Machine-Checked Exact-Rational Menu-Score, Knapsack Impossibility, and Bernstein Threshold-Guarantee Certificates for Downward-Closed Delegated Choice

Lluis Eriksson*

November 25, 2026

Abstract

Hamed Abdi, Kiarash Banihashem, MohammadTaghi Hajiaghayi, and Danny Mittal (`arXiv:2611.33095`, November 25, 2026) resolved two open frontiers in combinatorial delegated choice under downward-closed feasibility constraints: they designed a deterministic, agent-oblivious capped menu guaranteeing the principal strictly more than $29/70 \approx 0.41429$ of the expected offline optimum V (Theorem 1.1, improving the previous 0.1098 bound), established a short analytic proof that a single reported-value threshold guarantees $3/10$ (Theorem 4.1), and separated downward-closed constraints strictly from matroids ($1/2$) via a 5-item knapsack impossibility bounding every deterministic menu by $500/1011$ (Proposition 3.1) and every randomized utility-aware rule by $30/61$ (Theorem 3.2, Proposition 3.3; lowered to 0.48105 in Appendix A). We present a machine-checked, zero-floating-point formalization and exact-$\mathbb{Q}$ certificate auditor for the algebraic, combinatorial, Bernstein-positivity, and menu-score core of these results. First, an 18-theorem zero-dependency Lean 4 kernel (`DownwardClosedDelegation.lean`, Lean 4.35.0; 0 `sorry`, 0 warnings, 0 `Classical.choice`, 10 zero-axiom theorems) formalizes the 5-item knapsack optimum $V = 1011/50$, sub-$1/2$ ratios $500/1011, 5011/10110, 30/61$, normalized target $R = 7000000/2900007$, constructive Lemma 2.1 best-response inclusion-maximality under positive agent utilities, Section 5 seed reserves $b_P > 0$, exact Section 4 Taylor ($\log(13/2) < 15/8$) and small-b ($168034/70785 > 0$) positivity, universal positivity of $P(b) > 0$ over all rationals $b \in [2/13, 2/3]$ via exact homogeneous Bernstein expansion, and a general arbitrary-degree Bernstein convex lower bound. Second, `DELEGATION-CERT-1.0` audits four exact-$\mathbb{Q}$ modes across 16 canonical certificates (0.050–1.077 ms) and rejects 60/60 live adversarial forgeries.

1 Introduction

In combinatorial delegated choice [5, 3, 2], an agent observes independent pairs (X_i, Y_i) (with $X_i \geq 0, Y_i > 0$ arbitrarily correlated) and submits a feasible set $S \in \mathcal{F}$ to maximize $Y(S)$ subject to a principal's pre-committed acceptance rule, benchmarked against $V = \mathbb{E}[\max_{S\in\mathcal{F}} X(S)]$. Unlike combinatorial prophet inequalities—where downward-closed families admit no constant guarantee [6]—Abdi et al. [1] lifted the downward-closed delegation guarantee from 0.1098 [2] to strictly above $29/70$ via capped menus (Algorithm 2, Theorem 5.1), proved an analytic $3/10$ single-threshold bound (Theorem 4.1), and proved via a 5-item knapsack instance that downward-closed delegation is strictly harder than matroid delegation ($1/2$).

Three exact rational pillars anchor [1]:

1. *5-item knapsack separation (Section 3):* With downward closure of $\{A, B\}, \{A, E\}, \{C\}, \{J\}$ and jackpot probability $\varepsilon = 1/100$,

$$V = \frac{1}{100} \cdot 900 + \frac{99}{100}\left(\frac{1}{9} \cdot 18 + \frac{2}{9} \cdot 12 + \frac{2}{3} \cdot 10\right) = \frac{1011}{50}, \tag{1}$$

while every deterministic menu earns at most 10 (ratio $500/1011 < 1/2$), and letting $\varepsilon \to 0$ yields

*Independent researcher; manuscript prepared with transparent AI assistance via Google Antigravity. Source code, Lean 4 proofs, and exact certificates: https://github.com/lluiseriksson/downward-closed-delegation-certificates.

the universal randomized upper bound $30/61 < 1/2$.

