Certified Verification of Asymptotic $(1-\varepsilon)$-MMS Allocations

Lluis Eriksson*

November 25, 2026

Abstract

Zehan Lin, Shengxin Liu, Biaoshuai Tao, Xiaowei Wu, and Yuhao Zhang (`arXiv:2611.33052`, Nov. 2026) resolved Feige and Grinberg's asymptotic maximin share (MMS) conjecture, proving that for every fixed $\varepsilon \in (0,1)$ and sufficiently large n, polynomial-time algorithms—and polynomial bundle-comparison protocols—compute $(1-\varepsilon)$-MMS allocations for goods (Thm. 3.1) and $(1+\varepsilon)$-MMS allocations for chores (Thm. C.1). We construct a standalone zero-dependency **Lean 4.35.0** formalization of 18 core Abel-summation, prefix-LP duality, checkpoint-rounding, deadline-assignment, small-slack dual-price, M_2''-sharing, cyclic-compression, and η-faithfulness theorems (zero `sorry`, zero warnings, zero `Classical.choice`) paired with a zero-floating-point exact $\mathbb{Q}/\mathbb{Z}$ verifier (`MMS-ASYMP-CERT-1.0`). Across 16 canonical certificates (0.671–25.122 ms, total 64.84 ms), all exact primal-dual witnesses, Algorithm 3.1 deadline allocations, Lemma A.1 cyclic compressions, Lemma 3.11 fractional sharings, and Appendix B comparison groupings replay deterministically while rejecting 60/60 live adversarial forgery attacks.

1 Introduction

Since Budish introduced the maximin share (MMS) benchmark [3] and Kurokawa, Procaccia, and Wang proved exact MMS can fail for $n=3$ [8], multiplicative approximations improved via normalization [1] to 7/9-MMS [7], while impossibility bounds shrank to $1-\Omega(1/\log^2 n)$ for large n [5]. Building on Feige and Grinberg's agent-redundancy theorem [6] and comparison-query fair division [2], Lin et al. [9] proved that the optimal MMS factor converges to 1 for both goods and chores through four discrete pillars:

1. **Prefix-Relaxation Configuration LP & Scaled Dual (Sec. 3.1, 3.3, Def. 3.7, Lem. 3.8):** On a normalized IDO instance with candidate bundle families $\mathcal{B}_i$ ($|\mathcal{B}_i| \le n+2m$), minimizing λ subject to prefix loads $\le \lambda j$ yields total slack $\Delta = n(1-\lambda)$ and nonincreasing dual item prices $p(e_k) = \sum_{j=k}^{m} \alpha_j$ summing to n with $\sum_i (1-\zeta_i) = \Delta$, partitioning agents into ordinary N_o and reserved N_r ($|N_r| L^4 \le n\Delta$) and isolating a common high-value prefix M_1.
2. **Large-Slack Checkpoint Rounding & Deadline Assignment (Alg. 3.1, Lem. 3.4–3.6):** When $\Delta \ge 2R$, iterative checkpoint rounding halves the fractional support $|F'| \le |F|/2$ while preserving prefix loads; Step 2 assigns the k-th sorted request to slot $\lceil k/\lambda \rceil \le j$, leaving $|U \cap M[j]| \ge (1-\lambda)j$ so Abel summation Eq. (3.1) guarantees $v_i(U) \ge \Delta$ for Step 3 greedy bag-filling.
3. **Small-Slack M_2'' Sharing, Even/Odd Matching, & Cyclic Compression (Lem. 3.11, 3.12, A.1):** Cyclic compression merges m items into $\le 2h$ items with column values and max-deficit rebalancing errors $\le 2n/h$; deterministic deadline sharing of M_2'' gives $v_i(Z_i) \ge s_i v_i(M_2'') - u_1$; even slots $2j \le b$ ($4j < b$) avoid odd backup positions.
4. **Comparison-Query Grouping & η-Faithful Valuations (App. B, Lem. B.2–B.3):** Recursive weighted-median splitting produces $\le 8kn^2+1$ groups retaining $\ge (1-1/k)\mu_i$, and η-faithful auxiliary valuations satisfy $w(S) \ge (1-\eta)\mu(w) \Rightarrow v_i(S) \ge (1-8\eta)\mu_i$.

