The Power of Sperner's Lemma in Fair Division *

Anna Vakarova

Click here for the latest version.

Abstract

This paper develops a general approach to fair and efficient division that does not rely on equilibrium existence. First, I show that when the set of feasible allocations is compact, convex, and invariant to permutations of agents, and utilities are continuous and concave, there exists an allocation that is envy-free and weak Pareto efficient. The result applies to both problems with divisible resources and to indivisible-resource problems after appropriate convexification. In particular, it applies to problems with infinite-dimensional allocation spaces. The proof searches over Pareto weights and uses a regularized weighted welfare problem together with Sperner's lemma to identify weights at which all agents are simultaneously envy-free. Second, the same construction yields a computational counterpart: in finite-dimensional environments, stopping the Sperner procedure at finite resolution produces approximately envy-free and approximately weak Pareto-efficient allocations in polynomial time for a fixed number of agents, under suitable computational-access assumptions. Third, I combine the infinite-dimensional existence result with the finite-dimensional algorithm to allocate fairly, efficiently, and dynamically. The problem of fairness and efficiency of dynamic allocations arises in transportation, infrastructure, natural-resource management, computing, and has been extensively studied in computer science and engineering. Yet these problems cannot be tackled by present-day economic fair-division techniques. I approximate infinite-horizon environments by increasingly rich finite-dimensional truncations and apply the algorithm to each truncation. The approximation error separates into an algorithmic component and a truncation component, which can be controlled independently. Because the computational cost grows polynomially with the finite-dimensional problem size for a fixed number of agents, this approximation scheme remains computationally meaningful as the horizon expands. Applications include transportation capacity, renewable-resource extraction rights, and computing capacity.

*I am profoundly grateful to Federico Echenique for being the best advisor in the world. I am deeply indebted to Chris Shannon and Fabio Maccheroni for their generosity, insightful advice, and invaluable support. I am especially thankful to Quitze Valenzuela-Stookey and the participants of the Berkeley Theory Lunch for the many stimulating discussions, thoughtful questions, and helpful suggestions. I also thank the visiting speakers at Berkeley who generously shared their time and insights in discussing my research. I am grateful to Christian Bastek for his thoughtful discussion of this paper in Berlin, and to participants at conferences and seminars where this paper was presented for their valuable comments and engaging discussions.

To my Baby Boo,
who made me believe in myself.

1. Introduction

The allocation of scarce resources is a fundamental problem in economics. A central question is whether allocations satisfying desirable normative criteria can be guaranteed to exist in a variety of economic environments. The latter differ in the nature of the resources being allocated, the preferences of the agents, and the constraints governing feasible allocations. An important challenge is therefore to identify general principles that guarantee the existence of desirable allocations across a broad range of economic environments, without relying on arguments specific to each allocation problem. Beyond establishing existence, one may also ask whether these principles provide a way to compute and implement the allocations.

Two central normative criteria are efficiency and fairness. Efficiency requires that gains from trade not be left unrealized, but it is silent about how those gains are distributed. Fairness addresses the distributional issue. The two criteria are independent and may be in conflict: an efficient allocation need not be fair, and a fair allocation need not be efficient.

Who determines whether an allocation is fair? In this paper, I let the agents decide for themselves. An agent is **envy-free** at an allocation if, according to her own preferences, she weakly prefers her own consumption to that of any other agent. An allocation is envy-free if every agent is envy-free. Envy-freeness is the adopted notion of fairness.

The first result of the paper establishes that envy-freeness and efficiency can be achieved simultaneously in a broad class of allocation problems. Rather than relying on arguments tailored to particular environments, I provide a unified existence argument. On the one hand, my result encompasses existing fair-division settings, including standard problems involving divisible and indivisible goods. On the other hand, its generality allows me to extend the analysis of fair division to environments that have received little attention in this literature, including dynamic allocation problems in which scarce resources must be allocated repeatedly over time, and problems involving uncertainty and state-contingent allocations.

Establishing that a desirable allocation exists does not necessarily tell us how to find one. Indeed, even in the simplest static assignment problem of allocating n indivisible objects among n agents, finding an envy-free and Pareto-efficient random allocation is known to be PPAD-hard [Tröbst and Vazirani, 2025]. This means that the problem belongs to a class of computationally difficult problems for which no general efficient solution method is known. Existing algorithms obtain positive results by restricting the allocation environment, relaxing the desired normative criteria, or fixing the number of agents. My second result introduces a new efficient algorithmic approach to fair division that applies to a more general class of allocation environments than existing methods. Moreover, it provides strong approximation guarantees: for a fixed number of agents, the algorithm finds allocations that are arbitrarily close to envy-free and efficient. However, a fixed number of agents is necessary for the above properties to hold. The algorithm is a quantitative version of the existence proof.

For particular applications, the existence and computational results come together in a broad class of dynamic resource-allocation problems. In these environments, allocation decisions made today affect the resources available tomorrow, and feasibility must be maintained across time. Examples include allocating takeoff and landing slots at airports, railway capacity,

port infrastructure, extraction rights in fisheries, and computing capacity. My existence theorem applies directly to these potentially infinite-horizon environments, providing a common guarantee of ex-ante, or overall, envy-freeness and efficiency. To compute approximately desirable allocations, I consider finite-horizon approximations of the original problem and apply my algorithm to these approximations. The algorithm's polynomial dependence on the size of the problem makes this approach computationally meaningful. Together, these results provide a foundation for studying fair division in economic environments that have received relatively little attention in the literature, particularly dynamic resource-allocation problems.

A common approach in the literature of fair division is to come up with a notion of an equilibrium with equal incomes, establish its existence using fixed-point arguments, and then appeal to a version of the first welfare theorem to obtain efficiency.

Instead, the main tool behind my existence result is Sperner's lemma, the classical combinatorial precursor to Brouwer's fixed-point theorem. This approach allows me to dispense with monotonicity and local nonsatiation, assumptions that commonly accompany equilibrium-based arguments and that are violated in some allocation problems, for example, the ones with unit-demand. In particular, the framework treats goods, bads, and mixtures of the two without requiring a separate argument for each case.

This way of proving existence is particularly useful when the allocation space is infinite-dimensional. In such environments, establishing the existence of an equilibrium can itself become technically demanding. The direct argument based on Sperner's lemma bypasses this difficulty. The central idea is to search over the simplex of Pareto weights. For each vector of weights, I maximize a regularized weighted-welfare objective and ask which agents are envy-free at the resulting allocation. Applying Sperner's lemma on increasingly fine triangulations identifies, in the limit, a **special** vector of Pareto weights. Using this vector in the weighted-welfare problem yields an allocation that is both envy-free and weakly Pareto efficient.

While the proof of existence is non-constructive, the ideas in the proof can be used to obtain an algorithm. Indeed, the algorithm is a quantitative version of the same argument. Instead of taking increasingly fine triangulations all the way to the limit, I stop once the triangulation is sufficiently fine. The resulting allocation need not be exactly envy-free and efficient, but the approximation errors can be made arbitrarily small. This is important because computing an exactly envy-free and efficient allocation is intractable even in substantially simpler allocation environments. For a fixed number of agents, stopping short of the limit turns the existence argument into an algorithm that is polynomial in the dimension of the consumption space.

The paper makes two key contributions.

First, for problems with indivisible goods, the existence result extends the analysis from the sets of deterministic feasible allocations that are **finite** to the sets of deterministic feasible allocations that are **compact**. This extension is essential for dynamic allocation problems, in

which a deterministic allocation specifies an entire infinite sequence of consumptions. It allows me to establish the existence of envy-free and efficient allocations in dynamic environments with indivisibilities. Examples include allocating take-off and landing slots at airports over long horizons, storage and parking capacity at ports, railway capacity, and harvesting or extraction rights for renewable natural resources over time.

Second, the algorithmic result, while restricted to finite-dimensional environments, applies to a broader class of feasible sets than existing algorithms, and provides a stronger notion of approximate envy-freeness. These gains come with trade-offs. The algorithm guarantees only a weaker notion of approximate efficiency, and its polynomial-time guarantee holds when the number of agents is fixed.

Both the existence and computational results operate on a **convex** set of feasible allocations and are not intrinsically tied to whether the underlying resources are divisible or indivisible. With divisible goods, convexity is a natural and commonly satisfied assumption on the feasible set.

In the settings where goods are indivisible, there are several ways to convexify the set of feasible allocations, and the distinction between them matters. If the set of feasible 'deterministic' allocations is a compact subset of a metric space, I can pass to the space of probability measures over deterministic allocations. If deterministic allocations are embedded in a vector space, I can instead work with their closed convex hull. In computational applications, the deterministic problem often has additional linear structure, making it natural to work with a computationally tractable fractional relaxation.

These different convexifications also reflect the different requirements of existence and computational results. Working with probability measures allows the existence theorem to accommodate a rich class of preferences over lotteries while requiring very little structure from the underlying deterministic feasible set. Polynomial-time computation necessarily requires more structure: the finite-dimensional feasible set must admit a polynomial-time strong separation oracle. In some indivisible environments such computational access can be obtained directly for the set of admissible marginals; in others, the algorithm can instead be applied to a tractable fractional relaxation.

Working with a fractional relaxation raises a separate question of implementation, for rich and almost exhaustive results on which I refer the reader to Akbarpour and Nikzad [2020] and Budish et al. [2013b].

In dynamic problems, polynomial dependence on the dimension of the consumption space becomes substantive. Think, for example, about allocating airport capacity repeatedly over time. Whether divisible or randomized, the natural allocation is an infinite stream, so the underlying allocation problem is infinite-dimensional. My existence theorem applies directly to this original problem and establishes the existence of an envy-free and efficient allocation.

For computation, I approximate the dynamic problem by increasingly rich finite-dimensional problems and apply my algorithm to each one. There are therefore two distinct sources of approximation. Within each finite-dimensional problem, I compute an approximately envy-free and approximately efficient allocation; at the same time, the finite-dimensional problem itself approximates the original infinite-horizon allocation problem. These two approximations are controlled independently and enter the total error additively.

The paper is organized as follows. Part I develops the existence result. I introduce the general allocation framework, describe the main idea of the proof, establish the supporting results, and prove existence of an envy-free and weak Pareto efficient allocation. Part II turns the existence argument into a quantitative computational result. I develop the approximation algorithm and establish its polynomial-time complexity under suitable computational access to the feasible set. Part III applies the framework to dynamic allocation problems and develops finite-dimensional approximations that connect the infinite-dimensional existence result to the computational results of Part II. The Appendices A-C collect proofs and technical results omitted from the main text.

2. Literature Review

The modern mathematical and economic theory of fair division dates back to Steinhaus [1948]. Since then, the literature has grown into a vast interdisciplinary field, with a renewed surge of interest in recent decades, particularly in computer science.

The fair-division literature can be broadly divided into three strands:

(A) allocation of divisible goods;

(B) allocation of indivisible goods without monetary transfers;

(C) allocation of indivisible goods with monetary transfers.

The following review focuses on strands (A) and (B), which are most closely related to my framework and to which my results directly apply. Strand (C) also studies the existence of envy-free and Pareto-efficient allocations, notably Alkan et al. [1991] and Meertens et al. [2002]. However, these results exploit the presence of a divisible commodity, money, available in sufficient quantities, and rely on assumptions and techniques distinct from mine. In contrast, the approach in this paper introduces no monetary transfers and instead convexifies the allocation problem through randomization. I refer the reader to Meertens et al. [2002] for existence results and earlier references on envy-free and Pareto-efficient allocations with monetary transfers. More recent contributions investigate the role of monetary transfers and subsidies in achieving fairness, including their magnitude, computational complexity, and compatibility with incentive constraints [Aziz, 2021, Goko et al., 2024, Kawase et al., 2024].

2.1 Divisibles

The literature on fair division of divisible goods has a long history and is closely connected to general equilibrium theory.

In mathematics, the classical cake-cutting problem concerns the division of a measurable space among finitely many agents [Steinhaus, 1948, Dubins and Spanier, 1961, Weller, 1985, Akin, 1995, Berliant et al., 1992, Maccheroni and Marinacci, 2003, Hüsseinov and Sagara, 2013]. The set of deterministic feasible allocations is the space of partitions of the underlying space, which can be represented by the families of $\{0,1\}$-valued functions in L^∞ with respect to an appropriate finite measure and relaxed to weak* closed convex hull – the set of partitions of unity. Despite the infinite-dimensional feasible set, the feasibility constraint has a simple structure: the agents' allocations must sum to one almost everywhere.This structure is less suited to capturing the intertemporal feasibility constraints that arise in dynamic allocation problems, particularly over an infinite horizon.

In economics, fair division of divisibles is closely connected to competitive equilibrium with equal incomes (CEEI). Varian [1974] establishes the relationship between competitive equilibrium, envy-freeness, and Pareto efficiency. Under the standard assumptions ensuring equilibrium existence, equal incomes imply envy-freeness, while local nonsatiation guarantees Pareto efficiency. Unlike equilibrium-based approaches, my existence proof does not require monotonicity or local nonsatiation.

Building on the Eisenberg-Gale connection between competitive equilibrium and Nash welfare maximization [Eisenberg and Gale, 1959, Eisenberg, 1961], Bogomolnaia et al. [2017] prove the existence of fair and efficient allocation for goods, bads and mixed manna assuming continuity, concavity and homotheticity of utility representations. My result equally applies to the setting and does not rely on homotheticity.

2.2 Indivisibles

In contrast, the fair division of indivisible goods has attracted substantial attention more recently and presents distinct challenges. In particular, exact envy-freeness may fail to exist when goods must be allocated deterministically, without monetary transfers or randomization. These elementary nonexistence examples have motivated the development of alternative fairness criteria, including envy-freeness up to one good (EF1), envy-freeness up to any good (EFX), envy-freeness up to k goods (EFk), epistemic envy-freeness, proportionality, and maximin share (MMS) guarantees. For a comprehensive survey of this literature and its open problems, see Amanatidis et al. [2023].

Randomization provides an alternative approach to overcoming the nonexistence of envy-free deterministic allocations. In a seminal contribution, Hylland and Zeckhauser [1979] (HZ) study the random assignment of indivisible objects to agents with unit demand. They introduce the pseudomarket approach, in which agents purchase lotteries over objects using equal artificial budgets. Unlike standard competitive equilibria, however, pseudomarket equilibria need not be Pareto efficient due to violation of local non-satiation due to unit demand. HZ proved that there exists an equilibrium that is efficient and envy-free.

Budish et al. [2013a] extend the HZ framework to a broader class of assignment problems, allowing agents' consumption sets to be subject to constraints involving strict subsets of goods. Their equilibrium concept permits free disposal of goods at zero prices, and they impose a unique-bliss-point assumption to establish Pareto efficiency.

Echenique et al. [2021b] allow the set of feasible allocations to have constraints that include strictly more than one agent and strictly less than all agents. In their notion of equilibrium, they have different prices for different constraints, but the resulting equilibrium is efficient and envy-free. They also assume strict monotonicity of agents' preferences on the lower contour set of the set of feasible allocations.

Apart from assuming variations of monotonicity, the sets of feasible allocations that these papers consider require a lot of structure, and every mild relaxation of that structure requires a new fix. This is driven by the attempt to apply equilibrium tools and simultaneously ensure implementation. Equilibrium-based approach requires working with explicit characterization of the set of jointly feasible consumption bundles. Since I do not use it, implicit characterization which ensures implementation is sufficient for my argument. Not insisting on equilibrium-based approach is what allows me to generalize a finite set of allocations with specific polytope structure to a compact, permutation invariant subset of a metric space.

Echenique et al. [2021a] study fair division in the environments where agents do not enter the division on an equal footing – they have potentially different reservation utilities that must be respected to satisfy individual rationality. They come up with a notion of justified envy-freeness and prove that is compatible with efficiency. Even though their setting is different from the one considered in this paper, the proof of existence searches over the simplex of Pareto weights. My proof also searches over the simplex of Pareto weights, dimensionality of which only depends on the number of agents and allows me to study existence in large allocation spaces.

Cole and Tao [2021] generalized the existence of envy-free and efficient allocations even further, to the sets of feasible allocations that consist of arbitrary partitions of a given set of finitely many objects. Agent-specific permutation-invariance of this set is the only assumption that they make. Observe that the set of feasible deterministic allocations is nevertheless finite, which makes the set of feasible allocations finite-dimensional. Even though it is a significant generalization for literature, the environment cannot include many even finite-dimensional problems, for example, dynamic allocation problems of a finite horizon. Their proof is an application of a fixed-point theorem to the product of the feasible set of allocations and pareto weights, which can be technically challenging if the space of allocations is infinite-dimensional.

A few contributions establish existence results in settings that do not fall into any of the above three categories. Bai and Gölz [2022] study the allocation of indivisible goods without lotteries or monetary transfers, establishing the existence of envy-free and Pareto-efficient allocations with high probability when agents' cardinal valuations are randomly drawn from specified distributions. Their result thus relies on randomness in the preference profile, rather than randomization of allocations. In a different direction, Brams and Fishburn [2000] and Edelman and Fishburn [2001] investigate envy-freeness and Pareto efficiency for deterministic allocations of indivisible goods under identical or similar preferences. These contributions illustrate how restrictions on the preference domain can affect the existence of fair and efficient deterministic allocations.

2.3 Algorithms

A related CS literature studies the computational problem of finding fair and efficient random allocations.

Tröbst and Vazirani [2025] prove that computing an exactly envy-free and Pareto efficient lottery is PPAD-complete in standard unit-demand assignment problem with linear utilities. They provide a polynomial-time algorithm that produces a $(2+\varepsilon)$-approximately envy-free, and exactly Pareto efficient allocation. It is polynomial in n, which captures both the number of agents and items. Different from their result, my algorithm holds the number of agents fixed, but for a fixed number of agents, it's time is polynomial in the instance size for a general set of feasible allocations. It finds the allocation that satisfies weaker notion of approximate efficiency, but stronger notion of ϵ-envy-freeness. Their notion gives a constant-factor multiplicative guarantee, whereas I can approach exact envy-freeness with an arbitrary degree of precision.

Garg et al. [2025] study the computational trade-off between envy-freeness and Pareto efficiency. In the cardinal random-assignment model with linear utilities, they provide a polynomial-time algorithm (in n, for n items and agents) that computes an exactly envy-free and $(e^{1/e}-\epsilon)$-approximately Pareto-efficient allocation, extending the result to a more general setting with submodular feasibility constraints. They also consider Fisher markets with monotone concave utilities and establish a polynomial-time algorithm that computes a $(1+\epsilon)$-approximately envy-free and $(2+2\epsilon)$-approximately Pareto-efficient allocation.

My computational result differs in several respects. First, it applies to a more general class of feasible allocation set encompassing both divisible and indivisible goods, subject to the availability of a strong separation oracles. Second, it does not require monotonicity. Finally, for a fixed number of agents, my algorithm achieves arbitrarily small additive violations of envy-freeness and efficiency, rather than fixed multiplicative approximation factors.

Caragiannis et al. [2024] study the computation of envy-free and Pareto-efficient lotteries over a finite family of admissible partitions of indivisible goods, where the bundles of each partition can be assigned arbitrarily among agents and utilities are partition-based and cardinal, and linear in lotteries. They establish that the associated computational problem belongs to PPAD and, for a fixed number of agents, provide a polynomial-time algorithm that computes an exactly envy-free and Pareto-efficient lottery. Their algorithm exploits the structure of the feasible allocation set to obtain a computationally efficient representation of lotteries and uses the geometry of the utility possibility set under linearity of utilities to identify efficient allocations.

My computational result differs in both its guarantees and its domain. For a fixed number of agents, I compute an allocation that is approximately envy-free and approximately efficient over a more general class of feasible allocation sets, allowing concave rather than linear utilities, but not allowing for externalities. This greater generality also comes at the cost of approximation and the assumption of access to a strong separation oracle for the feasible allocation set. Both algorithms have running times polynomial in their respective input representations; in particular, neither guarantees polynomial time in the number of indivisible goods when the representation of admissible allocations is itself exponentially large.

