

PRELIMINARY VERSION

14th October 2026

*Comments are welcome. Most recent version, available at
<https://cgi.cse.unsw.edu.au/~haziz/BOBW.pdf>.*

From Subsidies to Lotteries: A Payment Bridge for Best-of-Both-Worlds Fairness

Haris Aziz¹ Xinhang Lu² Simon Mackenzie¹ Mashbat Suzuki¹

¹UNSW Sydney, Australia ²Kyushu University, Japan

{haris.aziz, simon.w.mackenzie, mashbat.suzuki}@unsw.edu.au
xinhang.lu@inf.kyushu-u.ac.jp

Abstract

Payments and randomization are two ways to remove envy when indivisible items are shared. We give a general payment bridge that transfers both existence and polynomial-time computability from allocations with subsidies to lotteries without payments. If, at every positive rescaling of the agents' utilities, some allocation with a scale-invariant property such as EF1 is envy-freeable, then an ex-ante envy-free lottery preserves that property in every outcome. Under standard access and rational-encoding assumptions, the bridge converts any polynomial-time deterministic algorithm for such allocations into a polynomial-time algorithm that outputs at most $n(n-1)$ allocations and their exact rational probabilities, where n is the number of agents. Applications give shorter proofs and explicit short lotteries for known best-of-both-worlds fairness guarantees, together with new results for unit-demand, transformed matroid-rank, and submodular valuations.

Contents

1	Introduction	1
2	Preliminaries	3
3	The payment bridge	4
3.1	Weighted tests	4
3.2	Envy tests are rescaled reassignments	5
3.3	An exact criterion	6
3.4	The payment bridge	7
3.5	Common payments and simultaneous benchmarks	8
3.6	The algorithmic reduction	9
4	Applications of certificates and weighted responses	9
4.1	Payment certificates	10
4.2	Common payments for submodular shares	11
4.3	Direct weighted responses	13
5	Discussion	13
	Guide to the appendices	15
A	Foundations: separation, support, and algorithms	15
A.1	Separation and support	15
A.2	A balancing example	16
A.3	Common-payment cancellation	16
A.4	Computing the lottery	17
A.5	Weighted comparisons and their bridge	18
B	Deterministic payment certificates	20
B.1	Unit-demand valuations	20
B.2	Multi-demand valuations	22
B.3	Transformed matroid-rank valuations	23
B.4	Fair team assignments that respect participants' choices	26
C	Nonadditive applications	29
C.1	The full submodular share curve	29
C.2	Subadditive valuations: fair chances without additivity	33
C.3	Expected envy and realized maximin shares for submodular goods	37
D	Limits of deterministic payment certificates	43
D.1	The deterministic obstruction	44
D.2	A fair lottery that avoids the certificate	45
D.3	What the example establishes about the bridge	46

AI methodologies	47
References	47

1 Introduction

Envy can be unavoidable when indivisible items are shared. If two agents each value a single item at one, the agent without it envies the other. There are two classical remedies. A lottery gives each agent the item with probability $1/2$, so neither agent envies in expectation. A payment gives the item to one agent and pays the other one unit, so neither agent envies after payments. Each remedy has its own literature. *Best-of-both-worlds (BoBW)* fairness seeks lotteries that are ex-ante envy-free and whose outcomes are envy-free up to one item (EF1) [FSV20, AFSV24]. Known BoBW lotteries are usually built by randomized algorithms designed for their domains. Fair division with subsidy asks when payments can remove the envy of an allocation, which is then called *envy-freeable*, and how small those payments can be [HS19, BDN⁺20].

We give a payment bridge that transfers both existence and efficient computation from the subsidy problem to the lottery problem: *if the envy of some EF1 allocation can be paid away at every utility scale, then envy can be randomized away without payments, keeping EF1 in every outcome.*

Payment bridge (Theorems 3.5 and 3.8). If every independent positive rescaling of a valuation profile admits an allocation that is both EF1 and envy-freeable, then the profile admits an ex-ante envy-free lottery supported on EF1 allocations. The lottery makes no payments, and no bound on the payments is needed. Under the stated access and rational-encoding assumptions, any polynomial-time deterministic algorithm finding such allocations at requested rescalings yields an explicit lottery with exact rational probabilities and at most $n(n - 1)$ outcomes in polynomial time.

We call an EF1 allocation that is envy-freeable at a given rescaling a **payment certificate** for that rescaling. The proof exploits an asymmetry. EF1 is unchanged when an agent’s utility is multiplied by a positive constant. Envy-freeability is not, because payments keep coefficient one. Suppose that no lottery over EF1 allocations is ex-ante envy-free. A separating hyperplane then gives a weighted envy test that every EF1 allocation fails. We derive utility units from the stationary distribution of a random walk on the agents defined by the test weights. At those units, the weighted sum of the payment inequalities cancels every payment, so an envy-freeable EF1 allocation passes the test, a contradiction. The units depend on the test, so the premise is needed at every scale. Figure 1 illustrates the response loop behind the construction. Theorem 3.8 gives the polynomial-time reduction, using polynomially many calls to the deterministic certificate algorithm at chosen rescalings.

The payment bridge is a simple corollary of a more general characterization result that we prove. We characterize exactly when a prescribed family of allocations, such as the EF1 allocations, supports an ex-ante envy-free lottery: at every rescaling of the utilities, some lottery over the family must be envy-freeable in expectation (Theorem 3.4). A payment certificate is the special case of a lottery with a single outcome, so the bridge follows at once.

Contributions.

- *A general algorithmic bridge.* Theorem 3.8 gives a polynomial-time reduction from deterministic payment-certificate algorithms to explicit ex-ante envy-free lotteries with exact rational probabilities and at most $n(n - 1)$ outcomes, under the stated access and encoding assumptions. The underlying existence theorem (Theorem 3.5) applies to every scale-invariant realized property and requires neither additivity nor monotonicity. The bridge also extends to unequal entitlements (Theorem A.3).

Domain	Ex ante	Ex post	Status	Result
<i>Payment bridge with deterministic certificates (Sections 3 and 4.1)</i>				
Additive goods	EF	Balanced EF1	Recovered [FSV20]; poly.	Cor. 3.6
Additive mixed manna	EF	EF1	Recovered [ABL ⁺ 26]; poly (New)	Cor. 3.6
Unit demand	EF	EFX ₀ , groupwise MMS, undominated by lotteries	New; poly.	Thm. 4.1
Multi-demand	EF	Balanced EF1	Recovered [HSV24]; poly.	Cor. 4.2
Transformed matroid ranks	EF	EF1, also PO for strict transformations	New; poly.	Thm. 4.3
Two-sided teams	EF for teams	Balanced EF1, no beneficial cyclic exchange	New; poly.	Thm. B.3
<i>Common payments (Section 4.2)</i>				
Submodular	PROP, MEF, $\frac{n}{2n-1}$ -MES	Balanced MEF1, hence PROP1 and $\frac{1}{2}$ -EF1	New	Thm. 4.4
<i>Direct weighted responses (Section 4.3)</i>				
Subadditive	$\frac{1}{2}$ -EF	EF1 and $\frac{1}{2}$ -EFX	Known [FMNP24]; explicit, poly.	Cor. 4.6
Submodular	$\frac{1}{2}$ -EF	EFL and $\frac{1}{8}$ -MMS	New	Thm. 4.7

Table 1: Principal lottery guarantees, grouped by method, for arbitrary numbers of agents unless stated, with equal entitlements. Recovered results follow from subsidy theorems through the bridge. “Poly.” means that the lottery is computable in polynomial time in the stated access model. Mixed manna means goods and chores. PROP is proportionality, MEF is marginal envy-freeness, and MES is the maximum expectation share. Teams are described in Section 4.1.

- *Characterization of ex-ante envy-freeness.* Every positive weighted envy test is a random reassignment of bundles, measured at suitable utility scales. So a family supports an ex-ante envy-free lottery if and only if each rescaled reassignment is defeated by some member; the bridge asks one allocation to defeat them all. Equivalently, at every scale some lottery over the family makes the identity a maximum-weight matching of agents to random bundles, randomizing the characterization of envy-freeability in [HS19] (Theorem 3.4).
- *Applications* (Table 1). Known additive lotteries for goods and for mixed manna [FSV20, AFSV24, ABL⁺26] follow in a few lines from deterministic subsidy theorems. In fact for mixed manna, we also establish polynomial-time computability. The subadditive guarantees of [FMNP24] receive an explicit polynomial-time lottery on at most $n(n-1)$ outcomes. As a consequence, our general tools provide a unified way to uncover several of the established BoBW results in the literature. We also obtain several new results including new lotteries cover unit demand, transformed matroid ranks, teams, and submodular goods. Their outcomes combine fairness with efficiency or shares, such as EFX₀ with groupwise MMS for unit demand, and EFL with 1/8-MMS for submodular goods. For heterogeneous multi-demand valuations, capped iterated matching supplies balanced EF1 payment certificates at every scale, recovering the BoBW guarantee of [HSV24].

Related work. Exact ex-ante EF with EF1 outcomes for additive goods is due to [FSV20, AFSV24]. Subsequent work treats binary valuations [Azi20, HPPS20], entitlements [HSV24, AGM23], mixed resources [BLL⁺26], efficiency and incentives [SC25, GSX⁺23], randomized envy-cycle elimination [GS24], mixed goods and chores [ABL⁺26], weighted mixed manna [LLTZ26], and

shares [BEF22, BF26]. Our bridge replaces a domain-specific randomized construction by a payment certificate at every rescaling.

Halpern and Shah characterize envy-freeability through envy cycles [HS19]. Subsidy bounds and algorithms appear in [BDN⁺20, CI21, BKNS22, KMS⁺24, GIK⁺24, LMS26, KMNP25]. That literature bounds the payments. The bridge uses only their existence, so a subsidy theorem that retains a scale-invariant property at every rescaling yields a BoBW lottery. For example, together with Theorem 3.5, the iterated matching theorem of [BDN⁺20] yields an ex-ante EF lottery with EF1 outcomes, as in [FSV20]. Lotteries whose outcomes include payments are studied for interim envy-freeness in [CKK25]. Here payments serve only as a certificate. For matroid ranks, random priority already has Lorenz-dominating and ex-ante EF support [BEF21]. Further deterministic results for ranks and teams appear in [BCIZ20, IKSS24]. Reassignment-based EF and efficient lotteries are studied in [CT21, CHR24]. Our characterization allows a prescribed family that need not be closed under permuting bundles among agents.

For subadditive valuations, Feldman et al. [FMNP24] obtain expected 1/2-EF with EF1 and 1/2-EFX support. EFX with charity and expected 1/2-EF for monotone valuations is studied in [KPV⁺25]. Marginal EF1 is due to [CKM⁺16]. The maximum expectation share is formulated in [FG25, BF24]. Feige and Fine [FF26] combine EFL with approximate MMS for submodular valuations.

Convex separation and response-to-mixture algorithms are classical [BV04, Bla56, PR08, JLB15, FP24]. The new ingredients here are the payment cancellation and the domain-specific certificates or responses. Further context appears in [BCL16, LLSW24].

Organization. Section 3 develops the exact characterization and the algorithmic payment bridge, Section 4 its applications, and Section 5 its limits. The appendices give the proofs, and a short guide at their start locates each one.

2 Preliminaries

Model. There are $n \geq 2$ agents $N = [n]$ and a finite set M of indivisible items, where $[n] = \{1, \dots, n\}$ and $m = |M|$. An **allocation** $A = (A_1, \dots, A_n)$ partitions M . Agent i has a valuation $v_i : 2^M \rightarrow \mathbb{R}$ with $v_i(\emptyset) = 0$. Unless stated otherwise, valuations need be neither additive nor monotone. Write $e_{ij}(A) = v_i(A_j) - v_i(A_i)$ for $i \neq j$, and $e_{ii}(A) = 0$. A positive value means that i envies j . The profile $s \cdot v$ has valuation $s_i v_i$ for each agent, with $s_i > 0$. Unequal entitlements satisfy $w_i > 0$ and $\sum_i w_i = 1$. We keep four kinds of numbers apart: utility scales s , entitlements w , payments p , and the test coefficients λ of Section 3, which are not lottery probabilities.

Fairness of an allocation. An allocation is **envy-free** (EF) if $e_{ij}(A) \leq 0$ for all i, j . It is **envy-free up to one item** (EF1) [LMMS04, Bud11] if, for every $i \neq j$, either there is no envy or some $a \in A_j$ satisfies $v_i(A_i) \geq v_i(A_j \setminus \{a\})$. For signed valuations we use the two-sided convention [ACIW22]: either there is no envy, or some $a \in A_i \cup A_j$ satisfies $v_i(A_i \setminus \{a\}) \geq v_i(A_j \setminus \{a\})$. The conventions agree for monotone goods valuations. EF1 is **scale invariant**: multiplying an agent's valuation by a positive constant preserves each comparison and its deletion witness.

Payments and lotteries. An allocation is **envy-freeable** [HS19] if there are payments $p_1, \dots, p_n \in \mathbb{R}$ with

$$v_i(A_i) + p_i \geq v_i(A_j) + p_j \quad (i \neq j), \quad (2.1)$$

or equivalently $e_{ij}(A) \leq p_i - p_j$. Payments attach to the bundles of this allocation. Adding a common constant changes no comparison, so payments may be nonnegative or centered to sum to zero. They need not be bounded. Envy-freeness is not scale invariant, because payments retain coefficient one as values rescale. By [HS19], A is envy-freeable if and only if no permutation of its bundles among the agents increases $\sum_i v_i(A_i)$ (Lemma 3.2). A lottery is **ex-ante envy-free** if

$$\mathbb{E}[v_i(A_i)] \geq \mathbb{E}[v_i(A_j)] \quad (i \neq j), \quad (2.2)$$

with both expectations taken over the same lottery. For nonadditive valuations, item marginals do not determine these expectations. A lottery is **ex-post EF1** if every allocation of positive probability is EF1.

Other outcome requirements. An allocation is **balanced** if bundle sizes differ by at most one. It is **Pareto optimal** (PO), or integrally Pareto optimal, if no complete allocation weakly increases every agent's own value and strictly increases one. For additive valuations, **fractional Pareto optimality** (fPO) also excludes improvements by fractional allocations, with utilities extended linearly to item fractions.

3 The payment bridge

We characterize when a family of allocations supports an ex-ante envy-free lottery, and obtain the payment bridge as the deterministic special case. Its computational form reduces explicit lottery construction to finding deterministic payment certificates in polynomial time (Theorem 3.8). The key facts are that every positive envy test is a rescaled random reassignment of bundles, and payments cancel against it.

3.1 Weighted tests

Ex-ante envy-freeness asks that the envy $e_{ij}(A) = v_i(A_j) - v_i(A_i)$ be nonpositive in expectation for every ordered agent pair (i, j) . In general, fix a finite nonempty outcome set X , a nonempty **permitted family** $S \subseteq X$ (e.g., the EF1 allocations), and $d \geq 1$ functions $g_k : X \rightarrow \mathbb{R}$. The value $g_k(x)$ is the **signed violation** of requirement k ; a positive value means x violates it. For envy-freeness, $g_{ij} = e_{ij}$ and $d = n(n-1)$. We seek a lottery on S with $\mathbb{E}[g_k] \leq 0$ for all k .

No single outcome need meet every requirement, so we ask for less. A vector $\lambda \in \mathbb{R}_{\geq 0}^d$ defines the **weighted test** $\lambda^\top g(x) \leq 0$; a permitted outcome passing it is a **response**, and different tests may have different responses. A fair lottery yields responses: $\mathbb{E}[\lambda^\top g] \leq 0$, so some outcome in its support passes. Indeed, the lottery meets each requirement on average and the weights are nonnegative, so it also meets their weighted sum on average; an average that is at most zero must include an outcome that is at most zero. Theorem 3.1 (LP duality; Appendix A) gives the converse: the only obstruction to a fair lottery is a weighted test that every permitted outcome fails. The g_k need no structure; duality applies because $\mathbb{E}[g_k]$ is linear in the lottery.

Theorem 3.1 (Weighted-test criterion). *There is a lottery on S satisfying $\mathbb{E}[g_k] \leq 0$ for every k if and only if*

$$\text{for every } \lambda \in \mathbb{R}_{\geq 0}^d \text{ there is } x \in S \text{ with } \lambda^\top g(x) \leq 0. \quad (3.1)$$

Whenever such a lottery exists, one exists with at most d supported outcomes.

Theorem 3.1 can also be viewed as a helpful formulation of the minimax theorem. For envy-freeness, consider the zero-sum game in which a designer picks an allocation $A \in S$ and an auditor picks an ordered pair (i, j) and wins the envy $e_{ij}(A)$ from the designer. A lottery is ex-ante envy-free exactly when the auditor's best response to it wins at most zero. Against a mixed audit with weights λ_{ij} on pairs, the designer's best response is an allocation in S minimizing $\sum_{i \neq j} \lambda_{ij} e_{ij}(A)$. Condition (3.1) says this minimum is at most zero, so the best response is a response in our sense. By the minimax theorem, the designer can hold the auditor to at most zero by committing to a lottery first exactly when her best response to every mixed audit does so.

3.2 Envy tests are rescaled reassignments

Our responses come from fair division with subsidies. Unlike EF1, envy-freeability depends on how each agent's utility is scaled, so we consider it at scaled utilities $s \cdot v = (s_1 v_1, \dots, s_n v_n)$ with $s \in \mathbb{R}_{>0}^n$; the scales will be chosen by the test λ . The payments serve only as a certificate. They cancel in the analysis, and no money changes hands in our lotteries.

A reassignment of the bundles of A is described by a doubly stochastic matrix P , where P_{ij} is the probability that agent i receives bundle A_j . By the Birkhoff–von Neumann theorem, P is a lottery over permutations. Its expected gain in scaled welfare is

$$T_A(s, P) = \sum_{i,j} s_i P_{ij} e_{ij}(A). \quad (3.2)$$

If A is envy-freeable at $s \cdot v$, no reassignment of its bundles has positive gain. If a permutation σ gives agent i the bundle $A_{\sigma(i)}$, envy-freeability gives $s_i e_{i\sigma(i)}(A) \leq p_i - p_{\sigma(i)}$. Summing over i , the right side vanishes, since each p_i appears once with each sign: payments travel with their bundles, so a reassignment only reshuffles the money. The lemma below proves the converse, so payments can be replaced by payment-free inequalities of the same form as a weighted test. We state it for any matrix with zero diagonal so that it also covers expected and weighted envy.

Lemma 3.2 (Payments cancel under reassignment). *Let E be a real $n \times n$ matrix with zero diagonal, and let $s \in \mathbb{R}_{>0}^n$. The following are equivalent.*

- (1) *Some $p \in \mathbb{R}^n$ satisfies $s_i E_{ij} \leq p_i - p_j$ for all $i \neq j$.*
- (2) *$\sum_{k=1}^r s_{i_k} E_{i_k i_{k+1}} \leq 0$ for every directed cycle of distinct agents $i_1, \dots, i_r$, where $r \geq 2$ and $i_{r+1} = i_1$.*
- (3) *$\sum_i s_i E_{i\sigma(i)} \leq 0$ for every permutation σ .*
- (4) *$\sum_{i,j} s_i P_{ij} E_{ij} \leq 0$ for every doubly stochastic P .*

In particular, A is envy-freeable at $s \cdot v$ if and only if $T_A(s, P) \leq 0$ for every doubly stochastic P . With $s = \mathbf{1}$ and $E = e(A)$, the equivalence of (1) and (3) is the characterization of [HS19]: A is envy-freeable if and only if no permutation of its bundles increases $\sum_i v_i(A_i)$.

Proof. If (1) holds, multiply each inequality by $P_{ij} \geq 0$ and sum. The right side is $\sum_i p_i \sum_j P_{ij} - \sum_j p_j \sum_i P_{ij} = 0$, because every row and column of P sums to one. This proves (4), which implies (3) because permutation matrices are doubly stochastic. A directed cycle, extended by fixed points, is a permutation whose fixed points contribute zero, so (3) implies (2). For (2) $\Rightarrow$ (1), give the edge $i \rightarrow j$ weight $s_i E_{ij}$ and let p_i be the maximum weight of a directed walk starting at i , including the empty walk of weight zero. Removing a cycle from a walk cannot decrease its weight, so the maximum is attained by one of the finitely many simple paths. Prepending the edge $i \rightarrow j$ to a maximizing walk from j gives $p_i \geq s_i E_{ij} + p_j$. $\square$

Condition (3) compares A only with permutations of its own bundles, not with allocations that regroup the items. So an allocation that is envy-freeable at $s \cdot v$ passes every test of the form $\lambda_{ij} = s_i P_{ij}$. The next lemma shows that every strictly positive test has this form at suitable scales s , which is why the bridge needs a certificate at every scale.

