Online Fair Allocation of Indivisible Goods

Recent Progress and Open Questions

Joseph Chuang-Chieh Lin

Department of Computer Science & Engineering
National Taiwan Ocean University

14 August 2026

Outline

Outline

Canonical Online Goods Model

At each time $t = 1, 2, \dots, T$:

1. one indivisible good g_t arrives;
2. values $(v_1(g_t), v_2(g_t), \dots, v_n(g_t))$ are revealed, estimated, or partially observed;
3. the algorithm assigns g_t to one agent;
4. in the irrevocable model, this decision cannot be changed.

Final allocation:

$$A_i^T = \{g_t : g_t \text{ was assigned to } i\}.$$

Four Axes of Online Fair Allocation

1. Arrival model

adversarial / stochastic / predicted

2. Commitment model

irrevocable / batched / reallocation

3. Information model

full values / binary likes / bandit feedback

4. Fairness horizon

final-time / prefix-wise / asymptotic

An online fair-allocation model is specified by choosing one alternative along each axis.

Fairness Horizons

Final-time fairness

Require fairness only for A^T.

Prefix-wise fairness

Require fairness for every A^t, $t \leq T$.

Asymptotic fairness

Require normalized unfairness to vanish, e.g.

$$\text{Envy}_T / T \to 0.$$

Commitment Models

**How strongly an online algorithm is bound by its earlier allocation decisions.*

irrevocable → batched → reallocation

stable but hard | delay decisions | fairer but disruptive

- irrevocable: $A_i^{t-1} \subseteq A_i^t, \quad \forall i \in N$; easy to implement; no reassignment cost.
- batched: waits until a collection of goods has arrived.
- reallocatin: an algorithm is allowed to repair earlier decisions; additional cost.

Commitment Models

**How strongly an online algorithm is bound by its earlier allocation decisions.*

irrevocable → batched → reallocation

irrevocable	batched	reallocation
stable but hard	delay decisions	fairer but disruptive

- irrevocable: $A_i^{t-1} \subseteq A_i^t, \quad \forall i \in N$; easy to implement; no reassignment cost.
- batched: waits until a collection of goods has arrived.
- reallocatin: an algorithm is allowed to repair earlier decisions; additional cost.

A major question: how much "changing the past" is necessary for fairness?

Outline

Online EF1: Two Very Different Worlds

Question

Can an irrevocable online allocation maintain EF1 for the accumulated bundles

$$A^t = (A_1^t, A_2^t, \dots, A_n^t)$$

after every arriving good?

general additive — impossible in general → **identical additive** — attainable by least-valued bundle

Key distinction on valuations

- Heterogeneous: different agents evaluate the same bundle on different scales.
- Identical: all bundle values are measured by one common function v.

A Three-Good Counterexample: Setup

We use two agents and one arriving good at a time. Fix any constant $K > 1$.

arriving good	v_1	v_2
g_1	1	1
g_2	K	$\frac{1}{K}$
g_3	K	K

Adaptive adversary convention

After g_1 is allocated, rename the recipient as agent 1 and the other agent as agent 2. Thus after the first arrival,

$$A_1^1 = \{g_1\}, \qquad A_2^1 = \emptyset.$$

This allocation is EF1.

Round 2: EF1 Forces the Choice

The second good has values

$$v_1(g_2) = K, \qquad v_2(g_2) = \frac{1}{K}.$$

Round 2: EF1 Forces the Choice

The second good has values

$$v_1(g_2) = K, \qquad v_2(g_2) = \frac{1}{K}.$$

g_2 to agent 1

$$A_1 = \{g_1, g_2\}, \quad A_2 = \emptyset.$$

Agent 2's own value is 0. Removing one good from A_1 leaves

$$v_2(\{g_1\}) = 1, \quad v_2(\{g_2\}) = \frac{1}{K}.$$

Both are positive, so EF1 fails.

