APPLIED MECHANISM DESIGN FOR SOCIAL GOOD

JOHN P DICKERSON

DUNCAN MCELFRESH

Lecture #1 – 01/28/2020

CMSC828M
Tuesdays & Thursdays
2:00pm – 3:15pm

COMPUTER SCIENCE
UNIVERSITY OF MARYLAND

Markets come in many forms …

… some of which don't conform to
conventional notions of markets …

… and some in which money may play little or no role.

– excerpt from *Who Gets What – and Why*

AESTHETICALLY-DISPLEASING ONE-SLIDE COURSE SUMMARY

Mechanism design is a field in economics and game theory that takes an engineering approach to designing economic mechanisms or incentives, toward desired objectives, in strategic settings, where players act rationally.

AKA "reverse game theory" – we're designing the rules of the game, not figuring out how to play it.

This course: can we design mechanisms for societal problems that perform well in practice, are computationally tractable, and whose workings and results are understandable by humans?

More info: https://marketdesign.github.io/

PREREQUISITE KNOWLEDGE

Aimed at Ph.D. students – but likely accessible to others with mathematical maturity and interest.

Assuming:

- Basic CS undergrad knowledge of theory (correctness proofs, NP-hardness, impossibility results)
- Basic CS/Math undergrad knowledge of optimization
  - Convex optimization, LPs, IPs – one lecture of primer, too
- Basic CS/Math undergrad knowledge of statistics and probability
- Ability to consume scientific papers (CS, Econ, OR)

All of this can be learned on the fly! Recommended books …

BOOK #1: "THE AGT BOOK"

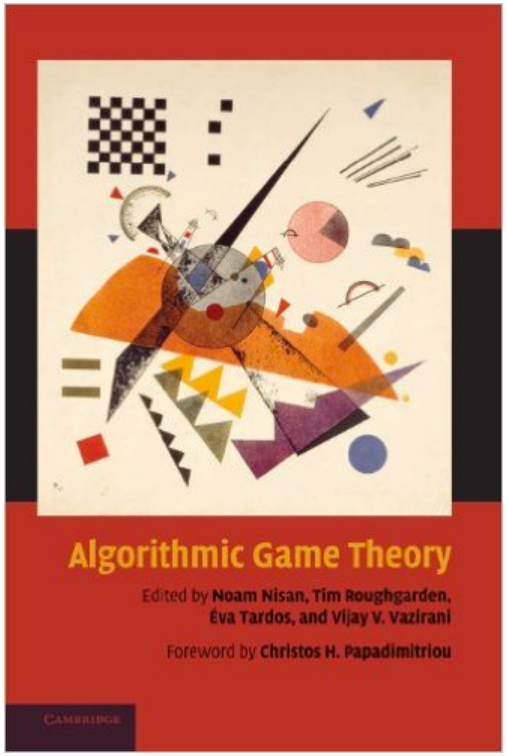

It's free online! Check the course webpage.

BOOK #2: “CSC HANDBOOK”

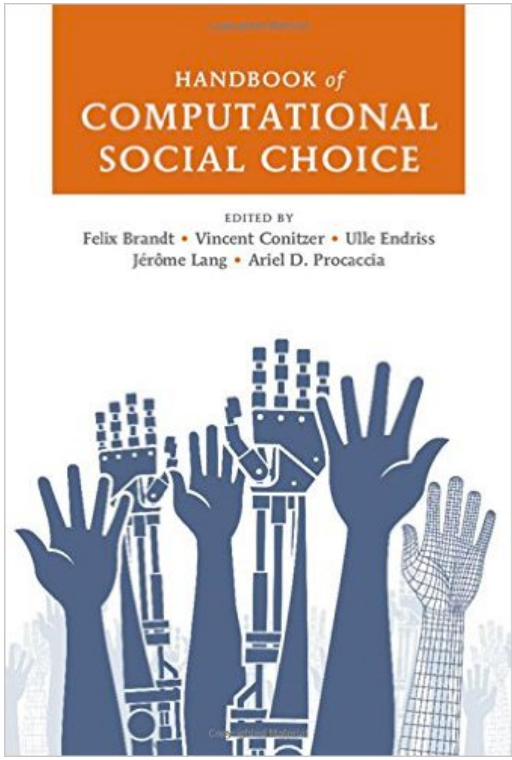

It’s free online! Check the course webpage.

BOOK #3: "MATCHING MARKET DESIGN"

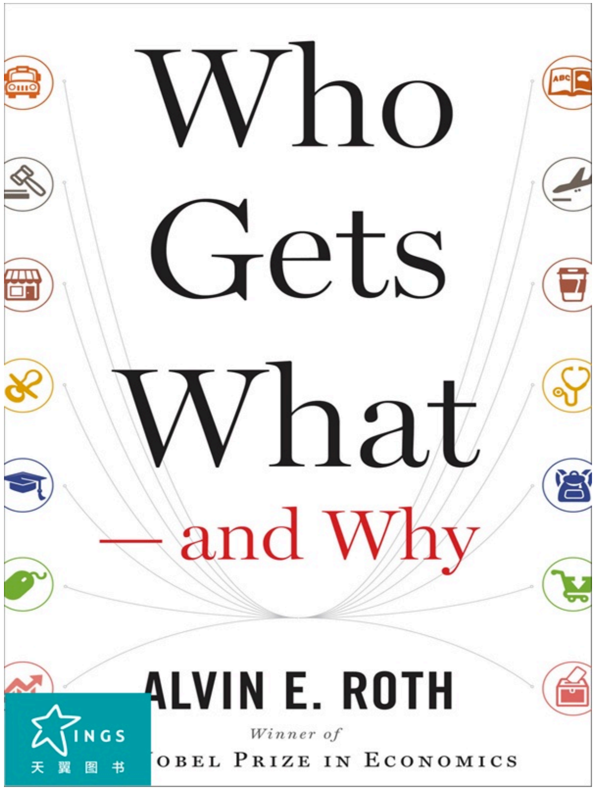

It's not free online! UMD has one copy.

