Fair and Efficient Online Allocations: Part I

Alexandros (Alex) Psomas
Purdue University

FAIR DIVISION TEXTBOOK TREATMENT

- INPUT:
  - The resources we are dividing
    - E.g. m indivisible items
  - The agents and their utility structure
    - E.g. n additive agents and a value $v_{i,j}$ for each agent i and item j
  - Constraints on the output (fairness, efficiency, etc)
    - E.g. EF1
- OUTPUT:
  - An allocation of the resources that (approximately) satisfies the constraints

FAIR DIVISION

- Standard real-world motivations:
  - Inheritance, Divorce settlements
  - Housing
  - Dividing land/airspace
  - Computational resources
  - Food donations
  - Kidney exchanges
  - Organ/blood donations

FAIR DIVISION

- Standard real-world motivations:
  - Inheritance, Divorce settlements
  - **Housing**
  - **Dividing land/airspace**
  - **Computational resources**
  - **Food donations**
  - **Kidney exchanges**
  - **Organ/blood donations**

} Not really one-shot problems

DYNAMIC FAIR DIVISION

Static resources, Dynamic Agents	Dynamic resources, Static Agents
Dividing land/airspace **Computational resources** **Housing**	**Food donations** **Blood donations**

Hybrids	**Kidney exchanges** **Organ/blood donations**

DYNAMIC FAIR DIVISION

Static resources, Dynamic Agents	Dynamic resources, Static Agents
Dividing land/airspace **Computational resources** **Housing**	**Food donations** **Blood donations**

This talk

Hybrids	**Kidney exchanges** **Organ/blood donations**

A FIRST PROBLEM

- There are n additive agents
- Indivisible items arrive over time
  - One in each stage for T stages
- Agent i has value $v_{it} \in [0,1]$ for item t that we learn when the item arrives

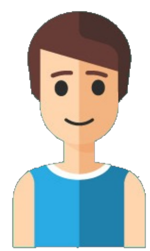

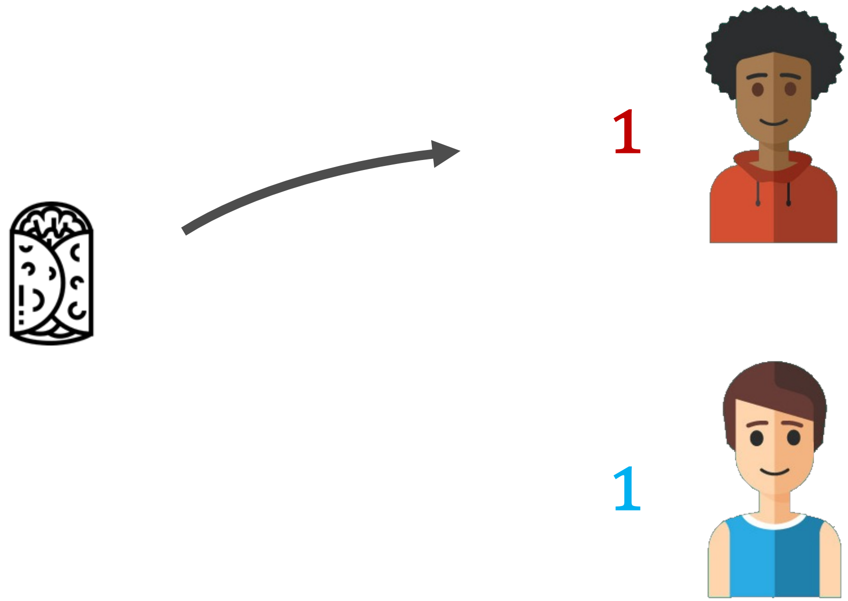
1
1

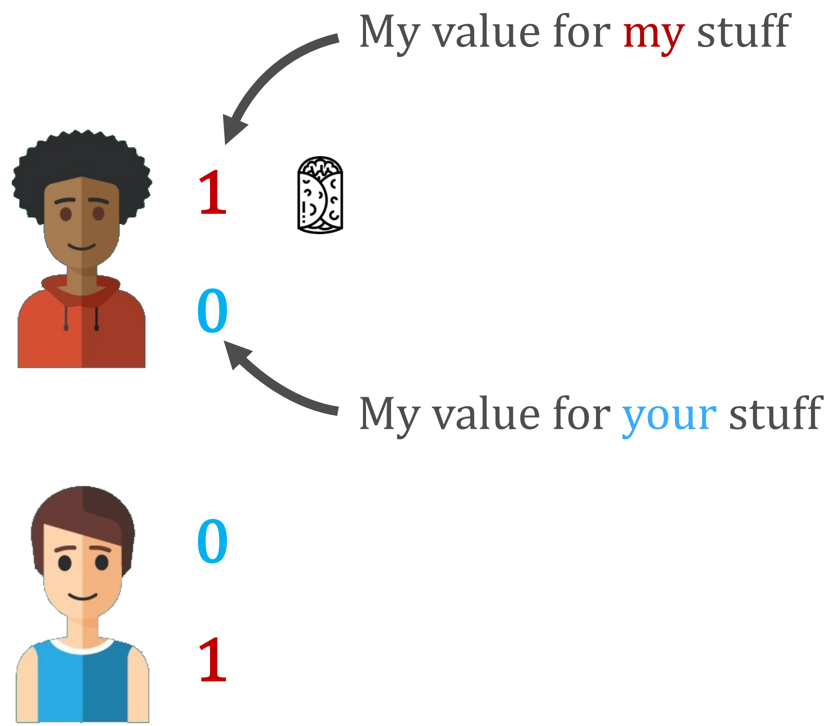
My value for my stuff
1
0
My value for your stuff
0
1

