CSC2556

Lecture 5

Facility Location
Stable Matching

Facility Location

Apprx Mechanism Design

1. Define the problem: agents, outcomes, values
2. Fix an objective function (e.g., maximizing sum of values)
3. Check if the objective function is maximized through a strategyproof mechanism
4. If not, find the strategyproof mechanism that provides the best worst-case approximation ratio of the objective function

Facility Location

- Set of agents N
- Each agent i has a true location $x_i \in \mathbb{R}$
- Mechanism f
  - Takes as input reports $\tilde{x} = (\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_n)$
  - Returns a location $y \in \mathbb{R}$ for the new facility
- Cost to agent i : $c_i(y) = |y - x_i|$
- Social cost $C(y) = \sum_i c_i(y) = \sum_i |y - x_i|$

Facility Location

- Social cost $C(y) = \sum_i c_i(y) = \sum_i |y - x_i|$
- Q: Ignoring incentives, what choice of y would minimize the social cost?
- A: The median location $\text{med}(x_1, \dots, x_n)$
  - n is odd → the unique "(n+1)/2"th smallest value
  - n is even → "n/2"th or "(n/2)+1"st smallest value
  - Why?

Facility Location

- Social cost $C(y) = \sum_i c_i(y) = \sum_i |y - x_i|$
- Median is optimal (i.e., 1-approximation)
- What about incentives?
  - Median is also strategyproof (SP)!
  - Irrespective of the reports of other agents, agent i is best off reporting x_i

Median is SP

No manipulation can help

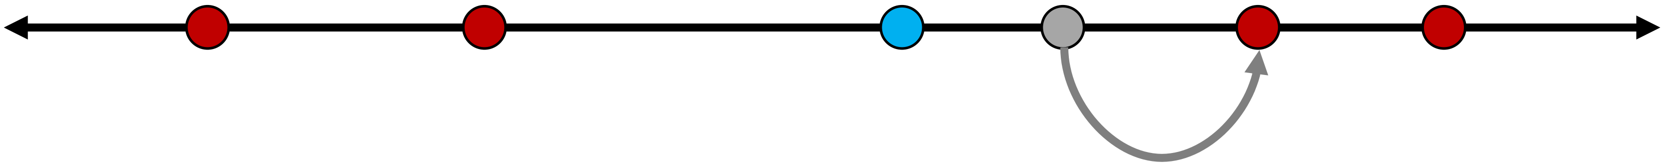

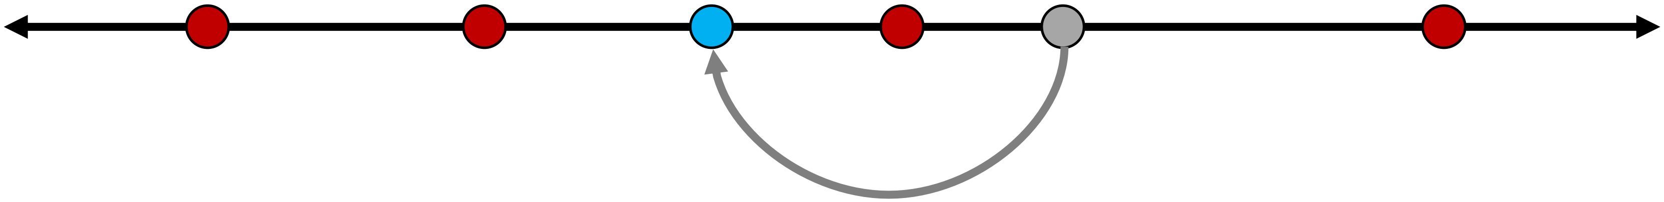

Max Cost

- A different objective function $C(y) = \max_i |y - x_i|$
- Q: Again ignoring incentives, what value of y minimizes the maximum cost?
- A: The midpoint of the leftmost ($\min_i x_i$) and the rightmost ($\max_i x_i$) locations
- Q: Is this optimal rule strategyproof?
- A: No!

