An Introduction to Matching, Assignment and Applications

Higher School of Economics

St. Petersburg

Onur Kesten

Carnegie Mellon University

2012 Nobel Laureates

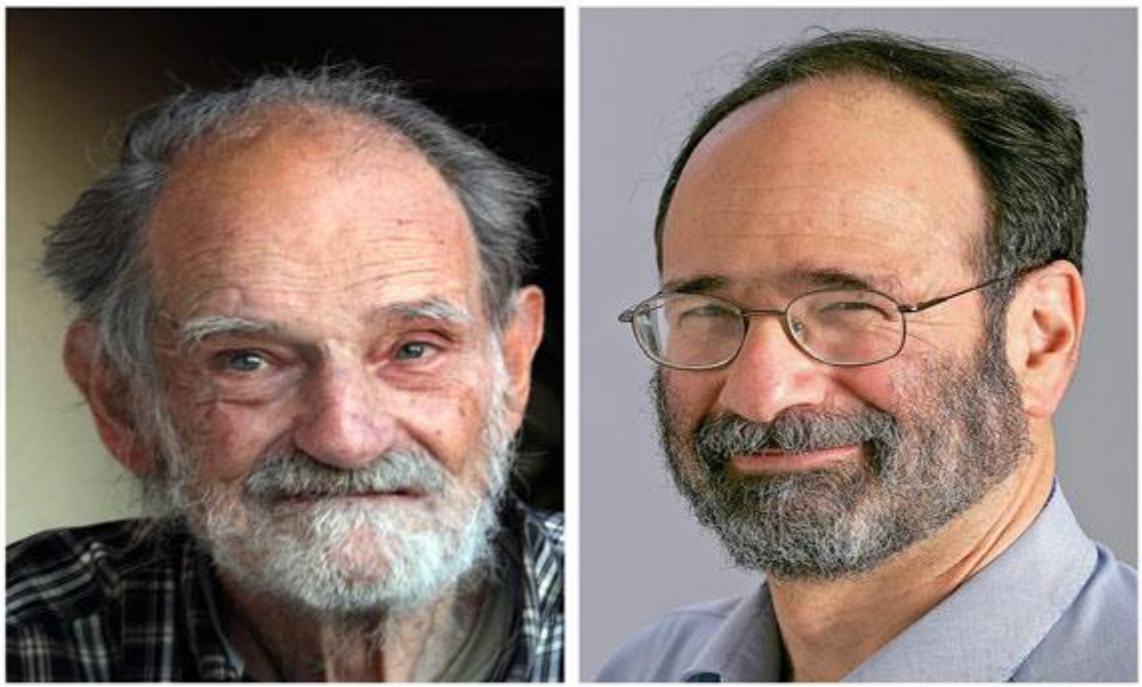

“for the theory of stable allocations and the practice of market design”

Matching & Assignment

- Statement of the problem
  - Two sides of the market to be matched.
  - Participants on both sides care about to whom they are matched.
  - Money can’t be used to determine the matching.

- Examples
  - Marriage & dating markets
  - Kidney exchange
  - School choice programs
  - College admissions
  - Child Adoption
  - Medical residencies
  - Sports drafts
  - Cadet-Branch matching
  - Job assignments in firms
  - Course Allocation

History of NRMP

- History turns out to be illuminating
  - Into the 1930s, medical students found residencies through a completely decentralized process.
  - But there were problems: students and hospitals made contracts earlier and earlier, eventually in the second year of med school!
- Hospitals decided to change the system by adopting a centralized clearinghouse.
  - National Resident Matching Program (NRMP) adopted, after various adjustments in 1952.
  - System has persisted, though with some modification in late 1990s to handle couples and some recent debate about salaries.
- Why might a centralized clearinghouse be useful? And how might the design of the clearinghouse matter?

From Theory to Practice

- Study of matching started as “pure” theory: first by David Gale and Lloyd Shapley (1962) who introduced DA algorithm, then others.

- In 1984, Al Roth made a surprising discovery:
  - Since the 1950s, US hospitals have used a clearinghouse to assign graduating medical students to residencies.
  - Students apply and interview at hospitals in the fall, then students and hospitals submit rank-order preferences in February.
  - A computer algorithm is used to assign students to hospitals, and matches are all revealed on a single day: “match day”.
  - Roth realized that the doctors have independently discovered and were using exactly the Gale and Shapley DA algorithm!