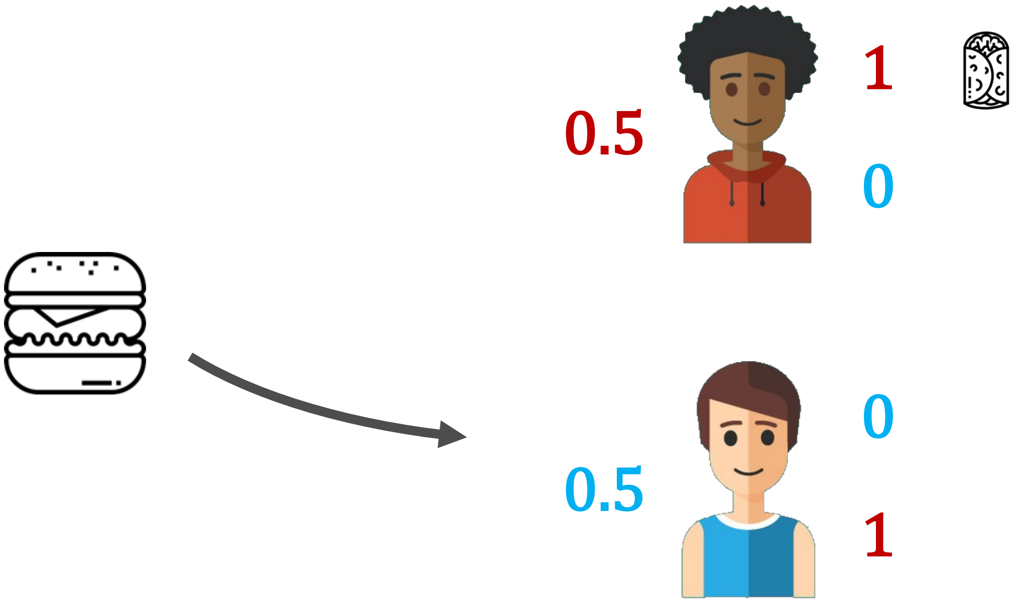
0.5
1
0
0.5
0
1

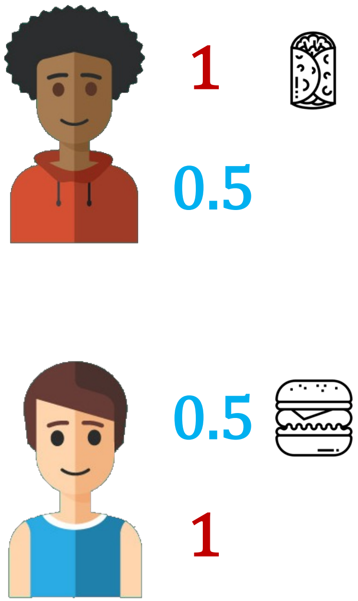
1
0.5
0.5
1

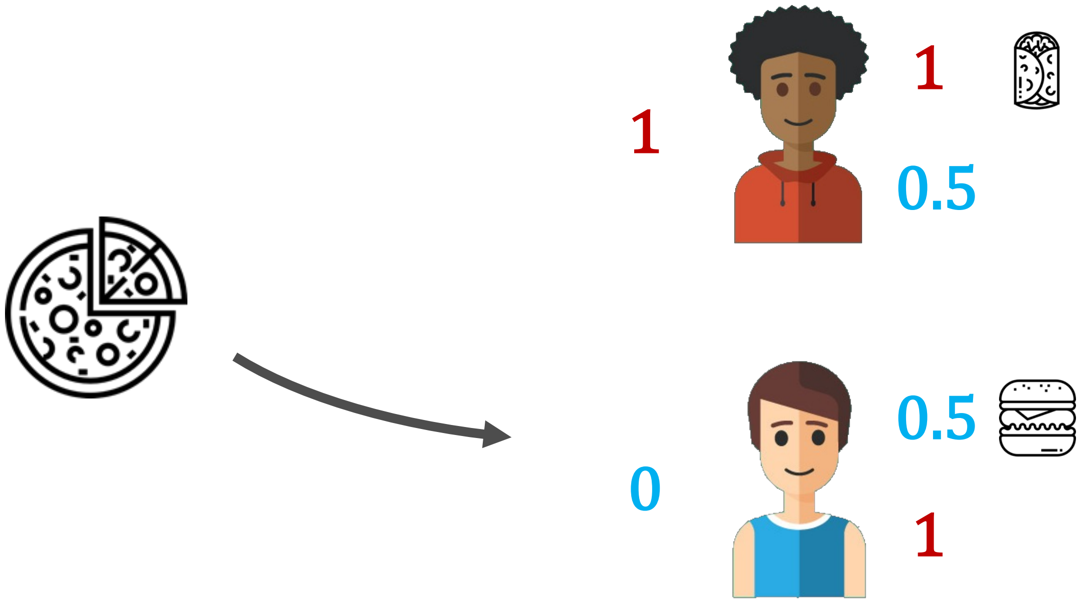
1
1
0.5
0.5
0
1

$$ENVY^{RB} = 1.5 - 1 = 0.5$$

$$ENVY^{BR} = 1 - 0.5 = 0.5$$

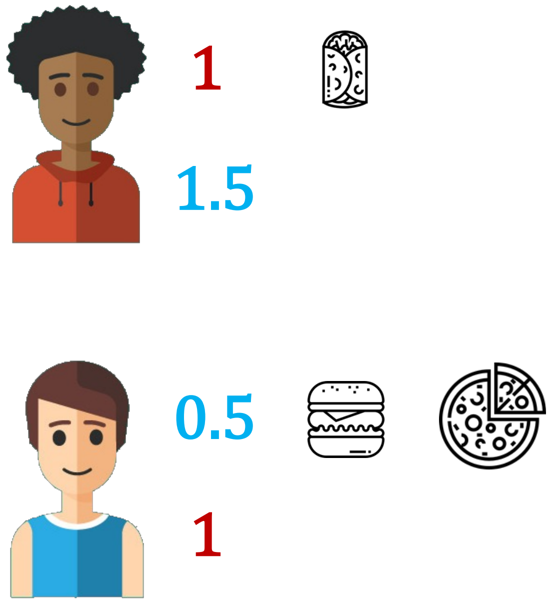

A FIRST PROBLEM

- For the static version, we can keep the maximum envy at most 1 (since $v_{i,t} \in [0,1]$)

- First goal:
  - Minimize the maximum envy at the final step T

A MODELING DECISION

- How is v_{it} generated?
  - Classic online algorithms: adaptive and non-adaptive adversary
  - Bayesian adversaries: values are drawn from a distribution

ADVERSARY MODELS

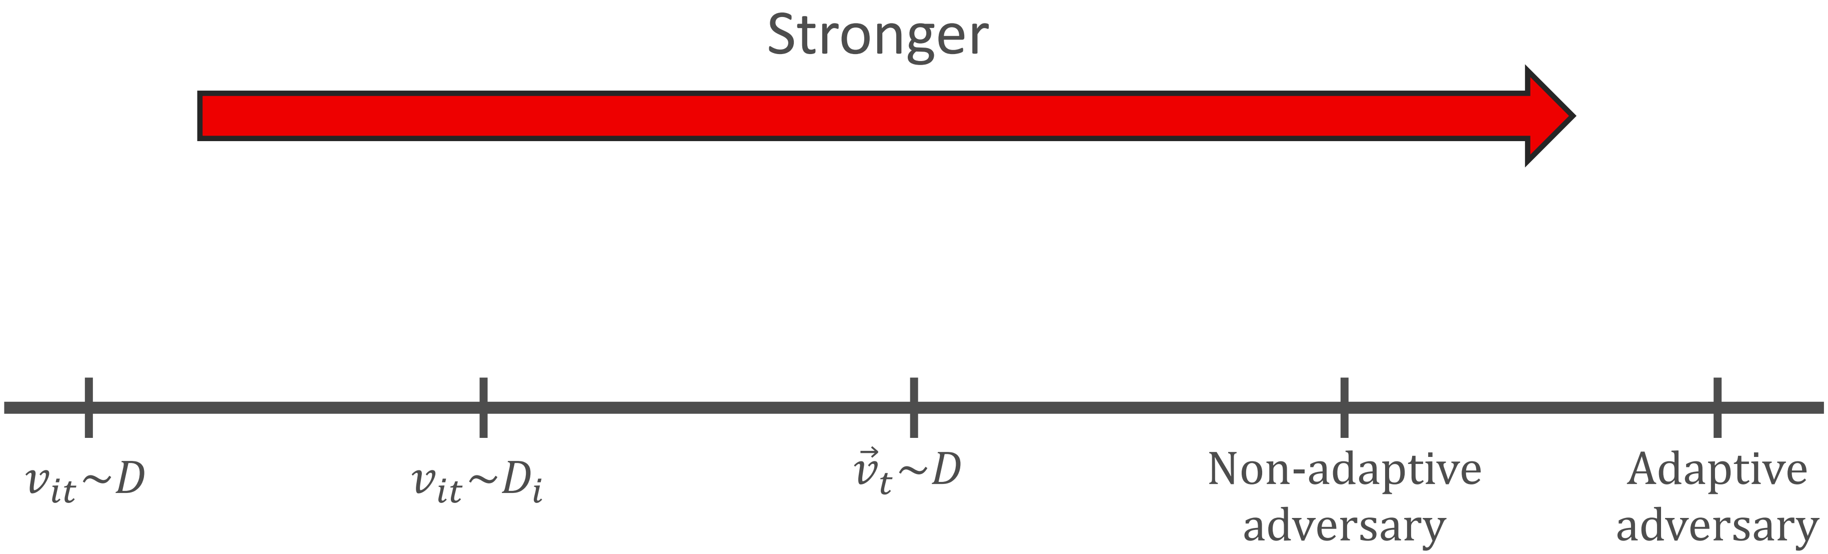

ADVERSARY MODELS

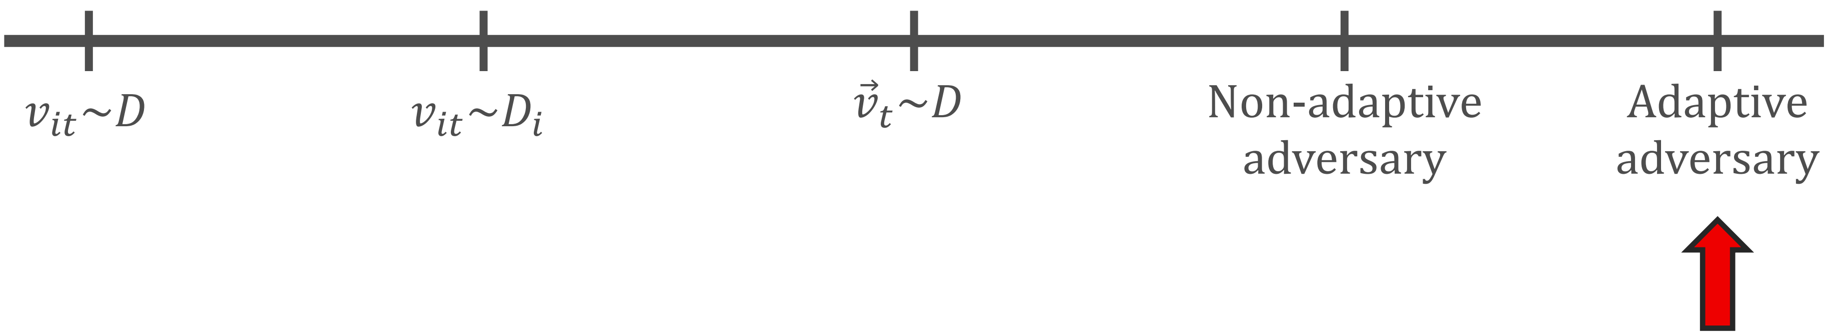

- We write down an algorithm
- The adversary decides the items' values after seeing our code, and the random outcomes of any coin flipping the algorithm does