Aziz et al. [2025] study the fair division of homogeneous divisible goods under valuations that are additive across goods but potentially nonlinear in quantities. They establish that deciding the existence of an envy-free and Pareto-efficient allocation is NP-hard for at least three goods, even under one-breakpoint piecewise-constant valuations. In contrast, for a single divisible good with piecewise-linear valuations, they provide a polynomial-time algorithm to decide whether such an allocation exists.

2.4 Incentives

Incentive compatibility is a central objective in market design. Ideally, economists want to design mechanisms that not only map every preference profile into an allocation with desirable properties (in this paper, envy-free and efficeint), but also incentivise agents to report truthfully.

Unfortunately, a substantial literature investigates the incompatibility of strategy-proofness, efficiency, and fairness in allocation mechanisms.

Consider first the random assignment of finitely many indivisible objects to finitely many agents. Zhou [1990] establishes the incompatibility of strategy-proofness, symmetry, and ex-ante Pareto efficiency. Bogomolnaia and Moulin [2001] strengthen this impossibility by replacing ex-ante Pareto efficiency with the weaker requirement of ordinal efficiency, while imposing ordinal equal treatment of equals. Nesterov [2017] further weakens efficiency to ex-post efficiency but strengthens fairness to envy-freeness in the stochastic-dominance order, establishing the incompatibility of these requirements with strategy-proofness. Shende and Purohit [2023] obtain an impossibility result under the still weaker requirement of contention-free efficiency. Similarly, Basteck and Ehlers [2023] establish the incompatibility of strategy-proofness, envy-freeness, and unanimity, a particularly weak efficiency requirement that restricts the mechanism only at preference profiles where agents have distinct first-choice objects.

For divisible goods, Serizawa and Weymark [2003] show that in a standard exchange economy, any strategy-proof and Pareto-efficient mechanism fails to guarantee a positive minimum level of consumption: for every positive bound, there exists a preference profile at which some agent receives a bundle arbitrarily close to the origin. Since in standard economy, preferences are assumed to be strictly monotone and endowment is fully allocated, it is most likely that agent is envious.

In cake-cutting environments, Bu et al. [2023] establish that strategy-proofness, Pareto efficiency, and approximate proportionality (as a notion of fairness) are incompatible restricted and favourable preferences domains.

As an application of the impossibility result of Zhou [1990], Appendix D establishes the incompatibility of strategy-proofness, ex-ante weak Pareto efficiency, and envy-freeness in the expected-utility framework.This impossibility motivates my focus on the joint existence of envy-free and efficient allocations, without imposing incentive compatibility.

A more optimistic perspective emerges when turning from finite to large economies. It is known that in large economies, envy-freeness can also provide approximate incentive guarantees. Hammond [1979] establishes a connection between envy-freeness and strategy-proofness in continuum exchange economies. Building on this insight, Azevedo and Budish [2019] show that envy-freeness implies strategy-proofness in the large for anonymous mechanisms. Their result

formalizes the intuition that, when an individual agent has negligible influence on aggregate allocations, misreporting her preferences is approximately equivalent to obtaining the allocation assigned to another reported type, which is undesirable under envy-freeness. These results suggest that envy-freeness mitigates incentives for strategic misreporting in sufficiently large allocation environments. That being said, in this project I do not provide any formal results on the subject matter.

2.5 Fair Division in Dynamic Allocation problems

In economics literature, the question of fair division in environments with time component has been barely touched upon, except by two papers on failures of fairness in overlapping generations models and a joint working paper by Federico Echenique, Albin Erlanson and me. Isaac and Piacquadio [2015] Establish an impossibility result of efficient distribution of resources while treating equally agents with identical preferences in different generations. Dubey and Ruscitti [2024] demonstrate that in the OLG framework with identical (stationary) well-behaved preferences and allocating a single divisible good, there are economies for which there is no efficient and envy-free allocation.

Allocating finitely many indivisible objects over a finite time horizon in a fair manner has been actively studied in CS. CS literature insists on preserving the indivisibility of the goods. Since an abundance of non-existence of EF and PE allocations in static environments, there was hope that repeated allocation of indivisibles might remedy that. Unfortunately, attempts to find allocations that satisfy both efficiency and envy-freeness witnessed fate similar to their counterparts in static environments. The only paper with positive results on existence of EF and PE allocations is Igarashi et al. [2024]. For specific environments, like the ones with two agents and even number of repetitions, or arbitrary finitely many agents but where the time horizon is a variable of choice, they prove existence of an allocation that is PE and overall EF for additive, both over time and across objects, utilities. Otherwise, the number of impossibility results is overwhelming and the rest of the literature gives up on efficiency and studies existence of allocations under different notions of fairness, which are often modified from their static counter-parts to respond to the time nature of the problem. [Caragiannis and Narang, 2024, Cookson et al., 2025] have mixed existence results in repeated assignments with full information where agents' valuations are over the sequences of items (so-called offline allocation problems). Another brunch studies examines various notions of fairness in online allocation problem, in which indivisible items arrive sequentially, and must be allocated immediately and irrevocably [Elkind et al., 2024, Thakral, 2019].

Part I
Existence

3. Existence theorem

Every allocation problem has two key ingredients: a set of feasible allocations and a domain of agents' preferences. I introduce a general set of feasible allocations that is assumed in the existence theorem. Later in the section, I derive the assumptions on feasible allocations from more basic primitives; I provide concrete examples, and connect my assumptions to the sets of feasible allocations studied in the literature.

Feasible allocations. Let $N = \{1, \dots, n\}$ denote a finite set of agents. Take E to be a metrizable locally convex topological vector space. Let $A \subseteq E^n$ be the set of feasible allocations.

Let $S_N = \{\sigma : N \to N \text{ such that } \sigma \text{ is a bijection}\}$ denote the set of permutations of N. For a permutation $\sigma \in S_N$ and an allocation $a \in A$, let σa denote the allocation obtained by relabeling agents according to σ, so that

$$(\sigma a)_{\sigma(i)} = a_i$$

for every $i \in N$.

I impose three assumptions on A:

1. A is compact;
2. A is convex;
3. A is permutation invariant: for every $a \in A$ and every permutation σ of N, $\sigma a \in A$.

Preferences. Let p_i denote an agent-coordinate projection. The preferences of agent i are defined on $p_i(A)$, and are represented by a continuous, **concave**, real-valued utility function

$$U_i : p_i(A) \to \mathbb{R}.$$

Note that I impose no monotonicity assumptions.

Definitions. An allocation $a = (a_1, \dots, a_n) \in A$ is **envy-free** if

$$U_i(a_i) \geq U_i(a_j) \qquad \text{for every } i, j \in N.$$

It is **weak Pareto efficient** if there is no feasible allocation $a' \in A$ such that

$$U_i(a'_i) > U_i(a_i) \qquad \text{for every } i \in N.$$

Our objective is to establish the existence of an allocation satisfying both properties.

Existence theorem.

Theorem 3.1. *Let $(N, E, A, (U_i : i \in N)$ be an allocation problem satisfying the assumptions above. Then there exists an allocation $a^* \in A$ that is envy-free and weak Pareto efficient.*

Remark 3.2. Before moving on, I would like to emphasize that while compactness and convexity are standard assumptions in many allocation problems, permutation invariance is a substantive symmetry assumption. It requires feasibility to be invariant to the identities of the agents: whenever an allocation is feasible, any relabeling of its recipients is feasible as well. In this sense, no agent is exogenously favored through the feasibility constraints. This symmetry makes envy-freeness a natural fairness criterion and plays an essential role in the existence argument.

In more complex environments, heterogeneity in agents' endowments or productivities may be modeled through agent-specific feasibility constraints that restrict the allocations available to different agents. Such constraints violate permutation invariance. At the end of this section, I provide an example of a setting in which permutation invariance does not hold and in which there is no weakly Pareto efficient and envy-free allocation.

Remark 3.3. Note that I assume that preferences are atomistic. I do not consider allocation problems with externalities.

4. Discussion of the assumptions in Theorem 3.1

The level of generality of the framework introduced above is more abstract than is customary in the recent fair-division literature. But even though the environments look different, the definitions of envy-freeness and weak Pareto efficiency that they use are unchanged. Consequently, the main argument applies to different environments once the relevant feasible allocation space A has been constructed.

Below, I demonstrate how existing models fit into this framework and satisfy its assumptions. I first consider the allocation of indivisible goods, constructing the feasible allocation set A by specifying a set of feasible deterministic allocations and then randomizing over them. I then show how allocation problems with divisible private goods also satisfy the assumptions.

4.1 Feasible allocations with indivisibles

Many allocation problems discussed in the literature concern models with indivisibilities. Here I connect those models with the assumptions in my existence theorem.

Let (Y, d_Y) denote the metric space of deterministic consumption bundles available to every agent in N, and let $Z \subseteq Y^n$ denote the set of feasible deterministic allocations. An element $z = (z_1, \ldots, z_n) \in Z$ specifies a consumption bundle z_i for every agent.

Observe that, even when Y is a vector space, Z may not be convex. When the goods to be allocated are indivisible, the deterministic allocation space is indeed non-convex. In applying my results, I need to convexify the space of allocations. There are different ways of convexifying a set. The one common in the literature and arguably that encodes a rich set of decision-relevant information is randomization. Let $\Delta(Z)$ denote the space of Borel probability

measures on Z. By the Riesz representation theorem, I identify $\Delta(Z)$ with a subset of $C(Z)^*$ and endow it with the weak* topology. A random allocation $\pi \in \Delta(Z)$ induces for every agent i a marginal probability measure over her own consumption bundles. This operation is defined via the marginalization map

$$M : \Delta(Z) \to \Delta(Y)^N \tag{1}$$
$$\pi \to ((p_i)_{\#}\pi)_{i \in N} \tag{2}$$

Define the set of feasible allocations to be the image of $\Delta(Z)$ under the marginalization map:

$$A = M(\Delta(Z))$$

Thus Z denotes the space of feasible deterministic allocations, while A denotes the induced space of feasible random allocations (feasible marginals). I will denote elements of A by $a = (a_1, \ldots, a_n)$.

I impose two assumptions on Z:

1. Z is compact;
2. Z is permutation invariant: for every $z \in Z$ and every permutation σ of N, $\sigma z \in Z$.

These assumptions on Z are needed to ensure that the assumptions on the set of feasible allocations are satisfied.

Proposition 4.1. *Z is compact and permutation invariant iff A is compact and permutation invariant.*

For the proof of Proposition 4.1, see Appendix A. Note that the constructed set A is convex.

4.1.1 Examples and scope

I include examples of the settings with indivisible commodities that have been considered in the literature. All of them are special cases of the previous section.

One-sided assignment In a classical one-to-one assignment problem, n items are to be assigned to n agents, and each of the agents is to receive one item. The set Z is the set of all $n \times n$ permutation matrices. Applications include assigning students to schools, households to housing, or workers to positions. The problem has been studied by Hylland and Zeckhauser [1979], Abdulkadiroğlu and Sönmez [1998], Bogomolnaia and Moulin [2001], Katta and Sethuraman [2006] and others.

Multi-unit assignment A commonly studied class of assignment problem concerns agents who demand multiple indivisible objects, subject to capacity and other feasibility constraints.

There are m indivisible objects to be assigned to n agents. Each object j has $c_j \in \mathbb{N}$ copies. Deterministic feasible allocations are represented by integer-valued assignment matrices satisfying object-capacity constraints and symmetric restrictions on individual consumption.

An important application is university course allocation, where students enroll in several courses subject to seat capacities, scheduling conflicts, and curriculum requirements. This class has been explored by Budish and Cantillon [2012], Budish et al. [2013a], Miralles and Pycia [2021] and others.

Allocation of indivisible bundles There are problems of allocating indivisible bundles that do not fall into the one- or multi-unit assignment category. Examples include dividing household assets and distributing equipment among research laboratories. Here each agent may receive an arbitrary bundle, and their preferences may depend on the entire bundle rather than on the values of individual objects separately. These problems are studied by Cole and Tao [2019] and Caragiannis et al. [2024].

There are m indivisible items to be allocated to n agents and a collection of feasible partitions of the items. Each feasible partition consists of n pairwise disjoint bundles. The bundles in each partition can be assigned to the agents in any order. Thus, the set Z consists of all allocations obtained by permuting the bundles of any feasible partition among the agents.

Differentiated commodities. In the setting of Mas-Colell [1975], commodities are available in integral amounts, while their characteristics vary over a compact metric space (K, d). An individual consumption bundle is a nonnegative, integer-valued Borel measure on K. For a positive integer α, let Y be the set of such measures μ satisfying $\mu(K) \leq \alpha$. The set of integer-valued Borel measures on a compact metric space K that are uniformly bounded by α is compact and metrizable in the weak* topology (see Appendix A). Given a total endowment measure ν, the set of feasible deterministic allocations is

$$Z = \left\{ (\mu^1, \ldots, \mu^n) \in Y^N : \sum_{i \in N} \mu^i \leq \nu \right\}.$$

4.1.2 Preferences

Agent i utility depends only on her marginal probability measure a^i. The preferences of agent i over random consumption bundles are assumed to be represented by a continuous, **concave**, real-valued utility function

$$U_i : \Delta(Y) \to \mathbb{R}.$$

In particular, when preferences admit an expected utility representation, there exists a continuous von Neumann–Morgenstern utility index $u_i : Y \to \mathbb{R}$, with

$$U_i(a_i) = \int_Y u_i(y)\, da_i(y).$$

The resulting utility is linear over $\Delta(Y)$, and therefore concave.

Another class of preferences satisfying the assumptions above is maxmin preferences under risk [Maccheroni, 2002, Cerreia-Vioglio et al., 2015]. Consider a decision maker who is uncertain about the appropriate von Neumann–Morgenstern utility index. Let $\mathcal{U}_i$ be a compact, convex subset of the Banach space $C(Y)$, endowed with the sup norm. Then

$$U_i(a_i) = \min_{u \in \mathcal{U}_i} \int_Y u(y) da_i(y).$$

4.2 Feasible allocations with divisibles

In this section, I discuss allocation problems with divisible private consumption.

Let E denote the agents' consumption space and let $A \subseteq E^n$ be the set of feasible allocations. An element $a = (a_1, \ldots, a_n) \in A$ directly specifies the consumption bundle of every agent without randomization. I assume that A is compact, permutation invariant and convex. Below are several examples of environments with divisible goods.

Divisible goods and mixed manna. In the classical fair division setting of Foley [1967], Varian [1974], or in the mixed manna setting of Bogomolnaia et al. [2017], there are m divisible goods with aggregate endowment $\omega \in \mathbb{R}^d_+$ to be allocated among n agents. The set of feasible allocations is

$$A = \left\{ (a_1, \ldots, a_n) \in (\mathbb{R}^d_+)^n : \sum_{i \in N} a^i = \omega \right\}.$$

In the classical setting, preferences are assumed to be **monotone** and convex, while mixed manna allows items to be a good for some agents and a bad for others. Hence, no monotonicity assumption is made. Preferences are represented by continuous, concave utility functions that are homogeneous of degree 1.

Cake cutting. In cake-cutting problems, the good to be divided is a heterogeneous divisible resource represented by a compact metric space L. Let λ be a finite Borel measure on L. Following Akin [1995], embed consumption bundles into $L^\infty(L, \lambda)$, endowed with the weak* topology, and let

$$Y = \{ f \in L^\infty(L, \lambda) : 0 \leq f \leq 1 \text{ a.e.} \}.$$

A fractional allocation is a partition of unity $a = (a_1, \ldots, a_n)$, and the feasible set is

$$A = \left\{ (a_1, \ldots, a_n) \in Y^N : \sum_{i \in N} a_i = 1 \text{ a.e.} \right\}.$$

With additive preferences, agent i's utility can be represented by

$$U_i(a_i) = \int_L a_i u_i \, d\lambda,$$

where $u_i \in L^1(L, \lambda)$ is her utility density.

4.3 Failure of existence without permutation invariance

The permutation-invariance assumption on the set A is substantive. Consider a setting from Pazner and Schmeidler [1974].[1] There are two consumers $N = \{1, 2\}$ and two goods, leisure l and consumption z. The set of deterministic consumption bundles is

$$Y = \left\{(l, z) \in \mathbb{R}^2_+ : 0 \leq l \leq 1\right\}.$$

The feasible set $A \subseteq Y^2$ is defined by the linear technological constraint

$$z_1 + z_2 - (1 - l_1) - \frac{1}{10}(1 - l_2) \leq 0.$$

That is,

$$A = \left\{((l_1, z_1), (l_2, z_2)) \in Y^2 : z_1 + z_2 - (1 - l_1) - \frac{1}{10}(1 - l_2) \leq 0\right\}.$$

The set A is not permutation invariant. For example,

$$\left((1, \tfrac{5}{10}), (0, \tfrac{5}{10})\right) \in A,$$

whereas

$$\left((0, \tfrac{5}{10}), (1, \tfrac{5}{10})\right) \notin A.$$

Assume that agents' preferences are represented by linear utilities

$$u_1(l, z) = l + \frac{11}{10}z, \qquad u_2(l, z) = l + 2z,$$

Proposition 4.2. *There is no weak Pareto efficient and envy-free allocation.*

Proof. Here is a sketch of the proof. The full proof is in the Appendix A, which builds on Pazner and Schmeidler [1974] who proved that there is no envy-free and Pareto efficient allocation.

The idea is to characterize the Pareto frontier, partition it into two subsets and then observe that every allocation in both of them generates envy. At every Pareto optimal allocation,

1. Either (agent 1 consumes zero leisure AND agent 2 consumes maximum leisure AND the consumption good is arbitrarily split between the two agents), making the two envy-free utility inequalities incompatible for every split of the consumption good.

2. Xor (agent 1 consumes zero leisure AND agent 2 consumes zero of consumption good), generating envy from agent 2 toward agent 1

[1] I am grateful to Shiran Rachmilevitch for the reference.

The rest of the proof shows that for every normalized vector of Pareto weights on the boundary of the simplex (which are $(1,0)$ and $(0,1)$), the agent with zero weight envies the agent with postive weight at the unique maximizer. □

5. Outline of the Argument

The proof of the theorem needs to overcome a basic obstacle. A natural way to obtain an efficient allocation is to maximize aggregate welfare. Consider the problem of maximizing utilitarian welfare:

$$\max_{a\in A} \sum_{i=1}^{n} U_i(a_i).$$

A solution is weak Pareto efficient, but there is no reason for it to be envy-free. Consider, for example, two agents and two indivisible items. For each agent, a and b denote the probabilities of receiving items a and b, respectively, and agents have unit demand. Their utilities are

$$u_1(a,b) = 7a + 3b, \qquad u_2(a,b) = 6a + 4b.$$

The program above has the unique solution $(1,0);(0,1)$. Hence the resulting allocation is weakly Pareto efficient, but agent 2 envies agent 1. Even though every mutually beneficial improvement has been exhausted, some agent may strictly prefer another agent's consumption bundle to her own.

Observe that in the utilitarian objective, all agents' utilities receive equal weight. The fact that the optimal allocation is not envy-free suggests adjusting these weights. Indeed, consider the family of weighted welfare problems

$$\max_{a\in A} \sum_{i=1}^{n} \lambda_i U_i(a_i), \qquad \lambda \in \Delta^{n-1}.$$

where I normalize the space of Pareto weights by taking the simplex $\Delta^{n-1} = \{\lambda \in R^n_+ : \sum_j \lambda_j = 1\}$.

Every solution is weakly Pareto efficient; if the weights are strictly positive, every solution is Pareto efficient. The question is whether I can choose the weights so that the resulting allocation is also envy-free.

Thus, I reduce the problem of existence to the search problem over the simplex of Pareto weights, for a vector of welfare weights that eliminates envy.

There is, however, an immediate difficulty. As the welfare weights λ vary, the set of maximizers can change discontinuously. I therefore regularize the welfare problem. For every $\delta > 0$, I consider a regularized weighted-welfare problem whose solution varies continuously with the weights. The regularization gives a continuous family of allocations indexed by the simplex of welfare weights.