Max Cost

- $C(y) = \max_i |y - x_i|$
- We want to use a strategyproof mechanism.
- Question: What is the approximation ratio of median for maximum cost?
  1. $\in [1,2)$
  2. $\in [2,3)$
  3. $\in [3,4)$
  4. $\in [4, \infty)$

Max Cost

- Answer: 2-approximation

- Other SP mechanisms that are 2-approximation
  - Leftmost: Choose the leftmost reported location
  - Rightmost: Choose the rightmost reported location
  - Dictatorship: Choose the location reported by agent 1
  - …

Max Cost

- Theorem [Procaccia & Tennenholtz, '09]
  No deterministic SP mechanism has approximation ratio < 2 for maximum cost.
- Proof:

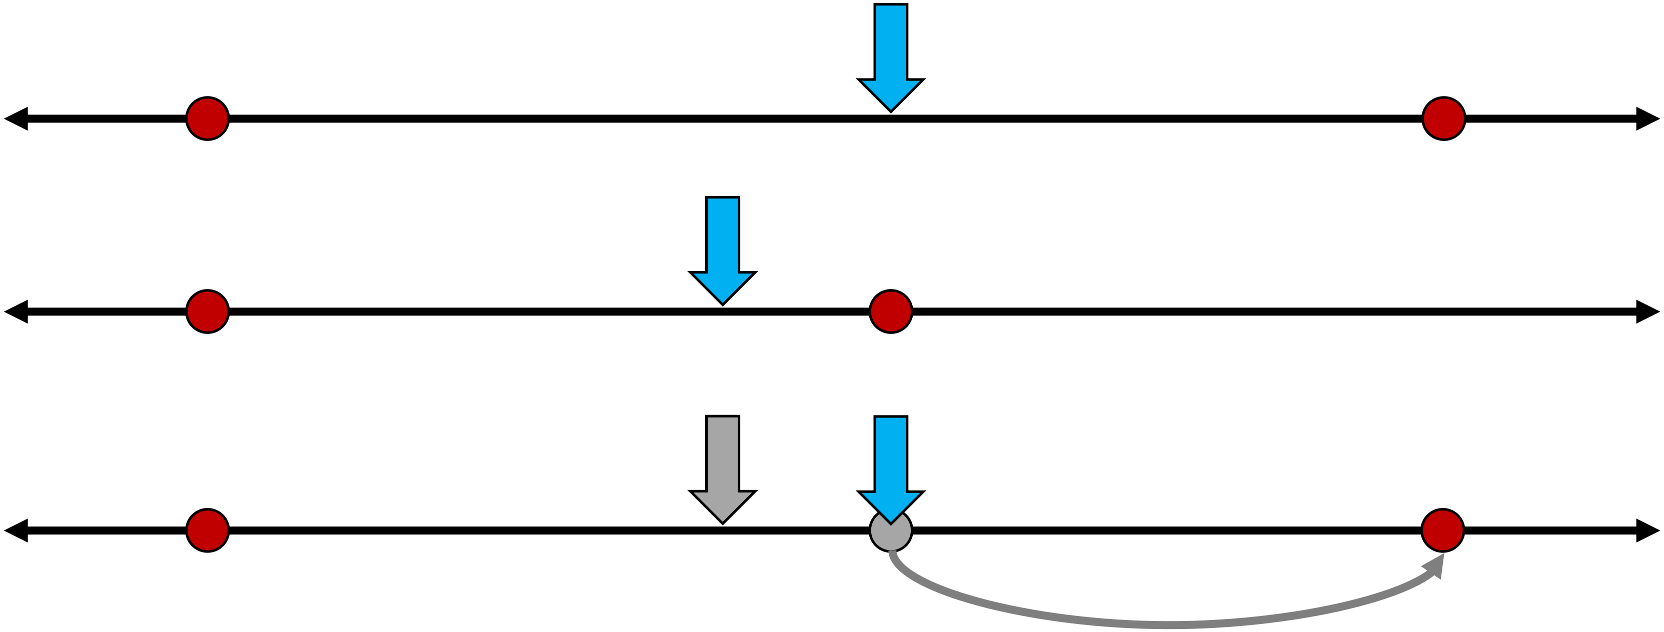

Max Cost + Randomized

- The Left-Right-Middle (LRM) Mechanism
  - Choose $\min_i x_i$ with probability ¼
  - Choose $\max_i x_i$ with probability ¼
  - Choose $(\min_i x_i + \max_i x_i)/2$ with probability ½

- Question: What is the approximation ratio of LRM for maximum cost?

