“Market Design and Computer-Assisted Markets: An Economist’s Perspective”

Simons Institute, Berkeley

May 31, 2013

Computer-assisted markets

- We're seeing increasingly intensive uses of computation in markets
  - **Markets can be run on computers (fast)**
  - Markets can be accessed over the internet (ubiquitous)
  - Markets can use computers as trusted intermediaries
  - **Market outcomes can be determined by (computationally intensive) algorithms**
  - Markets can be customized to a high degree
    - We're seeing increasingly interesting **auction markets**
    - E.g. google search words, packages of spectrum

Note that not all markets resemble auctions

- **Stanford and Berkeley don't raise tuition until just enough applicants remain to fill the freshman class.**
- They keep tuition low enough so that *many* students would like to attend, and then they admit a fraction of those who apply.
- Universities don't rely on prices alone to equate supply and demand
- **Labor markets and college/university admissions are more than a little like courtship and marriage**: each is a two-sided matching market that involves searching and wooing on both sides.

Matching markets

- *Matching markets* are markets in which **you can't just choose what you want** (even if you can afford it)**, you also have to be chosen.**
- You can't just inform Stanford that you're enrolling, or Google or Facebook that you're showing up for work. You also have to be *admitted* or *hired*. Neither can Stanford or Berkeley simply choose who will come, any more than one spouse can simply choose another: each also has to be *chosen*.

Some matching markets I've helped design

- **Medical labor markets**
  - Medical Residents: in the U.S.: NRMP in 1995
  - many Fellowship markets
- **School choice systems:**
  - New York City high schools, Boston public schools
  - 2012 Denver, New Orleans (including charter schools)
- Other labor markets
- Kidney exchange (which I'll talk about in EC 2013 in Philadelphia next month)

Deferred acceptance algorithm

- Gale and Shapley (1962)—stable matching
- Taught in algorithms classes
- Arises in markets
- Useful in some kinds of market design
  - Allows computerized clearinghouses to solve different kinds of problems
    - Some you might think are too easy to need a computer
    - Some you might think are too hard even with a computer

Many-to-one matching: The college admissions model

PLAYERS: Firms = $\{f_1,\ldots, f_n\}$ Workers= $\{w_1,\ldots,w_p\}$

 # positions $\mathbf{q_1,\ldots,q_n}$

Synonyms (sorry:-): F=Firms = C=Colleges = H=Hospitals = S=Schools
W=Workers = I = Individuals (sometimes students☺

PREFERENCES *over individuals* (complete and transitive), as in the marriage model:

$$P(f_i) = w_3, w_2, \ldots f_i \ldots \qquad [w_3 >_{fi} w_2]$$

$$P(w_j) = f_2, f_4, \ldots w_j \ldots$$

An OUTCOME of the game is a *MATCHING*:

$$\mu: F \cup W \rightarrow F \cup W$$

such that $\mu(f)$ contains w iff $\mu(w) = f$, and for all f and w

$|\mu(f)|$ is less than or equal to q_f

either $\mu(w)$ is in F or $\mu(w) = w$. so f is matched to the *set* of workers $\mu(f)$.

A matching μ is *individually irrational* if $\mu(w) = f$ for some worker w and firm f such that either the worker is unacceptable to the firm or the firm is unacceptable to the student. An individually irrational matching is said to be blocked by the relevant individual.

A matching μ is *BLOCKED BY A PAIR OF AGENTS* (f,w) if they each prefer each other to μ:

$[w >_f w'$ for some w' in $\mu(f)$ or $w >_f f$ if $|\mu(f)| < q_f]$

and $f >_w \mu(w)$

As in the marriage model, a matching is (pairwise) *stable* if it isn't blocked by any individual or pair of agents.

GS Deferred Acceptance Algorithm, with **workers applying**

- 0.0 Workers and firms *privately* submit preferences
- 0.1 *If some preferences are not strict*, arbitrarily break ties
- 1 a. Each worker applies to his/her top choice firm.
- b. Each firm f with q positions holds the top q applications among the acceptable applications it receives, and rejects all others.
- k a. Any worker rejected at step k-1 makes a new application, to its most preferred acceptable firm that hasn't yet rejected him/her. (If no acceptable choices remain, he/she makes no further offers.)
- b. Each firm holds its q most preferred acceptable applications to date, and rejects the rest.
- STOP: when no further proposals are made, and match each firm to the workers (if any) whose applications it is holding.

GS Deferred Acceptance Algorithm, with firms offering

- 0.0 Workers and firms *privately* submit preferences
- 0.1 *If some preferences are not strict*, arbitrarily break ties
- 1 a. Each firm f with q positions makes offers to its top q choices (if it has any acceptable choices).
- b. Each worker rejects any unacceptable offers and, if more than one acceptable offer is received, "holds" the most preferred and rejects all others.
- k a. Any firm rejected at step k-1 makes a new offer, for each rejection, to its most preferred acceptable worker who hasn't yet rejected it. (If no acceptable choices remain, it makes no further offers.)
- b. Each worker holds her most preferred acceptable offer to date, and rejects the rest.
- STOP: when no further proposals are made, and match each worker to the firm (if any) whose offer she is holding.

GS's 2 Remarkable Theorems

- **Theorem 1 (GS):** A stable matching exists for every marriage market.
- **Theorem 2 (GS):** When all firms and workers have **strict preferences**, there always exists an F-optimal stable matching (that every firm likes at least as well as any other stable matching), and a W-optimal stable matching. Furthermore, the matching μ_F produced by the deferred acceptance algorithm with firms proposing is the F-optimal stable matching. The W-optimal stable matching is the matching μ_W produced by the algorithm when the workers propose.

