CSC2556

Lecture 6

Kidney Exchange
Cake-Cutting

[Some illustrations due to: Ariel Procaccia]

Announcements

- Project proposal
  - Due: Mar 06 by 11:59PM
  - I'll soon put up a few sample project ideas.
  - If you have trouble finding a project idea, meet me.

- Structure
  - Problem space introduction
  - High-level research question
  - Prior work
  - Detailed goals

- Length: Ideally 1 page (2 pages max)

Kidney Exchange

Kidney Exchange

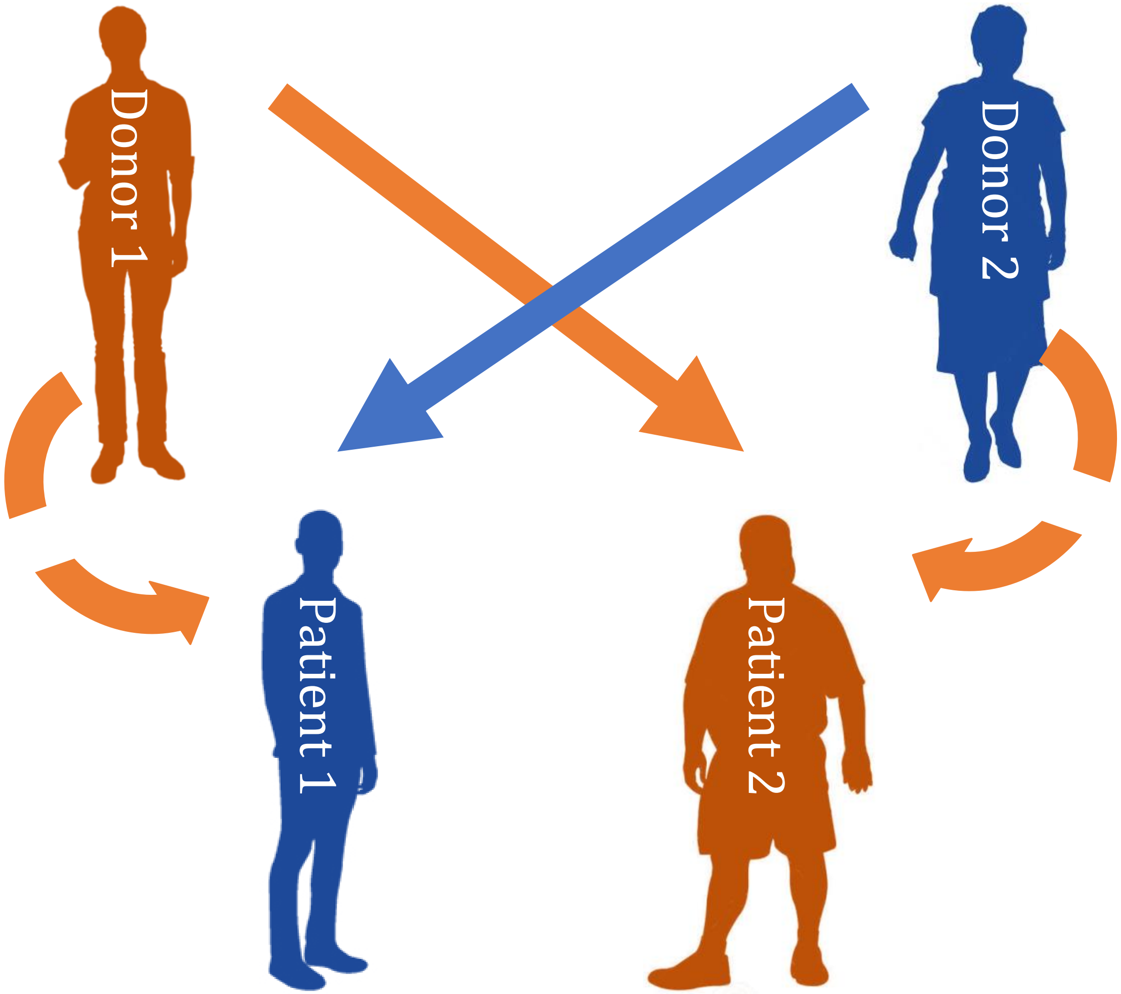

Incentives

- A decade ago kidney exchanges were carried out by individual hospitals
- Today there are nationally organized exchanges; participating hospitals have little other interaction
- It was observed that hospitals match easy-to-match pairs internally, and enroll only hard-to-match pairs into larger exchanges
- Goal: incentivize hospitals to enroll all their pairs

The strategic model

- Undirected graph, only pairwise matches
  - Vertex = donor-patient pair
  - Edge = compatibility

- Each agent controls a subset of vertices
  - Possible strategy: hide some vertices (match internally), and only reveal others
  - Utility of agent = # its matched vertices (self-matched + matched by mechanism)

The strategic model

- Mechanism:
  - Input: revealed vertices by agents (edges are public)
  - Output: matching
- Target: # matched vertices
- Strategyproof (SP): If no agent benefits from hiding vertices irrespective of what other agents do.