The key observation is simple. At every vector of welfare weights, at least one agent receiving positive weight does not envy anyone. I use such an agent to label the corresponding point of the weight simplex. The resulting labeling satisfies the boundary condition of Sperner's

lemma. A sufficiently fine simplicial subdivision must therefore contain a simplex carrying all n labels. As the mesh of the subdivision shrinks, continuity implies that we approach a single vector of welfare weights at which all agents are simultaneously envy-free.

This argument produces, for every $\delta > 0$, a vector of weights λ_δ whose regularized welfare maximizer is envy-free. Finally, I let the regularization vanish. Compactness and continuity allow me to pass to a limiting welfare weight and allocation. The limiting allocation remains envy-free and maximizes the unregularized weighted welfare objective. It is therefore weakly Pareto efficient.

The remainder of the section formalizes this argument.

6. Ingredients of the Existence Proof

6.1 The regularized welfare problem

Define the normalized space of Pareto weights, the simplex

$$\Delta^{n-1} = \left\{ \lambda \in \mathbb{R}^n_+ : \sum_{j=1}^{n} \lambda_j = 1 \right\}.$$

For a vector of welfare weights $\lambda \in \Delta^{n-1}$ and a regularization parameter $\delta > 0$, consider the regularized objective

$$Q(U, \lambda, \delta, a) = \sum_{i=1}^{n} \lambda_i U_i(a_i) - \delta G(a).$$

To ensure uniqueness of the maximizer, the concave objective is modified by adding a continuous, strictly concave map $-G : A \to \mathbb{R}$. The existence of a strictly concave map in metric spaces is a non-trivial exercise. Moreover, due to the nature of a key lemma, I also require this map to be invariant under permutations: $G(\pi a) = G(a)$ for any permutation π. See Appendix A for the construction of this map.

For the regularized weighted welfare program

$$\max_{a \in A} Q(U, \lambda, \delta, a), \tag{3}$$

denote the solution by $a(\lambda, \delta)$. For the rest of the text, for simplicity of exposition, the dependence of Q on the utilities will be omitted. The reader should keep in mind that all the results hold for a fixed family $(U_i)_{i \in n}$.

The role of regularization is not to alter the weighted welfare objective permanently. It is a device that allows me to search continuously over the simplex of welfare weights. In the proof, I eventually send δ to zero.

6.2 The Lucky Lemma

The next observation provides the labeling used later in the combinatorial argument.

Definition 6.1. The **carrier** of $\lambda \in \Delta^{n-1}$ is

$$\chi(\lambda) = \{i : \lambda_i > 0\}.$$

Lucky Lemma. *Fix $\lambda \in \Delta^{n-1}$ and $\delta > 0$. Let $a(\lambda, \delta) \in A$ solve the program (3). Then there exists an agent $i \in \chi(\lambda)$ who is envy-free at $a(\lambda, \delta)$.*

The original lemma is due Varian [1974]. Here, I present a version of the proof for the setup of this paper.

Proof. Fix $\lambda \in \Delta^{n-1}$ and $\delta > 0$, and write $a = a(\lambda, \delta)$. Suppose, toward a contradiction, that every agent in $\chi(\lambda)$ envies some other agent.

First, an agent in the carrier cannot envy an agent outside of the carrier. Indeed, suppose that $\lambda_i > 0$, $\lambda_j = 0$, and $U_i(a_j) > U_i(a_i)$. Let σ be the permutation that exchanges agents i and j. By permutation invariance, $\sigma a \in A$. Since $\lambda_j = 0$, exchanging the bundles of i and j strictly increases the weighted welfare term in the regularized objective. Moreover, G is permutation invariant, so $G(\sigma a) = G(a)$. Hence, σa yields a strictly larger value of the objective in (3) than a, contradicting the optimality of a.

Thus, every agent in $\chi(\lambda)$ envies another agent in $\chi(\lambda)$. Since $\chi(\lambda)$ is finite, there exists an envy cycle

$$j_1 \to j_2 \to \cdots \to j_k \to j_1$$

consisting entirely of agents in $\chi(\lambda)$. In particular, with $j_{k+1} = j_1$,

$$U_{j_\ell}(a_{j_{\ell+1}}) > U_{j_\ell}(a_{j_\ell}), \qquad \ell = 1, \ldots, k.$$

Let σ be the permutation that assigns to each agent j_ℓ the bundle $a_{j_{\ell+1}}$ and leaves all other agents' bundles unchanged. By permutation invariance, $\sigma a \in A$. Every agent in the cycle strictly prefers her bundle under σa to her bundle under a, while every other agent receives the same bundle. Since every $j_\ell \in \chi(\lambda)$,

$$\sum_{i=1}^{n} \lambda_i U_i((\sigma a)_i) > \sum_{i=1}^{n} \lambda_i U_i(a_i).$$

At the same time, permutation invariance of G implies $G(\sigma a) = G(a)$. Therefore, σa yields a strictly larger value of the regularized objective than a, again contradicting the optimality of a. Hence there exists an agent $i \in \chi(\lambda)$ who is envy-free at $a(\lambda, \delta)$. □

6.3 Simplicial subdivision

In search of the envy-free vector of weights, I employ repeated simplicial subdivision of the simplex of Pareto weights. I choose barycentric subdivision for its well-studied properties.

Recall that Δ^{n-1} is the standard closed $(n-1)$-simplex:

$$\Delta^{n-1} = \text{co}\{e_1, \ldots, e_n\}$$

where co denotes the convex hull and e_j is the jth vector of the standard basis of $\mathbb{R}^n$. A **nonempty face** of Δ^{n-1} is the convex hull

$$F = \mathrm{co}\{e_i : i \in I\}$$

for some nonempty subset $I \subseteq \{1, \ldots, n\}$. The dimension of this face is $|I| - 1$. For example, the vertices are the zero-dimensional faces.

For a nonempty face

$$F = \mathrm{co}\{e_i : i \in I\},$$

its **barycenter** is the point

$$b(F) = \frac{1}{|I|} \sum_{i \in I} e_i.$$

The first barycentric subdivision of Δ^{n-1} is constructed as follows. For every strictly increasing chain of nonempty faces

$$F_0 \subsetneq F_1 \subsetneq \cdots \subsetneq F_k \subseteq \Delta^{n-1},$$

form the simplex

$$\mathrm{co}\Big\{b(F_0), b(F_1), \ldots, b(F_k)\Big\}.$$

The collection of all such simplices covers Δ^{n-1}, and any two of them intersect, if at all, along a common face. Its $(n-1)$-dimensional simplices correspond to maximal chains

$$F_0 \subsetneq F_1 \subsetneq \cdots \subsetneq F_{n-1} = \Delta^{n-1}.$$

Figure 1 illustrates the first barycentric subdivision of Δ^2, with one simplex corresponding to a maximal chain of faces highlighted.

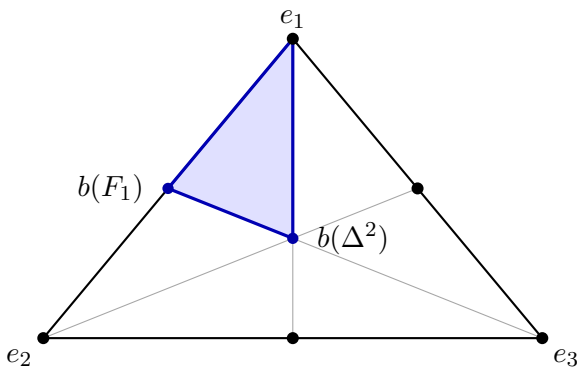


Figure 1: The first barycentric subdivision of Δ^2. The highlighted simplex corresponds to the maximal chain $\{e_1\} \subsetneq \mathrm{co}\{e_1, e_2\} \subsetneq \Delta^2$.

Denote the collection of these $(n-1)$-dimensional simplices by $\mathcal{B}_1(\Delta^{n-1})$, and call it **first barycentric subdivision** of Δ^{n-1}.

The iterated barycentric subdivisions are defined recursively. Set

$$\mathcal{B}_0(\Delta^{n-1}) = \{\Delta^{n-1}\}.$$

Having defined $\mathcal{B}_m(\Delta^{n-1})$, define

$$\mathcal{B}_{m+1}(\Delta^{n-1}) = \bigcup_{S \in \mathcal{B}_m(\Delta^{n-1})} \mathcal{B}_1(S).$$

where S denotes an $(n-1)$-dimensinal simplex and $\mathcal{B}_1(S)$ is its first barycentric subdivision. Notice the subdivision keeps track only of $(n-1)$-dimensional simplices.

Thus, the subdivision of order $m+1$ is obtained by barycentrically subdividing every $n-1$-dimensional simplex in the subdivision of order m. Figure 2 illustrates the second barycentric subdivision of Δ^2.

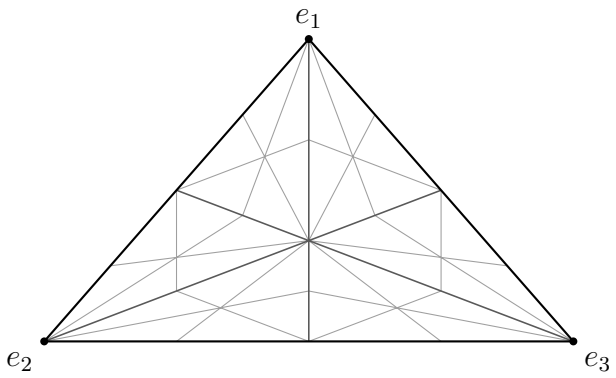


Figure 2: The second barycentric subdivision of Δ^2. The darker edges indicate the first subdivision.

For a simplicial subdivision, **mesh** is the largest diameter of all of the simplices. Under repeated barycentric subdivision, mesh converges to zero. In particular, if T' is a simplex in the first barycentric subdivision of an $n-1$-dimensional simplex T, then

$$\operatorname{diam}(T') \leq \frac{n-1}{n} \operatorname{diam}(T).$$

6.4 Proper labelings and Sperner's lemma

Let

$$\mathcal{B} = \mathcal{B}_m(\Delta^{n-1})$$

be a barycentric subdivision of arbitrary order m, and let $V(\mathcal{B})$ denote the union of the vertex sets of the simplices in $\mathcal{B}$.

Every point $\lambda \in \Delta^{n-1}$ has unique barycentric coordinates

$$\lambda = \sum_{i=1}^{n} \alpha_i e_i, \qquad \alpha_i \geq 0, \qquad \sum_{i=1}^{n} \alpha_i = 1.$$

As a reminder, the **carrier** of λ is the set

$$\chi(\lambda) = \{i \in \{1, \ldots, n\} : \alpha_i > 0\}.$$

Equivalently, the smallest face of Δ^{n-1} containing λ is

$$\text{co}\{e_i : i \in \chi(\lambda)\}.$$

A **labeling** of the subdivision is a map

$$\ell : V(\mathcal{B}) \longrightarrow \{1, \ldots, n\}.$$

The labeling is **proper** if

$$\ell(\lambda) \in \chi(\lambda)$$

for every $\lambda \in V(\mathcal{B})$. Thus, a subdivision vertex lying in a face of Δ^{n-1} may receive only the label of one of the vertices spanning that face. In particular, a proper labeling satisfies

$$\ell(e_i) = i, \qquad i = 1, \ldots, n.$$

Using the Lucky Lemma, construct the **lucky** labeling in the following way: to every $\lambda \in V(\mathcal{B})$, assign the index of an agent who is envy-free at the solution to the regularized weighted welfare program associated with λ. The Lucky Lemma guarantees that this agent can be chosen from $\chi(\lambda)$.

Corollary 6.2. *The lucky labeling is a proper labeling.*

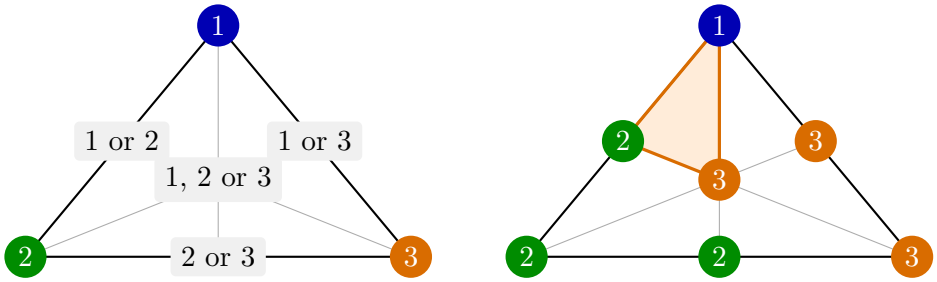
An $(n-1)$-dimensional simplex in $\mathcal{B}$ is **completely labeled** if the labels of its vertices are precisely $\{1, \ldots, n\}$.

Theorem 6.3 (Sperner's lemma)**.** *Let $\mathcal{B}$ be a barycentric subdivision of the simplex Δ^{n-1}, together with a proper labeling. Then $\mathcal{B}$ contains an odd number of completely labeled $(n-1)$-dimensional simplices. In particular, at least one completely labeled simplex exists.*

The theorem is illustrated for the first barycentric division of Δ^2 in Figure 3.

Thus, every barycentric subdivision of Δ^{n-1} contains a completely labeled simplex $\tilde{S}$. $\tilde{S}$ identifies n nearby vectors of Pareto weights, one for each agent. At the vertex labeled i, agent i is envy-free at the allocation that maximizes the corresponding regularized weighted-welfare. Although envy-freeness of each agent need not hold at the same vector of weights, repeated barycentric subdivision makes the diameter of the completely labeled simplex arbitrarily small. This observation is the key to the existence argument: continuity of the regularized optimizer allows the envy-freeness of every agent to be established jointly in the limit.

The procedure above explains why the search is conducted on the entire simplex of Pareto weights, including its boundary. On every face, the Lucky Lemma restricts the labels to agents whose weights are positive. In particular, at an original vertex e_i, the carrier is the singleton i, forcing the lucky labeling to assign label i. By contrast, every weight vector in a



(a) Carrier restrictions (b) Completely labeled simplex

Figure 3: Proper labeling and Sperner's lemma on Δ^2. Panel (a) illustrates the restrictions imposed by the carrier. Panel (b) shows an admissible proper labeling, with a completely labeled simplex highlighted.

simplex contained in the interior of Δ^{n-1} has full support, so the Lucky Lemma imposes no restrictions on the labels assigned to its vertices. Consequently, the Lucky Lemma alone would not guarantee the proper labeling required by Sperner's lemma. The use of the boundary is therefore essential to this argument.

It also helps explain why the result guarantees weak rather than strong Pareto efficiency: the limiting vector of Pareto weights may lie on the boundary of Δ^{n-1} and thus assign zero weight to some agents. Note that if the limiting vector of Pareto weights is in the interior, the resulting allocation is Pareto efficient.

7. Existence of Envy-Free and Efficient Allocation

Equipped with the above, I present a sketch of the proof of the main theorem 3.1. I include the statement of it below for convenience.

Theorem 7.1. *Let E be a metrizable locally convex topological vector space and let $A \subseteq E^n$ be compact convex and permutation invariant. Suppose that each agent's preferences on E are represented by a continuous and concave utility function. Then there exists an allocation $a^* \in A$ that is envy-free and weakly Pareto efficient.*

Proof sketch. Fix $\delta > 0$. Consider a sequence of simplicial subdivisions of Δ^{n-1} whose mesh converges to zero. At every vertex λ, solve the regularized welfare problem and assign as its label an envy-free agent in the support of λ. By the Lucky Lemma, this labeling is well defined and satisfies the Sperner boundary condition.

Sperner's lemma therefore gives a fully labeled simplex in every subdivision. Caution! At this point, a very tempting intuition is to assume that the sequence of fully labeled simplices is **nested**, but that need NOT be the case. Nevertheless, passing to a subsequence, these simplices "shrink" to a common welfare-weight vector λ_δ. By continuity of the regularized solution, the allocations associated with their vertices converge to $a^*(\lambda_\delta, \delta)$. Since every agent appears as a label in every fully labeled simplex, every agent is envy-free at the limiting allocation. Hence, for every $\delta > 0$, there exists $\lambda_\delta \in \Delta^{n-1}$ such that $a^*(\lambda_\delta, \delta)$ is envy-free.

It remains to remove the regularization. Take a sequence $\delta_k \downarrow 0$. By compactness, after passing to a subsequence, the weights λ_{δ_k} and the associated allocations converge to some λ^* and a^*, respectively. Continuity of utilities preserves envy-freeness in the limit. Moreover, as the regularization vanishes, the limiting allocation maximizes the unregularized weighted welfare objective

$$\sum_{i=1}^{n} \lambda_i^* U_i(a_i).$$

Since $\lambda^* \in \Delta^{n-1}$, any strict Pareto improvement over a^* would strictly increase this objective. Therefore, a^* is weakly Pareto efficient. Hence a^* is both envy-free and weakly Pareto efficient. □

The complete proof is deferred to Appendix A.

Part II
Computation

8. Intermezzo

In the next section, I provide a polynomial-time algorithm to compute an allocation that is approximately envy-free and approximately efficient. The algorithm presented there is a quantitative version of the existence proof.

8.1 Finite-dimentional allocation space

The existence result is established at the higher level of generality. Computation necessitates additional structure. Finite-dimensional allocation spaces arise in two ways. First, many economically relevant allocation problems are themselves finite-dimensional. Second, an infinite-dimensional allocation problem can be approximated by a sequence of finite-dimensional problems. The latter interpretation will be particularly useful for the dynamic allocation problems considered in Part III.

I therefore consider a finite-dimensional feasible allocation space

$$A \subseteq (\mathbb{R}^d)^n.$$

8.2 Convexity

The algorithm takes a **convex** feasible allocation space A as its primitive object. For divisible problems, A is the feasible allocation set.

For an indivisible allocation problem, however, there are several distinct ways of obtaining a convex feasible set. One possibility is to work with the set of feasible marginal lotteries $M(Z)$. Another is to work with a convex hull of Z. Yet another is to use fractional relaxation

of Z. These three methods produce different convex sets.

In particular, as it often is in the applications, suppose that Z is of the type $Z = X \cap P$, where P is a polytope that imposes feasibility constraints for some integral $X \subset (\mathbb{R}^d)^n$. For example, in a standard one-sided assignment problem, the deterministic consumption space is the set of vectors of length n with natural-number entries. Let $X = (\mathbb{N}^n)^n$, and define the polytope of constraints as

$$P = \left\{ v \in (\mathbb{R}^n)^n : \begin{array}{ll} \sum\limits_{i \in [n]} v_{ij} \leq 1 & \forall j \in [n], \\ \sum\limits_{j \in [n]} v_{ij} \leq 1 & \forall i \in [n], \\ v_{ij} \geq 0 & \forall i, j \in [n] \end{array} \right\}.$$

The **convex hull** of feasible deterministic allocations is

$$\overline{\text{co}}(X \cap P),$$

whereas the corresponding **fractional relaxation** is

$$\overline{\text{co}}(X) \cap P.$$

For the example above, substochastic extension of the Birkoff-von-Neumann theorem guarantees that $\overline{\text{co}}(X \cap P) = \overline{\text{co}}(X) \cap P$. But there are famous examples of the settings where no version of BvN holds, like multi-unit demand environments.

To illustrate, let G denote a finite set of indivisible objects and let $a_{ig} \in \{0, 1\}$ denote the integral assignment of object g to agent i in a constrained assignment problem. Let P describe the linear constraints

$$P = \Big\{ x \in \mathbb{R}^{N \times G} : \underline{q}_S \leq \sum_{(i,g) \in S} x_{ig} \leq \overline{q}_S, \\ \text{for collections of agent-good pairs in a constraint structure } \mathcal{H} \Big\}.$$

Let $X = \{0, 1\}^{N \times G}$. Then the set of integral feasible allocations is $X \cap P$. Its fractional relaxation, $\overline{\text{co}}(X) \cap P$, is

$$A = \left\{ a \in [0, 1]^{N \times G} : \underline{q}_S \leq \sum_{(i,g) \in S} a_{ig} \leq \overline{q}_S, \text{ for } S \in \mathcal{H} \right\}.$$

Therefore, a fractional allocation need not necessarily be implementable as a lottery over feasible deterministic allocations. In general, only the following weak inclusion holds

$$\overline{\text{co}}(X \cap P) \subseteq \overline{\text{co}}(X) \cap P.$$

The other direction is a question of **implementation**: can an allocation computed in a

fractional relaxation be implemented as a lottery over feasible deterministic allocations? It is a separate question from whether the allocation can be computed efficiently. I treat this as an implementation problem and return to it at the end of the section.

