Before Class…

Please sit in groups of 4 or more for lecture today!

We are practicing voting strategies.

Announcements

- Electronic assignment 12 due 4/30
- Programming assignment due 5/2

- Final exam 5/9 1-4pm (Rashid Auditorium)

- You're doing great!!!

AI: Representation and Problem Solving

Game Theory

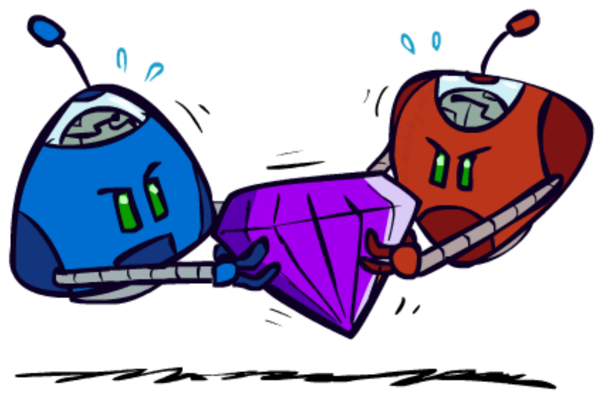

Instructors: Pat Virtue & Stephanie Rosenthal

Slide credits: Ariel Procaccia, Fei Fang

Mixed Strategy NE

P1

P2

GAME OF THRONES		NIGHT KING	
		FLEE	FIGHT
HUMANS	FLEE	90,10	20,80
	FIGHT	40,60	50,50

Other Properties of Strategies

Correlated Equilibrium

Pareto Optimal/Dominated

Pareto Optimal and Pareto Dominated

An outcome $u(\mathbf{s}) = \langle u_1(\mathbf{s}), \ldots, u_n(\mathbf{s}) \rangle$ is Pareto optimal if there is no other outcome that all players would prefer, i.e., each player gets higher utility

- At least one player would be disappointed in changing strategy

An outcome $u(\mathbf{s}) = \langle u_1(\mathbf{s}), \ldots, u_n(\mathbf{s}) \rangle$ is Pareto dominated by another outcome if all the players would prefer the other outcome

Joint vs Independent Strategies

A mixed strategy NE is one where each player chooses his/her action independently from the other players.

Joint vs Independent Strategies

A mixed strategy NE is one where each player chooses his/her action independently from the other players.

TRAFFIC		CAR 2: STOP	CAR 2: GO
CAR 1	STOP	0,0	0,1
CAR 1	GO	1,0	-100,-100

Joint vs Independent Strategies

A mixed strategy NE is one where each player chooses his/her action independently from the other players.

TRAFFIC		CAR 2: STOP	CAR 2: GO
CAR 1	STOP	0,0	0,1
CAR 1	GO	1,0	-100,-100

PSNE 1: (STOP,GO)

PSNE 2: (GO,STOP)

MSNE 3:

Joint vs Independent Strategies

A mixed strategy NE is one where each player chooses his/her action independently from the other players.

TRAFFIC		CAR 2: STOP	CAR 2: GO
CAR 1	STOP	0,0	0,1
CAR 1	GO	1,0	-100,-100

PSNE 1: (STOP,GO)

PSNE 2: (GO,STOP)

MSNE 3: Stop: 100/101
Go: 1/101

Joint vs Independent Strategies

A mixed strategy NE is one where each player chooses his/her action independently from the other players.

TRAFFIC		CAR 2: STOP	CAR 2: GO
CAR 1	STOP	0,0	0,1
CAR 1	GO	1,0	-100,-100

PSNE 1: (STOP,GO)

U(STOP,GO) = (0,1)

PSNE 2: (GO,STOP)

U(GO,STOP) = (1,0)

MSNE 3:

-0.0001= 0(.99)(.99)+0(.99)(.01) + 1(.01)(.99)-100(.01)(.01)

Joint vs Independent Strategies

A mixed strategy NE is one where each player chooses his/her action independently from the other players.

