CAP 5993/CAP 4993
Game Theory

Instructor: Sam Ganzfried

sganzfri@cis.fiu.edu

Schedule

- Project presentations on 4/18 and 4/20.
  - Today: Abdullah, Harold, Daniel
  - Thursday: Mario, Amir, Efrain, Farzana, Mai, Bingqian
- Project writeup due at beginning of class on 4/20.
  - 10 page limit, pdf preferred

- Final exam on 4/25, 2:15-4:15PM, ECS 235.

- Candidate C is a Condorcet winner. He would defeat A by a vote of 13 to 8 if the two of them were the sole candidates, and similarly would win by 13 votes to 8 votes against B.

No. committee members	First choice	Second choice	Third choice
1	A	B	C
7	A	C	B
7	B	C	A
6	C	B	A

- Who is the Condorcet winner now?

No. committee members	First choice	Second choice	Third choice
23	A	B	C
2	B	A	C
17	B	C	A
10	C	A	B
8	C	B	A

- Trick question, there isn't one!
- A defeats B by 33-27, B trounces C by 42-18, and C wins against A by 35-25.

No. committee members	First choice	Second choice	Third choice
23	A	B	C
2	B	A	C
17	B	C	A
10	C	A	B
8	C	B	A

Simple Majority Rule

- Suppose there are only two alternatives $A = \{a,b\}$. For each strict preference profile P^N we will denote the number of individuals who prefer a to b by $m(P^N)$. The simple majority rule is the social welfare function F defined by:
  - If $m(P^N) > n/2$ then society as a whole prefers a to b.
  - If $m(P^N) < n/2$ then society as a whole prefers b to a.
  - If $m(P^N) = n/2$ then society as a whole is indifferent between a and b.

Arrow’s Impossibility Theorem

- Theorem [Arrow 1951]: If |A| >= 3, then every social welfare function satisfying the properties of unanimity and independence of irrelevant alternatives is dictatorial.
  - Unanimity: if all individuals in society prefer a to b, then society also prefers a to b.
  - IIA: Whether a is preferable to b depends only on the way individuals compare a to b. E.g., Ann ranked higher than Dan, then Tanya’s grade changed (because she retook an exam), this should have no effect on relative ranking of Ann and Dan.
  - Dictatorship: a single voter has the power to always determine the group’s preferences.

Gibbard-Satterthwaite Theorem

- Theorem [1973,1975]: Let G be a nonmanipulable social choice function satisfying the unanimity property. If |A| >= 3 then G is dictatorial.

- If we wish to apply a nondicatorial social choice function, there are necessarily situations in which one (or more) of the individuals has an incentive to report a preference relation that is different from his or her true preference relation.

Borda method

- Every voter ranks the candidates, from most preferred to least preferred. A candidate receives k points (called Borda points) from a voter if that voter ranks the candidate higher than exactly k other candidates. The Borda ranking of a candidate is given by the total number of Borda points he receives from all the voters. The winning candidate (called the Borda winner) is then the candidate who has amassed the most Borda points.

Range Voting

- **Range voting** or **score voting** is a voting method for single-seat elections, in which voters give each candidate a score, the scores are added (or averaged), and the candidate with the highest total is elected.

- Sports such as gymnastics rate competitors on a numeric scale, although the fact that judges' ratings are public makes it less likely for them to engage in blatant tactical voting.

- The preferences of the women appear in the lower right-hand side of each cell and preferences of the men appear in the upper left-hand side. For example, Adam is 2nd on Anne's preference list, and Anne is fourth on Adam's preference list.

	Anne	Bess	Carol	Donna
Adam	4 ★ / 2	3 / 2	2 / ♣ 2	1 / 4
Ben	2 / ♣ 1	1 ★ / 4	3 / 4	4 / 3
Charles	4 / 3	1 / ♣ 1	3 ★ / 3	2 / 2
Dean	1 / 4	4 / 3	3 / 1	2 ★ / ♣ 1

- The matching depicted by the stars is not stable. This is because Carol and Adam have an objection to the matching. Carol prefers Adam (number 2 on her list) to Charles (number 3 on her list), and Adam prefers Carol (number 2 on his list) to Anne (number 4 on his list).

	Anne	Bess	Carol	Donna
Adam	4 ★ / 2	3 / 2	2 / ♣ 2	1 / 4
Ben	2 / ♣ 1	1 ★ / 4	3 / 4	4 / 3
Charles	4 / 3	1 / ♣ 1	3 ★ / 3	2 / 2
Dean	1 / 4	4 / 3	3 / 1	2 ★ / ♣ 1

- The matching depicted by the clubs is stable. To see this, note that Anne, Bess, and Donna are all matched with men who are number 1 on their lists, so that none of them will object to the matching with any man. Carol is matched with Adam, who is number 2 on her list, and therefore the only possible objection she may have is with Dean, who is number 1 on her list. But Dean prefers Donna (number 2 on his list) to Carol (number 3 on his list). This matching is thus stable, because no pair consisting of a man and a woman has an objection.