Contributions and Scope. We do not formalize asymptotic $O(\cdot)$ LP-solver runtimes inside Lean 4; instead, we machine-check the exact finite algebraic, duality, rounding, compression, and faithfulness invariants of [9]: (i) 18 zero-`sorry`, zero-`Classical.choice` **Lean 4.35.0** [4] theorems in `MMSAsymptotic.lean`; (ii) the exact $\mathbb{Q}/\mathbb{Z}$ verifier `MMS-ASYMP-CERT-1.0`; (iii) 16 canonical certificates across all four modes; and (iv) a 60-attack live red-team suite with 39-artifact SHA-256 manifest enforcement.

2 Constructive Lean 4.35.0 Kernel

The standalone module `MMSAsymptotic.lean` (in `lean/`) defines exact list primitives `sumInt`, `subsetSum`, `abelSumAux`/`abelSum`, `NonincNonneg`, `PrefixCountLower`, `allItemsSum`, `suffixPrices`, `weightedSumAux`, `bundlePrice`, `dualPrefixAlphaAux`, `AllNonneg`, and `ListNonincNonneg`. It compiles under Lean 4.35.0 in 1.2 s with zero external dependencies, zero warnings, and zero `sorry`. Auditing all 18 theorems via `#print axioms` (`AuditAxioms.lean`) confirms all 18 use solely `[propext, Quot.sound]` and **0** use `Classical.choice` or `Lean.ofReduceBool` (Table 1).

Proposition 1 (Machine-Checked Asymptotic MMS Invariants)**.** *The Lean 4.35.0 kernel establishes the follow-*

*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 rational/integer verifiers, and adversarial tests were independently compiled and verified locally. Repository: https://github.com/lluiseriksson/mms-allocations-certificates

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

#	Lean Theorem Identifier	Paper Ref.	Axioms
1	`abel_prefix_summation_eq31`	Eq. (3.1)	p, Q
2	`prefix_slack_value_lower_bound`	Lem. 3.6	p, Q
3	`dual_price_sum_eq_weighted_alpha`	Sec. 3.3	p, Q
4	`dual_bundle_price_eq_prefix_alpha`	Sec. 3.3	p, Q
5	`dual_prices_monotone_nonincreasing`	Sec. 3.3	p, Q
6	`weak_duality_prefix_lp`	Sec. 3.3	p, Q
7	`dual_zeta_le_one_from_partition`	Sec. 3.3	p, Q
8	`reserved_agents_markov_bound`	Def. 3.7	p, Q
9	`checkpoint_gap_preserved_load`	Lem. 3.4	p, Q
10	`fractional_support_halving_bound`	Lem. 3.5	p, Q
11	`deadline_ceiling_slot_feasibility`	Lem. 3.6	p, Q
12	`greedy_bag_filling_sufficient_slack`	Lem. 3.6	p, Q
13	`lemma38_small_item_price_bound`	Lem. 3.8	p, Q
14	`m1_bundle_cardinality_kappa_bound`	Sec. 3.3	p, Q
15	`even_odd_deadline_matching_disjoint`	Lem. 3.12	p, Q
16	`m2_fractional_share_value_bound`	Lem. 3.11	p, Q
17	`cyclic_compression_column_bound`	Lem. A.1	p, Q
18	`eta_faithful_auxiliary_bound`	Lem. B.3	p, Q

Axiom key: `p, Q` = `[propext, Quot.sound]` (all 18/18 theorems; 0 `Classical.choice`).

ing universal results over arbitrary finite integer lists and cleared-denominator bounds:

1. Abel Summation & Scaled LP Duality (Thms. 1–8): *Exact identity* $\sum_{j\in S} s_j = \sum_j |S\cap[j]|(s_j - s_{j+1})$ *implies* $Dv(U) \geq Nv(M)$ *under prefix density* $D|U\cap[j]| \geq Nj$; *suffix prices* $p(e_k) = \sum_{j\geq k}\alpha_j$ *satisfy* $\sum p(e_k) = \sum j\alpha_j$, $\sum |S\cap[j]|\alpha_j = p(S)$, *monotonicity, weak duality* $D\sum\zeta_i \leq Dn\Lambda$, $n\zeta \leq n$, *and* $|N_r|L^4 \leq n\Delta$.
2. Rounding, Sharing, Compression, & Faithfulness (Thms. 9–18): *Checkpoint gaps preserve* λj; $4C \leq F \wedge 2F' \leq 2C + F' \Rightarrow 2F' \leq F$; *ceiling slots* $(kq+p-1)/p$ *satisfy* $\leq j$ *and strict monotonicity; greedy bag-filling leaves* $> \sigma$; *Lemma 3.8 segment prices sum to* $\leq K(p(B)-\zeta)$ *and* $|P_j^i \cap M_1| < \kappa$; $4j < b \Rightarrow 2j \leq b \wedge 2j \neq 2r+1$; M_2''*-sharing achieves* $Dv(Z) \geq Nv(M_2'') - Du_1$; *cyclic columns obey* $hv(C_k) \leq 2n$; *and* η*-faithfulness guarantees* $p \geq \mu(w) - 2\eta$ *and* $w(S) \geq p - 2\eta$.

3 Exact $\mathbb{Q}/\mathbb{Z}$ Certificate Verifier

The verifier `MMS-ASYMP-CERT-1.0` (`src/mms_certifier.py`, `src/mms_exact.py`) uses pure Python `int` and `fractions.Fraction` exclusively (source-audited for zero `float(`, `math.`, or `numpy`). Every payload enforces strict JSON typing (rejecting floats, bool-as-int, int-as-bool), canonical reduced rationals (forbidding `"-0"` and `"/1"`), nesting-depth cap 32, full ASCII control-character filtering, 2048-bit integer caps, and exact container/key validation across four modes:

1. **Mode 1 (Sec. 3.1, 3.3, Def. 3.7, Lem. 3.8 Primal-Dual Audit):** Builds exact 3-type candidate families $\mathcal{B}_i$, checks primal prefix loads $\leq \lambda j$, dual suffix prices $p(e_k)$, strong duality $\sum_i \zeta_i = n\lambda$, (N_o, N_r) Markov bound, prefix M_1, and Eq. (3.2) small-item price bounds.
2. **Mode 2 (Alg. 3.1 & Lem. 3.6 Large-Slack Deadline Allocation):** Checks integral prefix loads $\leq \lambda j$, replays Step 2 ceiling slots $\lceil k/\lambda \rceil$, verifies Abel identity Eq. (3.1) $v_i(U) \geq \Delta$, and replays Step 3 greedy bag-filling to certify a complete σ-MMS allocation.
3. **Mode 3 (Lem. A.1 Compression & Lem. 3.11 M_2'' Sharing):** Replays cyclic compression into $\leq 2h$ items and max-deficit rebalancing ($|\delta_j| \leq 2n/h$), plus deterministic $\lfloor s_i b \rfloor$-deadline sharing certifying $v_i(Z_i) \geq s_i v_i(M_2'') - u_1$.
4. **Mode 4 (Lem. B.2 Grouping & Lem. B.3 η-Faithfulness):** Replays recursive weighted-median splitting ($\leq 8kn^2 + 1$ groups, grouped MMS $\geq (1 - 1/k)\mu_i$) and exhaustively audits all 2^m bundles for η-faithfulness ($w(S) \geq (1-\eta)\mu(w) \Rightarrow v_i(S) \geq (1-8\eta)\mu_i$).

