Multiagent Systems
Chapter 17: Game theoretic Foundations of Multi Agent Systems

http://mitpress.mit.edu/multiagentsystems

Edith Elkind
Evangelos Markakis

Outline

- Introduction
- Normal-form games
- Extensive-form games
- Bayesian games

What Does Game Theory Study?

Interactions of rational decision-makers (agents, players)

- Decision-makers: humans, robots, computer programs, firms in the market, political parties
- Rational: each agent has preferences over outcomes and chooses an action that is most likely to lead to the best feasible outcome
- Interactions: 2 or more agents act simultaneously or consequently

Why Study Game Theory?

- To understand the behavior of others in strategic situations
- To know how to alter one's own behavior in such situations to gain advantage
- Wikipedia: game theory attempts to mathematically capture behavior in strategic situations, in which an individual success in making choices depends on the choices of others

A Bit of History

- Early ideas:
  - Models on competition among firms: Cournot (≈1838), Bertrand (≈1883)
  - 0-sum games: end of 19th century (Zermelo) and early 20th century (Borel)
- Foundations of the field (1944):
  Theory of Games and Economic Behavior by
  John von Neumann and Oskar Morgenstern
- Key concept: Nash equilibrium
  (John Nash, 1951)
- Main applications:
  - microeconomics
  - political science
  - evolutionary biology

Normal-Form Games

Normal-Form Games

- Complete-information games
  - players know each other's preferences
- Simultaneous moves
  - All players choose their action at the same time (or at the time they make their own choice, they do not know or cannot observe the other players' choices)

Normal-Form Games

Formally:

- A normal-form game is given by
  - a set of players N
  - for each player i, a set of available actions A_i
  - for each player i, a utility function $u_i: A_1 \times \ldots \times A_n \rightarrow R$ (real numbers)
- Action profile: Any vector $(a_1, \ldots, a_n)$, with $a_i \in A_i$
  - each action profile corresponds to an outcome
  - u_i describes how much player i enjoys each outcome

Example: Prisoner's Dilemma

- Two agents committed a crime.
- The court does not have enough evidence to convict them of the crime, but can convict them of a minor offence (1 year in prison each)
- If one suspect confesses (acts as an informer), he walks free, and the other suspect gets 4 years
- If both confess, each gets 3 years
- Agents have no way of communicating or making binding agreements

Prisoner's Dilemma: the Model

- Set of players N = {1, 2}
- A_1 = A_2 = {confess (C), stay quiet (Q)}
- u_1(C, C) = -3 (both get 3 years)
- u_1(C, Q) = 0 (player 1 walks free)
- u_1(Q, C) = -4 (player 1 gets 4 years)
- u_1(Q, Q) = -1 (both get 1 year)
- $u_2(x, y) = u_1(y, x)$

Prisoner's Dilemma: Matrix Representation

P1 \ P2	quiet	confess
quiet	(-1,-1)	(-4, 0)
confess	(0, -4)	(-3, -3)

- Interpretation: the pair (x, y) at the intersection of row i and column j means that the row player gets x
and the column player gets y

Prisoner's Dilemma: the Rational Outcome

- P1's reasoning:
  - if P2 stays quiet, I should confess
  - if P2 confesses, I should confess, too

P1 \ P2	Q	C
Q	(-1,-1)	(-4, 0)
C	(0, -4)	(-3, -3)

- P2 reasons in the same way
- Result: both confess and get 3 years in prison.
  - note, however, if they chose to cooperate and stay quiet, they could get away with 1 year each.

Dominant Strategy: Definition

- Dominant strategy: a strategy that is best for a player no matter what the others choose
- <u>Definition</u>: a strategy a of player i is said to be a dominant strategy for i, if

  $$u_i\,(a_1, \ldots, a_{i-1}, a, a_{i+1}, \ldots, a_n) \geq$$
  $$u_i\,(a_1, \ldots, a_{i-1}, a', a_{i+1}, \ldots, a_n)$$

  for any $a' \in A_i$ and any strategies $a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n$ of other players.
- In Prisoner's Dilemma, Confess is a dominant strategy for each of the players

Dominant Strategy: Discussion

- Can a player have more than one dominant strategy?
  - It can happen if some actions result always in the same utility
- <u>Definition</u>: a strategy a of player i is said to be a dominant strategy of player i if

strictly

$u_i(a_1, \ldots, a_{i-1}, a, a_{i+1}, .., a_n)$ ~~$\geq$~~ $>$
$u_i(a_1, \ldots, a_{i-1}, a', a_{i+1}, \ldots, a_n)$

for any $a' \in A_i$ and any strategies $a_1, \ldots, a_{i-1}, a_{i+1}, .., a_n$ of other players.

- <u>Fact</u>: each player has at most one strictly dominant strategy

The Joint Project Game

- Two students are assigned a project
- If at least one of them works hard, the project succeeds
- Each student
  - wants the project to succeed (+5)
  - prefers not to make an effort (-2)
  - hates to be exploited, i.e., work hard when the other slacks (-8)

P1 \ P2	Work	Slack
Work	(3, 3)	(-5, 5)
Slack	(5, -5)	(0, 0)

Joint Project vs. Prisoner's Dilemma

	W	S
W	(3, 3)	(-5, 5)
S	(5, -5)	(0, 0)

	Q	C
Q	(-1, -1)	(-4, 0)
C	(0, -4)	(-3, -3)

- In JP, row player prefers (S, W) to (W, W) to (S, S) to (W, S)
- In PD, row player prefers (C, Q) to (Q, Q) to (C, C) to (Q, C)
  - column player has similar preferences
- These two games are equivalent!
- Game theory prediction: both students will slack

Battle of Sexes

P1 \ P2	Theatre	Football
Theatre	(2, 1)	(0, 0)
Football	(0, 0)	(1, 2)

- Charlie and Marcie want to go out, either to theatre or to a football game
- She prefers theatre, he prefers football
- But they will be miserable if they go to different places

Battle of Sexes

	T	F
T	(2, 1)	(0, 0)
F	(0, 0)	(1, 2)

- No player has a dominant strategy:
  - T is not a dominant strategy for Marcie:
    if Charlie chooses F, Marcie prefers F
  - F is not a dominant strategy for Marcie:
    if Charlie chooses T, Marcie prefers T
- However, (T, T) is a stable pair of strategies:
  - neither player wants to change his action given the other player's action
- (F, F) is stable, too

Notation

- Given a vector $\underline{\mathbf{a}} = (a_1, \ldots, a_n)$, let $(\underline{\mathbf{a}}_{-i}, a')$ be $\underline{\mathbf{a}}$, but with a_i replaced by a':

$$(\underline{\mathbf{a}}_{-i}, a') = (a_1, \ldots, a_{i-1}, a', a_{i+1}, \ldots, a_n)$$

- If $\underline{\mathbf{a}} = (3, 5, 7, 8)$, then $(\underline{\mathbf{a}}_{-3}, 4) = (3, 5, 4, 8)$

Nash Equilibrium (Nash'51)

- Definition: a strategy profile $\underline{\mathbf{a}} = (a_1, \ldots.., a_n)$ is a Nash equilibrium (NE) if no player can benefit by changing unilaterally his action: for each $i = 1, \ldots, n$ it holds that

  $$u_i(\underline{\mathbf{a}}) \geq u_i(\underline{\mathbf{a}}_{-i}, a') \text{ for all } a' \text{ in } A_i$$

- 2 player case: (a, b) is a NE if
  1. $u_1(a, b) \geq u_1(a', b)$ for every $a' \in A_1$
  2. $u_2(a, b) \geq u_2(a, b')$ for every $b' \in A_2$

Nash Equilibrium Pictorially

(,)	(,)	(x_1,)	(,)	(,)
(,)	(,)	(x_2,)	(,)	(,)
(,)	(,)	(x_3,)	(,)	(,)
(,y_1)	(,y_2)	(X, Y)	(,y_4)	(,y_5)
(,)	(,)	(x_5,)	(,)	(,)

X must be at least as big as any x_i in Y-column

Y must be at least as big as any y_j in X-row

Nash Equilibria in Battle of Sexes

P2

P1	Theatre	Football
Theatre	(2, 1)	(0, 0)
Football	(0, 0)	(1, 2)

Both (T, T) and (F, F) are Nash equilibria

Nash Equilibrium and Dominant Strategies

- Prisoner's dilemma:
  (C, C) is a Nash equilibrium

P1 \ P2	Q	C
Q	(-1,-1)	(-4, 0)
C	(0, -4)	(-3, -3)

<u>Theorem</u>: In any 2-player normal-form game, if

- a is a dominant strategy for player 1, and
- b is a dominant strategy for player 2,

then (a, b) is a Nash equilibrium

Best Response Functions

Towards an alternative way of defining equilibria:

- Given a vector $\underline{\mathbf{a}}_{-i}$ of other players' actions, player i may have one or more actions that maximize his utility
- Best response function:
  $B_i(\underline{\mathbf{a}}_{-i}) =$
  $\{a \text{ in } A_i \mid u_i(\underline{\mathbf{a}}_{-i}, a) \geq u_i(\underline{\mathbf{a}}_{-i}, a') \text{ for all } a' \text{ in } A_i\}$
- $B_i(\underline{\mathbf{a}}_{-i})$ is set-valued
- if $|B_i(\underline{\mathbf{a}}_{-i})| = 1$ for all i and all $\underline{\mathbf{a}}_{-i}$, we denote the single element of $B_i(\underline{\mathbf{a}}_{-i})$ by $b_i(\underline{\mathbf{a}}_{-i})$

Example

	L	C	R
T	(2*, 5*)	(3 , 3)	(6*, 3)
M	(2*, 7*)	(4 , 5)	(2 , 7*)
B	(1 , 4*)	(5*, 4*)	(2 , 1)

- $B_1(L) = \{T, M\}$
- $B_1(C) = \{B\}$
- $B_1(R) = \{T\}$
- $B_2(T) = \{L\}$
- $B_2(M) = \{L, R\}$
- $B_2(B) = \{L, C\}$

Best Responses and Nash Equilibria

- Recall:
- $\underline{\mathbf{a}} = (a_1, \ldots\ldots, a_n)$ is a Nash equilibrium if

  $u_i(\underline{\mathbf{a}}) \geq u_i(\underline{\mathbf{a}}_{-i}, a')$

  for all i and all a' in A_i
- In the language of best response functions:

  $\underline{\mathbf{a}} = (a_1, \ldots\ldots, a_n)$ is a Nash equilibrium if

  a_i is in $B_i(\underline{\mathbf{a}}_{-i})$ for all i

Example Revisited

	L	C	R
T	(2*, 5*)	(3 , 3)	(6*, 3)
M	(2*, 7*)	(4 , 5)	(2 , 7*)
B	(1 , 4*)	(5*, 4*)	(2 , 1)

- B_1(L) = {T, M}, B_1(C) = {B}, B_1(R) = {T}
- B_2(T) = {L}, B_2(M) = {L, R}, B_2(B) = {L, C}
- {T, L}, {M, L} and {B, C} are Nash equilibria

Infinite Action Spaces

- What if a player does not have a finite number of strategies?
- There are games where each player has to choose among infinitely many actions:
  - how much time to spend on a task?
  - how much to bid in an auction?
  - where to locate a new factory?
  - how much money to invest?
- The concept of best response functions turns out to be very useful here....

Example: Preparing for an Exam

- Two students are preparing together for a joint exam
- each player's effort level is a number in [0, 1]
- if player 1 invests x units of effort, and
  player 2 invests y units of effort,
  player 1's utility is $x(c + y - x)$,
  player 2's utility is $y(c + x - y)$,
  where c is a given constant, $0 < c < 1$
- When utility functions are differentiable, best responses can be found by simple calculus

Example: Preparing for an Exam

- Here:
- For a given y, u_1 is a quadratic function of x
- Similarly for u_2
- player 1's best response to y is (c+y)/2
- player 2's best response to x is (c+x)/2

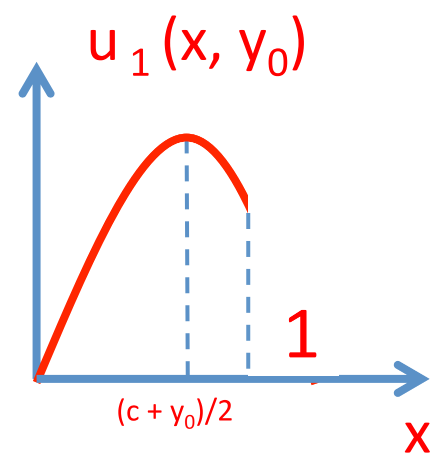

Joint Exam Preparation, Continued

- player 1's best response to y is (c+y)/2
- player 2's best response to x is (c+x)/2
- (c, c) is a Nash Equilibrium