ADVERSARY MODELS

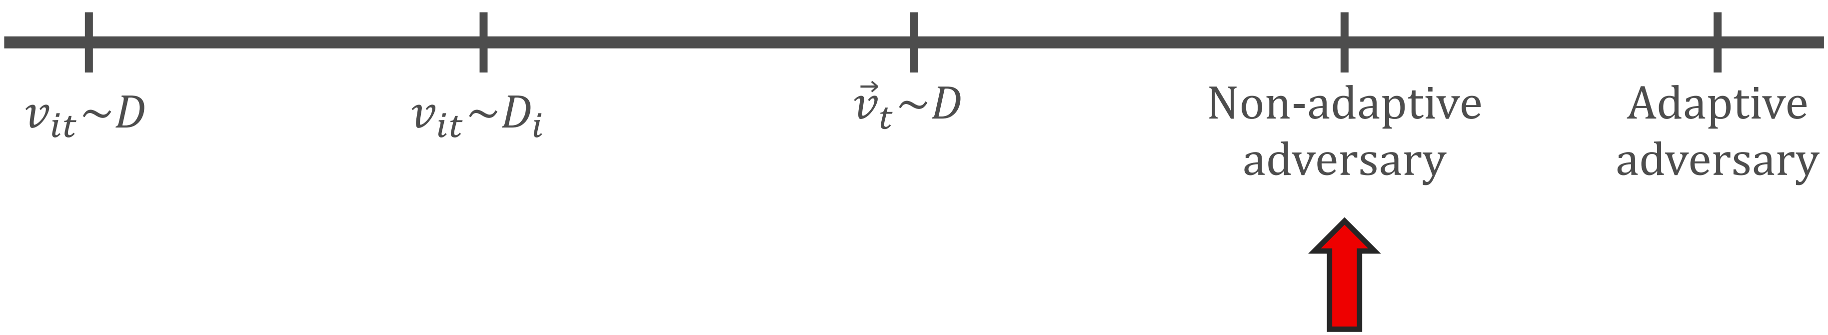

- We write down an algorithm
- The adversary decides the items' values after seeing our code, but **not** the random outcomes of the coin flipping the algorithm does

ADVERSARY MODELS

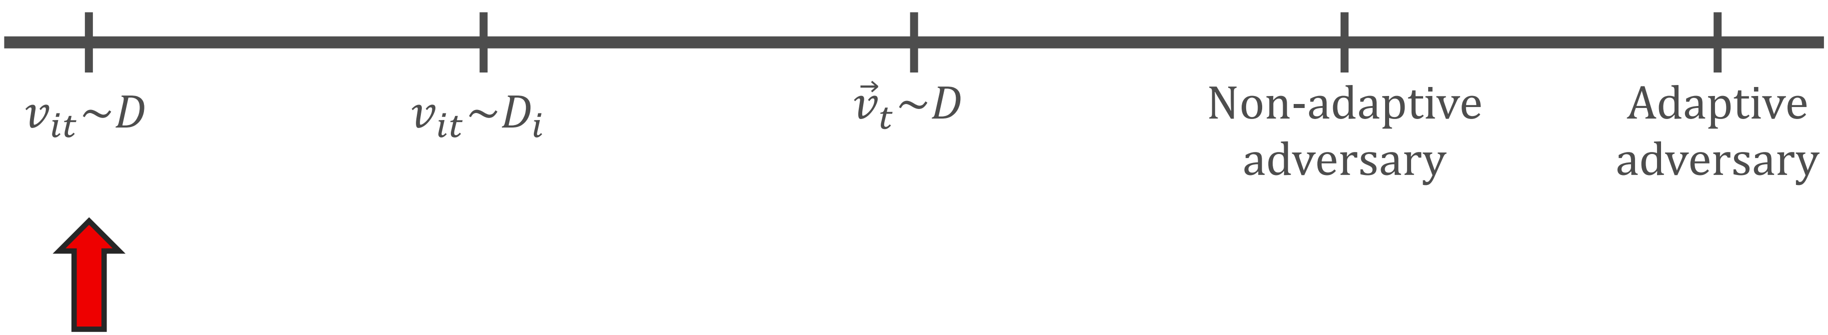

- Items' values are drawn independently and identically from a known distribution D, the same for all agents and all items

ADVERSARY MODELS

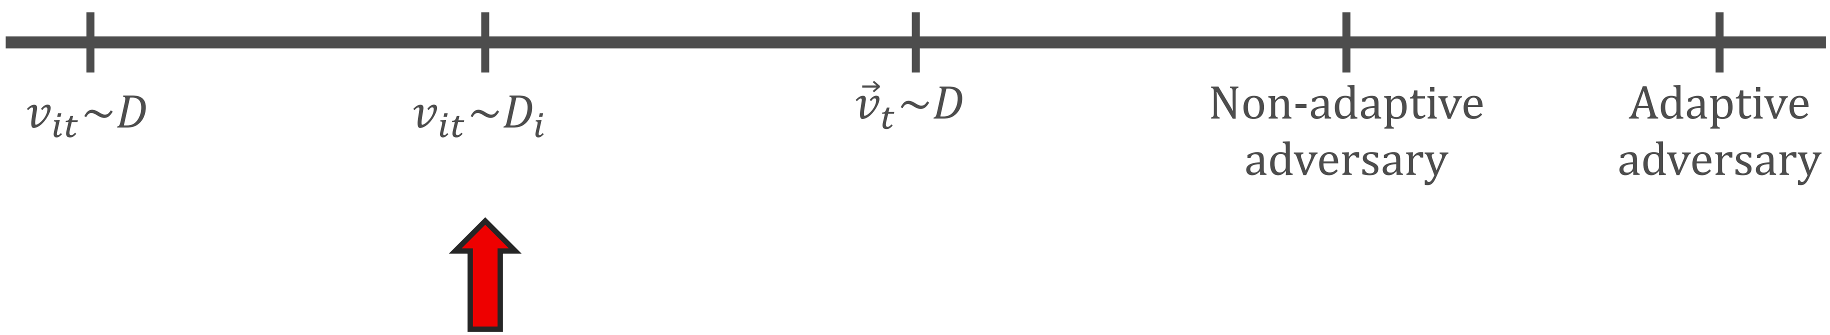

- Agent i's values are drawn independently and identically from a known, agent specific distribution D_i

ADVERSARY MODELS

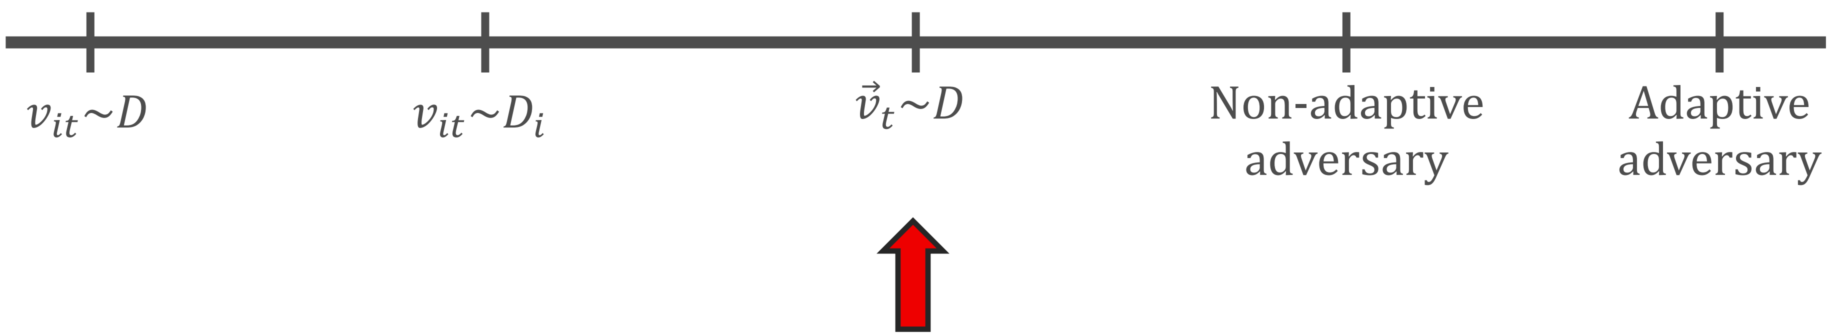

- At each time step t, a vector of values $\vec{v}_t = (v_{1,t}, \dots, v_{n,t})$ is drawn from a known distribution D
- Values can be correlated in a given step (but independent over different time steps)

WHAT TO EXPECT: FAIRNESS

- Linear ($\Theta(T)$) envy is trivial
  - E.g. giving all items to the same agent
- Vanishing envy: $\lim_{T\to\infty} \frac{\mathbb{E}[\max_{i,j} ENVY_T^{ij}]}{T} = 0$
  - $\mathbb{E}\left[\max_{i,j} ENVY_T^{ij}\right] \in o(T)$
- Envy free up to one item (EF1) with probability 1
- Envy free with high probability

WHAT TO EXPECT: EFFICIENCY

- Pareto efficiency:
  - An allocation is Pareto efficient if there is no allocation where all agents get more utility (with at least one agent getting strictly more utility)
- α-Pareto efficiency (Ruhe and Fruhwirth, 1990):
  - An allocation is α-Pareto efficient if no allocation improves the utility of all agents by a factor of $1/\alpha$
    - E.g., a dictatorship is $\frac{1}{n} + \epsilon$ Pareto efficient

