CMSC 858F
Introduction to Game Theory*

Mohammad T. Hajiaghayi
University of Maryland

*Some of these slides are originally prepared by Professor Dana Nau.

What is Game Theory?

- Game Theory is about interactions among self-interested agents (players)
- Different agents have different preferences (i.e. like some outcomes more than others)
- Note that game theory is not a tool; it is a set of concepts.
- Goals of this course:
  - Formal definitions and technicality of the algorithms
  - Better understanding of real-world games

Algorithmic Game Theory

- Algorithm Game Theory is often viewed as “incentive-aware algorithm design’’
- Algorithm design often deals with dumb objects though Algorithmic Game Theory often deals with smart (self-interested) objects
- Combines Algorithm Design and Game Theory
- Also known as Mechanism Design
- Goal of Mechanism Design
  - Encourage selfish agents to act socially by designing rewarding rules such that when agents optimize their own objective, a social objective is met

Some Fields where Game Theory is Used

- Economics, business
  - Markets, auctions
  - Economic predictions
  - Bargaining, fair division

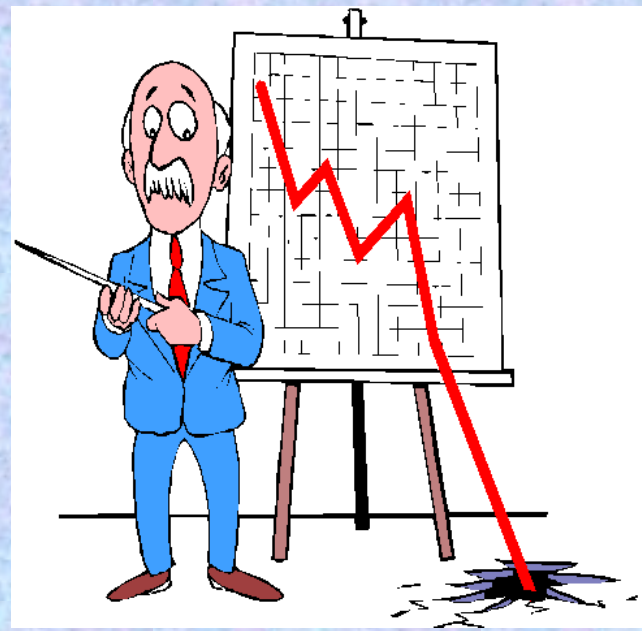

Some Fields where Game Theory is Used

- Government, politics, military
  - Negotiations
  - Voting systems
  - International relations
  - War
  - …

