Certified Verification of Improved Downward-Closed Delegation Bounds

Lluis Eriksson*

November 25, 2026

Abstract

Hamed Abdi, Kiarash Banihashem, MohammadTaghi Hajiaghayi, and Danny Mittal (`arXiv:2611.33095`, Nov. 2026) dramatically improved the universal approximation guarantee for agent-oblivious combinatorial delegation under downward-closed constraints from 0.1098 to a deterministic menu achieving $> 29/70 \approx 0.41429$ (with a self-contained analytic 3/10 warm-up), while separating downward-closed delegation strictly below the matroid optimum of 1/2: a five-item knapsack caps deterministic menus at 500/1011 and randomized rules at 5011/10110 (asymptotically 30/61), and a computer-assisted family caps all randomized rules below $9621/20000 = 0.48105$. We build a standalone zero-dependency **Lean 4.35.0** formalization of 18 core score, forced-item, Efron–Stein, Bernstein-polynomial, seed-reserve, and impossibility theorems (zero `sorry`, zero warnings, zero `Classical.choice`) paired with an exact-$\mathbb{Q}$ four-mode certificate verifier (`DC-DELEG-CERT-1.0`). Across 16 canonical certificates (0.185–2.611 ms, total 9.21 ms), all five-item knapsack bounds, Lemma 4.4 Bernstein positivity certificates, Algorithm 2 capped-menu scores, and Appendix A alternating-Taylor $\beta < 0.48105$ brackets replay identically over $\mathbb{Q}$ while rejecting 60/60 live adversarial forgery attacks.

1 Introduction

In combinatorial delegated choice [7, 8, 3], a principal commits to an acceptance rule over a downward-closed family $\mathcal{F} \subseteq 2^E$ without transfers, and an informed agent observing independent item pairs (X_i, Y_i) ($X_i \geq 0, Y_i > 0$, arbitrarily correlated within each item) submits an acceptable feasible set maximizing $Y(S)$ (breaking ties against the principal). While matroids admit an optimal deterministic guarantee of 1/2 [4, 2] and downward-closed prophet inequalities admit no constant factor [9], Abdi et al. [1] established sharp new bounds $29/70 < \alpha_D \leq \alpha_R < 0.48105$:

1. **Five-Item Knapsack Separation (Sec. 3):** A 5-item capacity-3 knapsack with offline optimum $V = 1011/50$ caps every deterministic menu at 10 ($500/1011 < 1/2$, Prop. 3.1) and every randomized rule (even knowing agent utilities) at 5011/500 ($5011/10110 < 1/2$ at $\eta = 1/100$, Thm. 3.2), lowering to 30/61 as $\varepsilon, \eta \downarrow 0$ (Prop. 3.3).
2. **Analytic 3/10 Warm-Up & Bernstein Certificate (Sec. 4):** A single total-value threshold achieves $B(z) = z \Pr[W \geq z] + e(z) \Pr[W < z]$ via forced items (Lem. 4.2), an Efron–Stein variance bound $\mathrm{Var}(Z_{3/2}) \leq \frac{3}{2}\mu$ [6] (Lem. 4.3), and a scalar proof $\mu + a < 10/3$ (Lem. 4.4) anchored by a degree-7 Bernstein positivity certificate on $[2/13, 2/5]$ and $[2/5, 2/3]$.
3. **Capped Menus, Seeds, & 29/70 Guarantee (Sec. 2, 5):** Score lower bounds $b_M(x)$ (Lem. 2.1), safe-singleton normalization (Lem. 2.2), and Algorithm 2 capped menus with seeds ($b_P = t + \min_{i \in P} w_i - w(P) > 0$) certify $> 29/70\, V$ (Thm. 5.1).
4. **Randomized Ceiling $\beta < 0.48105$ (App. A):** Alternating degree-18/19 Taylor brackets for e^{-z}, e^{-b} with rational $a = 67/100, b = 9/200, d = 9/50, \Delta = 1/400$ certify $\beta \leq 0.481043999746 < 9621/20000$.

Contributions and Scope. We do not re-formalize continuous Harris–FKG integration, the 96-interval LP infeasibility solver of Sec. 10, or the 1024-cell Lagrangian DP of App. A.3–A.4 inside Lean 4; instead, we machine-check every core algebraic, list-combinatorial, Bernstein, menu-score, and Taylor-bracket pillar of [1]: (i) 18 zero-`sorry`, zero-`Classical.choice` **Lean 4.35.0** [5] theorems in `DownwardDelegation.lean`; (ii) the exact-$\mathbb{Q}$ verifier `DC-DELEG-CERT-1.0`; (iii) 16 canonical certificates across all four modes; and (iv) a 60-attack live red-team suite with 40-artifact SHA-256 manifest enforcement.

