APPLIED MECHANISM DESIGN FOR SOCIAL GOOD

JOHN P DICKERSON

Lecture #2 – 01/30/2020

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

COMPUTER SCIENCE
UNIVERSITY OF MARYLAND

ANNOUNCEMENTS

Everyone who is registered for the course should have received an invitation to the CMSC828M Slack channel

If you did not:

- Check an email you don't normally check; or
- https://join.slack.com/t/cmsc828m/signup

Should auto-join if you have a @cs.umd.edu email (feel free to add your friends outside of the course!)

WRAPPING UP FROM LAST LECTURE …

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

Rangers Use Artificial Intelligence to Fight Poachers

Emerging technology may help wildlife officials beat back traffickers.

How do we compute these strategies?

What if the adversary isn't rational?

EXAMPLE: KIDNEY TRANSPLANTATION

- **US waitlist: over 100,000**
  - 35-37k added each year
- **4,537 people died while waiting**
- **11,559 people received a kidney from the deceased donor waitlist**

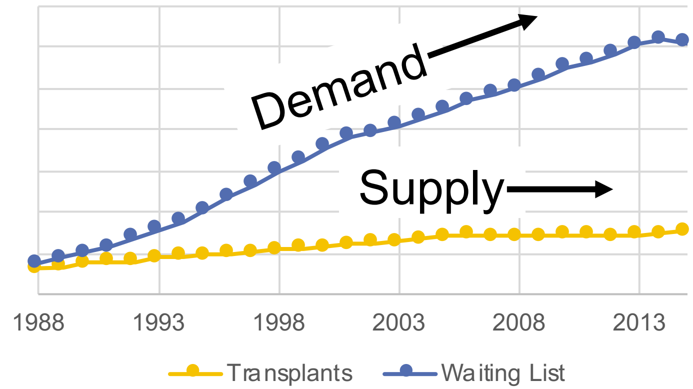

- **5,283 people received a kidney from a living donor**
  - Some through **kidney exchanges** [Roth et al. 2004]
  - (We work extensively with the UNOS exchange.)

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

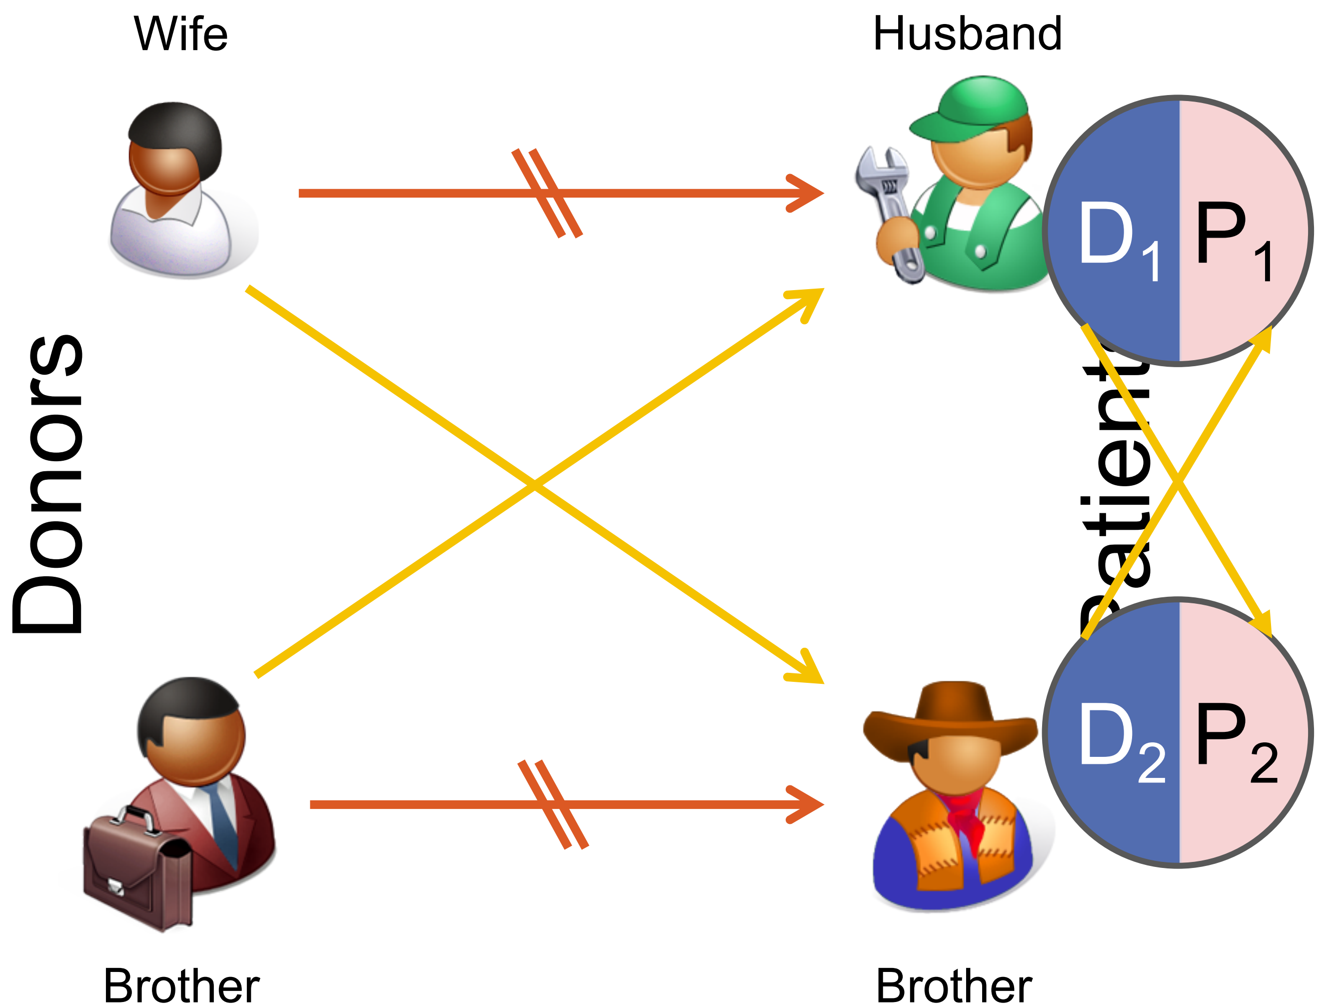

(2- and 3-cycles, all surgeries performed simultaneously)

NON-DIRECTED DONORS & CHAINS

[Rees et al. 2009]