Incentives

- There don't exist any stable mechanisms that make it a dominant strategy for everyone to state their true preferences
- **When preference lists are private**, in a many-to-one matching market (e.g. firms and workers) the deferred acceptance algorithm with the workers (the side who needs 1 position) proposing makes it **a dominant strategy for workers to state their true preferences.**

Stability turns out to be important for a successful 2-sided market clearinghouse

Market	Stable	Still in use (halted unraveling)
• NRMP	yes	yes (new design in ’98)
• *Edinburgh ('69)*	*yes*	*yes*
• *Cardiff*	*yes*	*yes*
• *Birmingham*	*no*	*no*
• *Edinburgh ('67)*	*no*	*no*
• *Newcastle*	*no*	*no*
• Sheffield	no	no
• Cambridge	no	yes
• London Hospital	no	yes
• Medical Specialties	yes	yes (~30 markets, **1 failure**)
• Canadian Lawyers	yes	yes (Alberta, no BC, Ontario)
• Dental Residencies	yes	yes (5) (no 2)
• Osteopaths (< '94)	no	no
• Osteopaths (≥ '94)	yes	yes
• Pharmacists	yes	yes
• Reform rabbis	yes (first used in ‘97-98)	yes
• **Clinical psych**	**yes (first used in ‘99)**	**yes**
• Lab experiments	yes	yes
(Kagel&Roth *QJE* 2000)	no	no

Lab experiments fit nicely on the list, just more of a variety of observations that increase our confidence in the robustness of our conclusions, the lab observations are the smallest but most controlled of the markets on the list…

Stable Clearinghouses based on deferred acceptance algorithm (blue -> today use the Roth Peranson Algorithm)

NRMP / SMS:

Medical Residencies in the U.S. (NRMP) (1952)
Abdominal Transplant Surgery (2005)
Child & Adolescent Psychiatry (1995)
Colon & Rectal Surgery (1984)
Combined Musculoskeletal Matching Program (CMMP)

- Hand Surgery (1990)

Medical Specialties Matching Program (MSMP)

- Cardiovascular Disease (1986)
- Gastroenterology (1986-1999; rejoined in 2006)
- Hematology (2006)
- Hematology/Oncology (2006)
- Infectious Disease (1986-1990; rejoined in 1994)
- Oncology (2006)
- Pulmonary and Critical Medicine (1986)
- Rheumatology (2005)

Minimally Invasive and Gastrointestinal Surgery (2003)
Obstetrics/Gynecology

- Reproductive Endocrinology (1991)
- Gynecologic Oncology (1993)
- Maternal-Fetal Medicine (1994)
- Female Pelvic Medicine & Reconstructive Surgery (2001)

Ophthalmic Plastic & Reconstructive Surgery (1991)
Pediatric Cardiology (1999)
Pediatric Critical Care Medicine (2000)
Pediatric Emergency Medicine (1994)
Pediatric Hematology/Oncology (2001)
Pediatric Rheumatology (2004)
Pediatric Surgery (1992)

Primary Care Sports Medicine (1994)
Radiology

- Interventional Radiology (2002)
- Neuroradiology (2001)
- Pediatric Radiology (2003)

Surgical Critical Care (2004)
Thoracic Surgery (1988)
Vascular Surgery (1988)

Postdoctoral Dental Residencies in the United States

- Oral and Maxillofacial Surgery (1985)
- General Practice Residency (1986)
- Advanced Education in General Dentistry (1986)
- Pediatric Dentistry (1989)
- Orthodontics (1996)

Psychology Internships in the U.S. and CA (1999)

Neuropsychology Residencies in the U.S. & CA (2001)
Osteopathic Internships in the U.S. (before 1995)
Pharmacy Practice Residencies in the U.S. (1994)
Articling Positions with Law Firms in Alberta, CA(1993)
Medical Residencies in CA (CaRMS) (before 1970)

****

British (medical) house officer positions

- Edinburgh (1969)
- Cardiff (197x)

New York City High Schools (2003)
Boston Public Schools (2006)
Other cities (2012)

The Market for Clinical Psychology Interns

In 1998, this market converted to a centralized match using the Roth-Peranson algorithm (run for the first time in the academic year '98-99 for jobs beginning in June 1999.) For approximately 25 years prior to that, a decentralized market was run, under a changing set of rules.

Part of market design for an existing market involves understanding the problems the market is encountering. The decentralized psychology market was studied in

Roth, A.E. and X. Xing "Turnaround Time and Bottlenecks in Market Clearing: Decentralized Matching in the Market for Clinical Psychologists," *Journal of Political Economy*, 105, April 1997, 284-329.

In the early 1990's, transactions in this market were supposed to all be made by telephone on "Selection Day," the second Monday in February, from 9:00 AM to 4:00 PM Central Standard Time. That is, the market was supposed to operate for seven hours.

Subject to many modifications of its rules, this kind of decentralized but uniform timing regime was used in this market since 1973.

One kind of modification has concerned the length of the market. In the early 1970's the market lasted *five days*, and was subsequently shortened to *three*, and for most of the 1980's the rules specified that the market would take place from 8:00 AM Monday until Noon the following day, i.e. for a *day and a half*.

Once again, by the 1990's this time had been shortened to *seven hours*, and in the late 90's it was shortened further (I think to *four* hours).

(This concern with the amount of time which the market, and individual offers, should remain open is one that has been observed in many markets.)

<u>APPIC Policy: Internship Offers and Acceptances (5/91)</u>
Adherence to these policies is a condition of membership in APPIC

"Selection day" currently begins at 9:00 am Central Standard Time on the second Monday in February, and ends at 4:00 pm that same day. This definition is subject to change.