Round 2: EF1 Forces the Choice

The second good has values

$$v_1(g_2) = K, \qquad v_2(g_2) = \frac{1}{K}.$$

g_2 to agent 1

$$A_1 = \{g_1, g_2\}, \quad A_2 = \emptyset.$$

Agent 2's own value is 0. Removing one good from A_1 leaves

$$v_2(\{g_1\}) = 1, \quad v_2(\{g_2\}) = \frac{1}{K}.$$

Both are positive, so EF1 fails.

g_2 to agent 2

$$A_1 = \{g_1\}, \quad A_2 = \{g_2\}.$$

EF1 holds:

$$v_1(A_2 \setminus \{g_2\}) = 0 \le 1,$$

$$v_2(A_1 \setminus \{g_1\}) = 0 \le \frac{1}{K}.$$

Therefore every EF1-preserving online rule must give g_2 to agent 2.

Round 3: No EF1-Preserving Choice Remains

Before g_3 arrives, any EF1-preserving algorithm is forced to have

$$A_1 = \{g_1\}, \qquad A_2 = \{g_2\}.$$

Now g_3 arrives with

$$v_1(g_3) = v_2(g_3) = K.$$

Round 3: No EF1-Preserving Choice Remains

Before g_3 arrives, any EF1-preserving algorithm is forced to have

$$A_1 = \{g_1\}, \qquad A_2 = \{g_2\}.$$

Now g_3 arrives with

$$v_1(g_3) = v_2(g_3) = K.$$

g_3 to agent 1

Agent 2 has value $\frac{1}{K}$. For agent 2,

$$v_2(A_1) = 1 + K.$$

Even after removing one good:

$$K > \frac{1}{K}, \qquad 1 > \frac{1}{K}.$$

So EF1 fails.

Round 3: No EF1-Preserving Choice Remains

Before g_3 arrives, any EF1-preserving algorithm is forced to have

$$A_1 = \{g_1\}, \qquad A_2 = \{g_2\}.$$

Now g_3 arrives with

$$v_1(g_3) = v_2(g_3) = K.$$

g_3 to agent 1

Agent 2 has value $\frac{1}{K}$. For agent 2,

$$v_2(A_1) = 1 + K.$$

Even after removing one good:

$$K > \tfrac{1}{K}, \qquad 1 > \tfrac{1}{K}.$$

So EF1 fails.

g_3 to agent 2

Agent 1 has value 1. For agent 1,

$$v_1(A_2) = 2K.$$

Removing either good leaves value K:

$$K > 1.$$

So EF1 fails.

Consequence: No General Online EF1 Guarantee

Theorem (Impossibility for heterogeneous additive valuations)

For general additive valuations, no deterministic irrevocable online algorithm can guarantee that the accumulated allocation is EF1 after every arriving good.

Keys to the Impossibility

The obstruction is not indivisibility alone. It is the combination of

- heterogeneous values,
- irrevocability, and
- lack of future information.

Positive Result for **Identical** Additive Valuations

Assume all agents share **the same additive valuation** v.

Least-valued-bundle rule

When g_t arrives, allocate it to an agent with minimum current bundle value:

$$r \in \arg\min_{i \in N} v(A_i^{t-1}).$$

Then update

$$A_r^t = A_r^{t-1} \cup \{g_t\}, \qquad A_i^t = A_i^{t-1} \quad (i \neq r).$$

Theorem

For identical additive valuations, the least-valued-bundle rule maintains EF1 after every arriving good.

Proof: EF1 Is Preserved (1/2)

We prove the theorem by induction on t.

- Initially, all bundles are empty, so EF1 holds.
- Assume A^{t-1} is EF1.
- Let $r \in \arg\min_i v(A_i^{t-1})$ receive the new good g_t.

Only agent r's bundle changes. Hence old EF1 witnesses still work except possibly for envy *toward* the new bundle of r.