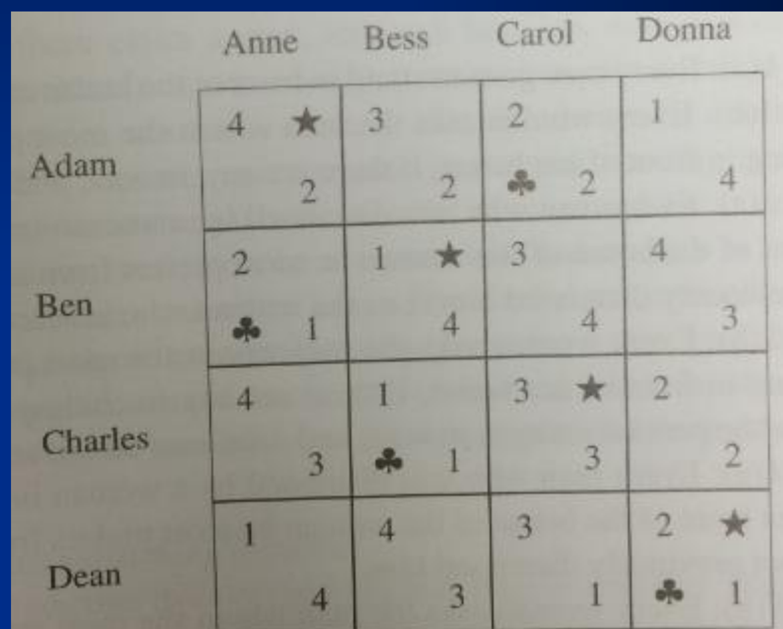

	Anne	Bess	Carol	Donna
Adam	4 ★ 2	3 2	2 ♣ 2	1 4
Ben	2 ♣ 1	1 ★ 4	3 4	4 3
Charles	4 3	1 ♣ 1	3 ★ 3	2 2
Dean	1 4	4 3	3 1	2 ★ ♣ 1

Matching problem

- Definition: a **matching problem** is given by:
  - A natural number n representing the number of men and the number of women in a population (thus, we assume that the number of women equals the number of men).
  - Every woman has a preference relation over the set of men.
  - Every man has a preference relation over the set of women.

- Preference relation is complete, irreflexive, and transitive

- Definition: A **matching** is a bijection from the set of men to the set of women.

- Definition: A man and a woman **object** to a matching if they prefer each other to the mates to whom they are matched under the matching. A matching is **stable** if there is no pair consisting of a man and a woman who have an objection to the mapping.

- Equivalent definition: A matching A is **stable** if in every case that a man prefers another woman to the woman to whom he is matched under A, that woman prefers the man to whom she is matched to him.

Gale-Shapley Algorithm [1962]

- Theorem: To every matching problem there exists a stable matching.

- Proof: Gale-Shapley Algorithm

“Men’s courtship algorithm”

- Stage 1(a): Every man goes to stand in front of the house of the woman he most prefers.
- Stage 1(b): Every woman asks the man whom she most prefers from among the men standing in front of her house, if there are any, to wait, and dismisses all the other men.
- Stage 2(a): Every man who was dismissed by a woman in the first stage goes to stand in front of the house of the woman he most prefers from among the women who have not previously dismissed him (i.e., the woman who is second on his list).
- Stage 2(b): Every woman asks the man whom she most prefers from amont the men standing in front of her house, if there are any (including the man whom she told to wait in the previous stage), to wait, and dismisses all the other men.

- In general:
- Stage k(a): Every man who was dismissed by a woman in the previous stage goes to stand in front of the house of the woman he most prefers from among the women who have not previously dismissed him.
- Stage k(b): Every woman asks the man whom she most prefers from among the men standing in front of her house, if there are any, to wait, and dismisses all the other men.
- The algorithm terminates when there is one man standing in front of every woman’s house.
- Can prove that the algorithm will always terminate (full proof in textbook), and every woman will have one man standing in front of her house. Can prove that the algorithm always terminates by finding a stable matching.

Example 22.6 (*Continued*) We apply the matching algorithm to Example 22.6. The table in Figure 22. describes the run of the algorithm.