Lemma 3.3 (Envy tests are rescaled reassignments). *For every strictly positive vector $(\lambda_{ij})_{i \neq j}$ there are scales $s \in \mathbb{R}_{>0}^n$ and a doubly stochastic matrix P with strictly positive entries such that $s_i P_{ij} = \lambda_{ij}$ for all $i \neq j$. Hence $\sum_{i \neq j} \lambda_{ij} g_{ij} = \sum_{i,j} s_i P_{ij} g_{ij}$ for every array g with zero diagonal; in particular $\sum_{i \neq j} \lambda_{ij} e_{ij}(A) = T_A(s, P)$. Conversely, every pair (s, P) defines the test $\lambda_{ij} = s_i P_{ij}$. For rational λ , s and P are computable exactly in polynomial time.*

Proof. Let $r_i = \sum_{j \neq i} \lambda_{ij}$, and let π be the stationary distribution of the random walk on agents that moves from i to $j \neq i$ with probability λ_{ij}/r_i . It is strictly positive because every transition is. Put $s_i = r_i/\pi_i$, $P_{ij} = \pi_i \lambda_{ij}/r_i$ for $j \neq i$, and $P_{ii} = 1 - \pi_i$. Then $s_i P_{ij} = \lambda_{ij}$ for $j \neq i$, and the diagonal terms vanish because $g_{ii} = 0$. Row i of P sums to $\pi_i + (1 - \pi_i) = 1$, and stationarity makes column j sum to $\pi_j + (1 - \pi_j) = 1$. Since $n \geq 2$, every $\pi_i < 1$, so P is strictly positive. For rational λ , π solves a rational linear system, which gives the encoding and running-time bounds. $\square$

These lemmas are the entire payment-cancellation mechanism: if A is envy-freeable at the scales of Lemma 3.3, then $\sum_{i \neq j} \lambda_{ij} e_{ij}(A) = T_A(s, P) \leq 0$. The scales come from the test, and payment magnitudes never matter (Appendix A.2 has an example).

3.3 An exact criterion

We now present a general characterization of ex-ante envy-freeness. It extends the characterization of envy-freeability in [HS19] from one allocation to lotteries over a prescribed family $\mathcal{F}$, which need not be closed under permuting bundles. Condition (iii) is their maximum-weight matching condition for random bundles. Unlike theirs, the conditions range over all scales, and this removes the payments: for $\mathcal{F} = \{A\}$, the theorem says that A is envy-free if and only if it is envy-freeable at every scale.

Theorem 3.4 (Exact criterion). *Fix a valuation profile v and a finite nonempty family $\mathcal{F}$ of complete allocations. The following are equivalent.*

- (i) *There is an ex-ante envy-free lottery supported on $\mathcal{F}$.*
- (ii) *For every positive scale vector $s \in \mathbb{R}_{>0}^n$, there are a lottery μ_s on $\mathcal{F}$ and payments $p^s \in \mathbb{R}^n$ such that*

$$s_i \mathbb{E}_{\mu_s}[v_i(A_i)] + p_i^s \geq s_i \mathbb{E}_{\mu_s}[v_i(A_j)] + p_j^s \quad (i \neq j). \quad (3.3)$$

- (iii) *For every $s \in \mathbb{R}_{>0}^n$, there is a lottery μ_s on $\mathcal{F}$ for which the identity is a maximum-weight perfect matching between agents i and random bundles j , with edge weights $M_{ij}^s = s_i \mathbb{E}_{\mu_s}[v_i(A_j)]$; that is, $\sum_i M_{ii}^s \geq \sum_i M_{i\sigma(i)}^s$ for every permutation σ .*
- (iv) *For every $s \in \mathbb{R}_{>0}^n$ and every doubly stochastic matrix P , some $A \in \mathcal{F}$ satisfies $T_A(s, P) \leq 0$. It suffices to test strictly positive P .*

Whenever such a lottery exists, one exists on at most $n(n-1)$ allocations. If every member of $\mathcal{F}$ is EF1, this is an exact characterization of ex-ante EF together with ex-post EF1.

Proof. (i) $\Rightarrow$ (ii): use the fair lottery at every scale with $p^s = 0$. For the next two steps, fix s and a lottery μ_s , and apply Lemma 3.2 to the expected envy matrix $E_{ij} = \mathbb{E}_{\mu_s}[e_{ij}(A)]$. Its condition (1) is (ii), and its condition (3) is (iii), because $\sum_i M_{i\sigma(i)}^s - \sum_i M_{ii}^s = \sum_i s_i E_{i\sigma(i)}$; hence (ii) $\Leftrightarrow$ (iii). Under

(ii), its condition (4) gives $\mathbb{E}_{\mu_s}[T_A(s, P)] \leq 0$ for every doubly stochastic P , so some allocation in the support of μ_s has $T_A(s, P) \leq 0$, which is (iv).

(iv) $\Rightarrow$ (i): by Theorem 3.1 it suffices to answer each nonnegative envy test λ within $\mathcal{F}$. If λ is strictly positive, Lemma 3.3 writes its score as $T_A(s, P)$ with P strictly positive, and the allocation given by (iv) answers it. Otherwise, since $\mathcal{F}$ is finite, one allocation answers $\lambda + k^{-1}\mathbf{1}$ for infinitely many k , and hence λ in the limit. Theorem 3.1 also gives the support bound. $\square$

We call (μ_s, p^s) a **randomized payment certificate**. Condition (iv) asks for one response per test, condition (ii) for one randomized certificate per scale, and the payment bridge below for a deterministic one. For nonadditive valuations, M_{ij}^s is not the value of a fractional set, and no allocation in the support of μ_s need be envy-freeable.

3.4 The payment bridge

Let $\mathcal{P}$ be a property of complete allocations and let $S_{\mathcal{P}}(v)$ denote the allocations satisfying it at profile v . The property is **scale invariant** if $S_{\mathcal{P}}(t \cdot v) = S_{\mathcal{P}}(v)$ for every positive t . A valuation class $\mathcal{V}$ is **closed under rescaling** if $t \cdot v \in \mathcal{V}$ whenever $v \in \mathcal{V}$. Closure allows us to invoke a deterministic theorem after changing utility units. Invariance ensures that the returned allocation is still permitted at the original profile. An allocation that satisfies $\mathcal{P}$ and is envy-freeable at u is a **payment certificate** at u .

Theorem 3.5 (Payment bridge). *Let $\mathcal{V}$ be a class of normalized valuation profiles closed under positive rescaling, and let $\mathcal{P}$ be a scale-invariant property of complete allocations. Suppose every $u \in \mathcal{V}$ admits an allocation that satisfies $\mathcal{P}$ and is envy-freeable at u . Then every $v \in \mathcal{V}$ admits an ex-ante envy-free lottery supported on at most $n(n-1)$ allocations, all satisfying $\mathcal{P}$. The lottery makes no payments. The implication requires neither additivity nor monotonicity and uses no bound on the payments.*

Proof. Fix $v \in \mathcal{V}$ and let $\mathcal{F} = S_{\mathcal{P}}(v)$. This family is finite, and nonempty by the premise at $u = v$. Fix $s \in \mathbb{R}_{>0}^n$. The premise at $u = s \cdot v \in \mathcal{V}$ supplies an allocation A that satisfies $\mathcal{P}$ at $s \cdot v$ and is envy-freeable there. Scale invariance places A in $\mathcal{F}$. The point mass on A , with its payments, satisfies condition (ii) of Theorem 3.4, which gives the lottery and its support bound. $\square$

For a fixed family $\mathcal{F}$, write $\Delta(\mathcal{F})$ for the lotteries on $\mathcal{F}$. By Lemma 3.2, the bridge premise and conditions (ii) and (iv) of Theorem 3.4 differ only in the order of quantifiers:

$$\begin{aligned} \text{payment bridge premise: } & \forall s \exists A \in \mathcal{F} \forall P : T_A(s, P) \leq 0, \\ \text{condition (ii): } & \forall s \exists \mu \in \Delta(\mathcal{F}) \forall P : \mathbb{E}_{\mu}[T_A(s, P)] \leq 0, \\ \text{condition (iv): } & \forall s \forall P \exists A \in \mathcal{F} : T_A(s, P) \leq 0. \end{aligned}$$

The premise implies (ii), which is equivalent to (iv). Theorem 5.1 shows that the premise is strictly stronger, so the criterion is strictly more permissive than the payment bridge.

In the one-item example of the introduction, rescaling gives item values s_1, s_2 . Giving the item to an agent with the higher value and paying the other her value supplies an EF1, envy-freeable allocation at every rescaling. The bridge therefore recovers the fair lottery. The certificate must be a single allocation: separate allocations witnessing EF1 and envy-freeability do not meet the premise. Since EF1 is scale invariant, known subsidy theorems give the following.




Figure 1: The payment bridge (Theorem 3.5) as an algorithm; grey tags justify each step, and in step 2 the test weights are perturbed to be positive. Each round adds an EF1 allocation not yet in $\mathcal{S}_t$, so the loop stops, and only with a fair lottery. The polynomial-time reduction is given in Theorem 3.8. Right ($n = 2$): the line $\lambda^\top y = 0$ separates the hull K_t of current envy vectors from the nonpositive quadrant; the new envy vector crosses it.

Corollary 3.6 (Recovering additive EF–EF1 lotteries). *Every signed additive profile with equal entitlements admits an ex-ante envy-free lottery over at most $n(n - 1)$ complete EF1 allocations. For nonnegative additive valuations, the outcomes can also be balanced, and the lottery is computable in polynomial time for rational input.*

Proof. For signed additive valuations, [ALMS25, Theorem 4.2] supplies an EF1, envy-freeable allocation at every positive rescaling. Apply Theorem 3.5 with $\mathcal{P} = \text{EF1}$, using the two-sided deletion convention. For nonnegative additive valuations, the iterated matching theorem [BDN⁺20, Theorem 1.3] gives a balanced, EF1, envy-freeable allocation. Apply Theorem 3.5 with $\mathcal{P} = \text{EF1} + \text{balance}$. Its computational version is Theorem 3.8.¹ $\square$

Unequal entitlements. Define weighted envy by $e_{ij}^w(v, A) = v_i(A_j)/w_j - v_i(A_i)/w_i$. Weighted envy-freeability is $e_{ij}^w \leq r_i - r_j$, where $r_i = p_i/w_i$ is a normalized payment. The same balancing argument gives exact ex-ante WEF whenever the permitted family contains a weighted envy-freeable allocation at every rescaling. This extension requires neither additivity nor monotonicity (Theorem A.3).

3.5 Common payments and simultaneous benchmarks

The final extension is needed for our submodular result. In addition to pairwise comparisons, suppose agent i has individual benchmark violations b_{ik} , and one centered payment vector bounds all of them. Balancing must now make the residual coefficient of every p_i the same constant, rather than zero. Centering then cancels that constant.

Proposition 3.7 (Common-payment cancellation). *Fix v and a finite nonempty family $\mathcal{F}$ of complete allocations, held fixed while v is rescaled. If $\mathcal{F}$ is specified by an outcome property, that property must be scale invariant. Let a_{ij} ($i \neq j$) and b_{ik} ($k = 1, \dots, K$) be homogeneous in agent i 's*

¹In the extended version of [ALMS25], we also prove that an envy-freeable and EF1 allocation can be computed in polynomial time for mixed goods and chores so the BoBW result can be achieved in polynomial time.

valuation:

$$a_{ij}(s \cdot v, A) = s_i a_{ij}(v, A), \quad b_{ik}(s \cdot v, A) = s_i b_{ik}(v, A).$$

Fix $c_1, \dots, c_K > 0$. Suppose that for every positive scale vector s there are $A \in \mathcal{F}$ and centered payments $\sum_i p_i = 0$ satisfying

$$a_{ij}(s \cdot v, A) \leq p_i - p_j, \quad b_{ik}(s \cdot v, A) \leq c_k p_i.$$

Then some lottery on $\mathcal{F}$ satisfies $\mathbb{E}[a_{ij}(v, A)] \leq 0$ and $\mathbb{E}[b_{ik}(v, A)] \leq 0$ for every i, j, k .

The proof in Appendix A.3 balances one combined test, so a single lottery meets every comparison and benchmark. Applying the bridge to each inequality separately would not ensure this. With no individual benchmarks, we recover the payment bridge.

3.6 The algorithmic reduction

The computational bridge takes a deterministic payment-certificate algorithm as a subroutine and returns an explicit fair lottery. Under the assumptions below, it makes polynomially many calls at positive rational rescalings and outputs the supported allocations and their exact rational probabilities in polynomial time.

The weighted-test criterion has a computational form: a polynomial-time algorithm returning one passing response to each rational test yields an exact rational lottery, under uniform encoding bounds (Proposition A.2). For the bridge, balancing and a deterministic allocation algorithm supply these responses, as in steps 3–5 of Figure 1. The loop in that figure terminates because each round adds a new allocation, but it may run for as many rounds as there are permitted allocations; the proposition needs only polynomially many responses.

Theorem 3.8 (Computational payment bridge). *In the setting of Theorem 3.5, suppose exact bundle values are evaluable in polynomial time with a known uniform polynomial bound on their rational encoding lengths, and positive rational rescalings are accessible in polynomial time. If a polynomial-time deterministic algorithm returns a complete, envy-freeable $\mathcal{P}$ -allocation at each such rescaling, an exact ex-ante EF lottery with rational probabilities and at most $n(n-1)$ permitted outcomes is computable in polynomial time.*

The proof in Appendix A handles zero test weights by a rational perturbation and invokes the established response-to-mixture machinery of [FP24, PR08, JLB15]. Its hypotheses distinguish an explicit distribution from a procedure that samples one outcome.

4 Applications of certificates and weighted responses

All of our applications follow the same pattern. We first fix a family of permitted complete allocations, namely those with the realized guarantees we want. We then find, within that family, a deterministic certificate or a response to each test. Finally, one of the general theorems of Section 3 turns this into a lottery with the desired expectations. The applications are grouped by the kind of certificate they use, from the most structured to the most flexible: payment certificates in Section 4.1, common payments in Section 4.2, and direct weighted responses in Section 4.3.

4.1 Payment certificates

Unit demand. For $v_i(S) = \max(\{a_{ig} : g \in S\} \cup \{0\})$, with $a_{ig} \geq 0$, the allowed outcomes combine fairness and efficiency. Here EFX_0 requires $v_i(A_i) \geq v_i(A_j \setminus \{g\})$ for every $g \in A_j$, including goods of zero marginal value. A deterministic allocation is **undominated by lotteries** if no lottery weakly improves all own utilities in expectation and strictly improves one. The **maximin share** $\text{MMS}_i(S, k)$ is the largest value agent i can guarantee by partitioning S into k bundles and accepting the least valuable.

Theorem 4.1 (Complete unit-demand allocation). *For arbitrary heterogeneous nonnegative unit-demand valuations, there is an ex-ante envy-free lottery over complete allocations, every one of which is EFX_0 , gives each agent her maximin share, and is undominated by lotteries. More strongly, in every outcome, every nonempty group $H \subseteq N$ and every $i \in H$ satisfy*

$$v_i(A_i) \geq \text{MMS}_i \left(\bigcup_{j \in H} A_j, |H| \right).$$

A polynomial-time matching algorithm produces the required deterministic allocation at each requested rescaling, so the computational payment bridge, Theorem 3.8, applies to rational singleton values.

At each rescaling, a maximum-weight matching gives at most one good to each agent. Assign every leftover good to an unenvied agent. Matching optimality ensures that no observer values a leftover above her matched good. The completed allocation retains all the stated properties and maximizes scaled welfare, hence is envy-freeable. The payment bridge then supplies the lottery. Appendix B.1.1 gives the proof. The efficiency guarantee concerns each outcome, not the expected utilities.

Multi-demand. Multi-demand valuations generalize unit demand: agent i values a bundle by the sum of its k_i largest nonnegative item weights, or all its weights if it has fewer than k_i goods. Capacities may differ, allowing arbitrary mixtures of unit-demand ($k_i = 1$) and additive ($k_i \geq m$) agents. Capped iterated matching provides balanced EF1, envy-freeable allocations at every rescaling (Appendix B.2). The bridge therefore recovers the guarantee of [HSV24, Theorem 6 and Corollary 3].

Corollary 4.2 (Multi-demand recovery). *Every heterogeneous multi-demand profile admits an ex-ante envy-free lottery over at most $n(n-1)$ complete, balanced EF1 allocations. For rational item weights, the lottery is computable in polynomial time.*

Transformed matroid ranks. Agents often value only goods they can use together (at most c per category, one per task), so a matroid rank $r_i(S)$ counts the usable goods in S . Transforming it covers binary submodular, cardinality-based, concave and threshold (non-submodular) valuations. A matroid rank $r_i(S)$ counts the largest independent subset of S . Agents may have different matroids and apply different nondecreasing transformations to these usable counts.

Theorem 4.3 (Transformed matroid ranks). *Let $v_i(S) = \phi_i(r_i(S))$, where each r_i is a matroid rank and each ϕ_i is a nondecreasing function on $\{0, \dots, m\}$ with $\phi_i(0) = 0$. A complete allocation that is simultaneously EF1 and envy-freeable exists. If the transformations are strictly increasing on attainable ranks, it can also be integrally Pareto optimal. Consequently there is an exact ex-ante envy-free lottery with these ex-post guarantees. For explicitly tabulated rational ϕ_i and independence or rank oracles, the allocation and its payments are computable in polynomial time.*

The allowed outcomes are EF1, together with integral PO when the transformations are strict. At each rescaling the construction selects a transformed-welfare maximizer among Lorenz-dominating independent allocations and completes it without losing the rank comparisons. The exchange argument proves envy-freeability for this same allocation (Appendix B.3). Rescaling replaces ϕ_i by $s_i\phi_i$, and the permitted properties are scale invariant, so Theorems 3.5 and 3.8 apply. The result permits complementarities from nonlinear transformations. It does not cover arbitrary element-weighted matroid ranks. The uniform-matroid subclass is covered by the separate multi-demand construction in Appendix B.2.

Teams. Here the allocated items have preferences too, as when players are drafted to clubs or students join project groups. The teams are the agents: team i values participant p at $v_i(p) \geq 0$ and a roster additively. Each participant weakly orders the teams; participants are not owed fairness, only the absence of trades that all affected parties accept. A cyclic exchange has each team on a cycle pass one participant to the next, keeping team sizes fixed. It is beneficial if no team on the cycle is worse off (it values the participant it receives at least as much as the one it sends), no moved participant is worse off, and someone is strictly better off. All these conditions survive rescaling the teams' values, so the bridge gives a lottery that is exactly envy-free in expectation for the teams, whose every outcome is balanced, EF1 for the teams and admits no beneficial cyclic exchange. A maximum-weight round-table construction supplies the certificate the bridge needs, an envy-freeable such outcome at every rescaling (Theorem B.3), in polynomial time (Corollary B.7).

4.2 Common payments for submodular shares

A normalized monotone valuation is **subadditive** if $v_i(S \cup T) \leq v_i(S) + v_i(T)$, and **submodular** if marginal values $\Delta_i(g \mid S) = v_i(S \cup \{g\}) - v_i(S)$ decrease as S grows. Submodularity implies subadditivity. An allocation is **proportional up to one item** (PROP1) if adding at most one unowned good lets every agent reach $v_i(M)/n$.

Write $U_i = \mathbb{E}[v_i(A_i)]$. We next show why the common-payment extension is useful. Agent i 's **maximum expectation share** is

$$\text{MES}_i = \max_{\mathcal{D}: \Pr_{S \sim \mathcal{D}}(g \in S) = 1/n \ \forall g} \mathbb{E}_{S \sim \mathcal{D}}[v_i(S)] \geq \text{MMS}_i. \quad (4.1)$$

An allocation is **marginally envy-free up to one good** (MEF1) [CKM⁺16] if, for every $i \neq j$ with $A_j \neq \emptyset$, deleting some $g \in A_j$ makes the marginal value of the rest of A_j on top of A_i at most $v_i(A_i)$. A lottery is **marginally envy-free** if $\mathbb{E}[v_i(A_i \cup A_j)] \leq 2U_i$. These notions coincide with EF1 and ex-ante EF for additive valuations. For submodular valuations they imply 1/2-EF1 and expected 1/2-envy-freeness.

Theorem 4.4 (Marginal envy-freeness with simultaneous expected shares). *Every normalized monotone submodular profile admits a lottery over complete, balanced, MEF1 allocations such that, simultaneously for all i, j ,*

$$U_i \geq \max \left\{ \frac{v_i(M)}{n}, \frac{n}{2n-1} \text{MES}_i \right\}, \quad (4.2)$$

$$\mathbb{E}[v_i(A_i \cup A_j)] \leq 2U_i. \quad (4.3)$$

Every supported allocation is also PROP1 and 1/2-EF1. A support of at most $n^2 + n + 1$ allocations suffices.

