CSC200: Lecture 9

- Today
  - Continuing game theory: equilibria, multiple equilibria, mixed strategies (Ch.6.4-6.8)
  - **NOTE: We will only cover up to slide 20 today but I am adding a few extra slides anticipating Lecture 10.**
- Next few lectures
  - game theory: Ch.6, Ch.8, Ch.9. Next set of lectures by Omer Lev and Tyrone Strangway
- Announcements
  - Asst.1: Due this Friday, Oct.16 ; due beginning of tutorial (not at end of tutorial so that you can ask any questions during tutorial.
  - Game theory question has been postponed to assignment 2. Thank you to the student(s) who pointed out typo(s).
  - Short questions (i.e. mainly clarifications) about assignment 1 can be sent to myself and TAs. Do not expect detailed answers.

More annoucements

- Quiz scope: chapter 4 (closures in a social-affiliation network) takes place Friday, October 23. The quiz will take place at the end of the tutorial (rather than the beginning).
- The quiz will be 20 minutes (although the paper says 15 minutes) and I believe you will find it easier than quiz 1.
- As per my standard policy, you are allowed one two-sided sheet of handwritten notes.
- Don't forget to vote. Voting theory next term. Omer Lev has written papers about voting theory and he might say some things about the subject given what I hope is your interest.
- Finally, Go Jays Go. I will end 10 minutes early today so that you can watch the game (if you are interested) and so that I can catch a flight.

Recap

- Last time we introduced games in matrix form
  - basic ingredients: players, strategies (actions), payoffs
  - discussed strategies, strategy profiles, partial strategy profiles
  - dominant strategies
  - *what are some "best responses" in these games?*

Bob

Alice	Uptown	Downtown
Uptown	600 / 400	400 / 300
Downtown	0 / 200	200 / 500

Alice payoff/Bob payoff

Firm 1 small \ Firm 2 large	A	B	C
A	4, 4	0, 2	0, 2
B	0, 0	1, 1	0, 2
C	0, 0	0, 2	1, 1

Dominant Strategy: Recap

- In a two player game: a *dominant strategy* for player 1 is any strategy that is a best response to *every* strategy of player 2.
  - One’s reasoning about how to act is straightforward
  - Our prediction of one’s action may also be “straightforward” (if dominant strategy is unique)
  - We’ve seen examples already (e.g., Alice moving uptown)
- Generalization to *n* players is easy: a *dominant strategy* for player *k* is any strategy that is a best response to *every partial strategy profile* of the other players.

Another Game with An Unappealing Outcome: Recap

- Alice and her friends like to hang out with Bob and his friends (both have a large group of friends)
- There are two clubs
  - Club 1: great place if a big crowd, lots of friends; otherwise horrible
  - Club 2: exciting, but tiny: If too crowded (e.g., if both groups show up), not so great
  - payoffs: see payoff matrix
- Alice and Bob (their groups) *both* have dominant strategies
  - no matter what the other does, going to Club 2 is better
- Equilibrium *(Alice: Club2; Bob: Club2)* has payoff (4,4) that is *worse for both of them* than the payoff (6,6) if they both go to Club 1

Alice \ Bob	Club 1	Club 2
Club 1	6 / 6	2 / 7
Club 2	7 / 2	4 / 4

Prisoner's Dilemma

	Quiet	Confess
Quiet	-1 / -1	-10 / 0
Confess	0 / -10	-4 / -4

- The last game is similar to what is known as the Prisoner's Dilemma
  - two suspects interrogated in separate rooms
  - not enough evidence for conviction unless one confesses
  - offer to both: confess (alone), go free; otherwise lesser charge
- Dominant strategy (confession) leads to a "socially undesirable" (for prisoners) outcome (both better off if they could "coordinate" on keeping quiet)
- Other examples:
  - performance enhancing drugs by athletes
  - the "arms race
- Ways around "dilemma" (all require different model or matrix)
  - reputation among peers? retaliation? future interactions?
- How sensitive to precise payoffs?
  - what if quiet/quiet was 0/0?

Best Response (Two Players)

- Let's formalize the analysis
- First for two players:
  - suppose player 2 plays strategy *t*
  - each strategy *s* of player 1 has a payoff $P_1(s,t)$
  - the *s* in S_1 with highest payoff is a *best response (BR) to t*
  - i.e., *s* is a BR for player 1 iff $P_1(s,t) \geq P_1(s',t)$ *for all s'* in S_1
  - similar definition applies to player 2

- Sometimes a player can have multiple best responses to the strategy of her opponent. If the BR is unique (i.e., strictly greater than the payoff of any other action), we call it a *strict best response*.