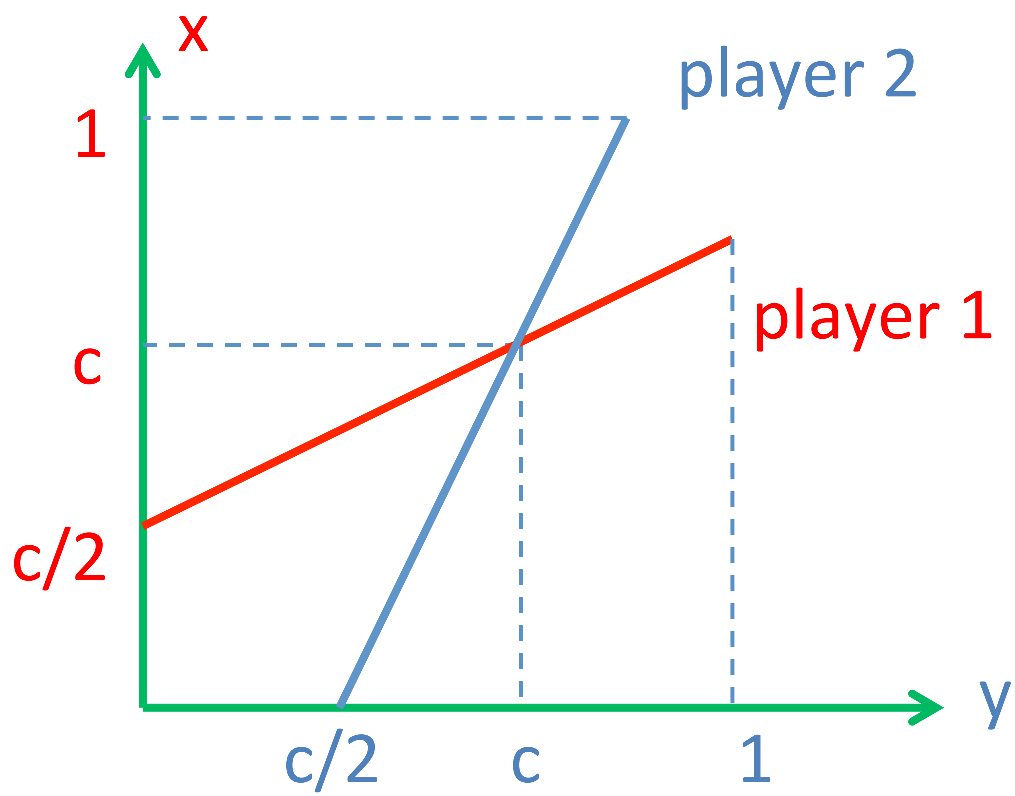

Joint Exam Preparation, Algebraically

- player 1's best response to y is (c+y)/2
- player 2's best response to x is (c+x)/2
- (x, y) is a Nash Equilibrium if
  - x is 1's best response to y
  - y is 2's best response to x
- x = (c+y)/2, y = (c+x)/2
- 2y = c+(c+y)/2 $\Rightarrow$ 4y = 3c+y
- Solution: x = c, y = c

Nash Equilibrium: Caution

1. The definition does not say that each game has a Nash equilibrium
   – some do not
2. The definition does not say that Nash equilibrium is unique
   – some games have many Nash equilibria
3. Nash equilibrium outcomes need not be strictly better than the alternatives, what matters is that they are not worse (to a deviation)
4. Not all equilibria are equally good
   – they can differ both in individual utilities and in total welfare

Non-existence of NE: Matching Pennies

	Heads	Tails
Heads	(1, -1)	(-1, 1)
Tails	(-1, 1)	(1, -1)

- Two players have 1 coin each
- They simultaneously decide whether to display their coin with Heads or Tails facing up
- If the coins match, player 1 gets both coins
- Otherwise player 2 gets them

no Nash equilibrium!

Non-existence of NE: Matching Pennies

	Heads	Tails
Heads	(1, -1)	(-1, 1)
Tails	(-1, 1)	(1, -1)

no Nash equilibrium!

Q: How would we play this game in practice?

A: Toss a coin

Matching Pennies: Randomization

	1/2 H	1/2 T
H	(1, -1)	(-1, 1)
T	(-1, 1)	(1, -1)

P[win]=P[loss]=1/2
E[utility] = 0

- Main idea: players may be allowed to play non-deterministically
- Suppose column player plays
  - H with probability 1/2
  - T with probability 1/2
- If we play H, the outcome is

  (H, H) w.p. 1/2 (+1);
  (H, T) w.p. 1/2 (-1)
- If we play T, the outcome is

  (T, H) w.p. 1/2 (-1);
  (T, T) w.p. 1/2 (+1)

Matching Pennies: Randomization

	1/2 H	1/2 T
H	(1, -1)	(-1, 1)
T	(-1, 1)	(1, -1)

No matter what we do,
P[win]=P[loss]=1/2

- If we play H w.p. p, T w.p. 1-p, we get

(H, H) w.p. p/2,

(T, H) w.p. (1-p)/2,

(H, T) w.p. p/2,

(T, T) w.p. (1-p)/2

Pr [+1] = Pr [(H, H) or (T, T)] = 1/2

Pr [-1] = Pr [(H, T) or (T, H)] = 1/2

How Should We Play?

- Suppose we (the row player) are playing against an opponent who mixes evenly: (H w.p. 1/2, T w.p. 1/2)
- Any strategy gives the same chance of winning (1/2)
- However, if we play H, the opponent can switch to playing T and win all the time
- Same if we play T
- If we play any action w.p. $p < 1/2$, the opponent can switch to this action and win w.p. $1-p > 1/2$
- Thus, the only sensible choice is for us to mix evenly, too

Mixed Strategies

- A mixed strategy of a player in a strategic game is a probability distribution over the player's actions
- If the set of actions is $\{a^1, \ldots, a^r\}$, a mixed strategy is a vector $\underline{\mathbf{p}} = (p^1, \ldots, p^r)$, where

  $p^i \geq 0$ for $i=1, \ldots, r$, $\quad p^1 + \ldots + p^r = 1$
- $p(a^i)$= probability that the player chooses action a^i
- Matching pennies: mixing evenly can be written as $\underline{\mathbf{p}} = (1/2, 1/2)$ or $p(H) = p(T) = 1/2$ or $1/2\ T + 1/2\ H$
- $P = (\underline{\mathbf{p}}_1, \ldots, \underline{\mathbf{p}}_n)$: mixed strategy profile
- Pure strategy: assigns probability 1 to some action

Mixed Strategies and Payoffs

- Suppose each player chooses a mixed strategy
- How do they reason about their utilities?
- Utilities need to be computed before the choice of action is realized
  - before the coin lands
- Mixed strategies generate a probability space
- Players are interested in their expected utility w.r.t. this space

Expected Utility (2 Players)

- Player 1's set of actions: $A = \{a^1, \ldots, a^r\}$
- Player 2's set of actions: $B = \{b^1, \ldots, b^s\}$
- Player 1's utility is given by $u_1: A \times B \rightarrow \mathbb{R}$
- If player 1 plays mixed strategy $\underline{\mathbf{p}} = (p^1, \ldots, p^r)$, and player 2 plays mixed strategy $\underline{\mathbf{q}} = (q^1, \ldots, q^s)$
- The expected utility of player 1 is

$$U_1(\underline{\mathbf{p}}, \underline{\mathbf{q}}) = \Sigma_{\,i=1, \ldots, r,\, j=1, \ldots, s}\; p^i\, q^j\, u_1(a^i, b^j)$$

- Similarly for player 2 (replace u_1 by u_2)

Expected Utility (n Players)

- Player i's set of actions: A_i
- Player i's utility is given by

  $u_i: A_1 \times \ldots \times A_n \to R$
- If player j plays mixed strategy $\underline{\mathbf{p}}_j$
- Then the expected utility of player i is

$$U_i(\mathbf{p}_1, \ldots, \mathbf{p}_n) = \sum_{(a_1,\ldots,a_n) \in A_1 \times \ldots \times A_n} \mathbf{p}_1(a_1) \ldots \mathbf{p}_n(a_n) u_i(a_1, \ldots, a_n)$$

Equilibria in Mixed Strategies

- Definition: A mixed strategy profile $P = (\underline{\mathbf{p}}_1, \ldots, \underline{\mathbf{p}}_n)$ is a mixed strategy Nash equilibrium if for any player i and any mixed strategy $\underline{\mathbf{p}}'$ of player i,

$$U_i(P) \geq U_i(P_{-i}, \underline{\mathbf{p}}')$$

- We refer to Nash equilibria in pure strategies as pure Nash equilibria

Equilibria in Mixed Strategies

- Theorem [Nash 1951]: Every n-player strategic game in which each player has a finite number of actions has at least one Nash equilibrium in mixed strategies

Properties of Mixed Nash Equilibria

1. Given a mixed strategy profile, can we verify that it is a mixed Nash equilibrium?

2. Given a strategic game, can we find all its mixed Nash equilibria?

Properties of Mixed Nash Equilibria

Checking if a profile is a mixed Nash equilibrium:

- Matching pennies: Can we easily verify that ((1/2, 1/2), (1/2, 1/2)) is a mixed equilibrium?
- We need to check all possible deviations:
  1. Deviations (p, 1-p) for player 1, for every p∈[0, 1]
  2. Deviations (q, 1-q) for player 2, for every q∈[0, 1]

- Inifinite number of possible deviations!

Properties of Mixed Nash Equilibria

- Is there an easier way?
- A mixed strategy is a convex combination of pure strategies:
- $\underline{\mathbf{p}} = (p^1, \ldots, p^r) =$
  $p^1(1,0,\ldots 0) + p^2(0,1,0,\ldots,0) \ldots + p^r(0,\ldots,1)$
- If a player has a profitable mixed deviation, there must be some pure strategy that is also profitable

Properties of Mixed Nash Equilibria

- Hence, it suffices to check only deviations to pure strategies
- Theorem: a mixed profile $P = (\underline{\mathbf{p}}_1, \ldots, \underline{\mathbf{p}}_n)$ is a mixed strategy Nash equilibrium if and only if for any player i and any pure strategy a of player i it holds that $U_i(P) \geq U_i(P_{-i}, a)$
- Corollary: If a profile P is a pure Nash equilibrium then it is also a mixed equilibrium

Example

	T	F
T	(3, 1)	(0, 0)
F	(0, 0)	(1, 3)

- $\underline{\mathbf{p}}$ = (4/5, 1/5), $\underline{\mathbf{q}}$ = (1/2, 1/2)
- $U_1(\underline{\mathbf{p}}, \underline{\mathbf{q}}) = .4 \times 3 + .1 \times 1 = 1.3$
- $U_2(\underline{\mathbf{p}}, \underline{\mathbf{q}}) = .4 \times 1 + .1 \times 3 = .7$
- To check whether ($\underline{\mathbf{p}}$, $\underline{\mathbf{q}}$) is a mixed NE, need to verify whether
  - $U_1(\underline{\mathbf{p}}, \underline{\mathbf{q}}) \geq U_1(F, \underline{\mathbf{q}})$
  - $U_1(\underline{\mathbf{p}}, \underline{\mathbf{q}}) \geq U_1(T, \underline{\mathbf{q}})$
  - $U_2(\underline{\mathbf{p}}, \underline{\mathbf{q}}) \geq U_2(\underline{\mathbf{p}}, F)$
  - $U_2(\underline{\mathbf{p}}, \underline{\mathbf{q}}) \geq U_2(\underline{\mathbf{p}}, T)$

YES $1.3 \geq .5 \times 1$?

NO $1.3 \geq .5 \times 3$?

Computing Mixed Nash Equilibria

- Support of a mixed strategy $\underline{\mathbf{p}}$:

$$\mathrm{supp}(\underline{\mathbf{p}}) = \{a \mid p(a) > 0\}$$

- Intuition: If an action is in the support of an equilibrium strategy, it should not be worse than any other pure strategy
- <u>Theorem</u>: suppose that P is a mixed Nash equilibrium, and $\underline{\mathbf{p}}$ is the strategy of player i.
If $p(x) > 0$ for some action $x \in A_i$,
then $U_i(P_{-i}, x) \geq U_i(P_{-i}, y)$ for any $y \in A_i$.
- <u>Corollary</u>: If $P = (P_{-i}, \underline{\mathbf{p}})$ is a mixed Nash equilibrium, and $x, y \in \mathrm{supp}(\underline{\mathbf{p}})$, then $U_i(P_{-i}, x) = U_i(P_{-i}, y)$.

Computing Mixed Nash Equilibria

- Consider a 2-player game, where
  - A = set of actions of the 1st player with $|A|=r$
  - B = set of actions of the 2nd player with $|B|=s$
- Let $A' \subseteq A$, $B' \subseteq B$
- We can find all mixed NE ($\underline{\mathbf{p}}$, $\underline{\mathbf{q}}$) with $\text{supp}(\underline{\mathbf{p}}) = A'$ and $\text{supp}(\underline{\mathbf{q}}) = B'$
  - By using previous theorem
- Main idea: Resort to solving a system of linear inequalities

Finding Mixed NE With Given Support

- Fix A', B', and let $p_1, \ldots, p_r$, $q_1, \ldots, q_s$ be variables
- Constraints:

(1) $\Sigma_{i=1, \ldots, r}\ p_i = 1$, $p_i \geq 0$ for each $i = 1, \ldots, r$

(2) $\Sigma_{j=1, \ldots, s}\ q_j = 1$, $q_j \geq 0$ for each $j = 1, \ldots, s$

(3) $p_i > 0$ for each $a^i \in A'$, $p_i = 0$ for each $a^i \notin A'$

(4) $q_j > 0$ for each $b^j \in B'$, $q_j = 0$ for each $b^j \notin B'$

(5) $\Sigma_{j=1, \ldots, s}\ q_j u_1(a^i, b^j) \geq \Sigma_{j=1, \ldots, s}\ q_j u_1(a^k, b^j)$
for each $a^i \in A'$ and each $a^k \in A$

(6) $\Sigma_{i=1, \ldots, r}\ p_i u_2(a^i, b^j) \geq \Sigma_{i=1, \ldots, r}\ p_i u_2(a^i, b^t)$
for each $b^j \in B'$ and each $b^t \in B$

- All constraints are linear $\Rightarrow$ can solve the system

Finding Mixed NE by Support Enumeration

- <u>Theorem</u>: (**p**, **q**) is a solution to the system (1)-(6) for given A', B' if and only it is a mixed Nash equilibrium and supp(**p**) = A', supp(**q**) = B'

- What if we want to find all mixed NE of this game?
- Go over all pairs A', B' such that A' $\subseteq$ A, B' $\subseteq$ B
- For each (A', B'), try to solve the system (1)-(6)
  - if the system does not have a solution, there is no mixed NE with support A', B'
  - Otherwise, every solution is a mixed NE with support A', B'

Finding Mixed NE by Support Enumeration

- What is the running time of this procedure?
  - Suppose |A|=r , |B|=s
- Then we need to solve 2^r x 2^s linear systems
  - for r = s = 3, this is 64 linear systems
- Infeasible by hand, and barely feasible by computer
- Other algorithms?

