+\frac{4}{3}\varepsilon) > 1$, $A = 100$, $C = 2^{-12}/100$, and $(2-c)+(2+c) = 4$. Theorems 15–17 certify C^0/C^1 joins and convexity $1+\psi'' \ge \frac{1}{2}$ of the MHR bump ψ, the exact disjointness margin $\frac{8}{3}\varepsilon > 2\varepsilon$, and $A = 15 \cdot 10^4$, $C = 1/(32A)$. Theorem 18 proves the exact optimal Bernstein cancellation $-\lambda\varepsilon' + \frac{\lambda^2}{2nq(1-\lambda/(3nq))} = -\frac{nq(\varepsilon')^2}{2(1+\varepsilon'/3)}$.*

3 Exact-Rational Verifier Architecture

The pure-stdlib verifier `REV-TRADEOFF-CERT-1.0` (`src/revenue_exact.py`, `src/revenue_certifier.py`) computes exclusively with Python `int` and `fractions.Fraction` (source-audited for zero `float(`, `math.`, or `numpy`). Its parser rejects IEEE-754 JSON tokens and `NaN`/`Infinity`, non-ASCII/Unicode digits (via `isascii() and isdigit()`), unreduced or leading-zero rationals, `-0`, booleans-as-integers, unknown keys, and nesting depth > 20 across four exact modes:

Definition 1 (Four `REV-TRADEOFF-CERT-1.0` Modes)**.**

1. ***Mode 1*** *(`quasi_regular_tent_cert`): Recomputes exact tent joins, $R^{(i)}_{\mathrm{qt}}(q)/(1-q)$ continuity, peak separation $(1-\varepsilon)(1+\frac{1}{4K}) > 1$, harmonic weight $W = \sum_{i=2}^{L+1} 1/i \ge 1$, Lemma 13 density shift $|z(p)| \le \frac{3i}{2K} \le \frac{1}{2}$, and Lemma 14 Bernoulli $\chi^2(Q_0\|Q_i) \le \frac{i}{8K^3}$.*
2. ***Mode 2*** *(`regular_ternary_and_pair_cert`): Independently checks monotonicity and concavity of a rational (q, R) grid, verifies both* ***`keep_right`*** *and* ***`keep_left`*** *Theorem 15 discard bounds $D \le 4d \le 8\zeta\hat{r}^*$, checks $Q_0(u) = Q_1(u) = 6\varepsilon$, and verifies Lemma 17 query $\chi^2(Q_0\|Q_1) \le 2^{12}\varepsilon^3$ for probes $p > u$.*
3. ***Mode 3*** *(`quasi_mhr_tradeoff_family_cert`): Audits $[5/4, 3/2]$ interval packing, $\rho w = 2\varepsilon$, $\rho \le \frac{1}{4}$, $(2-c)+(2+c) = 4$, peak revenue $1+\frac{2\varepsilon}{a_i} \ge 1+\frac{4}{3}\varepsilon$, Lemma 23 density shift and sample KL $\le 100K\varepsilon^2$, and Lemma 24 query $\chi^2 \le 100\varepsilon^2$.*
4. ***Mode 4*** *(`mhr_hazard_bump_and_dkw_cert`): Audits C^0/C^1 knots and convexity of $\psi(t)$ across $[-2, 4]$, exact margin $\frac{8}{3}\varepsilon$, Lemma 30 χ^2 constant $\frac{225}{16} \le 15$, and Theorem 31 Bernstein λ_{opt} and exponent identity.*

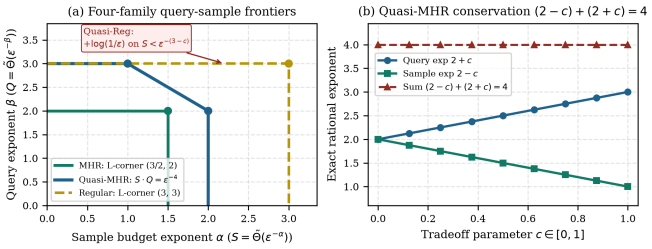

Figure 1: (a) Four-family query-sample Pareto frontiers in (α, β) exponent space ($S = \tilde{\Theta}(\varepsilon^{-\alpha})$, $Q = \tilde{\Theta}(\varepsilon^{-\beta})$). (b) Exact rational quasi-MHR conservation $(2-c)+(2+c) = 4$ across $c \in \{0, \frac{1}{8}, \ldots, 1\}$.