2. *Menu score & seeds (Lemma 2.1, Section 5):* Defining $b_M(x) = \min\{x(S) : S \text{ inclusion-maximal in } A_M(x)\}$ lower-bounds the principal's payoff under every agent best response; for clipped weights $w_i = \min(x_i, c)$, a seed P at level t has strictly positive reserve $b_P = t + \min_{i \in P} w_i - w(P) > 0$.
3. *Section 4 exact Bernstein positivity (Theorem 4.1):* Positivity of $\Gamma(\mu_0, b)$ on $[2/13, 2/3]$ reduces via $\sum_{k=0}^{7}(15/8)^k/k! - 13/2 = 3772963/234881024 > 0$ to certifying positivity of the degree-7 polynomial

$$P(b) = 1242b^7 - 3593b^6 + 1795b^5 - 2948b^4 + 7546b^3 - 5459b^2 + 1345b - 72 \quad (2)$$

via exact degree-7 Bernstein coefficients $\beta_0, \ldots, \beta_7 > 0$ on $[2/13, 2/5]$ and $[2/5, 2/3]$.

Contributions and Scope. We formally verify and audit the finite rational, Bernstein-positivity, knapsack-impossibility, and Algorithm 2 menu-score core of [1]. We explicitly scope our Lean 4 kernel away from the continuous-distribution stop-loss integrals of Sections 6–9 and the large-scale LP infeasibility dual of Section 10:

1. **Zero-dependency Lean 4 kernel (§2):** 18 constructive theorems in `DownwardClosedDelegation.lean` (Lean 4.35.0 [4]), 0 `sorry`, 0 warnings, 0 `Classical.choice`, and **10 zero-axiom theorems**.
2. **Pure-stdlib exact-$\mathbb{Q}$ verifier (§3):** `DELEGATION-CERT-1.0`, auditing knapsack-5 impossibility, exact degree-7 Bernstein positivity, Algorithm 2 capped menus, and single-threshold menus with zero floats.
3. **Empirical atlas & red-team audit (§4):** 16 canonical certificates (0.050–1.077 ms), exact $\varepsilon \to 0$ separation curves, full degree-7 Bernstein spectra, and 60/60 rejected live adversarial forgeries.

2 Zero-Dependency Lean 4 Kernel

We formalize 18 constructive theorems in `lean/DownwardClosedDelegation.lean` (Lean 4.35.0, zero external libraries), audited via `lean/AuditAxioms.lean` with 0 warnings, 0 `sorry`, 0 `Classical.choice`, and 0 `Lean.ofReduceBool` (Table 1). Ten of the 18 theorems depend on *zero axioms* (`none`); the remaining 8 depend only on `[propext, Quot.sound]`.

Proposition 1 (Constructive Delegation Kernel)**.**

1. Knapsack-5 impossibility & ordering (Theorems 1–6, 18): *Exact cleared* $V = 1011/50$, *sub-*$1/2$ *ratios* $500/1011$, $5011/10110$, $30/61$, $0.48105 < 30/61$, *strict ordering* $0.1098 < 3/10 < 29/70 < 500/1011 < 1/2$, *normalized target* $R = 7000000/2900007$, *and both Proposition 3.1 branch payoffs* ($= 10$).
2. Constructive Lemma 2.1 maximality, seeds, & closure (Theorems 7–10, 16–17): *Inductive proof over Boolean item masks that strictly positive agent utilities* ($y_i > 0$) *force every best response* S *to be inclusion-maximal in* $A_M(x)$ *and hence satisfy* $b_M \le x(S)$*; Lemma 2.2 singleton anchor* $b_M \le 1$*; Eq. (5.2) seed reserve* $b_P > 0$*; extension weight* $\ge t$*; downward-closure empty-set membership* ($\emptyset \in \mathcal{F}$) *and subset-chain monotonicity; and rescaling homogeneity.*
3. Section 4 universal Bernstein & Taylor positivity (Theorems 11–15): *Exact cleared small-*b *bound* $168034/70785 > 0$, *Taylor certificate* $3772963/234881024 > 0$ *for* $\log(13/2) < 15/8$, *universal Lean 4 proof that* $P(b) > 0$ *for* every *rational* $b = n/d \in [2/13, 2/5]$ *and* $[2/5, 2/3]$ *via exact homogeneous Bernstein expansion (certifying minima* $447966/78125$ *and* $6518/118125$*), and general inductive Bernstein convex lower bound.*

3 Exact-$\mathbb{Q}$ Verifier Architecture

`DELEGATION-CERT-1.0` (`src/delegation_certifier.py`) is a pure-stdlib zero-float verifier enforcing strict JSON types, reduced fractions, downward-closure checks, and exact product-distribution expectations across four modes:

1. **Mode 1 (`knapsack5_impossibility_exact`):** Evaluates the exact Section 3 5-item knapsack optimum $V(\varepsilon)$, deterministic ratio $10/V(\varepsilon)$, and strict gap $1/2 - 10/V(\varepsilon) > 0$.
2. **Mode 2 (`bernstein_sec4_positivity_exact`):** Computes all 8 exact rational degree-7 Bernstein coefficients $\beta_0, \ldots, \beta_7$ of $P(b)$ (2) on $[\ell, r] \subseteq [2/13, 2/3]$, certifies $\min_k \beta_k > 0$, and replays the exact Eq. (4.11) Taylor margin.