8.3 Computational complexity and approximate solutions

Exact existence does not imply efficient exact computation. Indeed, even in the considerably more structured environment of one-sided cardinal-utility matching, computing an envy-free and Pareto-optimal lottery is PPAD-complete, where computational complexity is parametrized by the number of agents = the number of items = n. To circumvent the issue, polynomial-time algorithms in the literature () hold the number of agents fixed.

A few environments for which polynomial-time algorithms have been developed are too restrictive to directly apply the algorithms to finite prefices of dynamic allocation problems.

To accomodate in particular finite prefices of dynamic allocation problems, but also other more general sets of feasible allocations, I develop a new algorithm that outputs ϵ-EF and ϵ-efficient allocation, where the degree of approximation is additive and arbitrarily small. The output of many of the existing algorithms provides approximate guarantee for at least one of the two notions. For a comprehensive review of the primitives and algorithmic results of related papers in CS and hightlighted differences with my result see Appendix **??**.

The algorithm searches over the simplex of Pareto weights,

$$\Delta^{n-1}.$$

Its running time is exponential in the number of agents, but polynomial in the remaining input parameters and in the desired accuracy when the number of agents is fixed. I therefore treat n as fixed throughout the computational analysis.

9. Computing Approximately Envy Free and Approximately Efficient Allocations

In this section, I study a fixed finite-dimensional allocation problem. The goal is to give an algorithm that outputs an allocation that is approximately envy-free and approximately efficient. The computational-complexity analysis is deferred to a later section.

9.1 Finite-dimensional Setup

There is a finite set of agents $N = \{1, \dots, n\}$. Let $Y \subset \mathbb{R}^d$ denote the agents' consumption space and let

$$A = \begin{cases} Z \subset Y^n, & \text{if } Y \text{ is convex,} \\ M(\Delta(Z)),\ Z \subset Y^n, & \text{if } Y \text{ finite.} \end{cases}$$

be the set of feasible allocations. It will be convenient to identify an allocation $a = (a_1, \dots, a_n) \in (\mathbb{R}^d)^n$ with the $d \times n$ matrix whose i-th column is a_i. I equip this space with the Frobenius

inner product

$$\langle a, b\rangle_F = \sum_{i=1}^{n} \langle a_i, b_i\rangle$$

and the associated Frobenius norm

$$\|a\|_F = \left(\sum_{i=1}^{n} \|a_i\|_2^2\right)^{1/2}.$$

Equivalently, under the natural identification of $(R^d)^n$ with $\mathbb{R}^{dn}$, these coincide with the usual Euclidean inner product and Euclidean ℓ_2-norm on $\mathbb{R}^{dn}$.

Throughout the section, I impose the following assumptions on A:

1. A is nonempty, compact, and convex;
2. A is permutation invariant;
3. $B < \infty$ is explicitely given such that

$$\|a\|_F^2 \leq B \qquad \text{for every } a \in A.$$

For each agent $i \in N$, preferences are represented by a utility function $U_i : Y \to \mathbb{R}$. I assume that each U_i is concave. The continuity assumption is stronger than for the existence result: utilities are assumed to be Lipschitz continuous, that is there exists $L < \infty$ such that for each $i \in N$,

$$|U_i(x) - U_i(y)| \leq L\|x - y\|_2 \qquad \text{for all } x, y \in Y.$$

9.2 Approximate fairness and efficiency

Definition 9.1. For $\varepsilon > 0$, an allocation $a \in A$ is **ε-envy-free** if envy is bounded by ϵ

$$U_i(a_j) - U_i(a_i) \leq \varepsilon \qquad \text{for all } i, j \in N.$$

Definition 9.2. For $\varepsilon > 0$, an allocation $a \in A$ is **ε-approximately efficient** if inefficiency is bounded by ϵ

$$\text{for all } a' \in A, \text{ there exists agent } i \text{ with } U_i(a'_i) - U_i(a_i) \leq \varepsilon\ .$$

It means that at any feasible allocation, there exists an agent i that improves by at most ϵ.

Remark 9.3. The definition above bounds inefficiency in the utility space. Another common way to bound inefficiency is to make the distance from an allocation to the set of weakly Pareto efficient allocations small. Bounding inefficiency in the utility space is natural for the argument of this paper. First, the desirable properties of the allocation are defined in the utility space. Second, regularized weighted welfare objective provides a bound on the size of a maximal utility improvement. Third, under regularization, it does not hold in general that the

resulting allocation is close to the set of maximizers of the unregularized weighted welfare in the allocation space.

9.3 The regularized welfare problem

Following along the lines of Part I, for a vector of welfare weights $\lambda \in \Delta^{n-1}$ and a regularization parameter $\delta > 0$, consider the regularized objective

$$Q_\delta(\lambda, a) = \sum_{i=1}^{n} \lambda_i U_i(a_i) - \delta G(a)$$

For the regularized weighted welfare program

$$\max_{a \in A} Q_\delta(\lambda, a), \tag{4}$$

denote the solution by $a_\delta(\lambda)$.

Throughout algorithmic part, I work with the specific choice

$$G(a) := \|a\|_F^2.$$

Since the regularized objective is additively separable, with the first term being concave and penalization being strongly concave, the mapping is strongly concave. Strong concavity of the objective allows me to quantitatively control the continuity of the maximizer. Quadratic form of the penalty is not necessary to ensure strong concavity, but will be used for computational tractability later.

In the sections below, I lay the groundwork for proving that the algorithmic procedure indeed provides the approximation guarantees. The ϵ errors are controled separately for EF and efficiency.

9.4 Lipschitz continuity of the regularized maximizer

Lemma 9.4. *For every $\delta > 0$, the maximizer map $a_\delta : \Delta^{n-1} \to A$ is Lipschitz continuous. In particular,*

$$\|a_\delta(\lambda) - a_\delta(\lambda')\|_F \leq C_\delta \|\lambda - \lambda'\|_2$$

where $C_\delta = \frac{L}{2\delta}$.

Proof. Denote the maximizers by $a = a_\delta(\lambda)$ and $b = a_\delta(\lambda')$. Since $Q_\delta(\lambda, a)$ is 2δ-strongly concave in a, using standard optimization argument I get

$$Q_\delta(\lambda, a) - Q_\delta(\lambda, b) \geq \delta \|a - b\|_F^2,$$

and similarly

$$Q_\delta(\lambda', b) - Q_\delta(\lambda', a) \geq \delta \|a - b\|_F^2,$$

Adding them, I obtain

$$2\delta\|a-b\|_F^2 \le \sum_i (\lambda_i - \lambda_i')(U_i(a_i) - U_i(b_i))$$

Using Lipschitz continuity of $\{U_i\}_{i\in N}$ and Cauchy-Schwartz, I can bound

$$|\sum_i (\lambda_i - \lambda_i')(U_i(a_i) - U_i(b_i))| \le \|\lambda - \lambda'\|_2 L\|a-b\|_F$$

Finally, dividing by $\|a-b\|_F$ on both sides, I conclude that

$$\|a_\delta(\lambda) - a_\delta(\lambda')\|_F \le \frac{L}{2\delta}\|\lambda - \lambda'\|_2$$

□

9.5 Approximate envy-freeness

Lemma 9.5 (Approximate envy)**.** *Let T' be a fully labeled simplex in the subdivision of Δ^{n-1}, and let $\widetilde{\lambda} \in T'$. Then, for every pair of agents $i, j \in N$,*

$$U_i\big(a_\delta(\widetilde{\lambda})_j\big) - U_i\big(a_\delta(\widetilde{\lambda})_i\big) \le K_\delta \operatorname{diam}(T'),$$

where $K_\delta = \frac{L^2}{\delta}$.

Proof. Since T' is fully labeled, it has a vertex λ^i labeled by i. By the construction of the labeling rule, agent i is envy-free at $a_\delta(\lambda^i)$; so

$$U_i\big(a_\delta(\lambda^i)_j\big) - U_i\big(a_\delta(\lambda^i)_i\big) \le 0 \qquad \text{for every } j \in N.$$

Using the triangle inequality, it follows that

$$\begin{aligned} &U_i\big(a_\delta(\widetilde{\lambda})_j\big) - U_i\big(a_\delta(\widetilde{\lambda})_i\big) \\ &\le \Big|U_i\big(a_\delta(\lambda^i)_j\big) - U_i\big(a_\delta(\widetilde{\lambda})_j\big)\Big| + \Big|U_i\big(a_\delta(\lambda^i)_i\big) - U_i\big(a_\delta(\widetilde{\lambda})_i\big)\Big|. \end{aligned}$$

Since U_i is L-Lipschitz and a_δ is C_δ-Lipschitz,

$$\Big|U_i\big(a_\delta(\lambda^i)_j\big) - U_i\big(a_\delta(\widetilde{\lambda})_j\big)\Big| \le LC_\delta\|\lambda^i - \widetilde{\lambda}\|_2$$

I conclude that

$$U_i\big(a_\delta(\widetilde{\lambda})_j\big) - U_i\big(a_\delta(\widetilde{\lambda})_i\big) \le \frac{2L^2}{2\delta}\operatorname{diam}(T').$$

□

Remark 9.6 (Continuity VS Lipschitz continuity). Note that I require stronger control on the continuity of the utilities and maximizer for the approximation result compared with existence result. Why so?

Mathematically, continuity of utilities and continuity of the argmax, guaranteed by Berge's theorem, is enough for the argument above (see Appendix B for the proof). Continuity of

envy is uniform on compact simplex Δ^{n-1}. Therefore, for every $\epsilon > 0$ there exists σ such that if $d(\lambda, \lambda_{EF(i)}) < \sigma$, then envy of agent i is bounded by ϵ. Taking λ to be in a fully labeled subsimplex of small enough diameter bounds envy of all agents by ϵ. Notice that for every regularization δ, that σ depends on a_δ in a non-explicit way. Asking for it as an algorithmic input is substantially stronger/more black-box than asking for a Lipschitz constant on the primitives. Computing σ using the Lipschitz constant is easy.

9.6 Approximate efficiency

The appropriate notion of approximate efficiency makes the argument particularly simple. Under the uniform-improvement criterion used below, the efficiency loss from regularization is controlled by the range of the penalty function on the feasible set. Define

$$\operatorname{range}_A G := \sup_{a \in A} G(a) - \inf_{a \in A} G(a).$$

The following lemma makes this precise.

Lemma 9.7. *Let $G : A \to \mathbb{R}$ be bounded. For every $\lambda \in \Delta^{n-1}$ and $\delta > 0$,*

$$\sup\{t \geq 0 : \exists z \in A \text{ such that } U_i(z_i) \geq U_i(a_\delta(\lambda)_i) + t \text{ for all } i\} \leq \delta \operatorname{range}_A G.$$

Proof. By optimality of $a_\delta(\lambda)$, for every $z \in A$,

$$\sum_i \lambda_i \Big(U_i(z_i) - U_i(a_\delta(\lambda)_i)\Big) \leq \delta \Big(G(z) - G(a_\delta(\lambda))\Big) \leq \delta \operatorname{range}_A G.$$

If

$$U_i(z_i) \geq U_i(a_\delta(\lambda)_i) + t \qquad \text{for all } i,$$

then, since $\lambda \in \Delta^{n-1}$,

$$t \leq \sum_i \lambda_i \Big(U_i(z_i) - U_i(a_\delta(\lambda)_i)\Big) \leq \delta \operatorname{range}_A G.$$

Taking the supremum over such t proves the claim. □

9.7 Approximation Algorithm

Algorithm 1: Approximate envy-free and efficient allocation

Input: Utilities $\{U\}_{i\in N}$, Lipschitz constant L, the radius bound B and precision parameter $\varepsilon > 0$.

Output: An allocation $a \in A$.

Parameter selection

(1) Compute

$$\delta = \frac{\varepsilon}{B}.$$

(2) Compute

$$\sigma = \varepsilon\frac{\delta}{L^2}.$$

Fixed-point computation

(3) Construct a simplicial subdivision of Δ^{n-1} with mesh σ.

(4) Compute a fully labeled simplex T'.

Allocation recovery

(5) Choose any $\widetilde{\lambda} \in T'$.

(6) Compute

$$a = a_\delta(\widetilde{\lambda})$$

by solving the regularized optimization problem.

(7) Return a.

Theorem 9.8. *For every $\varepsilon > 0$, Algorithm 1 returns an allocation that is ε-envy-free and ε-approximately efficient.*

Proof. Let $a = a_\delta(\widetilde{\lambda})$ be the allocation returned by Algorithm 1. The fully labeled simplex T' has diameter at most $\sigma = \varepsilon\frac{1}{K_\delta}$. Lemma 9.5 therefore gives

$$U_i\Big(a_\delta(\widetilde{\lambda})_j\Big) - U_i\Big(a_\delta(\widetilde{\lambda})_i\Big) \le K_\delta\sigma = \varepsilon,$$

for every $i, j \in N$. Hence, a is ε-envy-free.

Moreover, Algorithm 1 sets $\delta = \frac{\varepsilon}{B}$, where B is such that $\sup_{a\in A}\|a\|_F^2 \le B$. By Lemma 9.7, the uniform efficiency gap of a is at most $\delta \operatorname{range}_A G \le \delta B = \varepsilon$. Thus a is ε-approximately efficient. □

10. Computational Complexity

I now ask under what assumptions Algorithm 1 can be executed in polynomial time.

Definition 10.1. An **instance** of the allocation problem is a tuple

$$I = \Big(\{U_i\}_{i\in N}, A, P\Big),$$

where $A \subset (\mathbb{R}^d)^n$ is the feasible allocation set and P denotes any additional parameters. The **input size** of an instance denoted by s, is the number of bits to encode it.

Appendix B provides an example in which s is explicitely given for an H-representation of a polytope.

For the complexity result, I specialize preferences to linear utilities,

$$U_i(a_i) = \langle u_i, a_i \rangle, \qquad u_i \in \mathbb{Q}^d.$$

Linearity is imposed here because, together with the quadratic regularizer, it turns the regularized welfare problem into a Euclidean projection problem. Projection problems are a well-studied class of convex quadratic programs for which the exact solution can be computed. Exact solutions are necessary for Sperner's search: since envy-freeness is defined by weak inequalities, any approximation errors in allocation, translated into the errors in utility evaluations, can lead to mislabeling or empty set of envy-free agents.

Remark 10.2. Polynomial complexity of the algorithm can also be guaranteed in settings with non-linear utilities. For that, extra structure is imposed on the set of feasible allocations. The generality trade-offs between the utities and feasible set are dictated by the need to ensure that the regularized program admits an exact soluton in polynomial time.

For $\lambda \in \Delta^{n-1}$, define

$$c(\lambda) := (\lambda_1 u_1, \ldots, \lambda_n u_n)\,.$$

The regularized welfare problem is then

$$\max_{a \in A} \left\{ \langle c(\lambda), a \rangle_F - \delta \|a\|_F^2 \right\}. \tag{5}$$

Lemma 10.3. *For every $\lambda \in \Delta^{n-1}$ and $\delta > 0$,*

$$a_\delta(\lambda) = \mathrm{proj}_A \left(\frac{c(\lambda)}{2\delta} \right).$$

Proof. Completing the square gives

$$\delta \|a\|_F^2 - \langle c(\lambda), a \rangle_F = \delta \left\| a - \frac{c(\lambda)}{2\delta} \right\|_F^2 - \frac{\|c(\lambda)\|_F^2}{4\delta}.$$

The final term is independent of a, so maximizing $\langle c(\lambda), a \rangle_F - \delta \|a\|_F^2$ over A is equivalent to minimizing $\left\| a - \frac{c(\lambda)}{2\delta} \right\|_F^2$ over A. □

I impose the following computational assumptions.

Assumption 10.4 (Polynomial tractability)**.** *There exists a polynomial p such that for the instance of size s, the associated bounds B and L satisfy*

$$B(s) \le p(s), \qquad L(s) \le p(s),$$

and the strong separation oracle for A runs in time polynomial in s.

Under Assumption 10.4, the regularized program can be solved exactly in polynomial time (Grötschel et al. [1988], Kozlov et al. [1980]).

Size of the Pareto-weight search. For a requested accuracy $\varepsilon > 0$, Algorithm 1 chooses $\delta = \frac{\varepsilon}{B}$ and a mesh $\sigma = \frac{\varepsilon\delta}{L^2} = \frac{\varepsilon^2}{BL^2}$. A standard triangulation of the $(n-1)$-dimensional simplex with mesh σ has $O_n\left(\sigma^{-(n-1)}\right)$ vertices and simplices, with the constant that depends on n. Hence the number of regularized welfare problems that I need to run to label the triangulation is

$$O_n\left(\left(\frac{BL^2}{\varepsilon^2}\right)^{n-1}\right)$$

The total running time is the number of vertices in the triangulation times the cost of solving one regularized welfare problem and checking the n^2 envy inequalities at each vertex. Under the restriction to linear utilities and polynomial-time strong separation oracle, the regularized quadratic program can be solved exactly in time polynomial in s. Because the precision parameter ε enters the regularized objective linearly, the cost of solving the program "grows polynomially as ϵ goes to zero". As a way to formalize it. the cost of solving is polynomial in $\frac{1}{\epsilon}$. Finally, recovering the allocation at the selected Pareto-weight vector requires at most one further regularized welfare solve.

Combining the above, I conclude that the running time is bounded by

$$O\left(\left(\frac{B(s)L(s)^2}{\varepsilon^2}\right)^{n-1} \times \text{poly}\left(s, \frac{1}{\varepsilon}\right)\right)$$

Under Assumption 10.4, this is polynomial in s and $1/\varepsilon$. The theorem follows.

Theorem 10.5 (Polynomial running time)**.** *Fix the number of agents n. Suppose that Assumption 10.4 holds. Then, for $\varepsilon > 0$, Algorithm 1 computes an allocation that is ε-envy-free and ε-approximately weakly Pareto efficient, in the sense defined above, in time polynomial in the input size s and in $1/\varepsilon$.*

Assumption 10.4 deliberately leaves the representation of A abstract and relies on existence of a strong separation oracle. In the settings with indivisibles, a strong separation oracle that runs in polynomial time is a **strong assumption** for the set of admissible marginals which is characterized implicitely. Even though there are known examples when that assumption is satisfied (see Appendix B), it is not common.

Since convexity is necessary for the results, the algorithm may operate on a fractional relaxation. For comprehernsive (and almost exhaustive results) on exact and approximate universal implementation of bi-hierarchical constraint structures see Budish et al. [2013a] and Akbarpour and Nikzad [2020].

Part III
Applications

11. Dynamic Allocation Problems

11.1 Setup

I consider a finite set of agents $N = \{1, \dots, n\}$ and an infinite sequence of dates $t \in \mathbb{N}_0$. At each date t, an agent's consumption lies in a compact metric space Y_t. The space of consumption streams is

$$Y = \prod_{t=0}^{\infty} Y_t,$$

endowed with the product topology. Because Y is a countable product of compact metric spaces, it is itself compact and metrizable. An element $y = (y_0, y_1, \dots) \in Y$ describes an agent's consumption over time. A deterministic allocation is

$$z = (z_1, \dots, z_n) \in Y^N,$$

where $z_i = (z_{i,t})_{t \geq 0}$ denotes the consumption stream of agent i.

11.2 Dynamic feasibility

Dynamic feasibility is described by a state variable. At each date t, let S_t be a compact metric state space, and let

$$\Gamma_t \subseteq S_t \times Y_t^N \times S_{t+1}$$

be the set of feasible state-allocation-state transitions. Thus $(s_t, z_t, s_{t+1}) \in \Gamma_t$ means that allocation z_t is feasible at state s_t and can generate an admissible state s_{t+1}.