0.01% of the time, we risk death with such a strategy!

What if instead we have a mediator who chooses among joint strategies? Does this produce a higher expected utility and higher social welfare?

Correlated Equilibrium

Suppose a mediator computes the best joint strategy for p1 and p2, and shares a selected a_1 with p1 and a_2 with p2

TRAFFIC		CAR 2: STOP	CAR 2: GO
CAR 1	STOP	0,0	0,1
CAR 1	GO	1,0	-100,-100

Correlated Equilibrium

Suppose a mediator computes the best joint strategy for p1 and p2, and shares a selected a_1 with p1 and a_2 with p2

TRAFFIC		CAR 2: STOP	CAR 2: GO
CAR 1	STOP	0,0	0,1
CAR 1	GO	1,0	-100,-100

Mediator chooses:

50% (STOP,GO)

50% (GO,STOP)

Correlated Equilibrium

Suppose a mediator computes the best joint strategy for p1 and p2, and shares a selected a_1 with p1 and a_2 with p2

TRAFFIC		CAR 2: STOP	CAR 2: GO
CAR 1	STOP	0,0	0,1
CAR 1	GO	1,0	-100,-100

Mediator chooses:

50% (STOP,GO)

50% (GO,STOP)

If mediator tells C1 GO,
it knows C2 will STOP

Correlated Equilibrium

Suppose a mediator computes the best joint strategy for p1 and p2, and shares a selected a_1 with p1 and a_2 with p2

TRAFFIC		CAR 2: STOP	CAR 2: GO
CAR 1	STOP	0,0	0,1
CAR 1	GO	1,0	-100,-100

Mediator chooses:

50% (STOP,GO)

50% (GO,STOP)

Social welfare: 1

Each car goes ½ the time

The Game of Chicken

PSNE:

CHICKEN		PERSON 2	
		CHICKEN	DARE
PERSON 1	CHICKEN	6,6	2,7
	DARE	7,2	0,0

The Game of Chicken

MSNE: Chicken 2/3, Dare 1/3 for each player

Utility: $4/9*6 + 2/9*2 + 2/9*7 + 1/9*0 = 42/9 = 4.667$

CHICKEN		PERSON 2	
		CHICKEN	DARE
PERSON 1	CHICKEN	6,6	2,7
	DARE	7,2	0,0

The Game of Chicken

CE: Choose (C,C), (C,D), and (D,C) each with p=1/3

CHICKEN		PERSON 2	
		CHICKEN	DARE
PERSON 1	CHICKEN	6,6	2,7
	DARE	7,2	0,0

The Game of Chicken

CE: Choose (C,C), (C,D), and (D,C) each with p=1/3

CHICKEN		PERSON 2	
		CHICKEN	DARE
PERSON 1	CHICKEN	6,6	2,7
	DARE	7,2	0,0

If mediator tells P2 D,
he knows P1 plays C

If mediator tells P2 C,
½ the time P1 plays C
½ the time P1 plays D

The Game of Chicken

CE: Choose (C,C), (C,D), and (D,C) each with p=1/3

CHICKEN		PERSON 2	
		CHICKEN	DARE
PERSON 1	CHICKEN	6,6	2,7
	DARE	7,2	0,0

If mediator tells P2 D,
he get U=7

If mediator tells P2 C,
½ the time U=6
½ the time U=2

The Game of Chicken

CE: Choose (C,C), (C,D), and (D,C) each with p=1/3

Overall utility is

CHICKEN		PERSON 2	
		CHICKEN	DARE
PERSON 1	CHICKEN	6,6	2,7
	DARE	7,2	0,0

The Game of Chicken

CE: Choose (C,C), (C,D), and (D,C) each with p=1/3

Overall utility is 7/3+6/3+2/3 = 5 instead of 4.667

CHICKEN		PERSON 2: CHICKEN	PERSON 2: DARE
PERSON 1	CHICKEN	6,6	2,7
PERSON 1	DARE	7,2	0,0

Correlated Equilibrium

Suppose a mediator computes the best joint strategy for p1 and p2, and shares a selected a_1 with p1 and a_2 with p2

A correlated equilibrium is a distribution over action profiles $\vec{a}$ such that after a profile $\vec{a}$ is selected, playing a_i is a best response for player i conditioned on seeing a_i, given that everyone else will play according to $\vec{a}$.