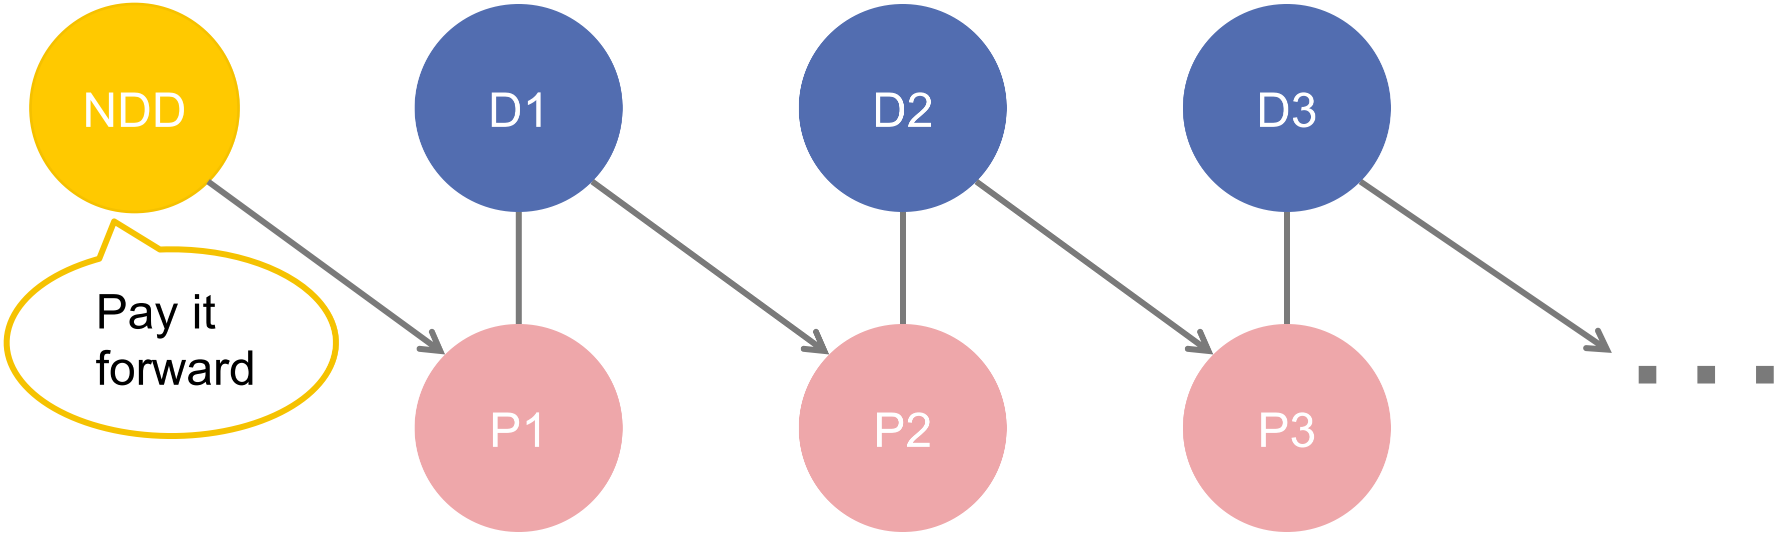

Not executed simultaneously, so no length cap required based on logistic concerns …

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

EXAMPLE: KIDNEY EXCHANGE

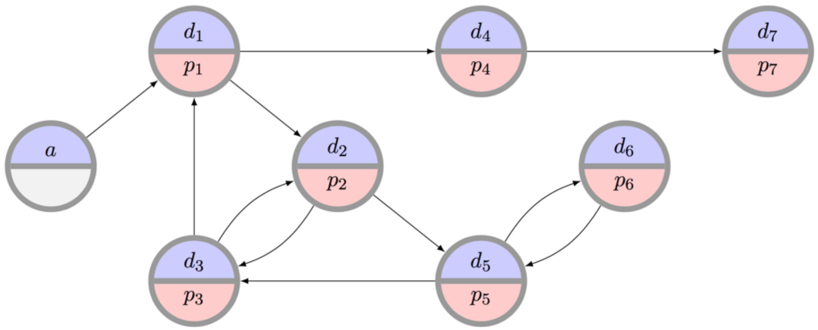

What is the "best" matching objective?

- 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

(THIS + NEXT TWO 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: REFRESHER ON PROBLEM

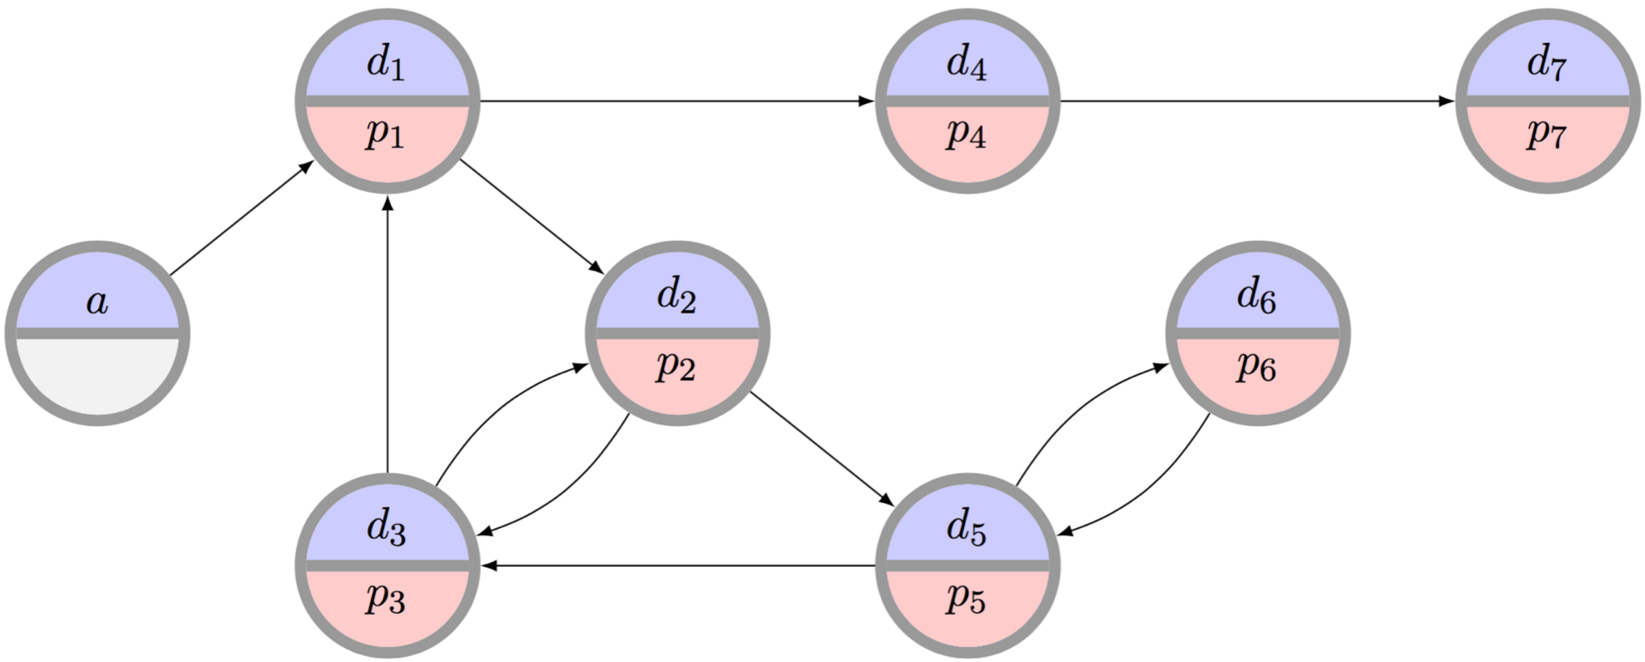

C.O. FOR KIDNEY EXCHANGE: THE EDGE FORMULATION

[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 (~about a month, possibly with Duncan) 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