- At most $\frac{(1/4)*2C+(1/4)*2C+(1/2)*C}{C} = \frac{3}{2}$

Max Cost + Randomized

- Theorem [Procaccia & Tennenholtz, '09]:
  The LRM mechanism is strategyproof.
- Proof:

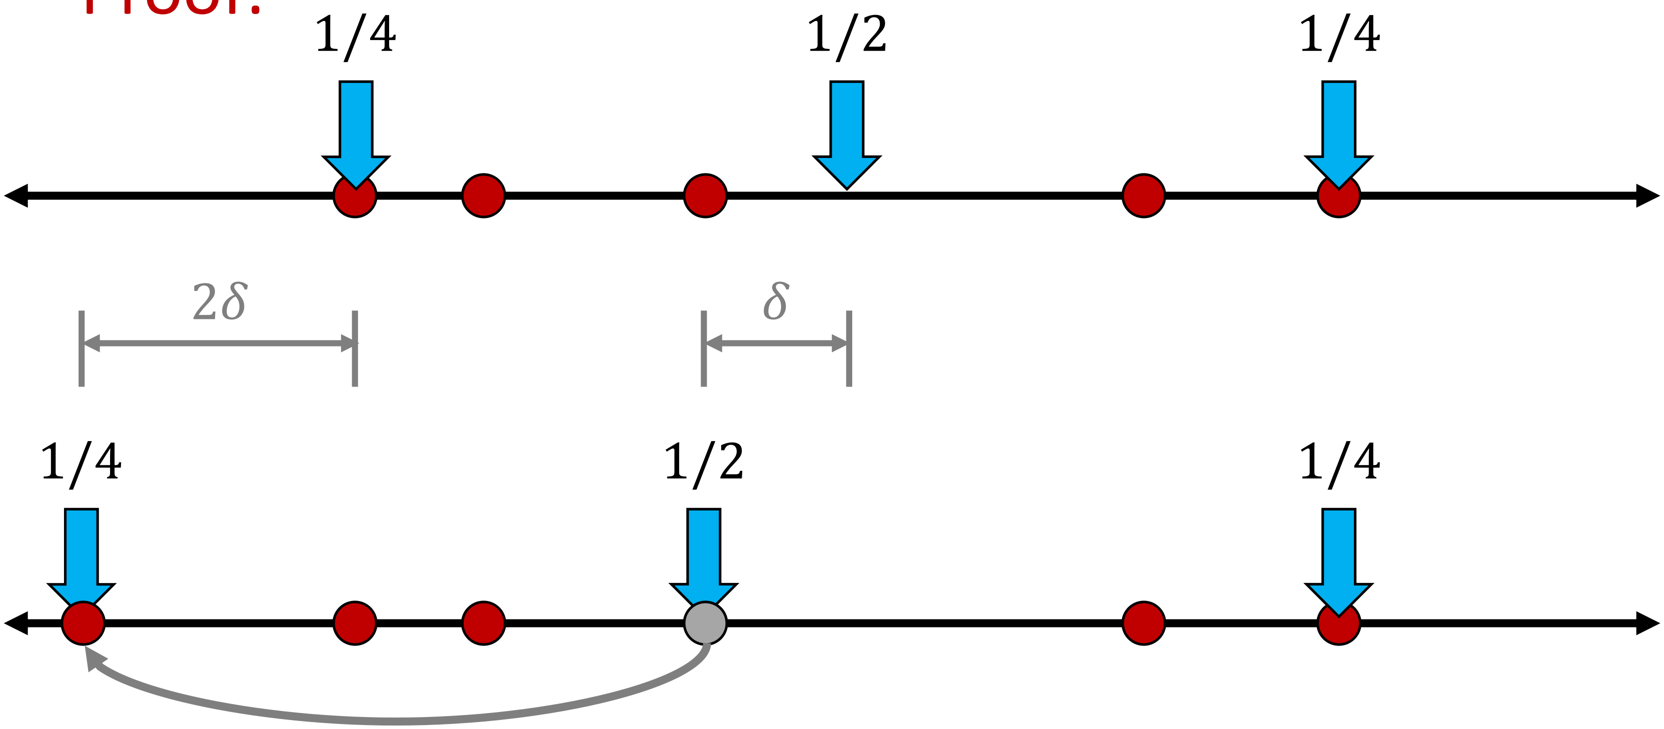

Max Cost + Randomized

- Exercise for you!
  Try showing that no randomized SP mechanism can achieve approximation ratio $< 3/2$.