Matching Theory

Gale-Shapley
(also see Sections 2.1-2.3, 4.1-4.4 of Roth-Sotomayor)

Marriage Model

- **Participants**
  - Set of men *M*, with typical man $m \in M$.
  - Set of women *W*, with typical woman $w \in W$.
  - One-to-one matching: each man can be matched to one woman, and vice-versa.
- **Preferences**
  - Each man has strict (complete, transitive, anti-symmetric) preferences over women and being unmatched, and vice versa.
  - A woman *w* is ***acceptable*** to man *m* if *m* prefers *w* to being unmatched.

Matching

- A ***matching*** is a set of pairs *(m,w)* such that each individual has one partner.
  - If the match includes *(m,m)* then m is unmatched.

- What are the desirable properties of a matching?

- How do we get a desirable matching?

Example 1

- Two men (*A)lan ,(B)ob* and two women (*C)athy,(D)ebra*

- *Alan* prefers *Cathy* to *Debra (P(A):C>D>A)*
- *Bob* prefers *Debra* to *Cathy (P(B):D>B>C)*
- *Cathy* prefers *Alan* to *Bob (P(C):A>B>C)*
- *Debra* prefers *Bob* to *Alan (P(D):B>A>D)*

- Possible matching: *(Alan, Debra)* and *(Bob, Cathy)?*

Example 2

- Two men (*A)lan ,(B)ob* and two women (*C)athy,(D)ebra*

- *Alan* prefers *Cathy* to *Debra (P(A):C>D>A)*
- *Bob* prefers *Debra* to *Cathy (P(B):D>C>B)*
- *Cathy* prefers *Alan* to *Bob (P(C):A>B>C)*
- *Debra* prefers *Bob* to *Alan (P(D):B>A>D)*

- Possible matching: *(Alan, Debra)* and *(Bob, Cathy)?*

Stability

- A matching is ***stable*** if
  - Every individual is matched with an acceptable partner or unmatched (***individual rationality***).
  - There is no man-woman pair, each of whom would prefer to match with each other rather than their assigned partner.
- If such a pair exists, they are a ***blocking pair.***
- A matching is ***unstable*** if it is not stable.

Example 2 continued

- Two men (*A)lan ,(B)ob* and two women (*C)athy,(D)ebra*

- *Alan* prefers *Cathy* to *Debra (P(A):C>D>A)*
- *Bob* prefers *Debra* to *Cathy (P(B):D>C>B)*
- *Cathy* prefers *Alan* to *Bob (P(C):A>B>C)*
- *Debra* prefers *Bob* to *Alan (P(D):B>A> D)*

- Unique stable matching: *(Alan, Cathy)* and *(Bob, Debra)*

Example 3

- Two men (*A)lan ,(B)ob* and two women (*C)athy,(D)ebra*

- *Alan* prefers *Cathy* to *Debra (P(A):C>D>A)*
- *Bob* prefers *Debra* to *Cathy (P(B):D>C>B)*
- *Cathy* prefers *Bob* to *Alan (P(C):B>A>C)*
- *Debra* prefers *Alan* to *Bob (P(D):A>B>D)*

- Two stable matchings *{(A,C),(B,D)}* and *{(A,D),(B,C)}*
- First match is better for the men, second for the women.

- *Is there always a stable matching? How to find one?*

Deferred Acceptance

- Men and women rank all potential partners
- **Algorithm**
  - (Step 1) Each man proposes to highest woman on his list. Women tentatively accept based on their preferred offer, and reject other offers, or all if none are acceptable.
  - (Step k) Each rejected man proposes to the next acceptable woman. Each woman considers the tentatively accepted offer (from step k-1) and the new offers, tentatively accepts the best acceptable offer and rejects the rest.
- This is the "man-proposing" version of the algorithm; there is also a "woman proposing" version.

