Private Pareto Optimal Exchange

By: Sampath Kannan, Jamie Morgenstern, **Ryan Rogers**, Aaron Roth

Exchange Market

Exchange Market

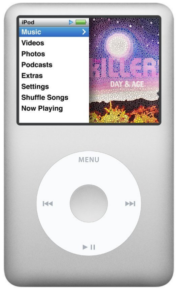

Exchange Market

- **Instance:** We have n players that each bring a good from a set G and each player i has preferences $\succ_i$
  - i.e. player i prefers good a to b iff $a \succ_i b$
- **Goal:** We want to compute an allocation of the goods to players that is
  - Individually Rational – players get a good no worse than what they start with
  - Pareto Optimal – if a player gets a more preferred good , then some player must be worse off.

Top Trading Cycles

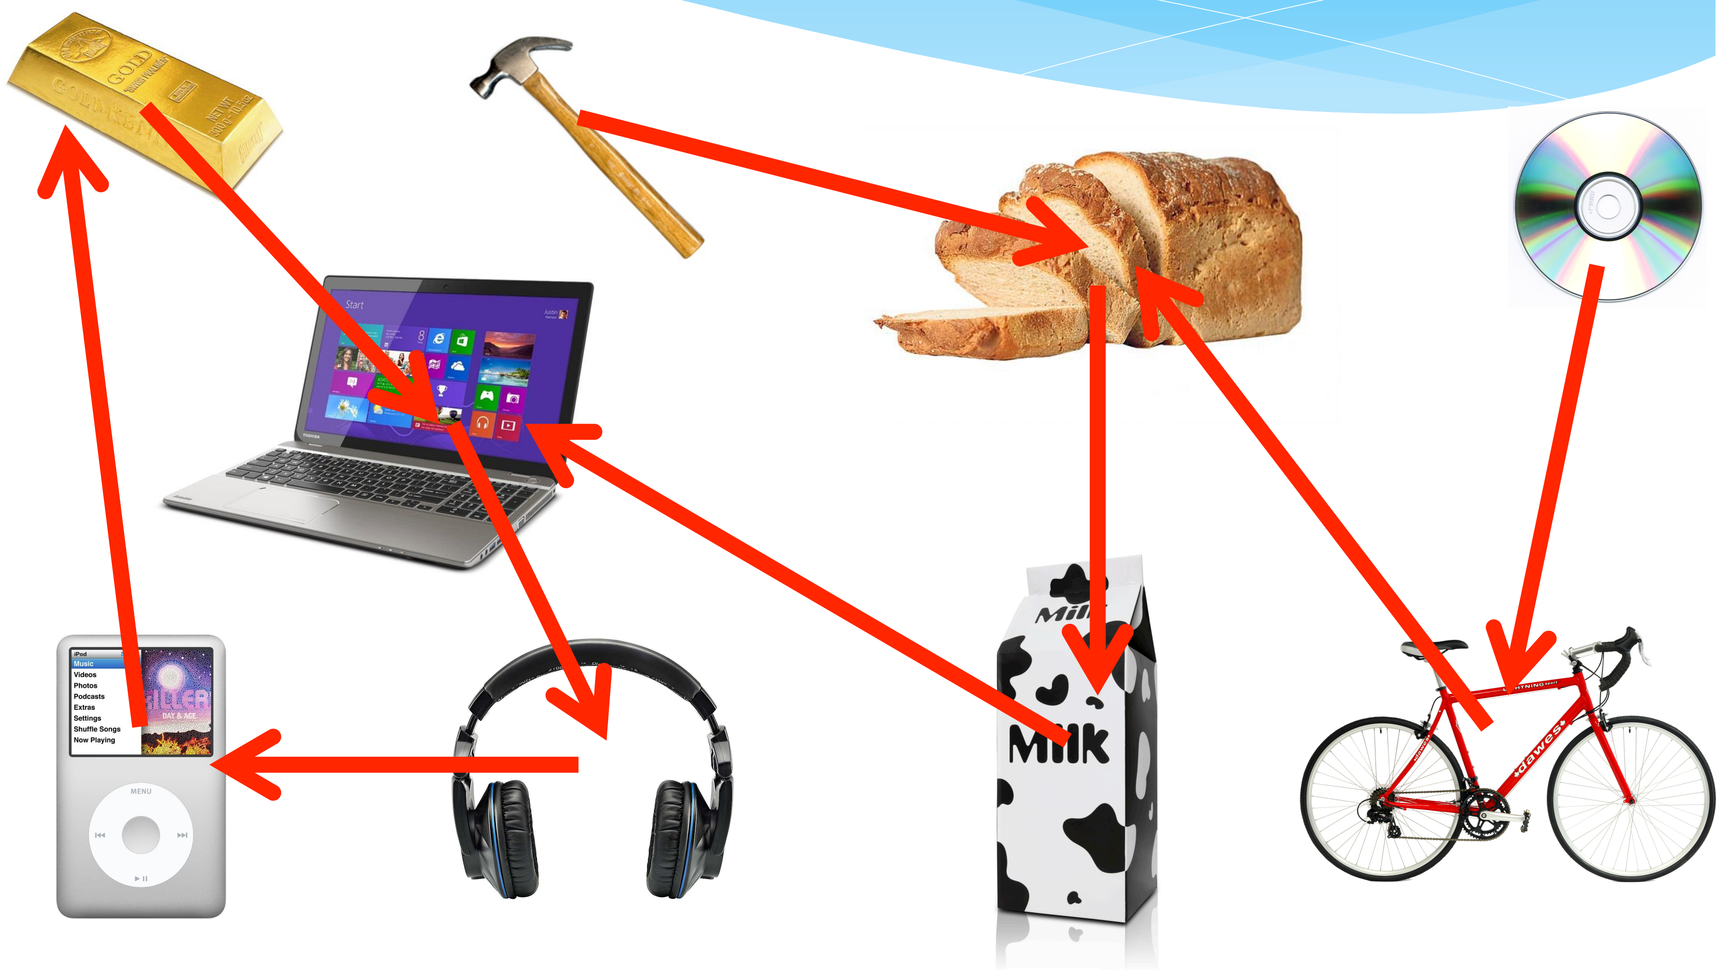

Top Trading Cycles

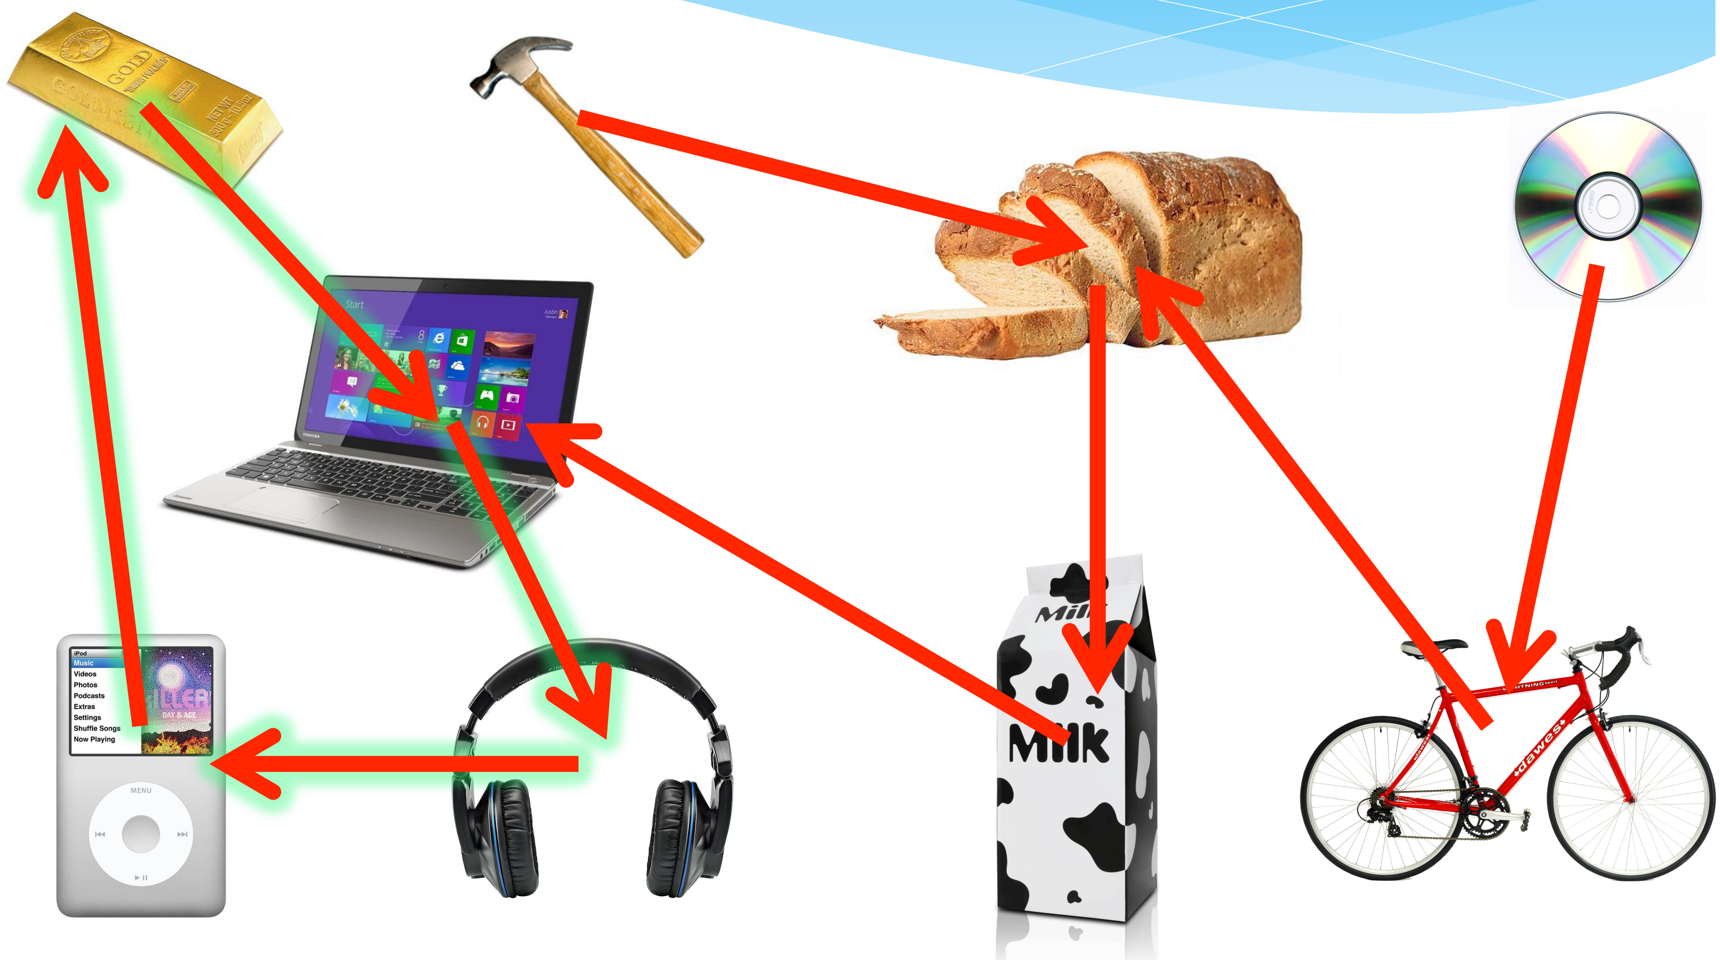

Top Trading Cycles