…

3. No internship offers in any form may be extended by agencies before the beginning of selection day.

a. The ***only*** information that agencies may communicate to applicants prior to this time is whether or not the applicant remains under consideration for admission (see item 2). The spirit of this item precludes any communication of an applicant's status prior to the time above, however, "veiled" or indirect such communication might be.

b. …

c. Internship programs may not solicit information regarding an applicant's ranking of programs or his/her intention to accept or decline an offer of admission until after that offer is officially tendered. *<ha!>*

4. Applicants must reply to all offers no later than the closing time on selection day.

a. This deadline applies to all offers including those to applicants who are initially considered "alternates" and are subsequently extended an offer any time prior to end of selection day.

b. Agencies may inquire as to the applicant's progress towards making a decision at any time after an offer is formally extended. *Under no circumstances, however, may an agency implicitly or explicitly threaten to rescind an offer if a decision is not made prior to the end of selection day* (except as noted in item 6).

c. It is in everyone's best interest that applicants make and communicate decisions to accept or reject each offer *as quickly as possible*.

d. Any offer that has not been accepted is void as of the ending hour of selection day.

Rules 5+6: “Firm-proposing deferred acceptance by telephone”

5. **An applicant must respond immediately to each offer tendered in one of three ways. The offer may be accepted, rejected or "held."**
    a. *Accepting* the offer constitutes a binding agreement between applicant and internship program.
    b. *Refusing* the offer terminates all obligations on either side and frees the internship program to offer the position to another applicant.
    c. *Holding* the offer means that the offer remains valid until the applicant notifies the program of rejection or acceptance, or until the end of selection day.

6. **Applicants may "HOLD" no more than one active offer at a time.**
    a. If an applicant is holding an offer from one program and receives an offer from a more preferred program, s/he may accept or "hold" the second offer provided that the less preferred program is notified *immediately* that the applicant is rejecting the previously held offer.
    b. If a program confirms that an applicant is holding more than one offer, the program is free to withdraw their previously tendered offer of acceptance, and to offer that position to another applicant *after* the offending applicant is notified of that decision.

7. An offer of acceptance to an applicant is valid only if the applicant has not already accepted an offer of admission to another program.

a. An applicant's verbal acceptance of an offer constitutes a *binding agreement* between the applicant and the program that may not be reversed unilaterally by either party.
b. Before programs extend an offer, they must first *explicitly* inquire whether the applicant has already accepted an offer elsewhere. If so, no offer may be tendered.
c. *A program may in no way suggest* that an applicant renege on previously accepted offers.
d. *If an applicant who has accepted an offer receives a second offer*, s/he is obligated to refuse the second offer and inform the agency that s/he is already committed elsewhere.
e. Any offer accepted subsequently to a prior commitment is automatically null and void, even if the offering agency is unaware of the prior acceptance and commitment.

8. When an applicant accepts an offer of admission, s/he is urged to immediately inform all other internship programs at which s/he is still under consideration that s/he is no longer available.

Rule 9: “Aftermarket”

9. **Applicants who have not accepted a position prior to the end of selection day may receive offers of admission after that deadline.**
   a. Applicants should be prepared to accept or reject such late offers quickly, since most other deliberations should have already taken place.
   b. Programs may legitimately place short but reasonable (☺) deadlines for responses to such late offers.

10. **Once a program has filled all available positions, all candidates remaining in their applicant pool must be notified that they are no longer under consideration.**
    a. Applicants who have not notified the agency that they have accepted a position elsewhere and who have not been selected by the agency should be notified by phone as soon as all positions are filled.
    b. If an applicant cannot be reached by phone, s/he should be so notified by letter postmarked no later than 72 hours after the end of selection day.

11. **Internship training directors should document their verbal agreement with each applicant in a letter postmarked no later than 72 hours following the end of selection day.**
    a. The letter should be addressed to the applicant, and should include confirmation of conditions of the appointment, such as stipend, fringe benefits, and the date on which the internship begins.
    b. A copy of that letter should be sent simultaneously to the applicant's academic program director.

12. **Applicants who receive offers which do not comply with these policies or who in other ways detect violations of these policies by an APPIC member program are urged to request compliance with APPIC policies from the program representative.**
    a. Applicants should immediately report any problems unresolved after such request to his/her academic program director.
    b. Academic program directors are urged to contact internship training directors immediately regarding such unresolved problems.
    c. Such compliance problems should be resolved through consultation among applicant, internship program, and academic training director whenever possible.
    d. Problems not amenable to resolution through such consultation should be reported as soon as possible to the APPIC Standards and Review Committee …

13. Internship directors who become aware of violations of policies on the part of students, academic training directors, or other internship directors are urged to immediately request compliance to the policies.

a. Internship directors are urged to contact academic training program directors immediately regarding problems that remain unresolved after such a request for compliance.
b. Internship program directors who become aware of violations of these policies by other internship programs should urge the applicant and academic training directors involved to follow the procedures outlined in 12 a-d above, and/or directly contact the other internship director.
c. Such compliance problems should be resolved through consultation among applicant, internship programs, and academic training director whenever possible.
d. Failure to resolve compliance problems through consultation should be reported to the APPIC Standards and Review Committee.

14. All reported violations of these policies will be considered by the APPIC Standards and Review Committee (SRC). SRC policies are described in the *APPIC Directory*. Violations of these policies should be reported to: Chair, APPIC Standards and Review Committee

(*These don't look like the rules of a trouble-free market…)*

Behavioral Observations on Selection Day

Transaction times were *FAST*:

Offers took about 5 minutes to deliver
Rejection of offers took about 1 minute
New offers were made immediately following a rejection

Surveys of students report that > 10% got early offers

There was a great deal of pressure on students to indicate a 1st choice (despite very explicit rules prohibiting such pressure)

There was considerable willingness by students to indicate a 1st choice. (And repeated game issues seemed to make these signals credible… "you see these people again…") **But couples had trouble making this kind of commitment**…

Employers paid serious attention to indications of first choice, in deciding to whom to give offers.