Correlated Equilibrium

Suppose a mediator computes the best joint strategy for p1 and p2, and shares a selected a_1 with p1 and a_2 with p2

A correlated equilibrium is a distribution over action profiles $\vec{a}$ such that after a profile $\vec{a}$ is selected, playing a_i is a best response for player i conditioned on seeing a_i, given that everyone else will play according to $\vec{a}$.

Given that P1 has seen action a_1,

$$\sum_{a_1 \in A_1} \sum_{a_2 \in \vec{a}} p(a_1, a_2) u_i(a_1, a_2) \geq \sum_{a'_1 \in A_1} \sum_{a_2 \in \vec{a}} p(a'_1, a_2) u_i(a'_1, a_2)$$

And the same for P2 for a_2.

AI: Representation and Problem Solving

Social Choice

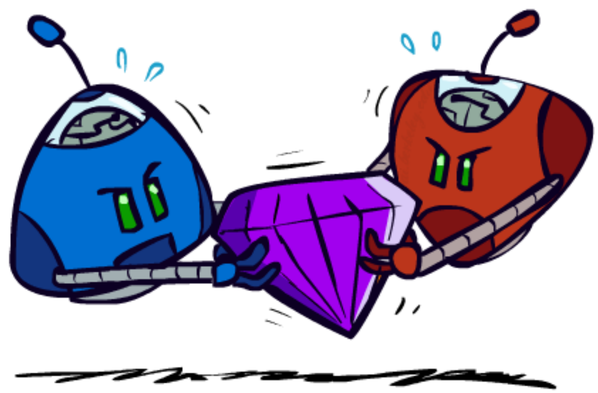

Instructors: Pat Virtue & Stephanie Rosenthal

Slide credits: Ariel Procaccia

Social Choice

Mathematical theory that deals with the aggregation of individual preferences

Origins in Ancient Greece
Formal foundations in 18th century – Condorcet and Borda
19th Century – Charles Dodgson
20th Century – Nobel prizes to Arrow and Sen

Voting Model

Set of N voters {1,2,…,N}
Set of A alternatives: $|A| = m$

Each voter has a ranking of alternatives

Preference profile: collection of all voter rankings

1	2	3
a	b	a
b	c	c
c	a	b

Voting Model

Set of N voters {1,2,…,N}
Set of A alternatives: |A| = m

1	2	3
a	b	a
b	c	c
c	a	b

Each voter has a ranking of alternatives

Preference profile: collection of all voter rankings

Voting rule: a function from a preference profile to an alternative (winner) of an election

Voting Rule: Plurality

Each voter gets one vote for their top-ranked preference.

Alternative with the most votes wins

1	2	3
a	b	a
b	c	c
c	a	b

Voting Rule: Plurality

Each voter gets one vote for their top-ranked preference.

Alternative with the most votes wins

a: 2 votes
b: 1 vote

1	2	3
a	b	a
b	c	c
c	a	b

Voting Rule: Borda Count

Each voter awards *m-k* points to their rank k.

Alternative with the most votes wins

Used in elections in Slovenia and Eurovision singing contest

1	2	3
a	b	a
b	c	c
c	a	b

Voting Rule: Borda Count

Each voter awards *m-k* points to their rank k.

