The Power of Sperner's Lemma in Fair Division

Anna Vakarova

UC Berkeley

Allocation under scarcity

How should scarce resources be allocated
among finitely many agents
with different preferences?

Efficiency and envy-freeness

Efficiency of $a \in A$

$\downarrow$

Weak Pareto efficiency

$$\nexists b \in A: \; U_i(b_i) > U_i(a_i) \quad \forall i.$$

Do not want to leave gains from trade on the table

Fairness of $a \in A$

$\downarrow$

Envy-freeness

$$U_i(a_i) \geq U_i(a_j) \quad \forall i, j.$$

How those gains are distributed

Envy-freeness and efficiency may conflict

Two agents divide one unit of each of two goods. An allocation is $a = (a_1, a_2)$, where $a_i(j)$ is agent i's share of good j.

$$U_1(a_1) = 7a_1(1) + 3a_1(2), \qquad U_2(a_2) = 6a_2(1) + 4a_2(2).$$

Weak Pareto efficient
NOT envy-free

$$a^{\mathrm{EW}} = ((1,0),(0,1)) \in \arg\max_{a \in A} \{U_1(a_1) + U_2(a_2)\}$$

	a_1	a_2
U_1	7	3
U_2	**6**	**4**

Envy-free
NOT efficient

$$a^{\mathrm{eq}} = \left(\left(\tfrac{1}{2}, \tfrac{1}{2}\right), \left(\tfrac{1}{2}, \tfrac{1}{2}\right)\right)$$

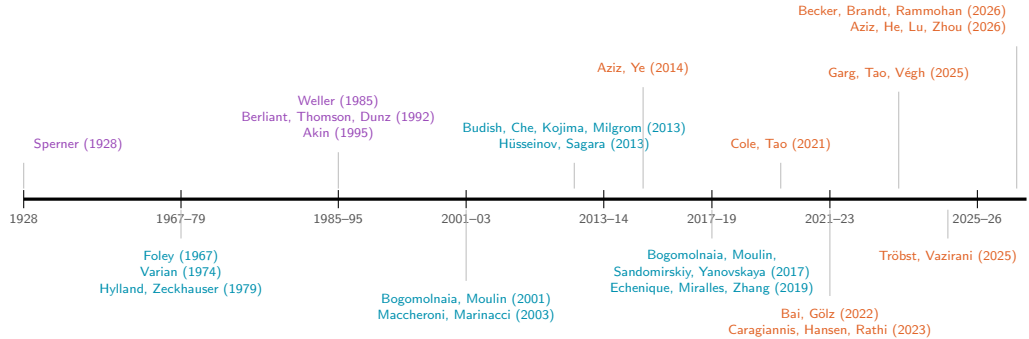
	a_1	a_2
U_1	5	5
U_2	5	5

$b = ((.7,.1),(.3,.9))$ gives
$(U_1(b_1), U_2(b_2)) = (5.2, 5.4) > (5,5)$.

$\Longrightarrow$ **Joint existence is nontrivial**

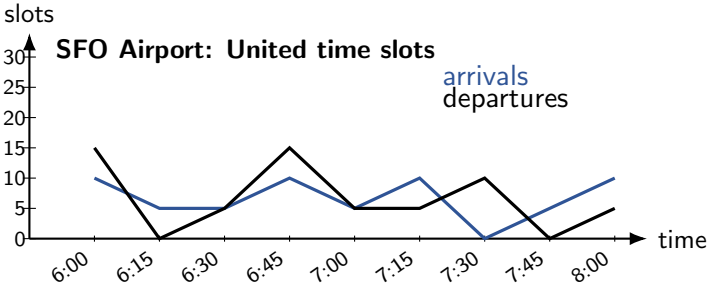
Is there an allocation that satisfies both?

Fairness and efficiency in the literature



Application: airport take-off/landing slots allocation

Example of a long-term consumption bundle stream



- Airlines book **take-off and landing slots in advance**.
- An **independent coordinator** respects neutrality, transparency, and non-discrimination:
  - FAA slot admin office in the US (statutory authority 49 U.S.C. §40103(b))
  - ACL, COHOR etc in the EU (EU Council Regulation No 95/93)
- Business models create **different preferences over schedules**.