BOOK #4: "COMBINATORIAL AUCTIONS"

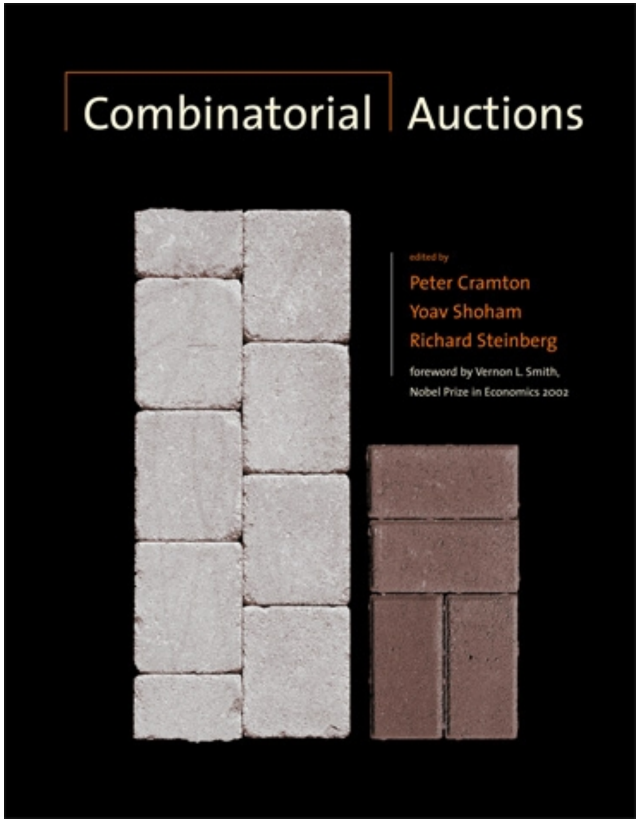

It's free online! Check the course webpage.

BOOK #5: "THE ETHICAL ALGORITHM"

It's not free online! UMD might have a copy?

SINGLE-SLIDE VERSION OF JOHN

UNOS DONATE LIFE®
UNITED NETWORK FOR ORGAN SHARING

COMPUTER SCIENCE
UNIVERSITY OF MARYLAND
2004 – 2008

Carnegie Mellon University
School of Computer Science
2010 – 2016

COMPUTER SCIENCE
UNIVERSITY OF MARYLAND
2016 – ∞

UMIACS
HCIL
CENTER FOR MACHINE LEARNING
APPLIED MATHEMATICS & STATISTICS, AND SCIENTIFIC COMPUTATION PROGRAM

- facebook research 2018 – 2020
- { Advisory roles in UMD alumni's startups } 2017 – now
- { Advisory role to a US regulatory agency } Late 2019 +

OPTIMIZED MARKETS

NATIONAL SECURITY AGENCY
UNITED STATES OF AMERICA

IBM

THE DIAN FOSSEY GORILLA FUND® INTERNATIONAL

WHO ARE YOU?

Area?

Interest in this course?

Goals in life? And goals for this course?

PhD/MSc/BSc?

Background in: CS, game theory, mechanism design, HCI?

Advisor?

I'm probably going to use Slack instead of Piazza – thoughts?

COURSE STRUCTURE

First 5-6 lectures: primers in GT/MD and optimization

Next 23 lectures: theory + application areas

- Security games
- Healthcare markets (organ allocation, financing R&D)
- Food allocation
- Combinatorial assignment (bidding on courses)
- Incentive auctions (FCC spectrum allocation)
- Fair Division (allocating rooms to schools)
- Voting (#YangGang)
- School choice (assigning kids to schools)
- Prediction markets (#YangGang?)
- 2-3 lectures of "floating topics" – ideas?

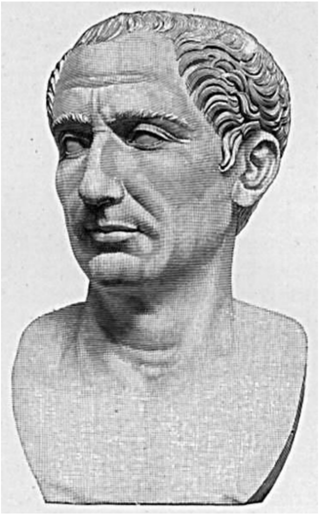

Ambitious …

GRADE #1: PROJECT

Students will complete a **semester-long course project** on something related to market and mechanism design.

- Individual or small group
- 100% theory, 100% applied, or convex combination

Goal: create something **publishable**!