Complexity Issues

- Suppose we simply want to find one mixed Nash equilibrium
- Even for n = 2 players, known algorithms have worst case exponential time [Kuhn '61, Lemke-Howson '64, Mangasarian '64, Lemke '65]
- The Lemke-Howson remains among the most practical algorithms till today for 2 players
- Bad news: algorithms that are guaranteed to have substantially better running time than support enumeration are not known
  - there are reasons to believe they do not exist

Complexity Issues: A Few More Details

- The problem is unlikely to be **NP**-hard
  [Megiddo, Papadimitriou ' 89]
- Proved to be **PPAD**-complete even for 2 players
  (hardness still holds for finding a sufficiently close approximation to an equilibrium)
  [Daskalakis, Goldberg, Papadimitriou ' 06, Chen, Deng, Teng ' 06]
- Main implication: the problem is equivalent to finding approximate fixed points of continuous functions on convex and compact domains
  - i.e., unlikely to admit a polynomial time algorithm
- Proved **NP**-hard if we add more constraints (e.g. find an equilibrium that maximizes the social welfare)
  [Gilboa, Zemel ' 89, Conitzer, Sandholm ' 03]

Strictly Dominated Actions

- Definition: for a mixed strategy **p** and an action $b \in A_i$, **p** strictly dominates b if $U_i(S_{-i}, \mathbf{p}) > U_i(S_{-i}, b)$ for any profile S_{-i} of other players' mixed strategies
  - We can define strict domination for a pair of mixed strategies, too
- Fact 1: It is possible that a strategy is not dominated by a pure strategy but only by a mixed strategy
- Fact 2: It suffices to consider profiles S_{-i} for the other players that consist only of pure strategies

Example: Actions Dominated by Mixed Strategies

	L	R
T	(5, 5)	(0, 0)
C	(0, 0)	(5, 5)
B	(2, 0)	(2, 0)

- Action B of player 1 is not strictly dominated by T or C
- However, it is strictly dominated by their even mixture, i.e., 0.5T + 0.5C:
  - fix any strategy $\underline{\mathbf{s}}$ = (s, 1-s) of player 2
  - $U_1((.5, .5, 0), \underline{\mathbf{s}}) = .5s \times 5 + .5(1-s) \times 5 = 2.5$
  - $U_1(B, \underline{\mathbf{s}}) = 2$

Strictly Dominated Actions: an Algorithmic Perspective

- How can we check if an action is strictly dominated?
- Suppose there are 2 players with action sets

  $A = \{a^1, \ldots, a^r\}$ and $B = \{b^1, \ldots, b^s\}$
- If we want to check whether an action a^i of the 1^{st} player is strictly dominated:
- We need to find values for probabilities $p_1, \ldots, p_r$ s.t.
  - for every b^j in B we have the constraint $u_1(a^i, b^j) < p_1 u_1(a^1, b^j) + \ldots + p_r u_1(a^r, b^j)$
  - also, $\Sigma_{i=1, \ldots, r}\, p_i = 1$, $p_i \geq 0$ for all $i = 1, \ldots, r$
- If the system of linear inequalities has a solution we have strict domination

Strictly Dominated Actions and Nash Equilibria

- Theorem: a strictly dominated action is not used with positive probability in any mixed NE
- Hence, we can eliminate strictly dominated strategies first, and then solve the remaining game
- In some cases this can lead to a much simpler game to work with

Eliminating Strictly Dominated Strategies: The Advantage

- To find a mixed NE in original game by support guessing: 2^3 x 2^2 = 32 systems of linear inequalities

	L	R
T	(5, 3)	(0, 0)
C	(0, 0)	(6, 8)
~~B~~	~~(2, 4)~~	~~(2, 3)~~

BUT:

- B is strictly dominated by (T+C)/2,
- Thus it remains to solve a 2 x 2 game

Iterated Elimination of Strictly Dominated Actions

- Action B of player 1 is dominated by T or C
- None of the actions of player 2 is dominated
- If player 1 is rational, she would never play B

	L	M	R
T	(4, 4)	(4, 1)	(3, 0)
C	(3, 1)	(3, 4)	(4, 0)
B	~~(2, 0)~~	~~(2, 0)~~	~~(2, 6)~~

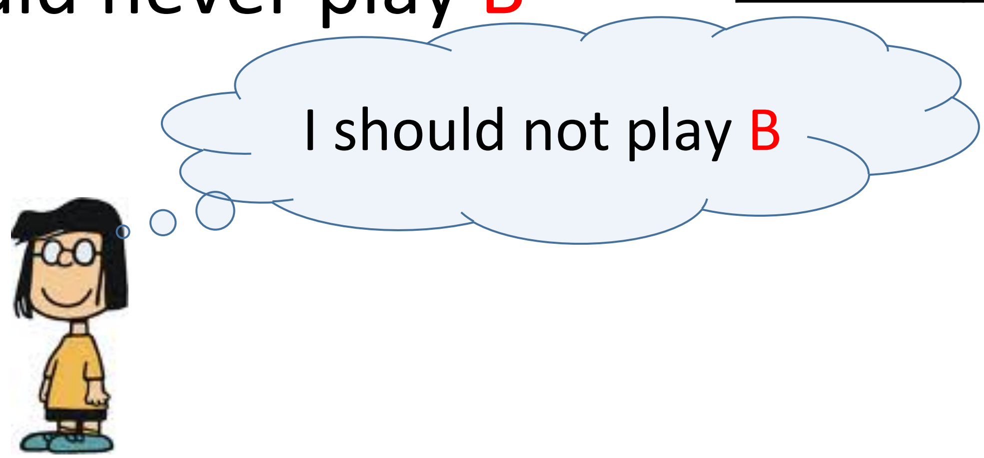

Iterated Elimination of Strictly Dominated Actions

- If player 2 knows player 1 is rational, he can assume player 1 does not play B
  - then player 2 should not play R

	L	M	R
T	(4, 4)	(4, 1)	(3, 0)
C	(3, 1)	(3, 4)	(4, 0)
B	~~(2, 0)~~	~~(2, 0)~~	~~(2, 6)~~

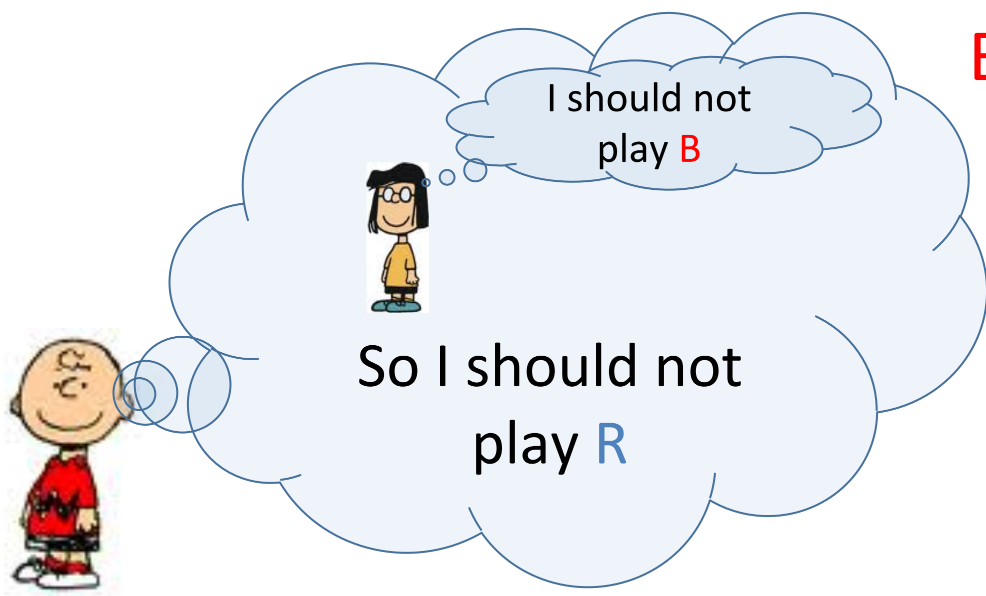

Iterated Elimination of Strictly Dominated Actions

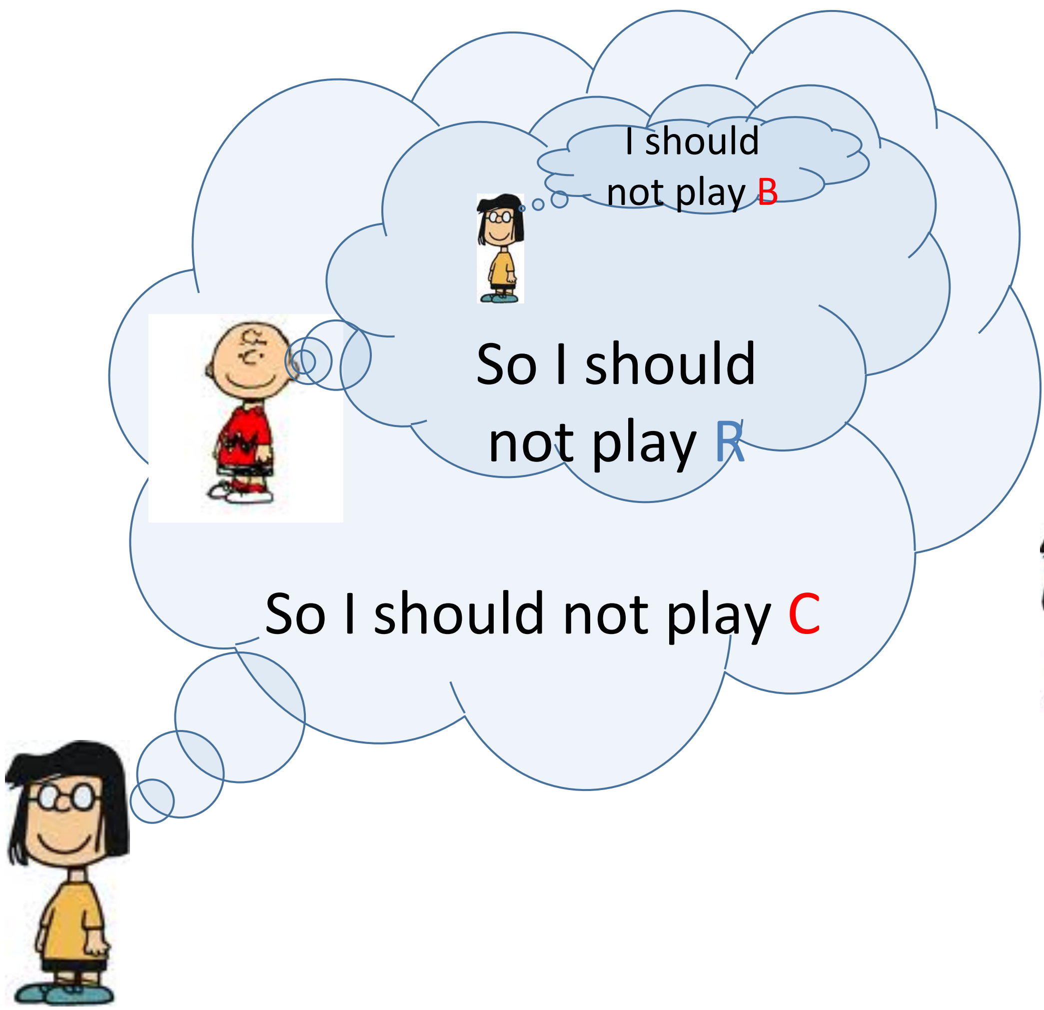

	L	M	R
T	(4, 4)	~~(4, 1)~~	~~(3, 0)~~
C	~~(3, 1)~~	~~(3, 4)~~	~~(4, 0)~~
B	~~(2, 0)~~	~~(2, 0)~~	~~(2, 6)~~

Iterated Elimination of Strictly Dominated Actions, Formally

- Given: an n-player game
  - pick a player i that has a strictly dominated action
  - remove some strictly dominated action of player i
  - repeat until no player has a strictly dominated action
- Fact: the set of surviving actions is independent of the elimination order
  - i.e., which agent was picked at each step

Iterated Elimination of Strictly Dominated Actions and Nash Equilibria

- <u>Theorem</u>: For a game G, suppose that after iterated elimination of strictly dominated actions the set of surviving actions of player i is A'_i. Then for any mixed Nash equilibrium $(\underline{\mathbf{p}}_1, \ldots, \underline{\mathbf{p}}_n)$ of G, $\text{supp}(\underline{\mathbf{p}}_i) \subseteq A'_i$ for all $i = 1, \ldots, n$.
  - in words: iterated elimination of strictly dominated actions cannot destroy Nash equilibria

Weakly Dominated Actions

- An action a of player i is weakly dominated by his mixed strategy **p** if
  - $U_i(\mathbf{s}_{-i}, a) \le U_i(\mathbf{s}_{-i}, \mathbf{p})$ for any profile $\mathbf{s}_{-i}$ of other players' actions
  - and $U_i(\mathbf{s}_{-i}, a) < U_i(\mathbf{s}_{-i}, \mathbf{p})$, for at least one profile $\mathbf{s}_{-i}$
- If we eliminate weakly dominated actions, we can lose Nash equilibria:
  - T weakly dominates B
  - L weakly dominates R
  - yet, (B, R) is a Nash equilibrium

	L	R
T	(2, 2)	(3, 0)
B	(0, 3)	(3, 3)

Iterated Elimination of Weakly Dominated Actions and Nash Equilibria

- The elimination order matters in iterated deletion of weakly dominated strategies
- Each order may eliminate a different subset of Nash equilibria
- Can we lose all equilibria of the original game?
- <u>Theorem</u>: For every game where each player has a finite action space, there is always at least one equilibrium that survives iterated elimination of weakly dominated strategies
  - thus: if we care for finding just one Nash equilibrium, no need to worry about elimination order

A special case: 0-sum games

- Games where for any actions $a \in A_1$, $b \in A_2$
  $u_1(a, b) = -u_2(a, b)$
- The payoff of one player is the payment made by the other
- Also referred to as strictly competitive
- It suffices to use only the matrix of player 1 to represent such a game

4	2
1	3

- How should we play in such a game?