Airport slot allocation: feasibility

An individual assignment is the infinite stream

$$z_i = (r_i, d_i) = \big(r_i(t), d_i(t)\big)_{t\geq 0}.$$

Individual feasibility: stock-flow

$r_{i,l}(t),\ d_{i,l}(t) \in \mathbb{Z}_+$	integrality and nonnegativity
$S_{i,l}(t+1) = S_{i,l}(t) + r_{i,l}(t) - d_{i,l}(t)$	stock flow
$S_{i,l}(t) \geq 0, \qquad \sum_l S_{i,l}(t) = F$	fleet feasibility
$\sum_l r_{i,l}(t) = \sum_l d_{i,l}(t) \leq Q_t$	operating limit

Aggregate capacity constraints

$$\sum_i r_{i,l}(t) \leq C^{\mathrm{arr}}_{t,l}, \qquad \sum_i d_{i,l}(t) \leq C^{\mathrm{dep}}_{t,l}.$$

Contribution: Part I – Existence

1. Indivisibles without transfers

Let (Y, d_Y) be a metric consumption space and let

$$Z \subseteq Y^n$$

be the set of feasible deterministic allocations.

This paper

Z may be any compact, permutation-invariant set.

finite feasible set $\longrightarrow$ compact feasible set

This step allows me to allocate infinite streams of consumption and work with intertemporal constraints.

Contribution: Part I – Existence

2. Divisibles

Let Y be a consumption space with a **linear structure** and let

$$Z \subseteq Y^n$$

be the set of feasible allocations.

This paper

Z may be any compact, permutation-invariant, **convex** set.

finite-dimensional feasible set $\longrightarrow$ infinite-dimensional feasible set

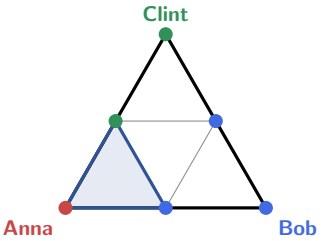
This step allows me to allocate natural resources (water, energy) over long-term horizons.

Contribution: Part I – Existence

At the core of the result is the application of a powerful tool

Sperner's lemma

to the simplex of **Pareto weights**



Goods, bads, mixtures

Contribution: Part II – Algorithm

This paper: a fixed number of agents

Output

- ε-EF;
- ε-efficient.

Running time

- polynomial in input size s[a];
- polynomial in $1/\varepsilon$;

[a] Input size is the number of bits required to encode the input.

Main contribution: applicability to a broad class of problems

Contribution: Part III – Dynamic Allocation Problems

I apply both the existence and algorithmic approximation results to **infinite-horizon dynamic allocation problems**:

1. Allocating takeoff and landing slots at airports
2. Allocating railroad capacity
3. Allocating natural-resource extraction rights

Existence

Deterministic allocations

Let $N = \{1, \ldots, n\}$ be a finite set of agents and (Y, d_Y) a metric consumption space.

A deterministic allocation is

$$z = (z_1, \ldots, z_i, \ldots, z_k, \ldots, z_n) \in Y^n.$$

The set of **feasible deterministic allocations** is denoted by

$$Z \subseteq Y^n.$$

Assumptions on Z:

- Z is compact;
- Z is permutation invariant: if $z = (z_1, \ldots, \mathbf{z}_i, \ldots, \mathbf{z}_k, \ldots, z_n) \in Z$, then

$$\sigma_{i,k}(z) = (z_1, \ldots, \mathbf{z}_k, \ldots, \mathbf{z}_i, \ldots, z_n) \in Z.$$

Random allocations

A **lottery** is a Borel probability measure over the set of feasible deterministic allocations:

$$\mu \in \Delta(Z).$$

A **random allocation** is a vector of marginal distributions on Y.

The set of feasible random allocations is

$$A = \{(a_1, \ldots, a_n) \in \Delta(Y)^n : \exists \mu \in \Delta(Z) \text{ with marginals } (a_1, \ldots, a_n)\}.$$

Examples

Multi-unit demand
There are m indivisible object types. Object j has c_j copies, and each agent receives at most k units in total.

Feasible partitions
Each feasible partition consists of n pairwise disjoint bundles. Once the bundles are formed, they may be assigned to agents in any order.

In both examples, Z is compact metric (finite) and permutation-invariant.

Divisibles

Y is a vector space, $A \subseteq Y^N$, and $a = (a_1, \ldots, a_n) \in A$ directly specifies every agent's consumption bundle: no randomization is required.

- **Divisible goods** Foley (1967); Varian (1974)
  There are d divisible goods with aggregate endowment $\omega \in \mathbb{R}^d_+$.
- **Mixed manna** Bogomolnaia et al. (2017)
  The same items may be goods for some agents and bads for others.

Assumptions on A

- compactness;
- permutation invariance;
- convexity.

