Human review packet: LMMS04FairDivision

Generated from byte-pinned source excerpts, the expanded PaperInterface specifications, and hash-bound raw-source-to-Spec screening records.

Paper: LMMS04FairDivision
Paper id: LMMS04FairDivision
Source version: ACM EC 2004 line-numbered canonical text extraction
Source record: audit/paper_statement_map.json
Rows: 6
Generated: 2026-09-06
Source URL: official source

How to use this packet

For each source claim, compare the exact source excerpts with the expanded PaperInterface specification. Mark up this PDF or fill the blank reviewer box in the TeX source; return the annotated file for reconciliation into the paper-local human-review log.

Open the interactive dashboard

The interactive dashboard is an optional alternative to using this PDF; it presents the same review surface rather than an additional review requirement. After forking and cloning (or pulling) AppliedModelingLib, run the following from the repository root. The cache command is needed only the first time in a new clone.

```
cd <your-repository-checkout>
lake exe cache get
python3 scripts/review_dashboard.py --paper LMMS04FairDivision --serve
```

Open http://127.0.0.1:8765/ in a browser and keep that terminal running while reviewing. The dashboard records saved annotations in this paper's local reviewer-owned review record.

Contents

Source claims

Review row 1

Source locator: cited publication, lines 170-174

Verbatim source input

Lemma 2.2. For any partial allocation A with envy graph
G, we can find another partial allocation B with envy graph
H such that:
• e(B) ≤ e(A)
• H is acyclic.

Expanded PaperInterface specification

```
∀ {Agent : Type u_3} {Item : Type u_4} [Fintype Agent] [Fintype Item] [DecidableEq Agent] [inst : DecidableEq Item]
  [Nonempty Agent] [Nonempty Item] (v : AppliedModelingLib.FairDivision.Valuation Agent Item) {alpha : ℝ}
  (goods : Finset Item) (A : AppliedModelingLib.FairDivision.Allocation Agent Item),
  LMMS04FairDivision.ProofInterface.isAllocationOf A goods →
    LMMS04FairDivision.ProofInterface.envyBoundedBy v A alpha →
      ∃ (B : AppliedModelingLib.FairDivision.Allocation Agent Item),
        LMMS04FairDivision.ProofInterface.isAllocationOf B goods ∧
          LMMS04FairDivision.ProofInterface.envyBoundedBy v B alpha ∧
            AppliedModelingLib.FairDivision.AcyclicEnvyGraph v B
```

Lean proof endpoint (built separately)

LMMS04FairDivision.PaperInterface.lemma2_2_acyclic_reductionSpec_proof

Recorded source-to-Spec screening **Verdict:** matches

Reason: The source lemma takes a partial allocation and returns an allocation of the same assigned goods with no larger maximum envy and an acyclic envy graph. The expanded target represents the partial assigned set explicitly as ‘goods‘; from any stated envy bound ‘alpha‘, it preserves that bound, preserves ‘isAllocationOf B goods‘, and concludes acyclicity. Instantiating ‘alpha‘ by the input maximum envy recovers the printed inequality.

Reviewer check: ☐ Matches source input

If left unchecked, explain the discrepancy or uncertainty below.

Reviewer annotation

Review row 2

Source locator: cited publication, lines 161-169

Verbatim source input

Theorem 2.1. For any set of goods and any set of players, there exists an allocation A such that the maximum envy
of A is bounded by the maximum marginal utility of the
goods, α. Furthermore, given oracle access for the utility
functions of the players, there is an O(mn3) time algorithm
for finding such an allocation.
Given an allocation A, we define the envy graph of A as
follows: every node of the graph represents a player and
there is a directed edge from p to q iff p envies q. The proof
of Theorem 2.1 is based on the following Lemma:

Expanded PaperInterface specification

```
∀ {Agent : Type u_3} {Item : Type u_4} [inst : Fintype Agent] [inst_1 : Fintype Item] [DecidableEq Agent]
  [inst_3 : DecidableEq Item] [inst_4 : Nonempty Agent] [inst_5 : Nonempty Item]
  (v : AppliedModelingLib.FairDivision.Valuation Agent Item) (goods : Finset Item),
  ∃ (A : AppliedModelingLib.FairDivision.Allocation Agent Item),
    LMMS04FairDivision.ProofInterface.isAllocationOf A goods ∧
      LMMS04FairDivision.ProofInterface.envyBoundedBy v A (LMMS04FairDivision.ProofInterface.maxMarginal v)
```

Lean proof endpoint (built separately)

LMMS04FairDivision.PaperInterface.theorem2_1_bounded_envy_allocation_existsSpec_proof

Recorded source-to-Spec screening **Verdict:** matches

Reason: This target is exactly the existence clause of printed Theorem 2.1: for a finite nonempty player/good model and specified finite goods it produces an allocation of those goods whose envy is bounded by the maximum marginal value. The separately inventoried oracle-time claim is an explicit partial-formalization boundary and is not represented by this existence-only Spec.

