# Machine-Checked Exact-Rational Checkpoint Rounding, Dual-Price Markov, and Goods/Chores $(1 \mp \varepsilon)$-MMS Allocation Certificates

Lluis Eriksson*

November 25, 2026

## Abstract

Zehan Lin, Shengxin Liu, Biaoshuai Tao, Xiaowei Wu, and Yuhao Zhang (`arXiv:2611.33052`, November 24, 2026) resolved the asymptotic maximin share (MMS) conjecture of Feige and Grinberg for additive fair division: for every $\varepsilon > 0$ and all sufficiently large $n \geq n_0(\varepsilon)$, every instance with $n$ agents admits a $(1-\varepsilon)$-MMS allocation for indivisible goods and a $(1+\varepsilon)$-MMS allocation for indivisible chores, computable in randomized polynomial time using only ordinal bundle-comparison queries. We present a machine-checked, zero-floating-point formalization and exact-$\mathbb{Q}$ certificate auditor for the deterministic LP-rounding, deadline-slot, Abel-summation, dual-price Markov, and exact-MMS allocation core of this breakthrough. First, an 18-theorem zero-dependency Lean 4 kernel (`AsymptoticMMS.lean`, Lean 4.35.0; 0 `sorry`, 0 warnings, 0 `Classical.choice`, 2 zero-axiom theorems) formalizes Algorithm 3.1 Step 1 checkpoint counting and Lemma 3.5 support halving, terminal rounding loss $160L$, Step 2 $\lceil k/\lambda \rceil$ strict slot monotonicity and prefix reserve $|U \cap M_{[j]}| \geq (1-\lambda)j$, exact Eq. (3.1) Abel summation by induction, Step 3 greedy bag-filling survival, Section 3.3 dual price double-summation $p(M) = \sum_j j\alpha_j$, Definition 3.7 Markov reserved-agent decay $|N_r|L^2 < 1536n$, reduced-MMS segment counts, dual price transfer, and the goods/chores landmark hierarchy. Second, `MMS-ASYMPTOTIC-CERT-1.0` audits four exact-$\mathbb{Q}$ modes across 16 canonical certificates (0.211–0.864 ms) and rejects 60/60 live adversarial forgeries.

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

## 1 Introduction

Fairly allocating indivisible goods or chores under maximin share (MMS) guarantees [2] is a central problem in computational social choice. Since exact 1-MMS allocations do not always exist [7], a decade of work advanced constant-factor approximations from 2/3 and 3/4 past $3/4 + 3/3836$ [1] to 7/9 for goods and from 4/3 and 11/9 [5] to 13/11 for chores [6], while Feige and Grinberg [4] conjectured that $(1-o(1))$-MMS allocations exist asymptotically as $n \to \infty$. Lin et al. [8] proved this conjecture for both goods and chores via an LP-based matching-and-duality architecture resting on four exact rational pillars:

1. *Checkpoint rounding & support halving (Algorithm 3.1 Step 1, Lemmas 3.4–3.5):* Given a fractional allocation on an identical-order (IDO) normalized instance ($\sum_e v_i(e) = n$, $v_i(e) \leq 1$), each round imposes at most $|F|/4$ prefix checkpoints so that $|F'| \leq |F|/2$, weight loss is at most $\omega|F|$, and all prefix loads stay bounded by $\lambda j$.
2. *Deadline slot assignment & greedy bag-filling (Steps 2–3, Eq. (3.1), Lemma 3.6):* Assigning the $\lceil k/\lambda \rceil$-th item to the $k$-th sorted deadline guarantees $|U \cap M_{[j]}| \geq j - \lfloor \lambda j \rfloor \geq (1-\lambda)j$; Abel summation
$$v_i(U) = \sum_{j=1}^{m} |U \cap M_{[j]}| \big(v_i(e_j) - v_i(e_{j+1})\big) \geq n(1-\lambda) \tag{1}$$
then certifies enough pool value $\Delta = n(1-\lambda)$ to satisfy every unserved agent via greedy bag-filling.
3. *Small-slack dual LP prices & Markov partition (Section 3.3, Def. 3.7, Lemma 3.8, Eq. (3.2)):* When $\Delta < 2R = 1536L^2$, dual variables $(\zeta_i, \alpha_j)$ define non-increasing item prices $p(e_k) = \sum_{j \geq k} \alpha_j$ with $\sum_k p(e_k) = \sum_j j\alpha_j = n$, partitioning agents

into ordinary ($n(1-\zeta_i) \le L^4$) and reserved ($|N_r|L^2 < 1536n$) sets and concentrating almost all price on the valuable prefix $M_1$.