Preferences

Agent i's utility depends only on her consumption bundle a_i.

Preferences over consumption bundles are represented by

$$U_i : A_i \to \mathbb{R},$$

where U_i is **continuous** and **concave**.

No local nonsatiation. No monotonicity.

Teaser: relaxation to ordinary convex preferences (quasi-concave utilities) is a work-in-progress with existence already extended.

Examples of preferences

Expected utility
For a continuous VNM utility index $u_i : Y \to \mathbb{R}$,

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

Linear in lotteries

Maxmin utility under risk Let $\mathcal{U}_i \subset C(Y)$ be compact and convex in the sup norm. Then

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

Concave in lotteries

Maccheroni (2002); Cerreia-Vioglio et al. (2015)

Main existence result

Theorem

Let E be a metrizable locally convex topological vector space and $A \subseteq E^n$ compact, convex, and permutation invariant. If each $U_i : E \to \mathbb{R}$ is continuous and concave, an envy-free and weak Pareto efficient $a^* \in A$ exists.

Envy-free

$$U_i(a_i^*) \geq U_i(a_j^*) \quad \forall i, j$$

Weakly efficient

$$\nexists b : \; U_i(b_i) > U_i(a_i^*) \quad \forall i$$

Teaser: relaxation to ordinary convex preferences (quasi-concave utilities) is a work-in-progress with existence already extended.

The central search problem

Two agents divide one unit of each of two goods. An allocation is $a = (a_1, a_2)$, where $a_i(j)$ is agent i's share of good j.

$$U_1(a_1) = 7a_1(1) + 3a_1(2), \qquad U_2(a_2) = 6a_2(1) + 4a_2(2).$$

Equal weights produce the weakly efficient allocation

$$a^{\mathrm{EW}} = ((1,0),(0,1)),$$

in which agent 2 envies.

Search over all Pareto weights:

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

Is some welfare maximizer also envy-free?

Pareto weights and regularization

$$\Delta^{n-1} = \{\lambda \in \mathbb{R}^n_+ : \sum_i \lambda_i = 1\}.$$

For $\delta > 0$ maximize

$$\sum_i \lambda_i U_i(a_i) - \delta G(a),$$

where G is continuous, strictly convex, and permutation invariant. The unique maximizer $a(\lambda, \delta)$ varies continuously with $(\lambda, \delta > 0)$.

The regularizer is a search device. It vanishes at the end.

Lucky Lemma

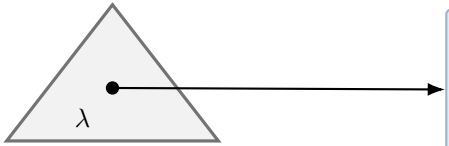
$$\chi(\lambda) = \{i : \lambda_i > 0\}.$$

Lucky Lemma

For $(\lambda, \delta > 0)$, at $a(\lambda, \delta)$ there exists $i \in \chi(\lambda)$ who is envy-free.

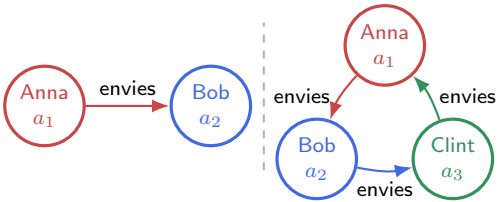
Lucky labeling: every weight vector can be labeled by an envy-free agent in its carrier.

Lucky Lemma: envy



$$\max_{a \in A} \left\{ \sum_i \lambda_i U_i(a_i) - \delta G(a) \right\} \longrightarrow a(\lambda, \delta)$$

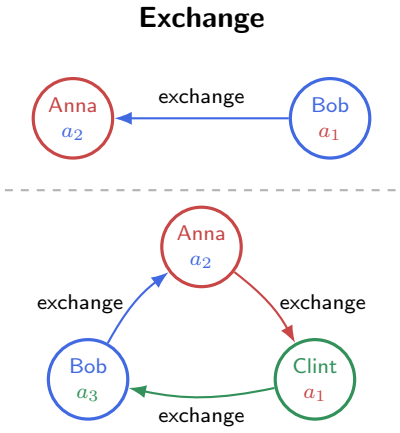
Envy graphs ⟵ **By contradiction**



Assume that every agent in the carrier envies someone. Let agents **point** to an agent of envy. Then either

1. A path is formed by an agent in the carrier to an agent outside it. If no path, then
2. There exists an envy cycle formed by a subset of agents in the carrier.