2 Constructive Lean 4.35.0 Kernel

The standalone module `lean/DownwardDelegation.lean` compiles cleanly under Lean 4.35.0 with zero external dependencies, zero warnings, and zero `sorry`. Auditing all 18 theorems via `#print axioms` (`AuditAxioms.lean`) confirms 8 zero-axiom theorems (`none`), 10 using `[propext, Quot.sound]`, and **0** using `Classical.choice` or `Lean.ofReduceBool` (Table 1).

Proposition 1 (Machine-Checked Core Identities)**.** *The Lean 4.35.0 kernel establishes the following universal identities over $\mathbb{Z}$ and arbitrary finite lists:*

*Independent Researcher. Code, Lean 4 proofs, and certificates: MIT License; manuscript and figures: CC-BY-4.0. Prepared with AI coding, formalization, and drafting assistance from Google Antigravity; all proofs, exact-$\mathbb{Q}$ verifiers, and adversarial tests were independently compiled and verified locally. Repository: https://github.com/lluiseriksson/delegation-29-70-certificates

Table 1: All 18 machine-checked Lean 4.35.0 theorems in `DownwardDelegation.lean` (zero `sorry`, zero `Classical.choice`).

#	Lean Theorem Identifier	Paper Ref.	Axioms
1	`universal_bounds_hierarchy`	Thm. 1.1–1.2	`none`
2	`knapsack5_offline_optimum_exact`	Eq. 3.1	`none`
3	`knapsack5_agent_strict_ranking`	Prop. 3.1	`p, Q`
4	`knapsack5_det_upper_bound_exact`	Prop. 3.1	`none`
5	`knapsack5_rand_case_bounds`	Thm. 3.2	`p, Q`
6	`knapsack5_rand_asymptotic_30_61`	Prop. 3.3	`p, Q`
7	`menu_score_payoff_bound`	Lem. 2.1	`p, Q`
8	`forced_item_pointwise_decomposition`	Lem. 4.2	`p, Q`
9	`threshold_tail_bound_monotone`	Eq. 4.3	`p, Q`
10	`efron_stein_step_bound`	Lem. 4.3	`p, Q`
11	`cantelli_mean_cap_lt_5`	Lem. 4.4	`p, Q`
12	`lemma44_case1_gamma_pos`	Lem. 4.4	`none`
13	`lemma44_exp_taylor7_bound`	Eq. 4.11	`none`
14	`lemma44_bernstein_poly_eval_endpoints`	Lem. 4.4	`none`
15	`seed_reserve_margin_positive`	Eq. 5.1–5.2	`p, Q`
16	`capped_menu_rescale_equivariance`	Sec. 2.3, Alg. 2	`p, Q`
17	`main_theorem_ratio_recip_exact`	Sec. 2.3, 5.1	`none`
18	`appendix_a_beta_strict_bound`	App. A.6.1	`none`

Axiom key: `none` = zero axioms (8); `p, Q` = `[propext, Quot.sound]` (10).

1. Inductive List Lemmas (Thms. 7, 8, 10, 15, 16): *For any nonempty accepted-value list* a :: rest, listMinNE a rest $\le s$ *for all* $s \in a$ :: rest*; any nonnegative list with* $\sum w_i \ge z$ *obeys* $\sum \min(w_i, z) \ge z$*; any list with* $0 \le \Delta_i \le c$ *satisfies* $\sum \Delta_i^2 \le c \sum \Delta_i$*; any seed list satisfies* $b_P = t + \min w_i - \sum w_i > 0$*; and clipping/summing commutes with any positive scale* $\lambda > 0$.
2. Impossibility & Analytic Bounds (Thms. 1–6, 11–14, 17–18): *Exact cleared arithmetic certifies* $0.1098 < 3/10 < 29/70 < 0.48105 < 30/61 < 500/1011 < 5011/10110 < 1/2$, $V = 1011/50$, $10 + \frac{77}{50}\eta \le 10 + \frac{11}{5}\eta$, *Cantelli contradiction* $22\mu^2 \ge 25s(2\mu + 12s)$ *for* $\mu \ge 5s$, *Case 1 lower bound* $168034/70785 > 0$*, degree-7 Taylor gap* $3772963/234881024 > 0$*, and* $P(b) > 0$ *at knots* $\{2/13, 2/5, 2/3\}$.