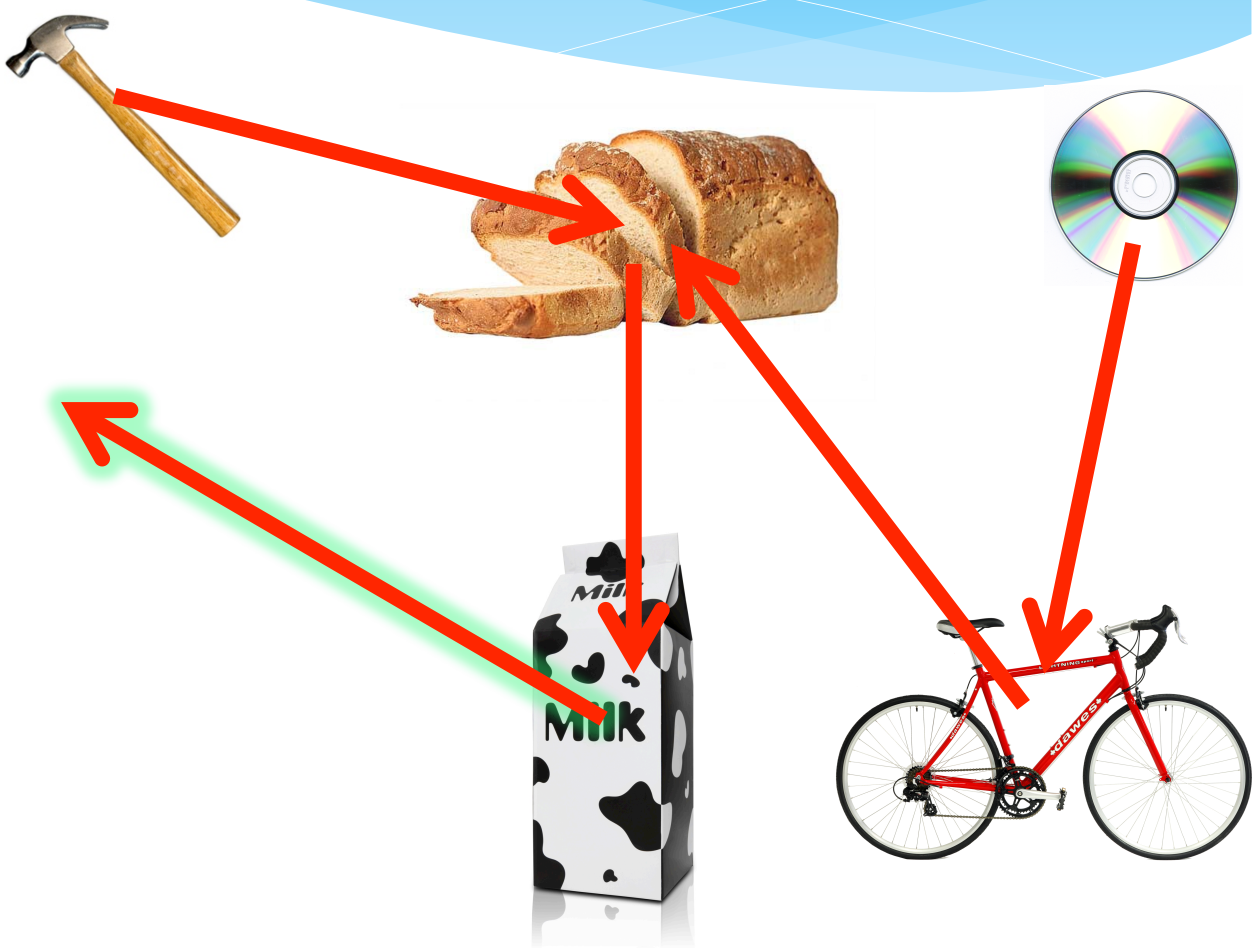

Top Trading Cycles

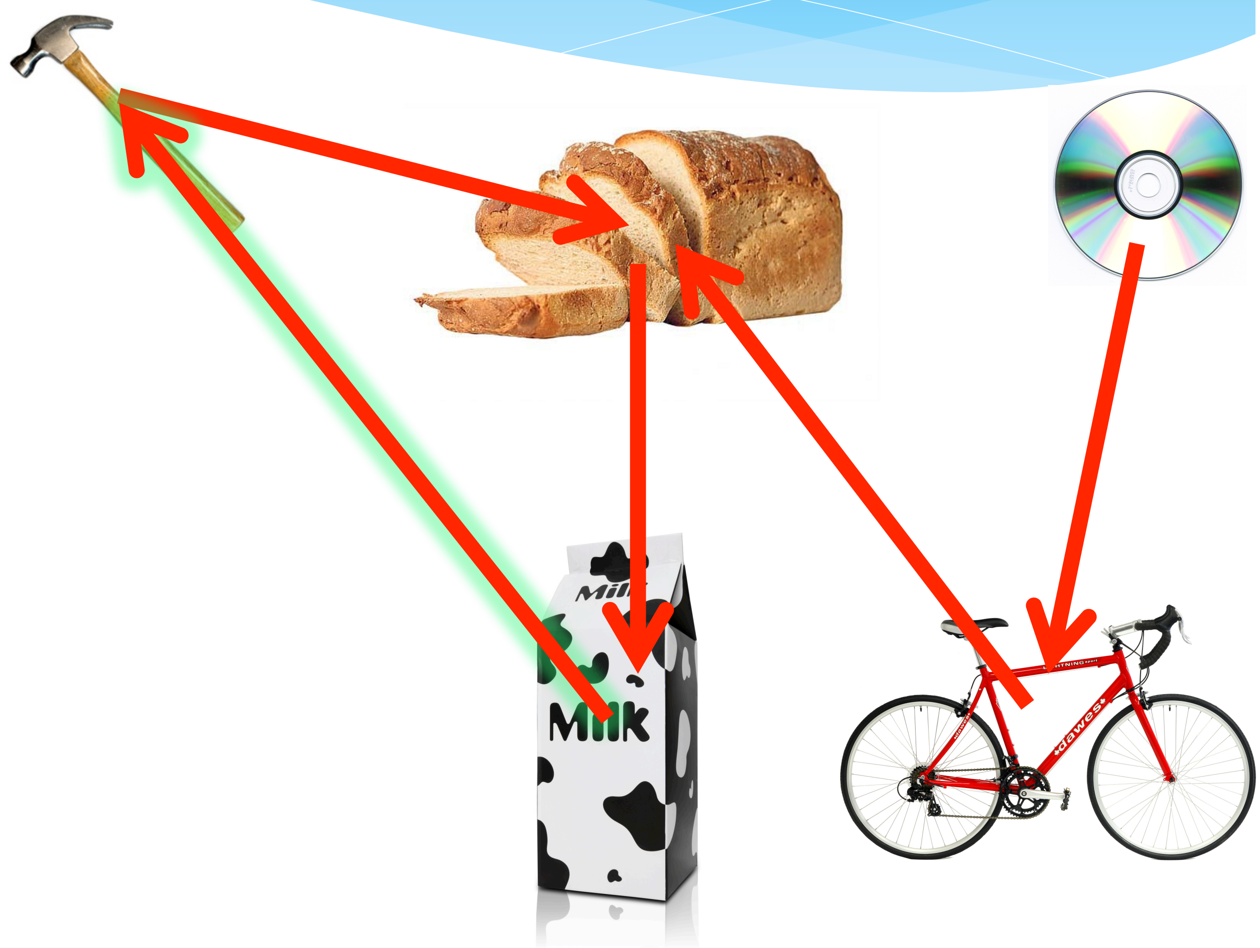

Top Trading Cycles

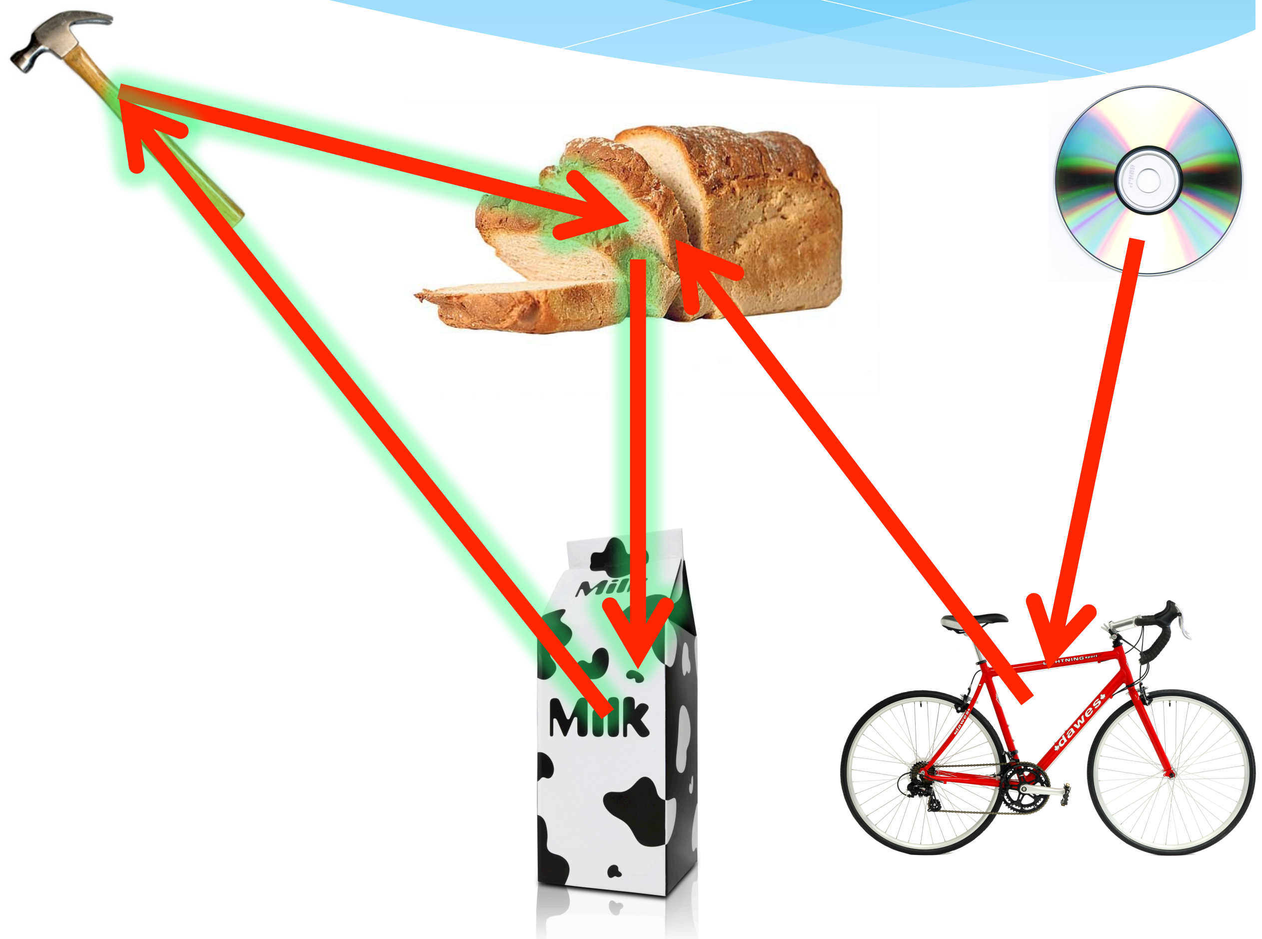

Top Trading Cycles

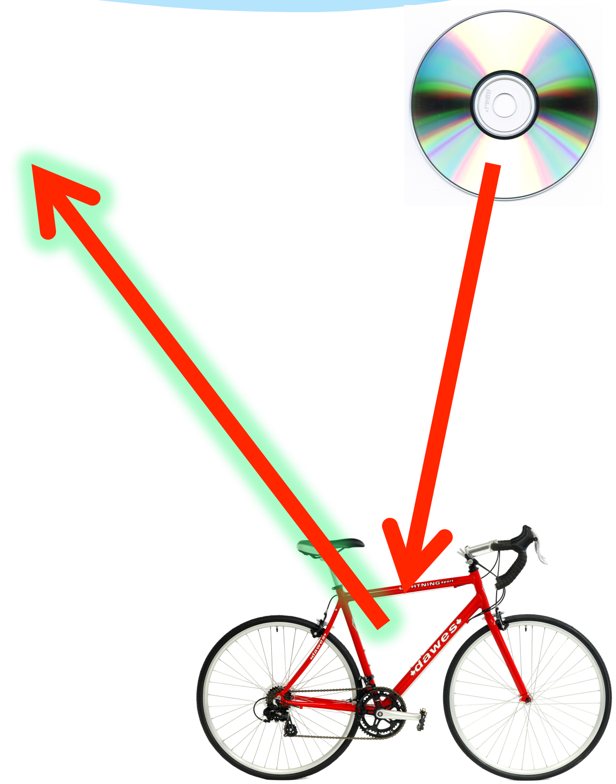

Top Trading Cycles

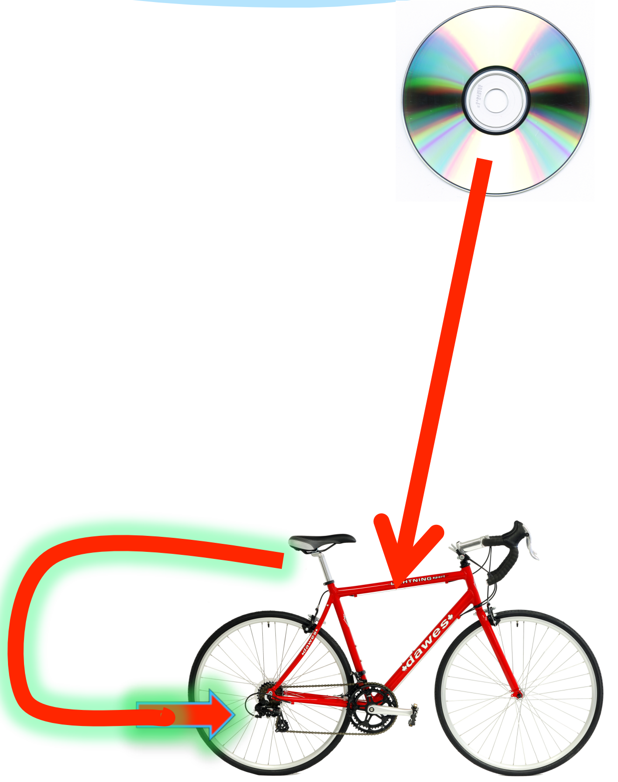

Top Trading Cycles

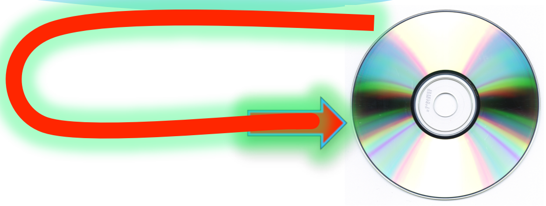

Model

- **Instance:** each player is endowed with a good that comes from a set of *k* types of goods and preferences are over the *k* types, which is sensitive information.
- **New Goal:** Obtain an "Approximately" PO Allocation that is also IR in a "private" way.
  - Definition: An allocation is α - **Approximate PO** if at most an α fraction of players can strictly improve without forcing another player to get a worse type of good.