ADAPTIVE ADVERSARY

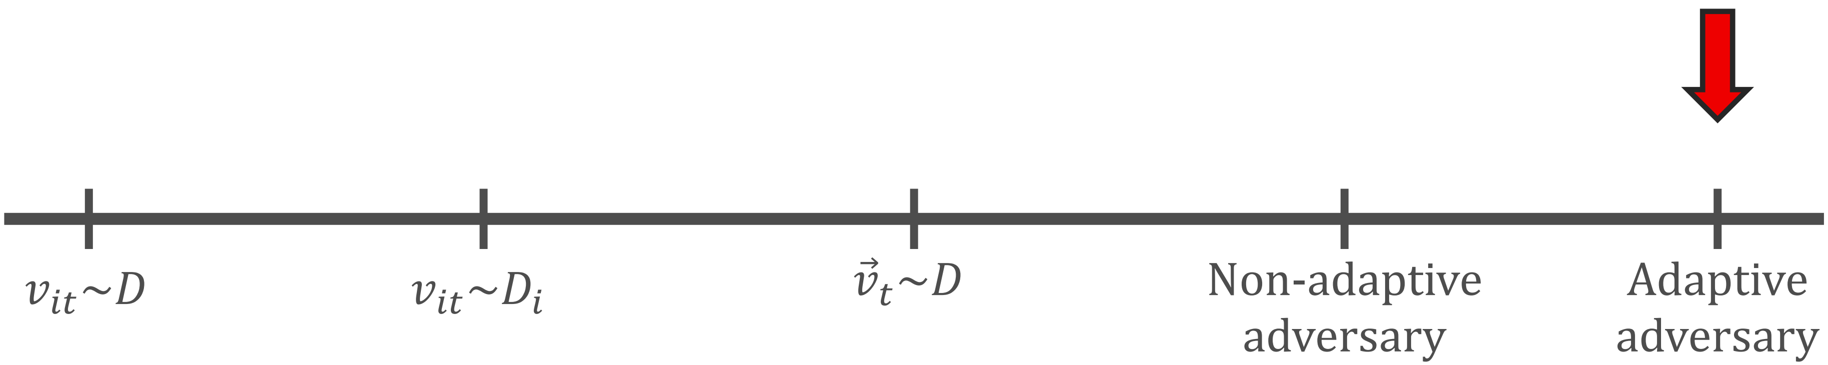

Algorithm : Random allocation

Fairness : $\mathbb{E}\left[\max_{i,j} Envy_{i,j}^{T}\right] \in \tilde{O}(\sqrt{T/n})$ [BKPP 2018]

Efficiency : $\frac{1}{n}$-Pareto efficient ex-ante

ADAPTIVE ADVERSARY

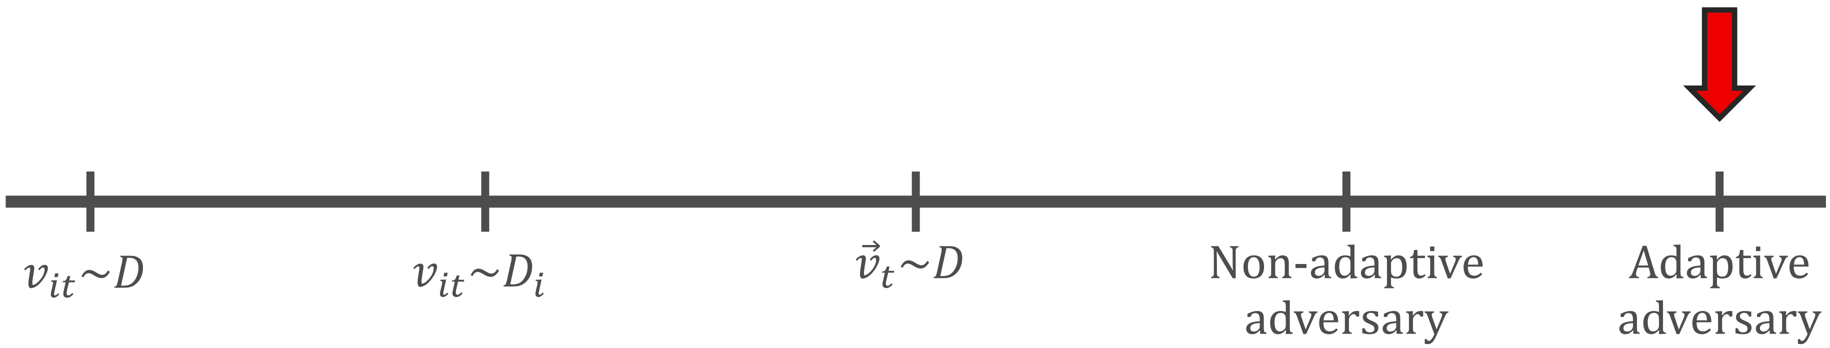

<u>Theorem</u>[BKPP 18]: An adaptive adversary can always ensure $\max_{i,j} \text{ENVY}_T^{ij} \in \Omega\left(\sqrt{T/n}\right)$.

- Thus, random allocation is asymptotically optimal!
- Good news: We can get the same guarantee with a deterministic algorithm!
  - Define a potential function $\phi(t)$
  - Allocate in a way that $\phi(t)$ is minimized
- *Question*: Can we improve the efficiency guarantee while maintaining optimal fairness?

WHAT ABOUT EFFICIENCY?

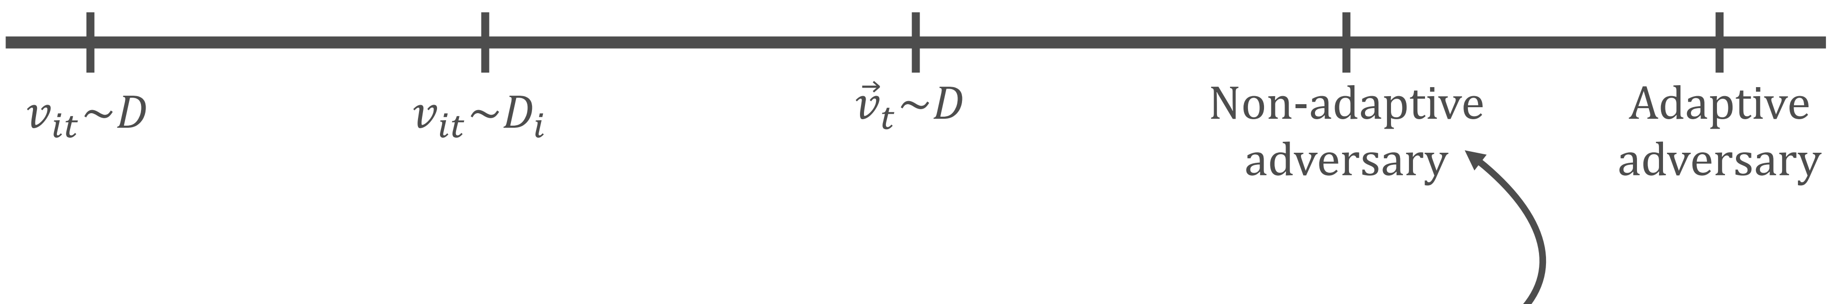

Theorem [ZP 20]: There is no algorithm that guarantees vanishing envy ($T^{1-\epsilon}$) and is $\left(\frac{1}{n} + \varepsilon\right)$-Pareto efficient for any $\varepsilon > 0$

PROOF SKETCH OF THEOREM FOR ADAPTIVE ADVERSARY

- Assume algorithm A had envy $f(T) \in o(T)$ on all inputs, and was $\frac{1}{n} + \epsilon$ Pareto efficient.
- Instance I^*: each agent i has values
  - $v_{i,t} = 1$ for the T/n-th segment
    - Items t $\in \left[\frac{T}{n}(i-1)+1, \dots, \frac{T}{n} i\right]$
  - $v_{i,t} = \epsilon$ for all other items
- An adaptive adversary can always stop showing I^* and make all remaining items worthless
- Therefore envy at step t must be at most $f(T)$ for all agents
- This implies that in each segment i, every agent must get $\frac{T}{n^2} \pm \frac{f(T)}{\epsilon}\left(1+\frac{2}{\epsilon}\right)^{i-1}$ items
- Thus, final utility for each agent is at most $\left(\frac{T}{n^2} + \frac{f(T)}{\epsilon}\left(1+\frac{2}{\epsilon}\right)^{n-1}\right) \cdot (1 + (n-1)\epsilon)$
- But, it is possible to give all agents utility T/n ■

PROOF SKETCH OF THEOREM FOR NON-ADAPTIVE ADVERSARY

- The non-adversary has n instances in their arsenal
- I_i's first $\frac{T}{n}i$ items follow I^*, and the rest have zero value
- Again, we bound the number of items the algorithm can allocate to each agent in each segment
- The new bound is looser and probabilistic, but gets the job done

WHAT ABOUT EFFICIENCY?