Stable Matching

Stable Matching

- Recap Graph Theory:

- In graph $G = (V, E)$, a matching $M \subseteq E$ is a set of edges with no common vertices
  - That is, each vertex should have at most one incident edge
  - A matching is perfect if no vertex is left unmatched.

- G is a bipartite graph if there exist V_1, V_2 such that $V = V_1 \cup V_2$ and $E \subseteq V_1 \times V_2$

Stable Marriage Problem

- Bipartite graph, two sides with equal vertices
  - n men and n women (old school terminology ☹)
- Each man has a ranking over women & vice versa
  - E.g., Eden might prefer Alice $\succ$ Tina $\succ$ Maya
  - And Tina might prefer Tony $\succ$ Alan $\succ$ Eden
- Want: a perfect, stable matching
  - Match each man to a unique woman such that no pair of man m and woman w prefer each other to their current matches (such a pair is called a "blocking pair")

Example: Preferences

Albert	Diane	Emily	Fergie
Bradley	Emily	Diane	Fergie
Charles	Diane	Emily	Fergie

Diane	Bradley	Albert	Charles
Emily	Albert	Bradley	Charles
Fergie	Albert	Bradley	Charles

Example: Matching 1

Albert	Diane	Emily	Fergie
Bradley	Emily	Diane	Fergie
Charles	Diane	Emily	Fergie

Diane	Bradley	Albert	Charles
Emily	Albert	Bradley	Charles
Fergie	Albert	Bradley	Charles

Question: Is this a stable matching?

Example: Matching 1

Albert	Diane	Emily	Fergie
Bradley	Emily	Diane	Fergie
Charles	Diane	Emily	Fergie

Diane	Bradley	Albert	Charles
Emily	Albert	Bradley	Charles
Fergie	Albert	Bradley	Charles

No, Albert and Emily form a **blocking pair**.

Example: Matching 2

Albert	Diane	Emily	Fergie
Bradley	Emily	Diane	Fergie
Charles	Diane	Emily	Fergie

Diane	Bradley	Albert	Charles
Emily	Albert	Bradley	Charles
Fergie	Albert	Bradley	Charles

Question: How about this matching?

Example: Matching 2

Albert	Diane	Emily	Fergie
Bradley	Emily	Diane	Fergie
Charles	Diane	Emily	Fergie

Diane	Bradley	Albert	Charles
Emily	Albert	Bradley	Charles
Fergie	Albert	Bradley	Charles

Yes! (Charles and Fergie are unhappy, but helpless.)

Does a stable matching always exist in the marriage problem?

Can we compute it in a strategyproof way?

Can we compute it efficiently?

Gale-Shapley 1962

- Men-Proposing Deferred Acceptance (MPDA):

1. Initially, no proposals, engagements, or matches are made.
2. While some man m is unengaged:
   - $w \leftarrow m$'s most preferred woman to whom m has not proposed yet
   - m proposes to w
   - If w is unengaged:
     - m and w are engaged
   - Else if w prefers m to her current partner m'
     - m and w are engaged, m' becomes unengaged
   - Else: w rejects m
3. Match all engaged pairs.

Example: MPDA

Albert	Diane	Emily	Fergie
Bradley	Emily	Diane	Fergie
Charles	Diane	Emily	Fergie

Diane	Bradley	Albert	Charles
Emily	Albert	Bradley	Charles
Fergie	Albert	Bradley	Charles

= proposed (orange) = engaged (green) = rejected (purple)

Running Time

- Theorem: DA terminates in polynomial time (at most n^2 iterations of the outer loop)

- Proof:
  - In each iteration, a man proposes to someone to whom he has never proposed before.
  - n men, n women $\rightarrow n \times n$ possible proposals
  - Can actually tighten a bit to $n(n-1)+1$ iterations

- At termination, it must return a perfect matching.

Stable Matching

- Theorem: DA always returns a stable matching.

- Proof by contradiction:
  - Assume (m, w) is a blocking pair.

  - Case 1: m never proposed to w
    - m cannot be unmatched o/w algorithm would not terminate.
    - Men propose in the order of preference.
    - Hence, m must be matched with a woman he prefers to w
    - (m, w) is not a blocking pair

Stable Matching

- Theorem: DA always returns a stable matching.

- Proof by contradiction:
  - Assume (m, w) is a blocking pair.