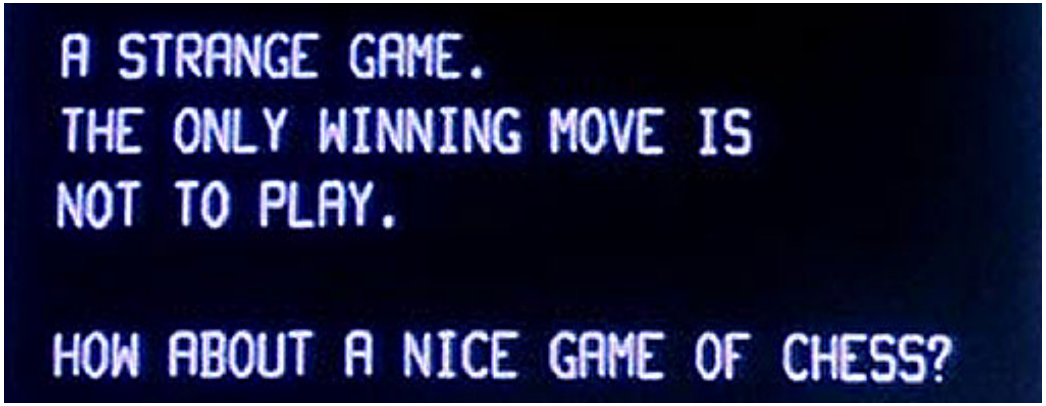

MACHINE LEARNING

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."

RANDOM GRAPH THEORY

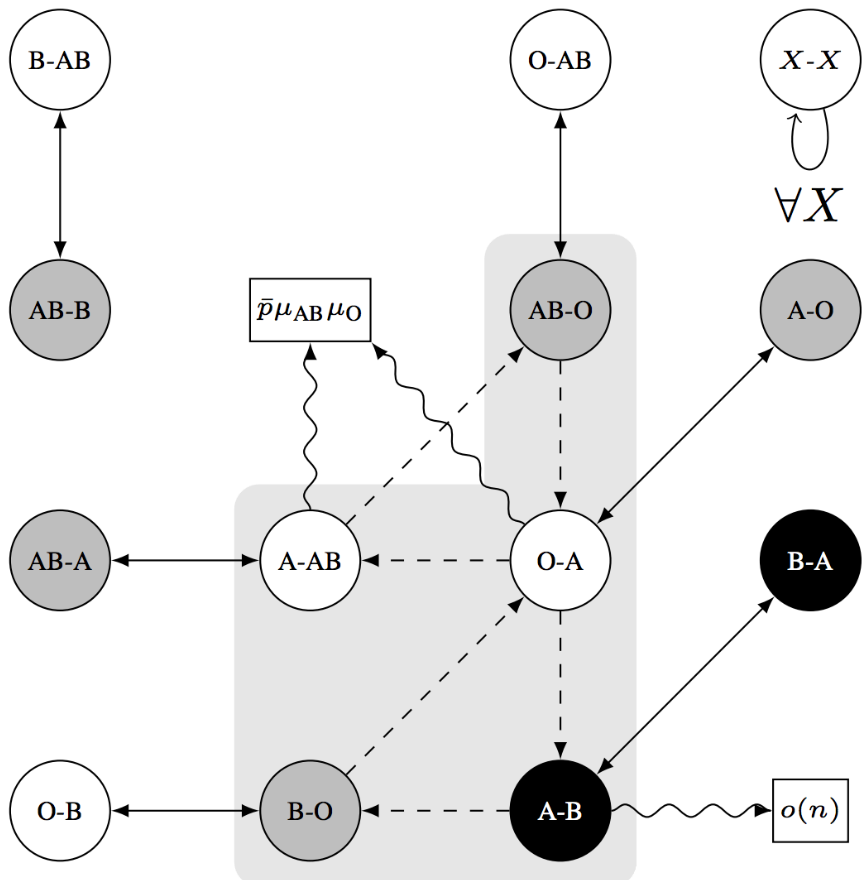

(Might cover a bit in the matching and barter exchange lectures; talk to me.)

~~NEXT~~ THIS CLASS: GAME THEORY PRIMER

Thanks to: AGT book, Blum (AB), Conitzer (VC), Sandholm (TS), Osborne&Rubinstein (OR)

WHAT IS GAME THEORY?

"… the study of mathematical models of conflict and cooperation between intelligent rational decision-makers."

"Intelligent rational decision-makers" = agents

- Have individual preferences specified by utility functions
- Can take different actions (or randomize over them)

UNIVERSITY OF MARYLAND

Utility of agents usually, but not always, depends on the actions of other agents

- What's best for me is a function of what's best for you …
  - … which is a function of what's best for me ...
    - ... which is a function of what's best for you ...
      - ... which is …

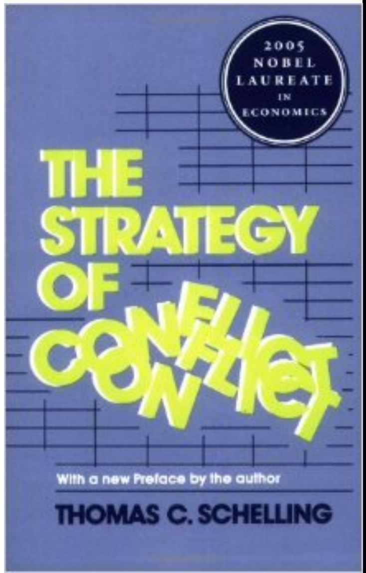

WHAT IS "UTILITY" …?

" … utility is a measure of preferences over some set of goods and services."

Formally:

- Let O be the set of outcomes
  (e.g., O = {{apple,orange}, {apple}, {orange}, { }})
- A utility function $u : O \rightarrow \Re$ ranks outcomes, and represents a preference relation $\preceq$ over the set of outcomes O

Example:

- u({apple,orange}) = 5
- u({apple}) = u({orange}) = 3
- u({ }) = 0 → { } $\prec$ {apple} $\preceq$ {orange} $\prec$ {apple,orange}

HOW DO WE MEASURE "UTILITY" …?

***u*({apple,orange}) = 5**

- 5 dollars? 5 clams? 5 days to live?
- Standard: 5 "utils" – it doesn't typically matter
- Agent's behavior under *u*(o) is typically the same as under *u*'(o) = a + b**u*(o)

***u*({apple}) = 3 < 5 = *u*({apple,orange})**

- Cardinal utility: 3 < 5
  - (We'll see this in security games and auctions)
- Ordinal utility: {apple,orange} $\prec$ {apple}
  - Doesn't encode strength of a preference, just ordering
  - (We'll see more of this in social choice)

RISK ATTITUDES

Which would you prefer?