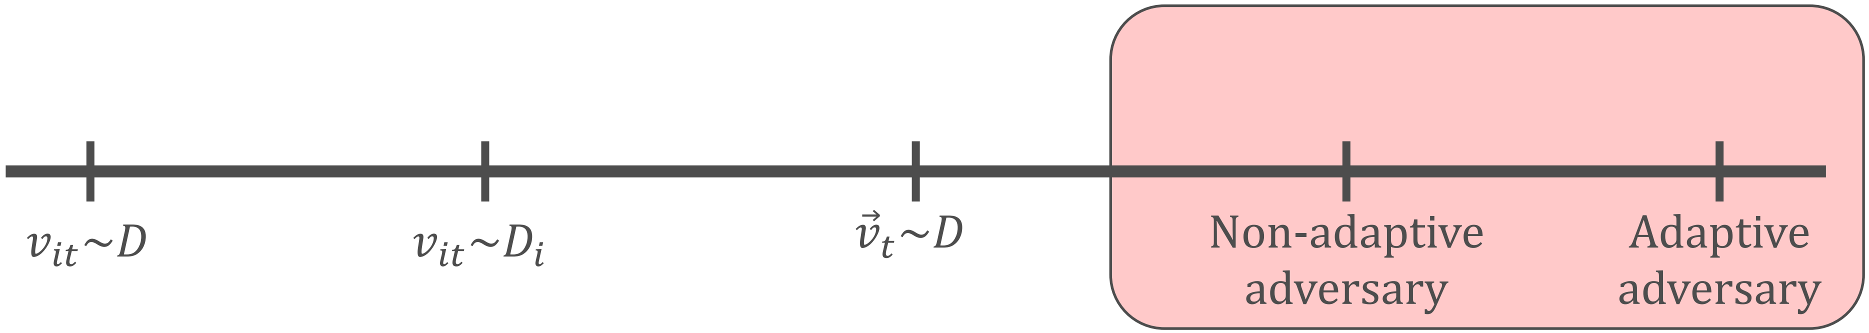

INDEPENDENT AND IDENTICAL DISTRIBUTION

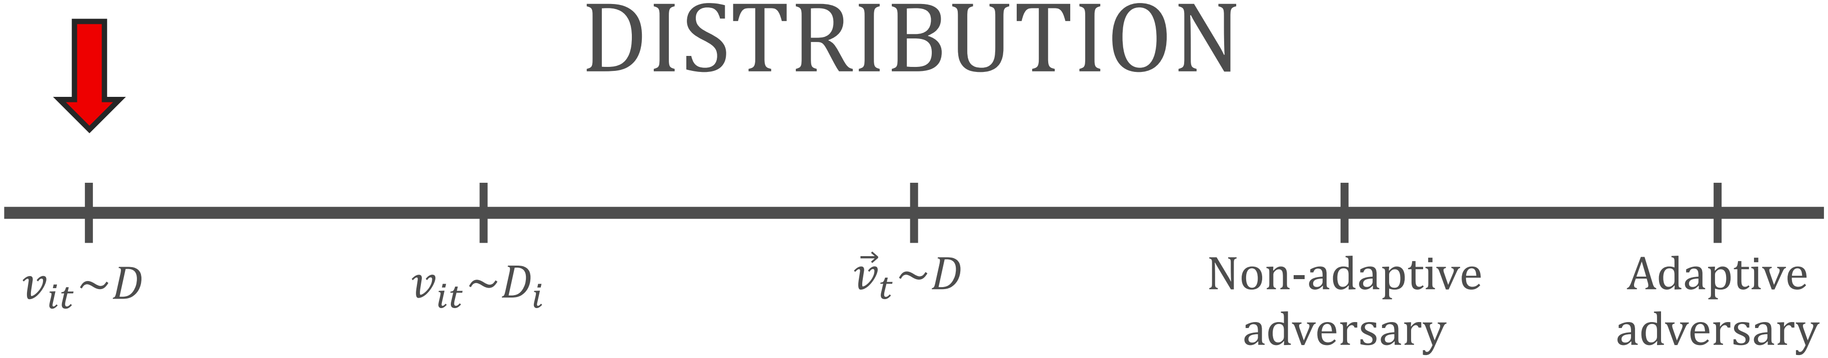

Algorithm : Give each item to the agent with the highest valuation

Guarantees (under mild conditions on D) [KPW, AAAI 16] :

- Pareto efficient (ex-post)
- Envy-freeness with high probability

GREEDY ALGORITHM

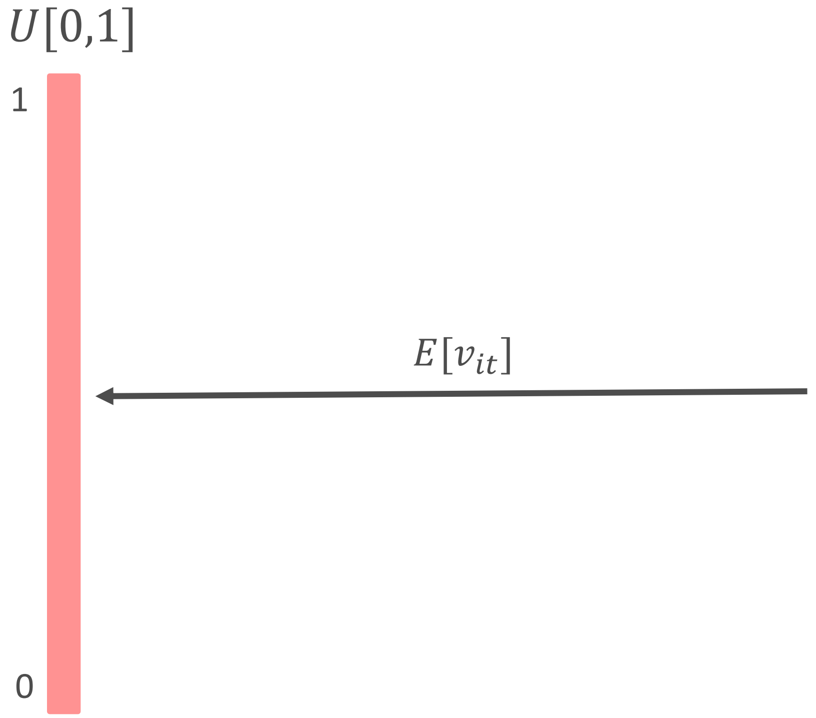

GREEDY ALGORITHM

- Everyone roughly receives the same number of items
    - But when i receives an item, it is more valuable
    - Chernoff Bounds
- Can replace $U[0,1]$ with any distribution with constant variance

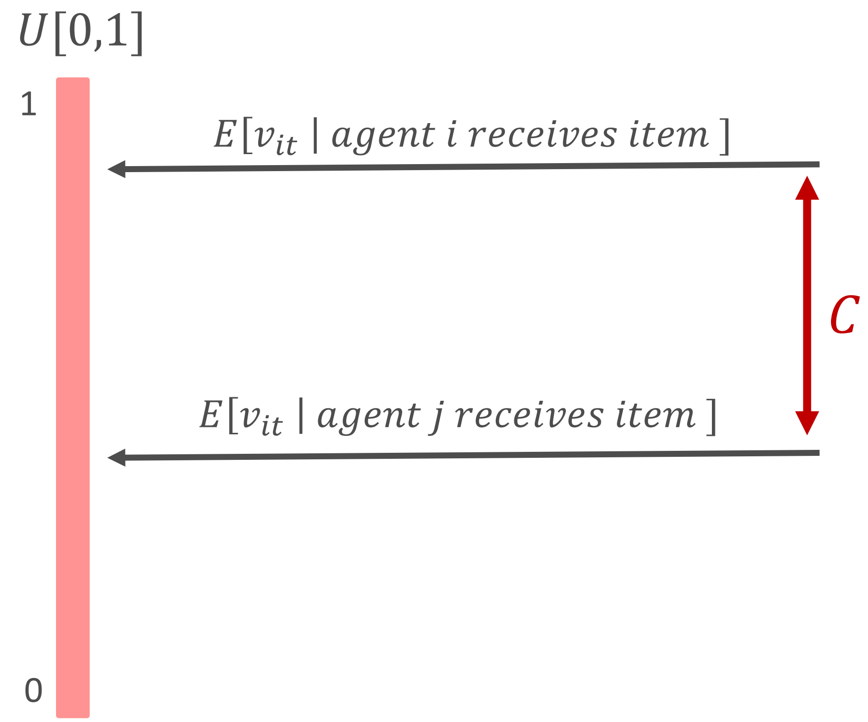

EXTENDING THE GREEDY ALGORITHM

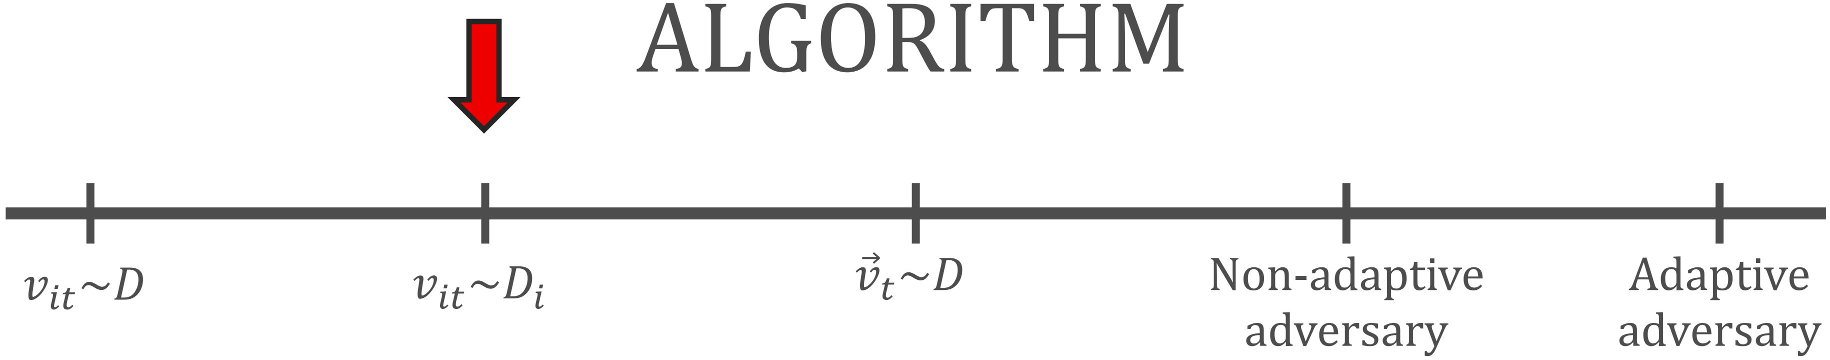

Proposed Algorithm :

Give each item to the agent with the highest quantile??