Important dates:

- Project proposals will be due in early March
- ~~Project presentations will be during the last 2-3 lectures~~
- Project writeups—formatted as, say, a NeurIPS conference paper or similar—will be due by the exam date for this course (Monday, May 18 at 10:30AM).

A FEW PROJECTS FROM RECENT YEARS …

Some projects from recent years that have been published:

- Faez Ahmed: diverse b-matching (IJCAI-17, arXiv)
- Neal Gupta: competition in dynamic matching markets (IJCAI-18)
- Duncan McElfresh: a new fairness metric for kidney exchange (AAAI-18, arXiv)
- Elissa Redmiles: behavioral economics and two-factor authentication (EC-18 & CCS-18, arXiv)
- Candice Schumann: diversity and multi-armed bandits (AAMAS-19 & NeurIPS-19, arXiv & arXiv)

GRADE #2: PRESENT A PAPER OR SCRIBE/SUMMARIZE

Students will present for at least ½ of a lecture on one of the assigned papers, or write a summary of that paper that is fit to be posted on the course webpage:

- (Welcome to do an entire class, but talk to me beforehand!)
- John or Duncan covers the basics of a topic in the first half of lecture

Check out the course webpage for topics

- Also: feel free to suggest a topic you like!

Logistics:

- We'll figure this out in second week
- I'll post some example presentations

GRADE #3: TAKE-HOME EXAMS

The majority of students want this to be an MSc/PhD qualifier

- So this course is an AI qualifier for CMSC grad students

What this means:

- **"…** the courses' written exam(s) which must account for at least 30% of the grade." – *UMD graduate handbook*

Implementation:

- One take-home midterm; 24 hours to do it
- One take-home final; 48 hours to do it
- I'll let you pick when you take exams (within reason; see site)

GRADE #4: CLASS PARTICIPATION

Please please please please please read the paper(s) before coming to class!

- I want to speak with you, not at you …
- Present ½ of the class on the paper – JPD gives intro

Participate on Slack, help brainstorm project ideas, ask questions in class, answer questions in class, etc.

(These should be easy points for everyone.)

GRADE BREAKDOWN | QUALIFIER APPROVAL

60% course project:

- **5% project proposal (early March)**
- **5% project checkup (April)**
- **50% course project writeup (mid-May)**

10% presenting/summarizing a paper

30% examinations:

- **Better of 10%/20% or 15%/15% midterm/final**

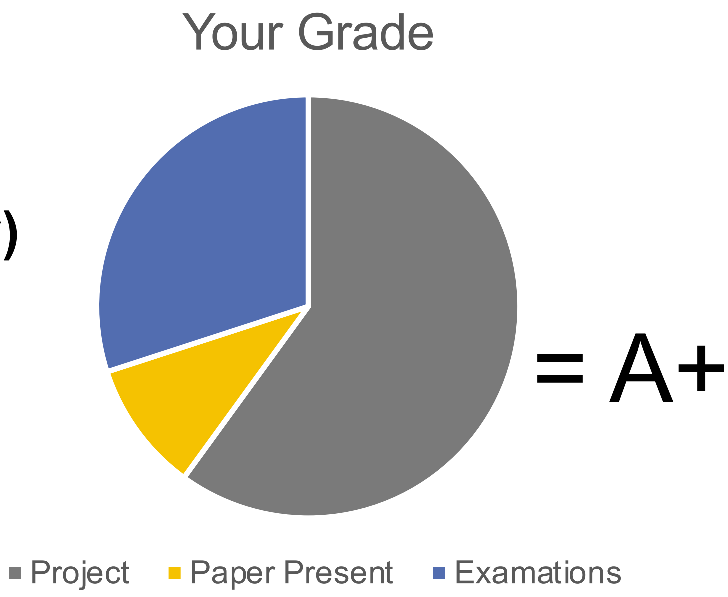

IMPORTANT WALLS OF TEXT

ANTI-HARASSMENT

(Adapted from ACM SIGCOMM's policies)

Common Sense!

The open exchange of ideas and the freedom of thought and expression are central to our aims and goals. These require an environment that recognizes the inherent worth of every person and group, that fosters dignity, understanding, and mutual respect, and that embraces diversity. For these reasons, we are dedicated to providing a harassment-free experience for participants in (and out) of this class.

Harassment is unwelcome or hostile behavior, including speech that intimidates, creates discomfort, or interferes with a person's participation or opportunity for participation, in a conference, event or program.

ACADEMIC INTEGRITY

Common Sense!

(Text unironically stolen from Hal Daumé III)

Any assignment or exam that is handed in must be your own work (unless otherwise stated). However, talking with one another to understand the material better is strongly encouraged. Recognizing the distinction between cheating and cooperation is very important. If you copy someone else's solution, you are cheating. If you let someone else copy your solution, you are cheating (this includes *posting solutions online in a public place)*. If someone dictates a solution to you, you are cheating.

Everything you hand in must be in your own words, and based on your own understanding of the solution. If someone helps you understand the problem during a high-level discussion, you are not cheating. We strongly encourage students to help one another understand the material presented in class, in the book, and general issues relevant to the assignments. When taking an exam, you must work independently. Any collaboration during an exam will be considered cheating. Any student who is caught cheating will be given an F in the course and referred to the University Office of Student Conduct. Please don't take that chance – if you're having trouble understanding the material, please let me know and I will be more than happy to help.