- A lottery ticket that pays out $10 with probability .5 and $0 otherwise, or
- A lottery ticket that pays out $3 with probability 1

How about:

- A lottery ticket that pays out $100,000,000 with probability .5 and $0 otherwise, or
- A lottery ticket that pays out $30,000,000 with probability 1

Usually, people do not simply go by expected value

RISK ATTITUDES – EXPECTED VALUE

An agent is risk-neutral if she only cares about the expected value of the lottery ticket

An agent is risk-averse if she always prefers the expected value of the lottery ticket to the lottery ticket

- Most people are like this 

An agent is risk-seeking if she always prefers the lottery ticket to the expected value of the lottery ticket

DECREASING MARGINAL UTILITY

Typically, at some point, having an extra dollar does not make people much happier (decreasing marginal utility)

"Typically"

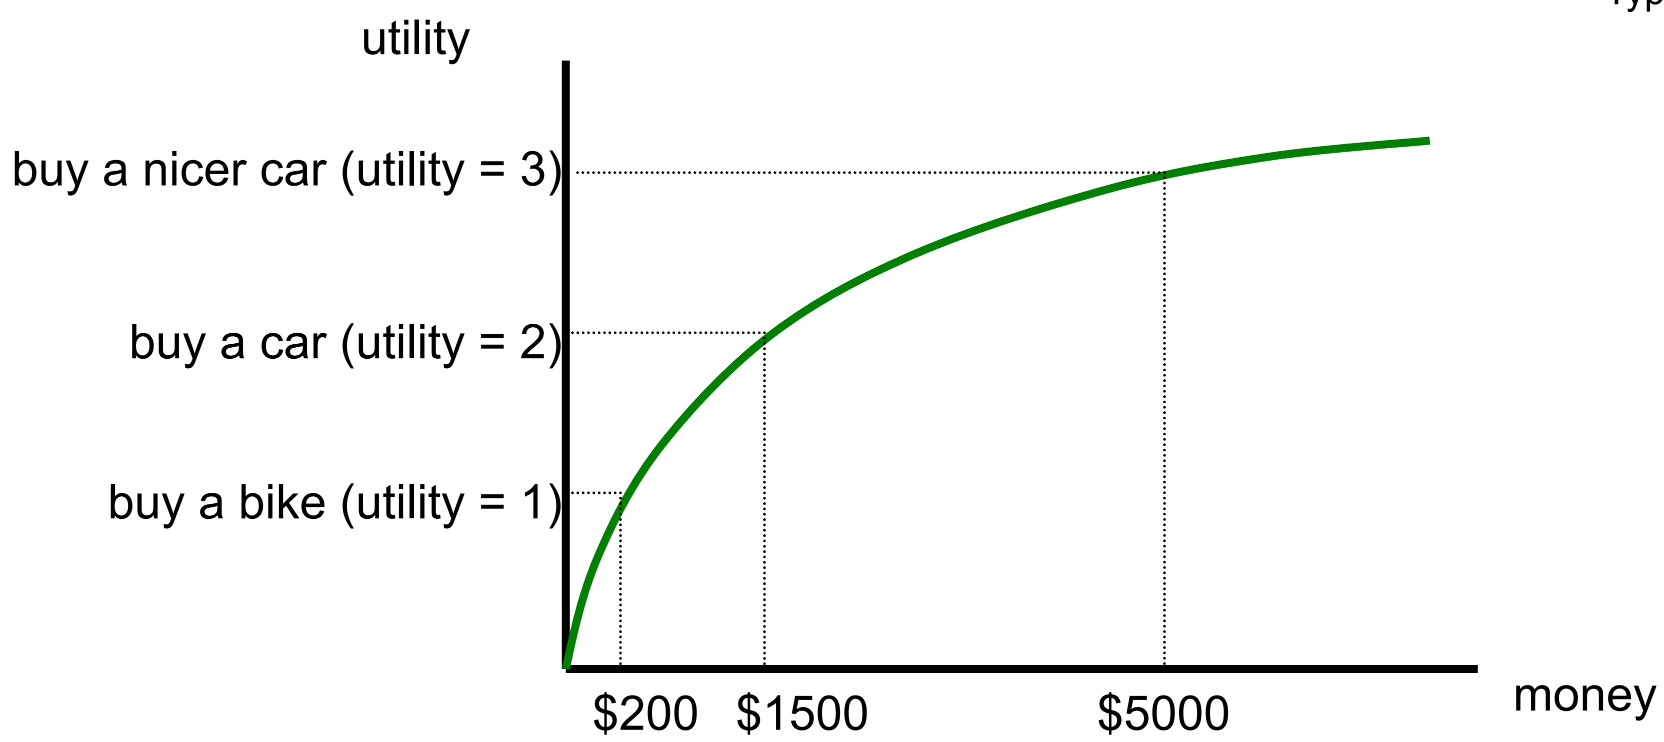

MAXIMIZING EXPECTED UTILITY

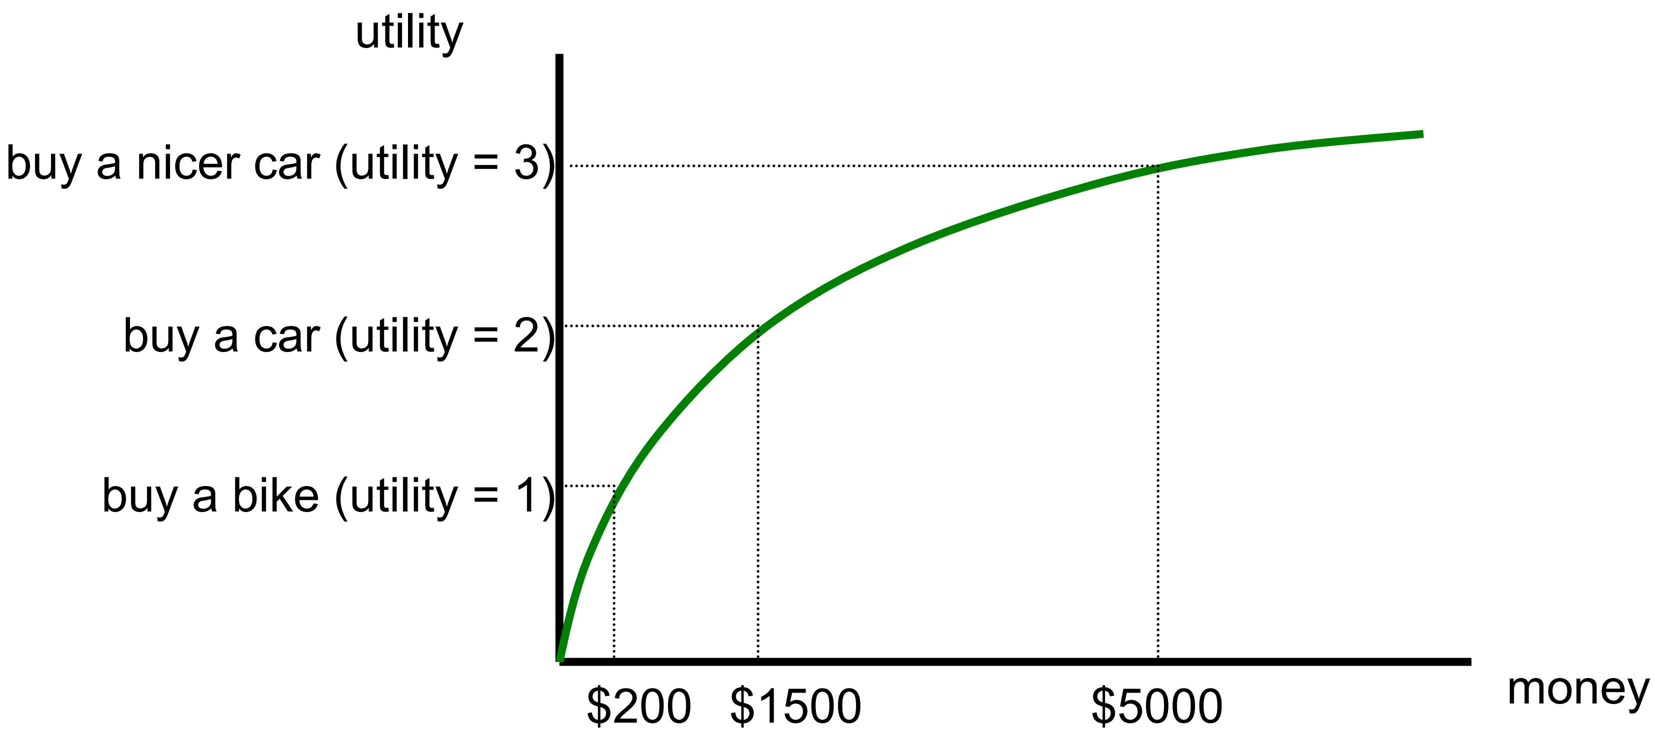

Lottery 1: get $1500 with probability 1 → gives expected utility 2

Lottery 2: get $5000 with probability .4, $200 otherwise

- → expected utility .4*3 + .6*1 = 1.8

$E_\$$[Lottery 2] = .4*$5000 + .6*$200 = $2120 > $1500 = $E_\$$[Lottery 1]

So: maximizing expected utility is consistent with risk aversion (assuming decreasing marginal utility)

RISK ATTITUDES ASSUMING EXPECTED UTILITY MAX'ING

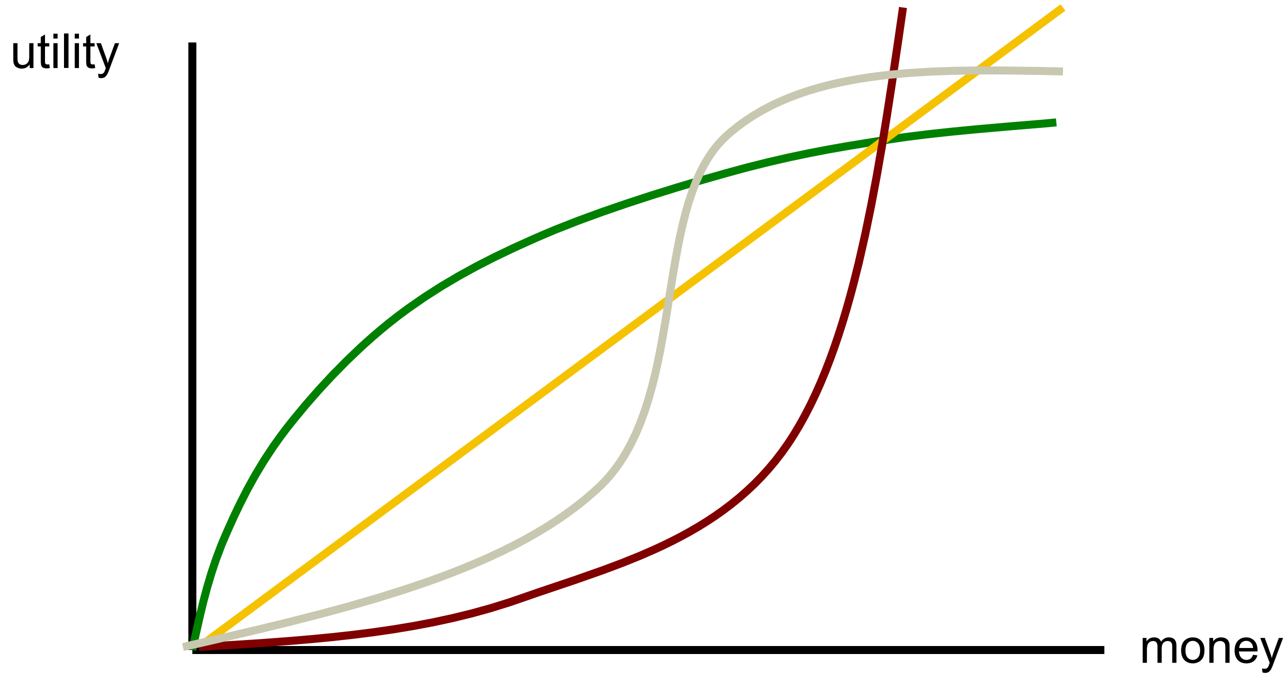

Green has decreasing marginal utility → risk-averse

Blue has constant marginal utility → risk-neutral

Red has increasing marginal utility → risk-seeking