- $U(0,1)$ *and* $U(0.49, 0.51)$
  - Agent 2 essentially only cares about the **number** of items
- This algorithm is envy-free whp, but **not efficient**

EXTENDING THE GREEDY ALGORITHM

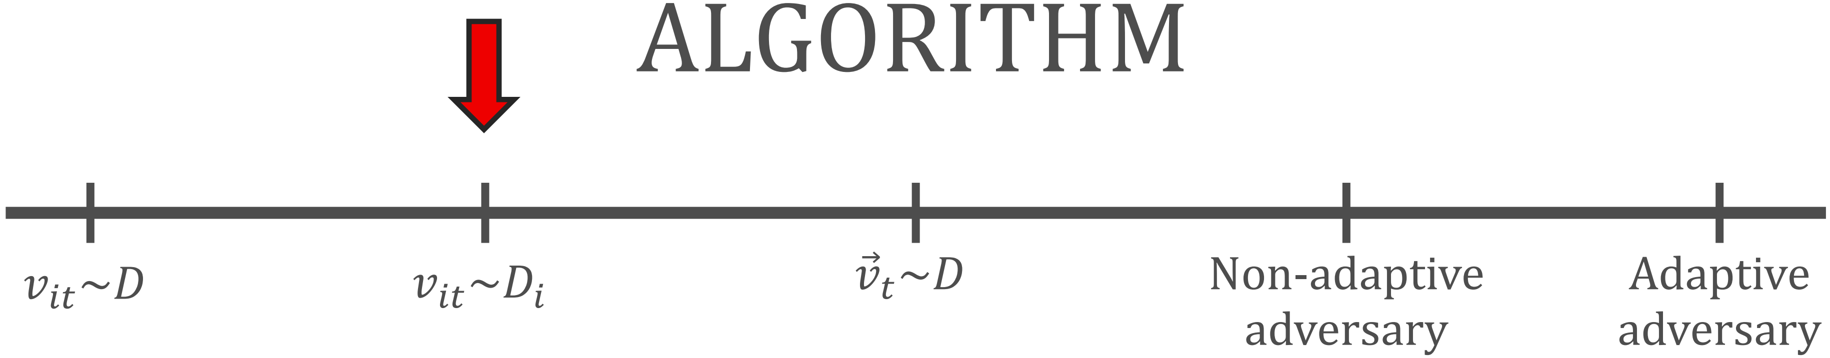

Algorithm [Bai, Gölz 2022]:

- Find β_i for each agent i, such that $\Pr\left[\beta_i v_{i,t} = argmax_j \beta_j v_{j,t}\right] = 1/n$
  - i.e. Allocating to $argmax_j \beta_j v_{j,t}$ gives i the next item with probability $1/n$

Properties:

- *Envy free with high probability*
- *Pareto optimal (since it maximizes weighted welfare)*

CORRELATED DISTRIBUTIONS

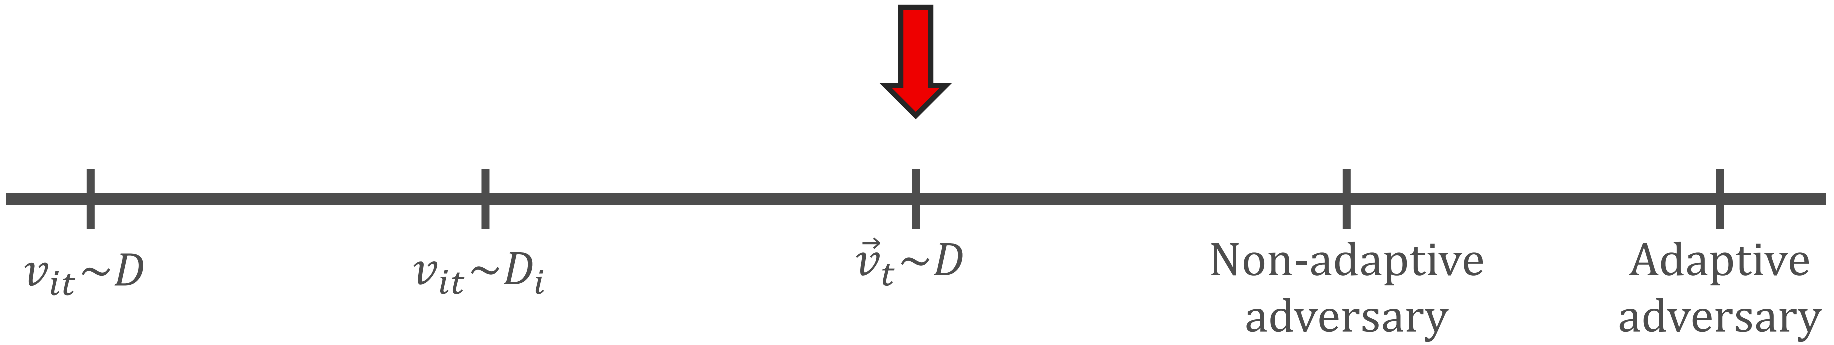

Theorem [ZP 20]: There is an ex-post Pareto optimal algorithm that guarantees to each pair of agents i, j:

- Either i does not envy j with high probability
- Or, i envies j by at most one item (with probability 1)

CORRELATED DISTRIBUTIONS

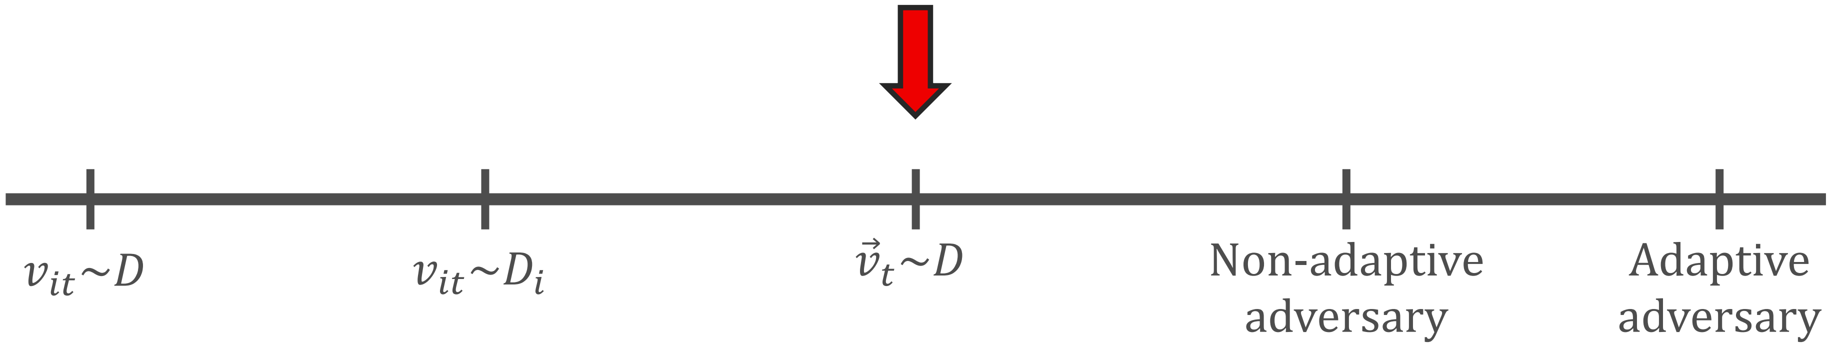

<u>Theorem</u> [ZP 20]: There is an ex-post Pareto optimal algorithm that guarantees to each pair of agents i, j:

- Either i does not envy j with high probability
- Or, i envies j by at most one item (with probability 1)

<u>Main structural result:</u>

Given n agents and m items, there is a Pareto efficient fractional allocation such that each agent i:

- Either strictly prefers her own bundle to the bundle of agent j
- Or i and j have identical allocations and the same value for all the items that are allocated to them

INTUITION

- How could you ever be ex-post Pareto?

INTUITION

- How could you ever be ex-post Pareto?