A special case: 0-sum games

- Idea: Pessimistic play
- Assume that no matter what you choose the other player will pick the worst outcome for you
- Reasoning of player 1:
  - If I pick row 1, in worst case I get 2
  - If I pick row 2, in worst case I get 1
  - I will pick the row that has the best worst case
  - Payoff = $\max_i \min_j R_{ij} = 2$
- Reasoning of player 2:
  - If I pick column 1, in worst case I pay 4
  - If I pick column 2, in worst case I pay 3
  - I will pick the column that has the smallest worst case payment
  - Payment = $\min_j \max_i R_{ij} = 3$

4	2
1	3

A special case: 0-sum games

- In general $\max_i \min_j R_{ij} \neq \min_j \max_i R_{ij}$
- Pessimistic play with pure strategies does not always lead to a Nash equilibrium
- Suppose we do the same with mixed strategies
- We would need then to compute the quantities:
  - $\max_{\mathbf{s}} \min_{\mathbf{t}} u_1(\mathbf{s}, \mathbf{t})$
  - $\min_{\mathbf{t}} \max_{\mathbf{s}} u_1(\mathbf{s}, \mathbf{t})$

A special case: 0-sum games

Back to the example:

- We deal first with $\max_{\mathbf{s}} \min_{\mathbf{t}} u_1(\mathbf{s}, \mathbf{t})$
- The maximum is achieved at some strategy $\mathbf{s} = (s_1, s_2) = (s_1, 1 - s_1)$
- Fact: Given **s**, the quantity $\min_{\mathbf{t}} u_1(\mathbf{s}, \mathbf{t})$ is minimized at a pure strategy for player 2

4	2
1	3

- Hence we need to compute:

$$\max_{s1} \min \{ 4s_1 + 1\text{-}s_1, 2s_1 + 3(1\text{-}s_1) \} = \max_{s1} \min \{ 3s_1 + 1, 3 - s_1 \}$$

A special case: 0-sum games

- Computing $\max_{s_1} \min \{ 3s_1 + 1, 3 - s_1 \}$:
- Just need to maximize the minimum of 2 lines

4	2
1	3

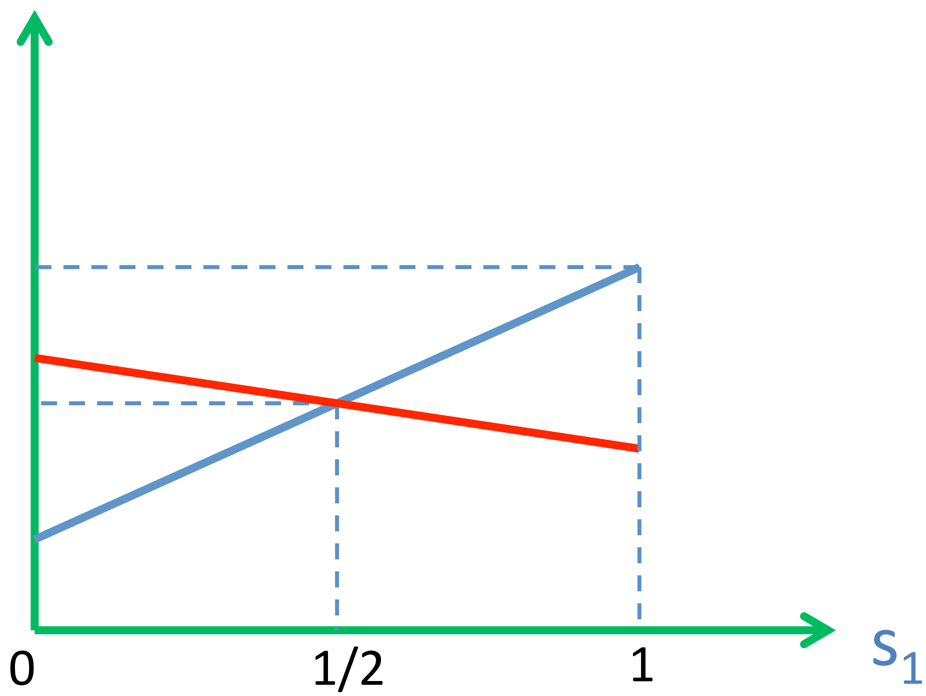

A special case: 0-sum games

- Computing $\max_{s1} \min\{ 3s_1 + 1, 3 - s_1 \}$:
- Just need to maximize the minimum of 2 lines

4	2
1	3

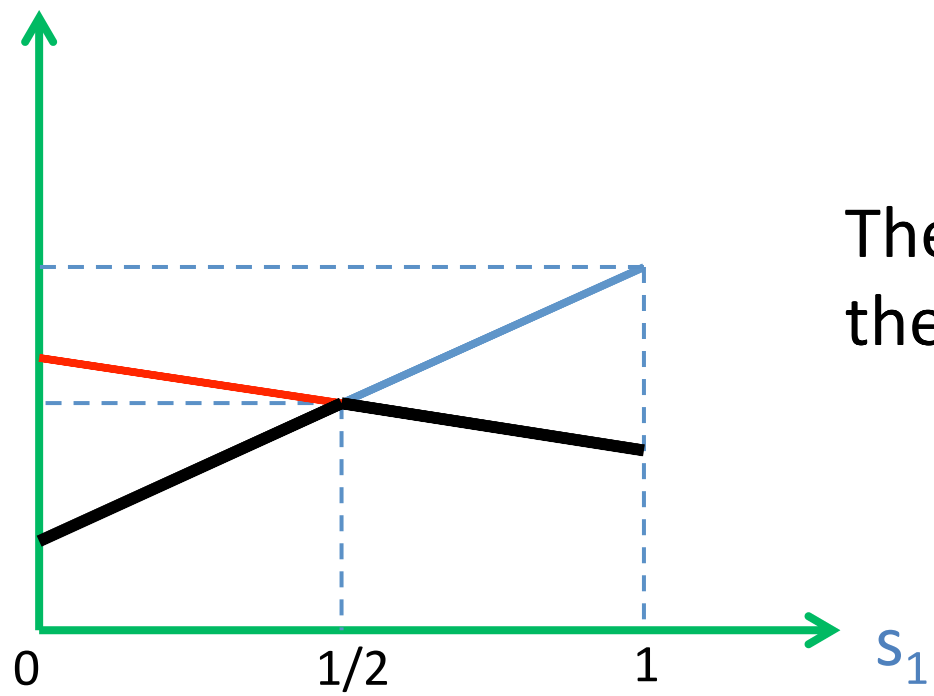

The min. is maximized at the intersection ➔ $s_1 = 1/2$

A special case: 0-sum games

Overall:

- $\max_{\mathbf{s}} \min_{\mathbf{t}} u_1(\mathbf{s}, \mathbf{t}) = \max_{s1} \min \{ 3s_1 + 1, 3 - s_1 \}$
$= 3*1/2 + 1 = 5/2$
- Player 1 should play $\mathbf{s} = (1/2, 1/2)$ to guarantee such a payoff
- By doing the same analysis for player 2, we have $\min_{\mathbf{t}} \max_{\mathbf{s}} u_1(\mathbf{s}, \mathbf{t}) = 5/2$
- Player 2 should play $\mathbf{t} = (1/4, 3/4)$ to guarantee such a payment
- Is it a coincidence that

  $\max_{\mathbf{s}} \min_{\mathbf{t}} u_1(\mathbf{s}, \mathbf{t}) = \min_{\mathbf{t}} \max_{\mathbf{s}} u_1(\mathbf{s}, \mathbf{t})$?

4	2
1	3

The Main Result for 0-sum games

<u>Theorem</u>: For any finite 0-sum game:

1. $\max_{\mathbf{s}} \min_{\mathbf{t}} u_1(\mathbf{s}, \mathbf{t}) = \min_{\mathbf{t}} \max_{\mathbf{s}} u_1(\mathbf{s}, \mathbf{t})$ (referred to as the value of the game)
2. The (mixed) strategy profile (**s**, **t**), where the value of the game is achieved, forms a Nash equilibrium
3. All Nash equilibria yield the same payoff to the players
4. If (**s**, **t**), (**s'**, **t'**) are Nash equilibria, then (**s**, **t'**), (**s'**, **t**) are also Nash equilibria

Extensive-Form Games

Simultaneous vs. Sequential Moves

- So far, we have considered games where players choose their strategies simultaneously

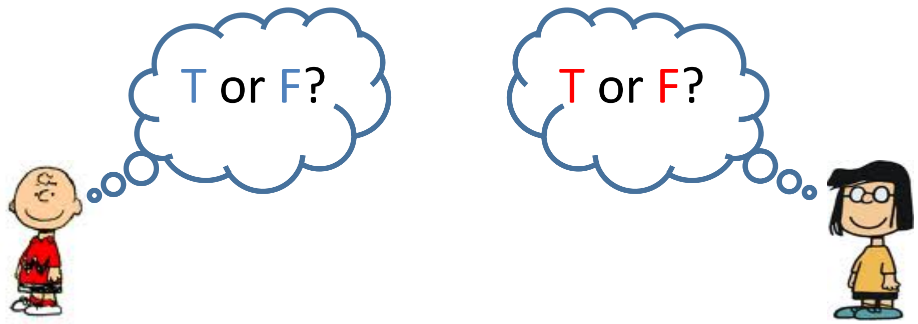

- What if players take turns choosing their actions?

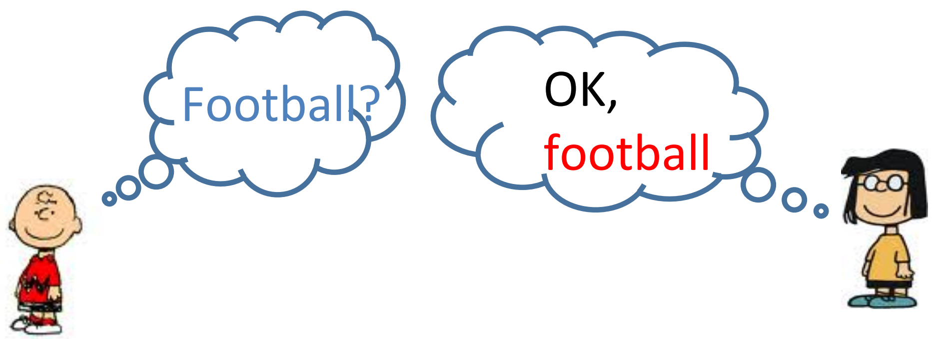

Games With Sequential Moves: More Examples

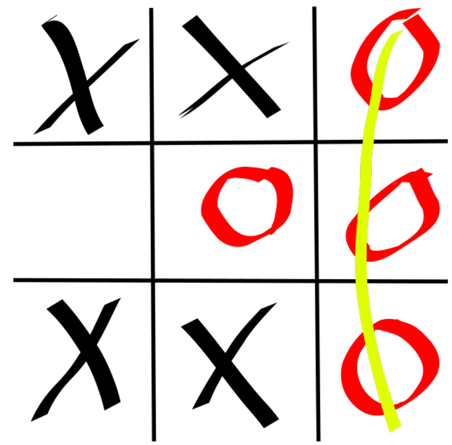

Chess and tic-tac-toe may differ in difficulty, but the underlying principle is the same: players take turns making moves, and eventually either one of the players wins or there is a tie

Another Example: Market Entry

- Suppose that in some country firm 1 is currently the only available fast food chain
- Firm 2 considers opening their restaurants in that country
- Firm 2 has 2 actions: enter (E), stay out (S)
- If firm 2 stays out, firm 1 need not do anything
- If firm 2 enters, firm 1 can either fight (F) (lower prices, aggressive marketing) or accept (A)

Market Entry: Payoffs

- If firm 2 stays out, its payoff is 0, and firm 1 has a payoff of 2
- If firm 2 enters and firm 1 fights, each gets a payoff of -1
- If firm 2 enters and firm 1 accepts, they share the market, so both get a payoff of 1

