CSCI 357: Algorithmic Game Theory

Lecture 13: Stable Matching

Shikha Singh


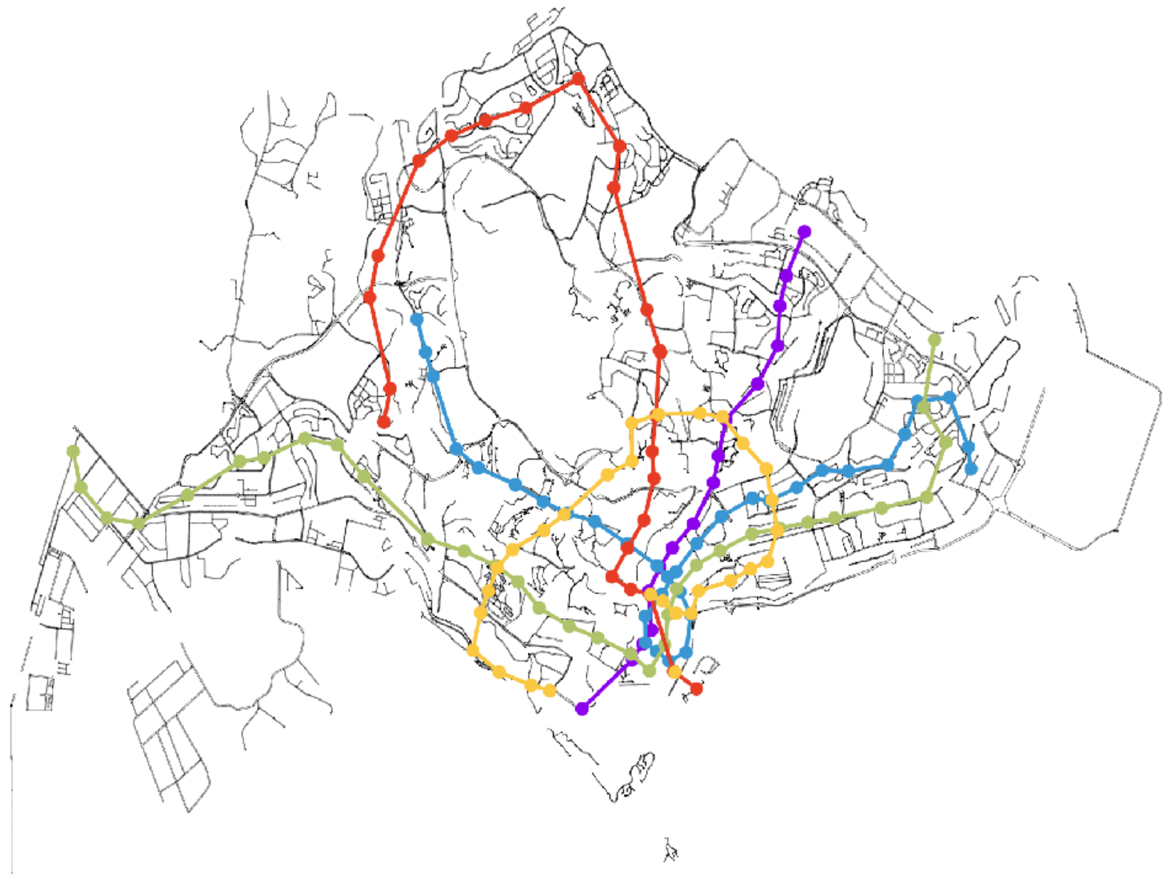

Announcements and Logistics

- No HW due this week 🎉
- Only 1 day to Spring Break! 🎉
- HW 4 budget agent competition results are in!
  - Will announce at the end of class (remind me!)
- HW 6 is out, due April 7 (Thurs after you return from break)
- Reminder to fill out TA feedback form
- Midterm grading: almost done, will return soon
- No office hours after lecture today

Questions?

Last Time

- Started mechanism design without money with one-sided market
- Discussed serial dictatorship:
    - Uniquely Pareto optimal, strategyproof algorithm
- Discussed top trading cycle as a way to run an exchange market (house allocation problem)

Today

- Wrap up discussion of TTC
    - Prove it is DSIC/ strategyproof
    - Stable allocation: no subset can exchange to make everyone better off and no one worse off
- Discuss applications of TTC: kidney exchange & school choice
- Discuss a stable matching algorithm for two-sided matching

Top-Trading Cycle [Gale & Shapley]

- Each agent report their overall preferences in the beginning
- **Step 1**. Each agent (simultaneously) points to its favorite house (among houses remaining)
  - Induces a directed graph G in which every vertex has outdegree 1
  - G must have at least 1 directed cycle (self loops count)
  - Pick directed cycles and make all trades on it (each agent gives its house to the agent that points to it)
  - Delete all agents and houses that were traded in Step 1
- While agents remain, go back to **Step 1**.

TTC is Strategyproof

- **Proof Overview**.
  - An agent's strategy is what preference ordering over n house to submit
  - What edges are formed is pre-determined by rankings submitted
- **Goal**: Fixing everyone else's strategy s_{-i} (their rankings), show that submitting a truthful ranking gives i the best possible item
  - For any preference order i may have
  - And for any ranking of others s_{-i}
- **Claim**. At any round t, pointing truthfully at the favorite remaining house gives the best possible outcome, fixing s_{-i}

TTC is Strategyproof

- **Proof**. Consider any round t. Fix everyone else's rankings s_{-i}
- What are the choices of items that agent i can possibly get at this round?
- Let N_i be i's choice set: of set of items that have a directed path to agent i
  - That is, if i were to point to any item in N_i : a directed cycle could form
- $|N_i|$ cannot go down in round $t+1$ if i is still unmatched
  - If agent j points to i at round t means i is their favorite among remaining items: this does not change as long as i is still unmatched
- Thus, pointing to favorite remaining item gets best possible outcome: truthful reporting is a dominant strategy

TTC is Stable

- A set $S \subseteq \{1,\ldots,n\}$ of agents form a **blocking coalition** for an assignment M if there is a way to reassign items $J = \{M(j) \mid j \in S\}$ within S in to make one of them better off without making anyone else worse off
- **(Stable Allocation)** An allocation is **stable** is there is no blocking coalition
- Stable allocations are also called "core" allocations in the literature
- Stronger condition than Pareto optimality!
  - Implies Pareto optimality when $S = N$ and $J =$ set of all houses
- A **minimal blocking** coalition S is one that does not have another blocking coalition S' such that $S' \subset S$
  - Goal show that some $i \in S$ has no incentive to trade with others in S

TTC is Stable

- **Proof**. Let S be a **minimal** blocking coalition wrt assignment M by TTC
- Let N_j denote the set of agents that get allocated in the jth round in TTC
- Let ℓ be the **first round** in an agent $i \in S$ receives their house
  - i gets their favorite house among those not obtained by $N_1, \ldots, N_{\ell-1}$
  - No member of S among these, why?
    - $N_j \cap S = \emptyset$ for $j = 1,\ldots,\ell - 1$: ℓ is the first round where anyone in S gets their house
- No reallocation within S can make i better off:
  - $S - \{i\}$ is a smaller blocking coalition ∎

Unique Stable Allocation

- **Theorem**. TTC algorithm outputs a **unique** stable allocation.
- **Proof**. Let N_j denote the set of agents who get allocated in round j by TTC
- All agents of N_1 receive their first choice: this must be true in any stable allocation
  - If not, the agents of N_1 can form a coalition for which internal reallocation can make everyone strictly better off
- Similarly, all agents of N_2 receive their top choice outside N_1
  - Given that every stable allocation agrees with TTC for agents in N_1, such an allocation must also agree for agents in N_2
- Inductively we can show that TTC allocation must be the unique stable allocation
  ∎