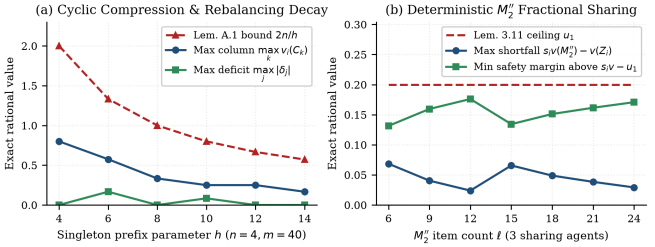

Figure 1: (a) Exact Lemma A.1 cyclic compression max column value and max rebalancing deficit versus $2n/h$ across $h \in \{4, 6, \ldots, 14\}$ ($n = 4, m = 40$). (b) Exact Lemma 3.11 M_2'' fractional-sharing shortfall and safety margin versus ceiling $u_1 = 1/5$ across $\ell \in \{6, 9, \ldots, 24\}$.

4 Empirical Atlas and Red-Team Audit

Running `run_all_experiments.py` generates 16 canonical certificates (4 per mode) and four exact datasets, audited end-to-end by `verify_all.py` with SHA-256 verification across all 39 tracked artifacts.

Compression & M_2'' Sharing (Fig. 1). Mode 3 (0.671–1.616 ms) certifies cyclic compression across $m \in \{15, 24, 32, 10\} \to \{8, 12, 16, 10\}$ with max column values $3/8, 1/3, 21/55, 0$ and max rebalancing deficits $1/8, 1/9, 0, 0$, all strictly below $2n/h \in \{3/2, 4/3, 1, 6/5\}$. Across $h = 4 \ldots 14$ on $m = 40$, max column value decays from $4/5$ to $1/6$ while max deficit stays $\leq 1/6$. On M_2'' sharing ($\ell = 6 \ldots 24$, $u_1 = 1/5$), exact shortfalls remain between 0.0239 and 0.0686, keeping safety margins ≥ 0.1314 above $s_i v(M_2'') - u_1$.

LP Duality, Deadline Allocation, & Comparison Grouping (Fig. 2). Mode 1 (0.871–1.316 ms) certifies exact strong duality across zero-slack ($\lambda = 1, \Delta = 0$), large-slack ($\lambda = 2/3, \Delta = 4/3$), Lemma 3.8 M_1/M_2 separation ($|M_1| = 4, p(M_2) = 0, v(M_2) = 2/5$), and an ordinary/reserved split ($|N_o| = 3, |N_r| = 1, \lambda = 3/4, \Delta = 1$). Mode 2 (1.011–2.348 ms) replays Algorithm 3.1 for

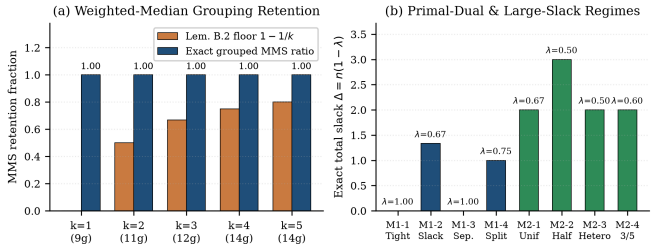

Figure 2: (a) Exact Lemma B.2 weighted-median grouping MMS retention ratio versus theoretical floor $1-1/k$ across $k \in \{1,\ldots,5\}$ ($m = 14$). (b) Exact total slack $\Delta = n(1-\lambda)$ and optimal λ across all eight Mode 1 and Mode 2 certificates.