<u>Question</u>: why isn't this fast, decentralized process inducing the behavior we'd expect from the (centralized) deferred acceptance procedure?

Decentralized deferred acceptance with random elements (with and without an aftermarket)

initial state: t=0, all positions are vacant, all workers are unmatched, no communication is underway.

Preferences: **P** = [$P(F_1),\ldots,P(F_n);P(w_1),\ldots,P(w_m)$] selected from some specified joint probability distribution.

1. Offers, deferred acceptances, and rejections:

a. All *available* firms, i.e. firms which are not currently engaged in communication and which have at least one position for which no offers are outstanding, attempt to make offers to their most preferred workers who haven't yet rejected them. Some subset of this set (containing no more than one firm seeking to make an offer to any given worker) succeeds in establishing communication with the worker to whom they wish to make an offer--this successful set is determined according to some specified probability distribution (which may depend on the current state of the system). Successful firms remain in communication with the workers they have contacted for some time period drawn from a specified distribution.

b. Any worker who receives an offer rejects it if it is unacceptable or if she has already received an offer from a more preferred firm. Otherwise she holds it (so that the firm in question has an offer outstanding for the position). [workers who have received an offer from the first choice among those remaining on their lists can now accept the offer, and inform all firms, and firms who have had all positions accepted can now inform all applicants that their positions are filled]

2 a. Is there any firm which has not already been rejected by all of its acceptable workers and which has a position not presently being held by any worker?

No

STOP. In this case the final outcome is the matching ì which matches each worker to the position (if any) that she is holding.

Yes

b. Set t = t+1. Has time expired, i.e. $t \geq t^*$?

No

Yes

(Has time expired, i.e. t $\geq$ t*?)

Yes

AFTERMARKET

3. exploding offers after time has expired.

a. Every worker who is holding an offer at time t* accepts it; any firms which (after t*) still have vacant positions proceed to make offers as in step 1a.

b. Every worker who has already accepted an offer rejects any new offer, and every worker who has not already accepted an offer accepts the first offer received from an acceptable firm

t=t+1

c. Check if there is at least one firm which has a position that is not being held by some worker and which has not yet offered a position to all of its acceptable workers (this includes firms which may be engaged in communication). If so, set t=t+1 and return to 3a.

No

d. Otherwise **STOP**, and let the final outcome be the matching Ì which matches each worker to the position (if any) that she has accepted.

Theorem 1: If the decentralized deferred acceptance procedure is run *without any fixed termination time (i.e. $t^* = \infty$),* then the outcome would be the same stable matching as that produced by the centralized deferred acceptance procedure. In particular, both procedures produce the firm-optimal stable matching with respect to the revealed preferences, μ_F.

Proof: familiar…

The need for computation: is 7 hours near or far from $t^* = \infty$, given how fast offers and replies are?

Basic Simulations (with many variations…)

"Medical model": deferred acceptance until natural termination (no time limit)

- 200 workers, each with uniform random preferences over 20 randomly selected firms
- 50 firms, each with 4 positions, and uniform random preferences over all workers who apply.

Each firm has two phones for outgoing calls, one for incoming

Actions take place each minute.

- Offers take 5 minutes
- Rejections take 1 minute
- Information calls (following acceptances, or all positions filled) take 1 minute.

<u>Table 1</u>: The Medical Model Telephone Market
Results of 100 simulations for each three turnaround times

Number of minutes required			
to make an offer	**5**	10	25
to reject an offer	**1**	2	5
Mean time to termination at a stable outcome	**18:18**	**36:32**	**91:14**
(standard deviation)	(8:10)	(16:20)	(40:52)
Mean time by which 90% of students have received an offer	**1:02**	2:03	5:04
Mean time by which 99% of students have received an offer	**5:19**	10:35	26:22
Longest time to termination	**39:25**	78:25	196:22
Shortest time to termination	**4:59**	9:55	25:00

A lot happens in the first hour, then things slow down. And busy signals aren't playing a role: when transaction times are increased, everything scales up proportionally.

Table 2: Hourly Progress of the Medical Model Telephone Market

Mean results based on 100 simulations.

Hour	# Students Who Have Received at Least One Offer	# Students Who Have Received an Offer From the Firm to Which They Will Ultimately Be Matched	# of Offers That Have Been Made	# of Offers That Have Not Been Rejected Immediately
0	0.00	0.00	0.00	0.00
1	**178.47**	**86.32**	**400.08**	**278.06**
2	**191.24**	**116.06**	**531.96**	**333.90**
3	**194.83**	**132.75**	**602.36**	**360.04**
4	196.50	143.81	648.58	375.70
5	197.41	152.14	681.79	386.80
6	198.02	158.48	707.38	395.01
7	**198.37**	**163.37**	**727.89**	**401.10**
8	198.54	167.66	745.23	406.29
9	198.68	171.46	761.06	410.70
10	198.84	174.77	775.07	414.65
..				
39	199.97	199.95	881.62	442.46
40	199.99	199.99	881.71	442.50

The market undergoes a kind of "phase change," from parallel processing in the first hour, to serial processing once most offers are being held…