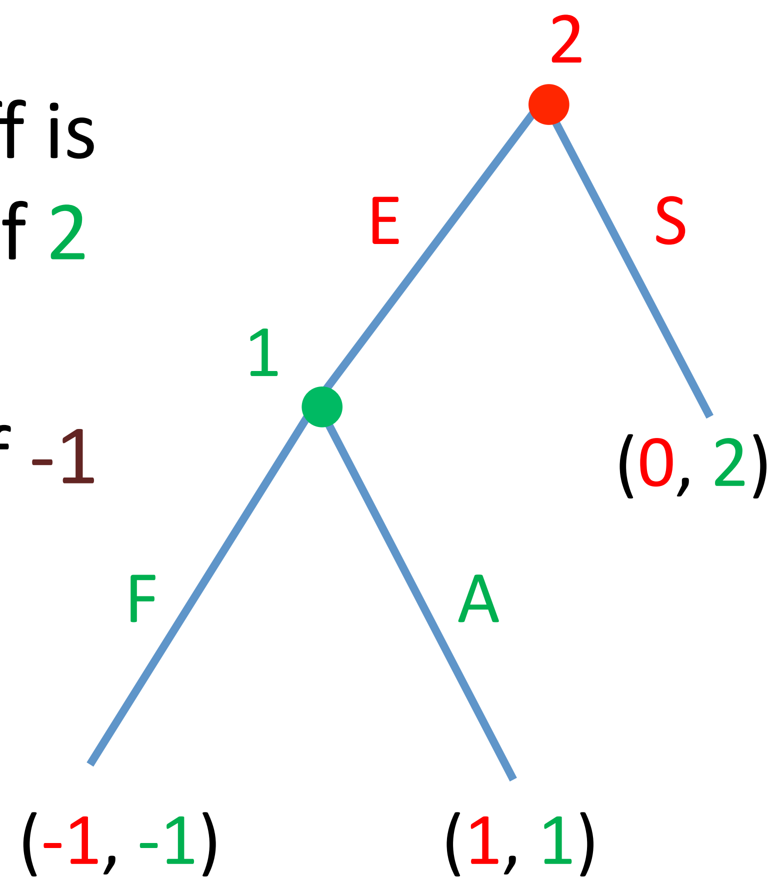

Extensive Form Games: General Case

- An extensive-form game is described by a game tree:
  - rooted tree, with root corresponding to the start of the game
  - each internal node of the tree is labeled by a player
  - each leaf is labeled by a payoff vector (assigning a payoff to each player in the game)
  - For a node labeled by a player X, all edges leaving the node are labeled by the actions of player X

Extensive Form Games: Playing The Game

- Let the label of the root be x
- Then the game starts by player x choosing an edge from the root; let y be the label of the endpoint of this edge
- Player y chooses next, etc.
- Players may appear more than once in the tree
- Not all players appear on all paths

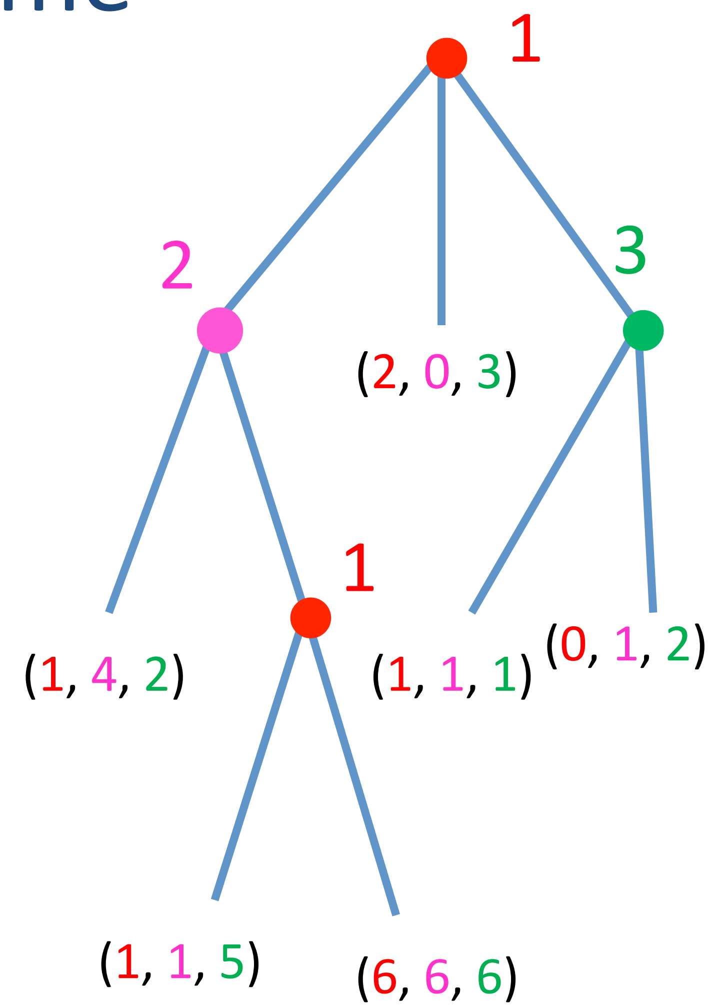

Extensive Form Games: Strategies

- Strategy of player x:
  a complete plan, i.e., which action would x choose for each node labeled with x
- Caution: Need to specify what to do even for nodes that seem unlikely to occur due to the players' choices
- Player 1 has 6 strategies:
  (L1, L4), (L1, R4), (C1, L4), (C1, R4), (R1, L4), (R1, R4)
- L4 looks redundant in (C1, L4) :
  if 1 chooses C1, he will not be able to choose L4
- But still (C1, L4) is a valid strategy
- If by mistake C1 is not played, then player 1 knows what to choose between L4, R4

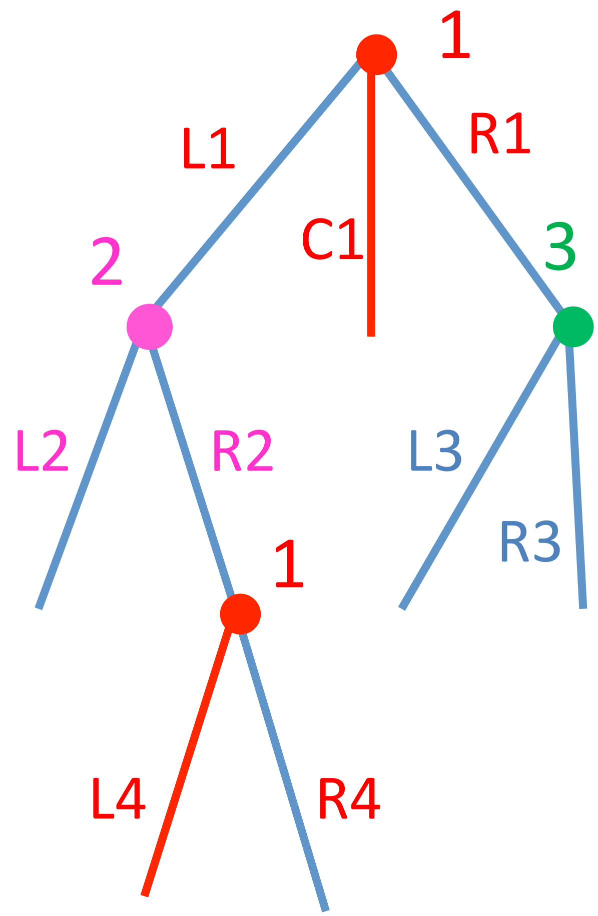

Market Entry: Predicting the Outcome

- How should players choose a strategy?
- Firm 1 can reason as follows:
  - If firm 2 enters, the best for me is to play A
- Firm 2 reasons as follows:
  - if I enter, firm 1 is better off accepting, so my payoff is 1
  - if I stay out, my payoff is 0
  - thus I am better off entering
- The only "rational" outcome is (E, A)
- Corresponds to a backward induction process

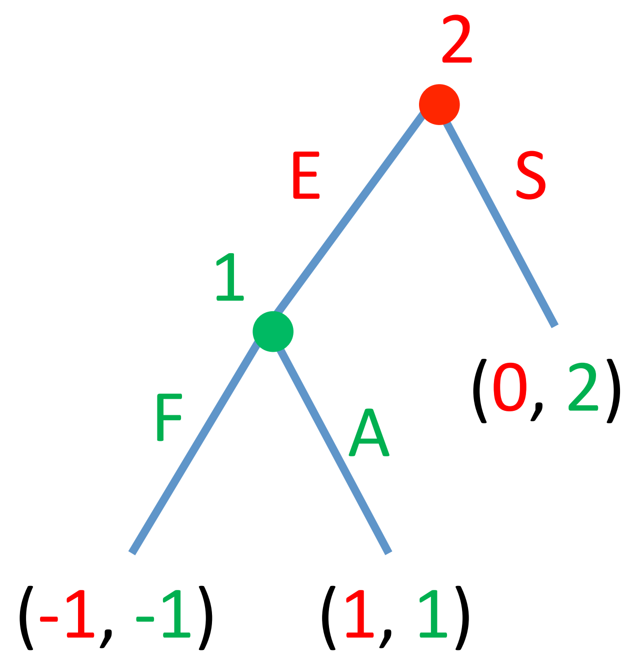

Predicting The Outcome: Backward Induction

- The outcome of the game can be predicted using backward induction:
  - Start with any node whose children are leaves only
  - For any such node, the agent who chooses the action will determine all payoffs including his own, so he will choose the action maximizing his payoff
    - breaking ties arbitrarily (we will come back to this)
  - Fix his choice of action, and delete other branches
  - Now his node has one outgoing edge, so it can be treated as a leaf
  - Repeat until the root's action is determined

Backward Induction: Example

- Player 1 prefers (6, 6, 6) to (1, 1, 5), so he chooses R4
- Player 2 prefers (6, 6, 6) to (1, 4, 2), so he chooses R2
- Player 3 prefers (0, 1, 2) to (1, 1, 1), so he chooses R3
- Player 1 prefers (6, 6, 6) to (2, 0, 3) and (0, 1, 2), so he chooses L1

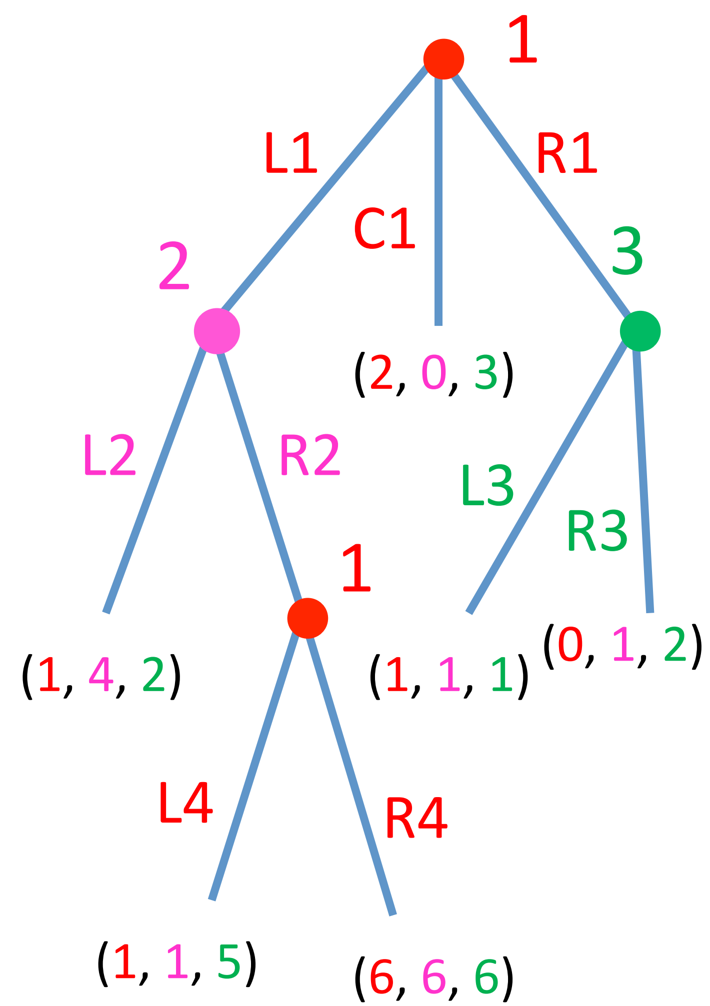

- Strategies for 1, 2, 3: (L1, R4), R2, R3

Converting Extensive Form Games Into Normal Form Games

- Given an extensive-form game G, we can list all strategies of each player
- Let N(G) be a normal-form game with the same set of players as G such that for each player i, {actions of player i in N(G)} = {strategies of player i in G}

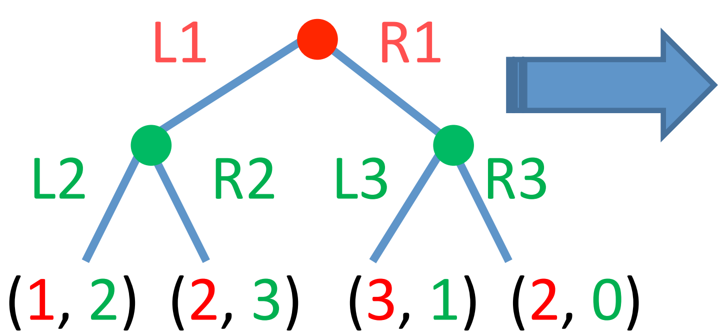

	(L2, L3)	(L2, R3)	(R2, L3)	(R2, R3)
L1	(1, 2)	(1, 2)	(2, 3)	(2, 3)
R1	(3, 1)	(2, 0)	(3, 1)	(2, 0)

Predicting the Outcome: Nash Equilibria of the Normal-Form Game

- Can we use the (pure) NE of N(G) as a prediction for the outcome of G?
- How do they relate to BI outcomes?
- Claim: any backward induction strategy profile in the extensive-form game G corresponds to a NE profile in the normal-form game N(G)
- Is the reverse true?