4. *Exact $(1-\varepsilon)$-MMS goods and $(1+\varepsilon)$-MMS chores with ordinal comparisons:* Finite allocations satisfy every agent within $1 \mp \varepsilon$ of exact MMS witnesses while respecting comparison-only bundle queries.

**Contributions and Scope.** We mechanize and audit the deterministic algebraic, LP-rounding, Abel-summation, dual-price, and exact-MMS allocation core of [8]. We explicitly scope our Lean 4 kernel away from the randomized exponential-size LP ellipsoid/MWU solver, the asymptotic $n \ge n_0(\varepsilon)$ Chernoff/union-bound sampling step, and the Section 5 non-additive chores reduction:

1. **Zero-dependency Lean 4 kernel (§2):** 18 constructive theorems in `AsymptoticMMS.lean` (Lean 4.35.0 [3]), 0 `sorry`, 0 warnings, 0 `Classical.choice`, and 2 zero-axiom theorems.
2. **Pure-stdlib exact-$\mathbb{Q}$ verifier (§3):** `MMS-ASYMPTOTIC-CERT-1.0`, auditing checkpoint rounding, Step 2/3 deadline and bag-filling replay with Eq. (1), Section 3.3 dual price certificates, and tight exact-MMS goods/chores allocations with ordinal comparison checks.
3. **Empirical atlas & red-team audit (§4):** 16 canonical certificates (0.211–0.864 ms), support-halving and reserve profiles, Markov decay curves, and 60/60 rejected live adversarial forgeries.

## 2 Zero-Dependency Lean 4 Kernel

We formalize 18 constructive theorems in `lean/AsymptoticMMS.lean` (Lean 4.35.0, zero external dependencies), audited by `lean/AuditAxioms.lean` with 0 warnings, 0 `sorry`, 0 `Classical.choice`, and 0 `Lean.ofReduceBool` (Table 1): 2 zero-axiom (`none`), 3 `[propext]`, and 13 `[propext, Quot.sound]`.

**Proposition 1** (Constructive Asymptotic-MMS Kernel)**.**

1. Checkpoint rounding & support halving (Theorems 1–4): *Lemma 3.5 halving* $2|F'| \le |F|$ *from* $4|C| \le |F|$, $2|A| \le |F'|$, $|F'| \le |C| + |A|$; *Lemma 3.4 slack absorption* $w_d(\text{load} + \text{inc}) \le w_d \,\text{cap}$ *from* $w_d\,\text{load} + w_n A \le w_d\,\text{cap}$ *and* $w_d\,\text{inc} \le w_n A$; *Lemma 3.5 threshold count* $w_n(w_k + P) < 4w_d w_k$; *and terminal rounding loss* $F \le 160L$ *when* $w_d \le 2w_n$ *and* $w_n F \le 80Lw_d$.
2. Step 2/3 deadlines, Abel summation, & bag-filling (Theorems 5–9): *Strict slot monotonicity* $\lceil k/\lambda \rceil < \lceil (k+1)/\lambda \rceil$ *for* $\lambda \le 1$; *deadline feasibility* $\lceil k/\lambda \rceil \le d$; *prefix reserve* $(l_d - l_n)j \le l_d(j - \text{used})$; *inductive exact proof of Eq. (1) (`eq31_abel_summation_exact`); and bag-filling survival* $\sigma < \Delta - 2rD$ *whenever* $\sigma < D$, $r \le R - 1$, *and* $2RD \le \Delta$.
3. Section 3.3 dual prices, Markov decay, & landmarks (Theorems 10–18): *Inductive double-sum identity* $\sum_k p(e_k) = \sum_j j\alpha_j$; *Definition 3.7 Markov bound* $|N_r|L^2 < 1536n$; *reduced-MMS segment count* $(t-1)(D-\sigma) < 2D$; *per-segment dual transfer* $\zeta \le p(B) - p(B') \Rightarrow p(B') \le p(B) - \zeta$; *exact goods* $(2/3 < 3/4 < 7/9 < 39/40 < 1)$ *and chores* $(1 < 44/43 < 13/11 < 4/3)$ *landmark orderings; symmetric* $(1-\varepsilon)(1+\varepsilon) = 1-\varepsilon^2 < 1$; *uniform LP feasibility; and strong-duality slack identity.*