Grey's marginal utility is sometimes increasing, sometimes decreasing → neither risk-averse (everywhere) nor risk-seeking (everywhere)

STRATEGIES & UTILITY

A strategy s_i for agent i is a mapping of history/the agent's knowledge of the world to actions

- Pure: "perform action x with probability 1"
- Randomized: "do x with prob 0.2 and y with prob 0.8"

A strategy set is the set of strategies available to agent i

- Can be infinite (infinite number of actions, randomization)

A strategy profile is an instantiation (s_1, s_2, s_3, …, s_N)

Abuse of notation: we'll use s_{-i} to refer to all strategies played other than that by agent i

- $i = 2$, then $s_{-i} = (s_1, s_3, ..., s_N)$

Utils awarded after game is played: $u_i = u_i(s_i, s_{-i})$

NATURE

Agents act strategically in the face of what they believe other agents will do, who act based on …

There may be other sources of non-strategic randomness

Included (when needed) in our models as a unique agent called nature, which acts:

- Probabilistically
- Without reasoning about what other agents will do

(Sometimes referred to as agent $i = 0$, often just nature.)

GAME REPRESENTATIONS

Column player aka. player 2 (simultaneously) chooses a column

Row player aka. player 1 chooses a row

		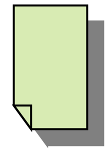	
	0, 0	-1, 1	1, -1
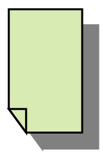	1, -1	0, 0	-1, 1
	-1, 1	1, -1	0, 0

A row or column is called an action or (pure) strategy

Row player's utility is always listed first, column player's second

Zero-sum game: the utilities in each entry sum to 0 (or a constant)

Three-player game would be a 3D table with 3 utilities per entry, etc.