Where the bridge is used. The permitted family is the complete, balanced MEF1 allocations. Repeated maximum-weight matching with marginal-value edges supplies one such allocation at every scaling. Recorded marginal values define an additive proxy H with $H_{ii} = v_i(A_i) = u_i$. Matching optimality gives one centered payment vector satisfying

$$u_i + p_i \geq H_{ij} + p_j, \quad \sum_i p_i = 0. \quad (4.4)$$

Submodularity turns this into a certificate for marginal envy,

$$v_i(A_i \cup A_j) - 2v_i(A_i) \leq p_i - p_j, \quad (4.5)$$

and for an entire curve of individual share benchmarks, all using the same payments (Lemma C.4). Proposition 3.7 therefore cancels every combined test and supplies one lottery with all the stated expectations. The original allocation need not be envy-freeable under v . The proxy carries the certificate. Appendix C.1 proves the full share curve and the support bound.

Curvature-sensitive guarantees. The same lottery improves smoothly as substitution weakens. Define agent i 's total curvature by

$$\kappa_i = 1 - \min_{g: v_i(\{g\}) > 0} \frac{v_i(M) - v_i(M \setminus \{g\})}{v_i(\{g\})} \in [0, 1],$$

with $\kappa_i = 0$ for the zero valuation. Write $F_i(q)$ for the maximum expected value of a random bundle in which every good has inclusion probability q , so $F_i(1/n) = \text{MES}_i$.

Corollary 4.5 (Curvature-sensitive fairness and shares). *The lottery in Theorem 4.4 can be chosen so that*

$$U_i \geq \frac{\mathbb{E}[v_i(A_j)]}{1 + \kappa_i}, \quad U_i \geq \max \left\{ \frac{v_i(M)}{n}, \frac{n}{n + (n - 1)\kappa_i} \text{MES}_i \right\}.$$

Every outcome is $1/(1 + \kappa_i)$ -EF1 from agent i 's perspective. More generally, the same lottery satisfies

$$F_i(q) \leq [nq + \kappa_i(1 - q)]U_i \quad (q \in [0, 1]). \quad (4.6)$$

At $\kappa_i = 1$ these are the general submodular bounds. At $\kappa_i = 0$ they become exact envy-freeness and the full MES benchmark. Appendix C.1 derives them by substituting curvature bounds into the same deterministic inequalities, without changing the lottery.

Comparison with prior work. Under subadditivity, ordinary EF1 implies MEF1, since an EF1 witness g gives

$$v_i(A_i \cup (A_j \setminus \{g\})) \leq v_i(A_i) + v_i(A_j \setminus \{g\}) \leq 2v_i(A_i).$$

Thus the EF1, $1/2$ -EFX outcomes of [FMNP24] are stronger, and for submodular valuations they are already PROP1 by Lemma C.1. Our added guarantees are the expected shares and balance in one lottery, together with the curvature refinement. At the additive endpoint this agrees with Corollary 3.6. Without restrictions on outcomes, independent uniform assignment of goods is exactly ex-ante EF and PROP and guarantees $(1 - 1/e) \text{MES}_i$ by the submodular correlation gap [ADSY10]. For $n \geq 3$ this exceeds $n/(2n - 1)$. The theorem's content is retaining the realized fairness and balance restrictions.

4.3 Direct weighted responses

The next applications construct responses to arbitrary weighted tests without establishing envy-freeability. For $\alpha \in (0, 1]$, expected α -EF means $\mathbb{E}[v_i(A_i)] \geq \alpha \mathbb{E}[v_i(A_j)]$. An outcome is α -EFX if $v_i(A_i) \geq \alpha v_i(A_j \setminus \{g\})$ for every $g \in A_j$, including zero-marginal deletions.

Corollary 4.6 (Explicit short lotteries for subadditive valuations). *Every normalized monotone subadditive profile admits an expected 1/2-EF lottery over complete EF1 and 1/2-EFX allocations, supported on at most $n(n-1)$ outcomes. If exact value queries return rationals whose numerators and positive denominators are bounded by 2^B for a known B , a deterministic algorithm polynomial in n, m, B and the query cost outputs such allocations and exact rational probabilities.*

The permitted family consists of EF1 and 1/2-EFX allocations. Theorem C.5 constructs one passing response to each test of $v_i(A_j)/2 - v_i(A_i)$: it chooses a starting matching and preserves a weighted potential through item additions and envy-cycle rotations. Proposition A.2 then computes the lottery. The fairness factors are those of [FMNP24]. The added algorithmic conclusion is an explicit polynomial-size distribution, rather than only an efficient sampler. The response proof and its oracle assumptions are in Appendix C.2.

Adding a realized MMS guarantee. For submodular goods, we also require a realized share guarantee. An allocation is **envy-free up to one less-preferred good** (EFL) if for every i, j either $|A_j| \leq 1$ or there is $g \in A_j$ with both $v_i(A_i) \geq v_i(A_j \setminus \{g\})$ and $v_i(A_i) \geq v_i(\{g\})$. Write $\mu_i = \text{MMS}_i(M, n)$.

Theorem 4.7 (Expected half-envy-freeness with realized EFL and MMS). *For arbitrary normalized monotone submodular valuations, there is a lottery over at most $n(n-1)$ complete allocations such that*

$$\mathbb{E}[v_i(A_i)] \geq \frac{1}{2} \mathbb{E}[v_i(A_j)] \quad (i, j \in N), \quad (4.7)$$

and every supported allocation is EFL and satisfies

$$v_i(A_i) \geq \frac{1}{8} \mu_i \quad (i \in N). \quad (4.8)$$

Every supported allocation is therefore also EF1 and 1/2-EFX, where the EFX comparison allows deletion of every good, including goods of zero value.

The response policy preserves its chosen weighted inequality under *every* remaining-item order. A completion argument based on [FF26] supplies an order that also meets all MMS thresholds. Choosing that completion answers the test within the smaller permitted family. Theorem 3.1 supplies a common lottery. This gives realized shares, whereas Theorem 4.4 controls expected shares. The full argument is in Appendix C.3. No polynomial-time construction with exact MMS thresholds is claimed.

5 Discussion

The payment bridge turns payment certificates at every rescaling into fair lotteries that preserve their realized properties. Under the access and encoding assumptions of Theorem 3.8, it also converts polynomial-time certificate algorithms into polynomial-time algorithms for explicit lotteries. Its premise is sufficient but not necessary: a weighted test needs only one response, whereas an envy-freeable allocation must defeat all bundle reassignments at once (Section 3.4).

Theorem 5.1 (Obstruction with positive additive goods). *There is an instance with three agents and five strictly positive additive goods with no allocation that is simultaneously EF1, Pareto optimal, and envy-freeable. Every EF1 and Pareto-optimal allocation admits a two-bundle exchange that strictly increases total welfare. Nevertheless the instance admits an ex-ante envy-free lottery over two EF1 and Pareto-optimal allocations.*

Appendix D gives the instance, lottery, and finite case analysis. Thus the general EF–EF1–PO question requires a weaker certificate or another method. For two agents the payment premise and lottery existence coincide for any fixed finite family, as the same appendix shows. Without lotteries, the instance shows that no EF1 and Pareto-optimal allocation of it can be made envy-free by payments.

Only scale-invariant realized properties cross the bridge. Equitability compares utilities across agents and is not scale invariant, so the cross-notion guarantees of [HPXZ26] lie outside the method.

Outlook. The bridge can be used in two ways. Any subsidy theorem that retains a scale-invariant property on a class closed under rescaling yields a BoBW lottery. Contrapositively, for a scale-invariant $\mathcal{P}$, if some profile admits no ex-ante envy-free lottery over $\mathcal{P}$ -allocations, then some rescaling of it admits no $\mathcal{P}$ -allocation that is envy-freeable. For example, EFX allocations support no ex-ante envy-free lottery for some two-agent submodular profile [FMNP24, Proposition A.3, arXiv version], so some rescaling of that profile has no envy-freeable EFX allocation.

The results leave some unresolved questions.

1. *EF1 beyond additivity.* Does every monotone submodular profile admit an EF1, envy-freeable allocation? The bridge would then give exact ex-ante EF with EF1 support, improving the expected envy factor $1/2$ of [FMNP24]. More generally, does this hold for all monotone valuations? The subsidy theorem [BDN⁺20, Theorem 1.4] does not retain EF1.
2. *Efficiency.* For three or more additive agents, can exact ex-ante EF coexist with EF1 and Pareto-optimal support [FSV20, Azi26]? Theorem 5.1 rules out the universal payment premise as a proof of this combination. By condition (ii) of Theorem 3.4, expected envy-freeability at every rescaling would suffice.

Guide to the appendices

Appendix A proves the general results of Section 3: separation and support, an example of balancing, common-payment cancellation, computation, and the weighted bridge. Appendix B proves the payment certificates of Section 4.1, Appendix C the nonadditive results of Sections 4.2–4.3, and Appendix D the obstruction of Section 5. Table 2 locates the statement and proof of each principal result.

Result	Content	Stated	Proved
Theorem 3.1	Weighted-test criterion	§3.1	§A.1
Theorem 3.4	Exact criterion, with payment and matching forms	§3.3	§3.3
Theorem 3.5	Payment bridge	§3.4	§3.4
Corollary 3.6	Additive goods and mixed manna	§3.4	§3.4
Theorem A.3	Weighted payment bridge	§A.5	§A.5
Proposition 3.7	Common-payment cancellation	§3.5	§A.3
Theorem 3.8	Computational payment bridge	§3.6	§A.4
Theorem 4.1	Unit demand	§4.1	§B.1
Corollary 4.2	Multi-demand	§4.1	§B.2
Theorem 4.3	Transformed matroid ranks	§4.1	§B.3
Theorem B.3	Two-sided teams	§B.4	§B.4
Theorem 4.4, Corollary 4.5	Submodular shares and curvature	§4.2	§C.1
Corollary 4.6	Explicit subadditive lottery	§4.3	§C.2
Theorem 4.7	Submodular EFL and MMS	§4.3	§C.3
Theorem 5.1	Obstruction	§5	§D

Table 2: Where each principal result is stated and proved.

A Foundations: separation, support, and algorithms

Section 3 proves the two cancellation lemmas, including the cycle characterization of envy-freeability, the exact criterion, and the payment bridge. This appendix supplies the remaining general proofs.

A.1 Separation and support

Proof of Theorem 3.1. If a lottery μ on S has $\mathbb{E}_\mu[g_k] \leq 0$ for all k , then $\mathbb{E}_\mu[\lambda^\top g] \leq 0$ for every $\lambda \geq 0$, so some supported outcome passes the test. Conversely, the expectations attainable by lotteries on S form the polytope $K = \text{conv}\{g(x) : x \in S\}$, and a fair lottery exists exactly when K meets $C = \mathbb{R}_{\leq 0}^d$. If they are disjoint, strong separation of the compact convex set K from the closed convex set C [BV04] gives λ with $\min_{y \in K} \lambda^\top y > \sup_{z \in C} \lambda^\top z$. The supremum is finite only if $\lambda \geq 0$, and then it is zero, so every $x \in S$ fails the test λ . The support bound is Lemma A.1. $\square$

Lemma A.1 (Support). *Every attainable expectation vector in $\mathbb{R}^d$ is attained by a lottery on at most $d + 1$ outcomes, and feasible expected inequalities $\mathbb{E}[g_k] \leq 0$ ($k \leq d$) are feasible on at most d outcomes.*

Proof. The first claim is Carathéodory’s theorem. For the second, take a feasible lottery μ on $d + 1$ outcomes. If their vectors are affinely dependent, Carathéodory’s reduction removes one. Otherwise solve $\sum_\ell h_\ell = 0$ and $\sum_\ell h_\ell g_\ell = -\mathbf{1}$; the solution is nonzero and sums to zero, so it has a negative

entry. Move to $\mu + \theta h$ with the largest $\theta > 0$ keeping all probabilities nonnegative. One probability vanishes, and every expected violation drops by θ . Both steps solve rational linear systems, so they preserve polynomial encoding length. $\square$

A.2 A balancing example

Lemma 3.3 rescales the rows of the comparison coefficients until every vertex has equal incoming and outgoing weight. For example, take

$$\lambda = \begin{pmatrix} 0 & 2 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix}, \quad b = (3, 5, 4), \quad q_{ij} = b_i \lambda_{ij}.$$

Then

$$q = \begin{pmatrix} 0 & 6 & 3 \\ 5 & 0 & 5 \\ 4 & 4 & 0 \end{pmatrix}, \quad \begin{array}{c|cc} & \text{outgoing weight} & \text{incoming weight} \\ \hline 1 & 9 & 9 \\ 2 & 10 & 10 \\ 3 & 8 & 8 \end{array}$$

At utility scales $s = (1/3, 1/5, 1/4)$, multiplying the paid inequalities by q_{ij} recovers the original test and cancels every payment. The stationary distribution is $(9, 10, 8)/27$, yielding the normalized row multipliers $b/27$. Its probabilities belong to an auxiliary walk on agents, not to the final allocation lottery. Lemma 3.3 turns the normalized circulation into the reassignment

$$P = \frac{1}{27} \begin{pmatrix} 18 & 6 & 3 \\ 5 & 17 & 5 \\ 4 & 4 & 19 \end{pmatrix}, \quad s = \left(9, \frac{27}{5}, \frac{27}{4}\right),$$

with s a multiple of the scales above. Every row and column of P sums to one, and $s_i P_{ij} = \lambda_{ij}$ for $i \neq j$. The test score of every allocation is therefore the expected welfare gain $T_A(s, P)$ of this random reassignment.

The premise must hold throughout the rescaling orbit because envy-freeability can change with utility units. Envy gains $(1, -1)$ permit payments with $p_1 - p_2 = 1$. Rescaling the first valuation by two gives gains $(2, -1)$ and incompatible payment inequalities. The rational perturbation needed for zero test coefficients is proved in the computational argument below.

A.3 Common-payment cancellation

Proof of Proposition 3.7. We must cancel payments in a test that combines pairwise comparisons and individual benchmarks. We first choose scales making every payment coefficient equal, then use centering to cancel their common value. By Theorem 3.1, fix arbitrary nonnegative test weights L_{ij} and μ_{ik} . Put $\nu_i = \sum_k c_k \mu_{ik}$ and $d_i = \sum_{j \neq i} L_{ij} + \nu_i$. For the moment add the same arbitrarily small positive amount to every L_{ij} . This makes every d_i positive and makes the stochastic matrix

$$W_{ij} = \frac{L_{ij} + \nu_i/n}{d_i} \quad (j \neq i), \quad W_{ii} = \frac{\nu_i/n}{d_i}$$

irreducible. Let π be its positive stationary distribution and choose the utility scales $s_i = d_i/\pi_i$. Apply the hypothesis at the scaled profile $s \cdot v$, obtaining A and centered p .

Homogeneity and the certificate inequalities give

$$\begin{aligned} & \sum_{i \neq j} L_{ij} a_{ij}(v, A) + \sum_{i, k} \mu_{ik} b_{ik}(v, A) \\ & \leq \sum_i p_i \left(\frac{d_i}{s_i} - \sum_{j \neq i} \frac{L_{ji}}{s_j} \right). \end{aligned}$$

Stationarity makes the expression in parentheses equal to the same number $\frac{1}{n} \sum_j \nu_j / s_j$ for every i . The right side is therefore zero because $\sum_i p_i = 0$. Thus A answers the combined test. Letting the perturbation vanish handles zero weights, and Theorem 3.1 gives the claimed lottery. $\square$

Approximate comparisons. With $a_{ij} = \alpha v_i(A_j) - v_i(A_i)$, Proposition 3.7 applies: a certificate $s_i v_i(A_i) + p_i \geq \alpha s_i v_i(A_j) + p_j$ at every scaling gives expected α -EF on the permitted family. Only the comparison value is discounted: discounting p_j as well would destroy the cancellation.

A.4 Computing the lottery

Proposition A.2 (Computing a lottery from weighted responses). *Consider a finitely encoded instance with a finite permitted set S and rational violation vectors $g(x) \in \mathbb{Q}^d$, where d is polynomial in the input size. Assume:*

1. *every permitted outcome has a polynomial-size representation;*
2. *each $g_k(x)$ is exactly evaluable in polynomial time, with a known uniform polynomial bound on its rational encoding length;*
3. *for every rational $\lambda \geq 0$, a polynomial-time algorithm returns $x \in S$ with $\lambda^\top g(x) \leq 0$.*

Then an exact rational lottery meeting all expected inequalities can be computed in polynomial time, using polynomially many responses and with support at most d .

Proof. In the zero-sum game where one player picks $x \in S$, with payoff row $(-g_1(x), \dots, -g_d(x), 0)$, and the other picks one of the $d + 1$ columns, the value is exactly zero by Theorem 3.1 and the zero column. Hypothesis 3 is a good-enough response to every mixed column choice, the column simplex has explicit inequalities, and payoff rows have uniformly bounded encoding length. The exact good-enough-response theorem of Farina and Pipis [FP24, Theorem 3.1], in the ellipsoid-against-hope line of [PR08, JLB15], then returns, after polynomially many responses, a rational mixture of at most $d + 1$ returned outcomes with nonnegative payoff against every column. Lemma A.1 removes one more outcome. There is no need to list all of S . $\square$

Proof of Theorem 3.8. We verify the hypotheses of Proposition A.2 with $d = n(n - 1)$ and $g = e$. An allocation is encoded by the recipient of each item, and its envy coordinates are differences of exactly evaluable bundle values. For a strictly positive rational test, Lemma 3.3 computes s and P , the deterministic algorithm runs on $s \cdot v$, and scale invariance and Lemma 3.2 give a permitted A with $\sum_{i \neq j} \lambda_{ij} e_{ij}(A) = T_A(s, P) \leq 0$.

Zero weights need care: the existence proof used the smallest positive score over all allocations, which we cannot compute. Encoding bounds replace it. Let 2^L bound the numerator and denominator of every reduced envy coordinate, uniformly over allocations, and let 2^K bound those of the query λ . A sum of d products $\lambda_{ij} e_{ij}(A)$ has denominator at most $2^{d(K+L)}$, so every nonzero score has magnitude at least $\gamma = 2^{-d(K+L)}$, for every allocation at once. Also $\sum_{i \neq j} |e_{ij}(A)| \leq d2^L$. Answer

the positive test $\hat{\lambda} = \lambda + \eta \mathbf{1}$ with $\eta = \gamma/(2d2^L)$, which has polynomial encoding length. The response satisfies $\lambda^\top e(A) \leq \eta \sum_{i \neq j} |e_{ij}(A)| \leq \gamma/2$, so its score for λ cannot be positive. The bound is polynomial in the rational input lengths, not strongly polynomial. $\square$

A.5 Weighted comparisons and their bridge

Fix entitlements $w_i > 0$, normalized so that $\sum_i w_i = 1$. They remain fixed when utilities are rescaled. Dividing v_i by w_i only rescales the observer's entire valuation. It does not insert the compared owner's factor $1/w_j$. Define weighted envy by

$$e_{ij}^w(v, A) = \frac{v_i(A_j)}{w_j} - \frac{v_i(A_i)}{w_i} \quad (i \neq j), \quad e_{ii}^w(v, A) = 0.$$

A lottery is **ex-ante weighted envy-free** (WEF) if $\mathbb{E}[e_{ij}^w(v, A)] \leq 0$ for every ordered pair. For signed additive valuations this implies weighted proportionality:

$$\mathbb{E}[v_i(A_i)] \geq w_i v_i(M).$$

Indeed, multiply each WEF comparison by w_j , sum over j , and use additivity and completeness. No assumption on the sign of $v_i(M)$ is needed.

An allocation is **weighted envy-freeable** if some real payments p_i satisfy

$$\frac{v_i(A_i) + p_i}{w_i} \geq \frac{v_i(A_j) + p_j}{w_j} \quad (i \neq j). \quad (\text{A.1})$$

This is the standard weighted subsidy convention [KAG⁺25]. Writing $r_i = p_i/w_i$ turns (A.1) into $e_{ij}^w \leq r_i - r_j$. Thus all normalized payments can be shifted by the same constant. Equivalently, p_i can be shifted by a common multiple of w_i . In particular, payments can be nonnegative, or centered with $\sum_i p_i = 0$, equivalently $\sum_i w_i r_i = 0$. Uniformly shifting the unnormalized payments is generally invalid.