EXAMPLE APPLICATION AREAS

MANY OF WHICH WE'LL COVER IN FUTURE CLASSES

EXAMPLE: MATCHING MARKETS

In matching problems, prices do not do all – or any – of the work

Agents are paired with other (groups of) agents, transactions, or contracts

- Workers to firms
- Children to schools
- Residents to hospitals
- Patients to deceased donors
- Advertisements to viewers
- Riders to rideshare drivers

UNCERTAINTY IN MATCHING MARKETS

- **Does a matched edge truly exist?**
- **How valuable is a match?**
- **Will a better match arrive in the future?**

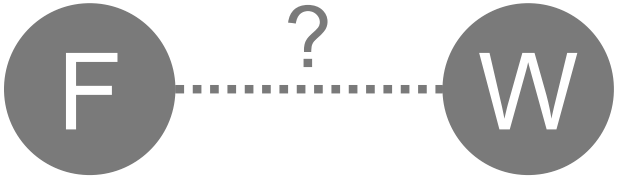

COMPETITION IN MATCHING MARKETS

Rival matching markets compete over the same agents

- **How does this affect global social welfare?**
- **How to differentiate?**

CADENCE OF MATCHING MARKETS

How quickly do new edges form?

How frequently does a market clear?

Is clearing centralized or decentralized?

Can agents reenter the market?

EXAMPLE: RESIDENT-HOSPITAL ASSIGNMENT

1940s: decentralized resident-hospital matching

- Market "unraveled", offers came earlier and earlier, quality of matches decreased

1950s: NRMP introduces hospital-proposing deferred acceptance algorithm

1970s: couples increasingly don't use NRMP

1998: matching with couple constraints

- (Stable matching may not exist anymore …)

Take-home message

Looks like: M.D.s aren't the only type of doctor who help people!

EXAMPLE: COMBINATORIAL COURSE ALLOCATION

[IMAGES FROM BUDISH ET AL. WORKING PAPER 2016]