Best Response (Example)

		Firm 2 A	Firm 2 B	Firm 2 C
Firm 1	A	4, 4	0, 2	0, 2
	B	0, 0	1, 1	0, 2
	C	0, 0	0, 2	1, 1

- Three strategy game
  - Two firms can approach three different clients, A, B, C
  - Client A is large (total business 8); B,C small (total 2)
  - Firm 1 is small; can only get *any* business if it approaches *same* client as Firm 2; splits the business with Firm 2 if both approach
  - Firm 2 can get B or C on its own, but can only get A (large client) if it does so with Firm 1
- Best responses:

	If F2 does: A	B	C
BR For F1	A	B	C

	If F1 does: A	B	C
BR For F2	A	C	B

A Three-Player Game

- Player 1 chooses the row (top, bottom), Player 2 the column (left, right), Player 3 the matrix (M1, M2)
  - ***Example:*** three MLAs: (closed) voting on salary increase. Vote passes if at least two vote for it. All want the raise, but would prefer to vote against it (and have the other two vote for it!)
    - P_k if raise passes, k votes against: 2
    - P_k if raise passes, k votes for: 1
    - P_k if raise fails, k votes against: 0
    - P_k if raise fails, k votes for: -1

e.g., P_2(for,agnst,for) = 2

For (M1)

	For (L)	Against (R)
For (U)	1 / 1 / 1	1 / 2 / 1
Against (D)	2 / 1 / 1	0 / 0 / -1

Against (M2)

	For (L)	Against (R)
For (U)	1 / 1 / 2	-1 / 0 / 0
Against (D)	0 / -1 / 0	0 / 0 / 0

Best Response (More than Two Players)

- A *strategy profile*: a collection of strategies, 1 per player
  - profile determines everyone's payoff
    - e.g., in semi-secret ballot *(for, for, against)*
  - *partial profile for k*: everyone's strategy *except k*'s
    - e.g., for player 2 *(for, ---, against)*
- For *n* players (players 1, 2, 3, … n)
  - suppose all players except *k* play partial profile ***t***
  - each strategy *s* of player *k* has a payoff $P_k(s,\mathbf{t})$
  - the *s* in S_k with highest payoff is a *best response (BR) for k to* ***t***
  - in other words, *s* is a BR iff $P_k(s,\mathbf{t}) \geq P_k(s',\mathbf{t})$ *for all s'* in S_k

Nash Equilibrium (Two Players)

- How do we predict behavior in a game?
  - When both players have a dominant strategy (DS), we expect these to be played by simple rationality. (What if there are ties?)
  - When one player has a DS, we expect it to be played; then we expect other player to "expect it" too, and play a best response
  - What if neither has a DS? How they behave depends on what they believe the other will do…
- *Nash equilibrium* ("deterministic" or "pure strategy"):
  - a pair of strategies (or profile) *s* for player 1 and *t* for player 2 such that *s* is a best response to *t*, and *t* is a best response to *s*
- Why is an NE a reasonable prediction?
  - if both playing their part of an NE, no reason to change strategy (in absence of some form of coordination)
  - if playing a profile that is not an NE, at least one wants to switch

Nash Equilibrium

- Most fundamental concept in game theory
  - many refinements, extensions, but this is the core
  - not a rationality concept per se, but ties rationality to “stable” predictions about behavior of your “opponents”
- Developed by John Nash in 1950
  - equilibria concepts existed for centuires
  - early analysis of von Neumann for zero-sum games
  - proved existence of “mixed strategy” equilibria for any “game” for which he was awarded a 1994 Nobel Prize (jointly with Harsanyi and Selten)