Model

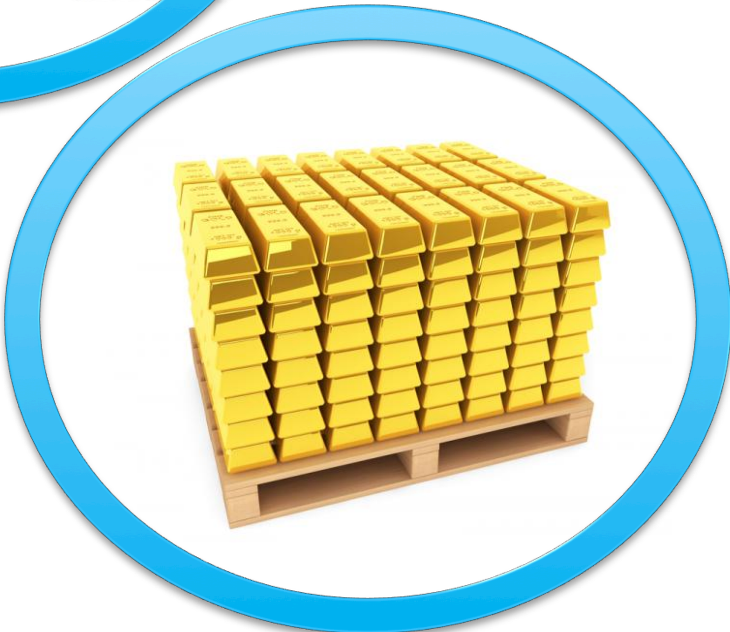

Motivation– Kidney Exchanges

Motivation– Kidney Exchanges

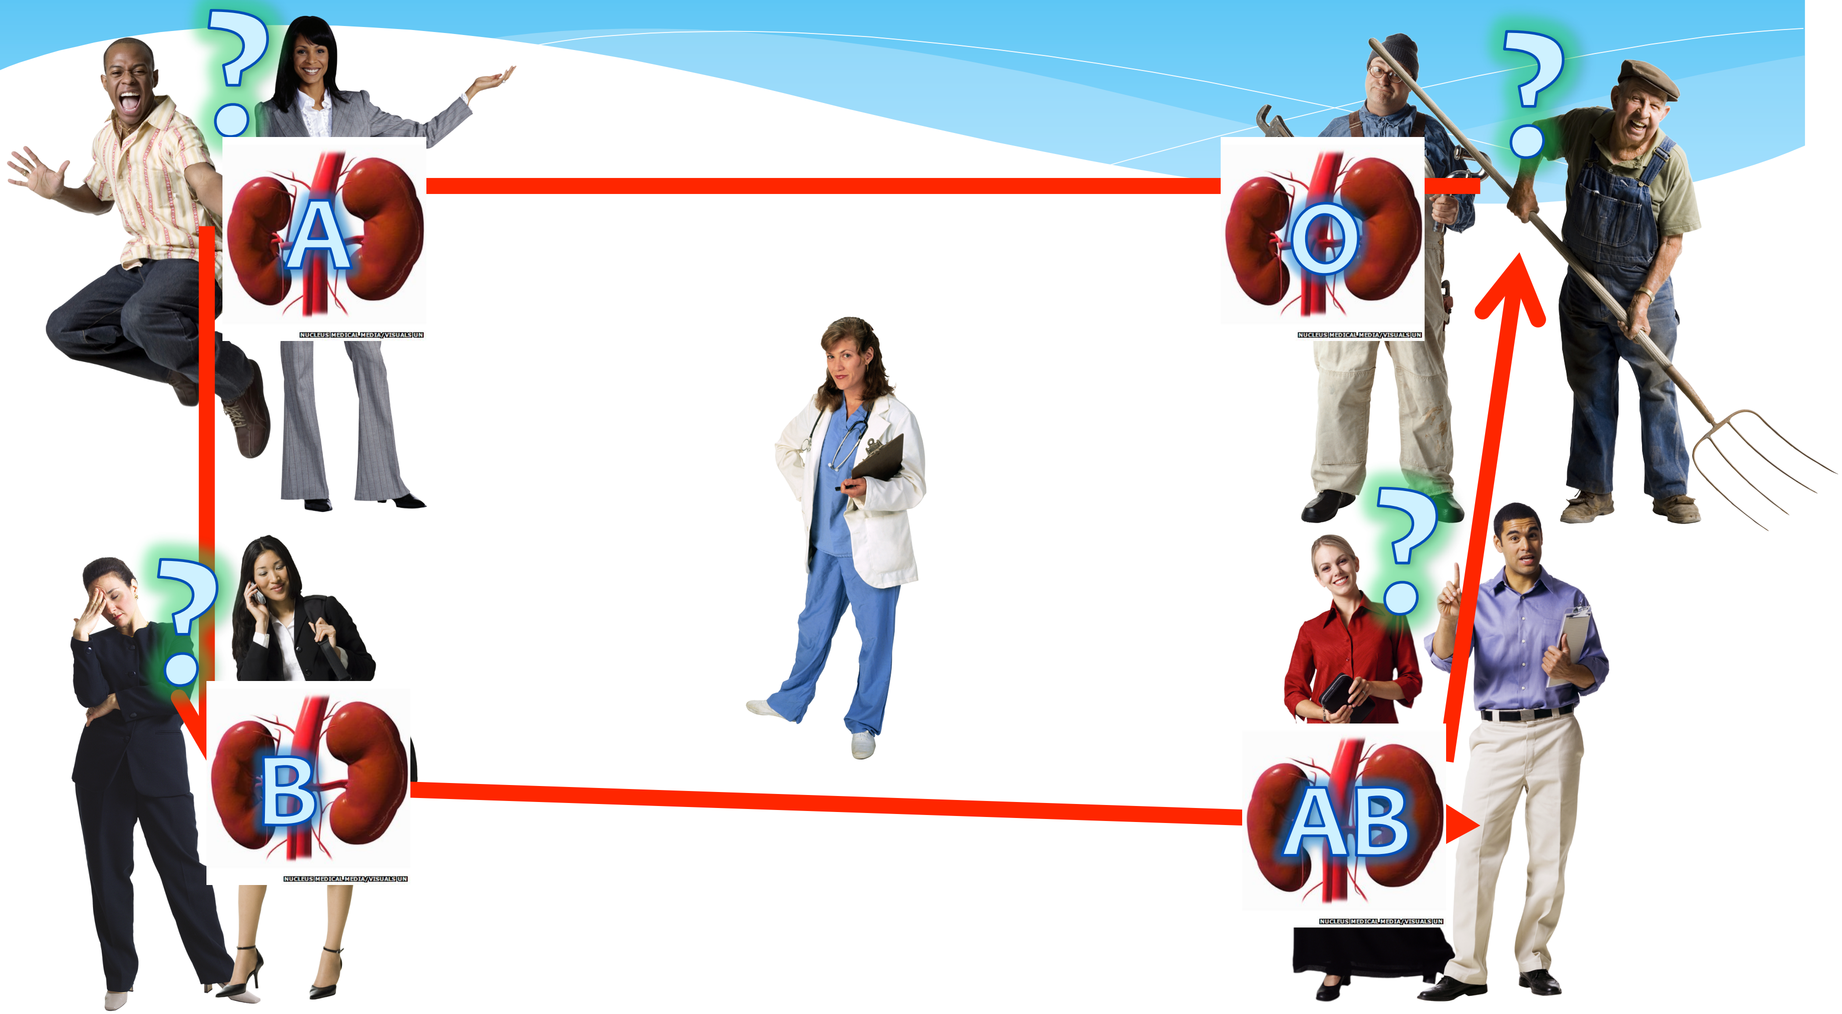

Motivation– Kidney Exchanges

Differential Privacy

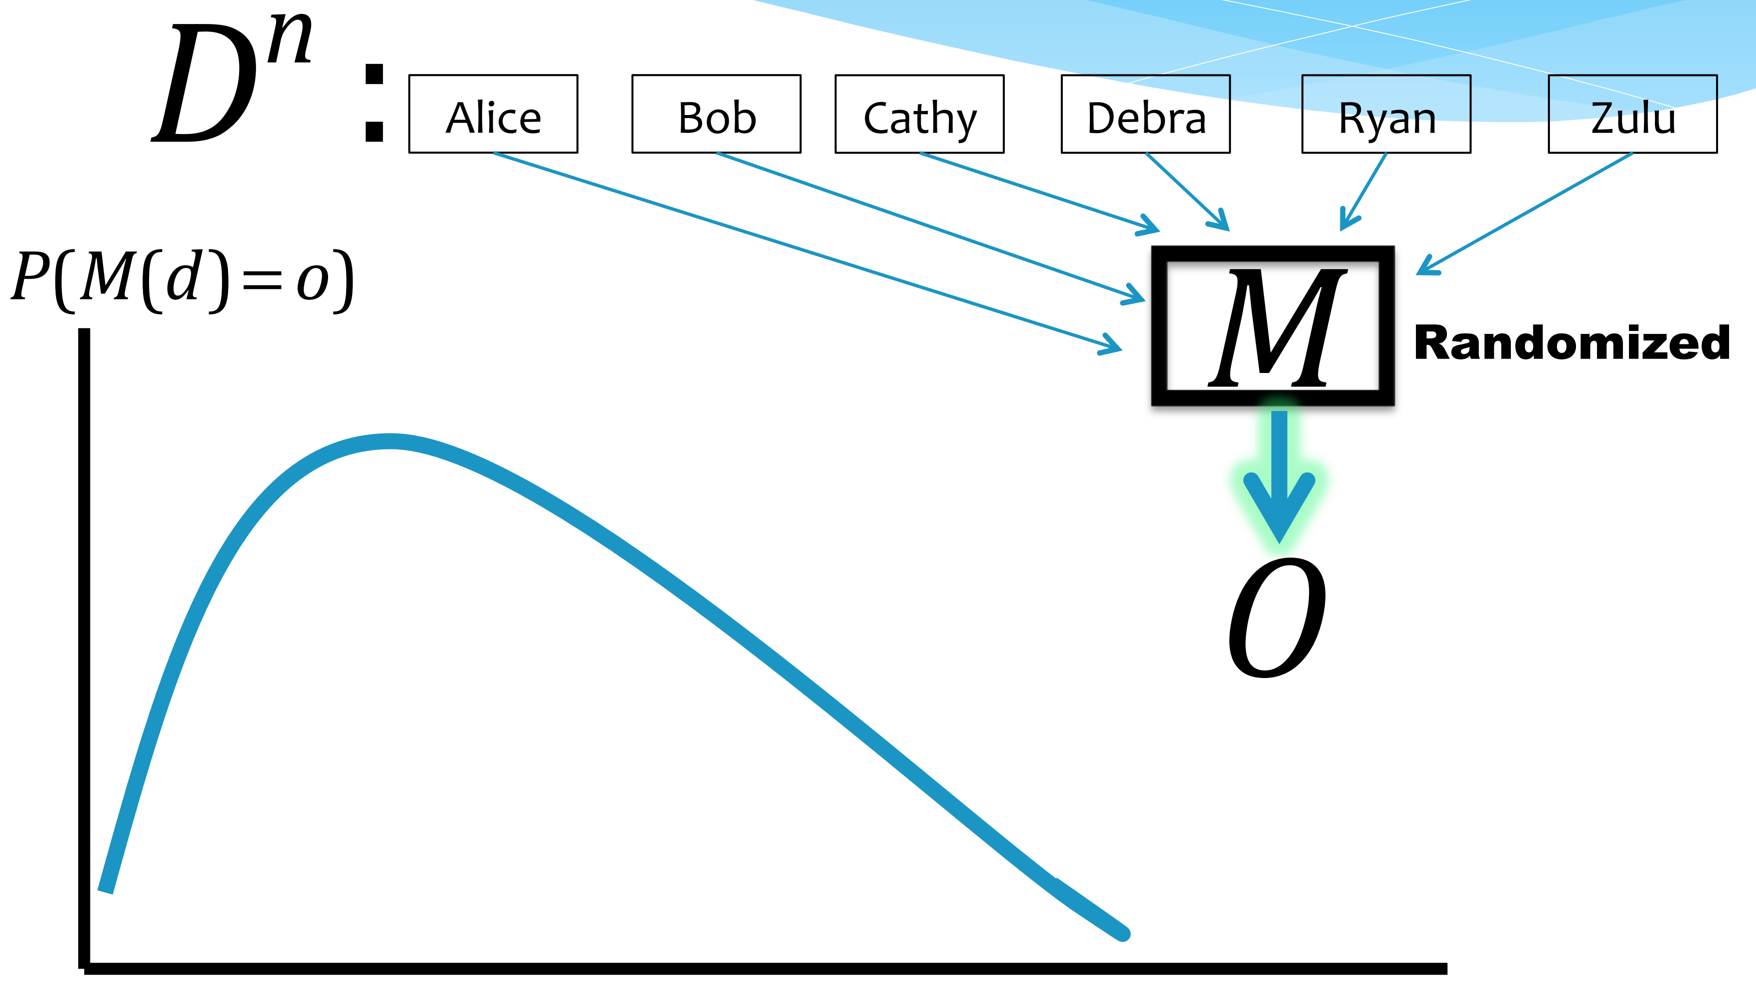

Differential Privacy

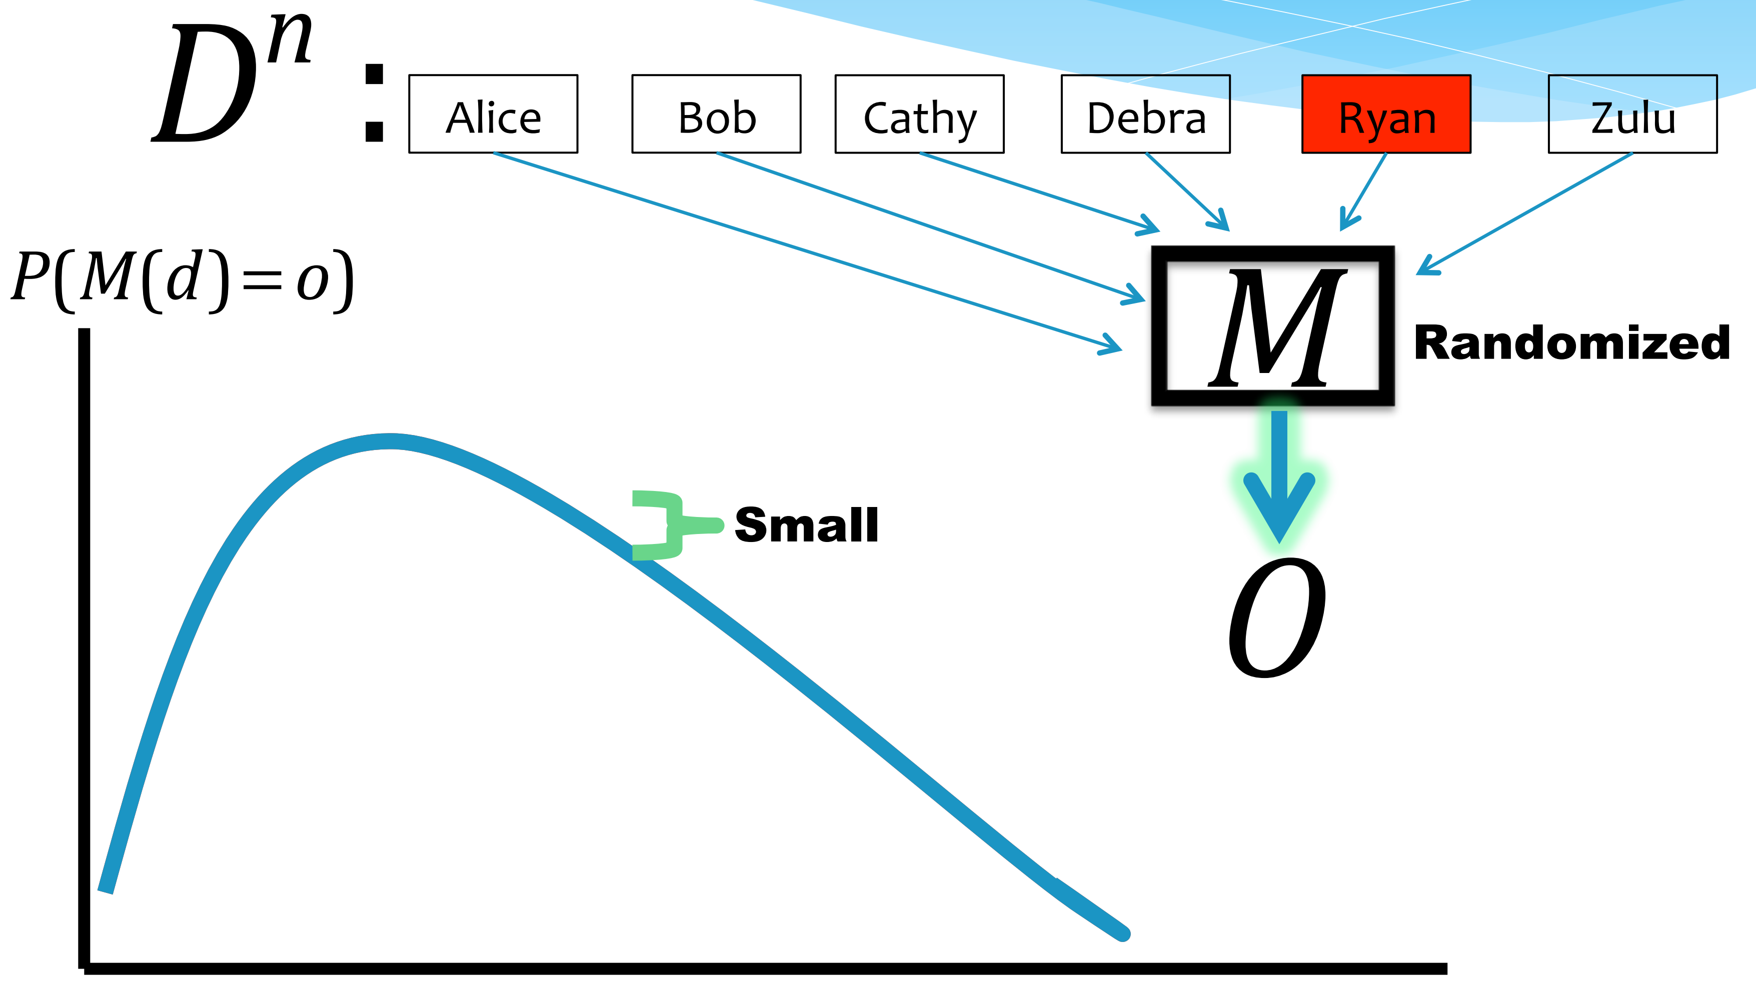

Differential Privacy

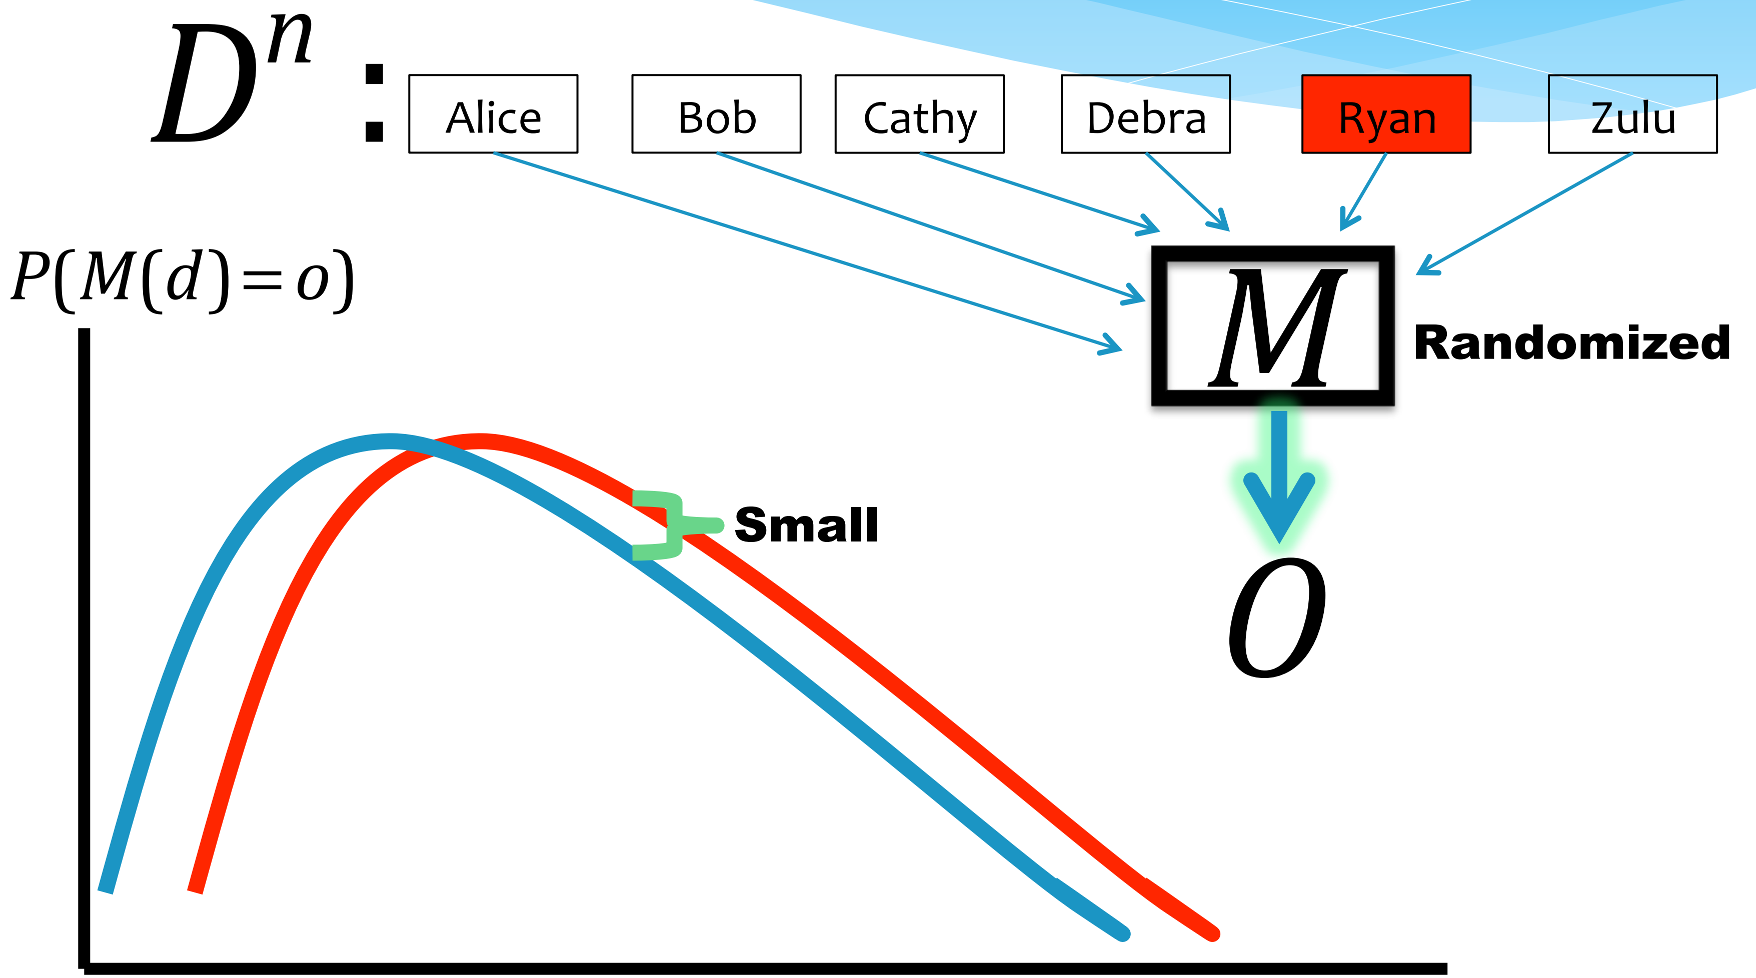

Differential Privacy

- A randomized algorithm $M : D^n \rightarrow O$ is ε-DP if for any $d_i, d_i' \in D$ and any $S \subseteq O$ we have

$(1+\varepsilon)$

$$P(M(d_i, d_{-i}) \in S) \leq e^{\varepsilon} P(M(d_i', d_{-i}) \in S)$$

Differential Privacy

- A randomized algorithm $M : D^n \rightarrow A^n$ is ε-DP if for any $d_i, d_i' \in D$ and any $S \subseteq A^n$ we have

$$P(M(d_i, d_{-i}) \in S) \leq e^{\varepsilon} P(M(d_i', d_{-i}) \in S)$$

Joint Differential Privacy

- A randomized algorithm $M : D^n \rightarrow A^n$ is ε-JDP if for any $d_i, d_i' \in D$ and any $S \subseteq A^{n-1}$ we have

$$P(\underbrace{M(d_i, d_{-i})_{-i}} \in S) \le e^{\varepsilon} P(\underbrace{M(d_i', d_{-i})_{-i}} \in S)$$

Everybody except player *i*

Is JDP Useful Here?