Table 3: The Telephone Market with 7 Hours Enforced Termination Time

Results of 100 simulations for each of the following cases

	The Psych Model	20 Students May Hold Two Offers Once for Two Hours	Every Student May Hold Two Adjacent Offers Until One Hour Before the Deadline	Every Student May Hold Two Adjacent Offers Until the Deadline	Every Firm First Issues Offers to Students Who Like It Best
Mean time to termination (standard deviation)	**7:43 (0:22)**	7:53 (0:10)	8:01 (0.10)	8:08 (0:07)	7:36 (0.37)
Mean time by which 90% of students have received at least one offer	**1:02**	2:11	2:22	2:33	0.57
Mean time by which 99% of students have received at least one offer	**5:07**	7:06	7:37	7:51	5:23
Mean # of blocking firms (standard deviation)	**1.58** (0.74)	**3.25** (1.26)	**6.32** (1.61)	**12.77** (2.27)	**2.34** (1.02)
Mean # of blocking students (standard deviation)	**16.67** (7.73)	**29.88** (9.80)	**48.74** (11.26)	**77.76** (11.57)	**15.74** (8.06)
Mean # of unmatched students	**0.88**	1.09	1.52	1.69	0.78
Mean # of unmatched firms	**0.87**	1.07	1.41	1.52	0.78

Table 4: Hourly Progress of the Psych Model Telephone Market

Mean results based on 100 simulations

Hour	# Students Who Have Received at Least One Offer	# Students Who Have Received an Offer from the Firm to Which They Will Ultimately Be Matched	# Offers that Have Been Made	# of Offers that Have Not Been Rejected Immediately
0	0.00	0.00	0.00	0.00
1	178.47	104.13	400.08	2.78.06
2	191.24	140.52	531.96	333.90
3	194.83	161.12	602.36	360.04
4	196.50	174.59	648.58	375.70
5	197.41	184.64	681.79	386.80
6	198.02	192.46	707.38	395.01
7	198.37	198.37	727.89	400.99
8	199.11	199.11	786.35	401.73
9	199.12	199.12	**786.79**	401.74

Recall that when the deferred acceptance algorithm was allowed to run its course, on average **882** offers were needed.

Conclusions

Markets in which offers must remain open for a specified time (even if it is short):

- Experience congestion
- Undergo phase changes—from parallel processing to serial processing
- Give firms an incentive to think about not only how much they like a worker, but how much the worker likes them.

Signaling can help this process work: students were asked for signals, and they influenced offers. (This is the opposite of private preferences…)

A critical element of a market is its *effective* length: how many possible transactions can be explored through the process of making offers. The effective length of the psychology market *increased* as its *duration* decreased from five days to one.

Postscript:

Since 1999, APPIC has run a centralized match, using the Roth-Peranson algorithm.

This fixed the congestion problem.

It also allows the market to address the problems faced by married couples on the market…a problem that was earlier experienced by medical doctors…

An initial "couples algorithm" in the 1970's medical match

- Couples (after being certified by their dean) could register for the match as a couple.
  - They had to specify one member of the couple as the "leading member."
  - They submitted a separate rank order list of positions for each member of the couple
- The leading member went through the match as if single.
- The other member then had his/her rank order list edited to remove positions not in the 'same community' as the one the leading member had matched to.
  - Initially the NRMP determined communities; in a later version, when couples were still defecting, couples could specify this themselves.

But this didn't work well for couples

- Why?
- The iron law of marriage: You can't be happier than your spouse.
- Couples consume *pairs* of jobs. So an algorithm that only asks for their preference orderings over *individual* jobs can't hope to avoid instabilities (appropriately redefined to include couples' preferences)
- But even if we ask couples for their preferences over pairs of jobs, we may still have a problem: Roth (1984) observed that the set of stable matchings may be *empty* when couples are present.

Why is the couples problem hard?

- Note first that the ordinary deferred acceptance algorithm won't in general produce a stable matching (even when one exists, and even when couples state preferences over pairs of positions)
  - In the worker proposing algorithm, if my wife and I apply to a pair of firms in Boston, and our offers are held, and I am later displaced by another worker, my wife will want to withdraw from the position in which she is being held (and the firm will regret having rejected other applications to hold hers)
  - In the firm proposing algorithm, it may be hard for a couple to determine which offers to hold.

And when couples are present…

- …the set of stable matchings may be empty
- And determining if so is NP-complete

Example--market with one couple and no stable matchings (motivated by Klaus and Klijn, and Nakamura (JET corrigendum 2009 to K&K JET 2005):

Let c=(s1,s2) be a couple, and suppose there is another single student s3, and two hospitals h1 and h2. Suppose that the acceptable matches for each agent, in order of preference, are given by

c: (h1,h2);	s3: h1, h2,
h1: s1, s3;	h2: s3, s2

Then no individually rational matching μ (i.e. no μ that matches agents only to acceptable mates) is stable. We consider two cases, depending on whether the couple is matched or unmatched.

Case 1: μ(c)=(h1,h2). Then s3 is unmatched, and s/he and h2 can block μ, because h2 prefers s3 to μ (h2)=s2.

Case 2: μ (c)=c (unmatched). If μ (s3)=h1, then (c, h1,h2) blocks μ. If μ (s3)=h2 or μ (s3)=s3 (unmatched), then (s3,h1) blocks μ.