Lucky Lemma: exchange along the cycle

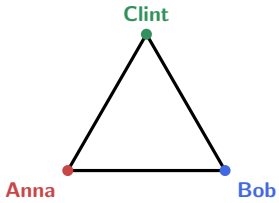


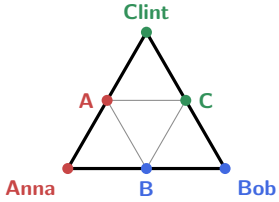
$$\sum_i \lambda_i U_i(a_i) - \delta G(a) \qquad \textbf{vs.}$$

$$\sum_i \lambda_i U_i\left(a_i^{\text{exchange}}\right) - \delta G\left(a^{\text{exchange}}\right)$$

1. Every supported agent weakly increases her utility, with at least one strict increase.
2. $G(a) = G\left(a^{\text{exchange}}\right)$.
3. The regularized objective strictly increases, a contradiction.

Simplex of Pareto weights





Lucky labeling l:

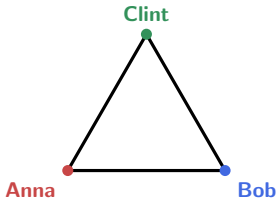
$$l(v) = \begin{array}{c}\text{name of the agent}\\ \text{who is envy-free at vertex } v\end{array}$$

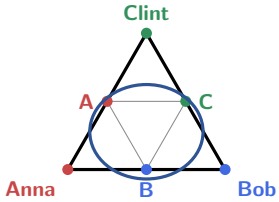
By the Lucky Lemma, the Lucky labeling is **proper**: the name of a vertex **inside a facet** is chosen **only** from the names of that facet's boundary vertices.

Sperner's Lemma

Every properly labeled subdivision contains a fully labeled subsimplex.

Simplex of Pareto weights





Lucky labeling l:

$$l(v) = \begin{array}{c}\text{name of the agent}\\ \text{who is envy-free at vertex } v\end{array}$$

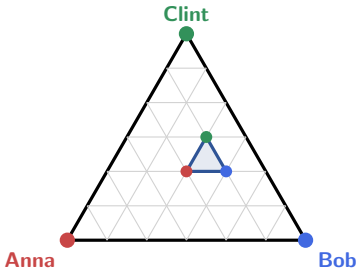
By the Lucky Lemma, the Lucky labeling is **proper**: the name of a vertex **inside a facet** is chosen **only** from the names of that facet's boundary vertices.

Sperner's Lemma

Every properly labeled subdivision contains a fully labeled subsimplex.

Finding envy-free Pareto weights

1. Subdivide Δ^{n-1}.
2. Label using the Lucky labeling.
3. Apply Sperner's lemma and find a completely labeled subsimplex.
4. Repeat ad infinitum for smaller and smaller diameter.



The limiting weight has an envy-free regularized maximizer.

Sending the perturbation to zero

Choose $\delta_k \downarrow 0$ and λ_k so $a_k = a(\lambda_k, \delta_k)$ is envy-free. Compactness gives

$$(\delta_k, \lambda_k, a_k) \longrightarrow (0, \lambda^*, a^*)$$

along a subsequence.

a_k EF $\xrightarrow{\text{continuity}}$ a^* EF

$a_k \in \arg\max\{W_{\lambda_k} - \delta_k G\}$ $\xrightarrow{\text{upper hemicontinuity}}$ $a^* \in \arg\max W_{\lambda^*}$

a^* is envy-free and weak Pareto efficient.

From existence to algorithmic approximation

Contribution

1. Broader set of applications
2. Relaxation from linear utilities to concave utilities

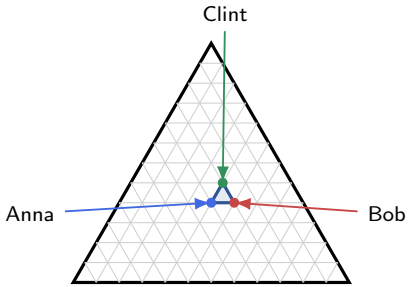
How?



Sperner's lemma gives me an algorithm to find an allocation that is ε-efficient and ε-EF.

How?

1. Find a fully labeled subsimplex of small diameter.
2. Pick a vector of Pareto weights from it.



Pareto weight changes a little $\Longrightarrow$ maximizing allocation changes a little $\Longrightarrow$ marginal bundles change a little $\Longrightarrow$ each agent's utility changes a little

What does the algorithm take?
Input

- requested precision $\varepsilon > 0$
- concave, L-Lipschitz utilities $U_i : Y \to \mathbb{R}$, where $Y \subseteq \mathbb{R}^d$

A nonempty feasible set $A \subseteq (\mathbb{R}^d)^n$

compact and convex | permutation invariant | $\|a\|_F^2 \leq B$ for every $a \in A$ | strong separation oracle[1]

Example: an H-representation[2]

$$A = \{a \in \mathbb{R}^{dn} : Ca \leq b\}.$$

[1] Computational complexity discussion at request.
[2] Input size of the instance.