DA in pictures

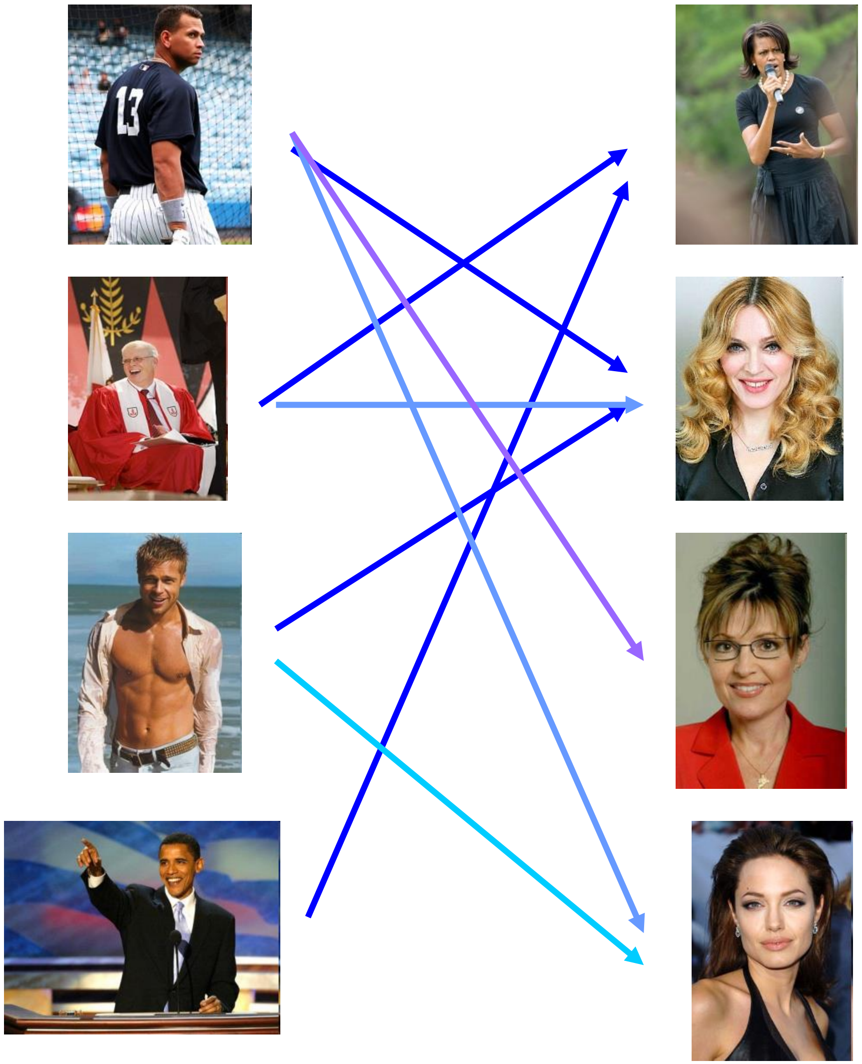

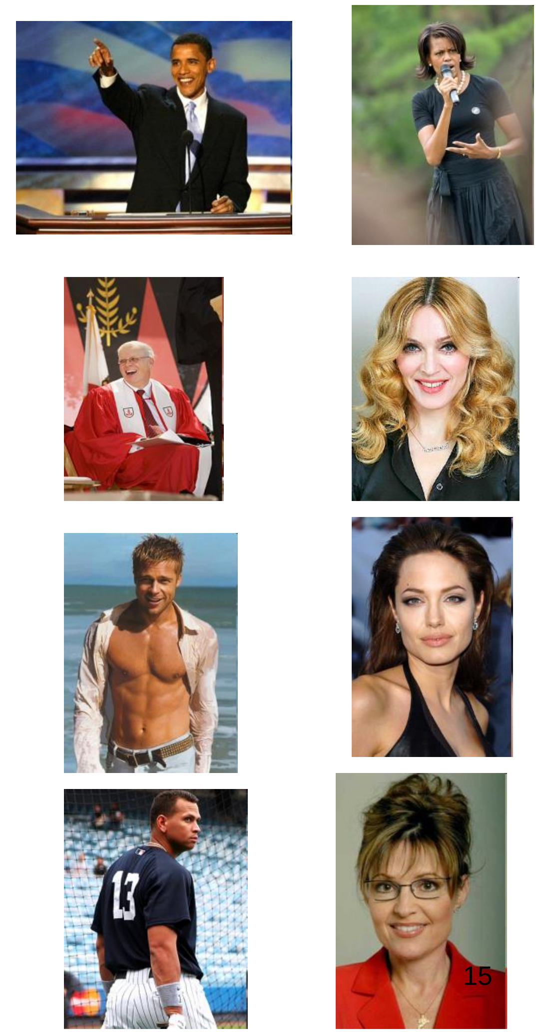

Stable matchings exist

Theorem. The outcome of the DA algorithm is a stable matching (so a stable matching exists).

Proof.

- Algorithm must end in a finite number of rounds.
- Suppose *m, w* are matched, but *m* prefers *w'*.
  - At some point, *m* proposed to *w'* and was rejected.
  - At that point, *w'* preferred her tentative match to *m*.
  - As algorithm goes forward, *w'* can only do better.
  - So *w'* prefers her final match to *m.*
- Therefore, there are *NO BLOCKING PAIRS.*

Aside: the roommate problem

- Suppose a group of students are to be matched to roommates, two in each room.
- Example with four students
  - A prefers B>C>D
  - B prefers C>A>D
  - C prefers A>B>D
  - No stable matching exists: whoever is paired with D wants to change and can find a willing partner.
- So stability in matching markets is not a given, even if each match involves just two people.