Market Entry Revisited

- Backward induction outcome of the extensive-form game is (E, A)
- Nash equilibria of the corresponding normal form game are (E, A) and (S, F)
- Thus, the converse is not true

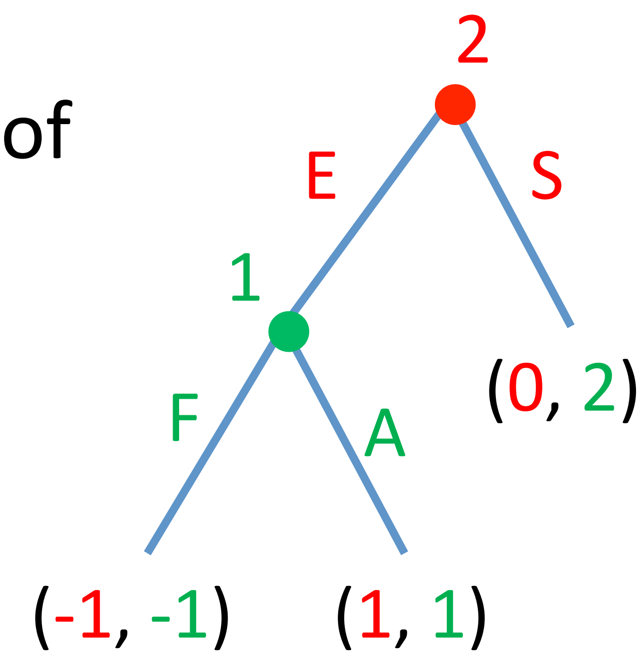

	A	F
E	(1, 1)	(-1, -1)
S	(0, 2)	(0, 2)

Market Entry Revisited

- The NE (S, F) is also not a “good” prediction
- (S, F) is a NE, because firm 1 promised to fight if firm 2 deviates to E
  - but this is an empty threat: it is irrational for firm 1 to fight!
- The matrix representation does not capture the fact that firm 2 moves first
- We need a different solution concept than just Nash equilibria!

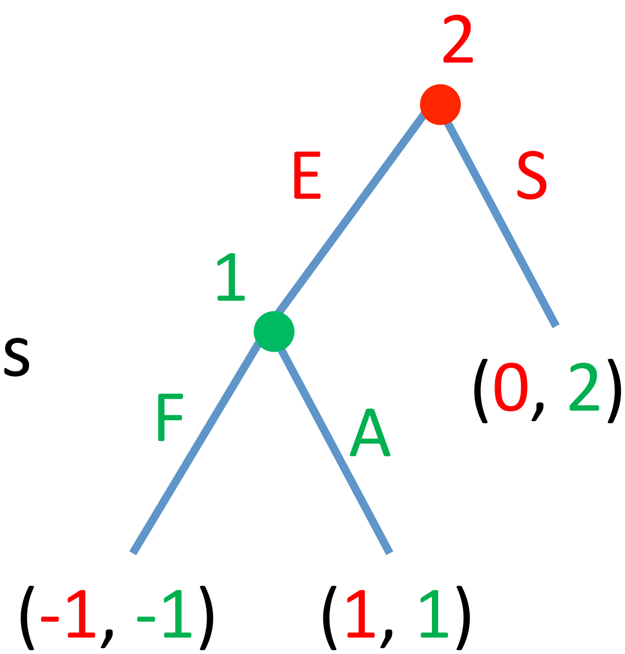

	A	F
E	(1, 1)	(-1, -1)
S	(0, 2)	(0, 2)

Subgames

- For any game G, each subtree defines a subgame G’
- In G’:
  - same set of players as G
  - set of actions of a player: subset of his actions in G
- Each strategy in G corresponds to a strategy in G’ (by projection)

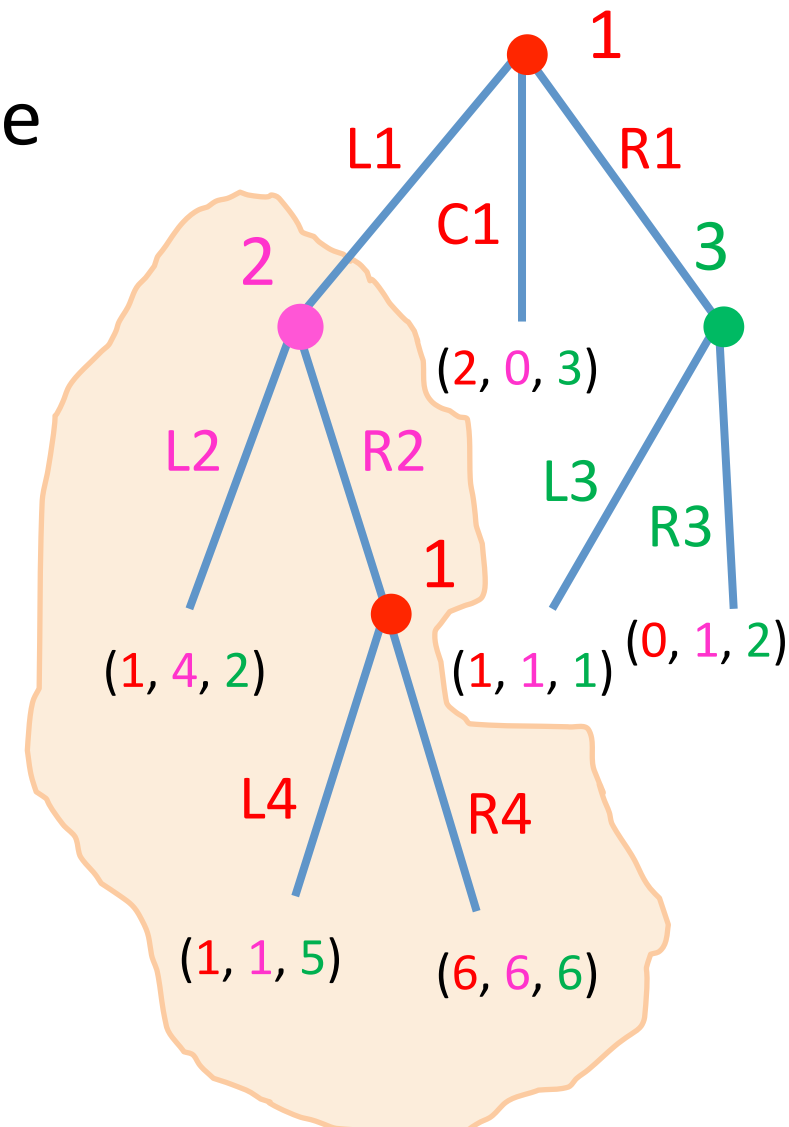

Strategies in Subgames

- Consider a strategy profile ((L1, R4), L2, R3) in G
- Its projection to G' is (R4, L2, Ø) and its projection to G'' is (Ø, Ø, R3)
- Generally, if $s = (a_1, \ldots, a_t)$ is a strategy of player i in G, and G' is a subgame of G, then the projection of s to G' consists of all actions in s associated with nodes of G'

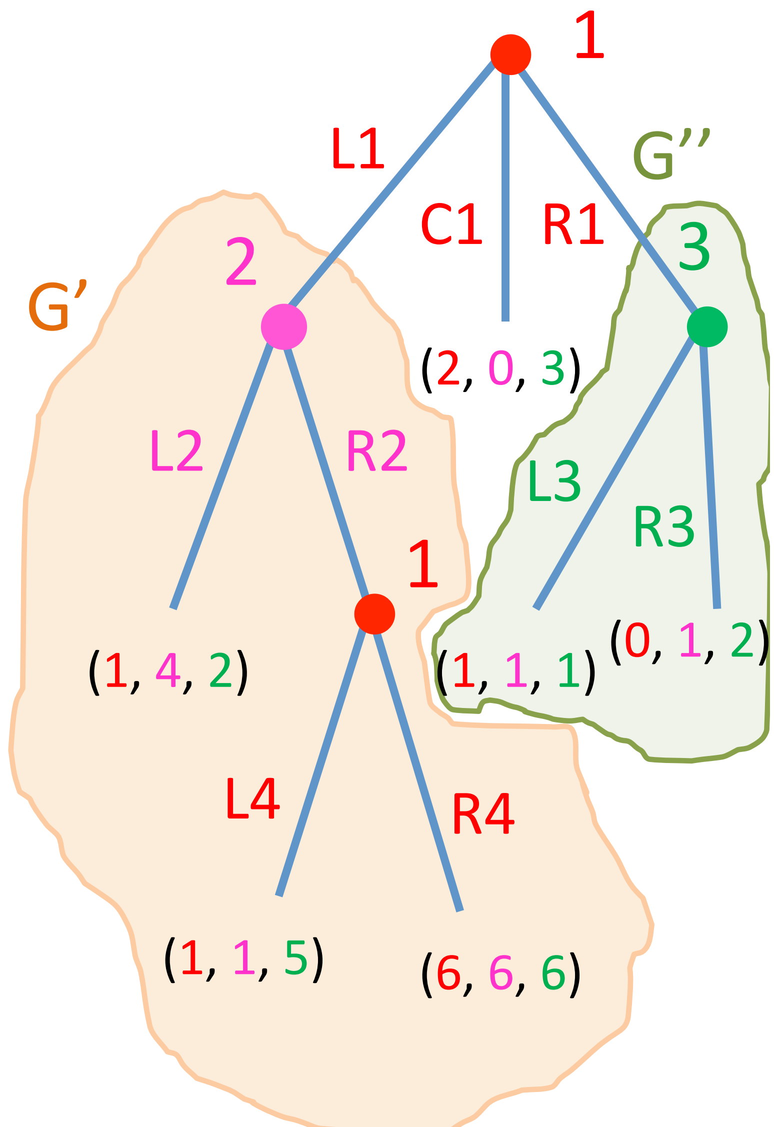

Subgames and Nash Equilibria

- Given an extensive-form game G, let
- N(G) be the associated normal-form game.
- **s** = (s_1, …, s_n) be a Nash equilibrium of N(G)
  - s_i is the strategy of player i
- Pick a subgame G' of G
- Let **s'** = (s'_1, …, s'_n) be projection of **s** on G'
- Is **s'** a NE in N(G')?
- Not always!

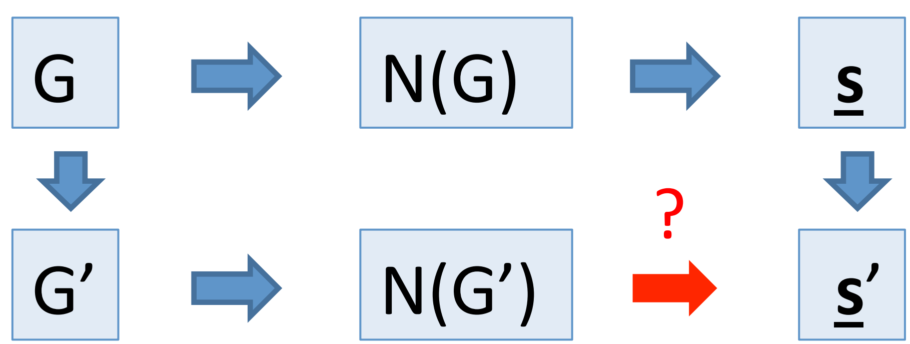

Equilibria in Subgames: an Example

- (S, F) is a Nash equilibrium in N(G)
- The projection of (S, F) to the left subgame G' is (Ø, F)
- (Ø, F) is not a NE in N(G'):
  - Firm 1 can profit by deviating to A

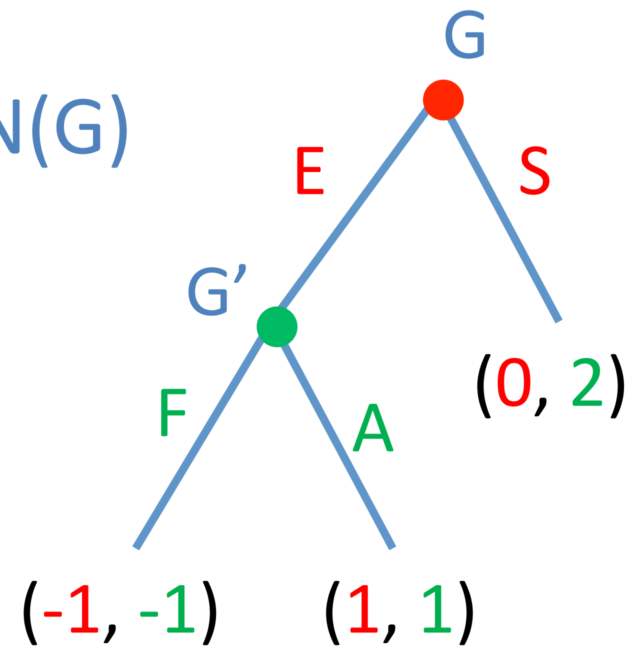

N(G)

	A	F
E	(1, 1)	(-1, -1)
S	(0, 2)	(0, 2)

N(G')

A	F
(1, 1)	(-1, -1)

Subgame-Perfect Equilibria

- Definition: Consider
  - G: an extensive-form game,
  - N(G): the corresponding normal-form game,
  - **s:** a NE of N(G).