Theorem A.3 (Weighted payment bridge). *Fix v, w and a finite nonempty family $\mathcal{F}$ of complete allocations. Suppose that for every positive scale vector c , some $A \in \mathcal{F}$ is weighted envy-freeable at $(c \cdot v, w)$. Then $\mathcal{F}$ supports an ex-ante WEF lottery with at most $n(n-1)$ outcomes. In particular, any scale-invariant realized property obtained together with weighted envy-freeability throughout a class closed under positive rescaling carries over to ex-ante WEF. Neither additivity nor monotonicity is required.*

Proof. Weighted envy-freeability at $(c \cdot v, w)$ is condition (1) of Lemma 3.2 for the matrix $E_{ij} = e_{ij}^w(v, A)$ at scales c , with normalized payments r . Fix a strictly positive test λ , and take c and P from Lemma 3.3. The premise at $(c \cdot v, w)$ supplies $A \in \mathcal{F}$, and condition (4) of Lemma 3.2 gives $\sum_{i \neq j} \lambda_{ij} e_{ij}^w(v, A) = \sum_{i,j} c_i P_{ij} e_{ij}^w(v, A) \leq 0$. Perturb zero test coordinates and pass to a constant subsequence of allocations, as in Theorem 3.4. The weighted-test criterion, Theorem 3.1, proves the conclusion and the support bound. $\square$

Computation and approximation. Under the rational-encoding and oracle hypotheses of Theorem 3.8, a polynomial-time algorithm for the stated deterministic certificate yields an exact rational lottery in polynomial time. Rational entitlements are part of the input. Division by them has polynomial encoding length. The proof of that theorem only uses homogeneity of the comparison coordinates and payment cancellation, so it applies verbatim to e^w . Similarly, replacing e_{ij}^w by

$\alpha_{ij}v_i(A_j)/w_j - v_i(A_i)/w_i$, for fixed $\alpha_{ij} \in (0, 1]$, carries over the same pair-dependent approximation factors. The paid inequality must have the form

$$\frac{c_i v_i(A_i) + p_i}{w_i} \geq \alpha_{ij} \frac{c_i v_i(A_j)}{w_j} + \frac{p_j}{w_j};$$

the comparison value, but not its payment, is discounted.

Weighted characterization. Lemma 3.2 holds for any matrix with zero diagonal, so it applies to $E_{ij} = e_{ij}^w(v, A)$ at scales c , with normalized payments r . Thus a fixed allocation A is weighted envy-freeable at $(c \cdot v, w)$ if and only if every directed cycle C satisfies $\sum_{(i,j) \in C} c_i e_{ij}^w(v, A) \leq 0$, if and only if

$$T_A^w(c, P) := \sum_{i,j} c_i P_{ij} e_{ij}^w(v, A) \leq 0 \quad \text{for every doubly stochastic } P. \quad (\text{A.2})$$

The proof of Theorem 3.4, with e^w in place of e , shows that $\mathcal{F}$ supports an ex-ante WEF lottery if and only if for every $c > 0$ and every strictly positive doubly stochastic P , some $A \in \mathcal{F}$ satisfies (A.2). The quantifier comparison after Theorem 3.5 therefore carries over to T^w and WEF. For two agents the premise of Theorem A.3 is also necessary: the only nontrivial cycle is $(1, 2, 1)$, and under a WEF lottery its score $c_1 e_{12}^w + c_2 e_{21}^w$ has nonpositive expectation, so some supported allocation is weighted envy-freeable at $(c \cdot v, w)$.

There is one change of interpretation: the comparison term after a permutation σ is $v_i(A_{\sigma(i)})/w_{\sigma(i)}$, retaining the original owner's denominator. It is a permutation of normalized bundle columns. It is not the welfare of a physical reassignment evaluated with each recipient's own entitlement. Ordinary welfare maximality or fPO does not establish this condition.

B Deterministic payment certificates

The four constructions below supply payment certificates for the permitted families in Section 4.1, in the same order. The bridge supplies the lottery in each case.

B.1 Unit-demand valuations

We use EFX_0 , maximin shares, and domination by lotteries as defined in Section 4.1. Requiring each support allocation to be undominated by lotteries does not make the final lottery Pareto optimal in expectation.

B.1.1 Completing a matching without sacrificing fairness


A unit-demand agent values a bundle by its best item. Every good must still be allocated, and an unused extra good can change how others value its recipient's bundle, so one-object assignment does not settle the complete-allocation problem.

The construction first maximizes the total value of assigning at most one good per agent. It then puts every leftover good with a carefully chosen recipient. The recipient is chosen to be unenvied, and matching optimality ensures that none of the leftovers creates new envy toward her. This preserves a particularly strong form of fairness while retaining maximum weighted welfare.

Example B.1 (Where the completed allocation comes from). Consider these singleton values, before any rescaling:

	a	b	c	d	e
1	9	10	1	5	4
2	1	12	2	4	3
3	2	3	8	2	1

The maximum matching gives a to agent 1, b to agent 2, and c to agent 3, for total value 29. To see optimality, if agent 2 does not get b , the total is at most $10 + 4 + 8 = 22$. If she does, the other agents can get at most 9 and 8, and the displayed matching attains both. Agent 1 envies agent 2, but nobody envies agent 1.



Give both leftovers to agent 1. Agent 2 values her enlarged bundle $\{a, d, e\}$ at $4 \leq 12$, and agent 3 values it at $2 \leq 8$. So agent 1 remains unenvied. Every other bundle is a singleton. Deleting its sole good removes any envy toward it. This explains both the completion rule and its EFX_0 guarantee. The proof below performs the same construction after every positive rescaling.

Proof of Theorem 4.1. Fix arbitrary positive scales c_i , and write $b_{ig} = c_i a_{ig}$. These scales are inputs to the deterministic construction, not the probabilities of the final lottery. Add n distinct dummy objects, all worth zero, so an agent can receive no real good while still being matched to something.

Find a maximum-weight matching assigning one real or dummy object q_i to each agent. Write $s_i = b_{iq_i}$ and let R be the unmatched real goods. Matching optimality gives the residual bound

$$b_{ih} \leq s_i \quad (i \in N, h \in R), \quad (\text{B.1})$$

because otherwise replacing q_i by h improves the matching.

We next locate the recipient of all leftovers. Draw an arc $i \rightarrow j$ when $b_{iq_j} > s_i$, meaning that i envies j 's matched object. This directed graph is acyclic: a cycle would allow its agents to rotate their matched objects, strictly increasing the matching weight. Choose a vertex r with no incoming arc. Give all of R to r , retain every matched real object with its matched agent, and remove the artificial dummies. This is a complete allocation A , and (B.1) shows that $c_i v_i(A_i) = s_i$. Moreover, for $i \neq r$,

$$c_i v_i(A_r) = \max(\{b_{iq_r}\} \cup \{b_{ih} : h \in R\}) \leq s_i.$$

Thus r is unenvied after completion. Every other bundle has at most one item. Envy toward such a bundle disappears after its sole item is deleted, while comparisons toward r are already envy-free. This proves EFX_0 , including zero-marginal deletions.

For any complete allocation B , select a highest-valued item from each nonempty B_i , using a distinct dummy for an empty bundle. The selected objects form a matching of weight $\sum_i c_i v_i(B_i)$. Consequently A globally maximizes this weighted welfare. It therefore maximizes welfare among permutations of its own bundles and is envy-freeable at the scaled profile. Because all c_i are strictly positive, no lottery can Pareto improve A : such an improvement would strictly increase the same weighted sum, whereas the expectation of a linear objective under any lottery is at most its maximum over deterministic allocations.

The share guarantee is stronger than the full-population MMS comparison: it also holds if a group recomputes its shares using only its own goods. Fix such a group $H \ni i$ and put $t = v_i(A_i)$. Every residual good is worth at most t to i . Among the goods in $\bigcup_{j \in H} A_j$, therefore, at most $|H| - 1$ goods are worth strictly more than t : only the other agents' matched goods can have that property. In any partition of this union into $|H|$ parts, one part contains none of these goods, and hence has unit-demand value at most t . This proves the groupwise maximin-share inequality, including empty bundles and ties. Taking $H = N$ gives ordinary MMS.

We can now identify the inputs to Theorem 3.5. The allowed outcomes are complete allocations satisfying EFX_0 , all the displayed groupwise share inequalities, and no domination by lotteries. Each requirement is unchanged by multiplying an agent's valuation by a positive constant. The construction has produced one such outcome, with envy-eliminating payments, at every choice of those constants. The theorem therefore gives an envy-free lottery over these allowed outcomes. One maximum-weight matching, $O(n^2)$ envy comparisons, and allocation of the leftovers compute each requested deterministic allocation. $\square$

Unit demand is the case $k_i = 1$ of the multi-demand valuations treated next, in Appendix B.2, for which ex-ante EF with balanced ex-post EF1 is already known [HSV24, Theorem 6 and Corollary 3]. The additional point in Theorem 4.1 is that every complete outcome can also satisfy EFX_0 , the share guarantees, and the stated efficiency property. Its proof uses the maximum of item values twice: residual items do not accumulate in value, and selecting one item from each bundle captures its entire welfare. Neither step extends to $k_i > 1$.

B.2 Multi-demand valuations

Let $a_{ig} \geq 0$ and let k_i be a positive integer. Agent i has valuation

$$v_i(S) = \max_{T \subseteq S: |T| \leq k_i} \sum_{g \in T} a_{ig}. \quad (\text{B.2})$$

These are element-weighted uniform matroid-rank valuations. We give a deterministic payment certificate for the known balanced EF1 lottery guarantee of [HSV24, Theorem 6 and Corollary 3], whose proof uses a picking-sequence decomposition. Their outcomes are already balanced.

The construction adapts iterated maximum-weight matching [BDN⁺20]. Every agent receives one good per full round, but her matching weights become zero after round k_i . Inactive agents still receive goods, so the allocation remains balanced. Dummies are introduced only in the final partial round. The key comparison assigns the selected goods of each rival bundle to earlier active rounds, where they were available.

Capped iterated matching

Input: Item weights $a_{ig} \geq 0$, capacities $k_i \geq 1$, and positive utility scales s_i .

Output: A complete, balanced EF1 allocation A that is envy-freeable at $s \cdot v$.

1. Set $A_i \leftarrow \emptyset$ for all i , $U \leftarrow M$, and $t \leftarrow 1$.
 2. While $U \neq \emptyset$:
 - (a) If $|U| < n$, add $n - |U|$ distinct zero-valued dummy goods to U .
 - (b) Find a maximum-weight matching assigning one good of U to every agent, with weights $s_i a_{ig} \mathbf{1}[t \leq k_i]$.
 - (c) Give each matched real good to its matched agent, remove all matched goods from U , and set $t \leftarrow t + 1$.
 3. Return A .
-

Figure 2: Capped iterated matching for multi-demand payment certificates.

Proposition B.2 (Multi-demand payment certificates). *For every multi-demand profile and positive scale vector s , the construction in Figure 2 returns a complete, balanced EF1 allocation that is envy-freeable at $s \cdot v$. For rational item weights and scales, the allocation and supporting payments are computable in polynomial time.*

Proof. The empty instance is immediate. Otherwise there are $R = \lceil m/n \rceil$ rounds. Write g_i^t for agent i 's round- t good, allowing a zero-valued dummy in the final round, and write $a_i(g) = a_{ig}$ and $b_i(g) = s_i a_{ig}$. Every full round gives one real good to every agent, so the final allocation is complete and balanced.

EF1. If $m < n$, each bundle has at most one good. Assume $m \geq n$ and fix agents i, j . Delete g_j^1 from A_j , and choose the at most k_i goods realizing $v_i(A_j \setminus \{g_j^1\})$. List them as $h_1, \dots, h_d$ in increasing order of their original rounds. There is at most one such good per round, and none comes from round 1, so h_ℓ was allocated in round at least $\ell + 1$. It was therefore available and unmatched in round ℓ . Agent i was active, since $\ell \leq d \leq k_i$. Replacing g_i^ℓ by h_ℓ cannot increase that round's matching weight, whence $a_i(g_i^\ell) \geq a_i(h_\ell)$. Thus

$$v_i(A_j \setminus \{g_j^1\}) = \sum_{\ell=1}^d a_i(h_\ell) \leq \sum_{\ell=1}^d a_i(g_i^\ell) \leq v_i(A_i).$$

In particular, g_j^1 is a deletion witness for every observer.

Envy-freeability. Fix a permutation σ . For each i , select $d_i = \min\{k_i, |A_{\sigma(i)}|\}$ highest- a_i goods of $A_{\sigma(i)}$, and list them as $h_i^1, \dots, h_i^{d_i}$ in increasing original-round order. The original round of h_i^t is at least t , so it was available in round t . At that round, assigning h_i^t to each i with $d_i \geq t$ is a feasible partial matching: the selected goods lie in distinct source bundles across agents, and every such agent is active. Extend it to a matching covering all agents. Nonnegative weights and round optimality give

$$\sum_{i: d_i \geq t} b_i(h_i^t) \leq \sum_{i: k_i \geq t} b_i(g_i^t).$$

Summing over rounds yields

$$\sum_i s_i v_i(A_{\sigma(i)}) \leq \sum_i \sum_{t=1}^{\min\{k_i, R\}} b_i(g_i^t) \leq \sum_i s_i v_i(A_i).$$

The last inequality holds because the active-round real goods form a subset of A_i of size at most k_i . Thus no permutation increases scaled welfare, and Lemma 3.2 gives envy-freeability. The construction uses R maximum-weight matchings; supporting payments are then obtained from the envy graph in polynomial time. $\square$

Proof of Corollary 4.2. Multi-demand profiles are closed under positive rescaling, and balance and EF1 are scale invariant. Proposition B.2 supplies their payment certificates. Theorems 3.5 and 3.8 give the lottery, its support bound, and polynomial-time computation. $\square$

B.3 Transformed matroid-rank valuations

A matroid on M is a nonempty family of **independent** subsets, closed under taking subsets, with the following augmentation property: if independent sets I, J satisfy $|I| < |J|$, some $g \in J \setminus I$ can be added to I while preserving independence. Its rank $r(S)$ is the largest size of an independent subset of S . The rank increments from adding one item are zero or one, and the augmentation property says that if $r(T) > r(S)$, some item of $T \setminus S$ increases the rank of S . Agents may have different matroids and may apply different nondecreasing functions ϕ_i to their usable counts. Thus one agent may care linearly about the count, another may value reaching two usable goods especially highly, and a third may stop benefiting beyond a threshold.

The challenge resembles completion under unit demand. A well-balanced allocation of independent bundles may leave redundant goods unassigned. Adding these goods does not help their recipients, but it can affect how others compare the bundles. We will complete such an allocation without losing EF1, and choose the initial allocation carefully enough that envy can also be eliminated by payments after any positive rescaling.

A partial allocation is **nonredundant** if each agent's bundle is independent in her own matroid: deleting any of its goods then lowers her rank. The vector $(r_i(A_i))_i$ records agents' own ranks. It is **Lorenz dominating** if, after sorting its coordinates in ascending order, the sum of its first h coordinates is at least that of every feasible own-rank vector, for every h . Such an allocation simultaneously maximizes total rank and balances rank as strongly as possible. The existence of nonredundant Lorenz-dominating allocations for heterogeneous matroid ranks is an established result [BEF21]. This is the deterministic existence theorem from which our construction starts.

The proof makes three choices for different reasons. It first selects a Lorenz-dominating allocation with maximum transformed welfare among such allocations. It then adds unused goods while

retaining each agent's own rank and an upper bound on everyone else's assessment of her bundle. Finally, exchanges of one rank unit show that every positive envy contribution around a cycle is offset by a negative one. The absence of a positive envy cycle is exactly the condition for envy-eliminating payments.

Proof of Theorem 4.3. Among nonredundant Lorenz-dominating allocations choose B maximizing $\sum_i \phi_i(r_i(B_i))$. Write $k_i = r_i(B_i) = |B_i|$ and $R = \sum_i k_i$, the maximum feasible total rank. From any bundle we can delete redundant items until an independent subset of the same rank remains. Thus any feasible vector of own ranks also has a nonredundant representation. The optimality of our choice of B applies to all such rank vectors, not only to the allocations that first represented them.

We first complete B while maintaining its own ranks, EF1, and the following bounds on comparison values:

$$r_i(A_j) \leq k_j \quad (i, j \in N). \quad (\text{B.3})$$

In words, nobody may assign j 's bundle a rank above j 's own rank. Initially these bounds follow from $|B_j| = k_j$. At any allocation with these bounds and the same Lorenz vector, if i envies j in rank, then augmentation moves an item from A_j to A_i and raises i 's rank by one. The donor must lose one rank unit. Otherwise total rank exceeds R . If $k_j \geq k_i + 2$, this transfer strictly improves the Lorenz vector. Thus

$$r_i(A_j) > k_i \implies k_j = k_i + 1 \text{ and } r_i(A_j) = k_j. \quad (\text{B.4})$$

In particular B is EFX₀: whenever envy occurs, deleting any item from the independent bundle B_j leaves at most k_i items.

Consider an unassigned good g . A recipient j would be unsuitable if somebody already envies her bundle, or if adding g would violate one of the comparison bounds. Record each objection from agent i by an edge $j \rightarrow i$: this means either $r_i(A_j) > k_i$ or $r_i(A_j \cup \{g\}) > k_j$. Every edge satisfies $r_i(A_j) = k_j$, by (B.4) or by the bound of one on a rank increment. An envy-only cycle would strictly decrease own-rank levels along its edges, which is impossible. If a cycle contains an edge of the second kind, rotate the current bundles along that cycle and give g to the receiver of one such edge. Each receiver obtains the former owner's rank, and that receiver gains one additional unit. Total rank increases, again impossible. Hence this graph is acyclic.

An acyclic directed graph has a sink, that is, a vertex with no outgoing edge. Give g to such a vertex. By the meaning of our edges, nobody envied that recipient beforehand, so deleting g supplies an EF1 witness against its enlarged bundle. No observer bound is violated, and the recipient's own rank cannot increase because R is maximum. Repeating gives a complete allocation A that is EF1 when evaluated by the ranks and satisfies (B.3). Applying each monotone ϕ_i preserves every EF1 witness.

We next prove that this same allocation admits envy-eliminating payments for the transformed valuations. Work first with the simpler rank differences $d_{ij} = r_i(A_j) - k_i$. Consider a path of distinct agents beginning with a positive d -edge and continuing through nonnegative d -edges. Thus the first agent envies the next in rank, and each later agent weakly prefers the next bundle to her own in rank. Inequality (B.3) makes the own-rank levels nondecreasing along this path. Augmentation can pass one unit of rank from its endpoint to its start: first move an item to raise the start's rank. The donor loses exactly one unit by maximality of R . Then an item from the next bundle restores that donor's rank. Continue in this fashion. Every donor must lose one unit, since a zero loss would give a feasible total rank of $R + 1$. All intermediate ranks are restored.

This exchange has a useful restriction. If the endpoint's own rank exceeded the start's by at least two, it would move a unit from a larger coordinate to a smaller one and strictly improve the Lorenz

comparison. That is impossible for the chosen vector. The endpoint ranks must therefore be k and $k + 1$. The exchange merely swaps them and remains Lorenz dominating. It cannot increase the transformed welfare we maximized when choosing B . Its gain at the start must therefore be no greater than its loss at the endpoint:

$$\phi_{\text{start}}(k + 1) - \phi_{\text{start}}(k) \leq \phi_{\text{end}}(k + 1) - \phi_{\text{end}}(k). \quad (\text{B.5})$$

This argument also shows that a nonnegative-rank-envy path cannot contain two positive edges: its rank levels would increase by at least two.

It remains to turn these path comparisons into the required cycle comparisons. Take any simple directed cycle of agents. If every rank-envy edge is nonnegative, the observer bounds force every rank level and comparison around it to be equal, giving zero transformed envy. Otherwise split the cycle into runs of nonnegative rank-envy edges, each followed by a negative edge. Here a run is just a consecutive stretch of edges for which $d_{ij} \geq 0$. Each has at most one positive edge. A run with none contributes zero transformed envy. If its positive edge starts at rank k , its transformed gain is $\phi_{\text{start}}(k + 1) - \phi_{\text{start}}(k)$. The endpoint has rank $k + 1$. The following negative edge evaluates the next bundle at rank at most k , so its transformed loss is at least $\phi_{\text{end}}(k + 1) - \phi_{\text{end}}(k)$. Inequality (B.5) shows that this loss cancels the preceding gain. Summing shows that every cycle has nonpositive transformed envy, so Lemma 3.2 provides supporting payments.

For strict transformations, an integral Pareto improvement in transformed values weakly raises every own rank and strictly raises one, contradicting maximum total rank. To apply Theorem 3.5, require each allowed complete allocation to be EF1, and also integrally Pareto optimal when the transformations are strict. These comparisons are unchanged by positive rescaling. Multiplying v_i by $c_i > 0$ merely replaces ϕ_i by another nondecreasing function $c_i \phi_i$, so the construction just proved supplies the required paid allocation for every such input. The payment bridge supplies exact ex-ante envy-freeness with the specified properties in every outcome.