BIDDING

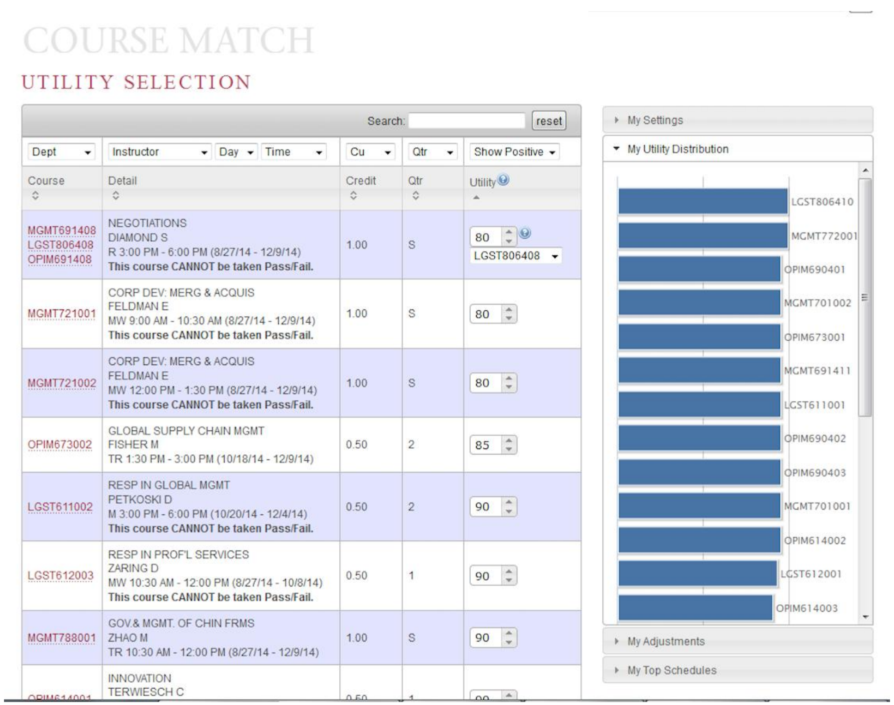

“Funny money” used for bidding

DYNAMIC EXCLUSIONS

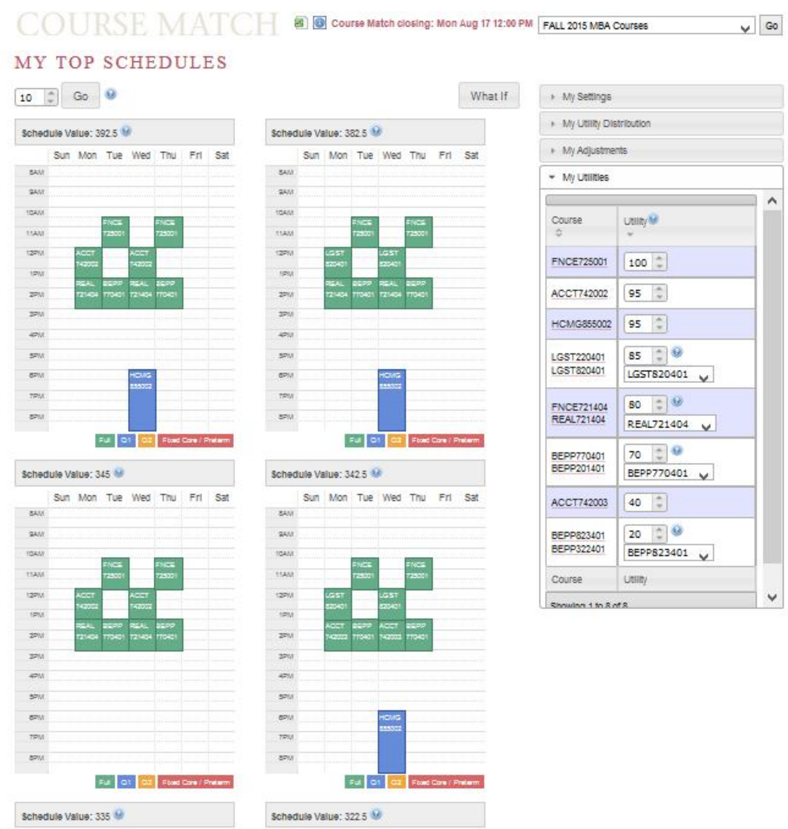

EXAMPLE: VOTING

Set of voters *N* and a set of alternatives:
{Hillary Clinton, The Donald, Gary Johnson}

Each voter ranks the candidates:

v_1: The Donald > HRC > Gary Johnson

v_2: HRC > Gary Johnson > The Donald

…

A preference profile is the set of all voters' rankings

Can we choose a voting rule – that is, a function that takes preference profiles and returns a winning alternative – that:

- "Behaves well"
- Isn't manipulable by strategic agents

EXAMPLE: FAIR ALLOCATION

http://spliddit.org

Divisible goods:

- Splitting land, cutting cake

Indivisible goods:

- Splitting up assets after divorce (house, cars, pets)

A chief concern: defining and guaranteeing the fairness of the final allocation

An allocation is envy free if each player values her own allocated set of goods at least as highly as any other player's allocated set

When do envy-free allocations exist? How can we compute them? What can we do when they don't exist?

EXAMPLE: FOOD BANK ALLOCATION

Food banks supply nutrition to the needy for free or at a reduced cost

- **Much of that food comes from donors (e.g. supermarkets, manufacturers)**

Distribution is overseen by a large non-profit organization, Feeding America

- **Previously: centralized allocation based on perceived need of food banks**
- **Currently: food banks bid in an online auction using a fake currency for loads of donated food.**

EXAMPLE: SECURITY GAMES

Where should we deploy security forces (checkpoints, cop cars, dogs, troops), assuming a rational adversary who can observe our deployment strategy?

- Checkpoints at airports
- Patrol routes on the water on the borders
- Anti-poacher teams near big game

How do we compute these strategies?

What if the adversary isn't rational?

Rangers Use Artificial Intelligence to Fight Poachers

Emerging technology may help wildlife officials beat back traffickers.

EXAMPLE: KIDNEY TRANSPLANTATION

- **US waitlist: over 100,000**
  - 36,157 added in 2014
- **4,537 people died while waiting**
- **11,559 people received a kidney from the deceased donor waitlist**

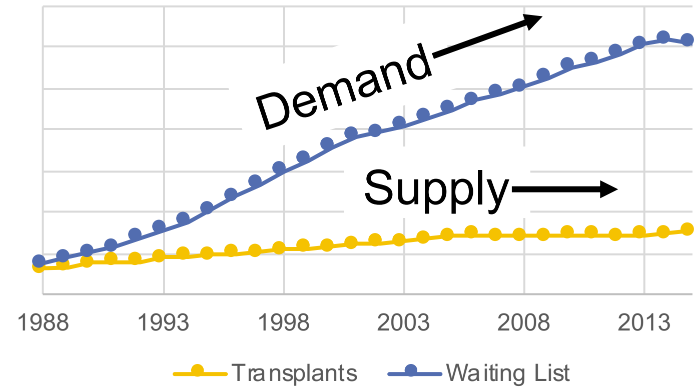

- **5,283 people received a kidney from a living donor**
  - Some through **kidney exchanges** [Roth et al. 2004]
  - (We work extensively with the UNOS exchange.)

UNOS DONATE LIFE

EXAMPLE: DECEASED-DONOR ALLOCATION

Online bipartite matching problem:

- Set of patients is known (roughly) in advance
- Organs arrive and must be dispatched quickly

Constraints:

- Locality: organs only stay good for 24 hours
- Blood type, tissue type, etc.

Who gets the organ? Prioritization based on:

- Age?
- QALY maximization?
- Quality of match?
- Time on the waiting list?