Current NRMP match
(Roth/Peranson algorithm)

- The algorithm starts as a **student (and couple)-proposing deferred acceptance algorithm**,
- Whenever a couple member is withdrawn from the position holding it's application, that position is put on a "hospital stack"as a possible source of blocking pairs
- then resolves instabilities with an algorithm modeled on the Roth-Vande Vate (1990) blocking-pair-satisfying algorithm
- Deals with major match complications
  - Married couples
    - They can submit preferences over pairs of positions
  - Applicants can match to pairs of jobs, PGY1&2
    - They can submit supplementary preference lists
  - Reversions of positions from one program to another

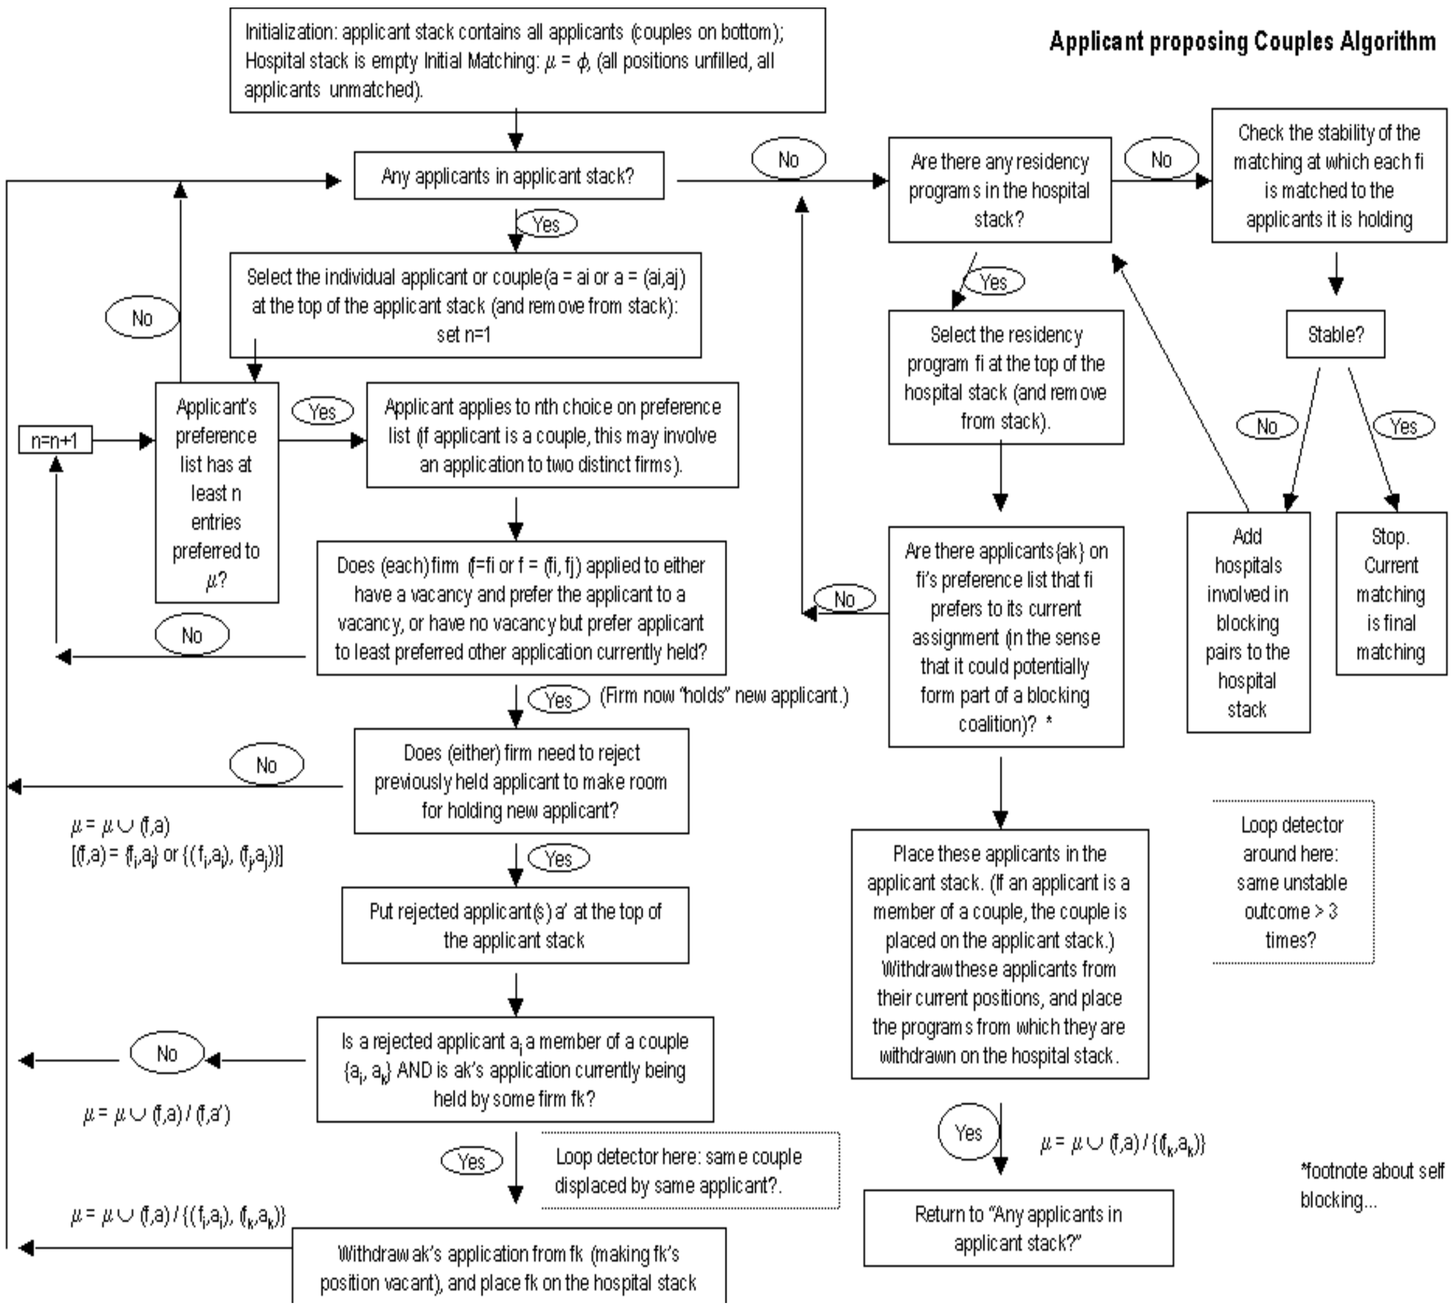

Applicant proposing Couples Algorithm
Initialization: applicant stack contains all applicants (couples on bottom); Hospital stack is empty Initial Matching: μ = ϕ, (all positions unfilled, all applicants unmatched).
Any applicants in applicant stack?
Yes
No
Select the individual applicant or couple(a = ai or a = (ai,aj) at the top of the applicant stack (and remove from stack): set n=1
Applicant's preference list has at least n entries preferred to μ?
n=n+1
No
Yes
Applicant applies to nth choice on preference list (if applicant is a couple, this may involve an application to two distinct firms).
Does (each) firm (f=fi or f = (fi, fj) applied to either have a vacancy and prefer the applicant to a vacancy, or have no vacancy but prefer applicant to least preferred other application currently held?
No
Yes
(Firm now "holds" new applicant.)
Does (either) firm need to reject previously held applicant to make room for holding new applicant?
No
μ = μ ∪ (f,a)
[(f,a) = (fi,ai) or {(fi,ai), (fj,aj)}]
Yes
Put rejected applicant(s) a' at the top of the applicant stack
Is a rejected applicant ai a member of a couple {ai, ak} AND is ak's application currently being held by some firm fk?
No
μ = μ ∪ (f,a) / (f,a')
Yes
Loop detector here: same couple displaced by same applicant?.
μ = μ ∪ (f,a) / {(fi,ai), (fk,ak)}
Withdraw ak's application from fk (making fk's position vacant), and place fk on the hospital stack
No
Are there any residency programs in the hospital stack?
No
Check the stability of the matching at which each fi is matched to the applicants it is holding
Yes
Select the residency program fi at the top of the hospital stack (and remove from stack).
Are there applicants {ak} on fi's preference list that fi prefers to its current assignment (in the sense that it could potentially form part of a blocking coalition)? *
No
Stable?
No
Yes
Add hospitals involved in blocking pairs to the hospital stack
Stop. Current matching is final matching
Place these applicants in the applicant stack. (If an applicant is a member of a couple, the couple is placed on the applicant stack.) Withdraw these applicants from their current positions, and place the programs from which they are withdrawn on the hospital stack.
Loop detector around here: same unstable outcome > 3 times?
Yes
μ = μ ∪ (f,a) / {(fk,ak)}
Return to "Any applicants in applicant stack?"
*footnote about self blocking...

Empirical puzzle

- Why do these algorithms virtually always find stable matchings, even though couples are present (and so the set of stable matchings *could* be empty)?