We finish by showing how to compute the initial allocation. This is separate from the completion and payment arguments. Three successive objectives must be implemented: maximize total rank, balance the ranks, then maximize transformed welfare. Encode them in one objective by choosing coefficients that make an improvement at an earlier stage dominate every possible change at later stages. Set

$$P = \sum_i \phi_i(m), \quad H = P + 1, \quad L = Hm^2 + P + 1, \quad f_i(t) = Lt - Ht^2 + \phi_i(t).$$

Maximizing $\sum_i f_i(k_i)$ over nonredundant allocations first maximizes total rank, then minimizes the integer sum of squares, then maximizes transformed welfare. A Lorenz-dominating vector exists, and strict convexity of the square function shows that the minimum square sum among maximum-total-rank vectors has exactly its sorted profile. Each f_i is strictly concave on integer ranks: the positive second difference of ϕ_i is at most P , whereas the quadratic contributes $-2H$.

Weighted matroid intersection optimizes weight over sets that are independent in each of two specified matroids, with a fixed feasible cardinality when required. We reduce the above optimization to that problem as follows. For agent i of matroid rank R_i , make an agent-specific copy of the real goods and add R_i dummy elements that create no independence constraints. Limit the combined set size for that agent to R_i . This operation is matroid truncation. The first matroid imposes these constraints separately for every agent, which is their direct sum. The second permits at most one copy of each real good, leaves the dummies unconstrained, and limits total size to $\sum_i R_i$. The all-dummy set is independent in both and has maximum size. Such a set is a **common base**. A common base with k_i real elements for agent i represents a nonredundant bundle of

size k_i . Give real elements weight zero and dummy s weight $-(f_i(s) - f_i(s - 1))$. Concavity makes these weights increasing in s , so the optimal $R_i - k_i$ dummies have indices $k_i + 1, \dots, R_i$, contributing $f_i(k_i) - f_i(R_i)$. A maximum-weight common base therefore implements exactly the desired optimization. The construction has at most $2nm$ elements and polynomially encoded rational weights. Weighted matroid intersection is polynomial [Sch03]. The completion needs polynomially many rank queries. Payments can be obtained from maximum path weights in the envy graph, whose absence of positive cycles was proved above, with polynomially many further arithmetic operations. $\square$

For linear transformations $\phi_i(t) = c_i t$, all support allocations maximize the same objective $\sum_i v_i(A_i)/c_i = \sum_i r_i(A_i)$. Consequently both each outcome and the final lottery are undominated by lotteries. This does not extend to arbitrary strict transformations: with two identical rank functions $r(S) = |S|$ on two goods and $\phi(t) = t^2$, the balanced allocation has utilities $(1, 1)$, while giving both goods to a uniformly random agent has expected utilities $(2, 2)$. The nonlinear theorem deliberately asserts integral PO.

The exact unit increments are essential: with unequal element weights, the receiving agent's gain and the donating agent's loss need not agree, so the argument does not extend to element-weighted matroid ranks. Weighted uniform matroid ranks are the multi-demand valuations of Appendix B.2.

B.4 Fair team assignments that respect participants' choices

In some allocation problems the items being allocated have preferences of their own, as when players are drafted to clubs, students are placed in project groups, or employees are assigned to departments. Let $N = [n]$ be the teams and M the set of m participants; an allocation partitions M among the teams. The teams are the agents: team i has a nonnegative additive valuation v_i over participants and compares bundles by v_i , so EF1 and balance are as in Section 2. In addition, each participant p weakly orders the teams by $\succeq_p$; these orders need not agree with one another or with the teams' values. Participants are not owed fairness. Instead, we ask that no exchange be available that every affected team and participant accepts.

A **cyclic exchange** selects $\ell \geq 2$ distinct teams $i_1, \dots, i_\ell$ and one participant $p_k \in A_{i_k}$ from each, and moves p_k to team i_{k+1} (indices modulo ℓ). Team i_k thus gives up p_k and receives p_{k-1} ; it weakly benefits if $v_{i_k}(p_{k-1}) \geq v_{i_k}(p_k)$, and participant p_k weakly benefits if $i_{k+1} \succeq_{p_k} i_k$. The exchange is **beneficial** if every selected team and every moved participant weakly benefits and at least one of them strictly benefits. For $\ell = 2$ this is a swap of two participants. Since each selected team sends and receives exactly one participant, an exchange preserves bundle sizes, and hence balance; exchanges in which a team sends several participants are not considered.

Theorem B.3 (Fair and consensually stable team lotteries). *For nonnegative additive team values and arbitrary weak participant preferences, there is an ex-ante envy-free lottery for the teams whose every allocation is balanced, EF1, and admits no beneficial cyclic exchange.*

The round-table construction. If $m = 0$, use the empty allocation. Otherwise let $K = \lceil m/n \rceil$ and arrange an assignment as a table with one row per team and K columns, or rounds, each entry holding one participant. Pad M with $nK - m$ distinct dummy participants, worth zero to every team, indifferent among teams, and allowed only in the last round. Write a_i^t for team i 's entry in round t . Represent each participant's weak order by integer ranks $r_p(i)$, lower being better, and

give dummies rank zero. Put

$$W_t = \sum_{i \in N} v_i(a_i^t), \quad R = \sum_{p \in M} r_p(\text{team receiving } p).$$

Among all padded tables choose one that lexicographically maximizes

$$(W_1, W_2, \dots, W_K, -R), \tag{B.6}$$

and remove the dummies. Each team receives $K - 1$ real participants in the full rounds and at most one more in the last round, so the allocation is balanced. The rounds only organize one final assignment; they are not dates of irrevocable choices.

Lemma B.4 (Envy-freeability). *The selected allocation is envy-freeable for the teams.*

Proof. Fix a round t and a permutation σ of the teams. Reassigning the participants of round t so that team i receives $a_{\sigma(i)}^t$ is feasible and leaves earlier rounds unchanged, so lexicographic optimality gives $\sum_i v_i(a_i^t) \geq \sum_i v_i(a_{\sigma(i)}^t)$. Summing over the rounds,

$$\sum_i v_i(A_i) \geq \sum_i v_i(A_{\sigma(i)}), \tag{B.7}$$

and Lemma 3.2 supplies payments. The participants need not approve these hypothetical permutations; they only prove an inequality. $\square$

Lemma B.5 (EF1 from consecutive rounds). *The selected allocation is EF1. For each team j , deleting its first-round participant a_j^1 removes every other team's envy toward it.*

Proof. Fix teams $i \neq j$ and a round $t < K$, so that a_i^t is real. If a_j^{t+1} is real, swapping it with a_i^t is feasible and changes no earlier round, so optimality gives

$$v_i(a_i^t) \geq v_i(a_j^{t+1}). \tag{B.8}$$

If a_j^{t+1} is a dummy, (B.8) holds because its value is zero. Summing over $t = 1, \dots, K - 1$ pairs team i 's first $K - 1$ positions with team j 's last $K - 1$ positions. Since values are nonnegative,

$$v_i(A_i) \geq v_i(A_i \setminus \{a_i^K\}) \geq v_i(A_j \setminus \{a_j^1\}),$$

where deleting a dummy deletes nothing. $\square$

Lemma B.6 (No beneficial cyclic exchange). *The selected allocation admits no beneficial cyclic exchange.*

Proof. Implement a proposed exchange by placing each incoming participant in the position vacated by the departing one. All moved participants are real, so the padding rule is respected, and each team weakly gains exactly when the value of its affected position weakly increases. If some team strictly gains, the first round with a positive change has a larger total while no earlier round changes, contradicting (B.6). If every team is indifferent, all round totals are unchanged, while some participant strictly gains and none loses. Then R strictly decreases, again a contradiction. $\square$

Proof of Theorem B.3. Apply Theorem 3.5 with teams as agents and participants as items, permitting the allocations that are balanced, EF1, and admit no beneficial cyclic exchange. Multiplying v_i by a positive constant preserves team i 's comparisons between participant sets, and participant preferences and team sizes do not change, so these requirements are scale invariant. At every rescaled profile, the construction run on the rescaled values returns a permitted allocation (Lemmas B.5 and B.6) that is envy-freeable there (Lemma B.4). The bridge gives the lottery. Different scales may produce entirely different tables; the bridge chooses both the assignments and their probabilities. $\square$

Corollary B.7 (Polynomial-time exact lotteries). *For explicitly given rational team values and weak participant orders, the lottery of Theorem B.3 is computable exactly in polynomial time, with rational probabilities and at most $n(n - 1)$ outcomes.*

Proof. A padded table is a perfect matching between nK positions and nK real or dummy participants, with the edges joining dummies to positions outside the last round deleted. The assignment polytope is integral, and so is each of its faces. Optimizing $W_1, \dots, W_K$ and then $-R$ in turn, each over the optimal face of the previous stage, by rational linear programming, returns a vertex implementing (B.6) without infinitesimal tie-breaking. There are $K + 1 \leq m + 1$ stages with polynomially encoded optima. Allocations are described by their m recipients and every envy coordinate is a sum of given rational values, so Theorem 3.8 applies. $\square$

The construction is related to the deterministic team-formation results of Igarashi, Kawase, Suksompong and Sumita [IKSS24], whose matching construction gives priority to individual positions. Prioritizing the total value of a round is what permits the within-round permutations behind (B.7), and hence envy-freeability. The result asserts neither Pareto optimality nor ordinary matching stability, and exchanges that move several participants per team are not excluded.

C Nonadditive applications

C.1 The full submodular share curve

We prove Theorem 4.4 and its curvature refinement using repeated matching. The recorded marginal values give one payment vector for the union comparisons and the full share curve.

C.1.1 Benchmarks, outcome guarantees, and the main conclusion

All valuations in this subsection are normalized, monotone, and submodular, and $U_i = \mathbb{E}[v_i(A_i)]$. The maximum expectation share is defined in (4.1), following the formulation of [FG25], which credits an equal-entitlement variant to [BF24]. Choosing a bundle uniformly from an MMS-optimal partition proves

$$\text{MES}_i \geq \text{MMS}_i. \quad (\text{C.1})$$

Recall that A is MEF1 if, for every $i \neq j$ with $A_j \neq \emptyset$, some $g \in A_j$ satisfies

$$v_i(A_i \cup (A_j \setminus \{g\})) - v_i(A_i) \leq v_i(A_i). \quad (\text{C.2})$$

By monotonicity, $v_i(A_j \setminus \{g\}) \leq 2v_i(A_i)$, so MEF1 implies 1/2-EF1. MEF1, balance, and PROP1 are scale invariant.

Lemma C.1 (From marginal fairness to proportionality). *Every complete MEF1 allocation for monotone submodular goods is PROP1.*

Proof. Fix agent i , put $u = v_i(A_i)$, and suppose there is an outside good. Let

$$\delta = \max_{g \in M \setminus A_i} (v_i(A_i \cup \{g\}) - u).$$

We bound the value of the whole set by adding the other bundles to A_i . The MEF1 witnesses bound each added marginal value. For each nonempty A_j with $j \neq i$, choose the witness in (C.2). Adding this witness back increases value by at most δ , by submodularity. Hence $v_i(A_i \cup A_j) - u \leq u + \delta$. The same bound holds when A_j is empty. Diminishing returns, applied while adding the other bundles, gives

$$v_i(M) - u \leq \sum_{j \neq i} (v_i(A_i \cup A_j) - u) \leq (n-1)(u + \delta).$$

Therefore $v_i(M) \leq nu + (n-1)\delta \leq n(u + \delta)$, so adding a good attaining δ reaches proportionality. If there is no outside good, the agent has the whole set and already meets the benchmark. $\square$

We now prove the stronger share-curve statement from which Theorem 4.4 and Corollary 4.5 follow.

C.1.2 A single lottery for the entire share curve

The proof gives a refinement of Theorem 4.4. For $q \in [0, 1]$, define

$$F_i(q) = \max_{\mathcal{D}: \Pr_{S \sim \mathcal{D}}(g \in S) = q} \mathbb{E}_{S \sim \mathcal{D}}[v_i(S)]. \quad (\text{C.3})$$

This is the MES benchmark at inclusion probability q . In particular, $F_i(0) = 0$, $F_i(1/n) = \text{MES}_i$, and $F_i(1) = v_i(M)$. The parameter q indexes benchmarks for the same equal-entitlement allocation problem. It does not change that problem's entitlements.

Some of a good's value may remain even when all its substitutes are present. Record this value by its smallest marginal contribution,

$$\ell_i(g) = v_i(M) - v_i(M \setminus \{g\}), \quad \ell_i(S) = \sum_{g \in S} \ell_i(g), \quad T_i = \mathbb{E}[\ell_i(A_i)]. \quad (\text{C.4})$$

Monotonicity and submodularity give $0 \leq \ell_i(S) \leq v_i(S)$: add the goods of S one at a time, and each marginal contribution is at least its value in (C.4). Thus $0 \leq T_i \leq U_i$.

Theorem C.2 (Simultaneous share curve). *For normalized monotone submodular valuations, there is one lottery over complete, cardinality-balanced MEF1 allocations satisfying (4.3) and*

$$F_i(q) \leq [1 + (n-1)q]U_i - (1-q)T_i \quad (i \in N, q \in [0, 1]). \quad (\text{C.5})$$

The support can have at most $n^2 + n + 1$ allocations. The same lottery satisfies every inequality in (C.5).

At $q = 1$, (C.5) gives full expected proportionality. At $q = 1/n$, it gives the more informative MES inequality

$$\left(2 - \frac{1}{n}\right)U_i - \left(1 - \frac{1}{n}\right)T_i \geq \text{MES}_i. \quad (\text{C.6})$$

Dropping the nonnegative T_i term from the resulting lower bound on U_i proves (4.2). Monotonicity turns (4.3) into expected 1/2-EF, and Lemma C.1 and (C.1) give the remaining assertions of Theorem 4.4. The retained term improves the guarantee when it is positive. It does not increase the uniform worst-case MES coefficient, since the term can vanish.

Proof of Corollary 4.5. Use the lottery of Theorem C.2. The definition of curvature and diminishing returns imply, for disjoint sets S, T ,

$$v_i(S \cup T) - v_i(T) \geq \ell_i(S) \geq (1 - \kappa_i) \sum_{g \in S} v_i(\{g\}) \geq (1 - \kappa_i)v_i(S). \quad (\text{C.7})$$

The first inequality follows by adding the goods of S one at a time. Every marginal is at least its terminal marginal. Goods of zero singleton value have all marginals zero and cause no exception. For $i \neq j$, taking $S = A_i$, $T = A_j$ in (C.7) and using (4.3) yields $\mathbb{E}[v_i(A_j)] + (1 - \kappa_i)U_i \leq 2U_i$, proving the expected envy bound. The self-comparison is automatic. Taking $T = A_j \setminus \{g\}$ for a MEF1 witness gives the same $1/(1 + \kappa_i)$ factor in every outcome.

Also, (C.7) implies $T_i = \mathbb{E}[\ell_i(A_i)] \geq (1 - \kappa_i)U_i$. Substitution into (C.5) gives

$$F_i(q) \leq [1 + (n-1)q - (1-q)(1 - \kappa_i)]U_i = [nq + \kappa_i(1-q)]U_i.$$

Set $q = 1/n$ for the MES factor and $q = 1$ for PROP. Balance, MEF1, PROP1 and the support bound are unchanged. This argument refines the common-payment lottery. It does not assert envy-freeability under the original submodular valuations. $\square$

C.1.3 The deterministic matching and its common payment certificate

We now prove Theorem C.2. The empty-goods case is immediate, so assume $m > 0$. Pad M by zero-valued dummy goods to a set $\widehat{M}$ of size nR , where $R = \lceil m/n \rceil$. Extend each valuation

by ignoring dummies. In round $r = 1, \dots, R$, assign one remaining good to every agent by a maximum-weight matching. If P_i^{r-1} is agent i 's current bundle, the edge to good g has weight

$$\Delta_i(g \mid P_i^{r-1}) = v_i(P_i^{r-1} \cup \{g\}) - v_i(P_i^{r-1}).$$

Among maximum-weight matchings, choose one assigning as many real goods as possible. Let g_i^r be agent i 's good in round r , $\hat{A}_i = P_i^R$, and $A_i = \hat{A}_i \cap M$.

Lemma C.3 (Fairness of the matching allocation). *The resulting real allocation is complete, cardinality-balanced, and MEF1.*

Proof. If at least n real goods remain and the selected matching uses a dummy, some real good is unselected. Replacing the dummy edge by that good has nonnegative weight gain. A positive gain contradicts maximal weight, and a zero gain contradicts the secondary choice. Thus only the final round can use dummies. Every agent consequently receives $\lfloor m/n \rfloor$ or $\lceil m/n \rceil$ real goods.

For $r < R$, the good g_j^{r+1} is available but unselected in round r . Replacing agent i 's selected edge by that good is a feasible matching. Optimality gives

$$\Delta_i(g_i^r \mid P_i^{r-1}) \geq \Delta_i(g_j^{r+1} \mid P_i^{r-1}). \quad (\text{C.8})$$

Let $B = \hat{A}_j \setminus \{g_j^1\}$. By diminishing returns and $P_i^{r-1} \subseteq \hat{A}_i$,

$$\begin{aligned} v_i(\hat{A}_i \cup B) - v_i(\hat{A}_i) &\leq \sum_{r=1}^{R-1} \Delta_i(g_j^{r+1} \mid \hat{A}_i) \\ &\leq \sum_{r=1}^{R-1} \Delta_i(g_i^r \mid P_i^{r-1}) \leq v_i(\hat{A}_i). \end{aligned}$$

If the first good g_j^1 is real, it is the required deletion witness. If it is a dummy, the bound already holds for the whole real bundle A_j , and removing any real good preserves it. Empty bundles impose no requirement. Removing the dummies proves MEF1 for A . $\square$

For analysis only, record each good's marginal value when it leaves the pool. If g is assigned in round r , put

$$h_i(g) = \Delta_i(g \mid P_i^{r-1}), \quad H_{ij} = \sum_{g \in \hat{A}_j} h_i(g), \quad u_i = v_i(A_i).$$

These recorded values form an auxiliary additive valuation for each observer. They do not replace the original valuation. The own entries telescope, so $H_{ii} = u_i$.

Lemma C.4 (One certificate for unions and all share benchmarks). *There are real payments $p_1, \dots, p_n$ with $\sum_i p_i = 0$ such that the matching allocation satisfies, for every $i \neq j$ and every $q \in [0, 1]$,*

$$v_i(A_i \cup A_j) - 2u_i \leq p_i - p_j, \quad (\text{C.9})$$

$$F_i(q) - [1 + (n-1)q]u_i + (1-q)\ell_i(A_i) \leq nq p_i. \quad (\text{C.10})$$

Proof. In every round, the selected matching defeats every permutation of its selected goods. Summing over rounds yields

$$\sum_i H_{ii} \geq \sum_i H_{i, \sigma(i)} \quad \text{for every permutation } \sigma.$$

In particular, the directed graph with edge weights $H_{ij} - H_{ii}$ has no positive cycle. Let p_i initially be the maximum weight of a walk starting at i , including the empty walk. Removing nonpositive cycles never lowers weight, so the maximum is finite. Prepending edge (i, j) to a maximizing walk from j proves $p_i \geq H_{ij} - H_{ii} + p_j$. Subtracting the mean from all payments therefore gives

$$u_i + p_i \geq H_{ij} + p_j \quad (i, j \in N), \quad \sum_i p_i = 0. \quad (\text{C.11})$$

This is an envy-freeability certificate for the recorded additive values. It does not assert envy-freeability under the original valuations.

Since an agent's prefix at every round is contained in her final bundle, diminishing returns implies

$$v_i(A_i \cup A_j) \leq u_i + \sum_{g \in A_j} \Delta_i(g \mid A_i) \leq u_i + H_{ij}.$$

Combining this with (C.11) proves (C.9). Summing the same payment inequalities over $j \neq i$ gives

$$\sum_{j \neq i} H_{ij} \leq (n-1)u_i + np_i. \quad (\text{C.12})$$

For any two real bundles X, S , submodularity gives the useful bound

$$v_i(S) \leq v_i(X) + \sum_{g \in S \setminus X} \Delta_i(g \mid X) - \ell_i(X \setminus S). \quad (\text{C.13})$$

Indeed, adding $X \setminus S$ to S increases value by at least $\ell_i(X \setminus S)$, whereas adding $S \setminus X$ to X increases value by at most the displayed marginal sum. Both operations end at $X \cup S$.