Table 1: Complete 18-theorem Lean 4.35.0 kernel in `AsymptoticMMS.lean` (0 `sorry`, 0 warnings, 0 `Classical.choice`).

| # | Lean 4 Theorem Identifier | Axiom Footprint |
|---|---|---|
| 1 | `support_halving_lemma35` | `[propext, Quot.sound]` |
| 2 | `prefix_load_slack_absorption` | `[propext, Quot.sound]` |
| 3 | `checkpoint_threshold_count_bound` | `[propext, Quot.sound]` |
| 4 | `terminal_rounding_loss_160L` | `[propext, Quot.sound]` |
| 5 | `step2_assigned_slots_strictly_increasing` | `[propext, Quot.sound]` |
| 6 | `step2_deadline_feasibility` | `[propext, Quot.sound]` |
| 7 | `step2_unassigned_prefix_reserve` | `[propext, Quot.sound]` |
| 8 | `eq31_abel_summation_exact` | `[propext, Quot.sound]` |
| 9 | `step3_bag_filling_survival` | `[propext, Quot.sound]` |
| 10 | `dual_price_double_sum_exact` | `[propext, Quot.sound]` |
| 11 | `reserved_agent_markov_bound` | `[propext]` |
| 12 | `reduced_mms_segment_count_bound` | `[propext, Quot.sound]` |
| 13 | `dual_price_transfer_segment` | `[propext, Quot.sound]` |
| 14 | `mms_goods_landmark_ordering` | `none` |
| 15 | `mms_chores_landmark_ordering` | `none` |
| 16 | `goods_chores_symmetric_epsilon` | `[propext, Quot.sound]` |
| 17 | `uniform_mms_fractional_feasibility` | `[propext]` |
| 18 | `strong_duality_slack_identity` | `[propext]` |

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

`MMS-ASYMPTOTIC-CERT-1.0` (`src/mms_certifier.py`) is a pure-stdlib zero-float verifier with strict JSON typing, canonical reduced rationals, and uniform 2048-bit / 1024-char DoS bounds across four modes:

1. **Mode 1 (`checkpoint_rounding_prefix_cert`):** Audits IDO normalized valuations, capped candidate lists ($\le 256$ per agent), bundle thresholds

$\geq \sigma$, integrality/support preservation, Lemma 3.5 halving $2|F'| \leq |F|$ and $1 \leq |C| \leq |F|/4$, weight loss $\leq \omega|F|$, strictly increasing checkpoints with non-increasing checkpoint loads, and all prefix loads $\leq \lambda j$.