This formulation encompasses the standard controlled dynamical system

$$s_{t+1} = F_t(s_t, z_t),$$

as well as inequality constraints and set-valued state transitions. In the polyhedral applications developed below, the transition constraints take the form

$$B_t s_{t+1} + C_t s_t + D_t z_t \leq h_t.$$

I take the initial state $s_0 = \bar{s}_0$ as given.

It is useful first to define the set of feasible allocation-state trajectories:

$$\widetilde{Z} := \left\{ (z, s) \in Y^N \times \prod_{t=0}^{\infty} S_t : s_0 \in S_0^0,\ (s_t, z_t, s_{t+1}) \in \Gamma_t \text{ for every } t \right\}.$$

The deterministic feasible allocation set is its projection onto consumption streams:

$$Z := \text{proj}_{Y^N} \widetilde{Z}.$$

Assume that every Γ_t is closed. Then $\widetilde{Z}$ is a closed subset of $Y^N \times \prod_{t=0}^{\infty} S_t$. The ambient space is compact and metrizable, and therefore so is $\widetilde{Z}$. Its projection $Z \subseteq Y^N$ is consequently compact and metrizable as well. I also impose permutation-invariance of the dynamic feasible set. Thus the dynamic problem generates exactly the type of compact deterministic feasible set as considered in Part I.

11.3 Random allocations

As in Part I, a random allocation is a Borel probability measure $\mu \in \Delta(Z)$.

For each agent i, let $p_i : Y^N \to Y, p_i(z) = z_i$ denote the agent-coordinate projection onto agent i's consumption space. For every agent i, her marginal distribution over consumption streams is $\mu_i := (p_i)_{\#}\mu \in \Delta(Y)$. As in Part I, define the marginalization map

$$M : \Delta(Z) \to \Delta(Y)^N, \qquad M(\mu) = (\mu_1, \ldots, \mu_n),$$

and let $A := M(\Delta(Z))$. Thus an allocation $a = (a_1, \ldots, a_n) \in A$ records the agents' marginal lotteries and ensures implementability: for every $a \in A$, there exists μ such that $a = M(\mu)$.

11.4 Examples of preferences over dynamic stochastic allocations

Let's consider the case of expected utility first. Let $v_i : Y \to \mathbb{R}$ be a continuous von Neumann-Morgenstern index over deterministic consumption streams. Agent i's utility from a marginal lottery $a_i \in \Delta(Y)$ is

$$U_i(a_i) = \int_Y v_i(y)\, da_i(y).$$

Observe that even when restricted to linear representations, the preference assumptions of Part I permit quite general intertemporal preferences over deterministic consumption streams. A particularly familiar class is given by separably additive preferences. Let m_i be a probability measure on $\mathbb{N}_0$, and let $u_{i,t} : Y_t \to \mathbb{R}$ be continuous period utilities. Assume that

$$\sum_{t=0}^{\infty} m_i(\{t\}) \|u_{i,t}\|_\infty < \infty.$$

Define

$$v_i(y) = \int_{\mathbb{N}_0} u_{i,t}(y_t)\, dm_i(t) = \sum_{t=0}^{\infty} m_i(\{t\}) u_{i,t}(y_t).$$

For each t, let $\pi_t : Y \to Y_t, \pi_t(y) = y_t$ denote the time-coordinate projection. Also, for convenience define $p_{i,t} := \pi_t \circ p_i$. For every agent i, her date-t marginal distribution is

$$\mu_{i,t} := (p_{i,t})_{\#}\mu \in \Delta(Y_t).$$

For a random dynamic allocation, Fubini's theorem gives

$$U_i(a_i) = \sum_{t=0}^{\infty} m_i(\{t\}) \int_{Y_t} u_{i,t}(y_t)\, d\mu_{i,t}(y_t).$$

Thus, under separably additive preferences, one may interpret consumers' reasoning as follows: first, for every date t, evaluate the lottery over feasible options $\mu_{i,t}$, and then average over time using the weights m_i.

11.5 Dynamic allocations of divisible goods

For dynamic allocation problems with divisible goods, I follow the exposition of Part I and impose additional vector space assumptions. Suppose that $Y_t \subset \mathbb{R}^m$, $S_t \subset \mathbb{R}^m$, and Γ_t are convex for every t. Since the initial state is fixed, the feasible allocation-state trajectory set $\widetilde{Z}$ is convex, and therefore its projection Z is convex as well. In this case, the space of feasible allocations is defined directly as

$$A := Z,$$

rather than through lotteries over deterministic allocations.

11.6 Existence

The class of dynamic allocation problems described above therefore falls directly within the framework of Part I. In both the random-allocation and divisible-goods formulations, the preceding construction generates a compact and permutation-invariant feasible allocation set A, and the preference specifications above satisfy the continuity and concavity assumptions of the existence theorem. Let $p_i(A)$ denote the agent-coordinate projection of the set of feasible allocations, whether $p_i(A) \subset \Delta(Y)$ for indivisible goods or $p_i(A) \subset Y$ for divisible goods.

Corollary 11.1. *Suppose that A is generated by a dynamic allocation problem as above. If each $U_i : p_i(A) \to \mathbb{R}$ is continuous and concave, then there exists an allocation $a \in A$ that is envy-free and weakly Pareto efficient.*

12. Intermezzo

The existence theorem applies directly to the infinite-dimensional allocation problem. Computation does not. The main idea is to approximate the infinite-horizon problem by a sequence of truncated, finite-horizon problems. I show that, for any desired degree of precision $\epsilon > 0$, there exists a finite horizon T such that applying the algorithm of Part II to the truncated problem and concatenating its output with an admissible continuation yields an ϵ-envy-free and ϵ-approximately efficient allocation.

The approximation guarantee follows from the fact that both envy-freeness and efficiency are evaluated in utility space. Under a mild continuity-at-infinity assumption, for a sufficiently large T, what happens beyond date T has an arbitrarily small effect on agents' utilities. Therefore, agents' infinite-horizon utilities can be uniformly approximated by utility functions defined on finite prefixes of their consumption streams. Consequently, the total approximation

error is bounded by the sum of two independently controlled terms: the error from truncating the infinite horizon and the error from applying the finite-dimensional algorithm.

A separate challenge is to ensure that the finite-horizon allocation produced by the algorithm admits a feasible infinite-horizon continuation. Satisfying the feasibility constraints through date T does not, in general, guarantee that the allocation can be lifted beyond T. To address this, I construct an extended formulation of the truncated feasible set that incorporates the **viability kernel**, restricting terminal states to those from which an admissible continuation exists.

Finally, applying the finite-dimensional algorithm requires the truncated feasible sets to satisfy the computational assumptions 10.4. Extended formulations can be useful in this regard, since they may admit efficient separation oracles even when the corresponding projections on the allocation coordinates have complicated representations. Establishing general sufficient conditions for computational tractability is beyond the scope of this paper. Instead, I impose the required computational assumptions and verify them on a case-by-case basis in the stylized applications below.

13. Finite-Dimensional Approximations

13.1 Finite prefices of the set of feasible allocaions

Notation. Before proceeding, I need to introduce unified notation for the divisible goods and indivisible goods setups. Fix a finite horizon $T \in \mathbb{N}_0$ and let

$$Y_{\leq T} := \prod_{t=0}^{T} Y_t.$$

Define the prefix projection

$$\pi_{\leq T} : Y \to Y_{\leq T}, \qquad \pi_{\leq T}(y) = (y_0, \ldots, y_T).$$

and

$$\pi_{\leq T} : Z \to (Y_{\leq T})^N, \qquad \pi_{\leq T}(z) = (z_0, \ldots, z_T).$$

With a slight abuse of notation, I use the same symbol for the induced projection on probability measures,

$$\pi_{\leq T} : \Delta(Y) \to \Delta(Y_{\leq T}), \qquad \pi_{\leq T}(\nu) := (\pi_{\leq T})_{\#}\nu.$$

and

$$\pi_{\leq T} : \Delta(Z) \to \Delta((Y_{\leq T})^N), \qquad \pi_{\leq T}(\mu) := (\pi_{\leq T})_{\#}\mu.$$

Thus, in either setup, agent i's finite-prefix consumption space is

$$A_i^T := \pi_{\leq T}(A_i) = \pi_{\leq T}(p_i A).$$

Therefore,

$$A_i^T \subseteq \begin{cases} Y_{\leq T}, & \text{if } A = Z, \\ \Delta(Y_{\leq T}), & \text{if } A = M(\Delta(Z)). \end{cases}$$

13.2 Prefix feasibility and continuation

A history can satisfy every constraint through date T and nevertheless admit no feasible continuation after T. For each date t, define the continuation set

$$K_t := \{s \in S_t : \text{there exists an infinite feasible continuation starting from } s\}.$$

This is the infinite-horizon viability kernel of the dynamic feasibility system. The extended finite-prefix formulation therefore includes the terminal condition $s_{T+1} \in K_{T+1}$. Define

$$P_T := \{(z_{0:T}, s_{0:T+1}) : s_0 = \bar{s}_0,\ (s_t, z_t, s_{t+1}) \in \Gamma_t \text{ for } t = 0, \ldots, T,\ s_{T+1} \in K_{T+1}\}$$

and

$$Z_T := \text{proj}_{z_{0:T}} P_T.$$

By construction,

$$Z_T = \pi_{\leq T}(Z).$$

Indeed, every infinite feasible allocation generates an element of P_T when truncated at T, while any element of P_T can be concatenated with the feasible continuation guaranteed by $s_{T+1} \in K_{T+1}$. Thus the terminal viability kernel restriction is exactly what defines prefixes of infinite allocations among allocation sequences that are feasible over a finite horizon.

Identifying classes of problems in which viability kernels admit a simple characterization is therefore important. In particular, to be able to apply the algorithm of Part 2 to finite-horizon prefices, I have to check whether the extended finite-prefix formulation satisfies the computational complexity assumption.

In some applications the viability restriction is trivial.

Proposition 13.1 (Idle continuation)**.** *Suppose $S_t = S$ for all t, and for every t and every $s \in S$ there exists an idle allocation $z_t(s) \in Y_t^N$ such that*

$$(s, z_t(s), s) \in \Gamma_t.$$

Then $K_t = S$ for every t.

Proof. Starting from any $s \in S$, choose the idle allocation at every subsequent date. The state remains equal to s forever, so the resulting trajectory is feasible at every date. □

Below, I will provide examples of applications that have this structure and the ones that do not. Regardless of whether the viability kernel is trivial or not, working with extended formulations is computationally useful. Even as convex compact subsets of finite-dimensional vector spaces, projections of polytopes need not have a convenient direct inequality description. Thus, the complexity of a projection Z_T can be much larger than the complexity of an extended representation P_T. Further discussion and an example are in Appendix C.

Indivisibles When Y_t is finite, I work with lotteries as in Part I. Let $\Delta(Z_T)$ denote the set of joint lotteries over feasible allocation prefixes.

For each agent i, let $p_i^T : Y_{\leq T}^N \to Y_{\leq T}$ denote the projection onto agent i's prefix, and define the marginalization map

$$M_T : \Delta(Z_T) \to \Delta(Y_{\leq T})^N, \qquad M_T(\mu^T) = \Big((p_1^T)_\# \mu^T, \ldots, (p_n^T)_\# \mu^T\Big).$$

The feasible marginal prefix allocations are therefore $A_{\leq T} := M_T(\Delta(Z_T))$. The extended-prefix formulation gives an equivalent description of the joint prefix lotteries. It formalizes treating $A_{\leq T}$ as the set of marginals on the allocation coordinates of probability measures on extended prefices P_T.

Lemma 13.2. *Let*

$$q_T : P_T \to Z_T, \qquad q_T(z_{0:T}, s_{0:T+1}) = z_{0:T}.$$

Then

$$(q_T)_\# \Delta(P_T) = \Delta(Z_T).$$

Consequently,

$$A_{\leq T} = M_T\Big((q_T)_\# \Delta(P_T)\Big).$$

The proof is in the Appendix C.

Divisibles When Z is convex and $A := Z$, the finite-prefix feasible allocation set is

$$A_{\leq T} := \pi_{\leq T}(A) = \pi_{\leq T}(Z) = Z_T.$$

Hence $A_{\leq T} = \mathrm{proj}_{z_{0:T}} P_T$, so P_T serves as an extended formulation of the finite-prefix allocation set $A_{\leq T}$.

Computational complexity Recall that the algorithm in Part II takes as part of an input the set of feasible allocations A that is a subset of $\mathbb{R}^{dn}$ for some integer d and satisfies the following:

- nonempty, compact, and convex;

- permutation invariant;
- explicitely provides $B < \infty$ such that

$$\|a\|_F^2 \leq B \qquad \text{for every } a \in A.$$

  and $B(s)$ is polynomial in the size of the instance
- supplies a strong separation oracle that runs in polynomial time in the size of the instance

For any T, compactness, convexity and permutation invariance of $A_{\leq T}$ follows from compactness, convexity and permutation invariance of A in both cases of divisible and indivisible goods. General set of sufficient conditions on a dynamic allocation problem for $A_{\leq T}$ to satisfy the rest of the assumptions remains an open problem. In the stylized applications section below they are verified case-by-case.

13.3 Continuity at infinity

Recall that every agent's preferences are represented by $U_i : A_i \to \mathbb{R}$ which is continuous and concave. On top of that, I assume that

Assumption 13.3 (Finite-prefix utilities)**.** *For each $i \in N$ and $T \in \mathbb{N}_0$, there exists a concave and Lipschitz continuous function $U_i^T : A_i^T \to \mathbb{R}$ such that*

$$\sup_{a_i \in A_i} \left| U_i(a_i) - U_i^T\Big(\pi_{\leq T}(a_i)\Big) \right| \leq \eta(T),$$

where $\eta(T) \to 0$. Observe that the requirement above is a Lipschitzian strengthening of a standard assumption of the continuity at infinity (see Appendix C).

Examples of preferences that satisfy uniform approximation by finite prefices Recall that I denote the agent-coordinate and time-coordinate projection by $p_{i,t} := \pi_t \circ p_i$. With a slight abuse of notation, I use the same symbol for the induced pushforward on probability measures. If $A = Z$, I define $a_{i,t} = p_{i,t}(a)$, while if $A = M(\Delta(Z))$, $a_{i,t} = (p_{i,t})_{\#}\mu$. Thus,

$$a_{i,t} \in \begin{cases} Y_t, & \text{if } A = Z, \\ \Delta(Y_t), & \text{if } A = M(\Delta(Z)). \end{cases}$$

Let $u_{i,t} : A_{i,t} \to \mathbb{R}$ be L-Lipschitz continuous, concave, and unformly bounded, $\sup_y |u_{i,t}(y)| \leq M_i$. Define

$$U_i(a_i) = (1 - \beta_i) \sum_{t=0}^{\infty} \beta_i^t u_{i,t}(a_{i,t}), \qquad 0 < \beta < 1$$

Define

$$U_i^T(a_{i,\leq T}) = \sum_{t=0}^{T} \beta_i^t u_{i,t}(a_{i,t}).$$

Then U_i^T is concave and, for the product norm on the finite-dimensional prefix space, L-Lipschitz. Moreover, the truncation error satisfies

$$\left|U_i(a_i) - U_i^T(\pi_{\leq T}(a_i))\right| \leq \sum_{t=T+1}^{\infty} \beta_i^t M_i = \frac{M_i \beta_i^{T+1}}{1-\beta_i}.$$

hence $\eta_i(T) \to 0$.

13.4 Approximation guarantees

The final step is to transfer the finite-prefix guarantees back to the infinite-horizon environment. Suppose $a^T \in A_{\leq T}$ is a finite-prefix allocation and that it admits a feasible lift $a \in A$. Let a_i^T denote the prefix corresponding to agent i. Assume that a^T is ε-envy-free for the truncated utilities:

$$U_i^T(a_j^T) - U_i^T(a_i^T) \leq \varepsilon \qquad \text{for every } i, j \in N$$

Assume that it is also ε-approximately efficient:

$$\sup_{b^T \in A^T} \min_{i \in N} \left\{U_i^T(b_i^T) - U_i^T(a_i^T)\right\} \leq \varepsilon$$

Lemma 13.4 (Truncation transfer)**.** *If $a^T \in A_{\leq T}$ is as above, then its feasible lift $a \in A$ is $\varepsilon + 2\eta(T)$-envy-free and $\varepsilon + 2\eta(T)$-approximately efficient. That is*

$$U_i(a_j) - U_i(a_i) \leq \varepsilon + 2\eta(T) \qquad \textit{for every } i, j \in N$$

and

$$\sup_{b \in A} \min_{i \in N} \{U_i(b_i) - U_i(a_i)\} \leq \varepsilon + 2\eta(T).$$

Proof. The bound follows easily from Assumption 13.3 For approximate envy-freeness, observe that

$$U_i(a_j) - U_i(a_i) \leq U_i^T(a_j^T) + \eta(T) - U_i^T(a_i^T) + \eta(T) \leq \varepsilon + 2\eta(T).$$

For efficiency, take any $b \in A$ and let b^T be its finite prefix. For every i

$$U_i(b_i) - U_i(a_i) \leq U_i^T(b_i^T) - U_i^T(a_i^T) + 2\eta(T).$$

Hence

$$\min_i \{U_i(b_i) - U_i(a_i)\} \leq \min_i \{U_i^T(b_i^T) - U_i^T(a_i^T)\} + 2\eta(T) \leq \varepsilon + 2\eta(T).$$

Taking the supremum over $b \in A$ proves the claim. □

The approximation result now follows directly. The only delicate part concerns the case of indivisibles. I make the following strong computational assumption: for every allocation μ in $\Delta(P_T)$, computing $M_T((q_T)_{\#}\mu)$ is executed in polynomial time. For divisibles, the result goes

through without extra assumptions since projecting onto a strict subset of coordinates for the representation of P_T above is computationally easy.

Theorem 13.5 (Approximate envy and efficiency in dynamic environments)**.** *Suppose the dynamic allocation problem is defined as above and:*

1. *utilities admit concave and Lipschitz-continuous finite-prefix approximations with*

$$\sup_{a_i \in A_i} |U_i(a_i) - U_i^T(\pi_{\leq T} a_i)| \leq \eta(T),$$

2. *feasible finite prefixes are represented by*

   - *for the settings with divisibles, by $A_{\leq T} = Z_T = \mathrm{proj}_{z_{0:T}} P_T$, where P_T satisfies the algorithmic assumptions of Part II*
   - *for the setting with indivisibles, by $A_{\leq T} = M_T\Big((q_T)_{\#}\Delta(P_T)\Big)$, where $\Delta(P_T)$ satisfies computational assumptions 10.4 of Part II and $L_T(\mu) = M_T((q_T)_{\#}\mu)$ operator can be evaluated in polynomial time*

Then, for every $\epsilon > 0$, for fixed n, there exists an ϵ-envy-free and ϵ-weakly Pareto efficient dynamic allocation with a polynomially computable finite prefix.

Proof. Set $\varepsilon_{\mathrm{alg}} := \frac{\epsilon}{2}$ and choose T such that $\eta(T) \leq \epsilon/4$. Apply 10.5 from Part II to A^T with accuracy $\varepsilon_{\mathrm{alg}}$. It returns a finite-prefix allocation a^T satisfying

$$U_i^T(a_j^T) - U_i^T(a_i^T) \leq \frac{\epsilon}{2} \qquad \text{for every } i, j \in N$$

and

$$\sup_{b^T \in A^T} \min_i \{U_i^T(b_i^T) - U_i^T(a_i^T)\} \leq \frac{\epsilon}{2}.$$

By assumption a^T admits a feasible lift $a \in A$. The truncation-transfer lemma gives $U_i(a_j) - U_i(a_i) \leq \frac{\epsilon}{2} + 2\frac{\epsilon}{4} = \epsilon$ and $\sup_{b \in A} \min_i \{U_i(b_i) - U_i(a_i)\} \leq \frac{\epsilon}{2} + 2\frac{\epsilon}{4} = \epsilon$. □

The theorem separates two conceptually different approximations. Part II introduces approximation because the finite-dimensional fairness-and-efficiency problem is solved only approximately. The second approximation comes from truncating the infinite allocation stream to a finite-horizon prefix. The two errors are controlled independently and then added by the transfer lemma.