Why stability?

- Stability seems to explain at least in part why some mechanisms have stayed in use.
  - If a market results in stable outcomes, there is no incentive for re-contracting.
  - We will see later in the class that the clearinghouses using the Deferred Acceptance algorithm have fared pretty well.
  - Other clearinghouses that use unstable matching mechanisms seem to have failed more often.
- But does one need an organized clearinghouse?

Decentralized market

- What if there is no clearinghouse?
  - Men make offers to women
  - Women consider their offers, perhaps some accept and some reject.
  - Men make further offers, etc..
- What kind of problems can arise?
  - Maybe *w* holds *m*'s offer for a long time, and then rejects it, but only after market has cleared.
  - Maybe *m* makes exploding offer to *w* and she has to decide before knowing her other options.
- No guarantee the market will be “orderly”…

Example: clinical psychology

- Clinical psychologists are employed as interns after they complete their doctoral degrees. About 500 sites offer 2,000 positions each year. Clears with one day market.
- On selection day, market opens at 9 am, closes at 4 pm.
  - While market is open, offers will be made and accepted according to a version of the DA algorithm, but a human version where people make phone calls.
  - Offers can be accepted early and programs often ask students to indicate in advance their willingness to accept an offer. (You'll see why.)
- Roth and Xing (1994) describe a site visit in 1993.

Example: clinical psychology

- Program had 5 positions, 71 applicants, 29 interviews,
  - Directors had ranked 20, and knew 6 would say yes if asked. Their strategy: “don’t tie up offers with people who will hold them”.
- Timeline on selection day
  - At 9:00, calls placed to candidates 1,2,3,5,12 ---- 3,5,12 accept.
  - Candidate 1 reached at 9.05, holds until 9.13, rejects.
  - In the interim, candidate 8 calls, says she will accept.
  - When 1 rejects, call placed to 8, who accepts.
  - While call is in progress, 2 calls to reject.
  - Call placed to 10 (who’d indicated acceptance), accepts at 9.21.
  - By 9.35, remaining candidates informed of non-offer.

Optimal stable matchings

- A stable matching is ***man-optimal*** if every man prefers his partner to any partner he could possibly have in a stable matching.

Theorem. The man-proposing DA algorithm results in a man-optimal stable matching.

- This matching is also ***woman-pessimal*** (each woman gets worst outcome in any stable matching).
- Note: the same result holds for the woman-proposing DA with everything flipped.

Proof

- Say that *w* is ***achievable*** for *m* if *(m,w)* in some stable matching.
- Method: show by induction that no man is ever rejected by a woman who is achievable for him.
  - Suppose this is the case through round *n.*
  - Suppose at round *n+1*, woman *w* rejects *m* in favor of *m'*.
  - Can there be a stable match that includes *(m,w)?*
    - If so, *m'* must be matched with some *w'* who he prefers to *w* and who is achievable for him (or else *w,m'* block).
    - But then *m'* could not be making an offer to *w* in round *n+1: m'* would have first extended an offer to *w'* and would not have been turned away*.*
- So in no round is a man rejected by an achievable woman.

Rural Hospitals Theorem

Theorem. The set of men and women who are unmatched is the same *in all* stable matchings.

Why "rural hospitals"? DA is used to assign doctors to hospitals – some hospitals wondered if changing the algorithm would help them fill positions.

Proof of RH Theorem

- Let M,W be the *sets* of men and women matched in the man-optimal stable matching (which is also woman-pessimal).
- Let M',W' be the *sets* of men and women matched in some other stable matching.
- Any man in M' must also be matched in the man-optimal stable matching, so $M' \subseteq M$ … and also $|M'| \leq |M|$.
- Any woman matched in the woman-pessimal stable matching must also be in W', so $W \subseteq W'$ … and also $|W| \leq |W'|$.
- In any stable matching, the number of matched men equals the number of matched women, so $|M|=|W|$ and $|M'|=|W'|$.
- Therefore $|M'| = |M| = |W| = |W'|$.
- And so we must have M=M' and W=W'.

Strategic Behavior

- The DA algorithm asks participants to report their preferences. Should they report truthfully or be strategic?