	2	1
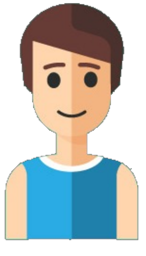	1	2
	w.p. 0.1	w.p. 0.9

INTUITION

- How could you ever be ex-post Pareto?
- Idea 1: every time item 1 comes, give it to red agent, o.w. give item to blue agent
  - Efficient, but not fair!

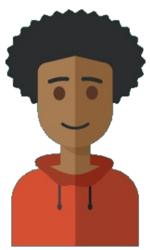

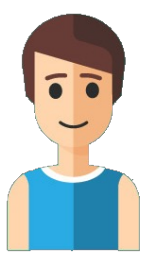

2	1
1	2
w.p. 0.1	w.p. 0.9

INTUITION

- How could you ever be ex-post Pareto?
- Issue: if you **ever** allocate item 2 to the red agent, you **cannot** allocate item 1 to the blue agent

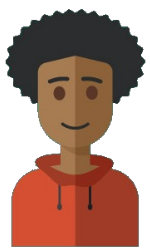	2	1
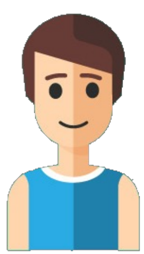	1	2
	w.p. 0.1	w.p. 0.9

INTUITION

- How could you ever be ex-post Pareto?
- Insight: we should Pareto efficient and fair in the instance where values are multiplied by probabilities

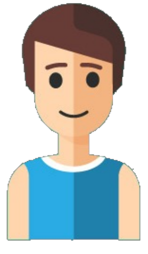

0.2	0.9
0.1	1.8

BLUEPRINT

- Construct this **static** instance I from the correlated distribution
- Find a fractional allocation x
- For the online problem, every time item k comes, allocate to agent i with probability x_{ik}

ALGORITHM

- Fact 1: Being Pareto efficient in I turns out to be enough for Pareto efficiency ex-post for the online problem!
- Fact 2: Being envy-free in I will give vanishing envy
- Question: Can we do better?

ALGORITHM

- The dream: Pareto efficiency and **strong** envy-freeness for *I*
  - Then, Chernoff would give EF w.h.p.
- Pretty much impossible
  - Agents could be identical…
- **CISEF:**
  - Either agent i strictly prefers her own bundle to the bundle of agent j
  - Or i and j have identical allocations and the same value (up to a scaling factor) for all the items that are allocated to either of them
- How?
  - Start from CEEI, and try to create strong-envy, without messing up efficiency

TAKE AWAYS

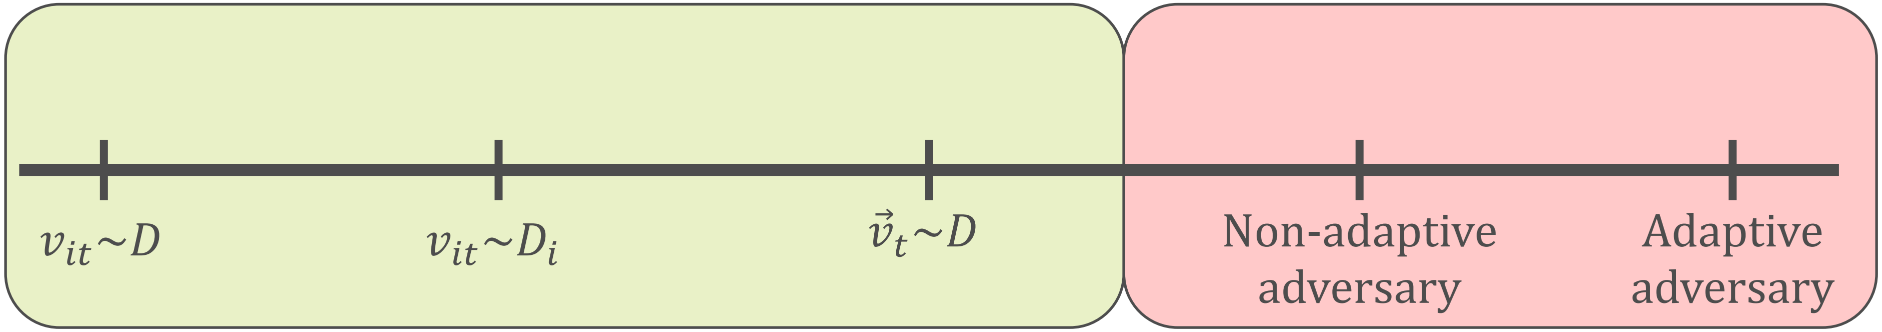
$v_{it} \sim D$
$v_{it} \sim D_i$
$\vec{v}_t \sim D$
Non-adaptive adversary
Adaptive adversary

TAKE AWAYS

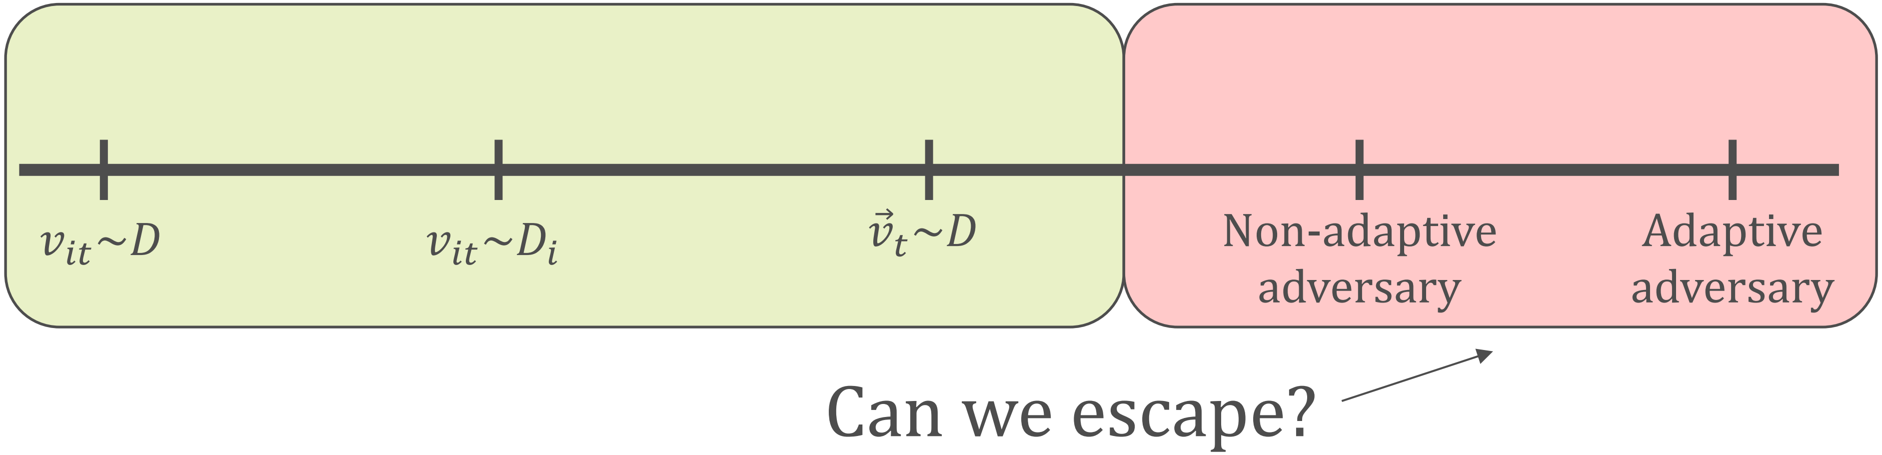
$v_{it} \sim D$
$v_{it} \sim D_i$
$\vec{v}_t \sim D$
Non-adaptive adversary
Adaptive adversary
Can we escape?

TAKE AWAYS

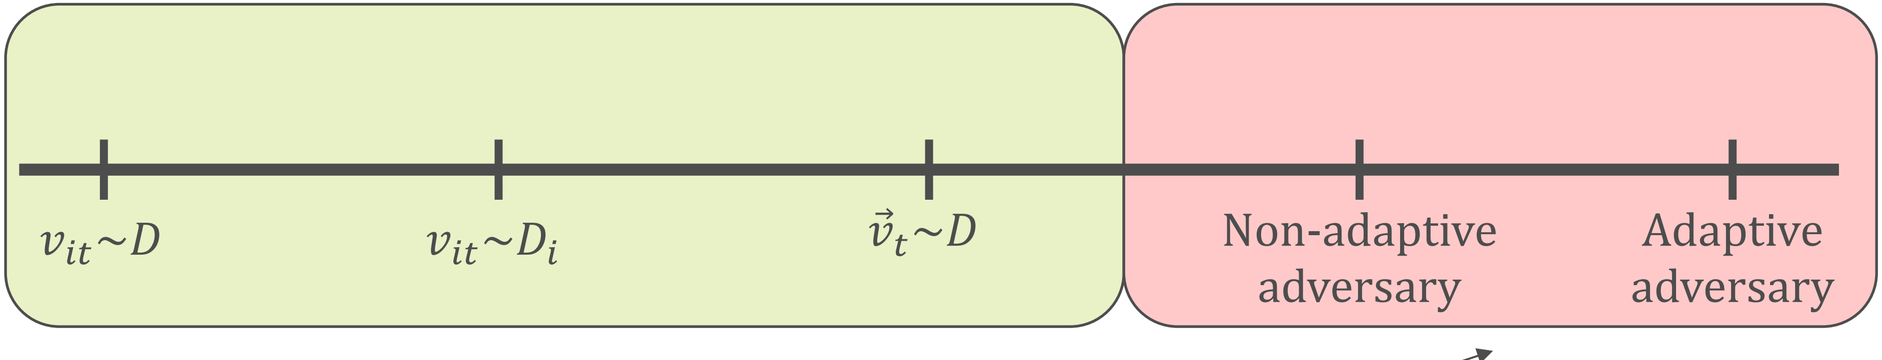

Can we escape?
See second half of tutorial.