  Then **s** is said to be a subgame-perfect NE (SPNE) if its projection onto any subgame G' of G is a NE of N(G')
- Why do we care for such strategy profiles?
- They are robust against any "change of plan"
  - At ANY node, every player is playing an optimal strategy against the projection of $\mathbf{s}_{-i}$ on the subgame starting from that node

Subgame-Perfect Equilibria

- Theorem: the output of backward induction is a subgame-perfect NE
- Intuition: backward induction proceeds subgame by subgame, finding an "optimal" solution in each

Corollaries

- A game is finite if the tree has finite depth and the outdegree of each node is finite
- Corollary 1: In a finite game a pure SPNE always exists
  - unlike pure NE in normal-form games

- Corollary 2: Consider a finite 2-player 0-sum extensive-form game with outcomes {win, lose, tie}. Then:
  - Either one of the players has a winning strategy
  - Or both have a strategy that can guarantee a tie
  - But it is often hard to tell which of the two applies!
  - Examples: tic-tac-toe (we can guarantee a tie), chess (open problem), …

Backward Induction: Handling Ties

- Suppose at some node, 2 or more branches lead to maximal payoff for the player who is choosing an action
- Then both lead to (distinct) SPNE
- If we want to find all SPNE, we need to explore all optimal choices at each node during the backward induction process

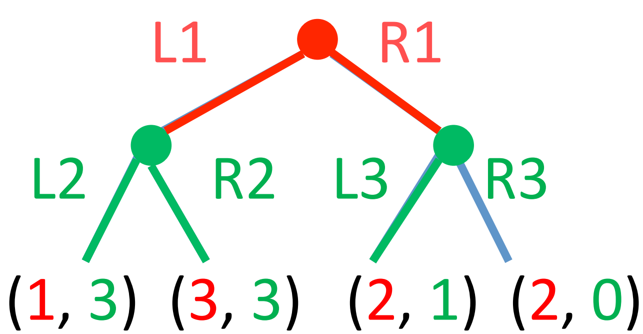

(L2, L3), R1

(R2, L3), L1

Bayesian Games

Games with Imperfect Information

- So far, we have assumed that the players know each others' payoffs for all strategy profiles
- However, this is not always the case:
  - in Battle of Sexes, one player may be uncertain that the other player enjoys their company
  - In an auction, bidders may be uncertain about the valuation of other participants

Example: Battle of Sexes

- Suppose P1 in uncertain whether P2 wants to go out with her:
  - w.p. 1/2, P2 enjoys P1's company
  - w.p. 1/2, P2 prefers to avoid P1's company
- P2 knows whether he wants to go out with P1

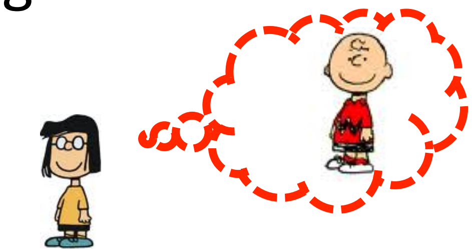

1

Pr = 1/2

2	T	F
T	(2, 1)	(0, 0)
F	(0, 0)	(1, 2)

Pr = 1/2

2	T	F
T	(2, 0)	(0, 2)
F	(0, 1)	(1, 0)

- 2 possible states of the world
  - P2 knows the state, P1 does not

Strategies

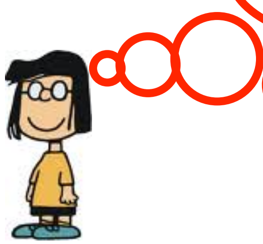

I believe that if Charlie wants to go out, he'll choose T, else he'll choose F, so if I choose T, my expected payoff will be 2 x 1/2 + 0 x 1/2 = 1

- Marcie's strategy: T or F
- Charlie's strategy: T or F
- However, when Marcie chooses her strategy, she needs to form a belief about Charlie's behavior in both states
  - in her mind Charlie's strategy is a pair (X, Y):
  - X is what would Charlie do if he wants to meet her (T or F)
  - Y is what would Charlie do if he wants to avoid her (T or F)

Strategies

- P2 (Charlie) can be of one of 2 types: "meet" or "avoid"
- When describing P2's strategy, we need to specify what each type of P2 would do
  - P2 knows what type he is, so he only needs one component of this description
  - P1 (Marcie) needs both components to calculate her payoffs
- Expected payoffs of P1 for each possible strategy of P2:

	(T, T)	(T, F)	(F, T)	(F, F)
T	2· ½ +2· ½ = 2	2· ½ +0· ½ = 1	0· ½ +2· ½ = 1	0· ½ +0· ½ = 0
F	0· ½ +0· ½ = 0	0· ½ +1· ½ = ½	1· ½ +0· ½ = ½	1· ½ +1· ½ = 1

Strategies

- Alternative interpretation:
  - before the game starts, P2 does not know his type
  - he needs to select his strategy for both types
  - then he learns his type
- In the table below, (X, Y, Z) indicates that
  - P1's expected payoff is X,
  - the payoff of the 1st type of P2 ("meet") is Y,
  - the payoff of the 2nd type of P2 ("avoid") is Z

	(T, T)	(T, F)	(F, T)	(F, F)
T	(2, 1, 0)	(1, 1, 2)	(1, 0, 0)	(0, 0, 2)
F	(0, 0, 1)	(½, 0, 0)	(½, 2, 1)	(1, 2, 0)

Which Strategy Profiles Are Stable?

- Strategy profile: a list of 3 actions (a, b, c), where
  - a is the action of P1
  - b is the action of type "meet" of P2
  - c is the action of type "avoid" of P2
- Intuitively, a strategy profile is stable if neither P1 nor any of the two types of P2 can increase their expected payoff by changing their action

	(T, T)	(T, F)	(F, T)	(F, F)
T	(2, 1, 0)	(1, 1, 2)	(1, 0, 0)	(0, 0, 2)
F	(0, 0, 1)	(½, 0, 0)	(½, 2, 1)	(1, 2, 0)

Stable Profiles

- (T, T, F) is stable:
  - if P1 deviates to F, her utility goes down to 1/2
  - if type 1 of P2 deviates to F, his utility goes down to 0
  - if type 2 of P2 deviates to T, his utility goes down to 0
- (T, T, T) is not stable:
  - type 2 of P2 can deviate to F and increase his payoff by 2

	(T, T)	(T, F)	(F, T)	(F, F)
T	(2, 1, 0) ☹	(1, 1, 2) ☺	(1, 0, 0)	(0, 0, 2)
F	(0, 0, 1)	(½, 0, 0)	(½, 2, 1)	(1, 2, 0)

Stable Profiles

- (F, T, T), (F, T, F), and (F, F, T) are not stable:
  - P1 can deviate to T and increase her payoff
- (T, F, T) and (T, F, F) are not stable:
  - type 1 of P2 can deviate to T and increase his payoff by 1
- (F, F, F) is not stable:
  - type 2 of P2 can deviate to T and increase his payoff by 1

	(T, T)	(T, F)	(F, T)	(F, F)
T	(2, 1, 0) ☹	(1, 1, 2) ☺	(1, 0, 0) ☹	(0, 0, 2) ☹
F	(0, 0, 1) ☹	(½, 0, 0) ☹	(½, 2, 1) ☹	(1, 2, 0) ☹

Tweaking the Game

- If Marcie thinks that "meet" or "avoid" are equally likely, her payoffs are as follows:

	(T, T)	(T, F)	(F, T)	(F, F)
T	$2\cdot \frac{1}{2} + 2\cdot \frac{1}{2} = 2$	$2\cdot \frac{1}{2} + 0\cdot \frac{1}{2} = 1$	$0\cdot \frac{1}{2} + 2\cdot \frac{1}{2} = 1$	$0\cdot \frac{1}{2} + 0\cdot \frac{1}{2} = 0$
F	$0\cdot \frac{1}{2} + 0\cdot \frac{1}{2} = 0$	$0\cdot \frac{1}{2} + 1\cdot \frac{1}{2} = \frac{1}{2}$	$1\cdot \frac{1}{2} + 0\cdot \frac{1}{2} = \frac{1}{2}$	$1\cdot \frac{1}{2} + 1\cdot \frac{1}{2} = 1$

- If she thinks that Charlie is of type "meet" w. p. 2/3 and "avoid" w. p. 1/3, her payoffs change:

	(T, T)	(T, F)	(F, T)	(F, F)
T	$2\cdot \frac{2}{3} + 2\cdot \frac{1}{3} = 2$	$2\cdot \frac{2}{3} + 0\cdot \frac{1}{3} = 1\frac{1}{3}$	$0\cdot \frac{2}{3} + 2\cdot \frac{1}{3} = \frac{2}{3}$	$0\cdot \frac{2}{3} + 0\cdot \frac{1}{3} = 0$
F	$0\cdot \frac{2}{3} + 0\cdot \frac{1}{3} = 0$	$0\cdot \frac{2}{3} + 1\cdot \frac{1}{3} = \frac{1}{3}$	$1\cdot \frac{2}{3} + 0\cdot \frac{1}{3} = \frac{2}{3}$	$1\cdot \frac{2}{3} + 1\cdot \frac{1}{3} = 1$

Stable Profiles in the Tweaked Game

- Charlie's payoffs are the same as before
- (T, T, F) is stable:
  - e.g., if P1 deviates to F, her utility goes down to 1/2
- (F, F, T) is stable, too:
  - e.g., if P1 deviates to T, her utility remains the same
- No other profile is stable

	(T, T)	(T, F)	(F, T)	(F, F)
T	(2, 1, 0)	(1⅓, 1, 2) ☺	(⅔, 0, 0)	(0, 0, 2)
F	(0, 0, 1)	(⅓, 0, 0)	(⅔, 2, 1) ☺	(1, 2, 0)

Example: both P1 and P2 can be of type “meet” or “avoid”