GAME REPRESENTATIONS

Extensive form
(aka tree form)

Matrix form
(aka normal form
aka strategic form)

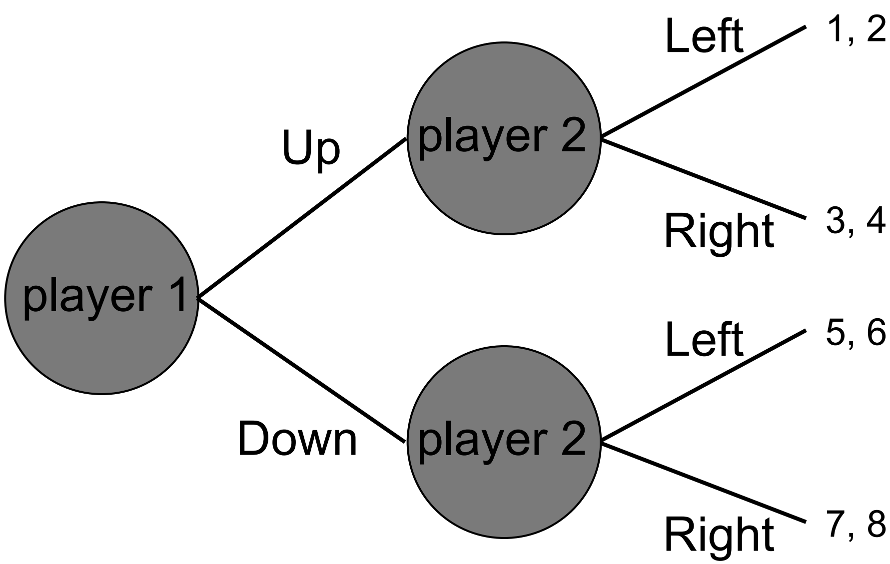

player 2's strategy

player 1's strategy	Left, Left	Left, Right	Right, Left	Right, Right
Up	1, 2	1, 2	3, 4	3, 4
Down	5, 6	7, 8	5, 6	7, 8

Potential combinatorial explosion →

SEINFELD'S ROCK-PAPER-SCISSORS

MICKEY: All right, rock beats paper!
(Mickey smacks Kramer's hand for losing)
KRAMER: I thought paper covered rock.
MICKEY: Nah, rock flies right through paper.
KRAMER: What beats rock?
MICKEY: (looks at hand) Nothing beats rock.

		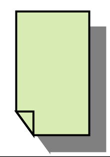	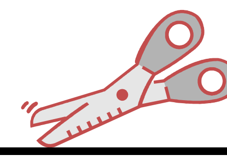
	0, 0	1, -1	1, -1
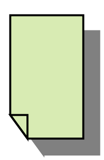	-1, 1	0, 0	-1, 1
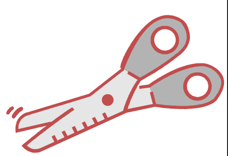	-1, 1	1, -1	0, 0

DOMINANCE

Player i's strategy s_i strictly dominates s_i' if

- for any s_{-i}, $u_i(s_i , s_{-i}) > u_i(s_i', s_{-i})$

s_i weakly dominates s_i' if

- for any s_{-i}, $u_i(s_i , s_{-i}) \geq u_i(s_i', s_{-i})$; and
- for some s_{-i}, $u_i(s_i , s_{-i}) > u_i(s_i', s_{-i})$

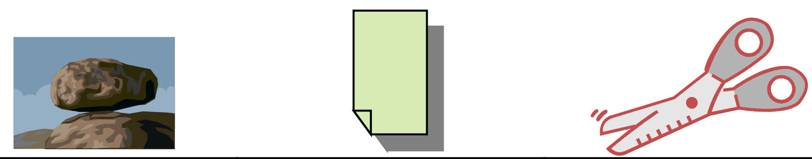

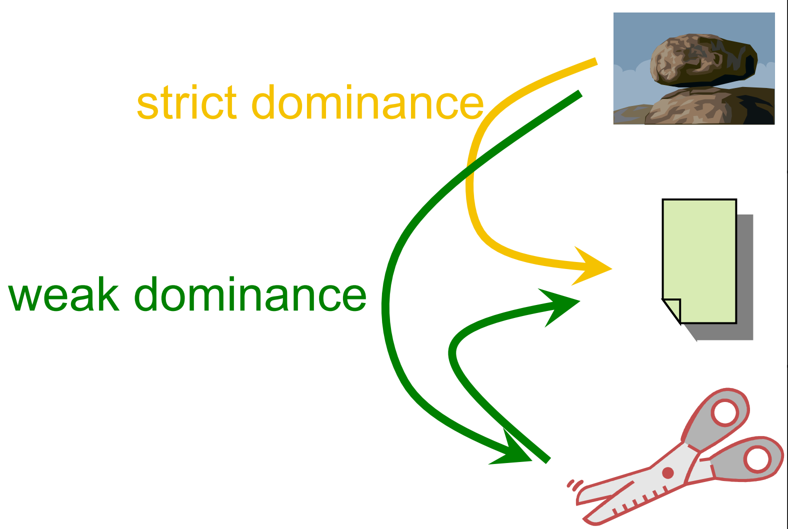

0, 0	1, -1	1, -1
-1, 1	0, 0	-1, 1
-1, 1	1, -1	0, 0

MIXED STRATEGIES & DOMINANCE

Mixed strategy for player i = probability distribution over player i's (pure) strategies

E.g.,1/3  **, 1/3** 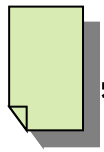 **, 1/3** 

Example of dominance by a mixed strategy:

1/2	3, 0	0, 0
1/2	0, 0	3, 0
	1, 0	1, 0

?????????

Usage:
σ_i denotes a mixed strategy, s_i denotes a pure strategy

BEST-RESPONSE STRATEGIES

Suppose you know your opponent's mixed strategy

- E.g., your opponent plays rock 50% of the time and scissors 50%

What is the best strategy for you to play?

Rock gives .5*0 + .5*1 = .5

Paper gives .5*1 + .5*(-1) = 0

Scissors gives .5*(-1) + .5*0 = -.5

So the best response to this opponent strategy is to (always) play rock

There is always some pure strategy that is a best response

- Suppose you have a mixed strategy that is a best response; then every one of the pure strategies that that mixed strategy places positive probability on must also be a best response

DOMINANT STRATEGY EQUILIBRIA (DSE)

Best response s_i^*: for all s_i', $u_i(s_i^*, s_{-i}) \geq u_i(s_i', s_{-i})$

Dominant strategy s_i^*: s_i^* is a best response for all s_{-i}

- Does not always exist
- Inferior strategies are called "dominated"

DSE is a strategy profile where each agent has picked its dominant strategy

- Requires no counterspeculation – just enumeration

	cooperate	defect
cooperate	3, 3	0, 5
defect	5, 0	1, 1

Pareto optimal?

Social welfare maximizing?