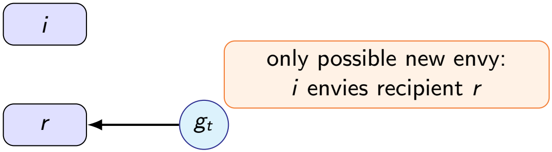

Proof: EF1 Is Preserved (2/2)

Fix any agent $i \neq r$. Remove the newly assigned good from r's bundle:

$$A_r^t \setminus \{g_t\} = A_r^{t-1}.$$

Since r was least-valued before the update,

$$v(A_r^{t-1}) \leq v(A_i^{t-1}).$$

Also, agent i's bundle did not change, so

$$v(A_i^{t-1}) = v(A_i^t).$$

Therefore,

$$v(A_i^t) \geq v(A_r^t \setminus \{g_t\}).$$

The good g_t itself eliminates every possible new envy toward r. Hence EF1 is preserved after the arrival of g_t.

Outline

Item-Centric Mechanisms

Offline algorithms often ask: *which good does this agent pick?*

Online algorithms often ask: *which agent should receive this arriving good?*

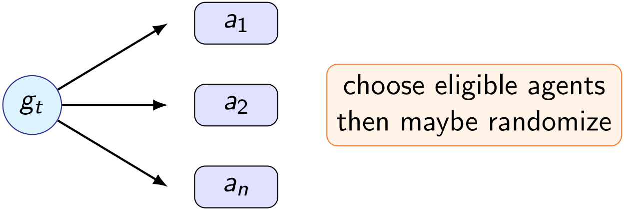

LIKE Mechanism ([Aleksandrov et al. 2015] & [Aleksandrov & Walsh 2020])

Definition (LIKE)

When g_t arrives, allocate it uniformly at random among agents who report positive value for it.

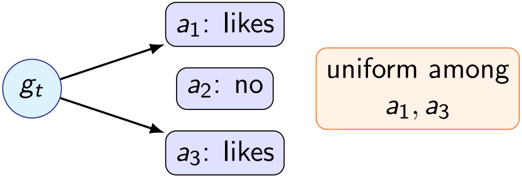

- simple;
- online strategy-proof in one-step binary settings;
- envy-free ex ante in some restricted domains.

Balanced LIKE

Definition (BALANCED LIKE)

Among agents who like the item, allocate uniformly to those with the ***fewest*** *previously allocated liked items.*

Design goal

Balance the number of allocated liked items over time.

Strategic caveat

An agent may **skip a current liked item** to become **more eligible for future items**.

One-Step Strategyproofness of LIKE

Proposition

In binary utilities, truthful reporting is weakly optimal for the current item under LIKE.

Proof.

If $v_i(g_t) = 0$, receiving the item gives utility 0, so reporting yes or no gives current utility 0.
If $v_i(g_t) = 1$, reporting no gives probability 0 of receiving g_t. Reporting yes gives probability $\frac{1}{k}$ for some $k \geq 1$, hence expected utility $\frac{1}{k} > 0$. $\square$

Outline

1. Online Models
2. Online Mechanisms & Problems
   - Impossibility and Possibility of an Online EF1 Allocation
   - LIKE-Type Online Mechanisms
   - Vanishing Envy
   - EF1 with Reallocation
   - Online MMS
   - Bandit Learning for Online Fair Allocation
   - Nash Welfare and the PACE Algorithm
   - Online Fair Division with Subsidies
3. Related Directions
4. Open Problems

Cumulative Envy

Definition (Envy after T rounds)

$$E_{ij}^{T} = v_i(A_j^T) - v_i(A_i^T),$$

$$Envy_T = \max_{i,j} \max\{0, E_{ij}^T\}.$$

Vanishing envy [Benade et al. (EC'18)]

An algorithm has vanishing envy if

$$\text{Envy}_T / T \to 0.$$

Random Allocation: A Clean Bound

Theorem ([Benade et al. (EC'18)])

If $v_i(g_t) \in [0,1]$ *and* **each item is allocated uniformly at random**, *then with probability at least* $1-\delta$,

$$Envy_T \leq \sqrt{2T\log\left(\frac{n(n-1)}{\delta}\right)}.$$

Consequence

$$\frac{\mathrm{Envy}_T}{T} = O\left(\sqrt{\frac{\log n}{T}}\right)$$

with high probability.