14. Stylized examples

This section illustrates the computational framework in three dynamic allocation environments: airports, railroads, fisheries.

Although economically distinct, these environments share several features that make them appealing candidates for dynamic fair division:

1. Agents care about long-run patterns of consumption
2. The set of feasible allocations is intertemporarily constrained, and the nature of the constraints is complex
3. Deterministic feasibility generates non-convex sets
4. Even though the allocation system evolves over time, fairness is evaluated for the same group of agents

The examples are organized according to the structure of continuation. In the airline and railway applications, every feasible finite history admits a trivial continuation obtained by leaving the physical state unchanged. Consequently, the finite-dimensional projection of the infinite-horizon feasible set is described entirely by the constraints imposed up to the truncation date. The fisheries application is different: zero activity need not generate a feasible continuation, and finite-prefix feasibility therefore requires a terminal viability condition.

The examples are deliberately stylized. Their purpose is not to model every institutional feature of the corresponding industries, but to exhibit economically natural environments in which the finite-dimensional projections required by the algorithm admit tractable extended formulations.

14.1 Applications with trivial continuation

Airline capacity. Let L be a finite set of airports and $N = \{1, \dots, n\}$ a fixed set of airlines. At date t, airline i receives arrival and departure capacity

$$z_{i,t} = (a_{i,t,l}, d_{i,t,l})_{l \in L} \in Y_t \subseteq \mathbb{N}_0^{2|L|},$$

where $a_{i,t,l}$ and $d_{i,t,l}$ denote the quantities of arrival and departure capacity allocated to airline i at airport l.

The state records the location of each airline's aircraft. Let $S_{i,t,l}$ be the number of aircraft of airline i available at airport l at the beginning of date t, and write $s_t = (S_{i,t,l})_{i \in N, l \in L}$.

Airlines have a common fleet size F, so

$$S_{i,t,l} \geq 0, \qquad \sum_{l \in L} S_{i,t,l} = F.$$

The stock-flow equations are

$$S_{i,t+1,l} = S_{i,t,l} + a_{i,t,l} - d_{i,t,l},$$

and total activity is bounded by

$$\sum_{l \in L} a_{i,t,l} = \sum_{l \in L} d_{i,t,l} \leq Q_t.$$

Airport capacity is shared across airlines:

$$\sum_{i \in N} a_{i,t,l} \leq C^{\text{arr}}_{t,l}, \qquad \sum_{i \in N} d_{i,t,l} \leq C^{\text{dep}}_{t,l}.$$

Fix an initial fleet distribution s_0, and let $\Gamma_t \subseteq S_t \times Y_t^N \times S_{t+1}$ be the set of triples satisfying the constraints above. The infinite-horizon deterministic feasible set is therefore

$$Z = \left\{ z \in Y^N : \exists (s_t)_{t \geq 0} \text{ such that } s_0 \text{ is fixed and } (s_t, z_t, s_{t+1}) \in \Gamma_t \text{ for every } t \geq 0 \right\}.$$

Because all individual technological restrictions are common across airlines and airport capacities depend only on aggregate use, Z is invariant under permutations of agents.

The important feature of this formulation is that there is no positive lower bound on future activity. Given any feasible state s_{T+1}, setting

$$a_{i,t,l} = d_{i,t,l} = 0, \qquad S_{i,t,l} = S_{i,T+1,l}, \qquad t > T,$$

for every i and l generates a feasible continuation. Thus every finite history satisfying the local constraints through date T is the prefix of an infinite feasible allocation.

Fix a horizon T. Define the finite-horizon extended set

$$P_T = \left\{ (a, d, S) : \begin{array}{ll} S_{i,t+1,l} = S_{i,t,l} + a_{i,t,l} - d_{i,t,l}, & t = 1, \ldots, T, \\ S_{i,t,l} \geq 0, \qquad \sum_l S_{i,t,l} = F, & \\ a_{i,t,l}, d_{i,t,l} \geq 0, & \\ \sum_i a_{i,t,l} \leq C^{\text{arr}}_{t,l}, & \\ \sum_i d_{i,t,l} \leq C^{\text{dep}}_{t,l}, & \\ \sum_l a_{i,t,l} = \sum_l d_{i,t,l} \leq Q_t, & \\ S_{T+1} \in K_{T+1} & \end{array} \right\}$$

Because of the availability of the idle action, $K_{T+1} = \{S : S \geq 0, \sum_l S_{i,l} = F \text{ for all } i \in N\}$.

Airlines preferences Consider two airlines with the same fleet capacity F, so that they face the same individual feasibility constraints but differ in their preferences over feasible schedules. Let W denote a low-cost carrier of the Wizz Air type and E a premium long-haul carrier of the Emirates type.

A low-cost carrier places a high value on aircraft utilization and traffic volume and may be relatively flexible about the precise temporal organization of its schedule. A simple specification is

$$U_W(x_W) = \sum_{t=0}^{\infty} \beta_W^t \sum_{l \in L} w_l^W \, u_W(a_{W,t,l} + d_{W,t,l}),$$

where $u_W : \mathbb{R}_+ \to \mathbb{R}$ is increasing and concave and $w_l^W \geq 0$ captures the profitability of operating at airport l. Concavity captures diminishing returns to additional activity at a given airport and date, while the weights can reflect the relative attractiveness of different airports;

for example, a low-cost carrier may place relatively high weight on profitable secondary airports. The absence of an adjustment-cost term reflects greater flexibility in reorganizing capacity across planning periods.

A premium long-haul carrier may instead value particular hub airports, plan over a longer horizon, and place greater value on schedule stability. A corresponding reduced-form specification is

$$U_E(x_E) = \sum_{t=0}^{\infty} \beta_E^t \left[\sum_{l \in L} w_l^E \, u_E(a_{E,t,l} + d_{E,t,l}) + \delta \sum_{l \in H} g^E(\min\{a_{E,t,l}, \kappa_l \sum_{s=t}^{t+K} d_{E,s,l}\}) - \rho \left\| x_{E,t+1} - x_{E,t} \right\|_1 \right],$$

$$\rho \geq 0,$$

where

$$x_{E,t} = \Big((a_{E,t,l}, d_{E,t,l}) \Big)_{l \in L}.$$

The weights w_l^E can be larger at the carrier's hub airports, while a discount factor β_E close to one represents a comparatively long planning horizon.

Define connection value as $\sum_{l \in H} g^E(\min\{a_{E,t,l}, \kappa_l \sum_{s=t}^{t+K} d_{E,s,l}\})$. It rewards an arrival at a hub when it is followed by many departure opportunities within the next K periods. $\kappa_l > 0$ converts departure opportunities into units comparable to arrivals. g^E is taken to be increasing and concave, to smooth the action of min.

The term

$$-\rho \left\| x_{E,t+1} - x_{E,t} \right\|_1$$

penalizes changes in the allocation of arrival and departure capacity across consecutive planning periods and therefore captures a preference for stable schedules.

Railroad application is in the Appendix C. It is different from the airline application as it explicitely models a network of routes of limited capacity that have to be allocated between the railway operators, not just stations-wide arrival and departure slots.

14.2 Applications with nontrivial continuation

The previous examples have the property that every state reachable at the end of a feasible prefix can be continued forever. This need not hold in dynamic environments in which current actions affect a common physical stock. Consider a divisible example of stylized fishery with common biological stock $s_t \in S = [\underline{s}, \overline{s}]$ and harvest $z_{i,t} = x_{i,t} \in [0, \overline{x}]$ by agent i. To preserve permutation invariance, all harvesters have the same technological effect on the resource. The stock evolves according to

$$s_{t+1} = a s_t + b - q \sum_{i \in N} x_{i,t}, \qquad a > 1, \quad b \geq 0, \quad q > 0,$$

and define

$$\Gamma = \left\{ (s, x, s') \in S \times [0, \overline{x}]^N \times S : s' = as + b - q \sum_{i \in N} x_i \right\}.$$

The parameter restriction $a > 1$ is a convenient specification of positive biological growth. Unlike in the transportation examples, zero activity does not leave the state unchanged:

$$x_{i,t} = 0 \ \forall i, t \qquad \Longrightarrow \qquad s_{t+1} = a^{t+1} s_0 + b\frac{a^{t+1} - 1}{a - 1}.$$

Thus a stock that is feasible today may fail to admit any infinite feasible continuation. Let

$$K = \{s \in S : \exists (x_t, s_{t+1})_{t \geq 0} \text{ with } s_0 = s \text{ and } (s_t, x_t, s_{t+1}) \in \Gamma \text{ for every } t \geq 0\}$$

denote the viability kernel. Since aggregate harvest $H_t = \sum_i x_{i,t}$ can take any value in $[0, n\overline{x}]$, the largest stock that can be held constant satisfies

$$s = as + b - qn\overline{x},$$

or

$$s^\star = \frac{qn\overline{x} - b}{a - 1}.$$

For any $s > s^\star$, even maximal harvesting cannot prevent the stock from increasing. Assume that $\underline{s} \leq s^\star < \overline{s}$.

Proposition 14.1. *The viability kernel of the fishery is $K = [\underline{s}, s^\star]$.*

Proof. If $s \in [\underline{s}, s^\star]$, choose aggregate harvest $H = ((a-1)s + b)/q$. Since $b \geq 0$, $H \geq 0$, and since $s \leq s^\star$, $H \leq n\overline{x}$; this aggregate harvest can be divided among the n agents with each $x_i \in [0, \overline{x}]$. The resulting next state is $s' = s$, so repeating the same action gives an infinite feasible continuation.

If $s > s^\star$, $s' - s \geq (a-1)s + b - qn\overline{x} = (a-1)(s - s^\star) > 0$. Using the fact that $s' - s + s - s^* = s' - s^* = a(s - s^*)$, observe that every feasible continuation from $s > s^\star$ satisfies $s_t - s^\star \geq a^t(s - s^\star)$, and therefore exceeds $\overline{s}$ after finitely many periods. Such a state is not viable. □

A finite sequence satisfying the stock equation and the state constraints through date T need not belong to $\Pi_{\leq T}(Z)$, because its terminal state may lie outside K. The finite-horizon extended set is therefore

$$P_T^{\text{fish}} = \{(z_0, \ldots, z_T, s_1, \ldots, s_{T+1}) : (s_t, z_t, s_{t+1}) \in \Gamma, \ t = 0, \ldots, T, \quad s_{T+1} \in K = [\underline{s}, s^*]\}.$$

Proposition 14.2. *For every $T \geq 0$,*

$$\Pi_{\leq T}(Z) = \text{proj}_{z_0, \ldots, z_T} P_T^{\text{fish}}.$$

Moreover, P_T^{fish} is a polyhedron that satisfies the assumptions 10.4.

Proof. If $z \in Z$, then its associated state sequence is feasible through date T, and $s_{T+1} \in K$ because the tail of the same allocation provides an infinite continuation; hence $\Pi_{\leq T}(z)$ belongs

to the projection of P_T^{fish}. Conversely, if a finite prefix belongs to P_T^{fish}, its terminal stock lies in K, so by definition there exists an infinite feasible continuation from s_{T+1}; concatenating this continuation with the prefix produces an element of Z. Finally, $K = [\underline{s}, s^\star]$ is described by two linear inequalities, and the remaining constraints are the $T+1$ linear stock-flow equations, individual harvest bounds, and state bounds. The dimension abd number of constraints therefore grow polynomially with T. □

The fisheries example isolates the role of viability in the finite-dimensional approximation. In the airline and railway models, trivial continuation makes every terminal state generated by a feasible prefix viable, so no terminal constraint is needed. In the fishery, the finite-prefix projection is obtained only after excluding terminal states from which future feasibility is impossible. More generally, if $K_{T+1} \subseteq S_{T+1}$ denotes the set of states admitting an infinite continuation, the natural finite-horizon extended formulation is

$$P_T = \{(z_0, \dots, z_T, s_1, \dots, s_{T+1}) : (s_t, z_t, s_{t+1}) \in \Gamma_t,\ t = 0, \dots, T, \quad s_{T+1} \in K_{T+1}\}\,.$$

The computational question is therefore not whether dynamic constraints can be truncated, but whether the terminal viability set itself admits a representation of polynomial complexity. The airline and railway models answer this question trivially because $K_t = S_t$; the fishery gives a nontrivial example in which $K_t \subsetneq S_t$ but remains explicitly and efficiently representable.

15. Conclusion

Read the abstract and/or introduction.

References

Atila Abdulkadiroğlu and Tayfun Sönmez. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3):689–701, 1998. doi: 10.2307/2998580.

Mohammad Akbarpour and Afshin Nikzad. Approximate random allocation mechanisms. *The Review of Economic Studies*, 87(6):2473–2510, 2020. doi: 10.1093/restud/rdz066.

Ethan Akin. Vilfredo pareto cuts the cake. *Journal of Mathematical Economics*, 24:23–44, 1995.

Erik M. Alfsen. *Compact Convex Sets and Boundary Integrals*. Springer, 1971. Springer Link.

Ahmet Alkan, Gabrielle Demange, and David Gale. Fair allocation of indivisible goods and criteria of justice. *Econometrica*, 59(4):1023–1039, 1991. doi: 10.2307/2938172.

Georgios Amanatidis, Haris Aziz, Georgios Birmpas, Aris Filos-Ratsikas, Bo Li, Hervé Moulin, Alexandros A. Voudouris, and Xiaowei Wu. Fair division of indivisible goods: Recent progress and open questions. *Artificial Intelligence*, 322:103965, 2023. doi: 10.1016/j.artint.2023.103965.

Eduardo M. Azevedo and Eric Budish. Strategy-proofness in the large. *The Review of Economic Studies*, 86(1):81–116, 2019. doi: 10.1093/restud/rdy042.

Haris Aziz. Achieving envy-freeness and equitability with monetary transfers. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 5102–5109, 2021. doi: 10.1609/aaai.v35i6.16645.

Haris Aziz, Zixu He, Xinhang Lu, and Kaiyang Zhou. Fair allocation of divisible goods under non-linear valuations. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 170–178. International Foundation for Autonomous Agents and Multiagent Systems, 2025. URL `https://arxiv.org/abs/2607.15613`.

Yushi Bai and Paul Gölz. Envy-free and pareto-optimal allocations for agents with asymmetric random valuations. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence (IJCAI)*, pages 53–59, 2022. doi: 10.24963/ijcai.2022/8.

Christian Basteck and Lars Ehlers. Strategy-proof and envy-free random assignment. *Journal of Economic Theory*, 209:105618, 2023. doi: 10.1016/j.jet.2023.105618.

Marcus Berliant, William Thomson, and Karl Dunz. On the fair division of a heterogeneous commodity. *Journal of Mathematical Economics*, 21(3):201–216, 1992. doi: 10.1016/0304-4068(92)90001-N.

Anna Bogomolnaia and Hervé Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295–328, 2001. doi: 10.1006/jeth.2000.2710.

Anna Bogomolnaia, Hervé Moulin, Fedor Sandomirskiy, and Elena Yanovskaya. Competitive division of a mixed manna. *Econometrica*, 85(6):1847–1871, 2017.

Steven J. Brams and Peter C. Fishburn. Fair division of indivisible items between two people with identical preferences: Envy-freeness, pareto-optimality, and equity. *Social Choice and Welfare*, 17(2):247–267, 2000. doi: 10.1007/s003550050019.

Xiaolin Bu, Jiaxin Song, and Biaoshuai Tao. On existence of truthful fair cake cutting mechanisms. *Artificial Intelligence*, 319:103904, 2023. doi: 10.1016/j.artint.2023.103904.

Eric Budish and Estelle Cantillon. The multi-unit assignment problem: Theory and evidence from course allocation at harvard. *American Economic Review*, 102(5):2237–2271, 2012. doi: 10.1257/aer.102.5.2237.

Eric Budish, Yeon-Koo Che, Fuhito Kojima, and Paul Milgrom. Designing random allocation mechanisms: Theory and applications. *American Economic Review*, 103(2):585–623, 2013a. doi: 10.1257/aer.103.2.585.

Eric Budish, Yeon-Koo Che, Fuhito Kojima, and Paul Milgrom. Designing random allocation mechanisms: Theory and applications, 2013b. Working paper.

Ioannis Caragiannis and Shivika Narang. Repeatedly matching items to agents fairly and efficiently. *Theoretical Computer Science*, 981:114246, 2024. doi: 10.1016/j.tcs.2023.114246.

Ioannis Caragiannis, Kristoffer Arnsfelt Hansen, and Nidhi Rathi. On the complexity of Pareto-optimal and envy-free lotteries. In Natasha Alechina, Virginia Dignum, Mehdi Dastani, and Jaime Simão Sichman, editors, *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '24, pages 244–252. International Foundation for Autonomous Agents and Multiagent Systems, 2024. URL `https://www.ifaamas.org/Proceedings/aamas2024/pdfs/p244.pdf`.

Simone Cerreia-Vioglio, David Dillenberger, and Pietro Ortoleva. Cautious expected utility and the certainty effect. *Econometrica*, 83(2):693–728, 2015. doi: 10.3982/ECTA11733.

Richard Cole and Xin Tao. Mechanism design for fair division: Allocating divisible items without payments. pages 385–386, 2019.

Richard Cole and Yixin Tao. On the existence of Pareto efficient and envy-free allocations. *Journal of Economic Theory*, 193:105207, 2021. doi: 10.1016/j.jet.2021.105207.

Benjamin Cookson, Soroush Ebadian, and Nisarg Shah. Temporal fair division. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 13727–13734, 2025. doi: 10.1609/aaai.v39i13.33500.

Ram Sewak Dubey and Francesco Ruscitti. Fair allocations in an overlapping generations economy. *Studies in Microeconomics*, 12(2):172–199, aug 2024. doi: 10.1177/23210222221119362.

Lester E. Dubins and Edwin H. Spanier. How to cut a cake fairly. *The American Mathematical Monthly*, 68(1):1–17, 1961. doi: 10.1080/00029890.1961.11989615.

Federico Echenique, Antonio Miralles, and Jun Zhang. Fairness and efficiency for allocations with participation constraints. *Journal of Economic Theory*, 195:105274, 2021a. doi: 10.1016/j.jet.2021.105274.

Federico Echenique, Antonio Miralles, and Jun Zhang. Constrained pseudo-market equilibrium. *American Economic Review*, 111(11):3699–3732, 2021b. doi: 10.1257/aer.20201769.

Paul H. Edelman and Peter C. Fishburn. Fair division of indivisible items among people with similar preferences. *Mathematical Social Sciences*, 41(3):327–347, 2001. doi: 10.1016/S0165-4896(00)00065-2.

Edmund Eisenberg. Aggregation of utility functions. *Management Science*, 7(4):337–350, 1961. doi: 10.1287/mnsc.7.4.337.

Edmund Eisenberg and David Gale. Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics*, 30(1):165–168, 1959. doi: 10.1214/aoms/1177706369.

Edith Elkind, Alexander Lam, Mohamad Latifian, Tzeh Yuan Neoh, and Nicholas Teh. Temporal fair division of indivisible items, 2024.

Duncan K. Foley. Resource allocation and the public sector. *Yale Economic Essays*, 7:45–90, 1967.

Jugal Garg, Yixin Tao, and László A. Végh. Tight efficiency bounds for the probabilistic serial and related mechanisms, 2025. URL `https://arxiv.org/abs/2507.03359`.

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. doi: 10.1016/j.geb.2023.12.006.

Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric Algorithms and Combinatorial Optimization*, volume 2 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 1988. doi: 10.1007/978-3-642-97881-4.

Peter J. Hammond. Straightforward individual incentive compatibility in large economies. *The Review of Economic Studies*, 46(2):263–282, 1979. doi: 10.2307/2297050.