Summary

- TTC is a computationally efficient, strategyproof, Pareto optimal and stable allocation algorithm for exchange markets
- Given all its nice properties, we don't hear of it as much as lotteries
- How good is the algorithm for practical applications?
    - Paired-kidney donation markets
    - School assignment (even though it doesn't fit the exchange model)
        - We will come back to this after two-sided markets

Application: Kidney Exchange

- Kidney exchange is legal but compensation for organ donation is illegal in US (and every country except Iran)
- In the US in 2013, around 100,000 people were on a waiting list to receive kidneys

Efficient Kidney Exchange: Coincidence of Wants in Markets with Compatibility-Based Preferences

By ALVIN E. ROTH, TAYFUN SÖNMEZ, AND M. UTKU ÜNVER*

Patients needing kidney transplants may have donors who cannot donate to them because of blood or tissue incompatibility. Incompatible patient-donor pairs can exchange donor kidneys with other pairs only when there is a "double coincidence of wants." Developing infrastructure to perform three-way as well as two-way exchanges will have a substantial effect on the number of transplants that can be arranged. Larger than three-way exchanges have less impact on efficiency. In a general model of type-compatible exchanges, the size of the largest exchanges required to achieve efficiency equals the number of types. (JEL C78, I12)

Roth et al's papers studies "hypothetical prices"
and competitive equilibrium in kidney exchange markets

For Sale, By Owner: The Psychology Of Repugnant Transactions

March 4, 2019 · 5:05 PM ET

SHANKAR VEDANTAM · PARTH SHAH · TARA BOYLE

32-Minute Listen · + PLAYLIST

Ja_inter/Getty Images

Application: Kidney Exchange

- Incompatible donor-patient pairs can participate in a larger exchange that use sophisticated matching algorithms
- In 2004, Roth Sonmez and Unver advocated for the TTC for kidney exchange
- Patient, donor pairs: a total ordering over kidneys determined by the likelihood of the successful transplant
  - Longest exchange (2014) involved 35 patients 35 donors
- Biggest dealbreaker: **long trading cycles**
  - Transplants must occur simultaneously due to incentive issues (if surgeries for P1 and D2 happen first, there is a risk that D1 will renege on its offer)

Facilitating More Transplants with Kidney Exchanges and Chains

Why your patients should know about the NKR

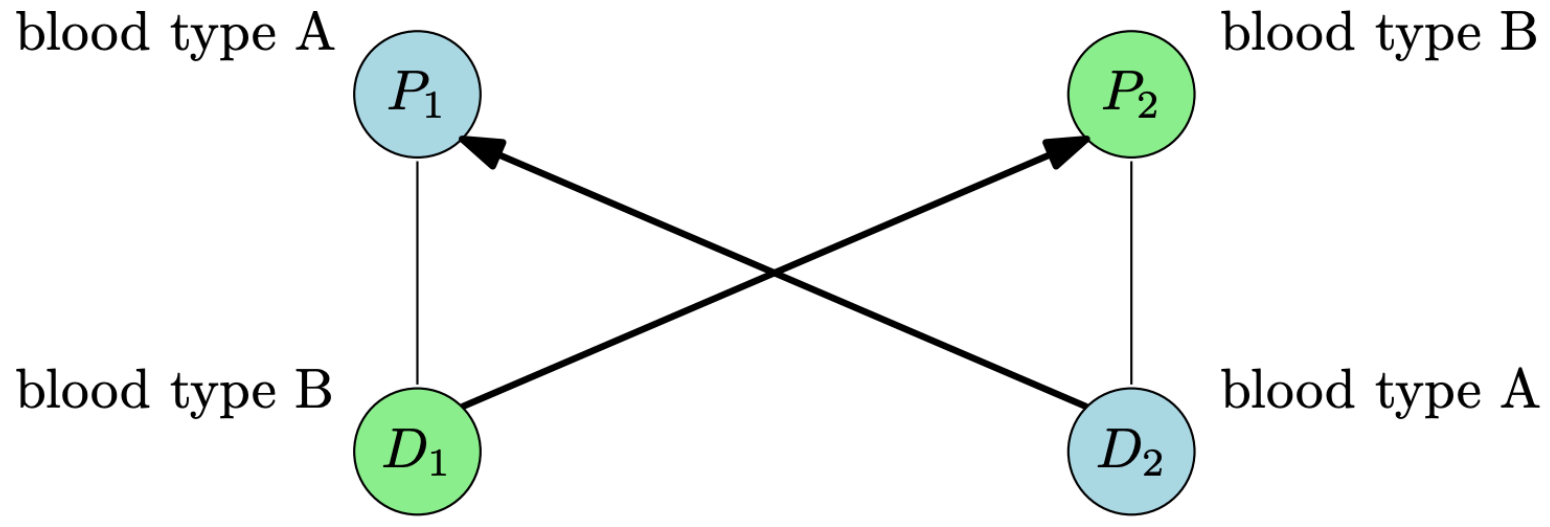

Using Max Matchings

- TCC model requires a total ordering over kidneys
    - In reality patients don't care which kidney they get as long as it is compatible with them
- In subsequent work, Roth et al propose using matchings
    - Matchings lead to 2-way swaps
- Nodes are now patient donor pairs, edges indicate compatibility
- Each agent i has a true edge set E_i and can report any subset of E_i (patients can refuse exchanges for any reason)
- **Goal**. Compute a maximum-cardinality matching
    - Use priority order on nodes for tie breaks: DSIC for individuals
- Full reporting at hospital level is still an issue