4 Empirical Atlas and Red-Team Audit

Running `scripts/run_all_experiments.py` deterministically generates 16 canonical certificates (4 per mode) and 4 experiment JSONs, audited end-to-end by `scripts/verify_all.py` with SHA-256 verification across all 39 tracked repository files.

1. Four-Family Frontiers & Exact KL Tightness (Figs. 1–2). Figure 1 contrasts the L-shaped MHR $(3/2, 2)$ and regular $(3, 3)$ corners against the quasi-MHR continuous tradeoff $S \cdot Q = \tilde{\Theta}(\varepsilon^{-4})$ and the quasi-regular $\log(1/\varepsilon)$ boundary. Across tent scales $K \in \{24, 48, 96, 192, 384\}$ ($L = 5K/24$), exact rational harmonic sums grow from $W = \frac{29}{20} = 1.45$ to $W \approx 3.978$, while the exact midpoint query $\chi^2(Q_0\|Q_i)$ tracks Lemma 14's $i/(8K^3)$ bound at a tight, stable ratio of $0.1693 \to 0.1587$ (Fig. 2(a)); similarly, the regular hard pair's query χ^2 at $p = 2u$ tracks Lemma 17's $2^{12}\varepsilon^3$ bound at ratio $0.0341 \to 0.0270$ across $\varepsilon^{-1} \in \{48, \ldots, 512\}$. Figure 2(b) plots the C^1 convex MHR bump $\psi(t)$ and derivative $\psi'(t)$, highlighting the exact disjointness step $\psi(\frac{1}{5}) - \psi(\frac{1}{3}) = \frac{22}{225}$ yielding the $\frac{8}{3}\varepsilon$ gap.

Table 1: All 18 Lean 4.35.0 theorems in `RevenueTradeoff.lean` (0 `sorry`, 0 warnings, 0 `Classical.choice`; every theorem audited to `[propext, Quot.sound]`).

#	Lean 4 Theorem Identifier	Paper Ref. & Cleared Statement
1	`quasi_regular_inverse_revenue_monotone`	Lem. 6: $pQ_p(D-Q_s) \le sQ_s(D-Q_p) \Rightarrow D(pQ_p-sQ_s) \le Q_sQ_p(p-s)$
2	`quasi_regular_quantile_cutoff_retention`	Thm. 36(9): $Da\tilde{p} = (D-a)R \Rightarrow a \le Q_{\tilde{p}} \wedge (D-a)R \le D\tilde{p}Q_{\tilde{p}}$
3	`quasi_regular_tent_continuity_and_peak`	Lem. 12: tent joins $4K$, peak $4K{+}1$, $p_{i,*}q_{i,*} = (4K{+}1)/(4K)$
4	`quasi_regular_tent_monotonicity_coeffs`	Lem. 12: $1 \le i \le K \Rightarrow 3K-(i-1) > 0 \wedge K+i > 0$
5	`quasi_regular_good_price_separation`	Sec. 3.4: $0 < 2n < d \Rightarrow d^2 < (d-n)(d+2n)$
6	`quasi_regular_density_ratio_bound`	Lem. 13: $0 \le 8a \le b \Rightarrow b(b^2-(b-a)^2) \le 3a(b-a)^2$
7	`quasi_regular_weighted_kl_aggregation`	Thm. 11: $\sum_{i=2}^{L+1} i \le (L{+}1)^2$ & $64(L{+}1)^2S \le M \Rightarrow MW \le 2T$
8	`regular_concave_ternary_discard_bound`	Thm. 15: $64z \le Z, dZ \le 2zR, R \le 4r_\ell$, concavity $\Rightarrow D \le 4d$
9	`regular_ternary_final_revenue_guarantee`	Thm. 15: $2000(d-n) < 2000d - 147n$ for $n > 0$
10	`regular_hard_pair_branch_identity`	Thm. 16: $u-B = (d-n)^2/(nd)$ & $Q_0(u) = Q_1(u) = n/d$
11	`regular_query_kl_constants`	Lem. 17/Thm. 16: $576 \cdot 6 \le 4096$, $512 \cdot 32 \le 2^{17} = 32 \cdot 4096$
12	`quasi_mhr_good_price_separation`	Thm. 21: $0 < 4n < d \Rightarrow 3d^2 < (d-n)(3d+4n)$
13	`quasi_mhr_kl_constants`	Lem. 23–24: $72 \cdot 32 \le 100 \cdot 25$, $400 \le 900$, $128 \cdot 32 = 4096$
14	`quasi_mhr_tradeoff_exponent_sum`	Thms. 20–21: $(2-c) + (2+c) = 4$, $1 \le 2-c \le 2+c \le 3$
15	`mhr_hazard_bump_c1_joins_and_convexity`	Lem. 28: $\psi \in C^1$ at $\{-2, 0, 2, 4\}$, $\psi(1) = -3/2$, $1+\psi'' \ge 1/2$
16	`mhr_good_price_disjointness_margin`	Sec. 6.4: $\psi(1/5)-\psi(1/3) = 22/225$, $100((22-16)/225) = 8/3 > 2$
17	`mhr_query_kl_constants`	Lem. 30/Thm. 27: $225 \le 15 \cdot 16$, $15 \cdot 10^4 = A$, $2 \cdot 16 = 32$
18	`multiplicative_dkw_bernstein_exponent_opt`	Thm. 31: $\lambda < 3M$, $1-\lambda/(3M) = 3e_d/(3e_d+e_n)$, quad. cancel