Alternative with the most votes wins

a: 2+0+2 = 4
b: 1+2+0 = 3
c: 0+1+1 = 2

1	2	3
a	b	a
b	c	c
c	a	b

Voting Rule: Single Transferable Vote (STV)

Each voter gets 1 vote per round

In each round, alternative with the least number of plurality votes is eliminated

Alternative left standing is winner

Used in Ireland, Malta, Australia, NZ

1	2	3
a	b	a
b	c	c
c	a	b

Voting Rule: Single Transferable Vote (STV)

Each voter gets 1 vote per round

In each round, alternative with the least number of plurality votes is eliminated

Round 1: a and b survive
Round 2: a wins

1	2	3
a	b	a
b	c	c
c	a	b

1	2	3
a	b	a
b	a	b

1	2	3
a	a	a

Voting Rule: Single Transferable Vote (STV)

Each voter gets 1 vote per round

In each round, alternative with the least number of plurality votes is eliminated

Nothing to eliminate?
Tie breaking strategies include borda count, having the most last place votes, having the most votes in the first round, etc

1	2	3
a	b	**c**
b	c	**b**
c	a	**a**

On your own, rank your favorite candies

Crunch
M&Ms
Reese's Cups
Snickers
Skittles
Milky Way
Almond Joy
Kit Kat

Compute the Plurality, Borda, STV winners

Crunch
M&Ms
Reese's Cups
Snickers
Skittles
Milky Way
Almond Joy
Kit Kat

Plurality Winners

Borda Count Winners

STV Winners

What patterns do you notice?

Voting Rule: Pairwise Election

Alternative *x* beats *y* in pairwise election if majority of voters prefer *x* to *y*

1	2	3
a	b	a
b	c	c
c	a	b

Voting Rule: Pairwise Election

Alternative *x* beats *y* in pairwise election if majority of voters prefer *x* to *y*

1	2	3
a	b	a
b	c	c
c	a	b

2 voters prefer a over b
2 voters prefer b over c
2 voters prefer a over c

Voting Rule: Plurality with Runoff

First Round: Top 2 plurality winners advance to second round

Second Round: Pairwise election between two winners

1	2	3
a	b	a
b	c	c
c	a	b

Voting Rule: Plurality with Runoff

First Round: Top 2 plurality winners advance to second round
Second Round: Pairwise election between two winners

1	2	3
a	b	a
b	c	c
c	a	b

Round 1: a and b move on
Round 2: 2 votes for a over b
Alternative a wins

Voting Rule: Condorcet Winner

Alternative *x* beats *y* in pairwise election if majority of voters prefer *x* to *y*

Condorcet winner *x* beats every other alternative *y* in pairwise election

2 voters prefer a over b and a over c
Alternative a is the Condorcet winner

1	2	3
a	b	a
b	c	c
c	a	b

Voting Rule: Condorcet Winner

Alternative *x* beats *y* in pairwise election if majority of voters prefer *x* to *y*

1	2	3
a	b	a
b	c	c
c	a	b

Condorcet winner *x* beats every other alternative *y* in pairwise election

Condorcet paradox is a cycle in majority preferences

1	2	3
a	b	**c**
b	c	**a**
c	a	b

Poll 2

Condorcet consistent – voting rule selects a Condorcet winner if one exists

Which rule is always Condorcet consistent?

a) Plurality
b) Borda Count
c) Both
d) Neither

Poll 2

Condorcet consistent – voting rule selects a Condorcet winner if one exists

Which rule is always Condorcet consistent?

a) Plurality
b) Borda Count
c) Both
d) Neither

3 voters	2 voters
a	b
b	c
c	a

Plurality: a
Borda: b
Condorcet: a

3 voters	2 voters	2 voters
a	b	c
b	c	b
c	a	a

Plurality: a
Borda: b
Condorcet: b

Compute Plurality w/Runoff, Condorcet Winners

Crunch
M&Ms
Reese's Cups
Snickers
Skittles
Milky Way
Almond Joy
Kit Kat

Plurality with Runoff Winners

Condorcet Winners

Fun Example

Plurality:

Borda:

STV:

Condorcet:

Plurality with runoff:

33 voters	16 voters	3 voters	8 voters	18 voters	22 voters
a	b	c	c	d	e
b	d	d	e	e	c
c	c	b	b	c	b
d	e	a	d	b	d
e	a	e	a	a	a