	Anne	Bess	Carol	Donna	Dismissed men
Stage 1(a)	Dean(1)	Ben(1), Charles(1)		Adam(1)	
Stage 1(b)	Dean(1)	Charles(1)		Adam(1)	Ben
Stage 2(a)	Dean(1), Ben(2)	Charles(1)		Adam(1)	
Stage 2(b)	Ben	Charles(1)		Adam(1)	Dean
Stage 3(a)	Ben(2)	Charles(1)		Adam(1), Dean(2)	
Stage 3(b)	Ben(2)	Charles(1)		Dean(2)	Adam
Stage 4(a)	Ben(2)	Charles(1)	Adam(2)	Dean(2)	

Figure 22.3 Men's courtship algorithm

Each stage of the run of the algorithm is described by two consecutive rows. In the row corresponding to part (a) we note the men who are standing in front of women's houses after the application of part (a) of that stage (prior to the dismissals of the women), and in the row corresponding to part (b) we note the men who are standing in front of women's houses after the application of part (b) of that stage (after the women have announced their dismissals). The number appearing by the name of each man represents the ranking of the woman in front of whose house he is standing, in his preference relation.

The algorithm ends with the matching:

(Adam – Carol, Ben – Anne, Charles – Bess, Dean – Donna),

which is the stable matching previously mentioned.

“Women’s courtship algorithm”

- Same algorithm with role of men and women reversed. By proof of the same theorem, this algorithm will also lead to a stable matching.

- Theorem: The men’s courtship matching is the best stable matching from the perspective of all the men, and the worst from the perspective of all the women; the women’s courtship matching is the best stable matching from the perspective of all the women, and the worst from the perspective of all the men. This holds true for every matching problem.
  - Every man who is matched to a different women in the two algorithms prefers the woman from the men’s courtship matching, and vice versa.

- The number of possible matchings equals n!, which rapidly grows large as n grows. In principle the number of stable matchings could still be rather small. However, Gusfield and Irving [1989] show that matching problems may have a very large number of stable matchings. Using their method, one can construct a matching problem with 8 men and 8 women, with 268 stable matchings; a matching problem with 16 men and 16 women, with 195,472 stable matchings; and a matching problem with 32 men and 32 women, with 104,310,534,400 stable matchings.

National Resident Matching Program

- United States-based private non-profit non-governmental organization created in 1952 to help match medical school students with residency programs.
- http://www.nrmp.org/match-a-to-z/video-tutorials/about-the-matching-algorithm-tutorial/

- The problem of matching hospitals to residents is a generalization of the stable marriage problem; and, as a result, the solutions to the two problems are very similar. A simplified version of the algorithm that is used to perform the match is described on the NRMP website. However, this description does not describe the handling of couples (pairs of applicants who wish to stay in the same geographic location), second-year positions, or special handling of residency positions that remain unfilled. The full algorithm is described in *Roth, Alvin; Elliott Peranson (September 1999). "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design"*.

- In order to understand how the algorithm works, it is helpful to begin by considering simpler case where there are no couples or secondary programs.
- As in the stable marriage problem, the basic goal in the simple case of the hospitals/residents problem is to match applicants to hospitals so that the final result is "stable". "Stability" in this case means that there is no applicant A and hospital H such that both of the following are true:
  - A is unmatched or would prefer to go to H over the hospital he is currently matched with
  - H has a free slot or would prefer A over one of the candidates currently filling one of its slots.

- It can be shown that for any instance of the problem, there is at least one valid solution. Under the old (pre-1995) NRMP algorithm, which favored hospitals over residents, in certain cases hospitals could benefit from lying about their preferences, but that is no longer true under the new system. In neither system can a resident or coalition of residents benefit simply by lying about their preferences, even if they have perfect knowledge of everyone's preferences. (Of course, both systems are susceptible to other forms of collusion. For example, if two applicants apply to the same program, the weaker is still capable of bribing the stronger into ranking the program lower on his list than he would otherwise.)
- Under the current system, it is impossible for an applicant to be harmed by including more residency programs at the bottom of his list if those programs are indeed preferable to not being matched.

- Adding couples who submit joint preference lists complicates the problem significantly. In some cases there exists no stable solution (with stable defined similarly to the way it is in the simple case). In fact, the problem of determining whether there is a stable solution and finding it if it exists has been proven NP-complete. As a result, the algorithm used by the NRMP is not strictly guaranteed to return a result in a reasonable amount of time, even if one exists. Also, while there is no randomization in the NRMP algorithm—so it will always return the same output when given exactly the same input—different outcomes can be produced by changing trivial features of the data such as the order in which applicants and programs are processed. However, in initial testing of the algorithm over 5 years of residency match data and a variety of different initial conditions, the current NRMP algorithm always terminated quickly on a stable solution. Testing also showed that "none of [the trivial] sequencing decisions had a large or systematic effect on the matching produced"—the maximum number of applicants ever observed to be affected in a single run was 12 out of 22,938.