Set $X = A_i$ and average (C.13) under any bundle distribution with item marginals q . Each outside good appears with probability q , and each own good is absent with probability $1 - q$. The final marginal of each outside good is at most its recorded marginal, so

$$\begin{aligned} \mathbb{E}[v_i(S)] &\leq u_i + q \sum_{g \notin A_i} \Delta_i(g \mid A_i) - (1 - q)\ell_i(A_i) \\ &\leq u_i + q \sum_{j \neq i} H_{ij} - (1 - q)\ell_i(A_i) \\ &\leq [1 + (n-1)q]u_i - (1 - q)\ell_i(A_i) + nq p_i. \end{aligned}$$

Taking the maximum over such bundle distributions proves (C.10). Every displayed inequality uses the same allocation and the same centered payments. $\square$

C.1.4 Scaling cancels all the payments at once

The deterministic construction is available after every independent positive scaling $v_i \mapsto t_i v_i$. Submodularity survives scaling. MEF1 and cardinality balance hold for the original profile whenever they hold for the scaled one. The share curve F_i , the marginal sum ℓ_i , and every value comparison scale by t_i . These are the precise invariance and homogeneity facts needed below.

Proof of Theorem C.2. We first reduce the continuum of benchmark parameters to finitely many tests. The polytope

$$\mathcal{K} = \left\{ (x, q) : x_S \geq 0, \sum_{S \subseteq M} x_S = 1, \sum_{S \ni g} x_S = q \ (g \in M), \ 0 \leq q \leq 1 \right\}$$

has finitely many vertices. Let Q consist of their positive q coordinates, together with 1. The set Q can be exponentially large. This is a finite existence reduction, not a polynomial-time procedure for listing or evaluating the benchmarks. To establish (C.5), it suffices to establish it for $q \in Q$. Indeed, the expression

$$\sum_{S \subseteq M} x_S v_i(S) - [1 + (n-1)q]U_i + (1-q)T_i$$

is affine in (x, q) , so it suffices to bound it at vertices of $\mathcal{K}$. At a positive- q vertex its first term is at most $F_i(q)$. At $q = 0$ it is $-U_i + T_i \leq 0$ for every lottery.

Let $\mathcal{A}$ be the fixed family of complete, cardinality-balanced MEF1 allocations. Apply Proposition 3.7 with

$$\begin{aligned} a_{ij}(v, A) &= v_i(A_i \cup A_j) - 2v_i(A_i), \\ b_{iq}(v, A) &= F_i(q) - [1 + (n-1)q]v_i(A_i) + (1-q)\ell_i(A_i), \quad c_q = nq \quad (q \in Q). \end{aligned}$$

All these expressions are homogeneous in agent i 's valuation, and $c_q > 0$. The matching construction and Lemma C.4 supply one centered payment vector satisfying these inequalities after every positive rescaling. MEF1 and balance keep $\mathcal{A}$ fixed. The proposition therefore gives one lottery satisfying every pairwise and finite share inequality. The vertex reduction extends the latter to all $q \in [0, 1]$. Self-comparisons are automatic.

Finally, preserve the n own-value coordinates $v_i(A_i)$, the n marginal-sum coordinates $\ell_i(A_i)$, and the $n(n-1)$ union-value coordinates $v_i(A_i \cup A_j)$ for $i \neq j$. Every claimed expectation is determined by these $n^2 + n$ coordinates. Carathéodory's theorem therefore compresses the law to at most $n^2 + n + 1$ allocations without changing any guarantee. $\square$

For $n \geq 3$, raising the factor in Theorem 4.4 to the $1 - 1/e$ of independent uniform assignment (Section 4.2) while retaining its realized guarantees is a natural next question.

C.2 Subadditive valuations: fair chances without additivity

We prove Corollary 4.6. All valuations are normalized, monotone, and subadditive, and approximate EFX includes every deletion.

C.2.1 A deterministic response to each test

The allowed outcomes are the complete allocations satisfying both EF1 and half-EFX. The comparisons to balance are $v_i(A_j)/2 - v_i(A_i)$: a positive number means that agent i falls short of the desired expected factor in this outcome. By Theorem 3.1, it is enough to find, for each set of nonnegative weights on these comparisons, one allowed allocation whose weighted sum is nonpositive. We call such an allocation a *response to those weights*.

Here is how to find it. First compute a short list of assignments giving each agent one good. That list provides each agent with a useful lower bound on her own value. Choose one of the assignments, then allocate the remaining goods. A good can safely be added to an agent whom nobody envies: deleting that newly added good preserves EF1. If everyone is envied, rotate bundles around an envy cycle, so the agents involved improve. The challenge is choosing among these possible moves so that the particular weighted sum we were asked to control also improves. We do this with a real-valued function of the current allocation, called a potential, and at every choice keep a successor whose potential is at least its current value.

The initial assignments and the distributions over envy cycles are ingredients of [FMNP24]. The potential below makes each choice deterministically from currently known values. It does not require calculating the expected value of the entire remaining random execution. Once the response algorithm is proved, the general computational reduction will choose the final lottery probabilities.

The proof also works under a slightly more general inequality than subadditivity. To expose exactly what it uses, suppose agent i is given a constant $\gamma_i \geq 1$ such that

$$v_i(S \cup \{g\}) - v_i(S) \leq \gamma_i v_i(\{g\}) \quad (S \subseteq M, g \notin S). \quad (\text{C.14})$$

An item can increase the value of an existing bundle by at most γ_i times its singleton value. Subadditivity gives $\gamma_i = 1$. Write $\beta_i = 1/(1 + \gamma_i)$. Readers interested only in subadditive values can set $\gamma_i = 1$ and $\beta_i = 1/2$ throughout the proof.

Theorem C.5 (Deterministic weighted response). *Under (C.14), for every nonnegative collection λ_{ij} there is a complete allocation that is EF1, is β_i -EFX for every evaluating agent i , and satisfies*

$$\sum_{i \neq j} \lambda_{ij} (\beta_i v_i(A_j) - v_i(A_i)) \leq 0. \quad (\text{C.15})$$

Such an allocation can be found using polynomially many value queries and arithmetic operations. With rational weights, rational supplied constants, and uniformly bounded rational value encodings, the algorithm has polynomial bit complexity.

The weights in (C.15) describe the comparisons that this particular call must answer. They are not lottery probabilities. We construct the allocation from an initial matching, then repeatedly add a good or rotate an envy cycle.

Step 1: obtain a list of starting allocations. If $m \leq n$, partition the goods into singleton and empty bundles and consider the n cyclic assignments of these bundles. Their uniform lottery has identical bundle distributions for every recipient, and every outcome is EFX. A minimum-score outcome answers the test.

Assume $m > n$. We need starting singletons that protect each agent against the goods left for later. Run probabilistic serial on singleton values for one unit of time: each agent consumes a currently most valuable available good at unit speed. The resulting fractional assignment x has row sums one and column sums at most one. Decompose x into a polynomial-size convex combination $\mathcal{D}$ of matchings assigning one good to each agent. This is the usual bipartite matching decomposition: it writes the fractional assignment as an explicit lottery over integral assignments. No initial allocation is assumed to exist without this construction.

Two properties of this initialization will be used. For every agent i , other agent j , and threshold a , her assigned singleton is at least as likely to have value at least a as j 's assigned singleton, when both are evaluated by i . To see this, agent i consumes the set of singletons of value at least a at full speed until that set is exhausted or time one is reached. Agent j cannot consume a greater amount of that set. Matching decomposition preserves these marginal probabilities.

Next, let R contain the goods whose column sums in x are less than one, and set

$$\delta_i = \max_{g \in R} v_i(\{g\}).$$

The set R is nonempty. Each of its goods remains available throughout the consumption process. Thus every initial singleton that i can receive has value at least δ_i . Every good left unassigned in a supported matching belongs to R : a column with sum one must be matched in every positive-probability matching. Consequently, all goods added later have singleton value at most δ_i to agent i .

Step 2: separate the last added good from the comparison. The bound δ_i has two roles: agent i starts with value at least δ_i , and every good added later has singleton value at most δ_i to her. We will keep her own value from decreasing. This lets us control the harm caused by one newly added good separately from the value of the rest of a comparison bundle.

At a partial allocation put $U_i = v_i(A_i)$. For each recipient j , record a bundle Y_j . After adding a good to A_j , let Y_j be the bundle just before that addition. After j receives a bundle in a cycle rotation, let $Y_j = A_j$. Leave records unchanged when the corresponding bundle does not change. Thus the current bundle is either Y_j itself or Y_j with one later-added good, so

$$v_i(A_j) \leq v_i(Y_j) + \gamma_i \delta_i. \quad (\text{C.16})$$

Own values will never decrease, and hence $U_i \geq \delta_i$ throughout.

The function we will maintain rewards own value above δ_i and subtracts any remaining envy of the recorded bundle. Define

$$F_{ij} = U_i - \delta_i - (v_i(Y_j) - U_i)_+, \quad (\text{C.17})$$

$$\Phi = \sum_{i \neq j} \lambda_{ij} \beta_i F_{ij}, \quad (\text{C.18})$$

where $z_+ = \max\{z, 0\}$. The first term measures the value gained above the initial floor. The second charges for any recorded comparison bundle still valued above one's own. For example, if $\delta_i = 2$, $U_i = 5$, and $v_i(Y_j) = 8$, the two terms are three and three, so $F_{ij} = 0$. The factor β_i in Φ converts the eventual bound on $(1 + \gamma_i)U_i - v_i(A_j)$ into the requested comparison $U_i - \beta_i v_i(A_j)$.

At initialization set $Y_j = A_j$. The singleton dominance just proved implies

$$\mathbb{E}_{\mathcal{D}}[(v_i(Y_j) - \delta_i)_+] \leq \mathbb{E}_{\mathcal{D}}[(U_i - \delta_i)_+] = \mathbb{E}_{\mathcal{D}}[U_i] - \delta_i.$$

Since $(v_i(Y_j) - U_i)_+ \leq (v_i(Y_j) - \delta_i)_+$, every $\mathbb{E}_{\mathcal{D}}[F_{ij}]$ is nonnegative. Choose a supported matching maximizing Φ . Its potential is therefore nonnegative. This is a finite deterministic choice from a computed list.

Step 3: choose moves that preserve this bound. We will temporarily consider small distributions over candidate moves, only to prove that at least one candidate has nondecreasing potential. The next lemma gives the property these distributions must have. For any observer i and threshold a , a newly recorded bundle given to j should reach value a no more often than i 's new own bundle reaches that value. Primes denote the values and records after the move.

Lemma C.6 (One-step potential inequality). *Consider a finite distribution of transitions from the current state. Suppose own values never decrease, unchanged bundles retain their records, and, for every i, j and every threshold $a > 0$,*

$$\Pr[v_i(Y'_j) \geq a, A'_j \neq A_j] \leq \Pr[U'_i \geq a]. \quad (\text{C.19})$$

Then $\mathbb{E}[F'_{ij}] \geq F_{ij}$ for every pair, and $\mathbb{E}[\Phi'] \geq \Phi$.

Proof. Fix a pair and abbreviate $U = U_i$, $Y = v_i(Y_j)$, and $\delta = \delta_i$. For $a > \delta$ define

$$f_a(U, Y) = \mathbf{1}[U \geq a] - \mathbf{1}[U < a \leq Y].$$

At a fixed threshold the score is one if own value reaches it, minus one if only the recorded comparison reaches it, and zero if neither does. This three-case description makes the one-step inequality direct.

If $U \geq a$, its value is one and stays one. If $U < a \leq Y$, its value is minus one and cannot decrease. In the remaining case $U < a$ and $Y < a$, the current value is zero. A new event $U' < a \leq Y'$ can then occur only if the comparison bundle changes, because its old record has value below a . Therefore

$$\mathbb{E}[f_a(U', Y')] \geq \Pr[U' \geq a] - \Pr[Y' \geq a, A'_j \neq A_j] \geq 0.$$

Finally,

$$F_{ij} = \int_{\delta_i}^{\infty} f_a(U_i, v_i(Y_j)) da.$$

Integration proves the pairwise claim. All distributions are finite, so these integrals are merely finite sums of lengths of intervals. Nonnegative weighted summation gives the scalar claim. $\square$

If an agent r is unenvied, give her any remaining good. Her newly recorded bundle is the old A_r , worth at most U_i to every observer $i \neq r$. All other records stay fixed. Thus (C.19) holds for this deterministic transition.

Otherwise every agent is envied. Represent strict envy by an edge $i \rightarrow j$ when $v_i(A_j) > v_i(A_i)$. Choose a source strongly connected component G , that is, a strongly connected component with no incoming edge from outside it. Inside such a component each agent can be reached from each other by an envy path. Nobody outside it envies a bundle inside. It has at least two vertices. We need a distribution over its directed simple cycles for which all edges entering a fixed vertex have the same marginal probability.

This equal-probability requirement has a concrete purpose. A vertex labels an existing bundle. If both i and j envy that bundle, we want the candidate rotation to offer it to them equally often. It is exactly the comparison used in the threshold lemma.

Such a distribution is obtained as follows. On the reversed graph, move uniformly from a vertex to one of its original incoming neighbors. A positive stationary distribution exists because the graph is strongly connected. It induces a circulation on the original graph, assigning equal flow to each edge entering the same vertex. Decompose this circulation into weighted simple cycles by repeatedly subtracting the minimum edge flow along a positive-flow cycle. Each subtraction removes an edge, so at most $|E(G)|$ cycles are needed. Normalize the cycle coefficients to sum to one. Edge marginals retain their equality.

Temporarily consider a cycle drawn from this distribution, rotating bundles along its envy edges. Each participating agent receives a strictly more valuable bundle. Reset participants' records to their new full bundles. For a threshold $a > U_i$, any old bundle A_k valued at least a by i is envied by i . If $k \in G$, the source property forces $i \in G$. Whenever another receiver j can get that bundle, the equal incoming-edge marginals imply

$$\Pr[j \text{ receives } A_k] = \Pr[i \text{ receives } A_k].$$

Summing over qualifying k proves (C.19): a simple cycle gives each receiver at most one bundle. The sum for j excludes $k = j$, whereas the sum for i may include it, which can only strengthen the inequality. If i lies outside G , no such high-valued bundle in G exists. Thresholds $a \leq U_i$ are immediate. Lemma C.6 therefore applies. Evaluate the successor potential for every cycle in the decomposition and choose a maximum. Its potential is at least the current one.

Repeat additions to unenvied agents and these maximum-potential cycle choices until all goods have been allocated. Between two additions the bundle partition is fixed. Every cycle participant strictly improves among at most n values of those bundles. Across all agents there can be at most $n(n-1)$ such strict improvements, and hence no more than this number of cycles, before the next addition. Termination is polynomial and does not depend on the sizes of utility gaps.

Step 4: check both kinds of fairness separately. We have maintained $\Phi \geq 0$. Two further checks are needed: all intermediate moves must preserve the restrictions on individual outcomes, and the final potential must imply the requested weighted inequality. Neither check follows just from increasing own values.

Initially each bundle is a singleton, so EF1 and EFX hold. A cycle rotation preserves each of these guarantees, including approximate EFX: every agent's own value weakly increases and the collection of comparison bundles is unchanged.

When a good g is added to an unenvied recipient r , deleting g proves EF1 toward the enlarged bundle. For every other agent i and every deletion h from that bundle, monotonicity and (C.14) give

$$v_i((A_r \cup \{g\}) \setminus \{h\}) \leq v_i(A_r) + \gamma_i v_i(\{g\}) \leq U_i + \gamma_i \delta_i \leq (1 + \gamma_i)U_i.$$

Thus β_i -EFX is preserved. The recipient's comparisons remain valid because her own value increases.

At termination, (C.16) and $U_i \geq \delta_i$ give

$$\begin{aligned} (1 + \gamma_i)U_i - v_i(A_j) &\geq 2U_i - v_i(Y_j) - \delta_i + (\gamma_i - 1)(U_i - \delta_i) \\ &\geq F_{ij}. \end{aligned}$$

Multiply by $\lambda_{ij}\beta_i$ and sum. Since $\Phi \geq 0$, the result is (C.15).

For the computation bound, probabilistic serial and bipartite matching decomposition use polynomially many rational operations. The cycle distribution uses a rational stationary system on at most n vertices and a polynomial-size flow decomposition. At each stage there are polynomially many successor potentials to compare. Standard determinant bounds give polynomial encodings for these rational systems. Together with the transition bound, this proves the bit-complexity claim in Theorem C.5.

Proof of Corollary 4.6. Allow exactly the complete allocations with the two ex-post guarantees proved above. Associate with each allocation its vector of shortfalls, $g_{ij}(A) = \beta_i v_i(A_j) - v_i(A_i)$. Theorem C.5 supplies a response with nonpositive weighted score to every nonnegative test. Thus the hypothesis of Theorem 3.1 is satisfied for these allowed allocations and these expected comparisons. It gives the lottery's existence. Proposition A.2 uses the polynomial-time response algorithm to compute the list of allocations and exact probabilities. For $\beta_i = 1/2$, a coordinate $a/(2b) - c/e$ has numerator $ae - 2bc$ and denominator $2be$, so the uniform coordinate encoding bound $2B + 3$ is sufficient. Supplied rational β_i give the same polynomial bound with their encoding lengths included. There are $n(n - 1)$ coordinates. $\square$

The same argument, with supplied rational γ_i , gives ex-ante β_i -EF with EF1 and β_i -EFX ex post under (C.14). That condition with $\gamma_i = 1$ is strictly weaker than subadditivity (take values 0, 1, 1, 2, 3 on sets of sizes 0, ..., 4), but it is already implicit in the proof of [FMNP24]. The fairness factors are those of [FMNP24]. The computational conclusion answers the explicit-distribution question in Section 5 of its arXiv version, whose sampler may describe an exponential-support law.

C.3 Expected envy and realized maximin shares for submodular goods

We prove Theorem 4.7 in three steps. The potential from Theorem C.5 answers a test under every remaining-item order. A completion argument, using the submodular underestimator and concentration machinery of [FF26], supplies an order with the MMS guarantee. Theorem 3.1 then gives one lottery. The initialization and completion must work for the same allocation policy, which we verify below.

C.3.1 The required response

Valuations are normalized, monotone, and submodular. Write $\mu_i = \text{MMS}_i(M, n)$.

Recall EFL from Section 4.3, in the singleton convention of [FF26]: for each ordered pair i, j , either $|A_j| \leq 1$, or some $g \in A_j$ satisfies

$$v_i(A_i) \geq v_i(A_j \setminus \{g\}), \quad v_i(A_i) \geq v_i(\{g\}). \quad (\text{C.20})$$

EFL implies EF1 and 1/2-EFX: an EFL witness splits a nonsingleton comparison bundle into two pieces, each worth at most one's own bundle. Subadditivity bounds the entire comparison bundle by twice one's own value. Removing any good can only lower this bound. A singleton becomes empty after deletion, and EF1 follows directly from (C.20).

Fix nonnegative test weights λ_{ij} . We will find one complete EFL and 1/8-MMS allocation satisfying

$$\sum_{i \neq j} \lambda_{ij} (2v_i(A_i) - v_i(A_j)) \geq 0. \quad (\text{C.21})$$

C.3.2 Why the chosen initial matching is a greedy allocation

Assume for now that $n \geq 2$ and $m > n$. Fix a strict refinement of each agent's ordering of singletons, choosing between utility ties once and for all. Run probabilistic serial for one unit of time: each agent consumes her favorite available good at unit speed. The resulting fractional matching has row sums one and column sums at most one. Decompose it into a finite lottery $\mathcal{D}$ of matchings that give one distinct good to every agent, retaining only positive coefficients.

Our weighted test will select one matching from this list. The share argument, however, starts with **serial greedy picking**: in some order, each agent takes her favorite remaining singleton. The next lemma shows that these descriptions are compatible for every matching that the test could select.

Lemma C.7 (Supported matchings can be generated greedily). *Every matching in $\mathcal{D}$ can be obtained by serial greedy picking in some agent order, using the fixed refinements of ties.*

Proof. For a good g , let $\tau(g)$ be its exhaustion time, setting $\tau(g) = 1$ if it is still available at time one. If agent i consumes a positive amount of g and prefers h to g , then h is already exhausted when this consumption begins. Positive consumption of g continues later, so $\tau(h) < \tau(g)$.

Fix a matching X in $\mathcal{D}$, and denote agent i 's assigned good by x_i . Every edge $i \rightarrow x_i$ has positive fractional mass. Draw an edge $i \rightarrow j$ if i prefers x_j to x_i . Exhaustion times strictly decrease along these edges, so this directed graph has no cycle. In addition, no agent prefers an unmatched good to her assigned good. Indeed, an unmatched good in a positive-coefficient matching has fractional column sum less than one, hence stays available throughout the consumption process.

Choose an agent with no outgoing edge. Her assigned good is her favorite among all goods. After she takes it, remove her vertex and repeat. The induced graph remains acyclic, and the statement about unmatched goods remains valid. This order reproduces X . Each pick is also utility maximizing before the refinement of ties. $\square$

The resulting greedy order may depend on the comparison weights, because the selected matching may do so. This causes no difficulty: both are fixed before the random completion begins.