Proof of the Random Allocation Bound

For a fixed pair (i, j), define

$$X_t = v_i(g_t)(\mathbf{1}\{g_t \to j\} - \mathbf{1}\{g_t \to i\}).$$

Then:

- $E_{ij}^T = \sum_{t=1}^T X_t$;
- $\mathbb{E}[X_t] = 0$;
- $X_t \in [-1, 1]$.

Hoeffding gives

$$\Pr\left[E_{ij}^T \geq r\right] \leq e^{-r^2/(2T)}.$$

Union bound over $n(n-1)$ ordered pairs and set

$$r = \sqrt{2T \log\left(\frac{n(n-1)}{\delta}\right)}.$$

Optimal Vanishing Envy

Theorem (Benade et al. (EC'18))

For adversarial online allocation of T goods with values in $[0,1]$, there is a deterministic polynomial-time algorithm with

$$Envy_T \in \widetilde{O}(\sqrt{T/n}),$$

and this is asymptotically optimal up to logarithmic factors.

Remark

Exact fairness may be impossible online, but normalized unfairness can vanish.

Outline

Changing the Past

- Requirement: after every time t, allocation A^t is EF1.
- To make this possible, old goods may be reassigned.
- **Objective: minimize the number of adjustments.**

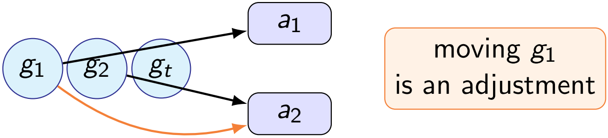

EF1 with Adjustments: Main Theorem

Theorem (He–Procaccia–Psomas–Zeng @IJCAI'19)

For online additive goods with EF1 required after every round:

1. *two agents: informed algorithms need no adjustments, but uninformed algorithms need $\Theta(T)$;*
2. *three or more agents: even informed algorithms need $\Omega(T)$ adjustments;*
3. *for $n \geq 3$, an uninformed algorithm uses $O(T^{3/2})$ adjustments.*

Interpretation of the EF1 Adjustment Result

- With two agents, future information can eliminate instability.
- Without future information, the adversary can force repeated repairs.
- With ≥ 3 agents, pairwise envy constraints interact cyclically.
- The online objective becomes bicriteria:

$$\text{fairness quality} \quad + \quad \text{stability of past allocations.}$$

Outline

Online MMS Benchmark

Offline MMS is defined using all goods:

$$\mu_i^n(M) = \max_{(B_1,B_2,\ldots,B_n)} \min_k v_i(B_k).$$

In online allocation, the algorithm does not know final M_T.

Definition (α-competitive online MMS (Zhou–Bai–Wu @ICML'23))

An online algorithm is α-competitive for MMS goods if

$$v_i(A_i^T) \geq \alpha \mu_i^n(M_T) \qquad \forall i.$$

Online MMS for Goods

Theorem (Zhou–Bai–Wu @ICML'23)

For online additive indivisible goods:

- *no positive competitive ratio is possible for MMS when $n \geq 3$;*
- *for two agents, there is an optimal $\frac{1}{2}$-competitive algorithm.*

MMS is a global benchmark. Adversarial online arrivals can make it unattainable for ≥ 3 agents.

Why MMS Is Fragile Online

early allocation → adversarial continuation → offline MMS benchmark

the future can make an early losing agent's MMS positive

Irrevocable decisions cannot be corrected after the final benchmark becomes clear.

Goods vs. Chores

Theorem (Zhou–Bai–Wu @ICML'23)

For online chores:

- *there is a* $(2 - \frac{1}{n})$*-competitive algorithm for* $n \geq 3$*;*
- *there is a* $\sqrt{2}$*-competitive algorithm for two agents;*
- *no algorithm can do better than* $\frac{15}{11}$ *for two agents.*

Reason for different behavior

For chores, load-balancing analogies provide stronger online structure.

Outline

Honor Among Bandits: Why Relevant?

Problem type

Online fair allocation of indivisible goods with unknown stochastic values.