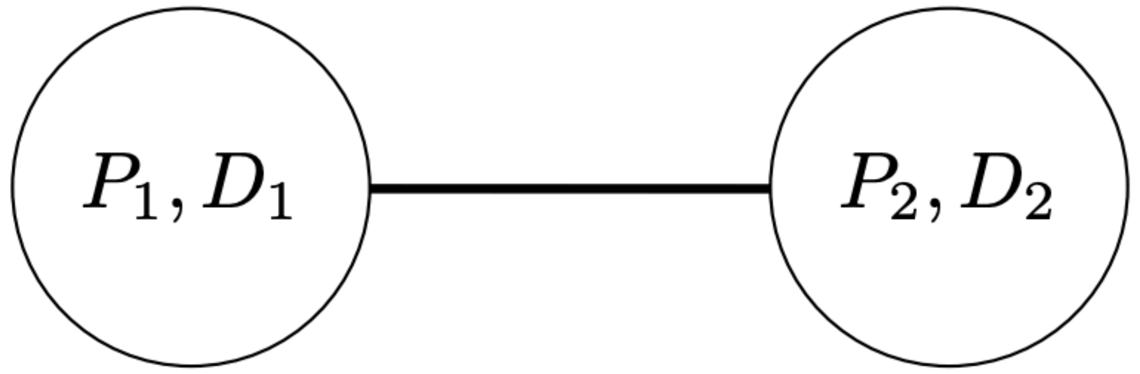

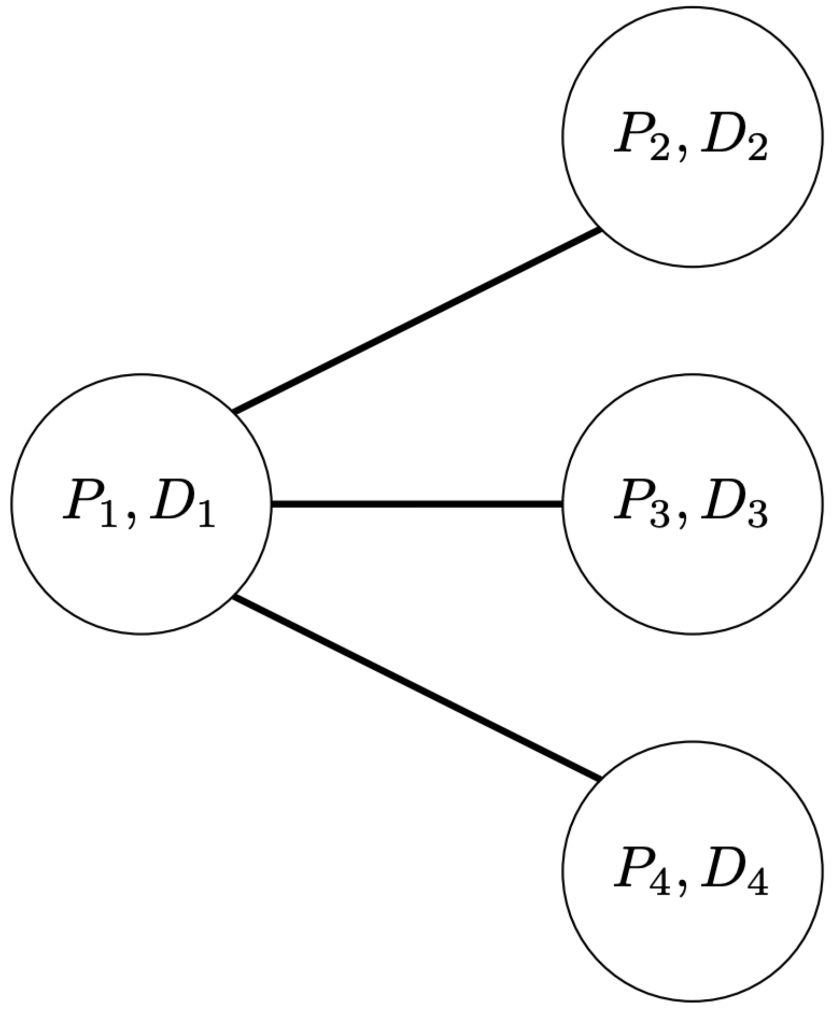

Incentive Challenges

- Need for full reporting at the hospital level
- Objective of individual hospitals: match as many of their patients as possible
- Objective of society: match as many patients as possible

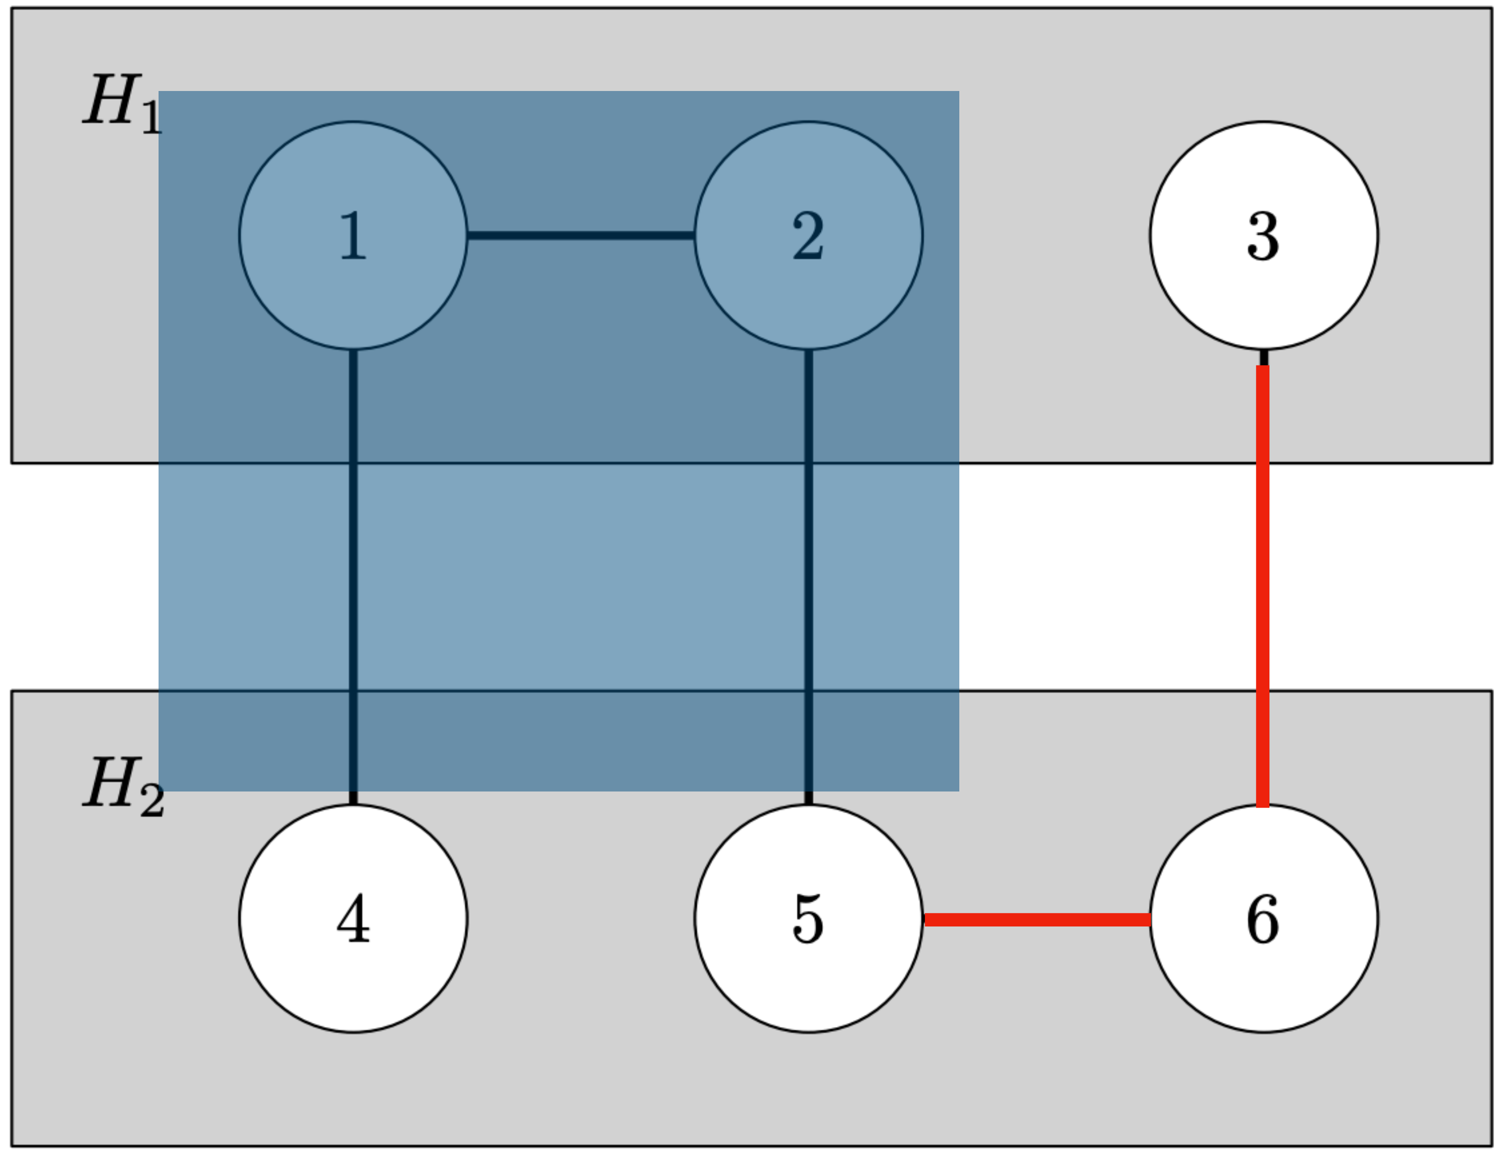

Incentive Challenges

- Need for full reporting at the hospital level
- Objective of individual hospitals: match as many of their patients as possible
- Objective of society: match as many patients as possible
- Need for **approximately optimal DSIC** mechanisms