C.3.3 A weighted response preserved under every item order

We record explicitly the feature of Theorem C.5 needed for the intersection with MMS. Let R be the goods whose fractional column sums are less than one, and put

$$\delta_i = \max_{g \in R} v_i(\{g\}).$$

Every initial good that i can receive from $\mathcal{D}$ is worth at least δ_i to her. Every good left after a supported matching belongs to R , and is therefore worth at most δ_i to her. These bounds hold for every matching on the list.

At a partial allocation write $U_i = v_i(A_i)$. Alongside each bundle A_j , maintain a record Y_j : after adding a good to A_j , record the bundle just before that addition. After j receives a bundle in a cycle rotation, record its entire new contents. Initially set $Y_j = A_j$, and leave the record unchanged when the bundle does not change. Subadditivity then gives

$$v_i(A_j) \leq v_i(Y_j) + \delta_i, \quad U_i \geq \delta_i. \quad (\text{C.22})$$

The second inequality will be maintained by never decreasing own values. Define

$$F_{ij} = U_i - \delta_i - (v_i(Y_j) - U_i)_+, \quad \Phi = \sum_{i \neq j} \lambda_{ij} F_{ij}, \quad (\text{C.23})$$

where $z_+ = \max\{z, 0\}$. This is the potential in the earlier response proof, with the common factor $1/2$ omitted.

Probabilistic serial makes one's own singleton stochastically dominate another agent's singleton under one's own ranking. In particular,

$$\mathbb{E}_{\mathcal{D}}[(v_i(Y_j) - \delta_i)_+] \leq \mathbb{E}_{\mathcal{D}}[U_i] - \delta_i.$$

Since $U_i \geq \delta_i$, it follows that $\mathbb{E}_{\mathcal{D}}[F_{ij}] \geq 0$. Choose a matching maximizing Φ over $\mathcal{D}$. Its potential is nonnegative, and Lemma C.7 gives a greedy order producing it.

An agent r is **unenvied** if $v_i(A_r) \leq v_i(A_i)$ for every other agent i . Complete the allocation by the following decisions, all made from the revealed history. If an unenvied agent exists, select one *before* the next good is revealed. Otherwise draw the strict envy graph, with an edge $i \rightarrow j$ when $v_i(A_j) > U_i$, and choose a strongly connected component with no incoming edge from outside. Rotate a directed envy cycle within this component that maximizes the next potential among the cycles described below. Repeat until an unenvied recipient exists. Then give her the next good. It is unnecessary to remove other cycles.

Here is why the cycle selection is available. On the reversed graph of the chosen component, take the stationary flow of the walk that moves uniformly to an original incoming neighbor. Reverse its edges back and decompose the resulting circulation into weighted simple cycles. After normalization, all edges entering any particular vertex have the same marginal probability. The threshold calculation in Lemma C.6 shows that this cycle distribution weakly increases each F_{ij} in expectation. Thus a cycle maximizing Φ weakly increases Φ deterministically.

For clarity, the comparison underlying that calculation is

$$\Pr[v_i(Y'_j) \geq a, A'_j \neq A_j] \leq \Pr[U'_i \geq a] \quad (a > 0).$$

An observer outside the chosen component envies none of its bundles. For an observer inside it, every sufficiently valuable bundle offered to j is offered to i with the same edge marginal. Summing

over these bundles proves the displayed inequality. Integrating over thresholds gives the expected potential inequality. This verifies that the selected cycle uses only current bundles and records.

Adding any remaining good to an unenvied recipient also weakly increases Φ : the changed record is the old unenvied bundle, while own values do not decrease. This fact does not depend on the identity of the added good. Between additions, each cycle strictly improves its participants among a fixed finite collection of bundles, so the cycle sequence terminates. Consequently, the recipient can indeed be selected before drawing the next good.

At the end, (C.22) implies

$$2v_i(A_i) - v_i(A_j) \geq 2U_i - v_i(Y_j) - \delta_i \geq F_{ij}.$$

Nonnegative weighted summation proves (C.21). We have established the stronger statement needed below: *every remaining-item order produces a passing response under this policy.*

C.3.4 A share guarantee for every admissible adaptive policy

We next show that the freedom to select an item order suffices for the share guarantee, even though the initial matching and subsequent cycle decisions have already been chosen to answer a weighted test.

Lemma C.8 (Robust completion from a greedy first round). *Fix any serial greedy allocation of one good to each agent. Before each subsequent addition, select an unenvied recipient, using a finite sequence of strict envy-cycle rotations if necessary. Make these decisions before drawing a uniformly random remaining good, and give that good to the selected recipient. Cycle and recipient choices may depend on the entire revealed history and on fixed external data. With probability greater than $1/2$, every agent receives at least $\mu_i/8$. Every execution produces a complete EFL allocation.*

Proof. First check EFL, which will also provide the EF1 inequalities used in the share analysis. Initial bundles are singletons. Agent i 's initial good is worth at least every good left after the greedy round, and her own value never decreases. If g is added to an unenvied bundle, deleting g removes envy toward that bundle, and $v_i(\{g\})$ is at most i 's current own value. These are exactly the two inequalities in (C.20). The recipient's own comparisons remain valid because her value increases. Cycle rotations preserve EFL because they weakly increase all own values and move comparison bundles intact.

We use $\rho = 1/8$ to obtain a uniform bound over the adaptive completion policies that preserve the prescribed weighted test. The constant 0.142 in [FF26] is proved for its allocation procedure. It cannot be imported without verifying that proof for these policies. Our elementary tail estimate and fixed-set union bound below give the smaller constant. We do not claim it is optimal or that the loss is necessary. Agents with zero MMS already meet the required bound. For each other agent, measure values in units of her MMS throughout this proof. This is a normalization for the analysis. It does not change the fixed completion policy or its choices. The target for each such agent is now ρ .

The agents that could fail form a fixed set. Let L contain the agents whose initial singleton is worth strictly less than ρ . Agents outside L already meet the target. For $i \in L$, call a good **heavy** if its singleton value to i is at least ρ , and **light** otherwise. Every heavy good was selected before i 's greedy pick: any heavy good still available then would have been preferable to her chosen singleton. Let H_i be the set of these heavy goods and put $k_i = n - |H_i|$.

It is useful to follow a bundle through changes of ownership. Each initial singleton starts a growing set. Later additions enlarge that set, and rotations move it intact. We refer to this history of a growing set as a **bundle lineage**. There are n such lineages, and k_i of them begin with light goods. All these numbers are fixed before any random completion draw.

If i is in position p of the greedy order, the $n - p + 1$ goods chosen at positions $p, \dots, n$ are light to her. Hence

$$k_i \geq n - p + 1, \quad |\{i \in L : k_i \leq k\}| \leq k \quad (k = 1, \dots, n). \quad (\text{C.24})$$

The second assertion follows because every agent it counts must lie among the last k greedy positions. This deterministic bound will control the sum of the failure probabilities.

Failure is impossible when $k_i \leq 8$. Fix $i \in L$ and suppose that her final value is less than ρ . Her own value was then below ρ throughout. A lineage containing a heavy good cannot receive a second good: immediately before that addition its owner would be unenvied, forcing i 's own value to be at least ρ . Thus every heavy lineage stays a singleton, and i 's final bundle belongs to a light lineage.

Write $k = k_i$. Delete the $n - k$ heavy singleton bundles and their current owners. The remaining allocation into k bundles is EF1. Deleting $n - k$ goods from an n -part MMS partition leaves at least k untouched parts, each worth at least one to i . By placing all remaining goods among those parts, we obtain a k -part partition of the reduced instance whose minimum value is at least one.

For a subadditive observer, any EF1 allocation into k bundles gives at least $1/k$ of that observer's k -part maximin share. To verify this elementary bound, remove one EF1 witness from each nonempty comparison bundle. Some part of a k -part MMS partition avoids all these at most $k - 1$ witnesses. Monotonicity and subadditivity bound its value by the sum of the values of the k allocation bundles after these deletions, which is at most k times one's own value. Consequently $k \leq 8$ contradicts the supposed failure.

Concentration on exactly k fixed MMS parts. For $k \geq 9$, delete every part of i 's original MMS partition that meets H_i , and select exactly k untouched parts $B_1, \dots, B_k$. Each is worth at least one and contains only light goods. This selection is fixed before the random draws. Let $T = \bigcup_{r=1}^k B_r$.

We state precisely the external technical input used here. Apply the submodular underestimator lemma of [FF26, Appendix C.2] to $S \mapsto \min\{v_i(S), 1\}$ on T , and extend the resulting function by ignoring goods outside T . It yields a normalized monotone submodular $w \leq v_i$ such that

$$w(B_r) = 1, \quad w(S) = \sum_{g \in S} w(\{g\}) \quad (S \subseteq B_r), \quad w(S) = w(S \cap T). \quad (\text{C.25})$$

In particular, every singleton has w -value less than ρ and $\sum_{g \in M} w(\{g\}) = k$. This verifies the no-large-item Assumption 5.1 of [FF26], with its parameter $\alpha = \rho$, for the restricted underestimator: goods in T are light and goods outside T have w -value zero. The original profile need not satisfy that assumption. The k disjoint parts have $w(B_r) = 1$, so the source's concentration calculation uses k here in place of its number of parts.

For a good g added to a bundle C , credit that addition by

$$c(g) = \begin{cases} w(C \cup \{g\}) - w(C), & w(C) < \rho, \\ w(\{g\}), & w(C) \geq \rho. \end{cases}$$

Initial singleton credits are $w(\{g\})$. Let W' be the total credit over all goods. The credit and drop-potential calculation of [FF26, Sections 4–5], restricted to these k parts, gives

$$\Pr[W' < kg(\rho) - t] \leq \exp\left(-\frac{t^2}{2\rho k + 2\rho t}\right), \quad g(\rho) = 1 - \rho - \rho \log(1/\rho), \quad t > 0. \quad (\text{C.26})$$

Here is the restriction and adaptation check, including the quantity to which concentration is applied. For $Q \subseteq B_r$, use the source's potential

$$P(Q) = \sum_{s=1}^{|Q|} \frac{1}{s} \mathbb{E}_{S \text{ uniform in } \binom{Q}{s}} [(w(S) - \rho)_+], \quad P(\emptyset) = 0.$$

Its bundle and drop bounds are $P(B_r) \geq g(\rho)$ and $0 \leq P(Q) - P(Q \setminus \{g\}) \leq w(\{g\})$. The running total consists of the credits already accrued plus $\sum_{r=1}^k P(Q_r)$, where Q_r is the unallocated portion of B_r . Initial singleton credits pay for the corresponding drops, so this total starts at least $kg(\rho)$ and ends at W' .

Condition on the revealed history, including all rotations and the chosen recipient. Conditional on the next draw lying in Q_r , it is uniform in Q_r . The source's one-round credit estimate therefore bounds its expected credit below by its expected potential drop. A draw outside T has both zero credit and zero drop. The running total is consequently a submartingale for the adapted policy. If the next good is g , its increment has absolute value at most $w(\{g\})$, and the sum of squared increments is at most

$$\sum_{g \in M} w(\{g\})^2 \leq \rho \sum_{g \in M} w(\{g\}) = \rho k.$$

The one-sided concentration inequality used in [FF26, Theorem 5.7] now gives (C.26). Full acyclicity is unnecessary: the recipient is fixed before the draw, and rotations leave the collection of bundle contents unchanged. Selecting exactly k parts, and giving zero w -value to everything outside them, justifies the parameter k in both the initial bound and the variation bound.

Failure forces the lower-tail event. Suppose again that i finishes below ρ . Every recipient's pre-addition bundle has original value at most i 's own value, hence less than ρ . Since $w \leq v_i$, every credit is then an actual w -marginal. These credits telescope along each lineage. Heavy lineages remain singletons outside T and contribute zero. All goods in the other k lineages are light. EF1 and subadditivity therefore give

$$W' \leq \sum_{C \text{ a final light lineage}} v_i(C) \leq kv_i(A_i) + (k-1)\rho < (2k-1)\rho. \quad (\text{C.27})$$

For a nonempty comparison bundle, its EF1 witness is light and costs less than ρ . An empty bundle needs no witness. We prove concentration without conditioning on failure, and only now include the failure event in this lower-tail event.

Set $t = k(g(\rho) - 2\rho) + \rho$ in (C.26). At $\rho = 1/8$, put $c = g(\rho) - 2\rho = (5 - \log 8)/8$. The elementary estimate $\log 8 < 21/10$ gives $c > 29/80$. Since $t \geq ck$ and $t^2/(k+t)$ is increasing for $t \geq 0$,

$$\frac{t^2}{2\rho k + 2\rho t} = \frac{4t^2}{k+t} \geq \frac{4c^2}{1+c}k > \frac{841}{2180}k > \frac{3}{8}k.$$

Thus

$$\Pr[i \text{ fails}] \leq e^{-3k_i/8} \quad (k_i \geq 9). \quad (\text{C.28})$$

Summing over a fixed set of agents. List the agents in L with $k_i \geq 9$ in increasing order of k_i . By (C.24), the r -th value is at least $\max\{9, r\}$. With $q = e^{-3/8}$, the union bound is at most

$$9q^9 + \sum_{r=10}^{\infty} q^r = q^9 \left(9 + \frac{q}{1-q} \right) < \left(\frac{7}{10} \right)^9 \frac{34}{3} = \frac{686011319}{1500000000} < \frac{1}{2}.$$

For exact numerical checks, the degree-three Taylor sum for $e^{3/8}$ exceeds $10/7$, proving $q < 7/10$. The degree-three Taylor sum for $e^{7/10}$ exceeds two, proving the logarithmic estimate above. Agents with $k_i \leq 8$ cannot fail, and agents outside L already meet their targets. This proves simultaneous success with probability greater than $1/2$. $\square$

The fixed initial set L and the deterministic count (C.24) are the all-agent step in this argument. We use the constituent estimates of Feige and Fine, rather than their all-agent theorem or a selection of agents made after seeing the run. The lemma also permits stopping cycle rotations as soon as an unenvied recipient is available, as required by the potential policy.

C.3.5 From successful completions to a common lottery

Proof of Theorem 4.7. For $n = 1$, give all goods to the sole agent. If $m \leq n$, form singleton and empty bundles and use their n cyclic assignments with equal probabilities. Every recipient then has the same bundle distribution, so the lottery is exactly ex-ante EF and every outcome is EFL. If $m < n$, every n -part partition has an empty part and $\mu_i = 0$. If $m = n$, every n -part partition either has an empty part or consists entirely of singletons. Thus $\mu_i = \min_{g \in M} v_i(\{g\})$, and every assigned singleton meets the full MMS bound. These cases satisfy the theorem.

Now assume $n \geq 2$ and $m > n$. Fix any nonnegative test λ . The construction in Section C.3.3 selects an initial matching and a policy that satisfy (C.21) for every item order. By Lemma C.7, this initial matching can be produced by a fixed greedy order. Its policy selects each recipient before the next draw, so Lemma C.8 applies to that same initialization and policy. Some remaining-item order therefore yields a complete EFL and $1/8$ -MMS allocation that passes the test.

Let $\mathcal{S}$ be the finite set of all complete EFL and $1/8$ -MMS allocations of the given instance. This set is independent of λ . Define the comparison shortfall

$$h_{ij}(A) = \frac{1}{2}v_i(A_j) - v_i(A_i).$$

For every nonnegative λ , we have produced $A \in \mathcal{S}$ with $\sum_{i \neq j} \lambda_{ij} h_{ij}(A) \leq 0$. Theorem 3.1 gives a lottery on $\mathcal{S}$ with $\mathbb{E}[h_{ij}(A)] \leq 0$ for every pair. These are precisely (4.7). Membership in $\mathcal{S}$ gives all the asserted realized guarantees. $\square$

Conditioning the random completion on success alone would not give the expected envy comparisons; the potential holds for every completion, and separation supplies the lottery, with support at most $n(n-1)$ by Lemma A.1.

D Limits of deterministic payment certificates

We prove Theorem 5.1; its obstruction and its lottery are both proved directly.

Pareto optimality in this section compares all complete integral allocations. Goods may be regrouped arbitrarily. Let the singleton values be

	a	b	c	d	e
v_1	38	20	1	20	20
v_2	19	7	12	12	7
v_3	5	3	5	4	3

(D.1)

Goods b and e are identical to every agent. In tables, a string such as ab denotes the bundle $\{a, b\}$. Bundles are always listed in agent order.

D.1 The deterministic obstruction

If an allocation A is envy-freeable, exchanging any two whole bundles cannot raise total welfare. Thus it necessarily satisfies

$$v_i(A_i) + v_j(A_j) \geq v_i(A_j) + v_j(A_i) \quad (i < j). \quad (\text{D.2})$$

This follows either by adding the two payment inequalities, whose payments cancel, or from Lemma 3.2. There are exactly eight EF1 allocations satisfying (D.2). Up to interchanging b and e , they have the five forms in Table 3. Every one is Pareto dominated.

A	Own utilities	Dominating allocation	Own utilities	Forms
(ab, ce, d)	(58, 19, 4)	(ab, de, c)	(58, 19, 5)	2
(ab, de, c)	(58, 19, 5)	(bde, a, c)	(60, 19, 5)	2
(ad, be, c)	(58, 14, 5)	(ab, de, c)	(58, 19, 5)	1
(bd, a, ce)	(40, 19, 8)	(be, a, cd)	(40, 19, 9)	2
(be, cd, a)	(40, 24, 5)	(be, ad, c)	(40, 31, 5)	1

Table 3: All EF1 allocations satisfying the three no-improving-swap inequalities, up to the symmetry $b \leftrightarrow e$. The final column counts labeled allocations. A strict utility improvement is visible in every row.

We give the finite classification explicitly so the obstruction can be checked from the five-good table. Positivity first implies that every agent must receive a nonempty bundle in an EF1 allocation. If an agent were empty, each of the other two bundles could contain at most one good, which cannot account for all five goods. The possible bundle sizes are therefore permutations of $(3, 1, 1)$ and $(2, 2, 1)$, giving 150 labeled allocations.

For every proposed ordered partition, EF1 is exactly the six tests

$$v_i(A_i) \geq v_i(A_j) - \max_{g \in A_j} v_i(\{g\}) \quad (i \neq j). \quad (\text{D.3})$$

The three additional tests (D.2) compare the sum on A_i with the sum on A_j in the corresponding difference row:

	a	b	c	d	e
$v_1 - v_2$	19	13	-11	8	13
$v_1 - v_3$	33	17	-4	16	17
$v_2 - v_3$	14	4	7	8	4

Table 4 lists every possible first bundle, up to $b \leftrightarrow e$, and the results of these tests. For a fixed representative A_1 , enumerate each nonempty proper subset of $M \setminus A_1$ as A_2 and let A_3 be its

First bundle A_1	Choices for A_2	EF1 choices	Also satisfying (D.2): (A_2, A_3)
a	14	6	none
b	14	6	none
c	14	0	none
d	14	6	none
ab	6	4	$(ce, d); (de, c)$
ac	6	1	none
ad	6	2	(be, c)
bc	6	4	none
bd	6	4	(a, ce)
be	6	5	(cd, a)
cd	6	2	none
$abc, abd, abe, acd,$ bcd, bce, bde	2 each	0 each	none

Table 4: An exhaustive classification using (D.3) and (D.2). The displayed first bundles represent all possibilities under $b \leftrightarrow e$. Counts are for the displayed first bundle, with the remaining goods still labeled.

complement. There are 14 choices when $|A_1| = 1$, 6 when $|A_1| = 2$, and 2 when $|A_1| = 3$. Thus each entry is obtained using only the displayed integer comparisons.

For example, when $A_1 = ab$, the six choices for A_2 are c, d, e, cd, ce, de . The choices c, d, ce, de pass EF1. The choice c fails the $(2, 3)$ comparison in (D.2), since the difference row sums are $7 < 12$. The choice d fails it because $8 < 11$. The surviving choices are ce and de , exactly as shown. All other rows use the same six or fourteen tests, or the two choices for a three-good first bundle.

The two tables prove that every EF1 allocation either has an improving two-bundle exchange, or is Pareto dominated. In particular, no EF1 Pareto-optimal allocation is envy-freeable. This establishes the deterministic part of Theorem 5.1.