- In general *once the hospitals' preference lists have been set*, there is no way for an applicant to match into a better position by deciding to match in a couple. For example, if a very strong applicant and a very weak applicant match as a couple, there is no mechanism *in the algorithm* that allows the stronger applicant to somehow improve the desirability of the weaker applicant.[(Of course, if the hospitals know that the stronger and weaker applicant are matching together prior to the run of the algorithm, they are always free to change their preference lists accordingly, which will obviously affect the final outcome.) Ensuring the members of the couple end up in compatible programs is essentially achieved by matching them individually and having them turn down programs, moving on to less desirable ones until their positions are acceptable, though the algorithm does not function exactly in this way. As a result, all else being equal, couples are relatively likely to be matched with less desirable programs than they would have been had they decided to match individually.

- However, it is not impossible for couples to match into better positions than they would have individually. First, since the algorithm does depend on some arbitrary factors (e.g. the order in which applicants are processed), one or both of the individuals could end up in a better position by chance alone, although this is extremely unlikely. Second, anything that affects the hospitals' preference lists prior to the run of the algorithm will obviously affect the final outcome, as in the case with the strong and weak applicant above. There is also some belief that being a part of a couple may be appealing in and of itself.

- It is possible for a medical student not to be matched to a program. Until the 2010 match, students who did not match went through a process called the Scramble. In this process, students were forced to apply en masse to whatever programs remained available, frequently having to change their intended specialty in the process. This worked in the following fashion: at noon the day after Match Day, the NRMP released a list of unfilled programs. Students would then apply both directly and through ERAS (Electronic Residency Application Service, the same process used for the Match) in substantial chaos. Four days after the Match, the Scramble ended. Most residencies filled within the first few hours of the Scramble, and nearly all in the first 48 hours. Scrambling was extremely competitive: in 2008, roughly 13,000 applicants, many of whom were foreign trained, scrambled for only 1,388 residencies.

- This process was widely seen as needlessly stressful and in need of improvement, and thus after the 2010 match the Scramble was replaced with the Supplemental Offer and Acceptance Program, or SOAP. SOAP functions with eight rounds of matching following the main Match Day match, and creates a systematic way for non-matched students to find residencies without the chaos of the Scramble. The primary changes were as follows: all matching takes place under the NRMP (no direct matching allowed, unlike the Scramble), all matching uses ERAS, and unmatched applicant data and unmatched program data is released at the same time, not program data a day later, as before, among other changes.

Kidney exchange

- **Kidney Paired Donation (KPD)** or **Paired Exchange**, is an approach to living donor kidney transplantation where patients with incompatible donors swap kidneys to receive a compatible kidney. KPD is used in situations where a potential donor is incompatible. Because better donor HLA and age matching are correlated with lower lifetime mortality and longer lasting kidney transplants, many compatible pairs are also participating in swaps to find better matched kidneys. In the United States, the National Kidney Registry organizes the majority of U.S. KPD transplants, including the largest swaps. The first large swap was a 60 participant chain in 2012 that appeared on the front page of the New York Times and the second, even larger swap, included 70 participants and was completed in 2014. Other KPD programs in the U.S. include the UNOS program which was launched in 2010 and completed its 100th KPD transplant in 2014 and the Alliance for Paired Donation.

- More than one-third of potential living kidney donors who want to donate their kidney to a friend or family member cannot donate due to blood type or antibody incompatibility. Historically, these donors would be turned away and the patient would lose the opportunity to receive a life-saving transplant. KPD overcomes donor-recipient incompatibility by swapping kidneys between multiple donor-recipient pairs. KPD is also being used to find better donor-recipient matches for compatible pairs who want a lower lifetime mortality and longer lasting transplant.

Matching students to classes

- Each student registers for more than one course
- Some courses meet at same time and/or have capacity constraints
- Initial approach: students bid on courses with fake money
  - Incentive issues and ultimately led to low student satisfaction
- Developed better approach that computes "approximate competitive equilibrium from Equal Incomes" (A-CEEI)
  - Strengths: deals with efficiency, fairness, and incentives
  - Issues: capacity constraints violated, not scalable (optimization problem is PPAD-hard), assumes students can report preferences *accurately*
  - Mixed-integer program to calculate market clearing price for each course
- CourseMatch solves all these issues:
  - https://www.youtube.com/watch?v=u9gwHW87ZJQ
  - Used for assignment of courses at UPenn's Wharton business school