- Definitions:
  - A *matching mechanism* is a mechanism that maps reported preferences into a matching.
  - A mechanism is *strategy-proof* if for each participant it is always optimal to be truthful.

DA is strategy-proof for men

Theorem (Dubins and Freedman; Roth). The men proposing deferred acceptance algorithm is strategy-proof for the men.

Proof.

- Fix the reports of all agents but man *m*.
- Show that whatever report man *m* starts with, he can make a series of (weak) improvements leading to a truthful report.

Proof

Suppose man *m* is considering a strategy that leads to a matching *μ* where he gets *w*. Each of the following changes improves his outcome

- Reporting that *w* is his only acceptable woman.
  - *μ* is still stable for the new problem.
  - By RH, in DA (for the new problem), *m* must get matched, and so must get *w.*
- Reporting honestly, but truncating at *w.*
  - *m* being unmatched is still blocked (because it was blocked if *m* reported just *w*), so *m* must do at least as well as *w.*
- Reporting honestly with no truncation.
  - This won't affect DA relative to above strategy.

DA manipulable by women

- **Example 3 continued** (two men, two women)
  - *A* prefers *C* to *D* *(P(A):C>D>A)*
  - *B* prefers *D* to *C* *(P(B):D>C>B)*
  - *C* prefers *B* to *A* *(P(C):B>A>C)*
  - D prefers *A* to *B* *(P(D):A>B>D)*
- Under man-proposing DA algorithm
  - If everyone reports truthfully: {*(A,C),(B,D)*}
  - If *C* reports that *A* is unacceptable, the outcome is instead {*(A,D),(B,C)*} --- better for *C!*

Strategic behavior

- The example on the previous slide can be used to establish the following result.

Theorem. There is *no* matching mechanism that is strategy-proof and always generates stable outcomes given reported preferences.

- Both versions of DA lead to stable matches, so neither version is strategy-proof for all participants!

Truncation strategies

- In the "man-proposing" DA, a woman can game the system by *truncating* her rank-order list, and stopping with the man who is the best achievable partner for her in any stable matching.

Theorem. Under the man-proposing DA, if all other participants are truthful, a woman can achieve her best achievable man using the above strategy.

- Question: how likely is it that one would have the information to pull off this kind of manipulation?

Proof

- The DA must yield a stable matching.
- If participants report as in Thm, one stable matching is the woman-optimal matching under the original true preferences.
- This gives the manipulator her best achievable man.
- By the RH theorem, she can't be unmatched in some other stable matching. She also can't get someone better because whatever would block under the true preferences will block under the reports.
- Therefore, she must get her best achievable man!

Example 3 continued

- Two men (*A)lan ,(B)ob* and two women (*C)athy,(D)ebra*

- *Alan* prefers *Cathy* to *Debra (P(A):C>D>A)*
- *Bob* prefers *Debra* to *Cathy (P(B):D>C>B)*
- *Cathy* prefers *Bob* to *Alan (P(C):B>A>C)*
- *Debra* prefers *Alan* to *Bob (P(D):A>B>D)*

- Man-proposing DA {(Alan,Cathy),(Bob,Debra)}
- Woman-proposing DA {(Alan,Debra),(Bob,Cathy)}

Example 3 continued

- Man-proposing DA {(Alan,Cathy),(Bob,Debra)} “truthful”

- *P(A):C>D>A*
- *P(B):D>C>B*
- *P’(C):B>C>A*
- *P(D):A>B>D*

- Man-proposing DA {(Alan,Debra),(Cathy,Bob)} “truncation” Better for Cathy!

How many stable matchings?

- Evidently, the incentives and scope for manipulation depend on whether preferences are such that there are many stable matchings.
- If there is a unique stable matching given true preferences, there is no incentive to manipulate if others are reporting truthfully.
- When might we have a unique stable matching?
  - Ex: if all women rank men the same, or vice-versa.
  - In “large” markets.

Stable Matching and Orderly Markets

Medical Residents, Law Clerks and College Admissions

Stability and Orderly Markets

- Some hypotheses to consider:
  - Centralized clearinghouse can lead to more “orderly” market than decentralized (eg clinical psychology)
  - Designing a clearinghouse to achieve a stable match might discourage re-contracting, or pre-contracting.
- How could one test these hypotheses?
  - Compare DA to alternative matching algorithms.
  - Compare centralized markets to decentralized.
  - Ideally, with some sort of experiment (lab? natural?)