- Kojima, Fuhito, Parag A. Pathak, and Alvin E. Roth, "Matching with Couples: Stability and Incentives in Large Markets," revised April 2013.

Some clues

- Empirically: Roth and Peranson (1999): as markets get large, # of interviews (and hence length of rank order lists) doesn't grow too much, and the set of stable matchings gets small.
- Theoretically: the set of stable matchings gets small and hard to manipulate,
  - Immorlica and Mahdian (2005), 1-1 matching
  - Kojima and Pathak (2009), many to one matching

Stylized facts

1. Applicants who participate as couples constitute a small fraction of all participating applicants.
2. The length of single applicants' rank order lists is small relative to the number of possible programs.
3. Applicants who participate as couples rank more programs than single applicants. However, the number of distinct programs ranked by a couple member is small relative to the number of possible programs.
4. The most popular programs are ranked as a top choice by a small number of applicants.
5. Even though there are more applicants than positions, many programs still have unfilled positions at the end of the centralized match.
6. A stable matching exists in all nine years in the market for clinical psychologists.

Random markets

- A random market is a tuple Γ=(H,S,C, $\succcurlyeq_{\{H^n\}}$, k,P,Q,ρ), where k is a positive integer (max length of ROL's), $P=(p_{\{h\}})_{\{h\in H\}}$ and $Q=(q_{\{h\}})_{\{h\in H\}}$ are probability distributions on H, and ρ is a function which maps two preferences over H to a preference list for couples.
- Hospitals' preference orderings are essentially arbitrary, and take account of their capacities, and couples preferences are formed from their individual preferences (drawn from probability distribution Q, different than P for singles), via an essentially arbitrary function ρ .

Random large markets

- A sequence of random markets is $(\Gamma^1,\Gamma^2,\ldots)$, where $\Gamma^n=(H^n,S^n,C^n,\succcurlyeq_{\{H^n\}},k^n,P^n,Q^n,\rho^n)$ is a random market in which $|H^n|=n$ is the number of hospitals.

Definition: A sequence of random markets $(\Gamma^1, \Gamma^2, \ldots)$ is **regular** if there exist $\lambda>0$, $a \in [0,(1/2))$, $b>0$, $r\geq 1$, and positive integers k and κ such that for all n,

1. $k^n=k$, (constant max ROL length, doesn't grow with n—could be bounded by clogn…)
2. $|S^n| \leq \lambda n$, $|C^n| \leq bn^a$, (singles grow no more than proportionally to positions—e.g. $\lambda>1$, and couples grow slower than root n)
3. $\kappa_h \leq \kappa$ for all hospitals h in H^n (hospital capacity is bounded)
4. $(p_h/p_{h'}) \in [(1/r),r]$ and $(q_h/q_{h'}) \in [(1/r),r]$ for all hospitals h,h' in H^n. (The popularity of hospitals as measured by the prob of being acceptable to docs does not vary too much as the market grows, i.e. no hospital is everyone's favorite (in after-interview preferences)

Stable matchings exist, in the limit

- Theorem: Suppose that $(\Gamma^1,\Gamma^2,\ldots)$ is a regular sequence of random markets. Then the probability that there exists a stable matching in the market induced by Γ^n converges to one as the number of hospitals n approaches infinity.