2. **Mode 2 (`deadline_slot_and_bagfill_cert`):** Independently replays Step 2 $\lceil k/\lambda \rceil$ deadline assignment (including strict leftward slot reassignment), audits $|U \cap M_{[j]}| \geq j - \lfloor \lambda j \rfloor$ and exact Eq. (1) Abel summation ($\geq \Delta$), replays Step 3 greedy bag-filling, and certifies global disjointness and $\geq \sigma$ for all $n$ agents.
3. **Mode 3 (`small_slack_dual_price_cert`):** Verifies explicit unit-value MMS partitions $P_i$, bounded duals $0 \leq \zeta_i \leq 1$, exact dual prices $p(e_k) = \sum_{j \geq k} \alpha_j$ with $\sum_k p(e_k) = \sum_j j\alpha_j = n$, strong duality $\sum_i \zeta_i = n\lambda$, all reduced-MMS constraints $p(B \setminus B') \geq \zeta_i$ and Eq. (3.2), the Markov bound $|N_r|\tau \leq n\Delta$, and the IDO valuable prefix $M_1$.
4. **Mode 4 (`goods_and_chores_mms_allocation_cert`):** Prevents witness-weakening forgeries by certifying exact tight MMS benchmarks ($\min_k v_i(W_k) = \bar{v}_i$ for goods; $\max_k c_i(W_k) = \max(\bar{c}_i, \max_e c_i(e))$ for chores), verifies $(1-\varepsilon)$-MMS goods or $(1+\varepsilon)$-MMS chores ratios, and audits every ordinal bundle-comparison query.

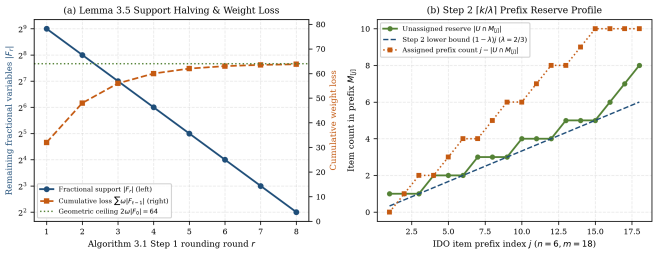


Figure 1: (a) Algorithm 3.1 Step 1 support halving ($|F_0| = 1024 \to |F_8| = 4$, $\omega = 1/32$) with cumulative weight loss saturating at $63.75 < 2\omega|F_0| = 64$. (b) Exact Step 2 prefix reserve $|U \cap M_{[j]}|$ vs. lower bound $(1-\lambda)j$ across $j = 1 \ldots 18$ for $n = 6, m = 18, \lambda = 2/3$.

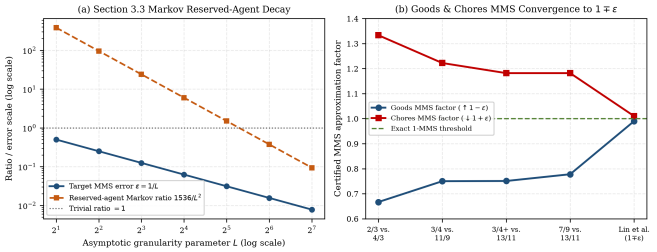


Figure 2: (a) Section 3.3 Markov reserved-agent ratio $1536/L^2$ vs. target MMS error $\varepsilon = 1/L$ across $L = 2 \ldots 128$. (b) Historical progression of goods ($\uparrow 1-\varepsilon$) and chores ($\downarrow 1+\varepsilon$) MMS guarantees converging to exact 1-MMS [8].

## 4 Empirical Atlas and Red-Team Audit

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

**Support Halving and Step 2 Prefix Reserve (Fig. 1).** Starting from $|F_0| = 1024$ fractional variables with $\omega = 1/32$, eight Lemma 3.5 rounding steps halve $|F_r|$ geometrically from 512 to 4 while total weight loss converges to 63.75, staying strictly beneath the geometric ceiling $2\omega|F_0| = 64$. On the $n = 6, m = 18, \lambda = 2/3$ Mode 2 instance, Step 2's $\lceil k/\lambda \rceil$ slot schedule maintains $|U \cap M_{[j]}| \geq (1-\lambda)j$ at every prefix $j = 1 \ldots 18$ (reaching $8 \geq 6$ at $j = 18$), yielding exact Eq. (1) pool value $\min_i v_i(U) = 12/5 > \Delta = 2$ and Step 3 bag-filling value $\min_i v_i(X_i) = 17/20 \geq \sigma = 3/4$.