Priority matching

- Men and women submit their preferences.
  - Each man-woman pair gets a priority based on their mutual rankings.
  - Algorithm matches the highest-priority pairs, takes them out of the market.
  - New priorities may be assigned and process iterates.
- Example of priority assignment:
  - Assign priority based on product of the two rankings, so that priority order is 1-1, 2-1, 1-2, 1-3, 3-1, 4-1, 2-2, 1-4, 5-1, etc…
  - Algorithm implements all “top-top” matches, then conditional top-tops, etc. When none remain, look for 2-1 matches, etc.
- Compare this to DA: will priority matching lead to a stable match? Should people be honest or strategic about their preferences?

Example

- $M = \{m_1, m_2, m_3\}$ and W $= \{w_1, w_2, w_3\}$.
- $$P(m_1): w_3 > w_2 > w_1$$
  $$P(m_2): w_1 > w_3 > w_2$$
  $$P(m_3): w_2 > w_3 > w_1$$
  $$P(w_1): m_1 > m_3 > m_2$$
  $$P(w_2): m_2 > m_1 > m_3$$
  $$P(w_3): m_2 > m_3 > m_1$$
- Priority mechanism outcome in which priorities are not reassigned?

Example Continued

- No couple with score 1.
- Only one couple with score 2, (m_2, w_3).
- From the remaining two have a score of 3: (m_1, w_1) and (m_3, w_2).
- $\Rightarrow \{(m_1, w_1), (m_2, w_3), (m_3, w_2)\}$.

- Not stable: (m_1, w_2) blocking pair.
- Not strategyproof: m_1 can submit w_2 as top choice.

$P(m_1): w_3 > w_2 > w_1$
$P(m_2): w_1 > w_3 > w_2$
$P(m_3): w_2 > w_3 > w_1$
$P(w_1): m_1 > m_3 > m_2$
$P(w_2): m_2 > m_1 > m_3$
$P(w_3): m_2 > m_3 > m_1$

Failure of priority matching

- Roth (1991, *AER*) studied residency matches in Britain, which are local and have used different types of algorithms --- a "natural experiment".
- Newcastle introduced priority matching in 1967.
  - By 1981, 80% of the preferences submitted contained only a single choice.
  - The participants had pre-contracted in advance!
- This is the type of "market unraveling" that plagued the US residency market prior to the NRMP.
- We'll have more to say about unraveling later.

Success of stable mechanisms

Market	Algorithm	Still Used?
NRMP	DA	yes
US Medical Specialties (about 30)	DA	yes
UK Residency matches (Roth, 1991)		
Edinburgh	DA	yes
Cardiff	DA	yes
Birmingham	Priority	no
Newcastle	Priority	no
Sheffield	Priority	no
Cambridge	Priority	yes
London hospital	Priority	yes
Canadian lawyers	DA	yes
Pharmacists	DA	yes
Reform rabbis	DA	yes

Stability and Market Participation

- Starting in the 1970s, an increasing number of *couples* graduated from medical school.
  - Typically couples want to be in the same city, but the DA algorithm doesn't account for this; it might put a husband in Boston and wife in Chicago.
  - So many couples started to go around the NRMP to find positions where they could be at the same hospital – match was threatened by a new form of unraveling.

- NRMP realized there was a problem and eventually asked Al Roth to re-design the match.

Couples: a problem!

- Couple *c1,c2* and single student *s*
- Two hospitals, each hiring one student
  - Hospital 1: *c1, s*
  - Hospital 2: s*, c2*
  - Single student: *H1, H2*
  - Couple: (H1, H2) or nothing.
- There is no stable match!

Re-Design of the NRMP

- No “clean” solution to the couples problem.
  - Student-proposing perhaps a bit better => switch from hospital to student proposing version of DA.
  - What if DA algorithm doesn’t find a stable match? Perturb the algorithm and keep going.
  - Not *guaranteed* to find a stable match, but seemed to work in simulations.

- These theory-guided “fixes” brought couples back into the match, stopping the unraveling.

The NRMP as a Case Study

- NRMP is an unusual but illuminating design of a matching market because it is *so* organized.
  - Motivation for moving to a clearinghouse was unravelling (and disorderly operation) of decentralized matching.
  - Design of the clearinghouse evidently quite important: systems with unstable matching seem to have fared poorly.
- Next, let’s consider some comparable markets with different approaches to matching and see whether and how these insights might extend.