EXAMPLE: KIDNEY EXCHANGE

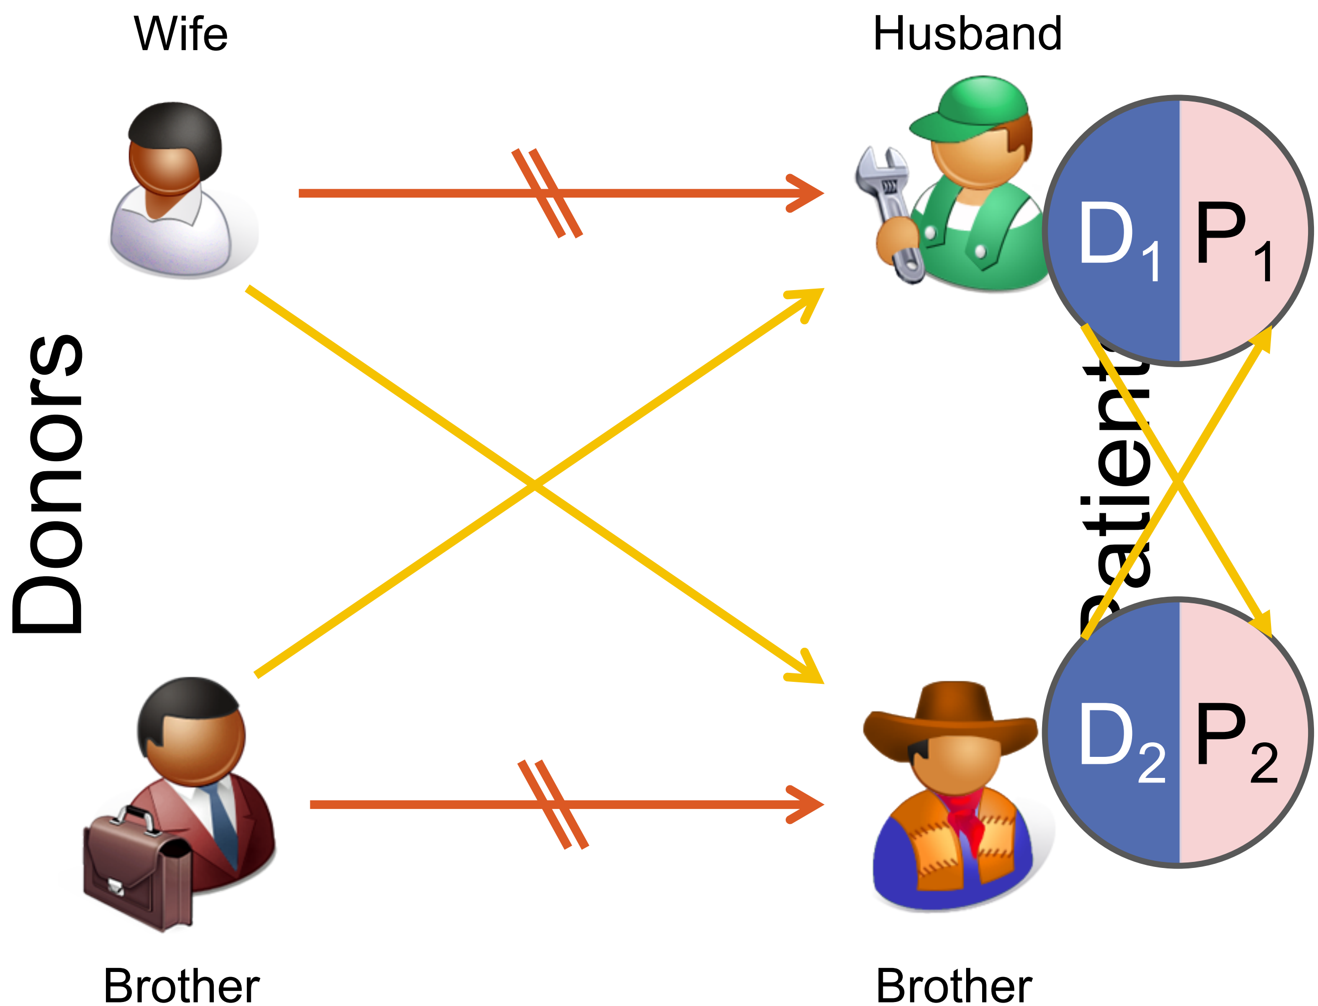

(2- and 3-cycles, all surgeries performed simultaneously)

NON-DIRECTED DONORS & CHAINS

[Rees et al. 2009]