Incentives of H1 and H2 are at odds: no DSIC mechanism that maximizes cardinality of matching

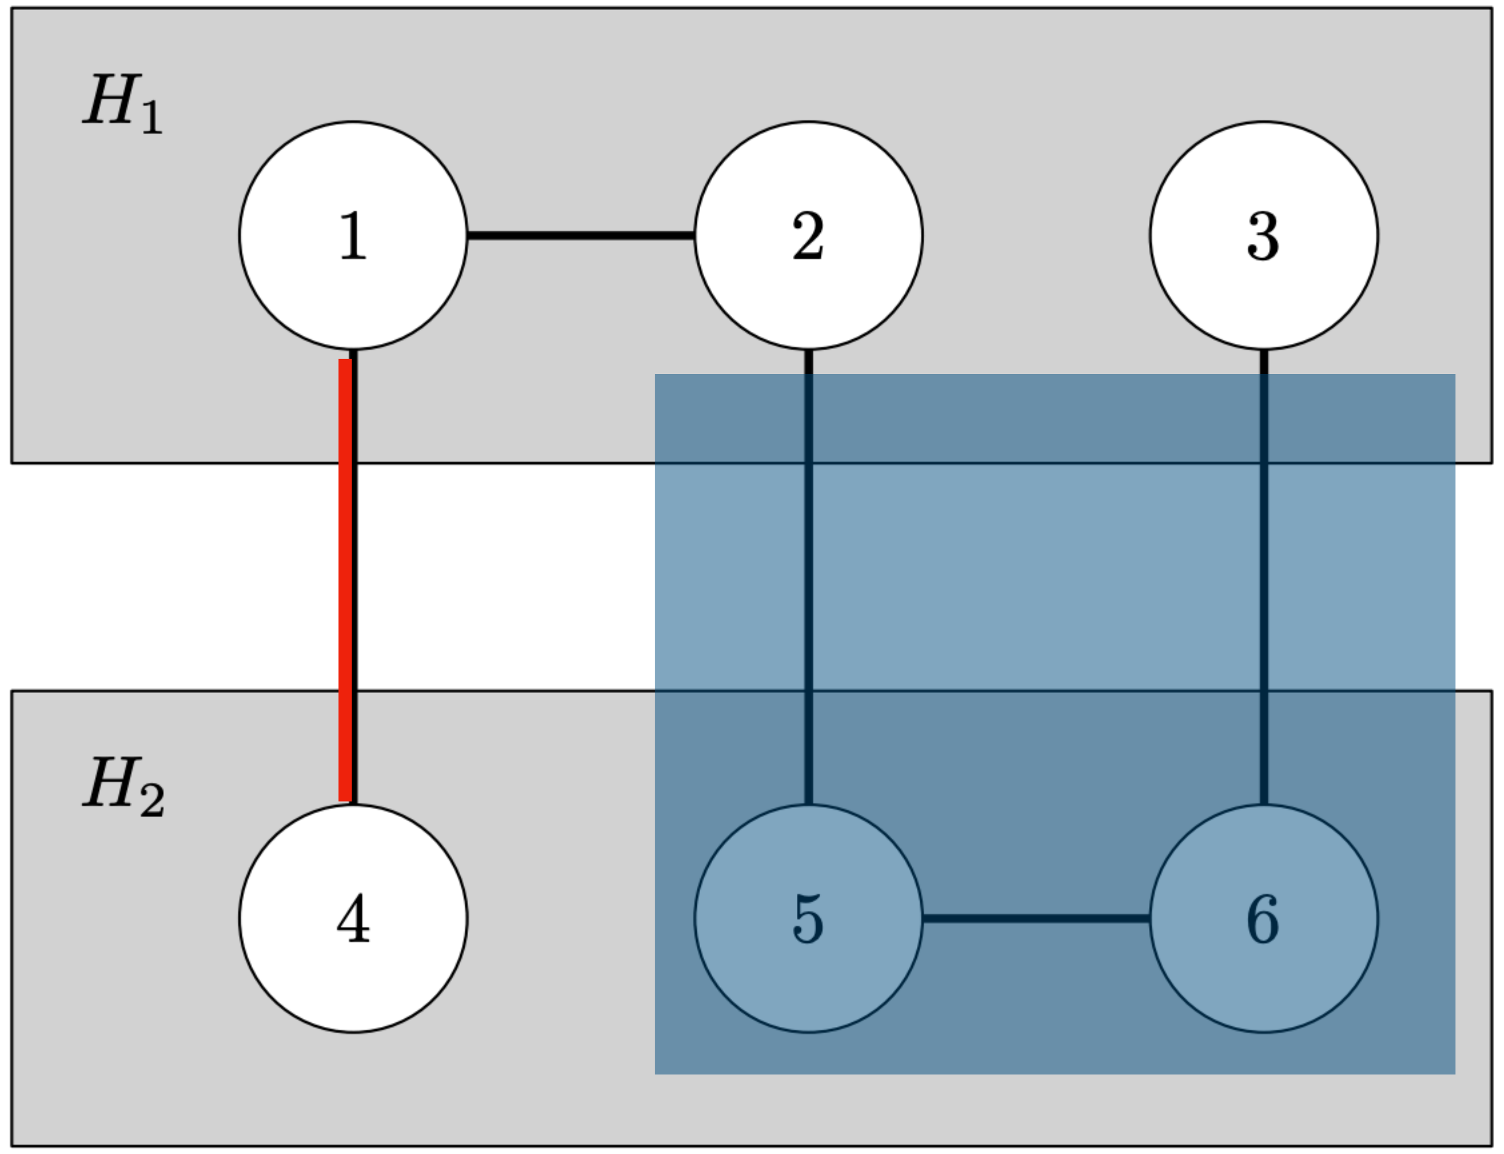

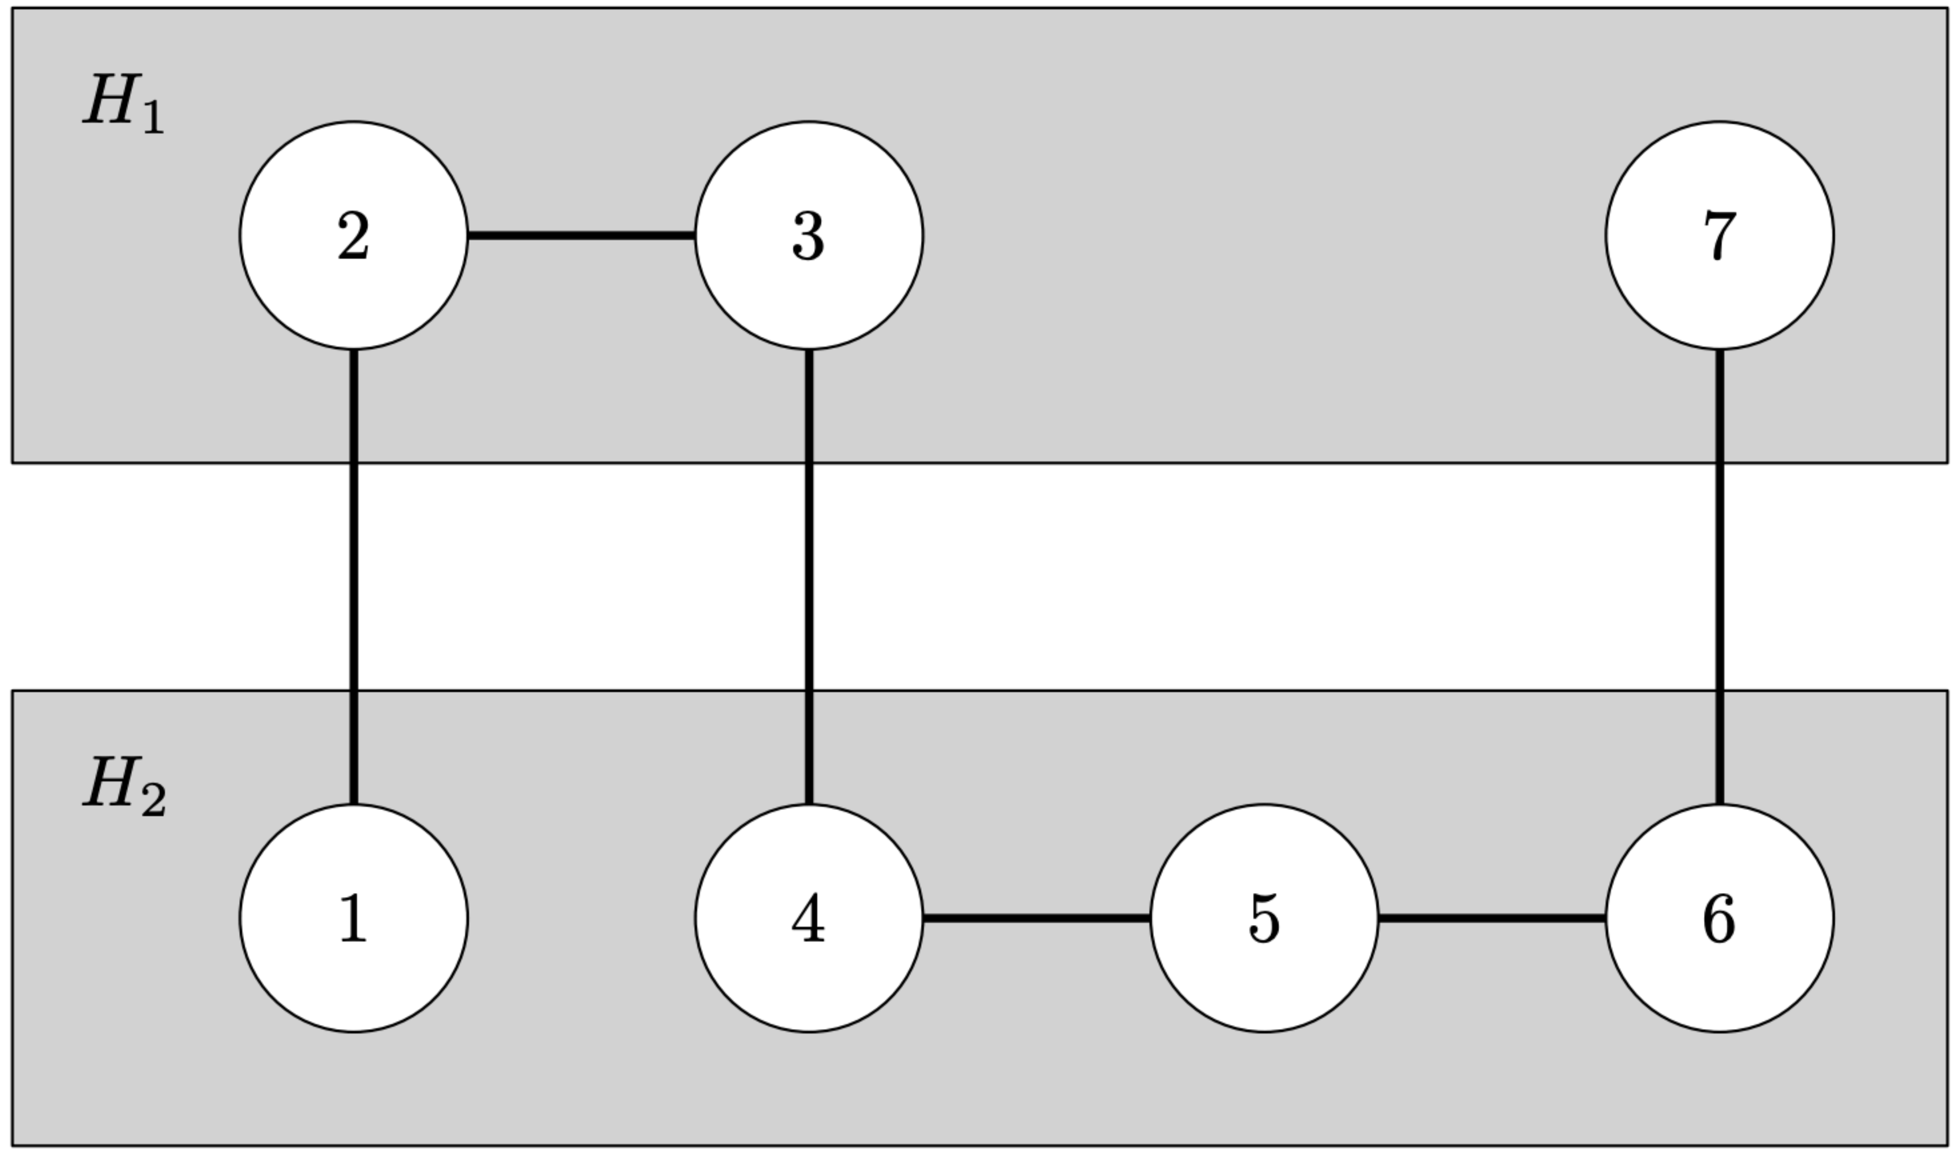

Two-Sided Matching Markets

Two-Sided Markets

- Consider a **two-sided market**:
  - A set H of n hospitals, a set S of n students
  - Each hospital has a complete and strict preference ranking of students
  - Each student has a complete and strict preference ranking of hospitals
- **Goal**: Find a **perfect matching** M (one where each student is matched to exactly one hospital and vice versa) that is **stable** (has no blocking pairs)
- A hospital h and student s form a **blocking pair** (h, s) in a matching M if
  - h prefers s to its current match in M
  - s prefers h to its current match in M

Stable Matching

- Fundamental problem:
  - How to match two sides and avoid opportunistic swapping
- Used to be called **"stable marriage/ dating problem"**
  - But these graphs are not bipartite
- **Centralized** direct-revelation mechanism:
  - Students and hospitals report preferences upfront
  - The algorithm is run based on these reported preferences
- All properties based on the **reported preference profile** and discuss incentive issues later
  - Stability guarantees are with respect to reported preference