TS

ZERO-SUM GAMES (2-P)

Two-player zero-sum games are a special – purely competitive – case of general games

- **Everything I win you lose, and vice versa**

+1, -1	-2, +2
+2, -2	0, 0

Example: heads-up poker (with no rake)

A minimax-optimal strategy is a strategy that maximizes the expected minimum gain

- **Guarantees the "best minimum" in expectation, no matter which strategy your opponent selects**

Theorem [von Neumann '28] – "Minimax Theorem":

- **Every 2-P zero-sum game has a unique value V**
- **Maximin utility:** $\max_{\sigma_i} \min_{s_{-i}} i(\sigma_i, s_{-i})$ $(= - \min_{\sigma_i} \max_{s_{-i}} u_{-i}(\sigma_i, s_{-i}))$
- **Minimax utility:** $\min_{\sigma_{-i}} \max_{s_i} u_i(s_i, \sigma_{-i})$ $(= - \max_{\sigma_{-i}} \min_{s_i} u_{-i}(s_i, \sigma_{-i}))$
- **Theorem:** $V = \max_{\sigma_i} \min_{s_{-i}} u_i(\sigma_i, s_{-i}) = \min_{\sigma_{-i}} \max_{s_i} u_i(s_i, \sigma_{-i})$

GENERAL-SUM GAMES (2-P)

You could still play a minimax strategy in general-sum games

- i.e., pretend that the opponent is only trying to hurt you

But this is not rational:

Col

Row		
	0, 0	3, 1
	1, 0	2, 1

- If Col were trying to hurt Row, Col would play Left, so Row should play Down
- In reality, Col will play Right (strictly dominant), so Row should play Up
- Is there a better generalization of minimax strategies in zero-sum games to general-sum games?

GENERAL-SUM GAMES: NASH EQUILIBRIA (2-P)

Nash equilibrium: a pair of strategies that are stable

Stable: neither agent has incentive to deviate from his or her selected strategy on their own

????????

Row / Col

Row	2, 2	-1, -1
	-1, -1	2, 2

Theorem [Nash 1950]: any general-sum game has at least one Nash equilibrium

- Might require mixed strategies (randomization)

Corollary for 2-P zero-sum games: Minimax Theorem!

- WLOG pick one of the NE, let V = value of Row player
- Assumed NE, so neither player can do better (even fully knowing the other player's mixed strategy!) → minimax-opt

EXAMPLE: CHICKEN

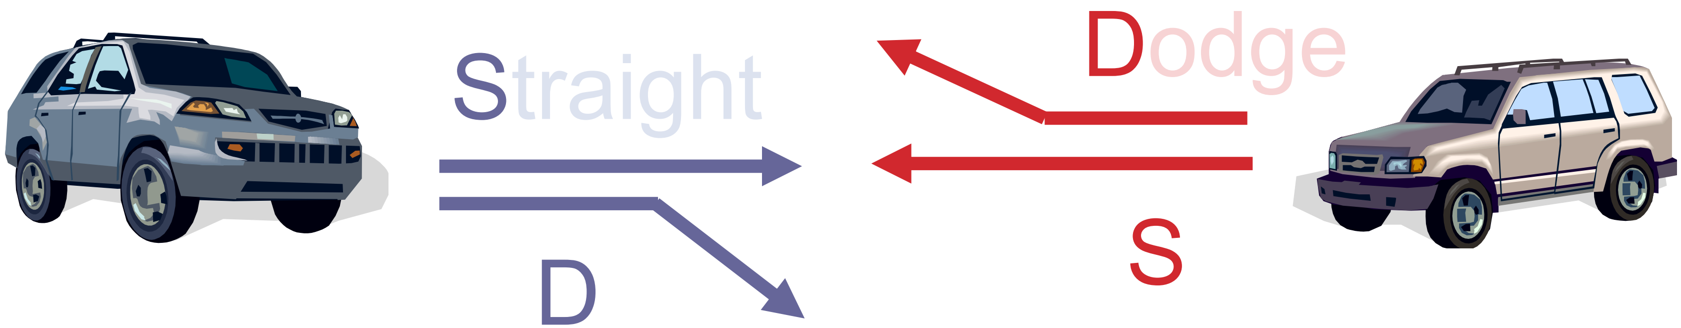

????????

????????

	D	S
D	0, 0	-1, 1
S	1, -1	-5, -5

- Thankfully, (D, S) and (S, D) are Nash equilibria
  - They are pure-strategy Nash equilibria: nobody randomizes
  - They are also strict Nash equilibria: changing your strategy will make you strictly worse off
- No other pure-strategy Nash equilibria

CHICKEN

	D	S
D	0, 0	-1, 1
S	1, -1	-5, -5

Is there an NE that uses mixed strategies?

- Say, where player 1 uses a mixed strategy?

- Note: if a mixed strategy is a best response, then all of the pure strategies that it randomizes over must also be best responses
- So we need to make player 1 indifferent between D and S
- Player 1's utility for playing D = $-p^c_S$
- Player 1's utility for playing S = $p^c_D - 5p^c_S = 1 - 6p^c_S$
- So we need $-p^c_S = 1 - 6p^c_S$ which means $p^c_S = 1/5$
- Then, player 2 needs to be indifferent as well
- Mixed-strategy Nash equilibrium: ((4/5 D, 1/5 S), (4/5 D, 1/5 S))
  - People may die! Expected utility -1/5 for each player

CRITICISMS OF NASH EQUILIBRIUM

Not unique in all games (like the example on Slide 31)

- Approaches for addressing this problem
  - Refinements (=strengthenings) of the equilibrium concept
    - Eliminate weakly dominated strategies first (IEDS)
    - Choose the Nash equilibrium with highest welfare
    - Subgame perfection … [see AGT book on course page]
  - Mediation, communication, convention, learning, …

Collusions amongst agents not handled well

- "No agent wants to deviate on her own"

Can be disastrous to "partially" play an NE

- (More) people may die!
- Correlated equilibria – strategies selected by an outsider, but the strategies must be stable (see Chp 2.7 of AGT)

CORRELATED EQUILIBRIUM

Suppose there is a trustworthy mediator who has offered to help out the players in the game

The mediator chooses a profile of pure strategies, perhaps randomly, then tells each player what her strategy is in the profile (but not what the other players' strategies are)

A correlated equilibrium is a distribution over pure-strategy profiles so that every player wants to follow the recommendation of the mediator (if she assumes that the others do so as well)

Every Nash equilibrium is also a correlated equilibrium

- Corresponds to mediator choosing players' recommendations independently

… but not vice versa

(Note: there are more general definitions of correlated equilibrium, but it can be shown that they do not allow you to do anything more than this definition.)

C.E. FOR CHICKEN

	D	S
D	0, 0 20%	-1, 1 40%
S	1, -1 40%	-5, -5 0%

Why is this a correlated equilibrium?

Suppose the mediator tells Row to Dodge

- From Row's perspective, the conditional probability that Col was told to Dodge is 20% / (20% + 40%) = 1/3
- So the expected utility of Dodging is (2/3)*(-1) = -2/3
- But the expected utility of Straight is (1/3)*1 + (2/3)*(-5) = -3
- So Row wants to follow the recommendation

If Row is told to go Straight, he knows that Col was told to Dodge, so again Row wants to follow the recommendation

Similar for Col

COMPLEXITY

Can compute minimax-optimal strategies in PTIME

Can compute 2-P zero-sum NE in PTIME

- (We'll see this as an example during the convex optimization primer lecture next week.)

Can compute correlated equilibria in PTIME

Unknown if we can compute a 2-P general-sum NE in PTIME:

- Known: PPAD-complete (weaker than NP-c, and different)
- All known algorithms require worst-cast exponential time

Our first "meaty" lectures will cover security games, which try to find Stackelberg equilibria:

- Varying complexity, will discuss during those lectures

DOES NASH MODEL HUMAN BEHAVIOR?