- There are n players and m item types.
- At each round t, an item of type $k_t \in [m]$ arrives.
- The item must be allocated immediately to one player.
- Player i's value for type k has unknown mean μ^*_{ik}.
- The algorithm observes only the reward of the selected player.

Procaccia, Schiffer, & Zhang @NeurIPS'24

online fair division $+$ unknown valuations $+$ bandit feedback.

Model: Fractional Allocation Policy

A fractional allocation is a matrix

$$X \in \mathbb{R}_{\geq 0}^{n\times m}, \qquad \sum_{i=1}^{n} X_{ik} = 1 \quad \forall k \in [m].$$

- X_{ik} is the probability that a type-k item is allocated to player i.
- If the true mean matrix is μ^*, the expected welfare is

$$\langle X, \mu^* \rangle = \sum_{i=1}^{n} \sum_{k=1}^{m} X_{ik}\mu^*_{ik}.$$

At round t:

$$k_t \text{ arrives}, \quad i_t \sim X_{:,k_t}, \quad \text{observe reward of } (i_t, k_t).$$

Fairness in Expectation

The paper studies welfare maximization subject to fairness constraints defined using the unknown mean matrix μ^*.

Envy-freeness in expectation

For every pair of players i, j,

$$\sum_{k=1}^{m} X_{ik}\mu^*_{ik} \geq \sum_{k=1}^{m} X_{jk}\mu^*_{ik}.$$

Player i does not prefer player j's allocation in expectation.

Proportionality in expectation

For every player i,

$$\sum_{k=1}^{m} X_{ik}\mu^*_{ik} \geq \frac{1}{n}\sum_{k=1}^{m} \mu^*_{ik}.$$

Regret Benchmark

If μ^* were known, one could solve the fair welfare maximization LP:

$$Y_{\mu^*} \in \arg\max_{X \in \mathcal{F}(\mu^*)} \langle X, \mu^* \rangle,$$

$\mathcal{F}(\mu^*)$: the set of feasible fractional allocations satisfying the fairness constraints.

Regret

For an online algorithm choosing $X_1, \ldots, X_T$,

$$R(T) = T\langle Y_{\mu^*}, \mu^* \rangle - \sum_{t=1}^{T} \mathbb{E}[\langle X_t, \mu^* \rangle].$$

Goal: low regret while satisfying fairness constraints.

Main Result: Explore-Then-Commit

Algorithmic idea

1. **Explore:** allocate items to estimate the unknown means μ^*_{ik}.
2. **Construct confidence intervals:** use the estimates to build a safe approximation of the fair feasible region.
3. **Commit:** maximize estimated welfare over allocation policies that are fair for every valuation matrix in the confidence region.

Guarantee

For envy-freeness in expectation or proportionality in expectation, the explore-then-commit algorithm achieves $\widetilde{O}(T^{2/3})$ regret while maintaining the fairness constraint.

The paper also proves a matching $\widetilde{\Omega}(T^{2/3})$ lower bound for this setting.

Bandit Learning Result

Theorem (Schiffer–Zhang @AAAI'2025)

There is a UCB-style online fair division algorithm that satisfies proportionality constraints with high probability and achieves

$$\widetilde{O}(\sqrt{T})$$

welfare regret.

Key improvement

This improves the previous best regret rate of $\widetilde{O}(T^{2/3})$.

Why UCB Can Work with Fairness Constraints

- UCB estimates unknown values optimistically.
- Linear optimization enforces proportionality constraints.
- Confidence intervals keep constraint violations controlled.
- The main challenge is many potentially tight fairness constraints.

observations → UCB estimates → fair LP → allocation

Outline

Nash Welfare in Online Allocation

Definition (Nash welfare)

$$NSW(A) = \left(\prod_{i=1}^{n} v_i(A_i) \right)^{1/n}.$$

- Offline MNW is EF1 and Pareto optimal.
- Online algorithms use accumulated utilities to guide allocation.
- Long-run guarantees are often asymptotic rather than exact.