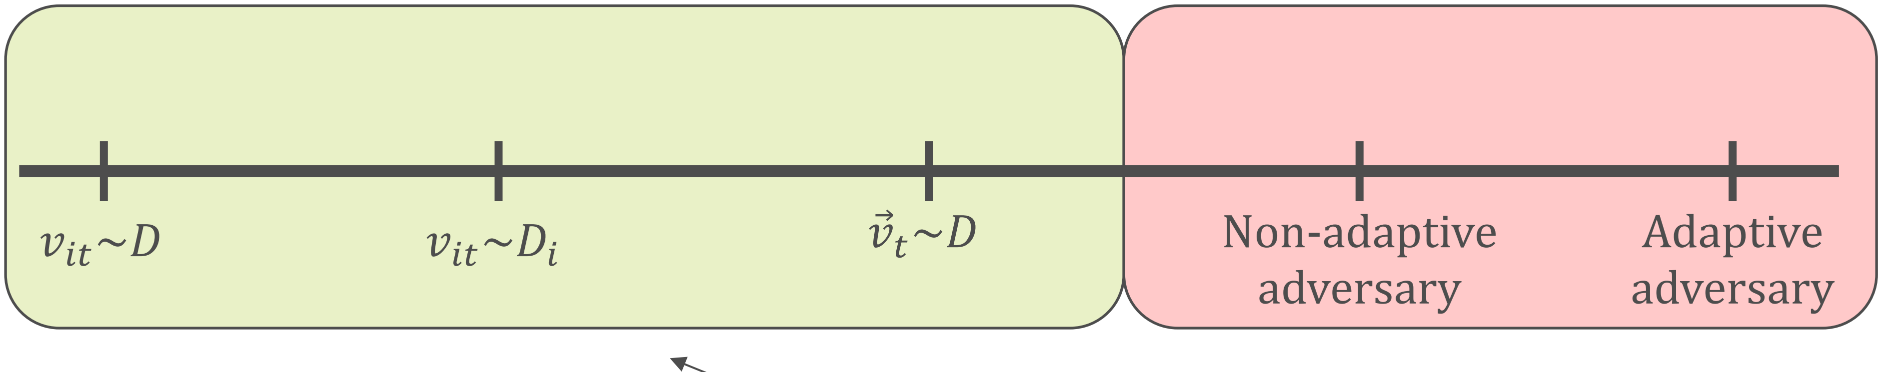

Can we get even more?

OPEN

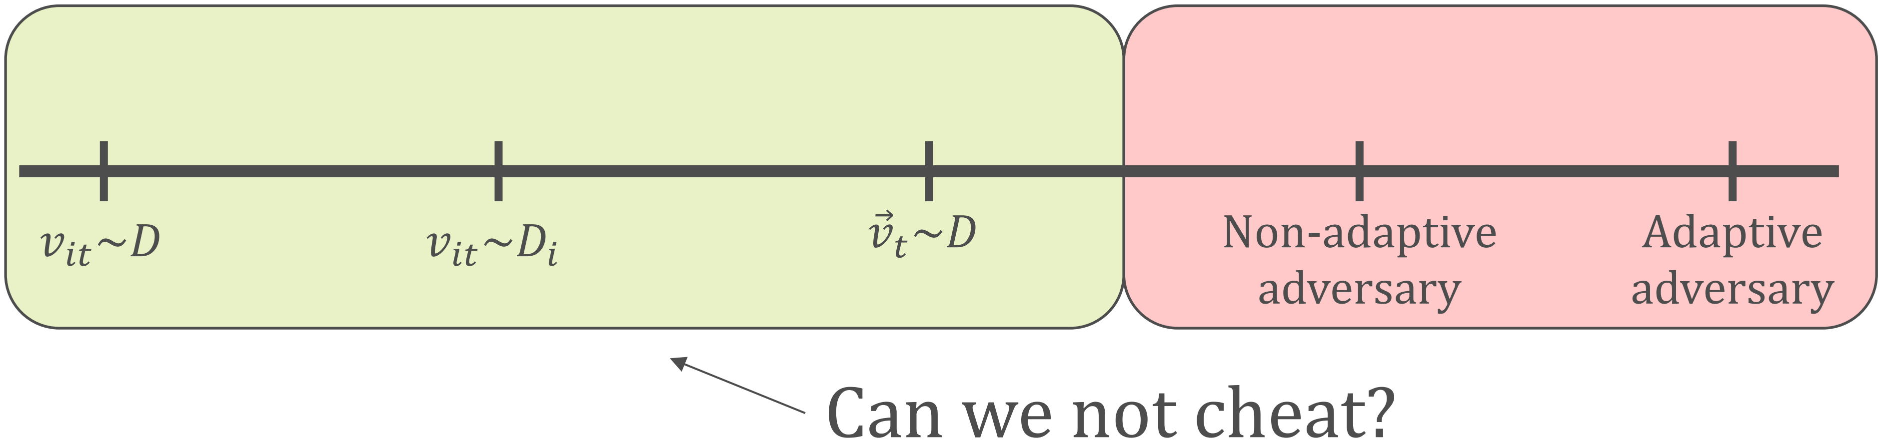

$v_{it} \sim D$
$v_{it} \sim D_i$
$\vec{v}_t \sim D$
Non-adaptive adversary
Adaptive adversary
Can we not cheat?

OPEN

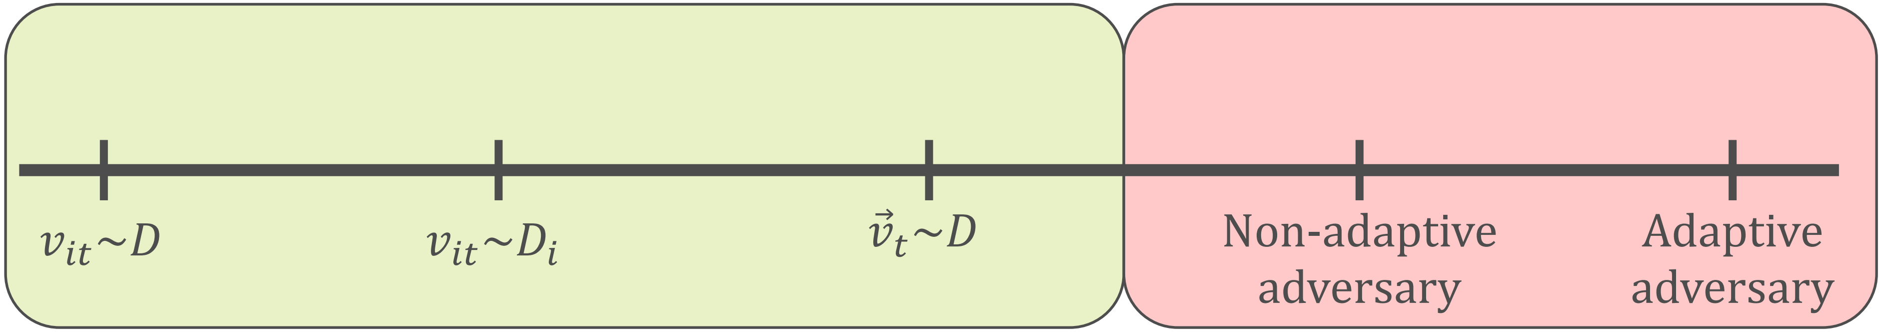

Can we not cheat?

- What if the adversary distribution can depend on T?
- Theorem [Bansal et al. 2020]: $O(\log T)$ envy w.h.p for two agents, against the correlated distribution adversary.

WHAT I DIDN'T TALK ABOUT

- Dynamic resources & static agents & incentives!
  - See references at the end of slides for a biased selection of papers
  - Highlights:
    - Infinite horizons, so more tricks available
    - Artificial currencies

REFERENCES

- How to Make Envy Vanish Over Time. Benade, Kazachkov, Procaccia, Psomas. EC 2018
- Fairness-Efficiency Tradeoffs in Dynamic Fair Division. Zeng, Psomas. EC 2020
- From monetary to nonmonetary mechanism design via artificial currencies. Gorokh, Banerjee and Iyer. MOR 2021
- Multiagent mechanism design without money. Balseiro, Gurkan, and Sun. OR 2019
- Dynamic mechanisms without money. Guo and Horner. 2009
- Competitive repeated allocation without payments. Guo, Conitzer, and Reeves. WINE 2019