Game: pick a number (let's say, integer) in
{0, 1, 2, 3, …, 98, 99, 100}

Winner: person who picks number that is
closest to 2/3 of the average of all numbers

Example: if the average of all numbers is 54, your best answer would be 36 (= 54 * 2/3)

LIVE  EXPERIMENT!

DOES NASH MODEL HUMAN BEHAVIOR?

What's the (Nash) equilibrium strategy?

"Level 0" humans: everyone picks randomly? E[v] = 50, choose 50 * 2/3

"Level 1" humans: everyone picks 50 * 2/3, I'll pick (50 * 2/3) * 2/3

"Level 2" humans: I'll pick ((50 * 2/3) * 2/3) * 2/3 …

N.E.: fixed point, "Level infinity", pick 0 or 1 depending on constraints

DOES NASH MODEL HUMAN BEHAVIOR?

Any guesses on behavior …?

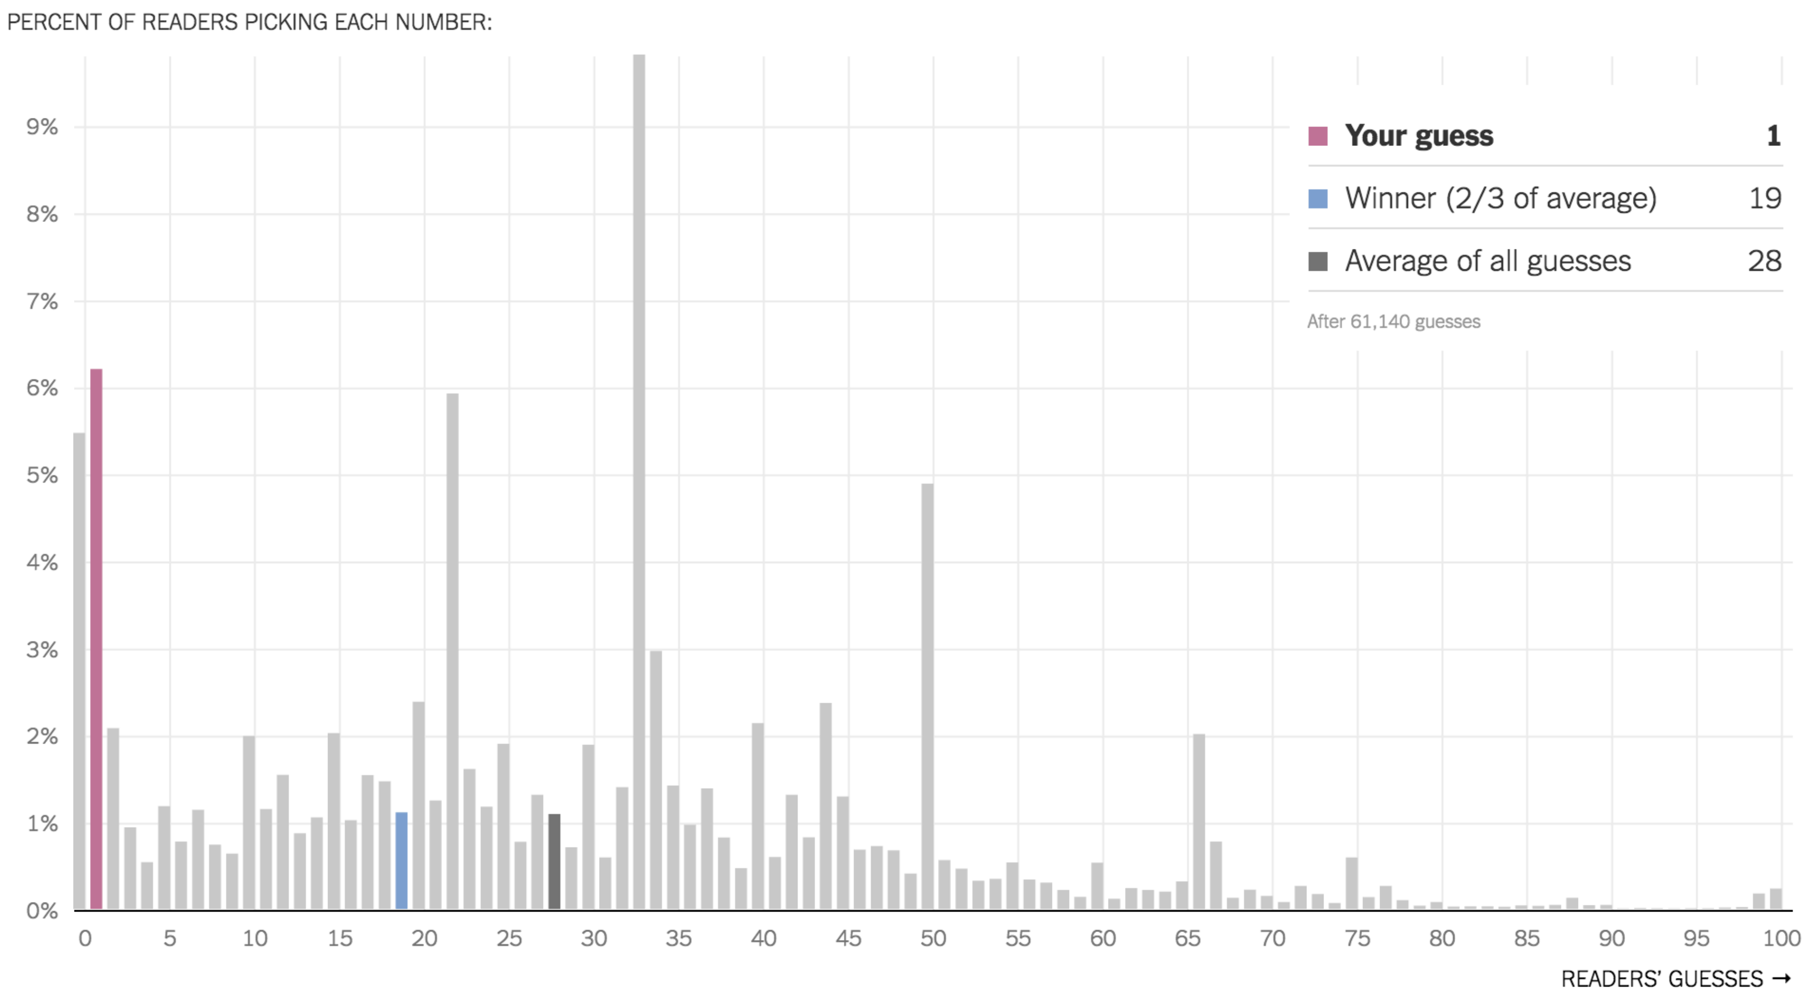
PERCENT OF READERS PICKING EACH NUMBER:
9%
8%
7%
6%
5%
4%
3%
2%
1%
0%
0
5
10
15
20
25
30
35
40
45
50
55
60
65
70
75
80
85
90
95
100
READERS' GUESSES →
Your guess 1
Winner (2/3 of average) 19
Average of all guesses 28
After 61,140 guesses

NY Times

NEXT CLASS: MECHANISM DESIGN PRIMER