What does the algorithm produce?

Output

An allocation $a \in A$ such that

1. ε**-envy-free**

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

Unequivocal in the literature

2. ε**-efficient**

$$\forall b \in A,\ \exists i \in N: \quad U_i(b_i) - U_i(a_i) \leq \varepsilon.$$

Debatable: two definitions

Alternative notion: $\mathrm{dist}(a, \mathrm{wPE}(A)) \leq \varepsilon$ (not guaranteed here).

Why this notion?

- Natural when searching over the set of normalized supporting hyperplanes (the "dual").
- Allocation-space approximation is delicate under regularization.

Approximation algorithm

Input: a feasible set $A \subseteq (\mathbb{R}^d)^n$ with a strong separation oracle; utilities $\{U_i\}_{i \in N}$; and L, B, $\varepsilon > 0$.

Parameter selection

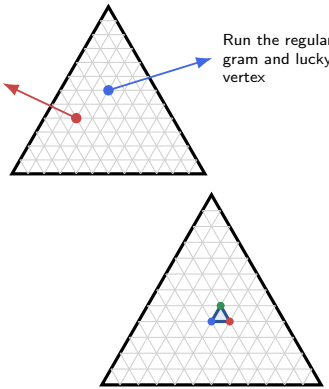
1. Compute $\delta = \varepsilon / B$.
2. Compute $\sigma = \varepsilon \delta / L^2$.

Fixed-point computation

3. Construct a simplicial subdivision of Δ^{n-1} with mesh at most 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

For $\epsilon > 0$, under the stated assumptions, the procedure above outputes $a \in A$ that is ϵ-EF and ϵ-efficient.

Conclusion

- **Scope:** One framework covers randomized indivisible allocations and divisible resources, including goods, bads, and mixtures.
- **Existence:** Compactness, convexity, permutation invariance, and continuous concave utilities deliver envy-freeness and weak Pareto efficiency without monotonicity.
- **Proof strategy:** Search over Pareto weights, regularize the welfare problem, and apply the Lucky Lemma with Sperner's lemma.
- **Computation:** For fixed n, a quantitative version of the same search produces ε-envy-free and ε-efficient allocations in polynomial time.

Infinite-horizon dynamic allocation problems

At every date $t \in \mathbb{N}_0$, an agent's consumption lies in Y_t. A consumption stream is

$$Y = \prod_{t=0}^{\infty} Y_t, \qquad z = (z_1, \ldots, z_n) \in Y^N, \qquad z_t = (z_{1,t}, \ldots, z_{n,t}) \in Y_t^N.$$

Dynamic feasibility. Let $s_t \in S_t$ be the state and

$$\Gamma_t \subseteq S_t \times Y_t^N \times S_{t+1}, \qquad (s_t, z_t, s_{t+1}) \in \Gamma_t.$$

This includes the standard controlled system $s_{t+1} = F_t(s_t, z_t)$.

$$\widetilde{Z} = \left\{ (z, s) \in Y^N \times \prod_{t=0}^{\infty} S_t : s_0 = \bar{s}_0,\ (s_t, z_t, s_{t+1}) \in \Gamma_t\ \forall t \right\},$$

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

Preferences

Let $a_i \in \Delta(Y)$, let $\mu_{i,t} = (\pi_t)_{\#} a_i \in \Delta(Y_t)$, and let m_i be a probability measure on $\mathbb{N}_0$.

Linear expected utility
For continuous $u_{i,t} : Y_t \to \mathbb{R}$,

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

Concave utility of marginals
For continuous, concave $u_{i,t} : \Delta(Y_t) \to \mathbb{R}$,

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

Assume $\sum_{t \geq 0} m_i(\{t\}) \|u_{i,t}\|_\infty < \infty$ so the series are well defined.
Different measures m_i allow agents to value dates differently.

Existence in dynamic environments

Corollary 12.1

Suppose that A is generated by a dynamic allocation problem as above. If each

$$U_i : A_i \to \mathbb{R}$$

is continuous and concave, then there exists an allocation $a \in A$ that is envy-free and weakly Pareto efficient.

Questions