  - Case 2: m proposed to w
    - w must have rejected m at some point
    - Women only reject to get better partners
    - w must be matched at the end, with a partner she prefers to m
    - (m, w) is not a blocking pair

Men-Optimal Stable Matching

- The stable matching found by MPDA is special.
- Valid partner: For a man m, call a woman w a valid partner if (m, w) is in *some* stable matching.
- Best valid partner: For a man m, a woman w is the best valid partner if she is a valid partner, and m prefers her to every other valid partner.
  - Denote the best valid partner of m by $best(m)$.

Men-Optimal Stable Matching

- Theorem: Every execution of MPDA returns the "men-optimal" stable matching: every man is matched to his best valid partner.
  - Surprising that this is a matching. E.g., it means two men cannot have the same best valid partner!

- Theorem: Every execution of MPDA produces the "women-pessimal" stable matching: every woman is matched to her worst valid partner.

Men-Optimal Stable Matching

- Theorem: Every execution of MPDA returns the men-optimal stable matching.
- Proof by contradiction:
  - Let S = matching returned by MPDA.
  - $m \leftarrow$ first man rejected by $best(m) = w$
  - $m' \leftarrow$ the more preferred man due to which w rejected m
  - w is valid for m, so (m, w) part of stable matching S'
  - $w' \leftarrow$ woman m' is matched to in S'
  - We show that S' cannot be stable because (m', w) is a blocking pair.

Men-Optimal Stable Matching

- **Theorem:** Every execution of MPDA returns the men-optimal stable matching.
- **Proof by contradiction:**