Claim: Any ε-JDP mechanism $M : D^n \rightarrow A^n$ that guarantees an α- PO allocation that is IR has

$$\alpha \geq 1 - \frac{e^{\varepsilon}}{1 + e^{\varepsilon}}$$

Other Definitions of Privacy

- A randomized algorithm $M : D^n \rightarrow A^n$ is ε-JDP if for any $d_i, d_i' \in D$ and any $S \subseteq A^{n-1}$ we have

$$P(M(d_i, d_{-i})_{-i} \in S) \leq e^{\varepsilon} P(M(d_i', d_{-i})_{-i} \in S)$$

Marginal Differential Privacy

- A randomized algorithm $M : D^n \rightarrow A^n$ is ε-MDP if for any $d_i, d_i' \in D$, any $S \subseteq A$, and we have $i \neq j$

$$P(M(d_i, d_{-i})_j \in S) \leq e^{\varepsilon} P(M(d_i', d_{-i})_j \in S)$$

Main Theorem

- Theorem: There exists an ε-MDP mechanism that gives an allocation that is IR and α-PO for

$$\alpha = \tilde{O}\left(\frac{k^{4.5}}{\varepsilon n}\right)$$

- Note that $\alpha \to 0$ as $n \to \infty$

Private-TTC

- Input: $x = \left(g_i, \succ_i\right)_{i=1}^{n}$
- Output: An allocation π

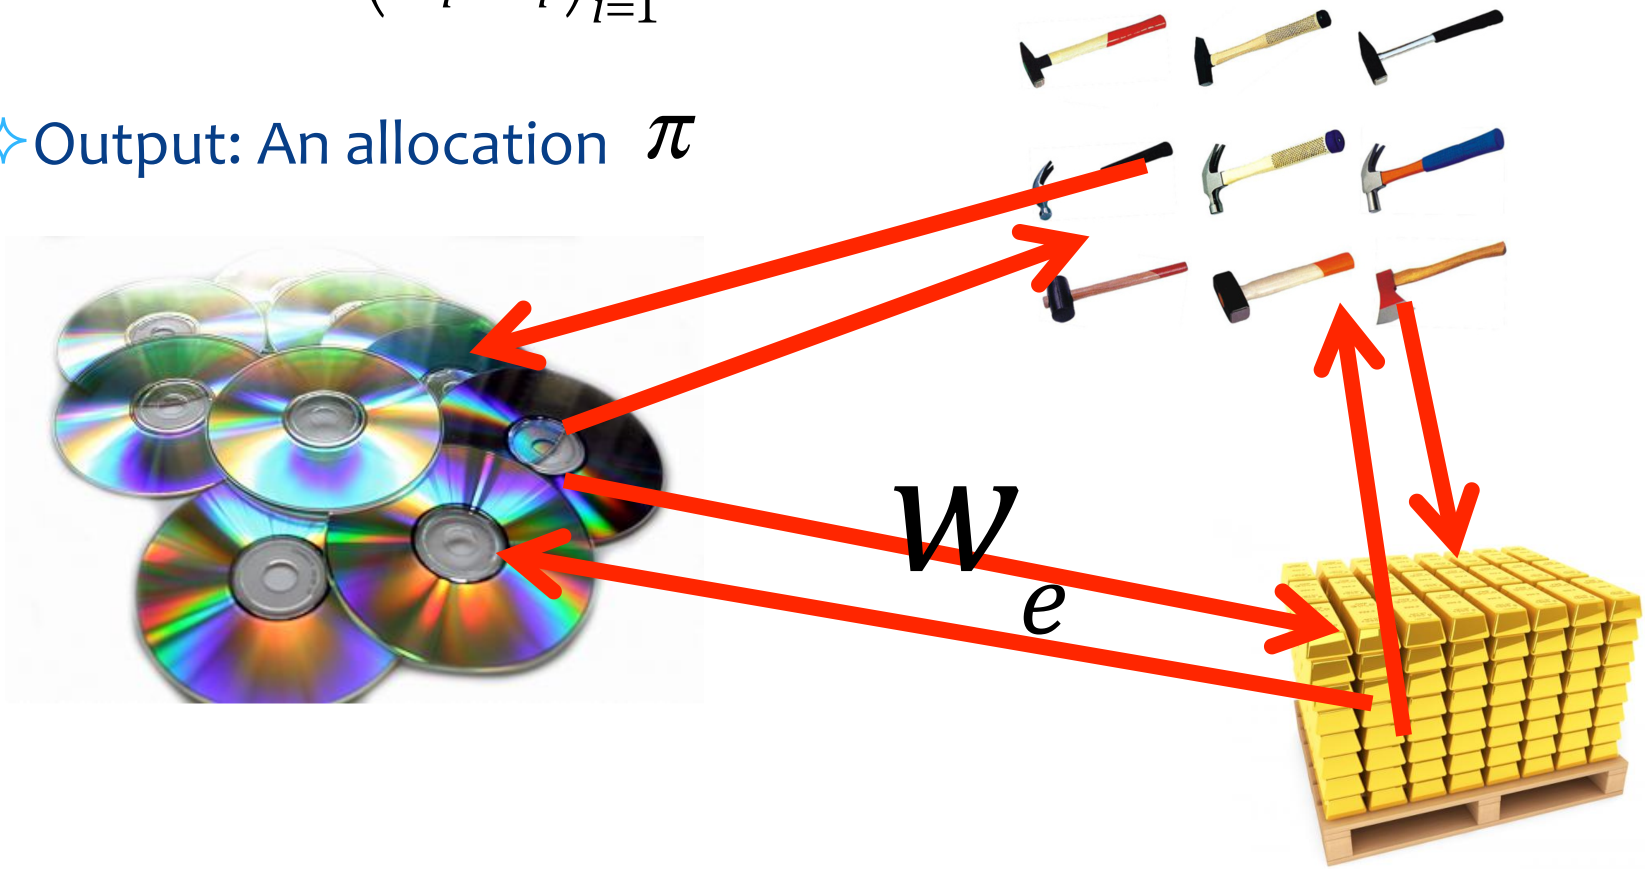

Private-TTC

1. Let $Z_e \sim Laplace(\sigma)$ and assign each edge weight $\hat{w}_e \leftarrow w_e + Z_e - 2E$ to each edge, where E is a high probability error bound.
2. Choose a positive weight cycle C w.r.t. $\hat{w}_e$
3. Select uniformly at random $\min_{e \in C}\{\hat{w}_e\}$ from the w_e that wanted to trade
4. Allocate to those selected the good they wanted – Return to 2.

Private-TTC

5. If there is no cycle, then there must be some node v with no outgoing edges, so delete v .
6. Have everyone that was pointing to v repoint to the good they most prefer of those that are left – Return to 1
7. When all nodes are deleted, all people not allocated a good gets the good they started with.

Observations

- The output allocation is always IR
- With high probability, we have

$$E \leq w_e - \hat{w}_e \leq 3E$$

- When a node gets deleted, there may be at most

$$D = O(kE)$$

goods left at that node

Open Problems

- We considered:

$$\varepsilon - MDP \;\leq\; ? \;\leq\; \varepsilon - JDP$$

- What if we allow players to lie about their preferences? Can we create a Mechanism that is also **Truthful**?

Questions?