Table 1: Complete 18-theorem Lean 4.35.0 kernel in `DownwardClosedDelegation.lean` (0 `sorry`, 0 warnings, 0 `Classical.choice`, 10 zero-axiom theorems).

#	Lean 4 Theorem Identifier	Axiom Footprint
1	`knapsack_expected_opt_exact`	none
2	`knapsack_det_ratio_lt_half`	none
3	`knapsack_rand_ratio_lt_half`	none
4	`knapsack_30_61_lt_half`	none
5	`main_guarantee_exceeds_prev`	none
6	`normalized_R_strict_margin`	none
7	`menu_score_best_response_lb`	[propext, Quot.sound]
8	`inclusion_maximal_singleton_bound`	[propext, Quot.sound]
9	`seed_reserve_strictly_positive`	[propext, Quot.sound]
10	`seed_extension_weight_exceeds_t`	[propext, Quot.sound]
11	`sec4_small_b_gamma_positive`	none
12	`sec4_exp_taylor_log_bound`	none
13	`sec4_bernstein_universal_2_13_to_2_5`	[propext, Quot.sound]
14	`sec4_bernstein_universal_2_5_to_2_3`	[propext, Quot.sound]
15	`bernstein_convex_comb_lb`	[propext, Quot.sound]
16	`downward_closure_empty_and_monotone`	none
17	`rescaling_menu_score_invariant`	[propext, Quot.sound]
18	`knapsack_branch_expected_payoffs`	none

3. **Mode 3 (`capped_menu_score_audit_exact`):** Verifies downward closure of $\mathcal{F}$, replays Algorithm 2 across all three body types (`threshold`, `seeds_upper`, `seeds_extensions`), computes exact $b_M(x)$ over all realizations, and certifies $\mathbb{E}[b_M]/V > 29/70$.
4. **Mode 4 (`single_threshold_menu_exact`):** Replays a single total-value threshold menu ($x(S) \geq z$) over exact finite distributions and certifies $\mathbb{E}[b_M]/V > 3/10$.

4 Empirical Atlas and Red-Team Audit

Running `scripts/run_all_experiments.py` produces 16 canonical certificates (4 per mode) and four exact-$\mathbb{Q}$ experiments, verified end-to-end by `scripts/verify_all.py` with SHA-256 verification across all 41 tracked artifacts.

Knapsack Separation and Bound Landscape (Fig. 1). Mode 1 (0.050–0.079 ms) certifies the exact paper instance at $\varepsilon = 1/100$ ($V = 1011/50$, ratio 500/1011, gap 11/2022 below 1/2) alongside $\varepsilon \in \{1/50, 1/200, 1/1000\}$ reaching $5000/10161 \approx 0.49208$ (converging monotonically to $30/61 \approx 0.49180$). Mode 3 (0.360–1.077 ms) audits Algorithm 2 across all three body types on the 5-item knapsack ($500/1011 > 29/70$), P_4 matchings (6/7),

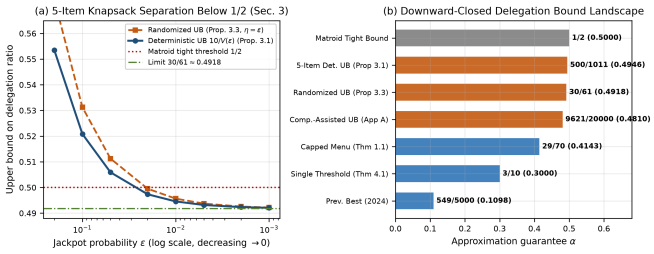

Figure 1: (a) Exact-$\mathbb{Q}$ Section 3 5-item knapsack deterministic ($10/V(\varepsilon)$) and Proposition 3.3 randomized upper bounds crossing strictly below 1/2 toward 30/61 as $\varepsilon \to 0$. (b) Complete downward-closed delegation bound landscape from 0.1098 to 1/2.