Key element of proof

- if the market is large, then it is a high probability event that there are a large number of hospitals with vacant positions (even though there could be more applicants than positions)
- So chains of proposals beginning when a couple displaces a single doc are much more likely to terminate in an empty position than to lead to a proposal to a hospital holding the application of a couple member.

Corollary

- Suppose that $(\Gamma^1, \Gamma^2, \ldots)$ is a regular sequence of random markets. Then the probability that the Roth-Peranson algorithm produces a stable matching in the market induced by Γ^n converges to one as the number of hospitals n approaches infinity.

Many open questions about couples

- Finite markets
- Large numbers of couples
- Average difficulty of computation (empirically, not hard…)

Some concluding market design observations

Stages and transitions observed in various markets (and the role of stable matchings and deferred acceptance algorithms…)

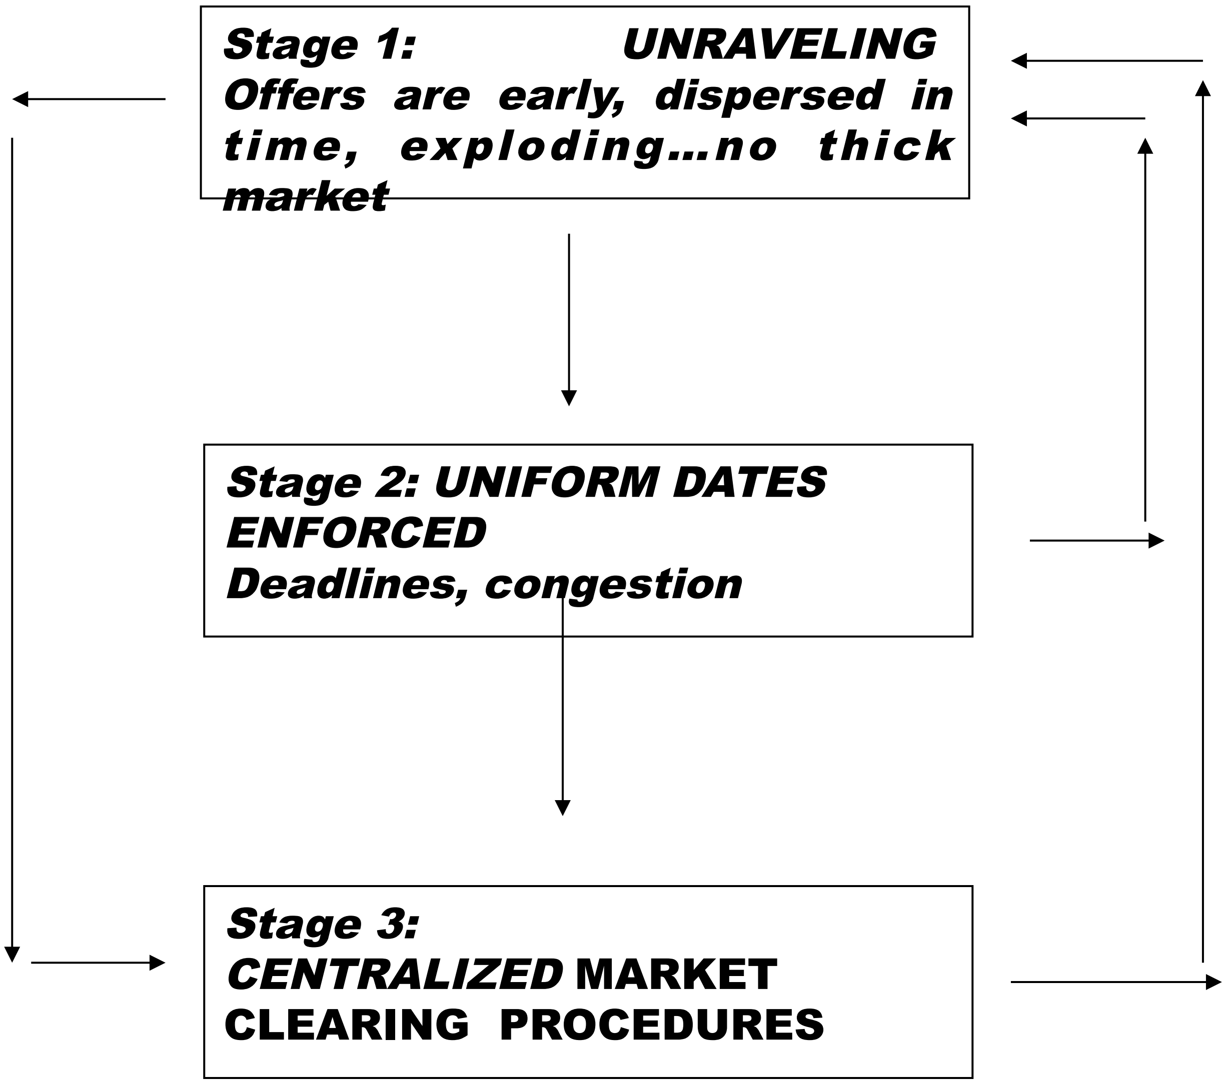

What have we learned from market design?

- To achieve efficient outcomes, marketplaces need make markets sufficiently
  - **Thick**
    - Enough potential transactions available at one time
  - **Uncongested**
    - Enough time for offers to be made, accepted, rejected, transactions carried out…
  - **Safe**
    - Safe to participate, and to reveal relevant information
- Computers can help with each of these.