Multiplicative Envy

Definition (Multiplicative envy)

A typical multiplicative envy ratio compares

$$\frac{v_i(A_j)}{v_i(A_i)}.$$

- Additive envy asks for differences.
- Multiplicative envy asks for ratios.

PACE: Pacing According to Current Estimated Utility

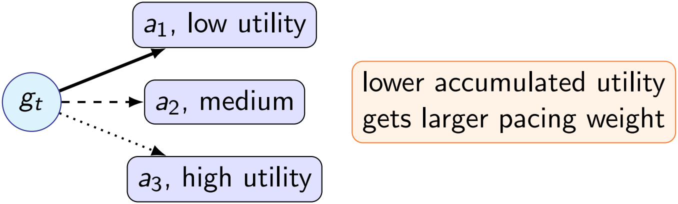

- Derived from dual averaging applied to a reformulation of the Eisenberg–Gale dual program.
- A market-equilibrium pacing method with a multiplicative-update flavor.

Outline

Online Fair Division with Subsidies

Online model

Indivisible goods $g_1, g_2, \ldots$ arrive sequentially. When g_t arrives:

- its value information becomes available;
- it must be allocated immediately and irrevocably;
- the number of future goods is unknown.

Let $X^t = (X_1^t, X_2^t, \ldots, X_n^t)$ denote the allocation after the first t goods.

Definition (Envy-freeness with subsidy)

Given a nonnegative subsidy vector $p^t = (p_1^t, p_2^t, \ldots, p_n^t) \in \mathbb{R}^n_{\geq 0}$, the outcome (X^t, p^t) is envy-free if

$$v_i(X_i^t) + p_i^t \geq v_i(X_j^t) + p_j^t \quad \forall i, j \in N.$$

An allocation X^t is **envy-freeable** if **such a subsidy vector exists**.
[Kulkarni et al. @AAMAS 2026.]

Envy-Freeability and Local Efficiency (1/2)

Definition (Local efficiency)

An allocation $X = (X_1, X_2, \ldots, X_n)$ is locally efficient if no permutation of its bundles **increases (utilitarian) social welfare***:*

$$\sum_{i \in N} v_i(X_i) \geq \sum_{i \in N} v_{\pi(i)}(X_i) \quad \text{for every permutation } \pi : N \to N.$$

Theorem (Envy-freeability characterization)

An allocation is envy-freeable if and only if it is locally efficient.

Online implication

Because the arrival process may *terminate after any round*, an online algorithm must maintain *local efficiency* after every *prefix*:

$$X^t \text{ must be locally efficient for every } t.$$

Why Local Efficiency Captures Envy-Freeability

Why envy implies local efficiency

If (X, p) is envy-freeable, agent $\pi(i)$ must not prefer X_i:

$$v_{\pi(i)}(X_{\pi(i)}) + p_{\pi(i)} \geq v_{\pi(i)}(X_i) + p_i.$$

Sum over i. Since $\sum_i p_{\pi(i)} = \sum_i p_i$, the subsidies cancel:

$$\sum_i v_i(X_i) \geq \sum_i v_{\pi(i)}(X_i).$$

$\Longrightarrow$ X is locally efficient.

The converse: the envy graph

Let $w(i,j) := v_i(X_j) - v_i(X_i)$. The subsidy constraints becomes $p_i - p_j \geq w(i,j)$.

$$X \text{ is locally efficient} \iff \text{no positive-weight cycle} \iff \exists \mathbf{p} : \; p_i - p_j \geq w(i,j) \\ \iff X \text{ is envy-freeable.}$$

Envy-Freeability and Local Efficiency (2/2)

Theorem (Additive valuations [Kulkarni et al. @AAMAS 2026])

For additive valuations, when good g_t arrives, assigning it to

$$i^* \in \arg\max_{i \in N} v_i(g_t)$$

preserves local efficiency and therefore preserves envy-freeability.

When Can Envy-Freeability Be Maintained Online?

Valuation class	Online envy-freeability
Additive	Yes
SPLC	Yes
k-demand	Yes
Budget-additive	No in general
Submodular with binary marginals	No in general
Supermodular with binary marginals	No in general

Subsidy cost may be large

Assuming every marginal value is at most 1, the total subsidy needed online can be as large as $\Omega\big(m(n-1)\big)$ even for additive valuations.