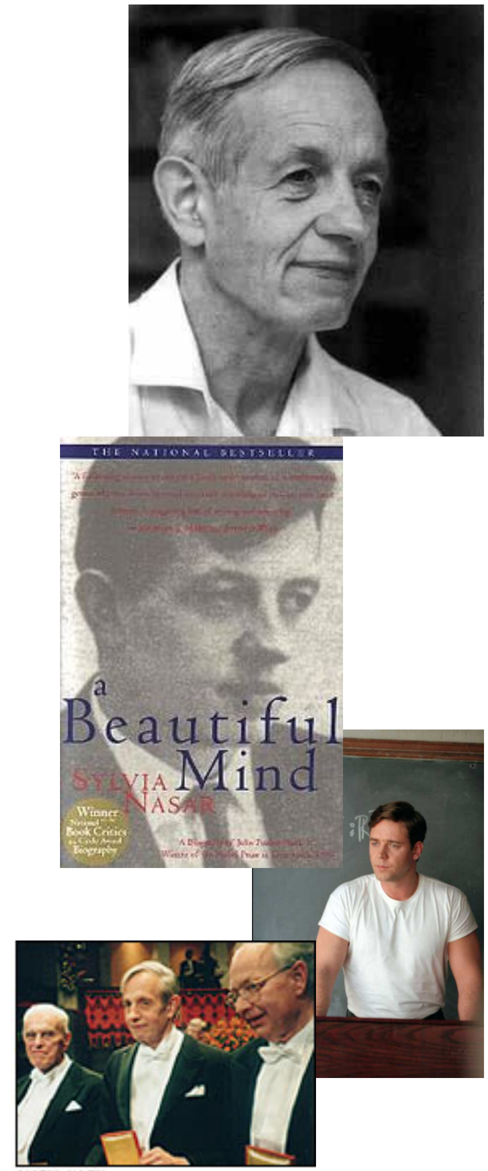

Equilibrium of the Three Action Game

- If F1 does B, F2 does C
- If F2 does C, F1 does C
- If F1 does C, F2 does B
- If F2 does B, F1 does B...
- *Cycle of BRs:* can't be a (deterministic) equilibrium with B or C!

- If F1 does A, F2 does A
- If F2 does A, F1 does A
- *(A,A) is a Nash equilibrium*
  - each company's strategy is a best response to that of its competitor

		Firm 2		
		A	B	C
	A	$4, 4$	$0, 2$	$0, 2$
Firm 1	B	$0, 0$	$1, 1$	$0, 2$
	C	$0, 0$	$0, 2$	$1, 1$

What are the Equilibria of these Games?

Alice \ Bob	Uptown	Downtown
Uptown	600 / 400	400 / 300
Downtown	0 / 200	200 / 500

- Bob-Up: Alice BR is Up
- Bob-Down: Alice BR is Up
  - Up is dominant for Alice
- Alice-Up: Bob BR is Up
- Alice-Down: Bob BR is Down
- NE: *($uptown_{Alice}$, $uptown_{Bob}$)*

Alice \ Bob	Uptown	Downtown
Uptown	600 / 300	400 / 400
Downtown	100 / 100	300 / 600

- Bob-Up: Alice BR is Up
- Bob-Down: Alice BR is Up
  - *Up is dominant for Alice*
- Alice-Up: Bob BR is Down
- Alice-Down: Bob BR is Down
  - *Down is dominant for Bob*
- NE: *($uptown_{Alice}$, $downtown_{Bob}$)*

Equilibria (Three-Player Example)

For (M1)

	For (L)	Against (R)
For (U)	1 / 1 / 1	1 / 2 / 1
Against (D)	2 / 1 / 1	0 / 0 / -1

Against (M2)

	For (L)	Against (R)
For (U)	1 / 1 / 2	-1 / 0 / 0
Against (D)	0 / -1 / 0	0 / 0 / 0

- Is $(for_1, against_2, for_3)$ a Nash equilibrium?
  - 1's best response to $(\text{--}, against_2, for_3)$ is for_1
  - 2's best response to $(for_1, \text{--}, for_3)$ is $against_2$
  - 3's best response to $(for_1, against_2, \text{--})$ is for_3
- So yes, outcome is stable (nobody wants to deviate)
- Is it the only equilibrium?
  - By symmetry: $(against_1, for_2, for_3)$ and $(for_1, for_2, against_3)$ must be NE too
  - What about $(against_1, against_2, against_3)$?

Multiple Equilibria (1)

- Previous example: multiple equilibria for a single game
  - problem if we want to use Nash eq. to predict behavior!
- Far from unusual, multiple equilibria arise in many games
- *Pure coordination game:*
  - Alice, Bob only want to live near each other
  - only issue is how to coordinate (no conflicting interests)
  - two (deterministic) NE: *(up,up)* and *(down,down)*

		Bob: Uptown	Bob: Downtown
Alice	Uptown	400 / 400	0 / 0
	Downtown	0 / 0	400 / 400

Multiple Equilibria (2)

- *Unbalanced coordination game:*
  - Alice, Bob want to live near each other, slightly prefer uptown
  - *still (!)* two (deterministic) NE: *(up,up)* and *(down,down)*
- In this case, choice of NE seems more obvious
  - select NE where *both* prefer the outcome over the other NE

Alice \ **Bob**	Uptown	Downtown
Uptown	400 / 400	0 / 0
Downtown	0 / 0	300 / 300

Multiple Equilibria (3)

- *Battle of the Sexes:*
  - Alice slightly prefers uptown; Bob slightly prefers downtown
    - *named for husband, wife trying to meet at ballet or boxing match (Luce and Raiffa 1957)*
  - *still (!)* two (deterministic) NE: *(up,up)* and *(down,down)*
- Choice of NE seems less obvious

Alice \ Bob	Uptown	Downtown
Uptown	400 / 200	0 / 0
Downtown	0 / 0	200 / 400

Multiple Equilibria (4)

- *Stag Hunt:*
  - If two hunter's team up, they will bag a stag; but if one tries alone, sure to fail
  - Safer option: hunt alone and try to bag a hare (no interaction, no cooperation needed)
  - two (deterministic) NE: *(stag,stag)* and *(hare,hare)*
- Choice of NE seems less obvious
  - *(s,s)* gives better payoff, but poor outcome if your partner opts out
  - *(h,h)* gives lesser payoff, but guaranteed!