Farhad Hüsseinov and Nobusumi Sagara. Existence of efficient envy-free allocations of a heterogeneous divisible commodity with nonadditive utilities. *Social Choice and Welfare*, 41 (4):923–940, 2013. doi: 10.1007/s00355-012-0714-y.

Aanund Hylland and Richard Zeckhauser. The efficient allocation of individuals to positions. *Journal of Political Economy*, 87(2):293–314, 1979. doi: 10.1086/260757.

Ayumi Igarashi, Martin Lackner, Oliviero Nardi, and Arianna Novaro. Repeated fair allocation of indivisible items. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 9792–9799, 2024. doi: 10.1609/aaai.v38i9.28837.

Tanguy Isaac and Paolo G. Piacquadio. Equity and efficiency in an overlapping generation model. *Social Choice and Welfare*, 44(3):549–565, mar 2015.

Akshay-Kumar Katta and Jay Sethuraman. A solution to the random assignment problem on the full preference domain. *Journal of Economic Theory*, 131(1):231–250, 2006. doi: 10.1016/j.jet.2005.05.001.

Yasushi Kawase, Kazuhisa Makino, Hanna Sumita, Akihisa Tamura, and Makoto Yokoo. Towards optimal subsidy bounds for envy-freeable allocations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 9824–9831, 2024. doi: 10.1609/aaai.v38i9.28842.

M. K. Kozlov, S. P. Tarasov, and L. G. Khachiyan. The polynomial solvability of convex quadratic programming. *USSR Computational Mathematics and Mathematical Physics*, 20 (5):223–228, 1980. doi: 10.1016/0041-5553(80)90098-1.

Fabio Maccheroni. Maxmin under risk. *Economic Theory*, 19(4):823–831, 2002.

Fabio Maccheroni and Massimo Marinacci. How to cut a pizza fairly: Fair division with decreasing marginal evaluations. *Social Choice and Welfare*, 20(3):457–465, 2003. doi: 10.1007/s003550200192.

Andreu Mas-Colell. A model of equilibrium with differentiated commodities. *Journal of Mathematical Economics*, 2(2):263–295, 1975. doi: 10.1016/0304-4068(75)90028-2. Also available via RePEc / ScienceDirect.

Marc Meertens, Jos Potters, and Hans Reijnierse. Envy-free and Pareto efficient allocations in economies with indivisible goods and money. *Mathematical Social Sciences*, 44(3):223–233, 2002. doi: 10.1016/S0165-4896(02)00064-1.

Antonio Miralles and Marek Pycia. Foundations of pseudomarkets: Walrasian equilibria for discrete resources. *Journal of Economic Theory*, 196:105303, 2021. doi: 10.1016/j.jet.2021.105303.

Alexander S. Nesterov. Fairness and efficiency in strategy-proof object allocation mechanisms. *Journal of Economic Theory*, 170:145–168, 2017. doi: 10.1016/j.jet.2017.05.004.

Manfred W. Padberg and M. R. Rao. Odd minimum cut-sets and b-matchings. *Mathematics of Operations Research*, 7(1):67–80, 1982. doi: 10.1287/moor.7.1.67.

Elisha A. Pazner and David Schmeidler. A difficulty in the concept of fairness. *The Review of Economic Studies*, 41(3):441–443, 1974. doi: 10.2307/2296762. URL `https://hdl.handle.net/10.2307/2296762`.

Shigehiro Serizawa and John A. Weymark. Efficient strategy-proof exchange and minimum consumption guarantees. *Journal of Economic Theory*, 109(2):246–263, 2003. doi: 10.1016/S0022-0531(03)00019-X.

Priyanka Shende and Manish Purohit. Strategy-proof and envy-free mechanisms for house allocation. *Journal of Economic Theory*, 213:105712, 2023. doi: 10.1016/j.jet.2023.105712.

Hugo Steinhaus. The problem of fair division. *Econometrica*, 16:101–104, 1948.

Neil Thakral. Matching with stochastic arrival. *AEA Papers and Proceedings*, 109:209–212, 2019. doi: 10.1257/pandp.20191032.

Thorben Tröbst and Vijay V. Vazirani. Cardinal-utility matching markets: The quest for envy-freeness, Pareto-optimality, and efficient computability. *Mathematics of Operations Research*, 2025. doi: 10.1287/moor.2024.0770. URL `https://doi.org/10.1287/moor.2024.0770`. Published online September 26, 2025.

Hal R. Varian. Equity, envy, and efficiency. *Journal of Economic Theory*, 9:63–91, 1974.

Dietrich Weller. Fair division of a measurable space. *Journal of Mathematical Economics*, 14 (1):5–17, 1985. doi: 10.1016/0304-4068(85)90023-0.

Lin Zhou. On a conjecture by gale about one-sided matching problems. *Journal of Economic Theory*, 52(1):123–135, 1990. doi: 10.1016/0022-0531(90)90070-Z.

A. Appendix A

A.1 Properties of the set of admissible marginals

Let $\Delta(Z)$ denote the space of Borel probability measures on Z. By the Riesz representation theorem, I identify $\Delta(Z)$ with a subset of $C(Z)^*$ and endow it with the weak* topology. A random allocation $\pi \in \Delta(Z)$ induces for every agent i a marginal probability measure $p^i \in \Delta(Y)$ over her own consumption bundles.

Define

$$A = \left\{(p^1, \ldots, p^n) \in \Delta(Y)^n : \exists\, \pi \in \Delta(Z) \text{ whose marginals are } (p^1, \ldots, p^n)\right\}.$$

For each $i \in N$, let $\mathrm{pr}_i : Z \to Y$ denote the coordinate projection, and let

$$M_i : \Delta(Z) \to \Delta(Y)$$

be the induced marginal map,

$$M_i(\pi) = \pi \circ \mathrm{pr}_i^{-1}.$$

Define

$$M : \Delta(Z) \to \Delta(Y)^n, \qquad M(\pi) = \Big(M_1(\pi), \ldots, M_n(\pi)\Big).$$

By definition,

$$A = M(\Delta(Z)).$$

I first record the continuity property used below.

Lemma A.1. *Let X and W be metric spaces, and let $T : X \to W$ be continuous. The pushforward map*

$$T_\# : \Delta(X) \to \Delta(W), \qquad T_\#\mu = \mu \circ T^{-1},$$

is continuous in the weak topology.

Proof. Let $\mu_k \to \mu$ weakly in $\Delta(X)$. For every bounded continuous function $f : W \to \mathbb{R}$, the composition $f \circ T$ is bounded and continuous on X. Hence

$$\int_W f\, d(T_\#\mu_k) = \int_X f \circ T\, d\mu_k \longrightarrow \int_X f \circ T\, d\mu = \int_W f\, d(T_\#\mu).$$

Therefore $T_\#\mu_k \to T_\#\mu$ weakly. □

Proposition A.2. *Suppose that $Z \subseteq Y^n$ is compact and permutation invariant. Then the induced feasible random allocation space*

$$A = \Big\{(p^1, \dots, p^n) \in \Delta(Y)^n : \exists\, \pi \in \Delta(Z) \text{ whose marginals are } (p^1, \dots, p^n)\Big\}$$

is compact and permutation invariant.

Proof. Since Z is a compact metric space, $\Delta(Z)$ is compact in the weak topology. Each coordinate projection

$$\mathrm{pr}_i : Z \to Y$$

is continuous, so Lemma A.1 implies that each marginal map M_i is continuous. Therefore

$$M : \Delta(Z) \to \Delta(Y)^n, \qquad M(\pi) = \Big(M_1(\pi), \dots, M_n(\pi)\Big),$$

is continuous. Since

$$A = M(\Delta(Z)),$$

the set A is the continuous image of a compact set and is therefore compact.

It remains to prove permutation invariance. Fix $\sigma \in S_N$. Define the relabeling map

$$T_\sigma : Z \to Z$$

by

$$\Big(T_\sigma z\Big)_{\sigma(i)} = z_i$$

for every $i \in N$. Since Z is permutation invariant, $T_\sigma(Z) = Z$. Moreover, T_σ is a homeomorphism.

Take $a = (a^1, \dots, a^n) \in A$. Then there exists $\pi \in \Delta(Z)$ whose marginals are $a^1, \dots, a^n$. Let

$$\pi_\sigma = (T_\sigma)_{\#}\pi.$$

The marginal of π_σ corresponding to agent $\sigma(i)$ is a^i. Hence the vector of marginals of π_σ is precisely σa. Since $\pi_\sigma \in \Delta(Z)$, it follows that

$$\sigma a \in A.$$

Thus A is permutation invariant. □

The reverse direction is also valid for the allocation spaces considered here. Let

$$D(Z) = \{(\delta_{z_1}, \dots, \delta_{z_n}) : z \in Z\} \subseteq A.$$

The map

$$z \mapsto (\delta_{z_1}, \dots, \delta_{z_n})$$

is a topological embedding of Z into $\Delta(Y)^n$. Moreover,

$$D(Z) = A \cap \{\text{vectors of Dirac measures}\}.$$

Since the set of Dirac measures is closed under weak convergence, compactness of A implies compactness of $D(Z)$, and hence compactness of Z. Similarly, if A is permutation invariant and $z \in Z$, then the permuted vector

$$(\delta_{z_{\sigma^{-1}(1)}}, \ldots, \delta_{z_{\sigma^{-1}(n)}})$$

belongs to A. Any joint probability measure having only Dirac marginals must itself be concentrated on the corresponding deterministic allocation. Hence the permuted deterministic allocation belongs to Z. Thus permutation invariance of A implies permutation invariance of Z.

A.2 Differentiated Commodities

Let (K, d) be a compact metric space and let $\alpha \in \mathbb{N}$. Define

$$Y = \{\mu : \mu \text{ is a nonnegative integer-valued Borel measure on } K, \quad \mu(K) \leq \alpha\}.$$

Proposition A.3. *The set Y is compact and metrizable in the weak* topology.*

Proof. First observe that every $\mu \in Y$ is purely atomic and has at most α atoms. Indeed, since μ is integer valued and $\mu(K) \leq \alpha$, every point in its support has integer mass at least one.

Consider the set

$$\mathcal{M}_\alpha = \{\mu \geq 0 : \mu(K) \leq \alpha\}.$$

Identifying finite Borel measures on K with positive elements of $C(K)^*$, the set $\mathcal{M}_\alpha$ is weak* compact. Since K is compact metric, $C(K)$ is separable, so the weak* topology on $\mathcal{M}_\alpha$ is metrizable. It therefore suffices to show that Y is weak* closed in $\mathcal{M}_\alpha$.

Take a sequence $\mu_m \in Y$ such that

$$\mu_m \xrightarrow{w^*} \mu.$$

Write

$$\mu_m = \sum_{k=1}^{r_m} q_{m,k} \delta_{x_{m,k}},$$

where $r_m \leq \alpha$, the points $x_{m,k} \in K$ are distinct, and $q_{m,k} \in \mathbb{N}$. Since

$$\sum_{k=1}^{r_m} q_{m,k} \leq \alpha,$$

there are only finitely many possible values of r_m and of the integer vector $(q_{m,1}, \ldots, q_{m,r_m})$.

Passing to a subsequence and relabeling if necessary, we may therefore assume that

$$r_m = r$$

and

$$q_{m,k} = q_k$$

for every m and every $k = 1, \dots, r$.

Since K is compact, after passing to a further subsequence we may assume that

$$x_{m,k} \longrightarrow x_k$$

for every $k = 1, \dots, r$. The limit points x_k need not be distinct.

For every $f \in C(K)$,

$$\int_K f \, d\mu_m = \sum_{k=1}^{r} q_k f(x_{m,k}) \longrightarrow \sum_{k=1}^{r} q_k f(x_k).$$

On the other hand, weak* convergence gives

$$\int_K f \, d\mu_m \longrightarrow \int_K f \, d\mu.$$

Hence

$$\int_K f \, d\mu = \sum_{k=1}^{r} q_k f(x_k)$$

for every $f \in C(K)$. By uniqueness in the Riesz representation theorem,

$$\mu = \sum_{k=1}^{r} q_k \delta_{x_k}.$$

If several of the limit points x_k coincide, their coefficients add. Since each q_k is an integer, the resulting mass at every point is again an integer. Moreover,

$$\mu(K) = \sum_{k=1}^{r} q_k \leq \alpha.$$

Thus $\mu \in Y$.

Therefore Y is weak* closed in the compact metrizable set $\mathcal{M}_\alpha$. Hence Y is compact and metrizable. $\square$

A.3 Failure of existence without permutation invariance

Consider Example 1 of Pazner and Schmeidler [1974]. There are two consumers and two goods, leisure l and consumption z, with utilities

$$u_1(l,z) = l + \frac{11}{10}z, \qquad u_2(l,z) = l + 2z.$$

The feasible set is

$$Z = \left\{((l_1,z_1),(l_2,z_2)) \in \mathbb{R}^4_+ : 0 \le l_j \le 1, \quad z_1 + z_2 \le (1-l_1) + \frac{1}{10}(1-l_2)\right\}.$$

Pazner and Schmeidler show that there is no allocation in Z that is both envy-free and Pareto efficient.

In fact, there is no allocation that is both envy-free and weakly Pareto efficient. Since Z is convex and compact and utilities are linear, every weakly Pareto efficient utility vector maximizes

$$\lambda_1 u_1 + \lambda_2 u_2$$

for some $\lambda \in \mathbb{R}^2_+ \setminus \{0\}$. If $\lambda_1, \lambda_2 > 0$, every maximizer is Pareto efficient, and hence cannot be envy-free by the argument of Pazner and Schmeidler [1974]. It therefore remains only to consider the boundary weights.

For $\lambda = (1,0)$,

$$u_1(l_1,z_1) \le l_1 + \frac{11}{10}\left((1-l_1) + \frac{1}{10}(1-l_2)\right) = \frac{121}{100} - \frac{1}{10}l_1 - \frac{11}{100}l_2.$$

Hence the unique maximizer of u_1 is

$$((l_1,z_1),(l_2,z_2)) = \left(\left(0, \frac{11}{10}\right), (0,0)\right).$$

This allocation is not envy-free, since consumer 2 strictly prefers consumer 1's bundle:

$$u_2\left(0, \frac{11}{10}\right) = \frac{11}{5} > 0 = u_2(0,0).$$

For $\lambda = (0,1)$,

$$u_2(l_2,z_2) \le l_2 + 2\left((1-l_1) + \frac{1}{10}(1-l_2)\right) = \frac{11}{5} - 2l_1 + \frac{4}{5}l_2.$$

Hence the unique maximizer of u_2 is

$$((l_1,z_1),(l_2,z_2)) = ((0,0),(1,1)).$$

This allocation is not envy-free, since consumer 1 strictly prefers consumer 2's bundle:

$$u_1(1,1) = \frac{21}{10} > 0 = u_1(0,0).$$

Thus no weakly Pareto efficient allocation in this economy is envy-free. Consequently, Example 1 of Pazner and Schmeidler [1974] shows that dropping permutation invariance can destroy existence of an envy-free and weakly Pareto efficient allocation even in a convex environment with linear utilities.

A.4 Strictly concave regularizer

To ensure uniqueness of the maximizer, I modify the concave objective by adding a continuous, strictly concave term $-G : A \to \mathbb{R}$. The existence of a strictly convex continuous function on a compact convex set is nontrivial. Moreover, I require G to be permutation invariant:

$$G(\sigma a) = G(a) \qquad \text{for every } a \in A \text{ and } \sigma \in S_N.$$

I now show that such a function exists.

For each $i \in N$, let

$$Y_i := \mathrm{pr}_i(Z)$$

denote the set of deterministic consumption bundles that agent i can receive in some feasible deterministic allocation. Since Z is permutation invariant,

$$Y_i = Y_j \qquad \text{for every } i, j \in N.$$

Denote this common compact metric space by Y_Z. Every marginal a^i of a feasible random allocation is supported on Y_Z, so $A \subseteq \Delta(Y_Z)^n$.

Since Y_Z is compact and metrizable, $\Delta(Y_Z)$ is compact, convex, and metrizable in the weak* topology. By a result of Hervé Alfsen [1971], a compact convex set admits a continuous strictly convex real-valued function if and only if it is metrizable. Hence there exists a continuous strictly convex function

$$F : \Delta(Y_Z) \to \mathbb{R}.$$

Define

$$G(a) := \sum_{i \in N} F(a^i), \qquad a = (a^1, \ldots, a^n) \in A.$$

The function G is continuous and strictly convex. Indeed, if $a \neq b$, then $a^i \neq b^i$ for at least one i, and strict convexity of F, together with convexity of F in the remaining coordinates, implies

$$G(ta + (1-t)b) < tG(a) + (1-t)G(b) \qquad \text{for every } t \in (0,1).$$

Finally, G is permutation invariant: for every $\sigma \in S_N$,

$$G(\sigma a) = \sum_{i \in N} F\Big((\sigma a)^i\Big) = \sum_{i \in N} F(a^i) = G(a).$$

Thus $-G$ is continuous, strictly concave, and permutation invariant.

A.5 Existence theorem

Theorem A.4 (Berge's maximum theorem)**.** *Let X and Θ be topological spaces. Let $f : X \times \Theta \to \mathbb{R}$ be continuous on the product space $X \times \Theta$. Let $C : \Theta \to\to X$ be a non-empty compact-valued correspondence. Define the argmax correspondence by*

$$C^*(\theta) = \{x \in X : f(x, \theta) \geq f(\tilde{x}, \theta) \ \forall \tilde{x} \in C(\theta)\}$$

If C is continuous at θ, then $C^*(\theta)$ is upper hemicontinuous at θ and non-empty compact-valued. Note: a set-valued map that is singleton-valued is upper-hemicontinous if and only if the corresponding function is continuous.

Theorem A.5. *Let E be a metrizable locally convex topological vector space and let $A \subseteq E^n$ be compact and permutation invariant. Suppose that each agent's preferences on E are represented by a continuous and concave utility function. Then there exists an allocation $a^* \in A$ that is envy-free and weakly Pareto efficient.*

The proof proceeds in two steps. The main step is the following lemma. Recall that I denote the unique maximizer of $\max_{a \in A} Q(\lambda, \delta, a)$ by $a^*(\lambda, \delta)$.

Lemma A.6. *For every $\delta > 0$, there exists $\lambda_\delta \in \Delta^{n-1}$ such that $a^*(\lambda_\delta, \delta)$ is envy-free.*

Proof. Fix a sequence of real numbers $c_k \downarrow 0$.

For every $k \in \mathbb{N}$, simplicially subdivide the simplex Δ^{n-1} so that the diameter of every subsimplex is less than c_k. By Sperner's lemma, for every k, there exists a completely labeled subsimplex, denoted by Δ_k^{n-1}. Choose $\lambda_k \in \Delta_k^{n-1}$.

This procedure produces a sequence such that, for each k, λ_k belongs to a completely labeled subsimplex of diameter less than c_k. By sequential compactness of Δ^{n-1}, there exists a convergent subsequence λ_{k_l}. Denote its limit by λ_δ.

I claim that $a^*(\lambda_\delta, \delta)$ is envy-free. Toward a contradiction, suppose that it is not. Then there exists a pair of agents (i, j) such that agent i envies agent j, that is,

$$g_{ij}(\lambda_\delta) = U_i\Big(a_i^*(\lambda_\delta, \delta)\Big) - U_i\Big(a_j^*(\lambda_\delta, \delta)\Big) < 0.$$

By Berge's Maximum Theorem above, $a^*(\cdot, \delta)$ is continuous for every $\delta > 0$. Since U_i are continuous, $g_{ij}(\cdot)$ is continuous. It follows that there exists $\epsilon > 0$ such that, for every λ satisfying

$$d(\lambda, \lambda_\delta) < \epsilon,$$

we have $g_{ij}(\lambda) < 0$. Thus, there is a neighborhood of λ_δ such that, for every welfare weight in that neighborhood, agent i envies agent j at the maximizing allocation.