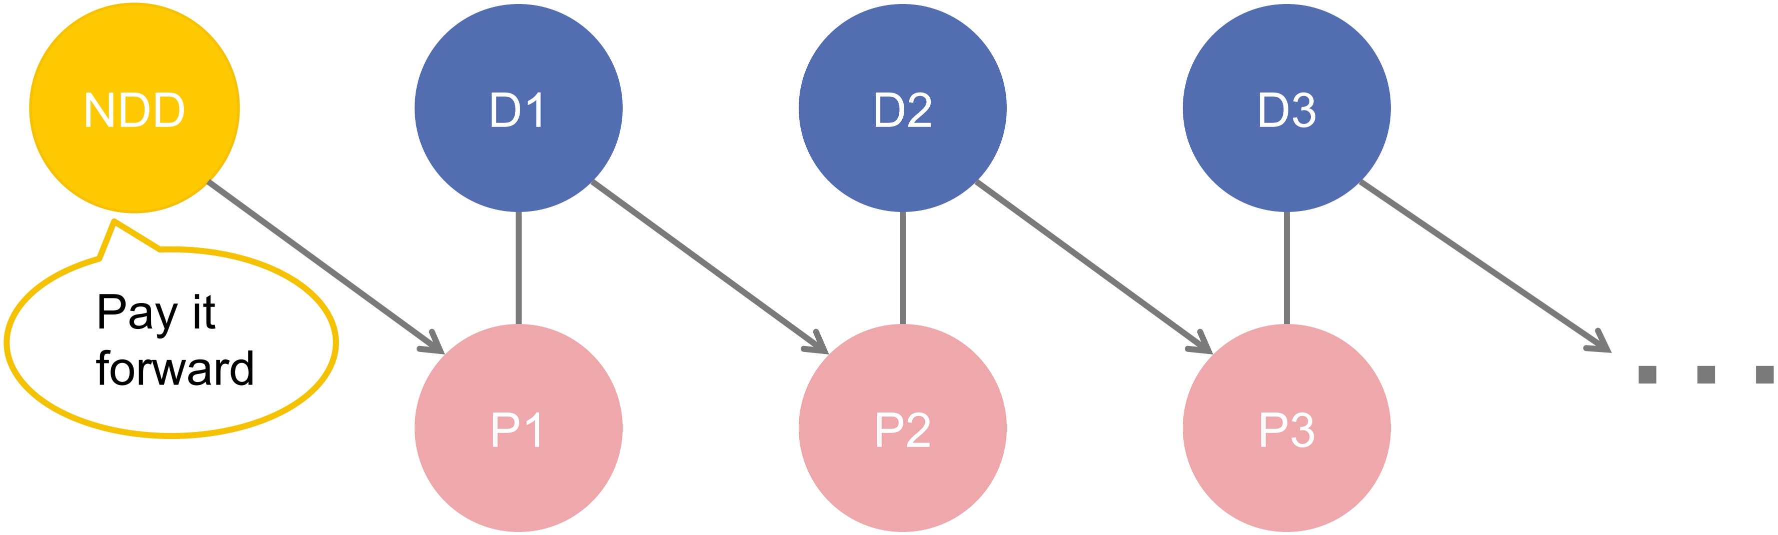

Not executed simultaneously, so no length cap required based on logistic concerns …

… **but in practice edges fail, so *some* finite cap is used!**

EXAMPLE: KIDNEY EXCHANGE

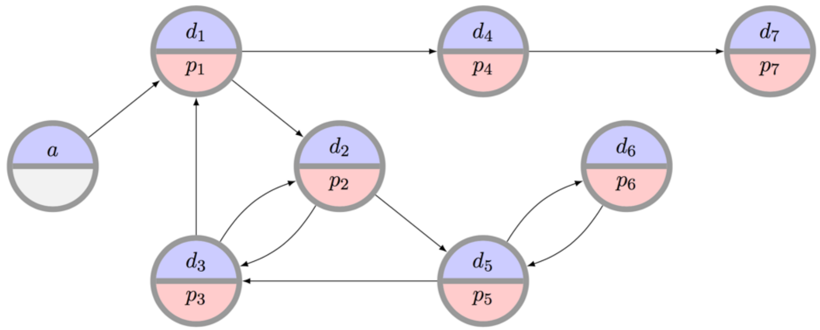

What is the "best" matching objective?

- Maximize matches right now or over time?
- Maximize transplants or matches?
- Prioritization schemes (i.e. fairness)?
- Modeling choices?
- Incentives? Ethics? Legality?

Can we design a mechanism that **performs well in practice**, is **computationally tractable**, and is **understandable by humans**?

TECHNIQUES WE'LL USE

(NEXT THREE LECTURES WILL COVER THESE, IN THE CONTEXT OF MECHANISM DESIGN)

COMBINATORIAL OPTIMIZATION

Combinatorial optimization lets us select the "best element" from a set of elements.

Some PTIME problems:

- **Some forms of matching**
- **2-player zero-sum Nash**
- **Compact LPs**

Some PPAD- or NP-hard problems:

- **More complex forms of matching**
- **Many equilibrium computations**

Some > NP-hard problems:

- **Randomizing over a set of all feasible *X*, where all feasible *X* must be enumerated (#P-complete)**

C.O. FOR KIDNEY EXCHANGE: THE EDGE FORMULATION

[Abraham et al. 2007]

Binary variable x_{ij} for each edge from *i* to *j*

Maximize

$u(M) = \Sigma\, w_{ij}\, x_{ij}$

Flow constraint

Subject to

$\Sigma_j\, x_{ij} = \Sigma_j\, x_{ji}$ for each vertex *i*

$\Sigma_j\, x_{ij} \leq 1$ for each vertex *i*

$\Sigma_{1 \leq k \leq L}\, x_{i(k)i(k+1)} \leq L\text{-}1$ for paths i(1)…i(L+1)

(no path of length L that doesn't end where it started – cycle cap)

C.O. FOR KIDNEY EXCHANGE: THE CYCLE FORMULATION

[Roth et al. 2004, 2005, Abraham et al. 2007]

Binary variable x_c for each feasible cycle or chain c

Maximize

$$u(M) = \Sigma\, w_c\, x_c$$

Subject to

$$\Sigma_{c\,:\,i \text{ in } c}\, x_c \leq 1 \text{ for each vertex } i$$

C.O. FOR KIDNEY EXCHANGE: COMPARISON

Tradeoffs in number of variables, constraints

- IP #1: $O(|E|^L)$ constraints vs. $O(|V|)$ for IP #2
- IP #1: $O(|V|^2)$ variables vs. $O(|V|^L)$ for IP #2

IP #2's relaxation is weakly tighter than #1's. Quick intuition in one direction:

- Take a length L+1 cycle. #2's LP relaxation is 0.
- #1's LP relaxation is $(L+1)/2$ – with ½ on each edge