Law Clerk Market

- Students graduating from law schools seek positions with federal or state judges (Avery et al., 2001).
  - A similar story of unraveling: market for clerkships to start in Fall 2003 cleared in September 2001.
  - Market is thin, fast and chaotic, with judges frequently making exploding offers with short deadlines.
- Attempts to enforce hiring dates and rules (there have been many) have not worked well.
  - In March 2002, Judicial Conference agreed to a one-year hiring moratorium, with hiring for Fall 2004 to start in Fall 2003. Moratorium was okay, but the start date was tricky…

Incentives

- **Students :** do not know what their other options may be and also because it is, quite simply, difficult and uncomfortable to hold off a federal judge.
- **Judges:** <u>Judge Kozinski:</u> *"From the judge's perspective, making an early offer allows him to … attract candidates who might not otherwise seriously consider him for a clerkship."*

<u>Another appellate judge:</u>

"I live in, and my office is located in, a country town . .. . [I]t is not every young man or woman who will come here to live; indeed, most won't. . . .[Initially] I did not employ law clerks until they had finished the first term of their senior year of law school. . . . I soon found out that it was more and more difficult to get law clerks from the top of the class. . . . But I have found that there are a few people in the top of the class at most law schools who had rather be assured of a job early, even in a town this size, than to wait and enter the contest in becoming clerks for judges in the larger cities with the larger and better-advertised reputations."

From the plan announcement

Q How is "Fall" determined under the Plan?

A **There is no fixed definition of Fall, nor is there any fixed starting date for the hiring season.** Under existing arrangements, some judges do their hiring in September, some in October, and others do it even later. The Plan does not change this.

Q Are judges forbidden from making "exploding offers," i.e., offers that require an applicant to respond promptly to an offer?

A **The Plan does not purport to address how an offer is given by a judge.** This is for each judge to determine. However, no applicant is obliged to act on an offer if the terms are unacceptable, nor is an applicant obliged to accept the first offer that he or she receives.

From the plan announcement..

Q Does the Plan endorse Summer interviewing?

A **No.** Many judges would have opposed the Plan had it endorsed Summer interviewing. There was a concern that Summer interviews would be very inconvenient for many people. The reasons are manifold: many judges are away on vacation during the Summer; law clerk applicants are otherwise occupied with Summer jobs, vacations, foreign travel, and bar examinations (for recent graduates); law professors often are away on vacation and thus unavailable to furnish references; and law schools do not release grades on any uniform schedule, so official student transcripts from some law schools are not available until near September. **However, the Plan does not forbid a law student who, say, is from Virginia and working in Tulsa during the Summer from talking with a judge who is otherwise available to chat.** This has happened under existing hiring arrangements and the judges saw no reason to prohibit it under the new Plan. The main point, however, is that the formal hiring process will take place in the Fall when applications will be submitted and materials and references from the law schools will be sent to the judges.

71

Exploding offers

- "I received the offer via voicemail while I was in flight to my second interview. The judge actually left three messages

  First, to make an offer.

  Second, to tell me that I should respond soon

  Third, to rescind the offer.

  It was a 35 minute flight.
- "I had 10 minutes to accept"
- "I asked for an hour to consider the offer. The judge agreed; however thirty minutes later [the judge] called back and informed me that [the judge] wanted to rescind my offer."

Market Unraveling

- The residency and law clerk markets are two examples of markets that have *unraveled* – clearing of the market has shifted earlier and earlier.
- Many matching markets, especially with fixed appointment dates, have suffered from this problem.
- Examples
  - Medical fellowships
  - Judicial clerkships
  - College admissions
  - College football bowls
  - High school prom
  - NBA/NCAA basketball recruiting
  - Baseball free agency
  - Political campaigns/primaries

Gastroenterology Fellowships

- Doctors completing residencies often continue training for 2-3 years as specialized fellows.
- Many fellowships, but not all, have adopted versions of the NRMP matching system.
- The gastroenterology (GI) fellowship is particularly interesting because in the mid-1990s the match collapsed, and only recently re-started.
  - Opportunity to do a “before and after” study of the effect of a centralized clearinghouse in an entry-level labor market.

What happened?