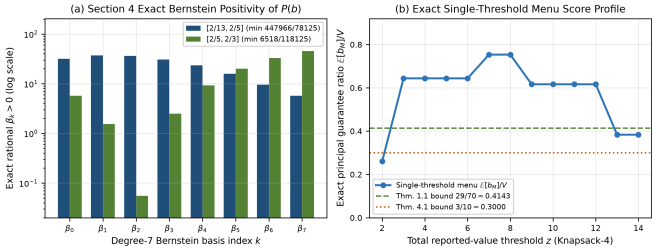

Figure 2: (a) All 16 exact rational degree-7 Bernstein coefficients $\beta_0, \ldots, \beta_7 > 0$ of $P(b)$ (2) on $[2/13, 2/5]$ and $[2/5, 2/3]$. (b) Exact single-threshold menu score profile $\mathbb{E}[b_M]/V$ on 4-item knapsack across $z = 2 \ldots 14$.

C_4 independent sets (7/11), and 4-item knapsack (55/73).

Bernstein Positivity and Threshold Profiles (Fig. 2). Mode 2 (0.275–0.348 ms) independently reproduces Abdi et al.'s exact page 12 Bernstein minima to the last digit: $\min_k \beta_k = 447966/78125$ on $[2/13, 2/5]$ (attained at β_7) and 6518/118125 on $[2/5, 2/3]$ (attained at β_2), plus sub-intervals $[2/13, 1/4]$ (249075/8192) and $[1/2, 2/3]$ (141/32). Mode 4 (0.265–0.732 ms) certifies single-threshold menus exceeding $3/10\,V$ on all four downward-closed families (ratios 500/1011, 6/7, 9/11, and 55/73).

Replay Atlas and Red-Team Audit (Fig. 3, Table 2). All 16 canonical certificates replay in 0.050–1.077 ms with zero floats. Our adversarial suite (`tests/test_suite.py`) tests 60 live forgery attacks—including float/bool confusion, unreduced fractions, oversized rational-string DoS, non-downward-closed families, invalid probability distributions, violated Algorithm 2 parameter invari-

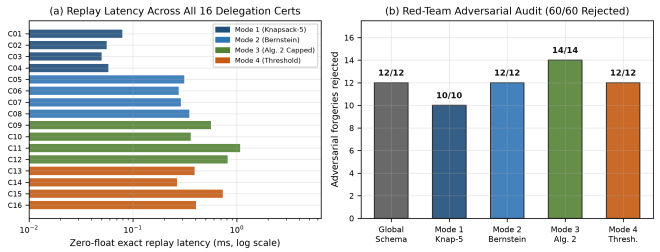

Figure 3: (a) Exact zero-float replay latency across all 16 canonical `DELEGATION-CERT-1.0` certificates (0.050–1.077 ms). (b) Red-team adversarial audit rejecting 60/60 live forgery attacks.

Table 2: Summary of all 16 canonical `DELEGATION-CERT-1.0` certificates and 60/60 live adversarial red-team audit.

Verification Mode	Certs	Latency (ms)	Attacks Rejected
Mode 1: Knapsack-5 Impossibility	4	0.050–0.079	10/10
Mode 2: Sec. 4 Bernstein Positivity	4	0.275–0.348	12/12
Mode 3: Alg. 2 Capped Menu Score	4	0.360–1.077	14/14
Mode 4: Single-Threshold Menu	4	0.265–0.732	12/12
Global Schema / Type / Rational	–	–	12/12
Total	**16**	**0.050–1.077**	**60/60**

ants ($t \leq c$, $a = 0$), tampered Bernstein coefficients, and menus falling below 29/70 or 3/10—achieving **60/60 rejections** with zero bypasses.

5 Conclusion

We delivered an 18-theorem zero-dependency Lean 4 kernel (10 zero-axiom theorems, 8 `[propext, Quot.sound]`, 0 `Classical.choice`) and the exact-$\mathbb{Q}$ verifier `DELEGATION-CERT-1.0` for the finite rational, universal Bernstein-positivity, knapsack-impossibility, and menu-score core of Abdi et al.'s downward-closed delegation bounds [1]. Mechanizing the full continuous stop-loss inequalities of Sections 6–9 and importing the Section 10 Farkas LP infeasibility dual into Lean 4 remain exciting targets for future formalized mechanism design.

References

[1] Hamed Abdi, Kiarash Banihashem, Mohammad-Taghi Hajiaghayi, and Danny Mittal. Improved downward-closed delegation bounds. *arXiv preprint arXiv:2611.33095*, 2026.

[2] Kiarash Banihashem, MohammadTaghi Hajiaghayi, Piotr Krysta, and Suho Shin. Delegated choice under combinatorial constraints. In *Proceedings of the 25th ACM Conference on Economics and Computation (EC)*, 2024.

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

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

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

[6] Aviad Rubinstein and Sahil Singla. Combinatorial prophet inequalities. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1671–1687, 2017.