Recent work focuses on balancing tight LP relaxations and model size [Constantino et al. 2013, Glorie et al. 2014, Klimentova et al. 2014, Alvelos et al. 2015, Anderson et al. 2015, Mak-Hau 2015, Manlove&O'Malley 2015, Plaut et al. 2016, …]**:**

- We will discuss (9/29) new compact formulations, some with tightest relaxations known, all amenable to failure-aware matching

GAME THEORY & MECHANISM DESIGN

We assume participants in our mechanisms are:

- Selfish utility maximizers
- Rational (typically – sometimes relaxed)

Game theory & M.D. give us the language to describe desirable properties of mechanisms:

- Incentive compatibility
- Individual rationality
- Efficiency

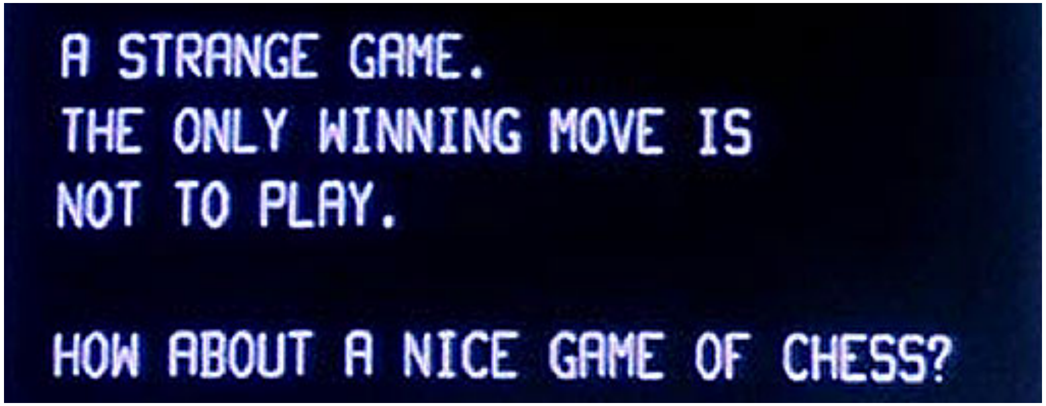

MACHINE LEARNING

Predicting supply and demand

Computing optimal matching/allocation policies:

- MDPs
- RL
- POMDPs, if you're feeling brave/masochistic

Aside: recent work looks at fairness and discrimination in machine learning – could be an interesting project.

- "… when a search was performed on a name that was "racially associated" with the black community, the results were much more likely to be accompanied by an ad suggesting that the person had a criminal record—regardless of whether or not they did."

CAN COMPUTERS BE RACIST?

Big data, the internet, and the law

RANDOM GRAPH THEORY

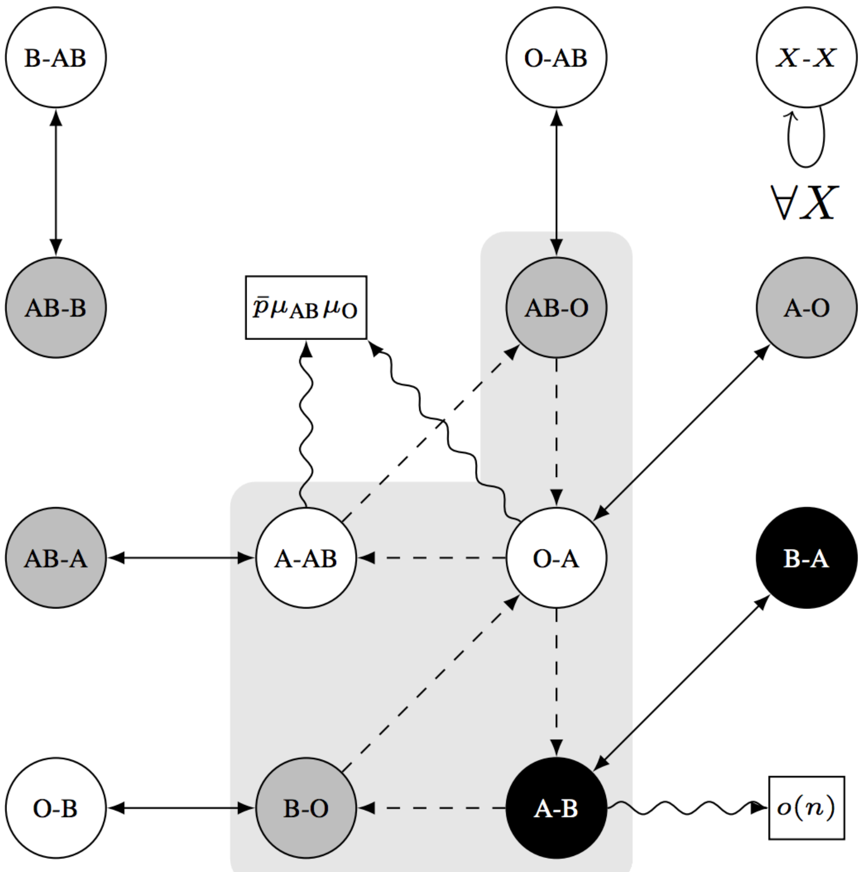

NEXT CLASS:
GAME THEORY PRIMER