Classes with item-independent subsidy bounds (Kulkarni et al., @AAMAS 2026))

Bounds independent of m are obtained for structured classes, including fixed-k demand, fixed-k valued, rank-one, restricted additive, and identical monotone valuations.

Outline

Divisible Online Allocation

- Some online fair allocation work studies **divisible** resources.
- Divisibility avoids single-good impossibilities but retains online uncertainty.
- With normalized values, nontrivial fairness–efficiency tradeoffs become possible.

Example

Gkatzelis–Psomas–Tan (AAAI'21) study immediately and irrevocably allocated divisible resources under normalized valuations.

Learning-Augmented Fairness (e.g., Aziz et al. @NeurIPS'25)

- Prediction may concern future arrivals, item types, values, or the optimal outcome.
- Desired guarantees:

Consistency

Good performance when predictions are accurate.

Robustness

Worst-case protection when predictions are wrong.

Open direction: consistency–robustness for online indivisible goods.

Outline

Open Problem: Online EF1

Open Problem

For which valuation domains and arrival models can an irrevocable online algorithm guarantee final-time or prefix-wise EF1?

- General additive valuations seem difficult.
- Binary and bivalued valuations are promising.
- Limited reallocation may give a useful middle ground.

Open Problem: Online MMS

Open Problem

After the no-competitive-ratio result for goods with $n \geq 3$, what relaxations make online MMS meaningful?

- stochastic arrivals;
- predictions of future types;
- resource augmentation;
- bounded reallocation;
- restricted valuations.

Open Problems: Online Fairness with Subsidies

Open Problem

For which valuation classes *can an online algorithm maintain envy-freeability with* **total subsidy independent of the number of arriving goods***?*

- Characterize the boundary between valuation classes that admit online envy-freeability and those that do not.
- Determine tight subsidy bounds for broader valuation classes.
- Study the trade-off between subsidy, welfare, and online fairness.
- Consider models in which the subsidy budget is fixed in advance.
- Study truthful online mechanisms when allocation and subsidy decisions are both strategic.

[Motivated by Kulkarni et al., @AAMAS 2026]

Open Problem: Best of Both Worlds

Open Problem

Can we obtain ex ante fairness and ex post approximate fairness simultaneously in online indivisible allocation?

References I

[1] Aleksandrov and Walsh: Online fair division: A survey. AAAI 2020.

[2] Amanatidis et al.: Fair division of indivisible goods: Recent progress and open questions. AIJ 2023.

[3] Benade, Kazachkov, Procaccia, Psomas: How to make envy vanish over time. EC 2018.

[4] He, Procaccia, Psomas, Zeng: Achieving a fairer future by changing the past. IJCAI 2019.

[5] Zhou, Bai, Wu: Multi-agent online scheduling: MMS allocations for indivisible items. ICML 2023.

[6] Yang, Liao, Kroer: Greedy-based online fair allocation with adversarial input. AAAI 2024. Newer version: Yang, Liao, Gao, Kroer: Online Fair Allocation with Best-of-Many-Worlds Guarantees. arXiv 2024.

References II

[7] Schiffer and Zhang: Improved regret bounds for online fair division with bandit learning. AAAI 2025.

[8] Wang and Wei: Online fair allocations with binary valuations and beyond. AAAI 2026.

[9] Gkatzelis, Psomas, Tan: Fair and efficient online allocations with normalized valuations. AAAI 2021.

[10] Aziz, Guo, Lam, Zhou: Learning-augmented facility location mechanisms for envy ratio. NeurIPS 2025.

[11] P. Kulkarni, R. Mehta, V. V. Narayan, and T. Ponitka: Online fair division with subsidy: When do envy-free allocations exist, and at what cost? AAMAS 2026.

[12] A. D. Procaccia, B. Schiffer, and S. Zhang: Honor among bandits: No-regret learning for online fair division. NeurIPS 2024.

Discussions