D.2 A fair lottery that avoids the certificate

Consider the following two allocations, each chosen with probability 1/2:

$$X = (ab, d, ce), \quad Y = (b, ad, ce). \quad (\text{D.4})$$

Only good a changes owner. For any allocation A , let $U(A)_{ij} = v_i(A_j)$ be the value agent i assigns to agent j 's bundle. Directly from (D.1),

$$U(X) = \begin{pmatrix} 58 & 20 & 21 \\ 26 & 12 & 19 \\ 8 & 4 & 8 \end{pmatrix}, \quad U(Y) = \begin{pmatrix} 20 & 58 & 21 \\ 7 & 31 & 19 \\ 3 & 9 & 8 \end{pmatrix}.$$

The resulting expected envy matrix, whose (i, j) entry is $\mathbb{E}[v_i(A_j) - v_i(A_i)]$, is

$$\begin{pmatrix} 0 & 0 & -18 \\ -5 & 0 & -5/2 \\ -5/2 & -3/2 & 0 \end{pmatrix}. \quad (\text{D.5})$$

All entries are nonpositive, so the lottery is ex-ante envy-free.

Both outcomes are EF1. At X , agent 2's envy of agent 1 disappears on deleting a , and her envy of agent 3 disappears on deleting c . All other comparisons at X already have no envy. At Y , agent 1 can delete a from agent 2's bundle and e from agent 3's bundle. Agent 3 can delete a from agent 2's bundle. Every other comparison at Y already has no envy.

Their Pareto optimality also has a direct proof. The own-utility vector at X is $(58, 12, 8)$. Any Pareto dominator must give agent 1 at least 58 and agent 3 at least 8. Neither floor can be met with a singleton. Agent 2 also needs a nonempty bundle, so the floors force bundle sizes $(2, 1, 2)$. Agent 1's pair must be ab , ad , or ae . If she receives ad , agent 2 must receive c to reach 12, leaving agent 3 with be , worth only 6. If she receives ab or ae , agent 2 must receive d : giving agent 2 c leaves agent 3 value 7. The only remaining allocations are X and its b/e relabeling, both with the same own utilities. A strict improvement is impossible.

At Y , the own-utility vector is $(20, 31, 8)$. Agent 3 needs at least two goods to reach 8, and agent 1 needs a nonempty bundle to reach 20. Agent 2 consequently receives at most two goods. The only pairs worth at least 31 to her are ac and ad . Giving her ac leaves b, d, e . After giving one of these to agent 1, the remaining two are worth at most 7 to agent 3. Giving her ad leaves b, c, e , and agent 1 must receive b or e . The two resulting allocations are Y and its b/e relabeling, again with identical own utilities. Thus Y is Pareto optimal. This completes the proof of Theorem 5.1.

Neither outcome in (D.4) is envy-freeable. At X , exchanging agents 2 and 3's bundles raises their total utility by $(19 + 4) - (12 + 8) = 3$. At Y , exchanging agents 1 and 2's bundles raises their total utility by $(58 + 7) - (20 + 31) = 14$. These are precisely the kinds of cycle that payments cannot eliminate. The lottery cancels expected complaints across outcomes instead.

D.3 What the example establishes about the bridge

The universal premise that every positively rescaled additive profile has an EF1, Pareto-optimal, envy-freeable allocation is false: the example fails it already at the identity scaling. Changing the tie-breaking of an EFability construction cannot repair this premise. A proof of exact expected EF with EF1 and Pareto-optimal support for unrestricted additive goods must therefore use a weaker or different certificate, or avoid requiring EFability at every scaling.

Weaker, constraint-relative efficiency does not restore the premise either. Cookson, Shah and Verma [CSV26] obtain balanced EF1 allocations that are fPO within the balanced polytope, which here is $P_{\text{bal}} = \{x \geq 0 : \sum_i x_{ig} = 1, 1 \leq \sum_g x_{ig} \leq 2\}$. All eight candidates in Table 3 are balanced. Four rows show balanced integral dominators, and (ab, de, c) , with utilities $(58, 19, 5)$, is dominated within P_{bal} by $(ab, d + \frac{3}{5}c, e + \frac{2}{5}c)$, with utilities $(58, 96/5, 5)$. Yet both lottery outcomes are balanced, EF1, and fPO in P_{bal} . Allocation Y maximizes welfare with weights $(9, 20, 60)$ even without balance. For X , weights $\alpha = (5, 5, 12)$, prices $p = (126, 36, 60, 60, 36)$ and a row multiplier $\xi = (64, 0, 0)$ satisfy $\alpha_i v_i(g) \leq p_g + \xi_i$, so weighted welfare on P_{bal} is at most $318 + 2 \cdot 64 = 446$, which X attains.

The example nevertheless satisfies the general weighted-response criterion. Fix any nonnegative weights λ_{ij} on the envy comparisons. Equation (D.5) implies that the average, over X and Y , of

$$\sum_{i \neq j} \lambda_{ij} (v_i(A_j) - v_i(A_i))$$

is nonpositive. At least one of those two allocations answers the test. Thus response feasibility can hold for the family of EF1, Pareto-optimal allocations even when the EFability premise fails. The example does not settle the existence problem for arbitrary additive instances.

With two agents the separation cannot occur. Indeed, fix an EF lottery on any finite permitted family unchanged by positive rescaling, and fix positive scales c_1, c_2 . Two-agent envy-freeability at

these scales is equivalent to the single cycle inequality

$$c_1(v_1(A_2) - v_1(A_1)) + c_2(v_2(A_1) - v_2(A_2)) \leq 0.$$

The expectation of this expression under the EF lottery is nonpositive. Some supported allocation therefore satisfies the inequality and is envy-freeable at these scales. The payment bridge gives the converse. This observation does not require additive valuations.

AI methodologies

The key insight concerning the payment bridge was obtained during extended interactions with ChatGPT (OpenAI) regarding the problem of BoBW results for monotone valuations. The insights were subsequently understood from multiple perspectives and significantly generalized and applied to the settings studied here. A first complete draft was shared internally on 5th September. Claude (Anthropic) and ChatGPT also assisted with restructuring, exposition and analyzing various valuation domains.

References

- [ABL⁺26] Haris Aziz, Xiaolin Bu, Xinhang Lu, Simon Mackenzie, Mashbat Suzuki, Biaoshuai Tao, and Toby Walsh. Best-of-both-worlds fairness for mixed goods and chores. arXiv:2607.10232, 2026. <https://arxiv.org/abs/2607.10232>. (Cited on p. 2.)
- [ACIW22] Haris Aziz, Ioannis Caragiannis, Ayumi Igarashi, and Toby Walsh. Fair allocation of indivisible goods and chores. *Autonomous Agents and Multi-Agent Systems*, 36(1):3:1–3:21, 2022. <https://doi.org/10.1007/s10458-021-09532-8>. (Cited on p. 3.)
- [ADSY10] Shipra Agrawal, Yichuan Ding, Amin Saberi, and Yinyu Ye. Correlation robust stochastic optimization. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1087–1096, 2010. <https://doi.org/10.1137/1.9781611973075.88>. (Cited on p. 12.)
- [AFSV24] Haris Aziz, Rupert Freeman, Nisarg Shah, and Rohit Vaish. Best of both worlds: Ex ante and ex post fairness in resource allocation. *Operations Research*, 72(4):1674–1688, 2024. <https://doi.org/10.1287/opre.2022.2432>. (Cited on pp. 1 and 2.)
- [AGM23] Haris Aziz, Aditya Ganguly, and Evi Micha. Best of both worlds fairness under entitlements. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems*, pages 941–948. IFAAMAS, 2023. <https://doi.org/10.5555/3545946.3598733>. (Cited on p. 2.)
- [ALMS25] Haris Aziz, Xinhang Lu, Simon Mackenzie, and Mashbat Suzuki. Fair division with indivisible goods, chores, and cake. arXiv:2511.04891, 2025. <https://arxiv.org/abs/2511.04891>. (Cited on p. 8.)
- [Azi20] Haris Aziz. Simultaneously achieving ex-ante and ex-post fairness. In *Proceedings of the 16th Conference on Web and Internet Economics*, volume 12495 of *Lecture Notes in Computer Science*, pages 341–355. Springer, 2020. https://doi.org/10.1007/978-3-030-64946-3_24. (Cited on p. 2.)
- [Azi26] Haris Aziz. Best-of-both-worlds fairness and Pareto optimality. arXiv:2608.08966, 2026. <https://arxiv.org/abs/2608.08966>. (Cited on p. 14.)

- [BCIZ20] Nawal Benabbou, Mithun Chakraborty, Ayumi Igarashi, and Yair Zick. Finding fair and efficient allocations when valuations don’t add up. In *Proceedings of the 13th International Symposium on Algorithmic Game Theory*, volume 12283 of *Lecture Notes in Computer Science*, pages 32–46. Springer, 2020.
https://doi.org/10.1007/978-3-030-57980-7_3. (Cited on p. 3.)
- [BCL16] Sylvain Bouveret, Yann Chevaleyre, and Jérôme Lang. Fair allocation of indivisible goods. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 12, pages 284–311. Cambridge University Press, 2016. (Cited on p. 3.)
- [BDN⁺20] Johannes Brustle, Jack Dippel, Vishnu V. Narayan, Mashbat Suzuki, and Adrian Vetta. One dollar each eliminates envy. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 23–39. ACM, 2020.
<https://doi.org/10.1145/3391403.3399447>. (Cited on pp. 1, 3, 8, 14, and 22.)
- [BEF21] Moshe Babaioff, Tomer Ezra, and Uriel Feige. Fair and truthful mechanisms for dichotomous valuations. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence*, volume 35, pages 5119–5126, 2021.
<https://doi.org/10.1609/aaai.v35i6.16647>. (Cited on pp. 3 and 23.)
- [BEF22] Moshe Babaioff, Tomer Ezra, and Uriel Feige. On best-of-both-worlds fair-share allocations. In *Proceedings of the 18th Conference on Web and Internet Economics*, pages 237–255. Springer, 2022. (Cited on p. 3.)
- [BF24] Moshe Babaioff and Uriel Feige. Share-based fairness for arbitrary entitlements. arXiv:2405.14575, 2024. <https://arxiv.org/abs/2405.14575>. (Cited on pp. 3 and 29.)
- [BF26] Moshe Babaioff and Gefen Frosh. Near-optimal best-of-both-worlds fairness for few agents. arXiv:2602.14668, 2026. <https://arxiv.org/abs/2602.14668>. (Cited on p. 3.)
- [BKNS22] Siddharth Barman, Anand Krishna, Yadati Narahari, and Soumyarup Sadhukhan. Achieving envy-freeness with limited subsidies under dichotomous valuations. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence*, pages 60–66, 2022. (Cited on p. 3.)
- [Bla56] David Blackwell. An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6(1):1–8, 1956.
<https://msp.org/pjm/1956/6-1/pjm-v6-n1-p01-p.pdf>. (Cited on p. 3.)
- [BLL⁺26] Xiaolin Bu, Zihao Li, Shengxin Liu, Xinhang Lu, and Biaoshuai Tao. Best-of-both-worlds fair allocation of indivisible and mixed goods. In Marios Mavronicolas, Qi Qi, and Grant Schoenebeck, editors, *Proceedings of the 20th Conference on Web and Internet Economics (WINE 2024)*, volume 15534 of *Lecture Notes in Computer Science*, pages 277–294. Springer, 2026.
https://doi.org/10.1007/978-3-032-08560-3_16. (Cited on p. 2.)
- [Bud11] Eric Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
<https://doi.org/10.1086/664613>. (Cited on p. 3.)
- [BV04] Stephen Boyd and Lieven Vandenbergh. *Convex Optimization*. Cambridge University Press, 2004. Section 2.5;
https://web.stanford.edu/~boyd/cvxbook/bv_cvxbook.pdf. (Cited on pp. 3 and 15.)
- [CHR24] Ioannis Caragiannis, Kristoffer Arnsfelt Hansen, and Nidhi Rathi. On the complexity of Pareto-optimal and envy-free lotteries. In *Proceedings of the 23rd International*

- Conference on Autonomous Agents and Multiagent Systems*, pages 244–252. IFAAMAS, 2024. <https://doi.org/10.5555/3635637.3662872>. (Cited on p. 3.)
- [CI21] Ioannis Caragiannis and Stavros Ioannidis. Computing envy-freeable allocations with limited subsidies. In *Proceedings of the 17th Conference on Web and Internet Economics*, pages 522–539. Springer, 2021. (Cited on p. 3.)
- [CKK25] Ioannis Caragiannis, Panagiotis Kanellopoulos, and Maria Kyropoulou. On interim envy-free allocation lotteries. *Mathematics of Operations Research*, 50(4):3257–3282, 2025. <https://doi.org/10.1287/moor.2023.0203>. (Cited on p. 3.)
- [CKM⁺16] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum Nash welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation*, pages 305–322. ACM, 2016. <https://doi.org/10.1145/2940716.2940726>. (Cited on pp. 3 and 11.)
- [CSV26] Benjamin Cookson, Nisarg Shah, and Paritosh Verma. Fair and efficient balanced allocations for additive valuations. arXiv:2608.06325v1, 2026. <https://arxiv.org/abs/2608.06325v1>. (Cited on p. 46.)
- [CT21] Richard Cole and Yixin Tao. On the existence of Pareto efficient and envy-free allocations. *Journal of Economic Theory*, 193:105207, 2021. <https://doi.org/10.1016/j.jet.2021.105207>. (Cited on p. 3.)
- [FF26] Uriel Feige and Assaf Fine. Simultaneous EF1 and approximate MMS allocations for submodular valuations. arXiv:2606.06451, 2026. <https://arxiv.org/abs/2606.06451>. (Cited on pp. 3, 13, 37, 38, 40, 41, and 42.)
- [FG25] Uriel Feige and Vadim Grinberg. Fair allocations with subadditive and XOS valuations. arXiv:2503.10513, 2025. <https://arxiv.org/abs/2503.10513>. (Cited on pp. 3 and 29.)
- [FMNP24] Michal Feldman, Simon Mauras, Vishnu V. Narayan, and Tomasz Ponitka. Breaking the envy cycle: Best-of-both-worlds guarantees for subadditive valuations. In *Proceedings of the 25th ACM Conference on Economics and Computation*, pages 1236–1266. ACM, 2024. <https://doi.org/10.1145/3670865.3673592>. (Cited on pp. 2, 3, 12, 13, 14, 34, and 37.)
- [FP24] Gabriele Farina and Charilaos Pipis. Polynomial-time computation of exact Φ -equilibria in polyhedral games. In *Advances in Neural Information Processing Systems 37*, pages 4207–4234, 2024. https://proceedings.neurips.cc/paper_files/paper/2024/hash/07e5845260f6f1e22c9b4e1471493056-Abstract-Conference.html. (Cited on pp. 3, 9, and 17.)
- [FSV20] Rupert Freeman, Nisarg Shah, and Rohit Vaish. Best of both worlds: Ex-ante and ex-post fairness in resource allocation. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 21–22. ACM, 2020. <https://doi.org/10.1145/3391403.3399537>. (Cited on pp. 1, 2, 3, and 14.)
- [GIK⁺24] Hiromichi Goko, Ayumi Igarashi, Yasushi Kawase, Kazuhisa Makino, Hanna Sumita, Akihisa Tamura, Yu Yokoi, and Makoto Yokoo. A fair and truthful mechanism with limited subsidy. *Games and Economic Behavior*, 144:49–70, 2024. <https://doi.org/10.1016/j.geb.2023.12.006>. (Cited on p. 3.)
- [GS24] Jugal Garg and Eklavya Sharma. Best-of-both-worlds fairness of the envy-cycle-elimination algorithm. arXiv:2410.08986, 2024. <https://arxiv.org/abs/2410.08986>. (Cited on p. 2.)
- [GSX⁺23] Xiaoxi Guo, Sujoy Sikdar, Lirong Xia, Yongzhi Cao, and Hanpin Wang. First-choice maximality meets ex-ante and ex-post fairness. In *Proceedings of the 32nd International*

- Joint Conference on Artificial Intelligence*, pages 2719–2727. IJCAI Organization, 2023. <https://doi.org/10.24963/ijcai.2023/303>. (Cited on p. 2.)
- [HPPS20] Daniel Halpern, Ariel D. Procaccia, Alexandros Psomas, and Nisarg Shah. Fair division with binary valuations: One rule to rule them all. In *Proceedings of the 16th Conference on Web and Internet Economics*, volume 12495 of *Lecture Notes in Computer Science*, pages 370–383. Springer, 2020. https://doi.org/10.1007/978-3-030-64946-3_26. (Cited on p. 2.)
- [HPXZ26] Hadi Hosseini, Shraddha Pathak, Lirong Xia, and Chengkai Zhang. Simultaneous envy and equitability guarantees. arXiv:2608.26410, 2026. <https://arxiv.org/abs/2608.26410>. (Cited on p. 14.)
- [HS19] Daniel Halpern and Nisarg Shah. Fair division with subsidy. In *Proceedings of the 12th International Symposium on Algorithmic Game Theory*, volume 11801 of *Lecture Notes in Computer Science*, pages 374–389. Springer, 2019. https://doi.org/10.1007/978-3-030-30473-7_25. (Cited on pp. 1, 2, 3, 4, 5, and 6.)
- [HSV24] Martin Hoefer, Marco Schmalhofer, and Giovanna Varricchio. Best of both worlds: Agents with entitlements. *Journal of Artificial Intelligence Research*, 80:559–591, 2024. <https://doi.org/10.1613/jair.1.15595>. (Cited on pp. 2, 10, 21, and 22.)
- [IKSS24] Ayumi Igarashi, Yasushi Kawase, Warut Suksompong, and Hanna Sumita. Fair division with two-sided preferences. *Games and Economic Behavior*, 147:268–287, 2024. <https://doi.org/10.1016/j.geb.2024.07.008>. (Cited on pp. 3 and 28.)
- [JLB15] Albert Xin Jiang and Kevin Leyton-Brown. Polynomial-time computation of exact correlated equilibrium in compact games. *Games and Economic Behavior*, 91:347–359, 2015. <https://www.sciencedirect.com/science/article/pii/S0899825613000249>. (Cited on pp. 3, 9, and 17.)
- [KAG⁺25] N. Klein Elmalem, H. Aziz, R. Gonen, X. Huang, K. Kimura, I. Saha, E. Segal-Halevi, Z. Sun, M. Suzuki, and M. Yokoo. Whoever said money won’t solve all your problems? weighted envy-free allocation with subsidy. arXiv:2502.09006, 2025. <https://arxiv.org/abs/2502.09006>. (Cited on p. 18.)
- [KMNP25] Pooja Kulkarni, Ruta Mehta, Vishnu V. Narayan, and Tomasz Ponitka. Online fair division with subsidy: When do envy-free allocations exist, and at what cost? arXiv:2510.13633, 2025. <https://arxiv.org/abs/2510.13633>. (Cited on p. 3.)
- [KMS⁺24] Yasushi Kawase, Kazuhisa Makino, Hanna Sumita, Akihisa Tamura, and Makoto Yokoo. Towards optimal subsidy bounds for envy-freeable allocations. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence*, volume 38, pages 9824–9831, 2024. (Cited on p. 3.)
- [KPV⁺25] Telikepalli Kavitha, Surya Panchapakesan, Rohit Vaish, Vignesh Viswanathan, and Jatin Yadav. Best-of-both-worlds guarantees with fairer endings. arXiv:2507.16209, 2025. <https://arxiv.org/abs/2507.16209>. (Cited on p. 3.)
- [LLSW24] Shengxin Liu, Xinhang Lu, Mashbat Suzuki, and Toby Walsh. Mixed fair division: A survey. *Journal of Artificial Intelligence Research*, 80:1373–1406, 2024. <https://doi.org/10.1613/jair.1.15800>. (Cited on p. 3.)
- [LLTZ26] Z. Lin, S. Liu, B. Tao, and S. Zhou. Almost envy-freeness for additive mixed manna with entitlements: Deterministic and randomized guarantees. arXiv:2609.02724v1, 2026. <https://arxiv.org/abs/2609.02724v1>. (Cited on p. 2.)
- [LMMS04] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM*

- Conference on Electronic Commerce*, pages 125–131. ACM, 2004.
<https://doi.org/10.1145/988772.988792>. (Cited on p. 3.)
- [LMS26] Xinhang Lu, Simon Mackenzie, and Mashbat Suzuki. Optimal subsidy bounds for goods and chores: One dollar each suffices. arXiv:2607.10089, 2026.
<https://arxiv.org/abs/2607.10089>. (Cited on p. 3.)
- [PR08] Christos H. Papadimitriou and Tim Roughgarden. Computing correlated equilibria in multi-player games. *Journal of the ACM*, 55(3):14:1–14:29, 2008.
<https://doi.org/10.1145/1379759.1379762>. (Cited on pp. 3, 9, and 17.)
- [SC25] Ankang Sun and Bo Chen. Randomized strategyproof mechanisms with best of both worlds fairness and efficiency. *European Journal of Operational Research*, 324(3):941–952, 2025. <https://doi.org/10.1016/j.ejor.2025.02.027>. (Cited on p. 2.)
- [Sch03] Alexander Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*. Springer, 2003. Chapter 41. (Cited on p. 26.)