Reviewer check: ☐ Matches source input

If left unchecked, explain the discrepancy or uncertainty below.

Reviewer annotation

Review row 3

Source locator: cited publication, lines 313-315

Verbatim source input

Theorem 2.3. When the utilities of the players are probability measures on (Ω, F) = ([0, 1], Borel sets) with atoms
of value at most α, there exists a partition A = (A1 , ..., An)
of Ω such that e(A) ≤ α.

Expanded PaperInterface specification

```
∀ {Agent : Type u_3} [inst : Fintype Agent] [DecidableEq Agent] [Nonempty Agent] {mu : Agent → MeasureTheory.Measure ℝ}
  [inst_3 : ∀ (agent : Agent), MeasureTheory.IsFiniteMeasure (mu agent)] {alpha a b : ℝ} (halpha_pos : 0 < alpha)
  (hatoms : ∀ (agent : Agent), LMMS04FairDivision.Lemma24.PaperAtomsBoundedBy (mu agent) alpha)
  (haggregate_support : (AppliedModelingLib.FairDivision.aggregateMeasure mu).real (Set.Ioc a b)ᶜ = 0),
  ∃ (H : Finset ℝ) (P :
    AppliedModelingLib.Probability.RealIntervalPartition
      ((AppliedModelingLib.FairDivision.aggregateMeasure mu).restrict (↑H)ᶜ) alpha a b)
    (A : AppliedModelingLib.FairDivision.Allocation Agent (↥H ⊕ Option P.Piece)),
    AppliedModelingLib.FairDivision.IsAllocationOf A Finset.univ ∧
      AppliedModelingLib.FairDivision.EnvyBoundedBy
        (AppliedModelingLib.FairDivision.measurePartitionCertificateOfFiniteSingletonsAndResidualAggregate mu H
            (LMMS04FairDivision.Lemma24.realIntervalResidualSet H P)
            (LMMS04FairDivision.PaperInterface.theorem2_3_real_interval_supported_atom_boundSpec._proof_2 H P)
            (LMMS04FairDivision.PaperInterface.theorem2_3_real_interval_supported_atom_boundSpec._proof_3 H P)
            (LMMS04FairDivision.PaperInterface.theorem2_3_real_interval_supported_atom_boundSpec._proof_4 H P)
            (LMMS04FairDivision.PaperInterface.theorem2_3_real_interval_supported_atom_boundSpec._proof_5 halpha_pos
              hatoms H)
            (LMMS04FairDivision.PaperInterface.theorem2_3_real_interval_supported_atom_boundSpec._proof_6 halpha_pos
              haggregate_support H P)).valuation
        A alpha
```

Lean proof endpoint (built separately)

LMMS04FairDivision.PaperInterface.theorem2_3_real_interval_supported_atom_boundSpec_proof

Recorded source-to-Spec screening **Verdict:** matches

Reason: The source's positive-atom branch reduces probability measures on the real interval to finitely many indivisible pieces of value at most 'alpha' and then obtains an allocation with envy at most 'alpha'. The expanded target makes that reduction explicit: finite measures, a positive atom bound, aggregate interval support, a finite high-point set, a residual interval partition, and the resulting finite allocation. Probability measures on '[0,1]' are a specialization of this finite-measure construction; the source's separate alpha-zero/cake-cutting route is retained as an explicit boundary rather than silently credited here.

Reviewer check: ☐ Matches source input

If left unchecked, explain the discrepancy or uncertainty below.

Reviewer annotation

Review row 4

Source locator: cited publication, lines 770-773

Verbatim source input

Theorem 4.1. Any mechanism that returns an allocation
with minimum possible envy cannot be truthful. The same is
true for any mechanism that returns an envy-free allocation
whenever there exists one.