Why Centralized? The Story of NRMP

- Medical residencies became widespread in the U.S. in1900s
- From 1900 to 1945, hospitals competed for doctors in an ad hoc and decentralized way
- As time went on, hospitals made offers to doctors earlier and earlier during their tenure at medical school
    - To get ahead of other hospitals
- Led to absurd trends: in 1945, it was standard to extend residency offers to medical students who had just finished their first year (i.e., two years before graduation)
    - Was this good for either side of the market?
- When a market reaches this point, it is said to have **unraveled**
    - Common in law graduates market and **CS job market!**

THE MATCH®
NATIONAL RESIDENT MATCHING PROGRAM®

Why Centralized? The Story of NRMP

- In1945 the situation was so bad that med schools decided they wouldn't release any student info until an appropriate date
- This stopped the unravelling but created other incentives
  - Mad dash to recruit top students
  - Hospitals started making **exploding offers**
- To resolve the chaos caused by exploding offers, hospitals did something radical: moved to a **central clearinghouse**
  - Led to the formation of NRMP
- A committee of students protested the process
  - Changes were made to resolve this
- What we discuss today is what "the Match" is all about

THE MATCH®
NATIONAL RESIDENT MATCHING PROGRAM®

The market for law school graduate is also known for these problems. **Roth** in this article **"Who Gets What And Why"** quotes a law school student who in 2005, on a flight from her 1st interview to 2nd interview, got 3 voicemail messages: the 1st extending an offer from where she just interviewed; the 2nd to urge her to return the call soon; and the 3rd to rescind the offer. Her flight was only 35 mins long!

Why Stability: The Story of NRMP

- Empirical evidence in support
- In UK in the 60s, residency programs decided to move from a decentralized system to a centralized clearinghouse
- The details of the implementation were left to individual regions
- Roth looked at data from 7 regions
    - Two followed a stable implementation; they remain in use today
    - Five regions implemented unstable variants, 3 of which did not survive long (due to poor participation and negotiations outside the system)

Finding Stable Matchings

- **Question**. Does such a stable matching always exist?
    - This does not seem obvious!
- We give a constructive proof of this through the deferred acceptance algorithm
    - Analyzed by Gale and Shapley in 1952 when NRMP was adopted
- Shapley & Roth (who extended his work) were awarded the **2012 Nobel Prize** (Gale did not share the prize, because he died in 2008.)
- We revisit Gale and Shapley's deferred acceptance algorithm (from CS256)

Deferred Acceptance Algorithm

Proceed in rounds until all hospitals matched. In each round,

- Each free hospital offers to its top choice among candidates it hasn't offered yet
- Each student *retains but defers accepting* **top offer**, rejects others (if a student receives a better offer than currently retained, they reject current and retain new offer: **trade up**)

	1st	**2nd**	**3rd**
MA	Aamir	Chris	Beth
NH	Aamir	Beth	Chris
OH	Chris	Beth	Aamir

	1st	**2nd**	**3rd**
Aamir	OH	NH	MA
Beth	MA	OH	NH
Chris	MA	NH	OH

Deferred Acceptance Algorithm

Proceed in rounds until all hospitals matched. In each round,

- Each free hospital offers to its top choice among candidates it hasn't offered yet
- Each student *retains but defers accepting* **top offer**, rejects others (if a student receives a better offer than currently retained, they reject current and retain new offer: **trade up**)

	1st	**2nd**	**3rd**
MA	Aamir	Chris	Beth
NH	Aamir	Beth	Chris
OH	Chris	Beth	Aamir

	1st	**2nd**	**3rd**
Aamir	OH	NH	MA
Beth	MA	OH	NH
Chris	MA	NH	OH

Deferred Acceptance Algorithm

Proceed in rounds until all hospitals matched. In each round,

- Each free hospital offers to its top choice among candidates it hasn't offered yet
- Each student *retains but defers accepting* **top offer**, rejects others (if a student receives a better offer than currently retained, they reject current and retain new offer: **trade up**)

	1st	2nd	3rd
MA	Aamir	Chris	Beth
NH	Aamir	Beth	Chris
OH	Chris	Beth	Aamir

	1st	2nd	3rd
Aamir	OH	NH	MA
Beth	MA	OH	NH
Chris	MA	NH	OH

Deferred Acceptance Algorithm

Proceed in rounds until all hospitals matched. In each round,

- Each free hospital offers to its top choice among candidates it hasn't offered yet
- Each student *retains but defers accepting* **top offer**, rejects others (if a student receives a better offer than currently retained, they reject current and retain new offer: **trade up**)

	1st	2nd	3rd
MA	Aamir	Chris	Beth
NH	Aamir	Beth	Chris
OH	Chris	Beth	Aamir

	1st	2nd	3rd
Aamir	OH	NH	MA
Beth	MA	OH	NH
Chris	MA	NH	OH

Deferred Acceptance Algorithm

Proceed in rounds until all hospitals matched. In each round,

- Each free hospital offers to its top choice among candidates it hasn't offered yet
- Each student *retains but defers accepting* **top offer**, rejects others (if a student receives a better offer than currently retained, they reject current and retain new offer: **trade up**)

	1st	2nd	3rd
MA	Aamir	Chris	Beth
NH	Aamir	Beth	Chris
OH	Chris	Beth	Aamir

	1st	2nd	3rd
Aamir	OH	NH	MA
Beth	MA	OH	NH
Chris	MA	NH	OH

Deferred Acceptance Algorithm

Proceed in rounds until all hospitals matched. In each round,

- Each free hospital offers to its top choice among candidates it hasn't offered yet
- Each student *retains but defers accepting* **top offer**, rejects others (if a student receives a better offer than currently retained, they reject current and retain new offer: **trade up**)

	1st	2nd	3rd
MA	Aamir	Chris	Beth
NH	Aamir	Beth	Chris
OH	Chris	Beth	Aamir

	1st	2nd	3rd
Aamir	OH	NH	MA
Beth	MA	OH	NH
Chris	MA	NH	OH

Stable matching!

Gale-Shapely Algorithm

GALE–SHAPLEY (*preference lists for hospitals and students*)

INITIALIZE M to empty matching.

WHILE (some hospital h is unmatched and hasn't proposed to every student)

  $s \leftarrow$ first student on h's list to whom h has not yet proposed.

  IF (s is unmatched)

    Add h–s to matching M.

  ELSE IF (s prefers h to current partner h')

    Replace h'–s with h–s in matching M.

  ELSE

    s rejects h.

RETURN stable matching M.

Analyzing the Algorithm

- **Running time**. In 256, we analyze this algorithm to be linear time
  - $O(n^2)$ running time, input size $O(n^2)$
  - Each hospital makes an offer to each student at most once, so the algorithm makes at most $O(n^2)$ iterations
  - Each iteration can be implemented in $O(1)$ time
- **Correctness**. Does it matching everyone (produce a perfect matching?)
  - Once a student receives an offer, always has a tentative match
  - In other words, if a student never receives an offer, means hospitals have not exhausted their preference list
- **Stability**. Does it produce a **stable matching**?

Stable Matching Proof

Lemma. The Gale Shapely Algorithm produces a stable matching.