Voting for Truth

Condorcet [1785]: the purpose of voting is not merely to balance subjective opinions but also a quest to find truth

Enlightened voters try to judge which alternatives best serve society

This is realistic in trials by jury, pooling expert opinions, and human computation

Crowdsourcing Molecule Designs

Developed by Adrien Treuille (CMU) and Stanford in 2010

Participants solve puzzles to find molecule designs

They vote on which 8 designs get synthesized, the votes aim to compare designs by true quality

Voting in Crowdsourcing

Amazon's Mechanical Turk (started in 2005)

Organizations can post HITs (Human Intelligence Tasks) for small amounts of money

e.g. identify content in image/video, write product description, or answer questions/surveys, etc

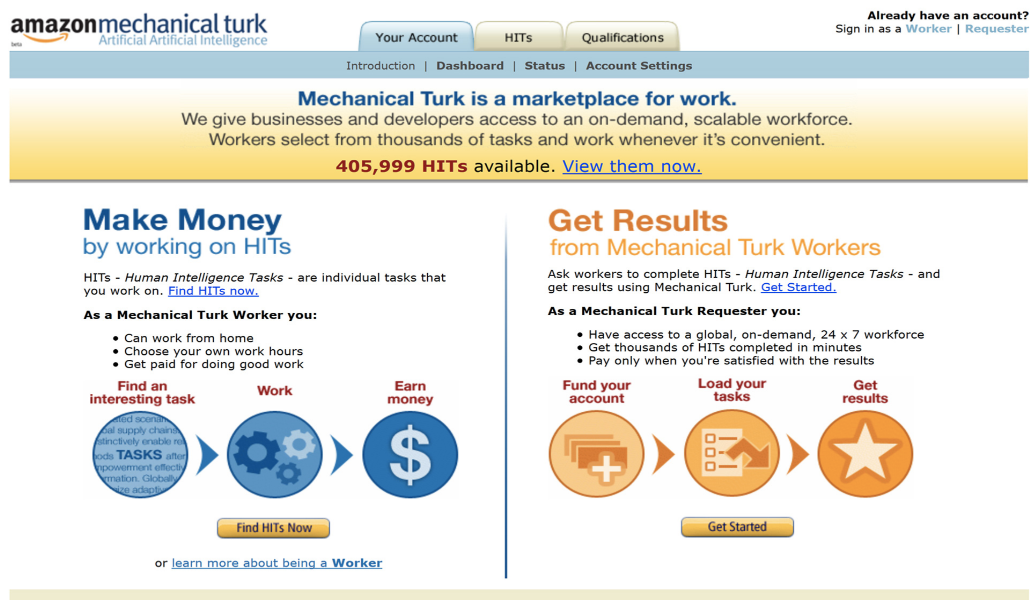

Common HIT Frameworks

1) An organization poses a question with a single right (but unknown) answer

Voting : They actually post the same question N times (often N=5). If a majority of the responses are the same, they can ensure that it is good/correct. They post more times if responses do not yield a majority.

Common HIT Frameworks

1) An organization poses a question with a single right (but unknown) answer

2) An organization poses a question with many answers and collect N responses

Voting: Once many responses are collected, they pose a new HIT asking new participants to a) pairwise rank responses, b) rank all responses for the best answer

How to rank many responses

[Mao, Procaccia, Chen 2013]

Compared ranking strategies Plurality, Borda, Condorcet

Found that Borda finds the winners most consistently even with noisy human responses, Plurality performs the worst

What are the consequences of this finding?

reCaptcha

Show participants 1 known and 1 unknown image

If they get the known one correct, assume unknown one is also reasonable

If 6 people request new, assume unreadable

Majority response wins (N=5)

Summary

Vocabulary

- Voting rules – Plurality, Borda Count, STV, Pairwise election, Condorcet winner, Plurality with runoff
- Crowdsourcing and Human Computation