OPT is manipulable

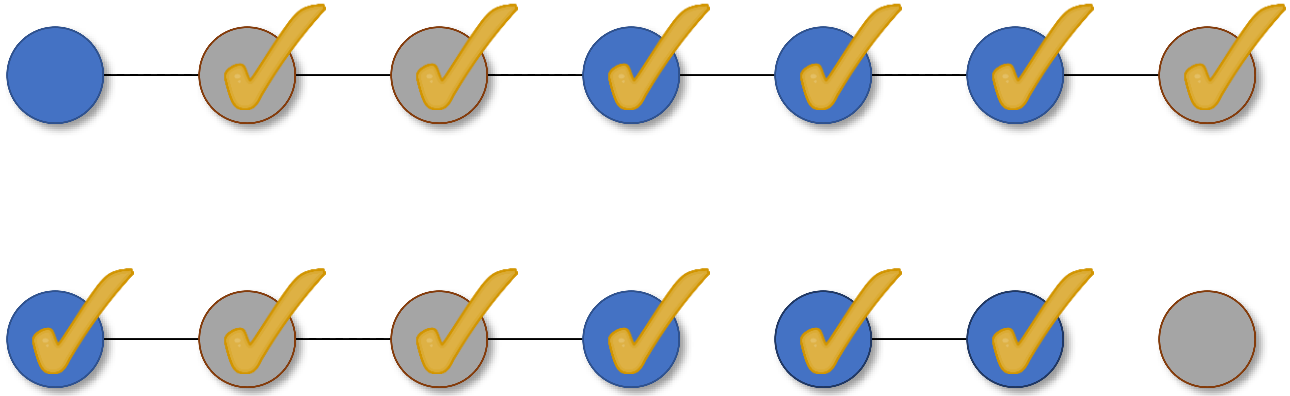

OPT is manipulable

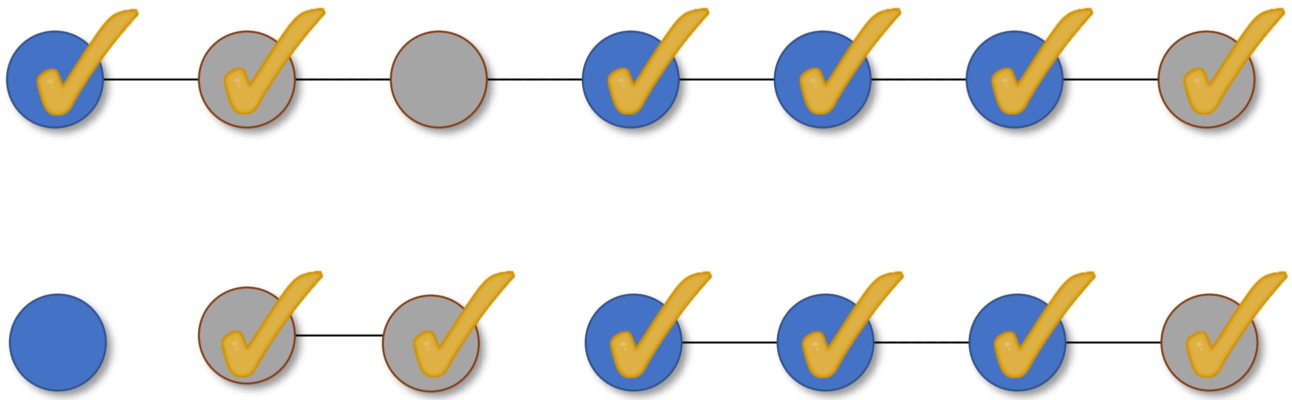

Approximating SW

- Theorem [Ashlagi et al. 2010]: No deterministic SP mechanism can give a $2 - \epsilon$ approximation
- Proof:

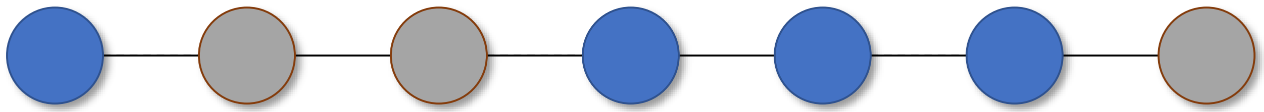

  - No perfect matching exists.
  - Any algorithm must either leave a blue node or a gray node unmatched.

Approximating SW

- Theorem [Ashlagi et al. 2010]: No deterministic SP mechanism can give a $2 - \epsilon$ approximation
- Proof:

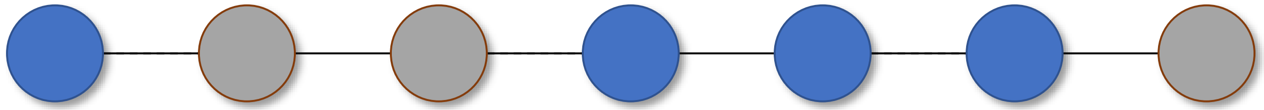

  - Suppose it leaves a blue node unmatched
    - If the blue agent hides two nodes as follows, the mechanism is forced to return a matching of size 1 when a matching of size 2 exists.

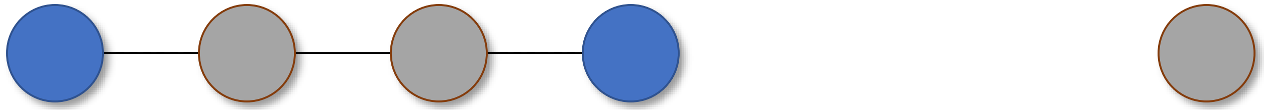

Approximating SW

- Theorem [Ashlagi et al. 2010]: No deterministic SP mechanism can give a $2 - \epsilon$ approximation
- Proof:

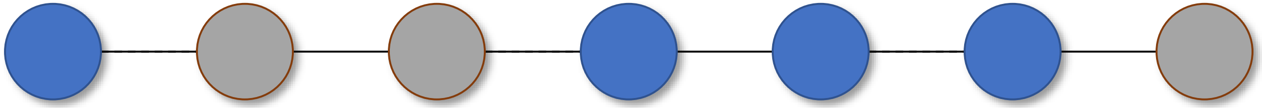

  - Suppose it leaves a gray node unmatched
    - If the gray agent hides two nodes as follows, the mechanism is forced to return a matching of size 1 when a matching of size 2 exists.

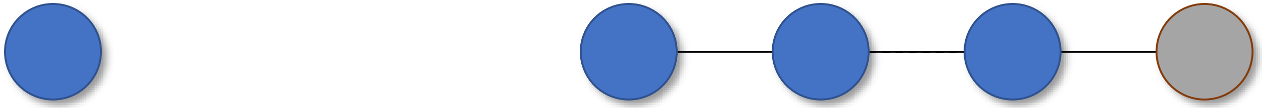

Approximating SW

- Theorem [Kroer and Kurokawa 2013]: No randomized SP mechanism can give a $\frac{6}{5} - \epsilon$ approximation.
- Proof: Homework!