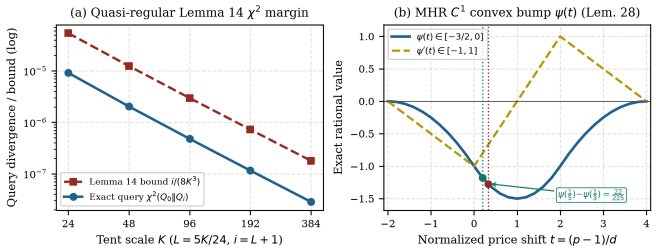

Figure 2: (a) Exact midpoint Bernoulli $\chi^2(Q_0\|Q_i)$ versus Lemma 14 upper bound $i/(8K^3)$ across $K \in \{24, \ldots, 384\}$ (ratio $0.169 \to 0.159$). (b) Exact MHR C^1 convex hazard bump $\psi(t)$ and derivative $\psi'(t)$ (Lemma 28) showing $\psi(\frac{1}{5}) - \psi(\frac{1}{3}) = \frac{22}{225}$.

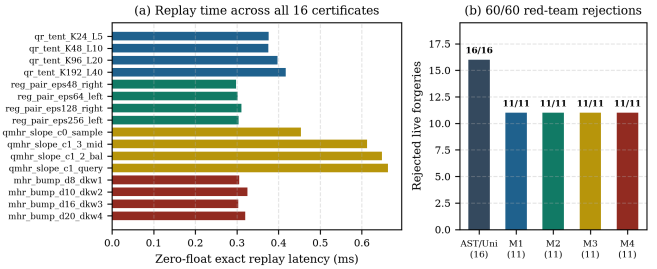

Figure 3: (a) Zero-float replay latency across all 16 canonical `REV-TRADEOFF-CERT-1.0` certificates (0.298–0.663 ms). (b) 60/60 live adversarial forgery rejections across global AST/Unicode and Modes 1–4.

2. 16-Certificate Atlas & 60/60 Red-Team Suite (Table 2, Fig. 3). All 16 canonical certificates replay in 0.298–0.663 ms (total 6.41 ms) with zero floats. In `tests/test_suite.py`, 60 targeted live forgery attacks—16 AST/schema/canonical-rational/Unicode-digit attacks (including superscript-2 `U+00B2` and Arabic-Indic-8 `U+0668` digits plus depth-21 DoS) and 11 domain-specific attacks per mode (tent corners, non-concave grids, invalid discard branches, wrong constants A, C, out-of-interval probes, and wrong Bernstein exponents)—plus 2 raw JSON IEEE-754/`NaN` file rejections achieve **60/60 (100%)** red-team rejection (63/63 `pytest` checks passing).