		Hunter 2: *Hunt Stag*	Hunter 2: *Hunt Hare*
Hunter 1	*Hunt Stag*	4, 4	0, 3
	Hunt Hare	3, 0	3, 3

Equilibrium Selection

- A very subtle topic, and important in practical settings
  - e.g., studied extensively in political science, behavioral economics, social psychology, etc.
- Schelling (yes, the same one) proposes *focal points* as one "heuristic" means of resolving the issue
  - factors external to the game may influence choices
    - *past history (met at Starbuck's last week)*
    - *social conventions (generally agree to pass on the right)*
  - factors internal to game structure may matter
    - *as in the unbalanced coordination game*
    - *or risk attitude (Stag Hunt)*
- *See also Harsanyi and Selten (Nash's co-laureates)*

But it gets worse: Matching Pennies

		Player 2	
		H	T
Player 1	H	$-1, +1$	$+1, -1$
	T	$+1, -1$	$-1, +1$

		Firm 2		
		A	B	C
Firm 1	A	$4, 4$	$0, 2$	$0, 2$
	B	$0, 0$	$1, 1$	$0, 2$
	C	$0, 0$	$0, 2$	$1, 1$

- What's the problem?

No Deterministic Equilibria

- Matching Pennies has no (deterministic) NE
  - 1 heads → 2 heads → 1 tails → 2 tails →1 heads …
  - ditto for "subgame" for Firm 1 and 2 (only clients are B and C)
- What are we to do?
  - allow players to *randomize* their choice of actions (or "flip coins")
- Suppose player 1 chooses heads/tails with probability 0.5, and player 2 chooses heads/tails with probability 0.5
  - does either have a reason to want to change their mind?

		Player 2: H	Player 2: T
Player 1	H	$-1, +1$	$+1, -1$
	T	$+1, -1$	$-1, +1$

Mixed Strategies

- Let S_1 be the set of *pure* strategies of player 1
- A *mixed strategy* σ_1 for player 1 is a probability distribution over S_1
  - assigns a probability to each s in S_1 (probabilities must sum to one, must be non-negative)
  - player 1 chooses each *s* with the corresponding probability
  - we say the player is *mixing* between his pure strategies
- Examples (note: pure strategies are a special case):
  - [0.5 heads, 0.5 tails]; [1.0 heads, 0 tails]; [0.01 heads, 0.99 tails]; etc.
  - Firm1: [0.5 A, 0.25 B, 0.25 C]; [1.0 A, 0 B, 0 C]; etc.
  - Bob: [1.0 uptown, 0 downtown]; [0.2 uptown, 0.8 downtown]; etc.
- In two-action games, often just state probability *p* of first action
  - probability of second action is $1-p$
  - so set of possible strategies σ is just a set of real numbers between 0 and 1.

Outcomes of Mixed Strategies

- Since players randomize, *outcome and payoffs are also random*
- Consider matching pennies
  - P1 plays heads with probability *p* (tails, *1-p*) *(e.g., p=0.3)*
  - P2 plays heads with probability *q* (tails *1-q*) *(e.g., q=0.8)*
  - Assume they randomize independently (can't coordinate!)
- Probability of various outcomes and payoffs

Outcome	Prob	e.g.	Payoff(1)	Payoff(2)
(H,H)	pq	0.24	-1	+1
(H,T)	p(1-q)	0.06	+1	-1
(T,H)	(1-p)q	0.56	+1	-1
(T,T)	(1-p)(1-q)	0.14	-1	+1

		Player 2	
		H	T
Player 1	H	$-1, +1$	$+1, -1$
	T	$+1, -1$	$-1, +1$

Payoff of Mixed Profile

- Since payoff to each player is random, the value to the player of a profile is her *expected payoff*
  - simply take the (probabilistic weighted) average of her payoffs

Outcome	Prob	e.g.	Payoff(1)	Payoff(2)
(H,H)	pq	0.24	-1	+1
(H,T)	p(1-q)	0.06	+1	-1
(T,H)	(1-p)q	0.56	+1	-1
(T,T)	(1-p)(1-q)	0.14	-1	+1

- Expected payoff (or expected value, EV) for player 1:
  - 0.24(-1) + 0.06(1) + 0.56(1) + 0.14(-1) = 0.24
- Expected payoff for player 2:
  - 0.24(1) + 0.06(-1) + 0.56(-1) + 0.14(1) = −0.24
- *Hardly a coincidence that one is minus the other*