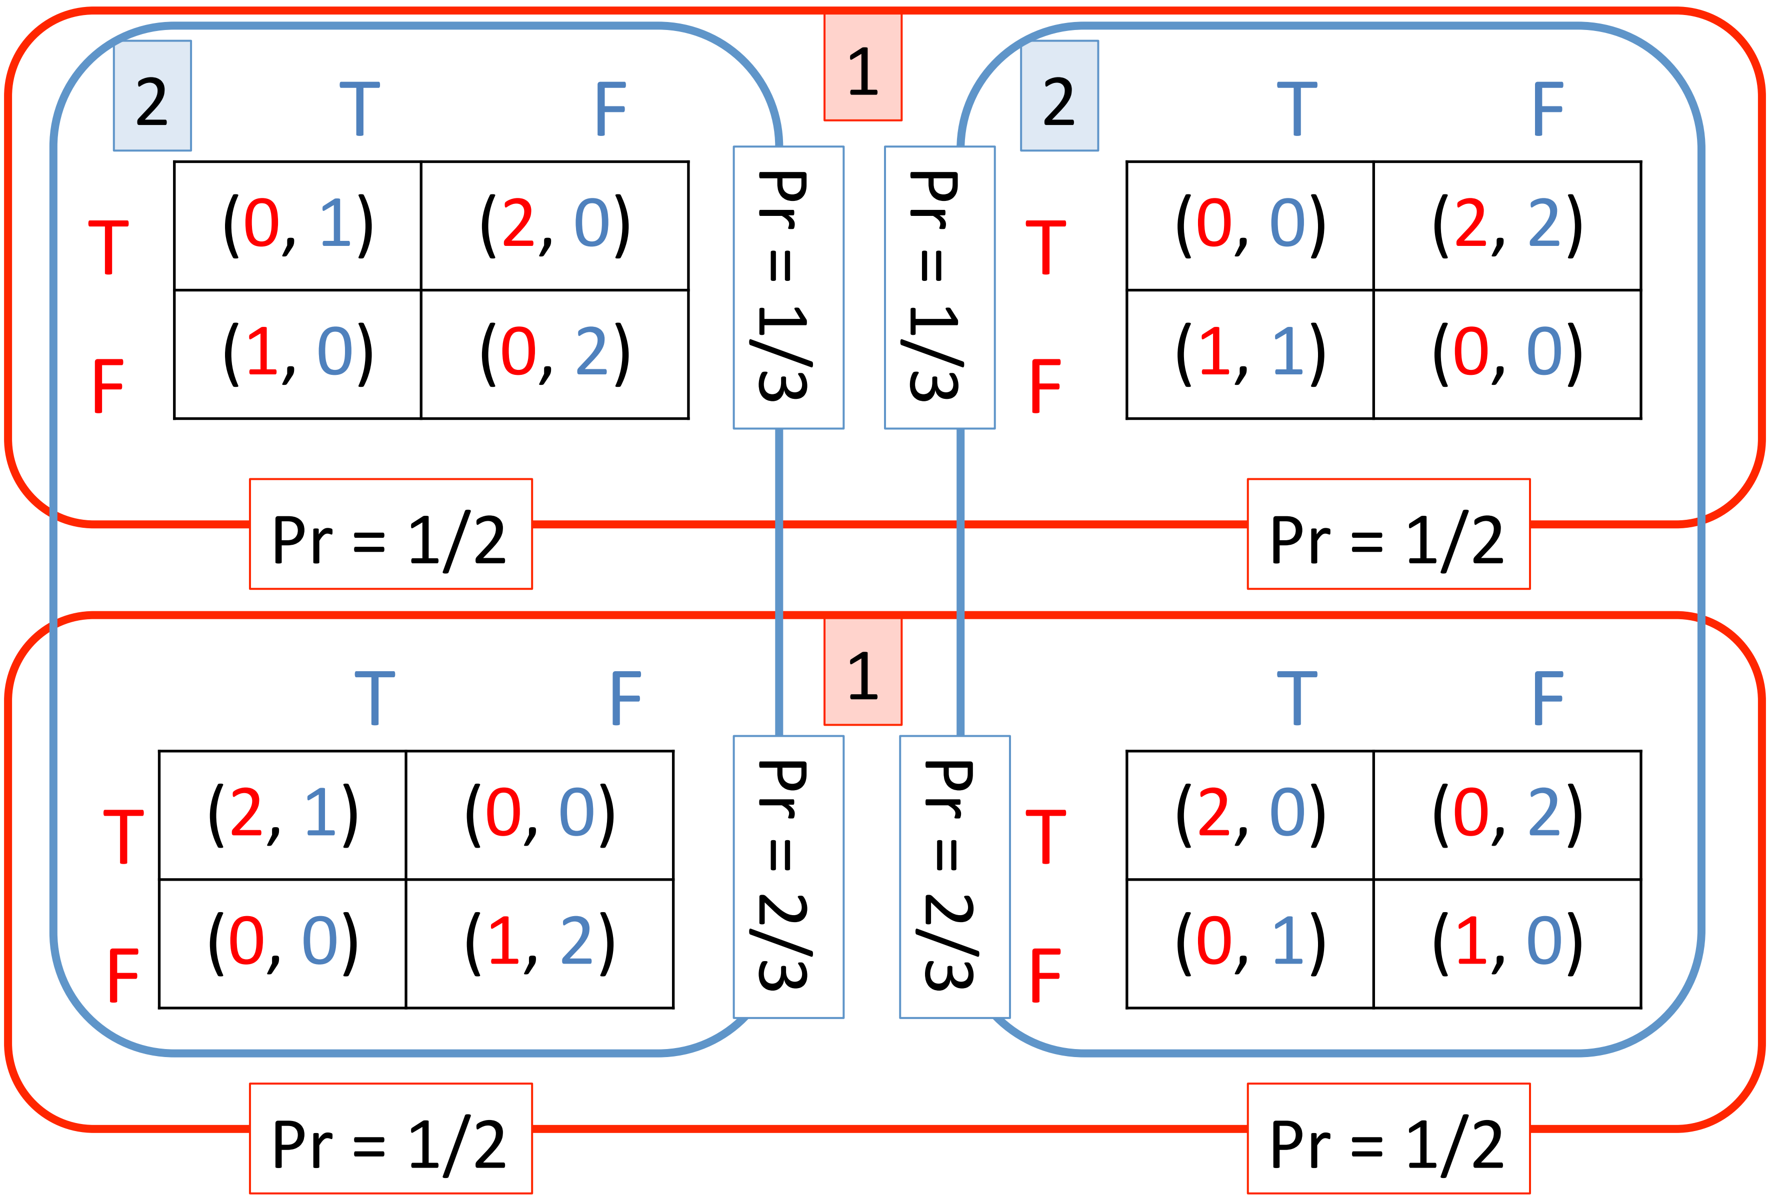

Interpretation

- Both P1 and P2 can be of type “meet” or of type “avoid”
- P1 knows her type, and believes that P2 is of type “meet” w.p. 1/2 and of type “avoid” w.p. 1/2
- P2 knows his type and believes that P1 is of type “meet” w.p. 2/3 and of type “avoid” w.p. 1/3
- P1 knows which of the red boxes she is in, but cannot determine the state within the box
- P2 knows which of the blue boxes he is in, but cannot determine the state within the box

Types and States

- In Bayesian games, each player may have several types
  - e.g., $Charlie^M$ or $Charlie^A$
- A player's type determines his preferences over action profiles
  - $Charlie^M$ prefers (T, T) to (T, F)
- A state is a collection of types (one for each player)
  - ($Marcie^A$, $Charlie^M$)
- in each state, each player's type is fixed
  - i.e., each state corresponds to a payoff matrix

Types and States, Continued

- Each player knows his type, and has a probability distribution over other players' types
  - Charlie knows he is of type "meet",
    and believes that Marcie is
    of type "meet" w.p. 2/3, and of type "avoid" w.p. 1/3
- Each player's strategy
  prescribes an action for each of his types
  - Charlie: T for Charlie^M, F for Charlie^A
- To compute expected payoffs, players take into account the probability of each type

Bayesian Game: Definition

- A Bayesian game G is given by
  - a set of players $N = \{1, ..., n\}$
  - for each player i, a set of actions A_i
  - for each player i, a set of types $T_{i,}$ $|T_i| = m$
  - for each type t of player i, a belief $p_{i, t}$ about all other agents' types
    - $p_{i, t}$ assigns probabilities to each vector $\underline{\mathbf{t}}_{-i}$ in $T_1 \times ... T_{i-1} \times T_{i+1} \times ... \times T_n$
  - for each type t of player i, a payoff function $u_{i, t}$ that assigns a payoff to each vector in $A_1 \times ... \times A_n$
- $G = (N, A_1, ..., A_n, T_1, ..., T_n, p_{1, 1}, ..., p_{n, m}, u_{1, 1}, ..., u_{n, m})$

Bayesian Games and Normal Form Games

- We can think of each type of each player as a separate player who chooses his own action
- $G \rightarrow E(G)$, where $E(G)$ is a normal-form game
- The set of players in $E(G)$ is $N' = \bigcup_{i \in N} T_i$
  - {$Charlie^M$, $Charlie^A$, $Marcie^M$, $Marcie^A$}
- For each player j in T_i, his set of actions is A_i
- What is the payoff of a player j in T_i for a given strategy profile?
  - it must take into account the action and the probability of each type of players in $\bigcup_{k \neq i} T_k$, but not the actions of other players in T_i

Computing Utilities in E(G): Example

- Consider the BoS game where both Marcie and Charlie can be of both types (M and A)
- Marcie believes that Charlie is of type M w.p. 1/3, A w.p. 2/3
- Charlie believes that Marcie is of type M w.p. 4/5, A w.p. 1/5
- Then in the normal form game E(G) we have
  - $N' = \{Marcie^M, Marcie^A, Charlie^M, Charlie^A\}$
  - $u_{Marcie^M}(T, F, T, F) = 1/3 \times 2 + 2/3 \times 0 = 2/3$
  - $u_{Marcie^M}(T, T, T, F) = 1/3 \times 2 + 2/3 \times 0 = 2/3$

Nash Equilibrium in Bayesian Games

- Definition: given a Bayesian game G, a strategy profile **s** (with one action for each type of each player in G) is said to be a Nash equilibrium of G if it is a Nash equilibrium of the respective normal-form game E(G)
  - no type of any player should want to change his action given the actions of all types of other players

Example: Battle of Sexes

- Marcie: "meet" w.p. 4/5, "avoid" w.p. 1/5
- Charlie: "meet" w.p. 3/4, "avoid" w.p. 1/4

- Charlie^M payoffs:

payoffs		best response:
u(T, T, T, *) = 1,	u(T, T, F, *) = 0	(T, T) → T
u(T, F, T, *) = 4/5,	u(T, F, F, *) = 2/5	(T, F) → T
u(F, T, T, *) = 1/5,	u(F, T, F, *) = 8/5	(F, T) → F
u(F, F, T, *) = 0,	u(F, F, F, *) = 2	(F, F) → F

- Charlie^A payoffs:

u(T, T, *, T) = 0,	u(T, T, *, F) = 2	(T, T) → F
u(T, F, *, T) = 1/5,	u(T, F, *, F) = 8/5	(T, F) → F
u(F, T, *, T) = 4/5,	u(F, T, *, F) = 2/5	(F, T) → T
u(F, F, *, T) = 1,	u(F, F, *, F) = 0	(F, F) → T

Example: Battle of Sexes

- Marcie: "meet" w.p. 4/5, "avoid" w.p. 1/5
- Charlie: "meet" w.p. 3/4, "avoid" w.p. 1/4

- Marcie^M payoffs:

		best response:
u(T, *, T, T) = 2,	u(F, *, T, T) = 0	(T, T) → T
u(T, *, T, F) = 3/2,	u(F, *, T, F) = 1/4	(T, F) → T
u(T, *, F, T) = 1/2,	u(F, *, F, T) = 3/4	(F, T) → F
u(T, *, F, F) = 0,	u(F, *, F, F) = 1	(F, F) → F

- Marcie^A payoffs:

u(*, T, T, T) = 0,	u(*, F, T, T) = 1	(T, T) → F
u(*, T, T, F) = 1/2,	u(*, F, T, F) = 3/4	(T, F) → F
u(*, T, F, T) = 3/2,	u(*, F, F, T) = 1/4	(F, T) → T
u(*, T, F, F) = 2,	u(*, F, F, F) = 0	(F, F) → T

Example: Battle of Sexes

- Best response:

Charlie^M :	Charlie^A :	Charlie:
(T, T) → T	(T, T) → F	(T, T) → (T, F)
(T, F) → T	(T, F) → F	(T, F) → (T, F)
(F, T) → F	(F, T) → T	(F, T) → (F, T)
(F, F) → F	(F, F) → T	(F, F) → (F, T)

Marcie^M :	Marcie^A :	Marcie:
(T, T) → T	(T, T) → F	(T, T) → (T, F)
(T, F) → T	(T, F) → F	(T, F) → (T, F)
(F, T) → F	(F, T) → T	(F, T) → (F, T)
(F, F) → F	(F, F) → T	(F, F) → (F, T)

Example: Battle of Sexes

- Marcie: "meet" w.p. 4/5, "avoid" w.p. 1/5
- Charlie: "meet" w.p. 3/4, "avoid" w.p. 1/4
- Best responses:

Charlie:	Marcie
(T, T) → (T, F)	(T, T) → (T, F)
(T, F) → (T, F)	(T, F) → (T, F)
(F, T) → (F, T)	(F, T) → (F, T)
(F, F) → (F, T)	(F, F) → (F, T)

- Nash equilibrium: (T, F, T, F), (F, T, F, T)

Illustration: First-Price Auctions With Incomplete Information

- First-price auction:
  - one object for sale, each bidder assigns some value to it
  - each bidder submits a bid
  - the bidder who submitted the highest bid wins the object and pays his bid
- Typically, bidders do not know each others' values; rather, they have beliefs about each others' values

First-Price Auction With Incomplete Information

- Alice and Bob bid for a painting
- Alice believes that Bob values the painting as \$100 w.p. 1/5, \$200 w.p. 4/5
- Bob believes that Alice values the painting as \$120 w.p. 2/5, \$150 w.p. 3/5
- Bayesian game:
  - types = values ({\$100, \$200} for Alice, {\$120, \$150} for Bob)
  - actions = bids (non-negative reals)
  - strategy: how much to bid for each type

Infinite Type Spaces

- So far, we considered games where each player has a finite number of types
- However, the number of types may be infinite:
  - Cournot oligopoly:
    the cost can be any real number
    in some interval $[c_1, c_2]$
  - first-price auction:
    Alice's value can be any real number
    between 100 and 200
- Warning: the associated normal-form game E(G) has infinitely many players, and we have not formally defined Nash equilibria for such games
  - however, the definition can be extended

Infinite Type Spaces: Strategies and Beliefs

- If a player's type space is a set T, and her action space is A, her strategy is a mapping $T \to A$
  - Battle of Sexes: {meet, avoid} $\to$ {T, F}
  - Cournot oligopoly with costs c_1, c_2: $\{c_1, c_2\} \to R$
  - Cournot oligopoly with costs in $[c_1, c_2]$: $[c_1, c_2] \to R$
- Players assign probabilities to other players' types: in a 2-player game
  - player 1 believes that player 2's type is drawn from T_2 according to a distribution F_2
  - player 2 believes that player 1's type is drawn from T_1 according to a distribution F_1

First Price Auction With Two Bidders

- First-price auction, 2 bidders
- each bidder's value is in [0, 1]
  - $T_1 = T_2 = [0, 1]$
- Each bidder knows his value and assumes that the other bidder's value is drawn from the uniform distribution on [0, 1]:
  - U[0, 1]: CDF $F(x) = x$, PDF $f(x) = 1$ for $x \in [0, 1]$
- <u>Proposition</u>: for each bidder, bidding half of his value is a NE strategy
  - i.e., assuming that bidder 2 bids $v_2/2$ (whatever v_2 is), bidder 1 maximizes his expected utility by bidding $v_1/2$, for every $v_1 \in [0, 1]$, and vice versa

Example:
First Price Auction With Two Bidders

- Proposition: for each bidder, bidding half of his value is a NE strategy
- Proof: suppose B1 has value v_1; let us compute his optimal bid b
  - suppose B2 bids $b_2 = v_2/2$
  - $b_2 \le 1/2$, so we can assume that $b \le 1/2$ as well
  - Pr $[b_2 \le x]$ = Pr$[v_2 \le 2x]$ = $2x$ for $x \le 1/2$
  - when B1 bids $b \le 1/2$, Pr [B1 wins] = Pr $[b_2 \le b]$ = $2b$
  - B1's expected utility = $2b(v_1 - b)$: maximized at $b = v_1/2$