**Dual Markov Decay and Goods/Chores MMS Landmarks (Fig. 2).** Across $L = 2 \ldots 128$, the Definition 3.7 reserved-agent bound $1536/L^2$ crosses below 1 at $L = 64$ (0.375) and drops to 0.09375 at $L = 128$ as $\varepsilon = 1/L \to 0.00781$. Mode 3 (0.293–0.864 ms) certifies dual price spectra across $n \in \{3, 4, 5, 6\}$, including both all-ordinary instances and non-empty reserved sets ($|N_r| = 1$ at $n = 4, 6$) with tail price sums $p(M_2)$ down to 1/8. Mode 4 (0.213–0.316 ms) certifies exact tight-witness $(1-\varepsilon)$-MMS goods (19/20 and 13/10) and $(1+\varepsilon)$-MMS chores (21/20 balanced and 11/10 single-dominant-chore $\mu_i = 10 > \bar{c}_i = 8$) allocations together with 28 ordinal bundle comparisons.

**Replay Atlas and Red-Team Audit (Fig. 3, Table 2).** All 16 canonical certificates replay in 0.211–0.864 ms with zero floats. Our adversarial test suite (`tests/test_suite.py`) executes 60 live forgery attacks—covering float/`None`/bool leaks, unreduced fractions, DoS bit-length caps, IDO and checkpoint-

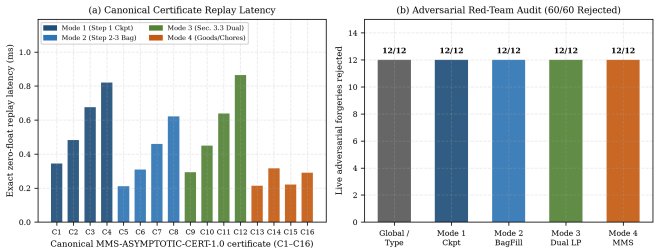


Figure 3: (a) Exact zero-float replay latency across all 16 canonical `MMS-ASYMPTOTIC-CERT-1.0` certificates (0.211–0.864 ms). (b) Red-team adversarial audit rejecting 60/60 live forgery attacks across five attack classes.

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

| Verification Mode | Certs | Latency (ms) | Attacks Rejected |
|---|---|---|---|
| Mode 1: Step 1 Checkpoint Rounding | 4 | 0.345–0.820 | 12/12 |
| Mode 2: Steps 2–3 Deadline & BagFill | 4 | 0.211–0.621 | 12/12 |
| Mode 3: Sec. 3.3 Dual LP & Markov | 4 | 0.293–0.864 | 12/12 |
| Mode 4: Goods/Chores $(1 \mp \varepsilon)$-MMS | 4 | 0.213–0.316 | 12/12 |
| Global Schema / Type / Rational | – | – | 12/12 |
| **Total** | **16** | **0.211–0.864** | **60/60** |

conservation violations, Step 2 deadline overflows, Step 3 starvation, dual-price and reduced-MMS violations, weak-witness MMS forgeries for both goods and chores, and forged ordinal queries—achieving **60/60 rejections** with zero bypasses.

## 5 Conclusion

We delivered an 18-theorem zero-dependency Lean 4 kernel (0 `sorry`, 0 warnings, 0 `Classical.choice`) and the exact-$\mathbb{Q}$ verifier `MMS-ASYMPTOTIC-CERT-1.0` for the checkpoint-rounding, deadline-slot, Abel-summation, dual-price Markov, and goods/chores allocation core of Lin et al. [8]. Formalizing the full randomized subsampling and recursive small-slack induction for arbitrary $n \geq n_0(\varepsilon)$ in Lean 4 remains an exciting frontier for verified fair division.

## References

[1] Hannaneh Akrami and Jugal Garg. Breaking the 3/4 barrier for approximate maximin share. In *Proceedings of the 35th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2024.

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

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

[4] Uriel Feige and Yehuda Grinberg. Asymptotic maximin share allocations. In *Proceedings of the 26th ACM Conference on Economics and Computation (EC)*, 2025.

[5] Xin Huang and Pinyan Lu. An algorithmic framework for approximating maximin share allocation of chores. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, 2021.

[6] Xin Huang and Erel Segal-Halevi. A reduction from chores allocation to job scheduling. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*, 2023.

[7] 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.

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