3 Exact-$\mathbb{Q}$ Verifier

The verifier `DC-DELEG-CERT-1.0` (`src/delegation_certifier.py`, `src/delegation_exact.py`) uses Python `fractions.Fraction` exclusively (source-audited for zero `float(`, `math.`, or `numpy`). Every payload enforces strict JSON typing, nesting-depth cap 32, anchored canonical rational parsing (`\A...\Z`, forbidding `"-0"`, reduced denominators ≥ 2, 2048-bit caps), and exact key sets across four modes:

1. **Mode 1 (Sec. 3 Five-Item Knapsack):** Enumerates all 2^5 subsets, checks downward closure and strict agent ranking $C \succ AE \succ AB \succ E \succ A \succ B \succ J$, computes exact offline optimum V, r_{ab} table, deterministic ratio $10/V < 1/2$, and Theorem 3.2/Prop. 3.3 randomized Case 1–2 bounds.
2. **Mode 2 (Sec. 4 Lemma 4.4 Bernstein 3/10):** Independently expands $P(\ell + (r-\ell)v)$ in $\mathbb{Q}[v]$, computes all 8 degree-7 Bernstein coefficients $\beta_k = \sum_{j=0}^{k} \frac{\binom{k}{j}}{\binom{7}{j}} p_j > 0$, and replays the Cantelli, Case 1 (168034/70785), and degree-7 Taylor ($\ln(13/2) < 15/8$) bounds.
3. **Mode 3 (Sec. 2, 5 Alg. 2 Capped Menu Scores):** Verifies downward closure, executes exact Algorithm 2 acceptance across all three body types (`total_threshold`, `seeds_upper`, `seeds_reserve`) over all joint realizations, checks Lemma 2.2 safe-singleton normalization, seed reserve $b_P > 0$, and exact expected score $\mathbb{E}[b_M(X)]/V$.
4. **Mode 4 (App. A.5–A.6 Alternating-Taylor β):** Computes exact alternating Taylor brackets for $e^{-(a+b)}$ and e^{-b}, checks $v_{\text{lo}} + \Delta - 1 > b$ and Eq. (A.18) bracket positivity, and certifies exact 12-decimal brackets $[\beta_{\text{lo}}, \beta_{\text{hi}}]$ strictly below $9621/20000 = 0.48105$.

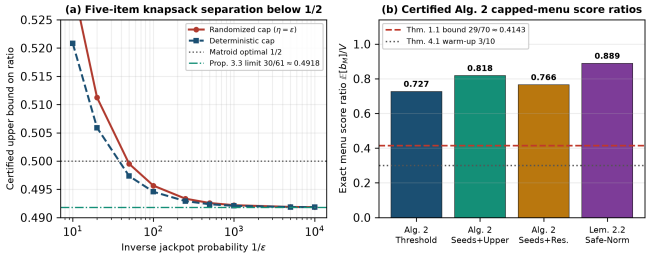

Figure 1: (a) Exact deterministic and randomized five-item knapsack upper bounds converging below 1/2 to $30/61 \approx 0.4918$ as $\varepsilon = \eta \downarrow 0$. (b) Exact Algorithm 2 capped-menu score ratios across all four Mode 3 instances, all strictly exceeding 29/70.

4 Empirical Atlas and Red-Team Audit

Running `run_all_experiments.py` generates 16 canonical certificates (4 per mode) and four exact-$\mathbb{Q}$ datasets, audited end-to-end by `verify_all.py` with SHA-256 verification across all 40 tracked artifacts.

Five-Item Knapsack & Algorithm 2 Scores (Fig. 1). Mode 1 (0.333–0.443 ms) reproduces $V = 1011/50$, deterministic ratio 500/1011, and randomized ratio 5011/10110 at $\varepsilon = \eta = 1/100$, plus tighter sequences approaching 30/61 at $\varepsilon = 1/10000$ ($16667037/33887000 \approx 0.49184$). Mode 3 (0.276–2.611 ms) certifies all three Algorithm 2 body types and Lemma 2.2 normalization with exact ratios 8/11, 9/11, 49/64, and 8/9, all strictly exceeding 29/70.