Expanded PaperInterface specification

```
∀ (M : LMMS04FairDivision.Theorem41.LMMS41Mechanism),
  AppliedModelingLib.FairDivision.ReturnsAllocationOf M LMMS04FairDivision.Theorem41.lmms41Goods →
    AppliedModelingLib.FairDivision.ReturnsEnvyFreeWheneverExists M LMMS04FairDivision.Theorem41.lmms41Goods →
      ¬AppliedModelingLib.FairDivision.DirectFairDivisionMechanism.Truthful M
```

Lean proof endpoint (built separately)

```
LMMS04FairDivision.PaperInterface.theorem4_1_source_not_truthful_envy_free_whenever_existsSpec_proof
```

Recorded source-to-Spec screening **Verdict:** matches

Reason: The target is the envy-free-when-possible clause of Theorem 4.1 on the concrete two-player finite source counterexample universe: a mechanism allocating those goods and returning an envy-free allocation whenever one exists cannot be truthful. The printed proof proves the same impossibility by a finite egg construction; Lean verifies the same counterexample structure with an eight-egg witness rather than the source's illustrative one-hundred-egg count.

Reviewer check: ☐ Matches source input

If left unchecked, explain the discrepancy or uncertainty below.

Reviewer annotation

Review row 5

Source locator: cited publication, lines 770-773

Verbatim source input

Theorem 4.1. Any mechanism that returns an allocation
with minimum possible envy cannot be truthful. The same is
true for any mechanism that returns an envy-free allocation
whenever there exists one.

Expanded PaperInterface specification

```
∀ (M : LMMS04FairDivision.Theorem41.LMMS41Mechanism),
  AppliedModelingLib.FairDivision.ReturnsMinimumReportEnvy M LMMS04FairDivision.Theorem41.lmms41Goods →
    ¬AppliedModelingLib.FairDivision.DirectFairDivisionMechanism.Truthful M
```

Lean proof endpoint (built separately)

LMMS04FairDivision.PaperInterface.theorem4_1_source_minimum_envy_not_truthfulSpec_proof

Recorded source-to-Spec screening **Verdict:** matches

Reason: The target states the minimum-envy clause on the concrete two-player finite source counterexample universe and concludes that every mechanism satisfying that allocation rule is not dominant-strategy truthful. The printed proof establishes its universal impossibility through such a finite egg witness; Lean uses a smaller eight-egg witness with the same additive values and profitable-deviation argument, which is a proof-witness reduction and does not alter the advertised impossibility conclusion.

Reviewer check: ☐ Matches source input

If left unchecked, explain the discrepancy or uncertainty below.

Reviewer annotation

Review row 6

Source locator: cited publication, lines 764-776

Verbatim source input

Theorem 4.2. Suppose that vp,i ≤ α ∀p ∈ N, j ∈ M .
Then for every [U+000F control character] > 0, and for large enough n, there exists
a truthful algorithm such that with high probability the allocation
√ output by the algorithm has maximum envy at most
O(α n1/2+[U+000F control character]).

Theorem 4.1. Any mechanism that returns an allocation
with minimum possible envy cannot be truthful. The same is
true for any mechanism that returns an envy-free allocation
whenever there exists one.

Proof. The proof is based on the probabilistic method.
Allocate each good independently to player p with probability 1/n. Clearly this is a truthful mechanism. We will show

Expanded PaperInterface specification

```
∀ {Agent : Type u_3} {Item : Type u_4} [inst : Fintype Agent] [inst_1 : Fintype Item] [inst_2 : DecidableEq Agent]
  [inst_3 : DecidableEq Item] [inst_4 : Nonempty Agent] [Nonempty Item]
  [inst_6 : Fintype (AppliedModelingLib.FairDivision.Allocation Agent Item)]
  [inst_7 : DecidableEq (AppliedModelingLib.FairDivision.Allocation Agent Item)],
  LMMS04FairDivision.ProofInterface.randomizedTruthful
    (LMMS04FairDivision.Theorem42.lmms42UniformRandomMechanism Agent Item)
```

Lean proof endpoint (built separately)

LMMS04FairDivision.PaperInterface.theorem4_2_uniform_random_mechanism_truthfulSpec_proof

Recorded source-to-Spec screening **Verdict:** matches

Reason: The selected source atom is the proof's assertion that independently assigning every good uniformly to a player is truthful. The target quantifies over finite nonempty agent and item types and proves expected-utility truthfulness of precisely that uniform randomized mechanism. The source theorem's separate high-probability asymptotic envy-rate conclusion is explicitly retained as an unproved boundary.

Reviewer check: ☐ Matches source input

If left unchecked, explain the discrepancy or uncertainty below.

Reviewer annotation