SP mechanism: Take 1

- Assume two agents

- $\text{M}\small{\text{ATCH}}_{\{\{1\},\{2\}\}}$ mechanism:
  - Consider matchings that maximize the number of "internal edges" for each agent.
  - Among these return, a matching with max overall cardinality.

Another example

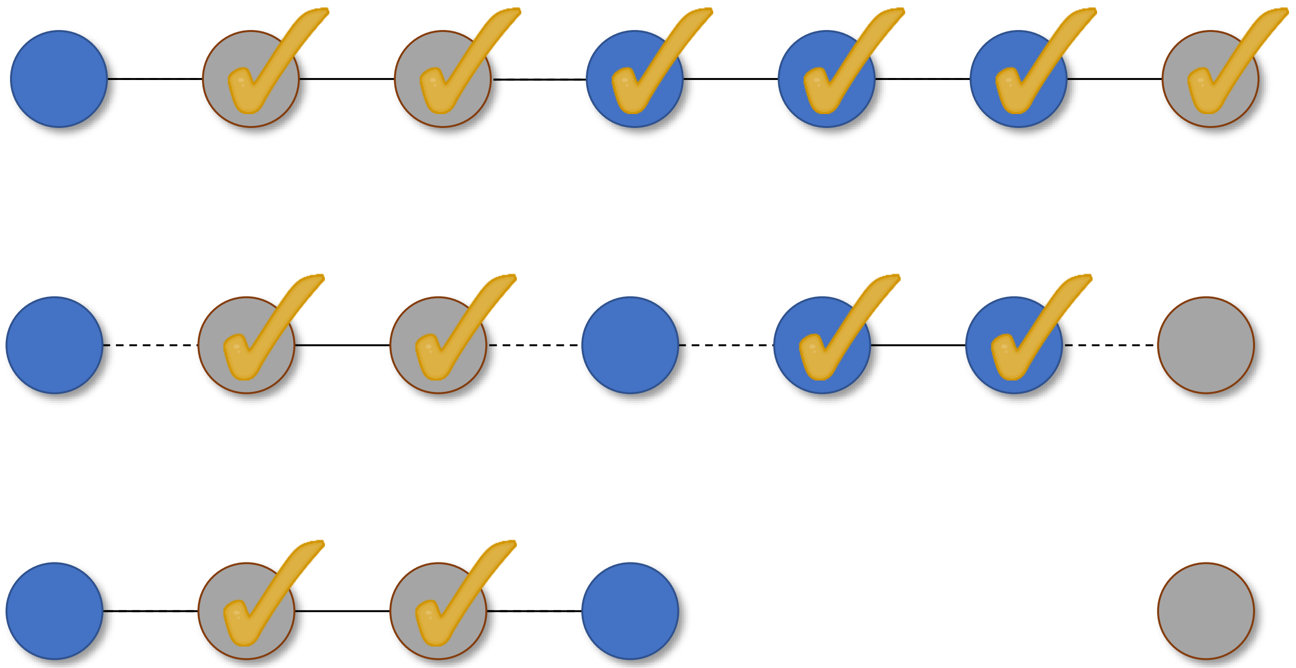

Guarantees

- $\textsc{Match}_{\{\{1\},\{2\}\}}$ gives a 2-approximation
  - Cannot add more edges to matching
  - For each edge in optimal matching, one of the two vertices is in mechanism's matching

- Theorem (special case): $\textsc{Match}_{\{\{1\},\{2\}\}}$ is strategyproof for two agents.

Proof

- M = matching when player 1 is honest, M' = matching when player 1 hides vertices
- $M \Delta M'$ consists of paths and even-length cycles, each consisting of alternating M, M' edges

What's wrong with the illustration on the right?

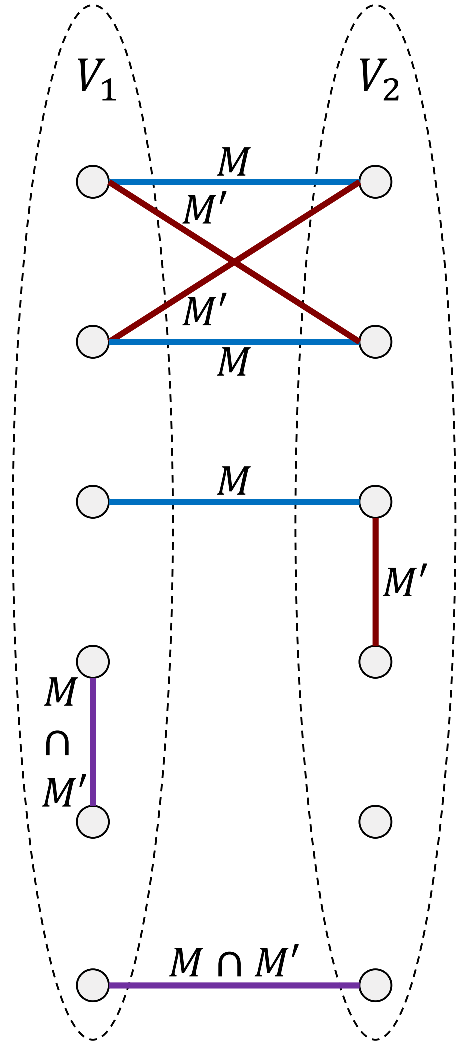

Proof

- Consider a path in $M \Delta M'$, denote its edges in M by P and its edges in M' by P'

- Consider sets P_{11}, P_{22}, P_{12} containing edges of P among V_1, among V_2, and between V_1- V_2
  - Same for $P'_{11}, P'_{22}, P'_{12}$