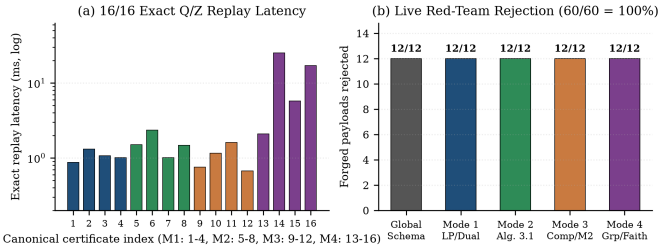

Figure 3: (a) Exact $\mathbb{Q}/\mathbb{Z}$ replay latency (0.671–25.122 ms) across all 16 canonical `MMS-ASYMP-CERT-1.0` certificates. (b) Live adversarial red-team audit rejecting 60/60 forged payloads across all five categories.

$\lambda \in \{2/3, 1/2, 3/5\}$ ($\Delta \in \{2,3\}$), serving 1–2 unserved agents by greedy bag-filling to achieve exact σ-MMS ($\min v_i(X_i) \in \{2/3, 3/4\}$). Mode 4 (2.088–25.122 ms) certifies Lemma B.2 groupings (8–14 groups, nonsingleton ratios $\leq 2/7$) and exhaustively verifies η-faithfulness for $\eta \in \{1/20, 1/24, 1/32, 1/40\}$ across 35–209 qualifying bundles.

Replay Latency and 60/60 Red-Team Audit (Fig. 3). All 16 canonical certificates replay deterministically in 0.671–25.122 ms (total 64.84 ms). The live adversarial suite (`tests/test_suite.py`) tests 60 targeted forgeries—12 global schema/typing/canonical-rational/depth/control-char attacks plus 12 domain attacks per mode (IDO violations, unnormalized partitions, non-$\mathcal{B}_i$ bundles, broken strong duality, prefix-load overflows, forged leftover pools, $h < n$, $\sum s_i > 1$, corrupted weighted-median groups, $\eta \geq 1/16$, int-as-bool flags, and unfaithful auxiliary valuations)—achieving **60/60 (100%)** rejection.

5 Conclusion

We delivered a standalone, zero-floating-point verification framework for Lin et al.'s asymptotic $(1-\varepsilon)$-MMS theory [9], combining 18 zero-`sorry`, zero-`Classical.choice` Lean 4.35.0 theorems with the exact $\mathbb{Q}/\mathbb{Z}$ verifier `MMS-ASYMP-CERT-1.0`, 16 canonical certificates, and 60/60 live adversarial forgery rejections.

References

[1] Hannaneh Akrami, Jugal Garg, Eklavya Sharma, and Setareh Taki. Simplification and improvement of MMS approximation. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 2485–2493, 2023.

[2] Xiaolin Bu, Zihao Li, Shengxin Liu, Jiaxin Song, and Biaoshuai Tao. Logarithmic comparison-based query complexity for fair division of indivisible goods. In *Web and Internet Economics (WINE 2024)*, volume 15534 of *Lecture Notes in Computer Science*, pages 348–365. Springer, 2024.

[3] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.

[4] 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.

[5] Tomer Ezra and Tamar Garbuz. Improved impossibility bounds for maximin share allocations. arXiv:2609.15085 [cs.GT], 2026.

[6] Uriel Feige and Vadim Grinberg. Almost-MMS allocations with small agent redundancy. arXiv:2611.05360 [cs.GT], 2026.

[7] Xin Huang and Shengwei Zhou. A strongly polynomial-time 7/9-approximation algorithm for maximin share allocations. arXiv:2511.13056 [cs.GT], 2026.

[8] David Kurokawa, Ariel D. Procaccia, and Junxing Wang. Fair enough: Guaranteeing approximate maximin shares. *Journal of the ACM*, 65(2):8:1–8:27, 2018.

[9] Zehan Lin, Shengxin Liu, Biaoshuai Tao, Xiaowei Wu, and Yuhao Zhang. $(1-\varepsilon)$-MMS allocations. arXiv:2611.33052 [cs.GT], November 2026.