Proof. (By contradiction) Let M be the resulting matching. Suppose $\exists (h, s)$ such that $(h, s'), (h', s) \in M$ and

- h prefers s over s' and s prefers h over h'

Thus h must have offered to s before s'

- Either s broke the match to h at some point for some h'', or s already had a match h'' that s preferred over h

But students always trade up, so s must prefer final match h' over h'', which they prefer over h. $(\Rightarrow\Leftarrow)$ ∎

Stable Matching Properties

- The deferred-acceptance algorithm does not specify the order in which the hospitals should make offers
- Do all orders produce the same unique matching?
- Given an input instance, there may be several stable matchings.
- **A Different Question.** Does Gale-Shapely produce the "best matching" for hospitals or students?
- Turns out hospital-proposing algorithm produces a unique matching that is **hospital optimal** and **student pessimal**
  - Matches hospital to "best achievable" student and student to "worst-achievable" hospital among all stable matchings

Best Achievable Partner

Let I be an *instance* of the stable marriage problem

- A student $s \in S$ is an **achievable partner** for hospital $h \in H$, if (h, s) *is* part of *some* stable matching of I.
  - We call the pair (h, s) an **achievable pair**
- For hospital $h \in H$, let **best(h)** denote the most preferred achievable partner of h
- **Lemma**. $M^* = \{(h, \text{best}(h)) \mid h \in H\}$ is the unique output of the hospital-proposing deferred-acceptance algorithm.
  - This is true regardless of the order in which different hospitals make offers

Hospital-Optimal Matching

- **Lemma**. $M^* = \{(h, \text{best}(h)) \mid h \in H\}$ is the unique output of the hospital-proposing deferred-acceptance algorithm.
- **Proof** (By Contradiction)**.** Let h be the **first hospital** to be rejected by $s^* = \text{best}(h)$
  - s^* instead holds on to offer from some h'
- s^* must be the best achievable partner for h', why?
  - If not h' has already been rejected by best(h'), violates condition that h is the first such hospital
- Let M be a stable matching s.t. $(h, s^*) \in M$
- **Claim**. (h', s^*) is a blocking pair for matching M, why?
  - s prefers h' to h, and h' prefers s^* to whoever they are matched to in M ($\Rightarrow\Leftarrow$) ∎

Takeaways

- The outcome of hospital-offering deferred acceptance is hospital-optimal, among all stable matching
    - There is no tradeoff to make in terms of who offers first!
- What about the accepting side?
    - The outcome of the hospital-offering deferred acceptance is **students-pessimal**, among all stable matchings
    - In particular, students get matched to their worst-achievable partner among all stable matchings
- **Incentive considerations**. Which side of the market has an incentive to misreport their preferences?
- Can misreports be beneficial? Is the mechanism strategyproof?

Stability and Strategyproofness

- **Lemma.** Truthful reporting is a weakly dominant strategy for hospitals in the hospital-proposing deferred acceptance mechanism
  - While intuitive, this is surprisingly annoying to prove
    - See Theorem 10.6.18 in http://www.masfoundations.org/mas.pdf
  - Stability is only wrt to **reported preferences**, if someone misreports then stability is defined with respect to reported preferences only
- Is truthful reporting a dominant strategy if you are on the other-side of the market: for students in a hospital-proposing DA?
  - This is not too difficult to see
  - Let’s take an example

Misreports from Students

- Consider the following truthful preference profile

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

	1st	2nd	3rd
Aamir	MA	OH	NH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

- Consider the following truthful preference profile

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

	1st	2nd	3rd
Aamir	MA	OH	NH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

- Consider the following truthful preference profile

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

	1st	2nd	3rd
Aamir	MA	OH	NH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

- Consider the following truthful preference profile

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

	1st	2nd	3rd
Aamir	MA	OH	NH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

- Consider the following truthful preference profile
  - Produces the following stable matching:
  - (MA, Beth), (NH, Chris), (OH, Aamir)

	1st	**2nd**	**3rd**
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

	1st	**2nd**	**3rd**
Aamir	MA	OH	NH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

- **Class exercise**. Can one of the students misreport their preferences to end up with a better match?
  - Does it every make sense to misreport about your top choice?
  - What about lower order misreports?

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

	1st	2nd	3rd
Aamir	MA	OH	NH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

- Suppose Aamir misreports (swaps NH and OH)

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

New Preference Profile

	1st	2nd	3rd
Aamir	MA	NH	OH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

- Suppose Aamir misreports (swaps NH and OH)

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

New Preference Profile

	1st	2nd	3rd
Aamir	MA	NH	OH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

- Suppose Aamir misreports (swaps NH and OH)

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

New Preference Profile

	1st	2nd	3rd
Aamir	MA	NH	OH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

- Suppose Aamir misreports (swaps NH and OH)

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

New Preference Profile

	1st	2nd	3rd
Aamir	MA	NH	OH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

- Suppose Aamir misreports (swaps NH and OH)

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

New Preference Profile

	1st	2nd	3rd
Aamir	MA	NH	OH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

- Suppose Aamir misreports (swaps NH and OH)

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

New Preference Profile

	1st	2nd	3rd
Aamir	MA	NH	OH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

- Suppose Aamir misreports (swaps NH and OH)

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

New Preference Profile

	1st	2nd	3rd
Aamir	MA	NH	OH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

- Suppose Aamir misreports (swaps NH and OH)

	1st	2nd	3rd
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

New Preference Profile

	1st	2nd	3rd
Aamir	MA	NH	OH
Beth	OH	MA	NH
Chris	MA	OH	NH

Misreports from Students

DA is not strategyproof (the receiving side can misreport and achieve a better match)

- Suppose Aamir misreports (swaps NH and OH)
- New matching: (MA, Aamir), (NH, Chris), (OH, Beth)
- Aamir improved from NH to top choice MA!

	1st	**2nd**	**3rd**
MA	Beth	Aamir	Chris
NH	Aamir	Chris	Beth
OH	Aamir	Beth	Chris

New Preference Profile

	1st	**2nd**	**3rd**
Aamir	MA	NH	OH
Beth	OH	MA	NH
Chris	MA	OH	NH

Can't Have Both

- Can there be a mechanism that is both strategy proof and stable?
  - Unfortunately, no
- **Theorem**. No mechanism for two-sided matching is both stable and strategyproof.
  - Proof developed in Homework 6
- Many interesting questions:
  - How much information is needed to find a useful manipulation?
  - What is the optimal manipulation cheating strategy
- Empirically manipulations do not play a large role
  - If not many stable partners, can't gain much

The Match and its Evolution

- **NRMP Revisited**. The original 1952 implementation of the DA algorithm was the **hospital-optimal version**
- Students protested that the match was favoring hospitals

MONDAY, OCTOBER 22, 1951. The New York Times MONDAY, OCTOBER 22, 1951. 25 L+

MEDICAL SENIORS HIT INTERNE PLAN

Delegates of 44 Schools Meet Here to Protest Selection by 'Matching Machine'

TEACHERS PRAISE SYSTEM

They Argue That It Bars Unfair Recruiting—Students Insist on Choosing Their Hospitals

Delegates representing seniors in nearly all of the country's leading medical schools met here yesterday to express overwhelming opposition to a proposed mathematically contrived plan to place medical students in hospitals as internes.

They indicated that a great majority of their classmates preferred the present system whereby the country's hospitals, which have 10,000 internships, scramble for the best of a year's 6,000 medical graduates.

The meeting was held at Bard Hall of the College of Physicians and Surgeons of Columbia University, but it was made clear that the university was not its sponsor.