- Note that $|P_{11}| \geq |P'_{11}|$
  - Property of the algorithm

Proof

- Case 1: $|P_{11}| = |P'_{11}|$

- Agent 2's vertices don't change, so $|P_{22}| = |P'_{22}|$
- M is max cardinality $\Rightarrow |P_{12}| \geq |P'_{12}|$
- $U_1(P) = 2|P_{11}| + |P_{12}|$

$$\geq 2|P'_{11}| + |P'_{12}| = U_1(P')$$

Proof

- Case 2: $|P_{11}| > |P'_{11}|$
- $|P_{12}| \geq |P'_{12}| - 2$
  - Every sub-path within V_2 is of even length
  - Pair up edges of P_{12} and P'_{12}, except maybe the first and the last
- $U_1(P) = 2|P_{11}| + |P_{12}|$
  $\geq 2(|P'_{11}| + 1) + |P'_{12}| - 2$
  $= U_1(P')$ ■

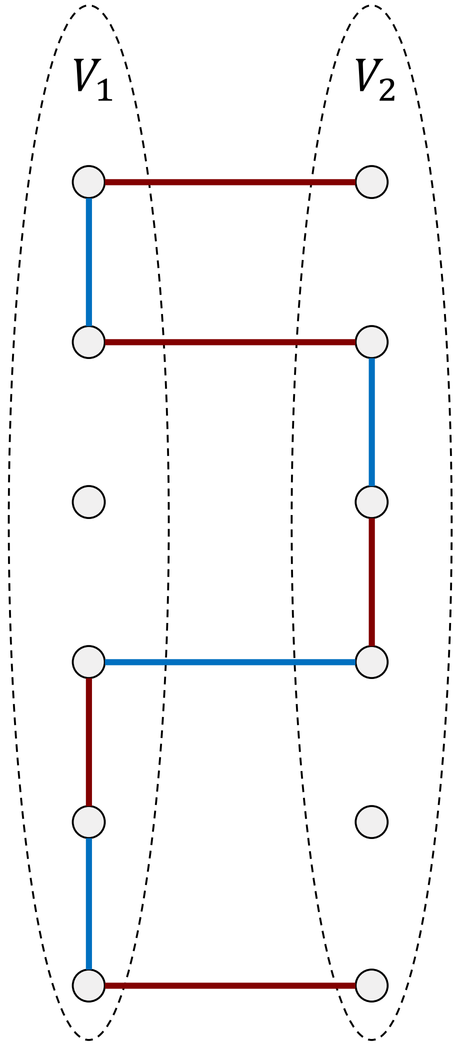

The case of 3 players

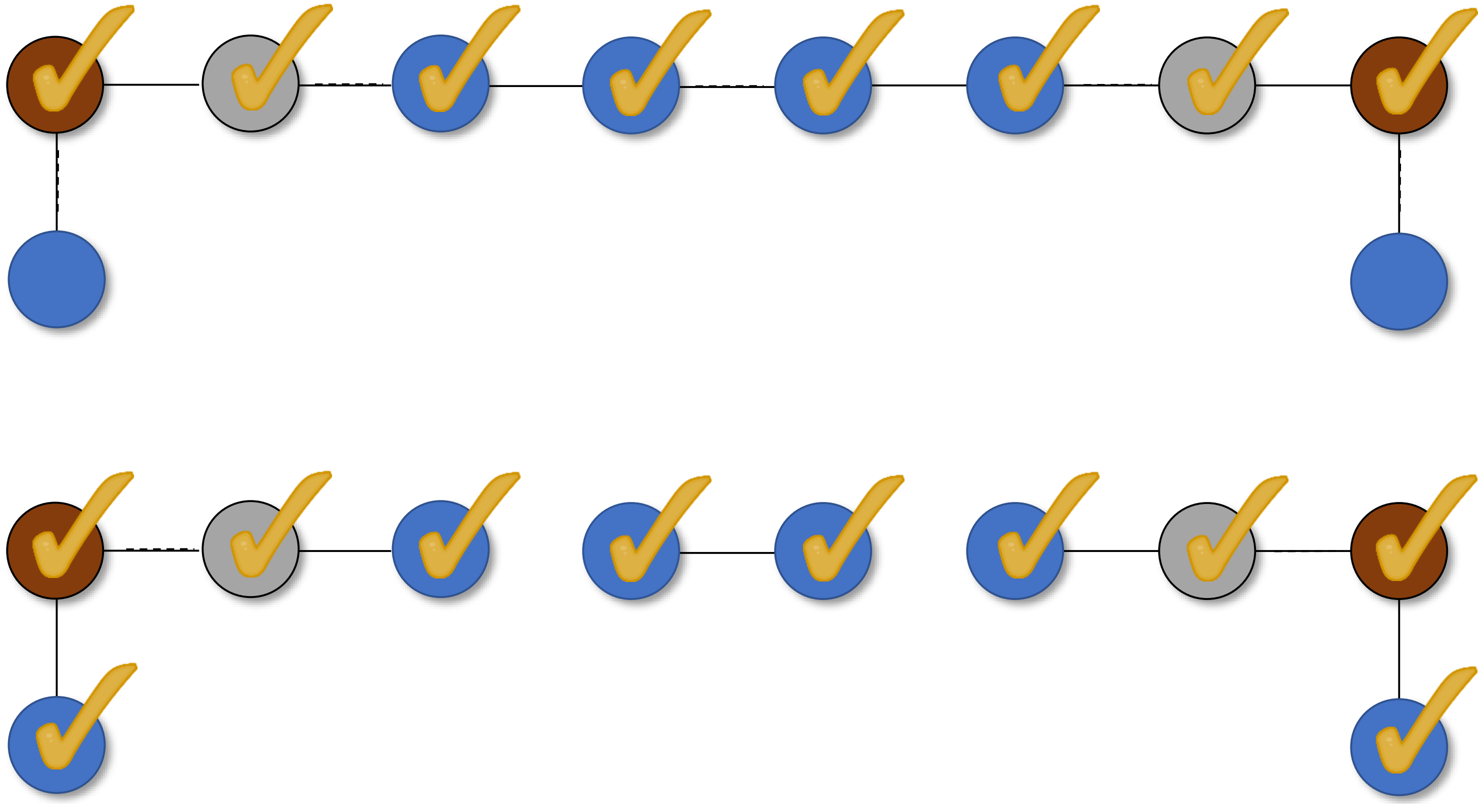

SP Mechanism: Take 2

- Let $\Pi = (\Pi_1, \Pi_2)$ be a bipartition of the players

- MATCH$_\Pi$ mechanism:
  - Consider matchings that maximize the number of "internal edges" and do not have any edges between different players on the same side of the partition
  - Among these return a matching with max cardinality (need tie breaking)

Eureka?

- Theorem [Ashlagi et al. 2010]: MATCH_{Π} is strategyproof for any number of agents and any partition Π.
- Recall: For $n = 2$, $\text{MATCH}_{\{\{1\},\{2\}\}}$ is a 2-approximation
- Question: $n = 3$, $\text{MATCH}_{\{\{1\},\{2,3\}\}}$ approximation?
  1. 2
  2. 3
  3. 4
  4. More than 4

The Mechanism

- The MIX-AND-MATCH mechanism:
  - Mix: choose a random partition Π
  - Match: Execute MATCH_{Π}

- Theorem [Ashlagi et al. 2010]: MIX-AND-MATCH is strategyproof and a 2-approximation.

- We only prove the approximation ratio.

Proof

- $M^* =$ optimal matching
- Claim: I can create a matching M' such that
  - M' is max cardinality on each V_i, and
  - $\sum_i |M'_{ii}| + \frac{1}{2} \sum_{i \neq j} |M'_{ij}| \geq \sum_i |M^*_{ii}| + \frac{1}{2} \sum_{i \neq j} |M^*_{ij}|$

  - $M^{**} =$ max cardinality on each V_i
  - For each path P in $M^* \Delta M^{**}$, add $P \cap M^{**}$ to M' if M^{**} has more internal edges than M^*, otherwise add $P \cap M^*$ to M'