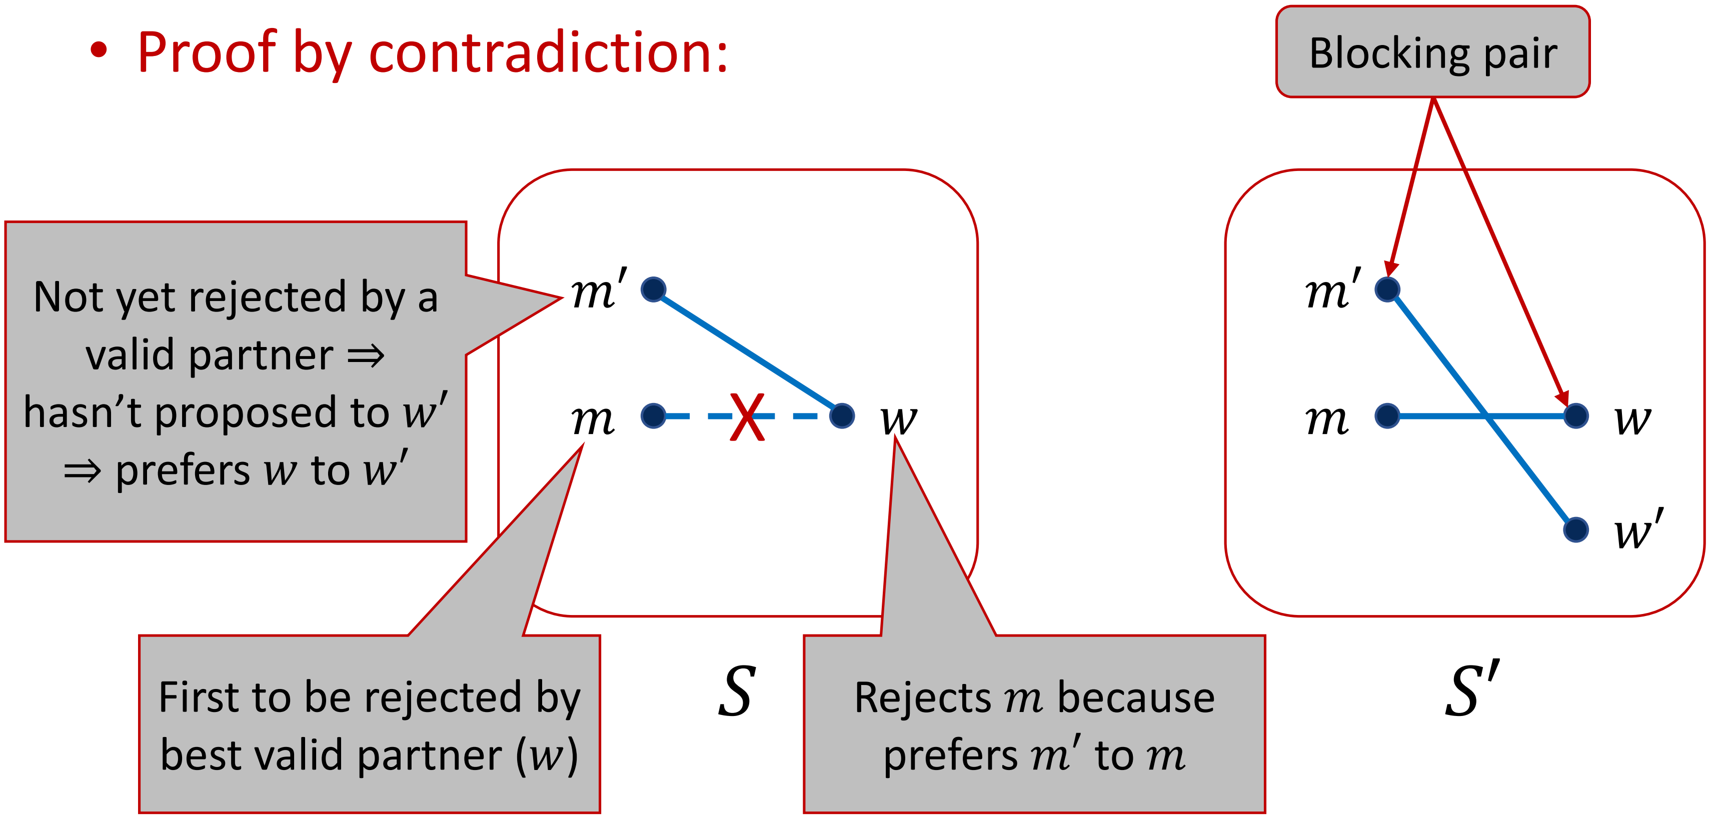

Strategyproofness

- Theorem: MPDA is strategyproof for men.
  - We'll skip the proof of this.
  - Actually, it is group-strategyproof.

- But the women might gain by misreporting.

- Theorem: No algorithm for the stable matching problem is strategyproof for both men and women.

Women-Proposing Version

- Women-Proposing Deferred Acceptance (WPDA)
  - Just flip the roles of men and women
  - Strategyproof for women, not strategyproof for men
  - Returns the women-optimal and men-pessimal stable matching

Extensions

- Unacceptable matches
  - Allow every agent to report a partial ranking
  - If woman w does not include man m in her preference list, it means she would rather be unmatched than matched with m. And vice versa.
  - (m, w) is blocking if each prefers the other over their current state (matched with another partner or unmatched)
  - Just m (or just w) can also be blocking if they prefer being unmatched than be matched to their current partner

- Magically, DA still produces a stable matching.

Extensions

- Resident Matching (or College Admission)
  - Men → residents (or students)
  - Women → hospitals (or colleges)
  - Each side has a ranked preference over the other side
  - But each hospital (or college) q can accept $c_q > 1$ residents (or students)
  - Many-to-one matching

- An extension of Deferred Acceptance works
  - Resident-proposing (resp. hospital-proposing) results in resident-optimal (resp. hospital-optimal) stable matching

Extensions

- For ~20 years, most people thought that these problems are very similar to the stable marriage problem

- Roth [1985] shows:
  - No stable matching algorithm is strategyproof for hospitals (or colleges).

Extensions

- Roommate Matching
  - Still one-to-one matching
  - But no partition into men and women
    - "Generalizing from bipartite graphs to general graphs"
  - Each of n agents submits a ranking over the other $n-1$ agents

- Unfortunately, there are instances where no stable matching exist.
  - A variant of DA can still find a stable matching *if* it exists.
  - Due to Irving [1985]

NRMP: Matching in Practice

- 1940s: Decentralized resident-hospital matching
  - Markets “unralveled”, offers came earlier and earlier, quality of matches decreased
- 1950s: NRMP introduces centralized “clearinghouse”
- 1960s: Gale-Shapley introduce DA
- 1984: Al Roth studies NRMP algorithm, finds it is really a version of DA!
- 1970s: Couples increasingly don’t use NRMP
- 1998: NRMP implements matching with couple constraints (stable matchings may not exist anymore…)
- More recently, DA applied to college admissions