- Key event may have been 1996 study in JAMA stating that there are “too many” GI docs, and calling for 25-50% reduction in fellowships that programs then endorsed.
- Following this, both sides appear to have felt that they were on the short side of the market -- although it seems demand for positions did not actually fall off.
- After the match died, interesting changes
  - Market became early and very rushed
  - Matches were made “locally”

Collapse of GI Match

	Posts withdrawn (%)	Posts in Match	Percent Matched	Applicants per position
1992	--	377	97	1.8
1993	--	399	94	1.6
1994	--	369	93	1.6
1995	4	337	89	1.3
1996	5	298	75	0.9
1997	16	213	85	1.1
1998	44	99	78	1.5
1999	60	14	--	--

Unraveling of Interviews

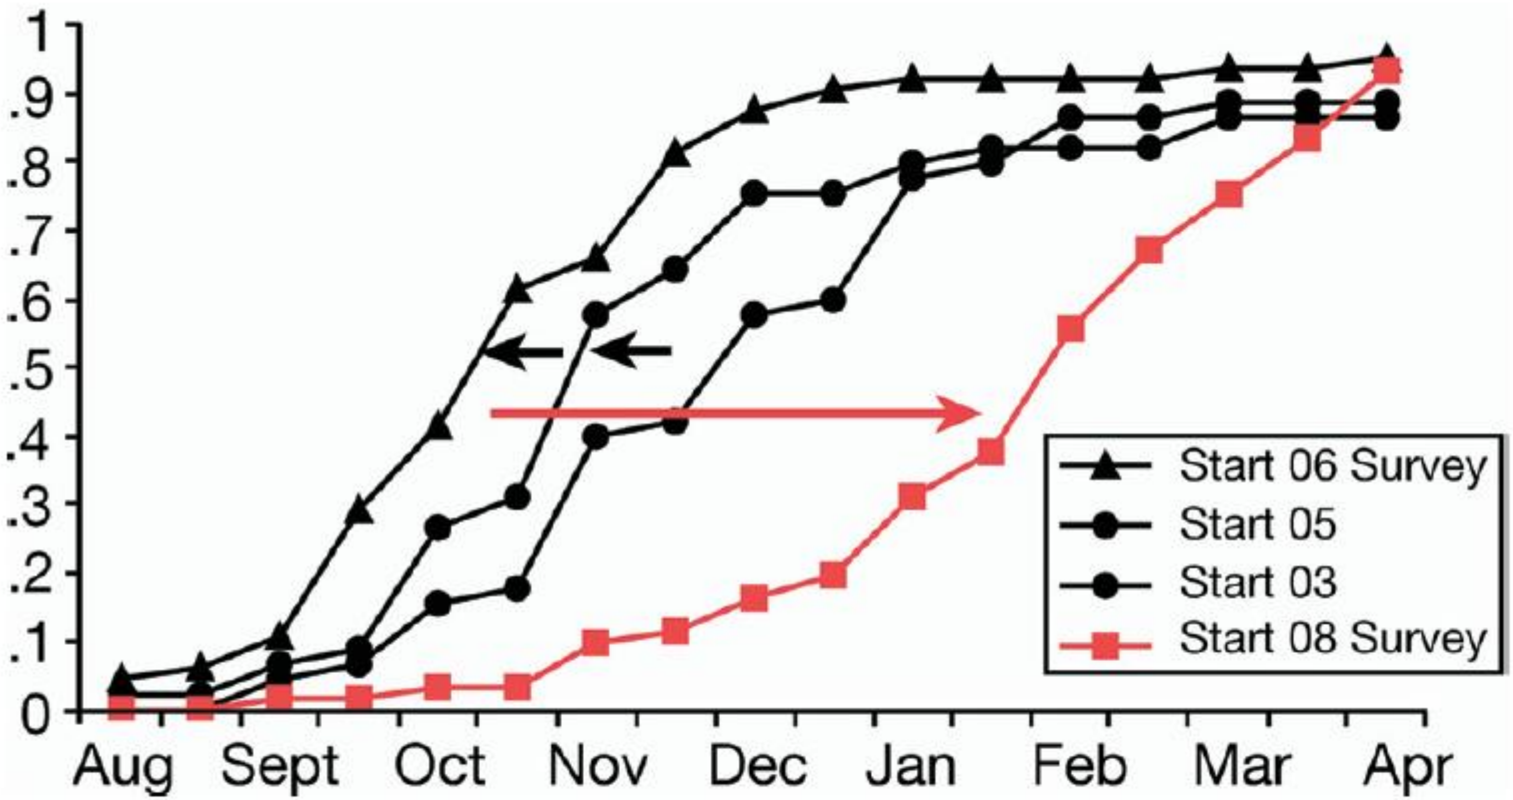

Some tentative conclusions

- Decentralized markets with fixed appointment dates can have timing issues that may create problems
  - Employers and workers may have an incentive to “jump the gun” in order to ensure a match or a good match.
  - Employers may be hesitant to leave offers outstanding, and may want to use “exploding” offers to rush decisions.
  - The market can clear in a disorderly fashion so that participants end up with a relatively limited set of choices.

- These issues *potentially* can translate into inefficient & unfair matching.