  - For every internal edge M' gains relative to M^*, it loses at most one edge overall ■

Proof

- Fix Π and let M^{Π} be the output of MATCH_{Π}

- The mechanism returns max cardinality across Π subject to being max cardinality internally, therefore

$$\sum_{i} \left|M_{ii}^{\Pi}\right| + \sum_{i \in \Pi_1, j \in \Pi_2} \left|M_{ij}^{\Pi}\right| \geq \sum_{i} \left|M'_{ii}\right| + \sum_{i \in \Pi_1, j \in \Pi_2} \left|M'_{ij}\right|$$

Proof

$$\mathbb{E}[|M^{\Pi}|] = \frac{1}{2^n} \sum_{\Pi} \left(\sum_i |M_{ii}^{\Pi}| + \sum_{i \in \Pi_1, j \in \Pi_2} |M_{ij}^{\Pi}| \right)$$

$$\geq \frac{1}{2^n} \sum_{\Pi} \left(\sum_i |M_{ii}'| + \sum_{i \in \Pi_1, j \in \Pi_2} |M_{ij}'| \right)$$

$$= \sum_i |M_{ii}'| + \frac{1}{2^n} \sum_{\Pi} \sum_{i \in \Pi_1, j \in \Pi_2} |M_{ij}'|$$

$$= \sum_i |M_{ii}'| + \frac{1}{2} \sum_{i \neq j} |M_{ij}'| \geq \sum_i |M_{ii}^*| + \frac{1}{2} \sum_{i \neq j} |M_{ij}^*|$$

$$\geq \frac{1}{2} \sum_i |M_{ii}^*| + \frac{1}{2} \sum_{i \neq j} |M_{ij}^*| = \frac{1}{2} |M^*| \quad \blacksquare$$

Cake-Cutting

Cake-Cutting

- A heterogeneous, divisible good
  - Heterogeneous: it may be valued differently by different individuals
  - Divisible: we can share/divide it between individuals
- Represented as $[0,1]$
  - Almost without loss of generality
- Set of players $N = \{1, \dots, n\}$
- Piece of cake $X \subseteq [0,1]$
  - A finite union of disjoint intervals

Agent Valuations

- Each player i has a valuation V_i that is very much like a probability distribution over $[0,1]$

- Additive: For $X \cap Y = \emptyset$, $V_i(X) + V_i(Y) = V_i(X \cup Y)$

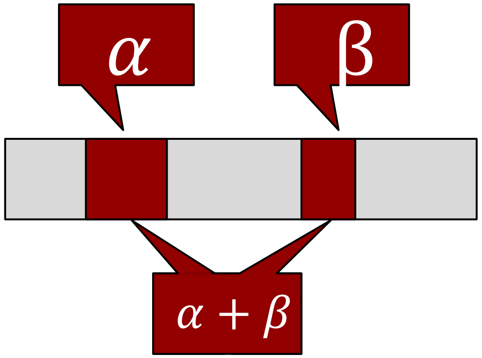

- Normalized: $V_i([0,1]) = 1$

- Divisible: $\forall \lambda \in [0,1]$ and X, $\exists Y \subseteq X$ s.t. $V_i(Y) = \lambda V_i(X)$

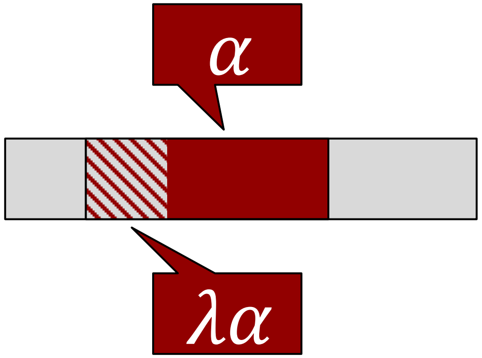

Fairness Goals

- An allocation is a disjoint partition $A = (A_1, \dots, A_n)$ of the cake
- We desire the following fairness properties from our allocation A:

- Proportionality (Prop):

$$\forall i \in N: \; V_i(A_i) \geq \frac{1}{n}$$

- Envy-Freeness (EF):

$$\forall i, j \in N: V_i(A_i) \geq V_i(A_j)$$

Fairness Goals

- Prop: $\forall i \in N: V_i(A_i) \geq 1/n$
- EF: $\forall i, j \in N: V_i(A_i) \geq V_i(A_j)$

- Question: What is the relation between proportionality and EF?
  1. Prop $\Rightarrow$ EF
  2. EF $\Rightarrow$ Prop
  3. Equivalent
  4. Incomparable

Cut-and-Choose

- Algorithm for $n = 2$ players

- Player 1 divides the cake into two pieces X, Y s.t.
$$V_1(X) = V_1(Y) = 1/2$$

- Player 2 chooses the piece she prefers.

- This is EF and therefore proportional.
  - Why?

Input Model

- How do we measure the "time complexity" of a cake-cutting algorithm for n players?
- Typically, time complexity is a function of the length of input encoded as binary.
- Our input consists of functions V_i, which requires infinite bits to encode.
- We want running time just as a function of n.

Robertson-Webb Model

- We restrict access to valuations V_i's through two types of queries:
  - $\text{Eval}_i(x, y)$ returns $V_i([x, y])$
  - $\text{Cut}_i(x, \alpha)$ returns y such that $V_i([x, y]) = \alpha$

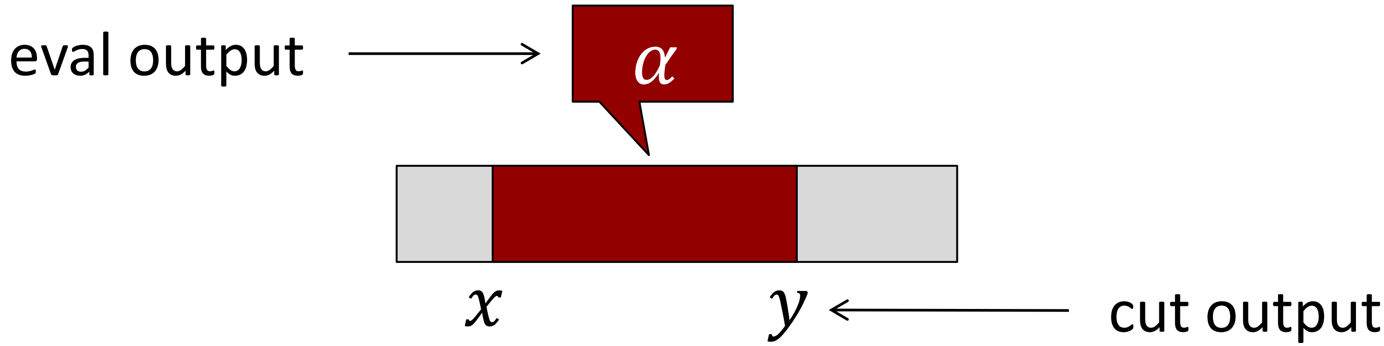

Robertson-Webb Model

- Two types of queries:
  - $\text{Eval}_i(x, y) = V_i([x, y])$
  - $\text{Cut}_i(x, \alpha) = y$ s.t. $V_i([x, y]) = \alpha$
- Question: How many queries are needed to find an EF allocation when $n = 2$?
- Answer: 2
  - Why?

Dubins-Spanier

- Protocol for finding a proportional allocation for n players

- Referee starts at 0, and continuously moves knife to the right.
- Repeat: when the piece to the left of knife is worth $1/n$ to a player, the player shouts “stop”, gets the piece, and exits.
- The last player gets the remaining piece.