Seventy students attended the meeting at which forty-four colleges were represented either by delegates or through communications giving the opinion of the medical school's seniors on "the matching plan for internship" organized by the National Interassociation Committee on Internship about two years ago.

A prospectus of the plan of leading hospitals calls it "the accepted procedure for 1951-52" in determining which medical graduate shall go to what hospital to complete his medical education. Medical men said yesterday that its success or failure would have much to do with the chances of a medical student properly to complete his education during a wartime emergency.

By Winston Churchill: *The Second World War*

Volume V—Closing the Ring

INSTALLMENT 15—TEHERAN: CONCLUSIONS

Book II—Teheran to Rome

ROOSEVELT, STALIN, CHURCHILL AGREE ON PLANS FOR WAR ON GERMANY IN TALKS AT TEHERAN; 1,500 MORE TONS OF BOMBS DROPPED ON BERLIN

AGREEMENT: The news on the morning of Dec. 4, 1943.

NEED FOR TEACHERS EXPECTED TO GROW

1,200 More a Year Required in State, Board Officials Say —Triple Sessions Feared

LACK IS WORST IN GRADES

Special Subjects Also Suffer—Syracuse Parley Cites High Birth Rate, Low Salaries

By LEONARD BUDER

The Match and its Evolution

- A new algorithm was adopted in 1997
  - Primary motivated was to give couples the option to get placed in geographically nearby programs
  - But in addition was made student-proposing
- Changes incentives for hospitals, but did it make a difference?
- Empirically, at least for the datasets arising in NRMP, less than 1% of the hospitals could have benefited by misreporting

Stable Matching Summary

- When choosing a matching in a two-sided market stability is important to ensure participants don't circumvent the algorithm
- When choosing between stable outcomes, you have to make trade-offs between the two sides of the market
  - Should it favor students or hospitals?
- Lots of generalizations:
  - Incomplete preferences with ties
  - Stable roommates problem
  - Many-to-one stable matchings
  - Approximately stable matchings

Classic Problem

Marriage, Honesty, and Stability

Nicole Immorlica* Mohammad Mahdian*

Abstract

Many centralized two-sided markets form a matching between participants by running a stable marriage algorithm. It is a well-known fact that no matching mechanism based on a stable marriage algorithm can guarantee truthfulness as a dominant strategy for participants. However, as we will show in this paper, in a probabilistic setting where the preference lists of one side of the market are composed of only a constant (independent of the the size of the market) number of entries, each drawn from an *arbitrary* distribution, the number of participants that have more than one stable partner is vanishingly small. This proves (and generalizes) a conjecture of Roth and Peranson [23]. As a corollary of this result, we show that, with high probability, the truthful strategy is the best response for a given player when the other players are truthful. We also analyze equilibria of the deferred acceptance stable marriage game. We show that the game with complete information has an equilibrium in which a $(1-o(1))$ fraction of the strategies are truthful in expectation. In the more realistic setting of a game of incomplete information, we will show that the set of truthful strategies form a $(1+o(1))$-approximate Bayesian-Nash equilibrium. Our results have implications in many practical settings and were inspired by the work of Roth and Peranson [23] on the National Residency Matching Program.

Stable Husbands

Donald E. Knuth, Rajeev Motwani, and Boris Pittel
Computer Science Department, Stanford University

Abstract. Suppose n boys and n girls rank each other at random. We show that any particular girl has at least $(\frac{1}{2}-\epsilon)\ln n$ and at most $(1+\epsilon)\ln n$ different husbands in the set of all Gale/Shapley stable matchings defined by these rankings, with probability approaching 1 as $n \to \infty$, if ϵ is any positive constant. The proof emphasizes general methods that appear to be useful for the analysis of many other combinatorial algorithms.

ON LIKELY SOLUTIONS OF A STABLE MARRIAGE PROBLEM[1]

BY BORIS PITTEL
Ohio State University

To the memory of Mikhail L'vovich Tsetlin

An (n men–n women) stable marriage problem is studied under the assumption that the individual preferences for a marriage partner are uniformly random and mutually independent. We show that the total number of stable matchings (marriages) is at least $(n/\log n)^{1/2}$ with high probability (whp) as $n \to \infty$ and also that the total number of stable marriage partners of each woman (man) is asymptotically normal with mean and variance close to $\log n$. It is proved that in the male (female) optimal stable marriage the largest rank of a wife (husband) is whp of order $\log^2 n$, while the largest rank of a husband (wife) is asymptotic to n. Earlier, we proved that for either of these extreme matchings the total rank is whp close to $n^2/\log n$. Now, we are able to establish a whp existence of an egalitarian marriage for which the total rank is close to $2n^{3/2}$ and the largest rank of a partner is of order $n^{1/2}\log n$. Quite unexpectedly, the stable matchings obey, statistically, a "law of hyperbola": namely, whp the product of the sum of husbands' ranks and the sum of wives' ranks in a stable matching turns out to be asymptotic to n^3, uniformly over all stable marriages. The key elements of the proofs are extensions of the McVitie–Wilson proposal algorithm and of Knuth's integral formula for the probability that a given matching is stable, and also a notion of rotations due to Irving. Methods developed in this paper may, in our opinion, be found useful in probabilistic analysis of other combinatorial algorithms.

Stable Matching Research

Jan 2022

The Short-Side Advantage in Random Matching Markets

Linda Cai * Clayton Thomas †

Two-sided matching markets with correlated random preferences have few stable pairs

HUGO GIMBERT, Université de Bordeaux, LaBRI, CNRS, F-33400 Talence, France
CLAIRE MATHIEU, Université de Paris, IRIF, CNRS, F-75013 Paris, France
SIMON MAURAS, Université de Paris, IRIF, CNRS, F-75013 Paris, France

2021

2021

Tiered Random Matching Markets: Rank Is Proportional to Popularity

Itai Ashlagi
Department of Management Science and Engineering, Stanford University, CA, USA
iashlagi@stanford.edu

Mark Braverman
Department of Computer Science, Princeton University, NJ, USA
mbraverm@cs.princeton.edu

Amin Saberi
Department of Management Science and Engineering, Stanford University, CA, USA
saberi@stanford.edu

Clayton Thomas
Department of Computer Science, Princeton University, NJ, USA
claytont@cs.princeton.edu

Geng Zhao
Department of Computer Science, Stanford University, CA, USA
gengz@stanford.edu

In which matching markets does the short side enjoy an advantage?*

Yash Kanoria† Seungki Min‡ Pengyu Qian§

2021

2018

Stable matching mechanisms are not obviously strategy-proof ☆

Itai Ashlagi [a,*], Yannai A. Gonczarowski [b,c,*]

[a] *Management Science & Engineering, Stanford University, United States of America*
[b] *Einstein Institute of Mathematics, Rachel & Selim Benin School of Computer Science & Engineering, and the Federmann Center for the Study of Rationality, The Hebrew University of Jerusalem, Israel*
[c] *Microsoft Research, Israel*

Received 30 July 2017; final version received 14 March 2018; accepted 9 July 2018

2022

On Fairness and Stability in Two-Sided Matchings

Gili Karni ✉
Weizmann Institute of Science, Rehovot, Israel

Guy N. Rothblum ✉
Weizmann Institute of Science, Rehovot, Israel

Gal Yona ✉
Weizmann Institute of Science, Rehovot, Israel

2019

A $(1+1/e)$-Approximation Algorithm for Maximum Stable Matching with One-Sided Ties and Incomplete Lists*

Chi-Kit Lam C. Gregory Plaxton