Since $\lambda_{k_l} \to \lambda_\delta$ and $c_{k_l} \to 0$, for sufficiently large l,

$$d(\lambda_{k_l}, \lambda_\delta) < \frac{\epsilon}{2} \qquad \text{and} \qquad c_{k_l} < \frac{\epsilon}{2}.$$

Then, for every $\lambda \in \Delta^{n-1}_{k_l}$,

$$d(\lambda, \lambda_\delta) \leq d(\lambda, \lambda_{k_l}) + d(\lambda_{k_l}, \lambda_\delta) < c_{k_l} + \frac{\epsilon}{2} < \epsilon.$$

Therefore, $\Delta^{n-1}_{k_l}$ cannot be completely labeled, a contradiction. □

I can now complete the proof of the theorem. A useful feature of envy-freeness is that, under continuity of utilities, it is preserved under limits, since its defining inequalities are weak.

Proof of the theorem. Let $\delta_l \downarrow 0$. By sequential compactness of Δ^{n-1}, passing to a subsequence if necessary, $\lambda_{\delta_l} \to \bar{\lambda}$ for some $\bar{\lambda} \in \Delta^{n-1}$. By sequential compactness of A, passing to a further subsequence if necessary, $a^*(\lambda_{\delta_l}, \delta_l) \to \bar{a}$ for some $\bar{a} \in A$.

By Berge's Maximum Theorem, $\bar{a} \in \arg\max_{a \in A} Q(\bar{\lambda}, 0, a)$. Since $\bar{\lambda} \in \Delta^{n-1}$, the allocation $\bar{a}$ is weakly Pareto efficient.

Moreover, for every l, $a^*(\lambda_{\delta_l}, \delta_l)$ is envy-free. Hence, for every pair of agents i, j,

$$U_i\Big(a_i^*(\lambda_{\delta_l}, \delta_l)\Big) \geq U_i\Big(a_j^*(\lambda_{\delta_l}, \delta_l)\Big).$$

By continuity of U_i, taking $l \to \infty$ yields

$$U_i(\bar{a}_i) \geq U_i(\bar{a}_j).$$

Therefore, $\bar{a}$ is envy-free. □

B. Appendix B

B.1 Continuity VS Lipschitz continuity

Fix $\delta > 0$ and define the envy of agent i at allocation a by

$$E_i(a) := \max_{j \neq i} U_i(a_j) - U_i(a_i).$$

Since U_i is continuous, E_i is continuous. Moreover, since $a_\delta : \Delta^{n-1} \to A$ is continuous and Δ^{n-1} is compact, $E_i \circ a_\delta$ is uniformly continuous. Since N is finite, for every $\epsilon > 0$ there exists $\gamma > 0$ such that, for every $i \in N$,

$$\|\lambda - \lambda'\|_2 < \gamma \quad \Longrightarrow \quad |E_i(a_\delta(\lambda)) - E_i(a_\delta(\lambda'))| < \epsilon.$$

Take a subdivision of Δ^{n-1} whose simplices have diameter less than γ, and let T' be a fully labeled simplex. Choose any $\widetilde{\lambda} \in T'$. For every $i \in N$, there is a vertex $\lambda^i \in T'$ labeled by i.

Hence

$$\|\widetilde{\lambda} - \lambda^i\|_2 < \gamma.$$

By the labeling rule, agent i is envy-free at $a_\delta(\lambda^i)$, so $E_i(a_\delta(\lambda^i)) \leq 0$. Therefore,

$$E_i(a_\delta(\widetilde{\lambda})) < E_i(a_\delta(\lambda^i)) + \epsilon \leq \epsilon.$$

Thus $a_\delta(\widetilde{\lambda})$ is ϵ-envy-free.

B.2 Input size of a polytope

For an integer z, let

$$\langle z \rangle := 1 + \lceil \log_2(|z| + 1) \rceil .$$

For a rational $q = p/r$ in lowest terms, with $r > 0$, let

$$\langle q \rangle := \langle p \rangle + \langle r \rangle.$$

For

$$A = \{a \in \mathbb{R}^{dn} : Ca \leq b\}, \qquad C \in \mathbb{Q}^{m \times dn}, \quad b \in \mathbb{Q}^m,$$

the binary input size is

$$s := \langle m \rangle + \langle dn \rangle + \sum_{r=1}^{m} \sum_{k=1}^{dn} \langle C_{rk} \rangle + \sum_{r=1}^{m} \langle b_r \rangle.$$

As a solid reference for standard rational Turing encoding; see Grötschel–Lovász–Schrijver (1993), Chapters 2 and 8.

B.3 Strong separation oracle for the set of admissible marginals with implicit characterization

Recall that I denote by $Z \subseteq Y^N$ the set of feasible deterministic allocations. Assume Y is finite. For a probability measure $\pi \in \Delta(Z)$, let $p = (p^1, \ldots, p^n)$ denote its vector of marginal distributions, where $p^i \in \Delta(Y)$ is the marginal distribution of agent i. Now A is the set of admissible marginals:

$$A = \left\{ (p^1, \ldots, p^n) \in \Delta(Y)^N : \exists\, \pi \in \Delta(Z) \text{ whose marginals are } (p^1, \ldots, p^n) \right\}.$$

As established earlier, A is a convex compact subset, hence admits a characterization via support functions that is useful computationally. A profile $p = (p^1, \ldots, p^n)$ belongs to A if and only if, for every collection of functions $\phi_i : Y \to \mathbb{R}$,

$$\mathbb{E}_p \left[\sum_{i \in N} \phi_i(z_i) \right] \leq \max_{a \in A} \sum_{z \in Z} a(z) \sum_{i \in N} \phi_i(z_i) = \max_{z \in Z} \sum_{i \in N} \phi_i(z_i) \tag{6}$$

Assume that $n|Y|$ is polynomially bounded in the input size. Then the standard (rational)

well-describedness conditions of the GLS Theorem (see Grötschel et al. [1988, Definition 6.2.2, Lemma 6.2.4, and Theorem 6.4.9]) are satisfied, and thus polynomial-time linear optimization and polynomial-time strong separation are equivalent for A. Hence polynomial-time optimization over the deterministic feasible set Z provides the computational access to A required by Theorem 10.5.

Strong separation oracle on a convex hull . Take a finite set of primitive goods G, and suppose that a single agent is only allowed to consume goods in *paired configurations*. For instance, the primitive objects could be $2m$ components that must be assembled into m two-component units, and a feasible consumption bundle is a perfect matching on complete non-bipartite graph of those $2m$ components. Then the single-agent deterministic consumption set is

$$Y_m = \left\{\chi^M : M \text{ is a perfect matching on } K_{2m}\right\} \subseteq \{0,1\}^{\binom{2m}{2}},$$

and its randomized, or convexified, version is the perfect-matching polytope

$$P_m = \operatorname{conv}(Y_m) \subseteq R^{E(K_{2m})}.$$

Its ambient dimension is only

$$|E(K_{2m})| = \binom{2m}{2} = O(m^2),$$

so the dimension is polynomial in the natural parameter m.

However, Edmonds' description contains, in addition to the degree equations and nonnegativity constraints, the odd-cut inequalities

$$x(\delta(S)) \geq 1 \qquad \text{for every odd } S \subseteq V,$$

with

$$3 \leq |S| \leq 2m - 3.$$

There are exponentially many such subsets S, and many of these inequalities are facet-defining. Hence the polytope has exponentially many facets as m grows.

Nevertheless, P_m admits a polynomial-time strong separation oracle. Given a rational point

$$x \in \mathbb{Q}^{E(K_{2m})},$$

one first checks the degree and nonnegativity constraints,

$$x_e \geq 0 \qquad \text{and} \qquad x(\delta(v)) = 1 \quad \text{for all } v \in V.$$

The remaining question is whether there exists an odd set $S \subseteq V$ such that

$$x(\delta(S)) < 1.$$

Finding a minimum-weight odd cut can be done in polynomial time via the Padberg-Rao separation procedure [Padberg and Rao, 1982]. If such a set S exists, the oracle returns the violated inequality

$$x(\delta(S)) \geq 1.$$

If no such odd set exists, then

$$x \in P_m.$$

Thus, the perfect-matching polytope is a standard example of a polytope with an exponentially large explicit facet description that is nevertheless polynomial-time separable.

C. Appendix C

C.1 Prefix representation and extended formulations

Even when Z_T is finite-dimensional, its projection onto allocation variables need not have a convenient direct inequality description. The state variables in P_T therefore serve as extended variables:

$$Z_T = \mathrm{proj}_{z_{0:T}} P_T.$$

This distinction is computationally useful because the complexity of a projection can be much larger than the complexity of a lifted representation. Nothing in Part II, however, requires the projected set itself to be given by a polynomial-size explicit H-representation. What matters is the computational access assumed there. A polynomial-size rational extended formulation is one sufficient route to such access, but it is not imposed as a primitive assumption here.

An exponential gap between a polytope and its extended formulation . Consider the n-dimensional cross-polytope

$$Z_n := \left\{ x \in \mathbb{R}^n : \sum_{i=1}^{n} |x_i| \leq 1 \right\}.$$

Its inequality representation is

$$Z_n = \left\{ x \in \mathbb{R}^n : \sum_{i=1}^{n} \sigma_i x_i \leq 1 \quad \forall \sigma \in \{-1, 1\}^n \right\}.$$

Each of these 2^n inequalities defines a distinct facet. Hence, any direct inequality representation of Z_n requires at least 2^n inequalities.

Now introduce auxiliary variables $u \in \mathbb{R}^n$ and define the extended polytope

$$P_n := \left\{ (x, u) \in \mathbb{R}^{2n} : \begin{array}{l} -u_i \leq x_i \leq u_i, \quad i = 1, \ldots, n, \\ \sum_{i=1}^{n} u_i \leq 1. \end{array} \right\}.$$

This formulation requires only $2n + 1$ inequalities.

To see that $Z_n = \mathrm{proj}_x(P_n)$, observe that every $(x, u) \in P_n$ satisfies

$$\sum_{i=1}^{n} |x_i| \leq \sum_{i=1}^{n} u_i \leq 1.$$

Conversely, for any $x \in Z_n$, choosing $u_i = |x_i|$ for every i yields $(x, u) \in P_n$.

Thus,

$$Z_n = \mathrm{proj}_x(P_n),$$

although Z_n has 2^n facets and P_n admits an extended formulation with only $2n+1$ inequalities.

C.2 Connection between $A_{\leq T}$ and $\Delta(P_T)$ for indivisibles

Lemma C.1. *Let*

$$q_T : P_T \to Z_T, \qquad q_T(z_{0:T}, s_{0:T+1}) = z_{0:T}.$$

Then

$$(q_T)_{\#}\Delta(P_T) = \Delta(Z_T).$$

Consequently,

$$A_{\leq T} = M_T\Big((q_T)_{\#}\Delta(P_T)\Big).$$

Proof. By construction, $q_T(P_T) = Z_T$. First, take any $\nu \in \Delta(P_T)$. Since $q_T(P_T) = Z_T$, its push-forward $(q_T)_{\#}\nu$ is a Borel probability measure supported on Z_T. Hence $(q_T)_{\#}\nu \in \Delta(Z_T)$. Therefore

$$(q_T)_{\#}\Delta(P_T) \subseteq \Delta(Z_T).$$

Conversely, take any $\mu^T \in \Delta(Z_T)$. Since $q_T : P_T \to Z_T$ is a continuous surjection between compact metric spaces, there exists a Borel measurable selection $r_T : Z_T \to P_T$ such that $q_T \circ r_T = \mathrm{id}_{Z_T}$. Define

$$\nu := (r_T)_{\#}\mu^T \in \Delta(P_T).$$

Then

$$(q_T)_{\#}\nu = (q_T)_{\#}(r_T)_{\#}\mu^T = (q_T \circ r_T)_{\#}\mu^T = \mu^T.$$

Thus

$$\Delta(Z_T) \subseteq (q_T)_{\#}\Delta(P_T),$$

which proves

$$(q_T)_{\#}\Delta(P_T) = \Delta(Z_T).$$

The final statement follows by applying the marginalization map M_T. □

C.3 Continuity at infinity

For a map $U_i : A_i \to \mathbb{R}$, continuity at infinity is defined by $\omega(T) \to 0$, where

$$\omega(T) = \sup\{|U_i(a_i) - U_i(a_i')| : a_i, a_i' \in A_i \text{ and } \pi_{\leq T}(a_i) = \pi_{\leq T}(a_i')\} \tag{7}$$

Definition C.2. U_i admits a uniform approximation by finite-prefices if there exists $U_i^T : A_i^T \to \mathbb{R}$ such that $\eta(T) \longrightarrow 0$ and

$$\sup_{a_i \in A_i} \left| U_i(a_i) - U_i^T\Big(\pi_{\leq T}(a_i)\Big) \right| \leq \eta(T), \tag{8}$$

Lemma C.3. *$U_i : A_i \to \mathbb{R}$ is continuous at infinity iff it admits a uniform approximation by finite-prefices.*

Proof. Assume (8). Then for any a_i, a_i' with $\pi_{\leq T}(a_i) = \pi_{\leq T}(a_i')$

$$|U_i(a_i) - U_i(a_i')| \leq \left| U_i(a_i) - U_i^T(\pi_{\leq T}(a_i)) \right| + \left| U_i(a_i') - U_i^T(\pi_{\leq T}(a_i')) \right|$$

Then $\omega(T) \leq 2\eta(T) \to 0$.

Assume (7). For $x \in \pi_{\leq T}(A_i)$, let

$$U_i^T(x) = \frac{1}{2}\Big(\sup_{a_i : \pi_{\leq T}(a_i) = x} U_i(a_i) + \inf_{a_i : \pi_{\leq T}(a_i) = x} U_i(a_i) \Big)$$

Then

$$\left| U_i(a_i) - U_i^T\Big(\pi_{\leq T}(a_i)\Big) \right| \leq \frac{1}{2} \left| U_i(a_i) - \sup_{b_i : \pi_{\leq T}(b_i) = \pi_{\leq T}(a_i)} U_i(a_i) \right| + \frac{1}{2} \left| U_i(a_i) - \inf_{b_i : \pi_{\leq T}(b_i) = \pi_{\leq T}(a_i)} U_i(a_i) \right|$$

hence $\eta(T) \leq \omega(T) \to 0$. □

C.4 Railroads example

A richer example with the same continuation structure is the allocation of railway capacity. Let $G = (V, E)$ be a finite directed network of stations and track segments, and let $\tau_e \in \mathbb{N}$ be the travel time on edge e. Each agent is a railway operator with a fixed stock of rolling stock. In the original notation, $d_{t,e}^i$ denoted the number of trains of operator i that start traversing edge e at date t, while $s_{t,u}^i$ denoted the number of its trains available at station u. In the notation of this paper, let $z_{i,t} = (z_{i,t,e})_{e \in E} \in Y_t$, where $z_{i,t,e} = d_{t,e}^i$. Let B^- and B^+ denote the origin and destination incidence matrices of G. The stock-flow equation can then be written in the delayed form

$$s_{t+1}^i = s_t^i - B^- z_{i,t} + B^+\Big(z_{i,t+1-\tau_e,e}\Big)_{e \in E}.$$

Thus a train that starts traversing $e = (u, v)$ at date t leaves the stock at u immediately and enters the stock at v only after τ_e periods. In the benchmark case $\tau_e = 1$ for every e,

$$s_{t+1}^i = s_t^i + B z_{i,t}, \qquad B := B^+ - B^-,$$

which is the usual network-flow conservation equation. For heterogeneous travel times, the lagged departures can equivalently be included in the state, so that the system takes the first-order form $(s_t, z_t, s_{t+1}) \in \Gamma_t$ used above.

The shared resource is not merely the entrance to an edge but its occupancy during the entire traversal. If edge e can simultaneously accommodate $C_{t,e}$ trains, feasibility requires

$$\sum_{i \in N} \sum_{r=0}^{\tau_e - 1} z_{i,t-r,e} \leq C_{t,e}.$$

The unit-capacity version considered above has $C_{t,e} = 1$. Departures must also be supported by rolling stock available at their origin,

$$\sum_{e \in \delta^+(u)} z_{i,t,e} \leq s_{i,t,u},$$

and the model can accommodate separate station, departure, and arrival capacities, for example

$$\sum_{i \in N} s_{i,t,u} \leq C^{\text{stat}}_{t,u}, \qquad \sum_{i \in N} \sum_{e \in \delta^+(u)} z_{i,t,e} \leq C^{\text{dep}}_{t,u}, \qquad \sum_{i \in N} \sum_{e \in \delta^-(u)} z_{i,t+1-\tau_e,e} \leq C^{\text{arr}}_{t+1,u}.$$

These constraints have an equivalent time-expanded-network representation: every physical edge $e = (u, v)$ generates an arc

$$(u, t) \longrightarrow (v, t + \tau_e),$$

while waiting is represented by arcs

$$(u, t) \longrightarrow (u, t + 1).$$

Flow conservation in this network simultaneously records the location of rolling stock and the delayed effect of travel decisions. In particular, the formulation uses local edge-time variables rather than variables indexed by complete train paths; an earlier path formulation, in which a railway path $p = ((e_1, t_1), \ldots, (e_k, t_k))$ was treated as a single object and conflicting paths competed for common resources, is unnecessary once the same constraints are represented locally on the time-expanded network.

The model still admits a canonical continuation. After any finite history, set all new physical-edge departures equal to zero. Trains already in transit complete the trips that have already been initiated, and after at most $\bar{\tau} := \max_{e \in E} \tau_e$ additional periods all rolling stock is located at stations; thereafter it can remain there on waiting arcs forever. Thus, unlike the airline model, the state need not become constant immediately at the truncation date, but no new allocation decisions are required after that date. A finite-prefix formulation need only retain the finite pipeline of previously initiated trips—equivalently, extend the time-expanded network by at most $\bar{\tau}$ dates to complete them. Hence deterministic travel times, multi-period track occupancy, and rolling-stock conservation enlarge the finite-dimensional representation only polynomially. Railway capacity therefore provides a more involved member of the same class of dynamic allocation problems with trivial continuation, without requiring a separate viability-kernel computation.

D. Appendix D: Incentives counter-example

Proposition D.1. *For three agents and three indivisible objects, no random assignment mechanism is simultaneously strategy-proof, envy-free, and weakly Pareto efficient under expected-utility preferences.*

Proof. Suppose such a mechanism exists. Restrict its domain to

$$\mathcal{D} = \{u_t = (1, t, 0) : 0 < t < 1/6\},$$

where the coordinates correspond to objects A, B, C. Let $\mathcal{B}_3$ denote the set of bistochastic matrices.

We first show that, on $\mathcal{D}^3$, envy-freeness (EF) and weak Pareto efficiency (wPE) imply Pareto efficiency (PE).

Fix a profile $u \in \mathcal{D}^3$ and an EF and wPE allocation $P \in \mathcal{B}_3$. By EF and bistochasticity, each agent i satisfies

$$u_i(P_i) \geq \frac{1}{3}\sum_{j=1}^{3} u_i(P_j) = \frac{1 + t_i}{3}.$$

Since $u_i(P_i) = p_{iA} + t_i p_{iB}$ and $p_{iB} \leq 1$,

$$p_{iA} \geq \frac{1 - 2t_i}{3} > 0 \qquad \text{for every } i.$$

By convexity of the feasible utility set, wPE implies that P maximizes

$$\sum_{i=1}^{3} \lambda_i u_i(P_i)$$

for some nonzero $\lambda \geq 0$.

We claim that $\lambda \gg 0$. Otherwise, take agents i, j with $\lambda_i = 0$ and $\lambda_j > 0$. Since $p_{iA} > 0$, agent j must receive positive probability of some $g \in \{B, C\}$. Transferring a small probability of A from i to j, in exchange for the same probability of g, strictly increases weighted welfare, a contradiction.

Thus $\lambda \gg 0$, and hence P is PE.

Consequently, the restricted mechanism is strategy-proof, Pareto efficient, and envy-free. Since envy-freeness implies symmetry (equal expected utilities for agents with identical preferences), this contradicts the three-agent impossibility argument of Zhou [1990, proof of Theorem 1], which uses only profiles in $\mathcal{D}^3$. □