Dubins-Spanier

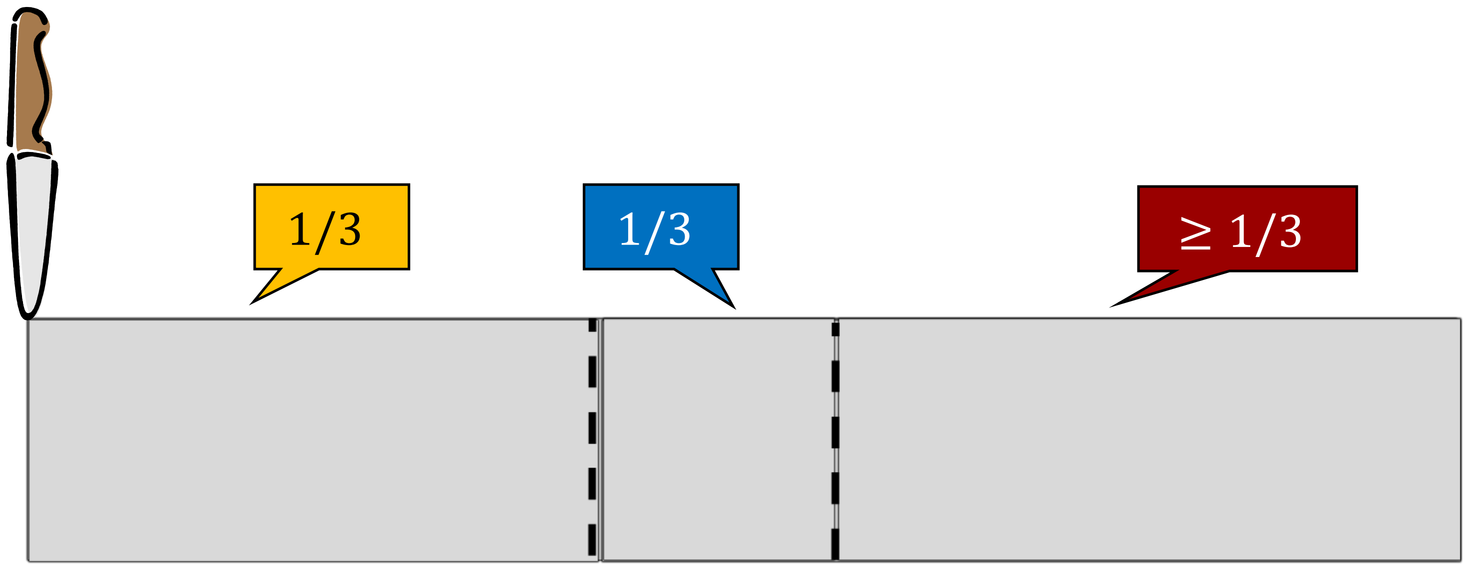

Dubins-Spanier

- Moving knife is not really needed.
- At each stage, we can ask each remaining player a cut query to mark his $1/n$ point in the remaining cake.
- Move the knife to the leftmost mark.

DUBINS-SPANIER

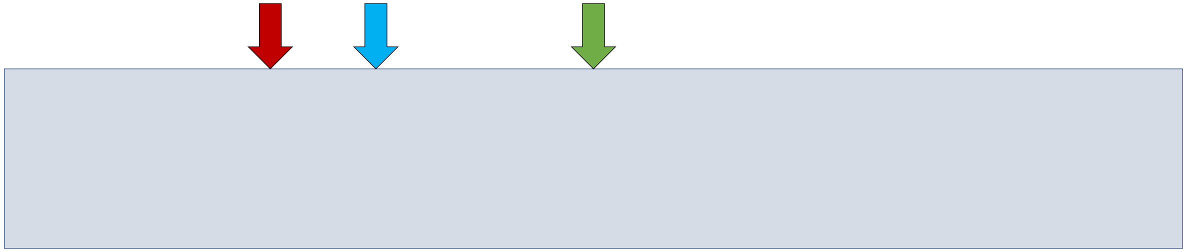

DUBINS-SPANIER

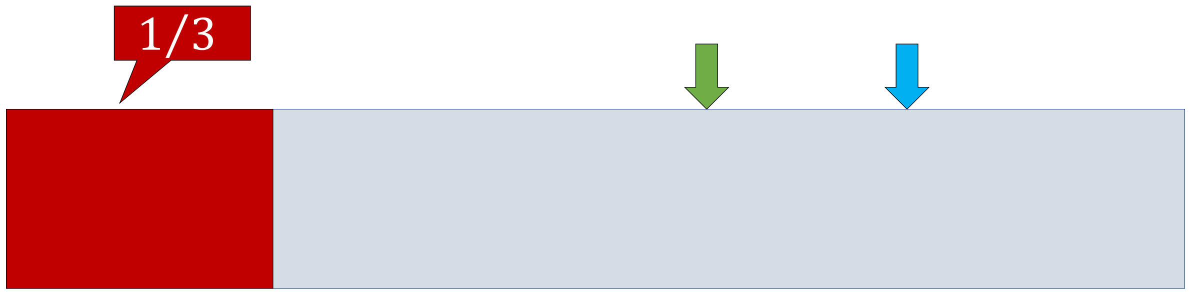

DUBINS-SPANIER

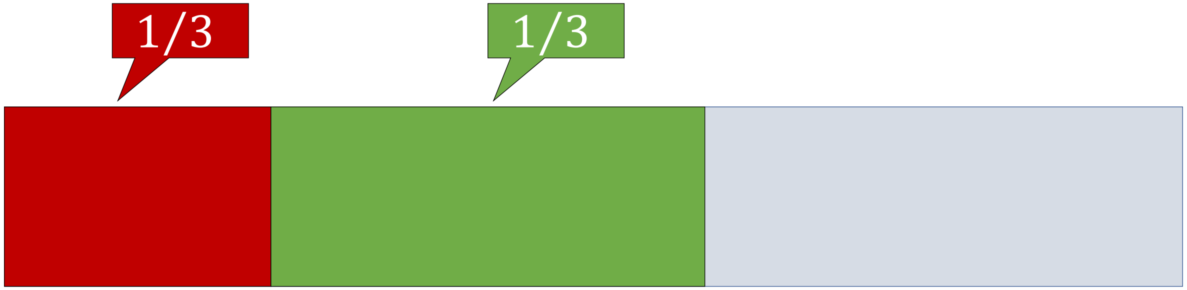

Dubins-Spanier

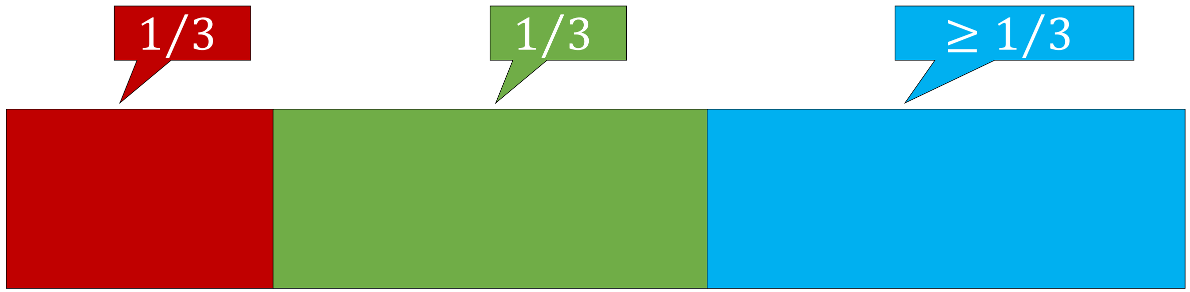

Dubins-Spanier

- Question: What is the complexity of the Dubins-Spanier protocol in the Robertson-Webb model?