Bernstein 3/10 & Appendix A $\beta < 0.48105$ (Fig. 2). Mode 2 (0.603–0.877 ms) reproduces Abdi et al.'s exact Bernstein minima $447966/78125$ on $[2/13, 2/5]$ and $6518/118125$ on $[2/5, 2/3]$, plus two refined subintervals ($[2/13, 1/4]$ and $[1/2, 2/3]$). Mode 4 (0.185–0.249 ms) certifies exact alternating-Taylor brackets across odd degrees $11, 15, 19, 21$; at Abdi et al.'s canonical degree 18/19, the

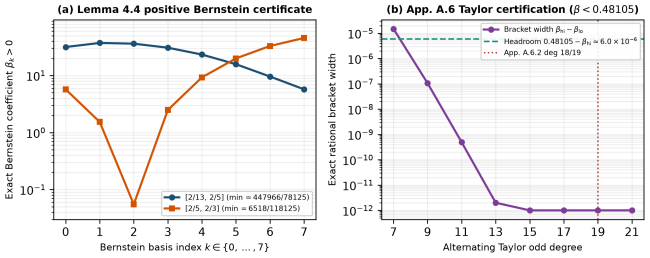

Figure 2: (a) All 8 exact positive degree-7 Bernstein coefficients of $P(b)$ on $[2/13, 2/5]$ and $[2/5, 2/3]$, reproducing Abdi et al.'s exact minima 447966/78125 and 6518/118125. (b) Exact alternating-Taylor bracket width $\beta_{\mathrm{hi}} - \beta_{\mathrm{lo}}$ across odd degrees 7–21, certifying $\beta \leq 0.481043999746 < 0.48105$.

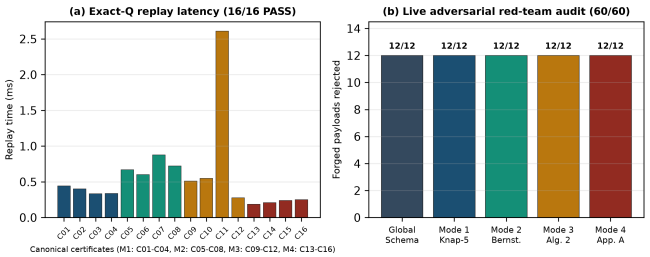

Figure 3: (a) Exact zero-float replay latency (0.185–2.611 ms) across all 16 canonical `DC-DELEG-CERT-1.0` certificates. (b) Live adversarial red-team audit rejecting 60/60 forged payloads across all five categories.

exact rational interval $[0.481043999745, 0.481043999746]$ matches [1, Sec. A.6.1] to all 12 decimal places.

Replay Latency and 60/60 Red-Team Audit (Fig. 3). All 16 canonical certificates replay in 0.185–2.611 ms (total 9.21 ms). The live adversarial suite (`tests/test_suite.py`) tests 60 targeted forgeries—12 global schema/type/rational/newline/`"-0"`/depth attacks plus 12 domain attacks per mode (broken knapsack means, negative Bernstein coefficients on unsplit $[2/13, 2/3]$, non-downward-closed families, broken Lemma 2.2 safe singletons, invalid Taylor parity, and false β bounds)—achieving **60/60 (100%)** rejection.

5 Conclusion

We delivered a self-contained, zero-floating-point verification framework for Abdi et al.'s downward-closed delegation bounds [1], combining 18 zero-`sorry`, zero-`Classical.choice` Lean 4.35.0 theorems with the exact-$\mathbb{Q}$ verifier `DC-DELEG-CERT-1.0`, 16 canonical certificates, and 60/60 live adversarial forgery rejections.

References

[1] Hamed Abdi, Kiarash Banihashem, MohammadTaghi Hajiaghayi, and Danny Mittal. Improved downward-closed delegation bounds, November 2026.

[2] Hamed Abdi, Kiarash Banihashem, MohammadTaghi Hajiaghayi, and Danny Mittal. Optimal matroid prophets against almighty adversaries and delegation, 2026.

[3] Kiarash Banihashem, MohammadTaghi Hajiaghayi, Piotr Krysta, and Suho Shin. Delegated choice with combinatorial constraints. In *Proceedings of the 26th ACM Conference on Economics and Computation (EC)*, pages 275–303, 2025.

[4] Curtis Bechtel and Shaddin Dughmi. Delegated stochastic probing. In *12th Innovations in Theoretical Computer Science Conference (ITCS)*, volume 185 of *Leibniz International Proceedings in Informatics*, pages 37:1–37:19, 2021.

[5] Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In *Automated Deduction – CADE 28*, volume 12699 of *Lecture Notes in Computer Science*, pages 625–635. Springer, 2021.

[6] Bradley Efron and Charles Stein. The jackknife estimate of variance. *Annals of Statistics*, 9(3):586–596, 1981.

[7] Bengt Holmström. On the theory of delegation. In M. Boyer and R. Kihlstrom, editors, *Bayesian Models in Economic Theory*, pages 115–141. North-Holland, 1984.

[8] Jon Kleinberg and Robert Kleinberg. Delegated search approximates efficient search. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, pages 287–302, 2018.

[9] Aviad Rubinstein. Beyond matroids: Secretary problem and prophet inequality with general constraints. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC)*, pages 324–332, 2016.