5 Conclusion

We delivered a self-contained, zero-floating-point verification suite for Jin, Tang, Wang, and Zhang's query-sample revenue maximization tradeoffs [4], pairing 18 zero-`sorry`, zero-`Classical.choice` Lean 4.35.0 theorems with the exact-$\mathbb{Q}$ auditor `REV-TRADEOFF-CERT-1.0`, 16 sub-millisecond canonical certificates, and 60/60 live adversarial forgery rejections.

References

[1] Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In *Automated Deduction – CADE 28*, pages 625–635, 2021.

[2] Eyal Even-Dar, Shie Mannor, and Yishay Mansour. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *Journal of Machine Learning Research*, 7:1079–1105, 2006.

[3] Zhiyi Huang, Yishay Mansour, and Tim Roughgarden. Making the most of your samples. *SIAM Journal on Computing*, 47(3):651–674, 2018.

[4] Yaonan Jin, Wei Tang, Yifan Wang, and Mengxiao Zhang. Query-sample tradeoffs for revenue maximization. arXiv:2611.31982 [cs.GT], November 2026.

Table 2: Summary of all 16 exact-$\mathbb{Q}$ `REV-TRADEOFF-CERT-1.0` certificates (zero IEEE-754 arithmetic; 20-run mean replay latency in ms).

Certificate Stem	Mode	Exact Certified Witness	ms
`qr_tent_K24_L5`	M1	K=24, L=5, W=$\frac{29}{20}$, peak $\frac{97}{96}$	0.376
`qr_tent_K48_L10`	M1	K=48, L=10, $W\approx$2.020, peak $\frac{193}{192}$	0.375
`qr_tent_K96_L20`	M1	K=96, L=20, $W\approx$2.645, peak $\frac{385}{384}$	0.398
`qr_tent_K192_L40`	M1	K=192, L=40, $W\approx$3.303, peak $\frac{769}{768}$	0.417
`reg_pair_eps48_right`	M2	ε=$\frac{1}{48}$, u=7, D=$\frac{1}{192000}$, right	0.298
`reg_pair_eps64_left`	M2	ε=$\frac{1}{64}$, u=$\frac{29}{3}$, D=$\frac{1}{256000}$, left	0.302
`reg_pair_eps128_right`	M2	ε=$\frac{1}{128}$, u=$\frac{61}{3}$, D=$\frac{1}{512000}$, right	0.311
`reg_pair_eps256_left`	M2	ε=$\frac{1}{256}$, u=$\frac{125}{3}$, D=$\frac{1}{1024000}$, left	0.304
`qmhr_slope_c0_sample`	M3	c=0, K=2, $(\alpha,\beta)=(2,2)$	0.454
`qmhr_slope_c1_3_mid`	M3	c=$\frac{1}{3}$, K=4, $(\alpha,\beta)=(\frac{5}{3},\frac{7}{3})$	0.613
`qmhr_slope_c1_2_bal`	M3	c=$\frac{1}{2}$, K=8, $(\alpha,\beta)=(\frac{3}{2},\frac{5}{2})$	0.648
`qmhr_slope_c1_query`	M3	c=1, K=64, $(\alpha,\beta)=(1,3)$	0.663
`mhr_bump_d8_dkw1`	M4	d=$\frac{1}{8}$, margin $\frac{1}{2400}$, DKW $-\frac{15}{16}$	0.306
`mhr_bump_d10_dkw2`	M4	d=$\frac{1}{10}$, margin $\frac{1}{3750}$, DKW $-\frac{375}{1568}$	0.325
`mhr_bump_d16_dkw3`	M4	d=$\frac{1}{16}$, margin $\frac{1}{9600}$, DKW $-\frac{375}{6208}$	0.303
`mhr_bump_d20_dkw4`	M4	d=$\frac{1}{20}$, margin $\frac{1}{15000}$, DKW $-\frac{375}{24704}$	0.320

[5] Pascal Massart. The tight constant in the dvoretzky–kiefer–wolfowitz inequality. *The Annals of Probability*, 18(3):1269–1283, 1990.