1. $\Theta(n)$
2. $\Theta(n \log n)$
3. $\Theta(n^2)$
4. $\Theta(n^2 \log n)$

Even-Paz

- Input: Interval $[x, y]$, number of players n
  - Assume $n = 2^k$ for some k
- If $n = 1$, give $[x, y]$ to the single player.
- Otherwise, let each player i mark z_i s.t.

$$V_i([x, z_i]) = \frac{1}{2} V_i([x, y])$$

- Let z^* be the $n/2$ mark from the left.
- Recurse on $[x, z^*]$ with the left $n/2$ players, and on $[z^*, y]$ with the right $n/2$ players.

Even-Paz

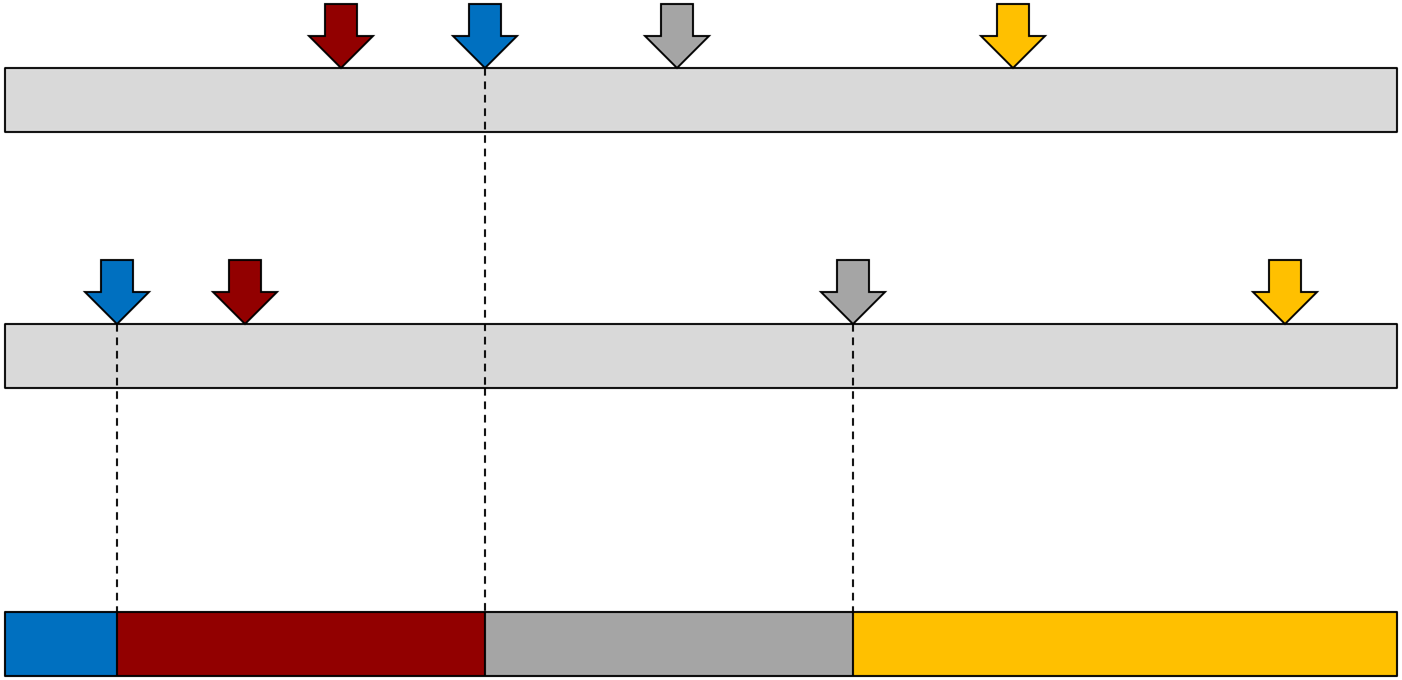

EVEN-PAZ

- Theorem: EVEN-PAZ returns a Prop allocation.
- Proof:
  - Inductive proof. We want to prove that if player i is allocated piece A_i when $[x, y]$ is divided between n players, $V_i(A_i) \geq (1/n)V_i([x, y])$
    - Then Prop follows because initially $V_i([x, y]) = V_i([0,1]) = 1$
  - Base case: $n = 1$ is trivial.
  - Suppose it holds for $n = 2^{k-1}$. We prove for $n = 2^k$.
  - Take the 2^{k-1} left players.
    - Every left player i has $V_i([x, z^*]) \geq (1/2)\, V_i([x, y])$
    - If it gets A_i, by induction, $V_i(A_i) \geq \frac{1}{2^{k-1}}\, V_i([x, z^*]) \geq \frac{1}{2^k}\, V_i([x, y])$

EVEN-PAZ

- Question: What is the complexity of the Even-Paz protocol in the Robertson-Webb model?

1. $\Theta(n)$
2. $\Theta(n \log n)$
3. $\Theta(n^2)$
4. $\Theta(n^2 \log n)$

Complexity of Proportionality

- Theorem [Edmonds and Pruhs, 2006]: Any proportional protocol needs $\Omega(n \log n)$ operations in the Robertson-Webb model.

- Thus, the EVEN-PAZ protocol is (asymptotically) provably optimal!

Envy-Freeness?

- "I suppose you are also going to give such cute algorithms for finding envy-free allocations?"

- Bad luck. For n-player EF cake-cutting:
  - [Brams and Taylor, 1995] give an unbounded EF protocol.
  - [Procaccia 2009] shows $\Omega(n^2)$ lower bound for EF.
  - Last year, the long-standing major open question of "bounded EF protocol" was resolved!
  - [Aziz and Mackenzie, 2016]: $O\left(n^{n^{n^{n^{n^{n}}}}}\right)$ protocol!
    - Not a typo!