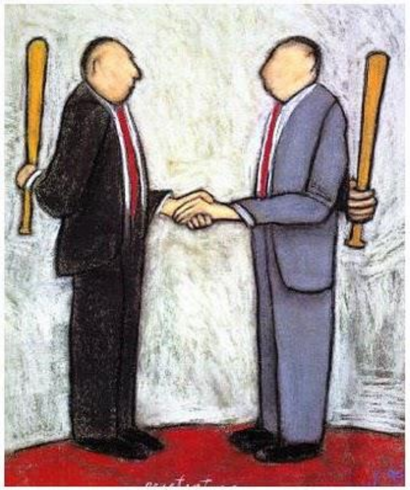

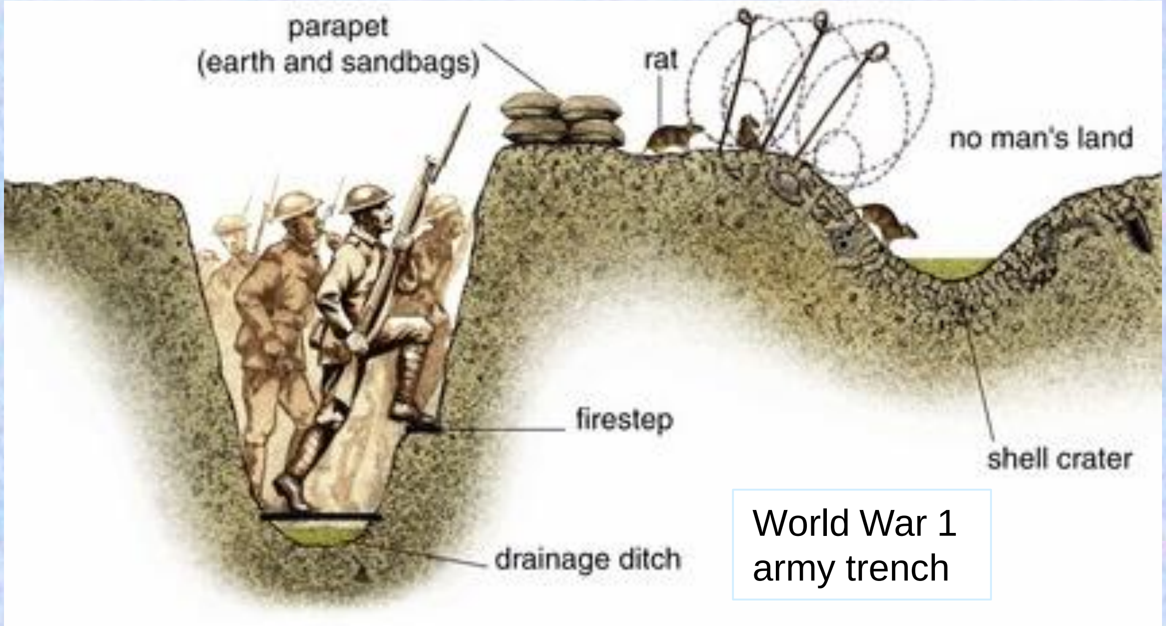

World War 1
army trench

Some Fields where Game Theory is Used

- Biology, psychology, sociology
  - Population ratios, territoriality
  - Social behavior
  - …

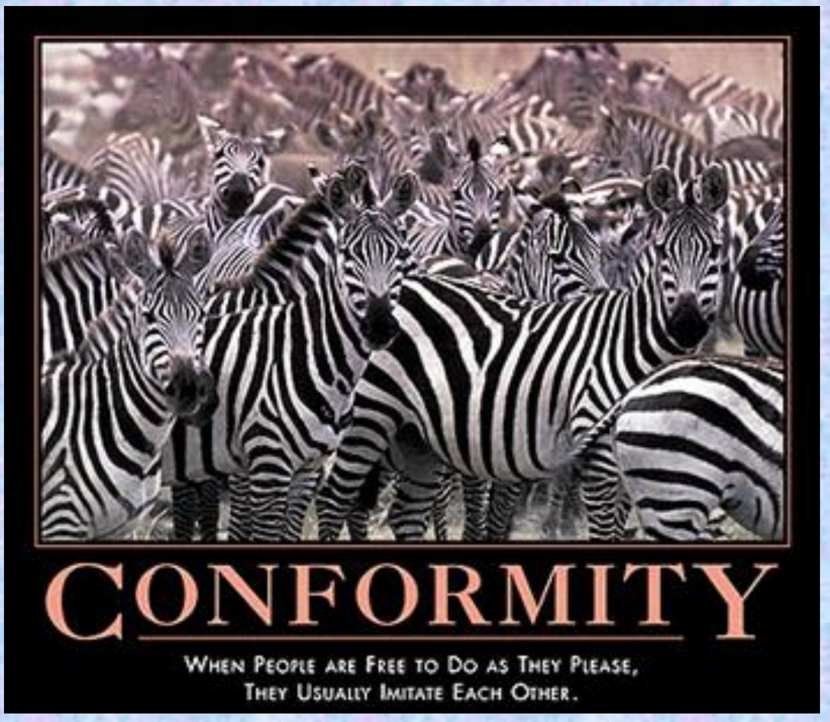

Some Fields where Game Theory is Used

- Engineering, computer science
  - Game programs
  - Computer and communication networks
  - Road networks
  - …

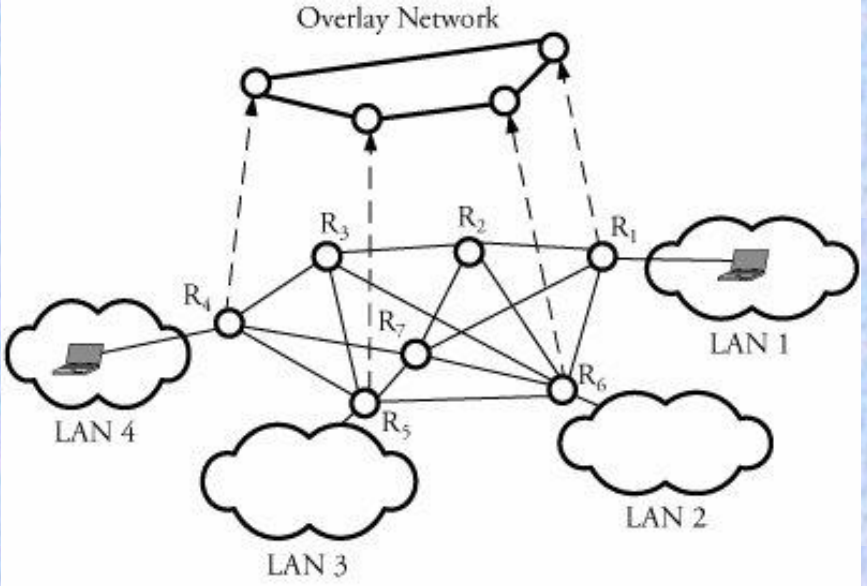

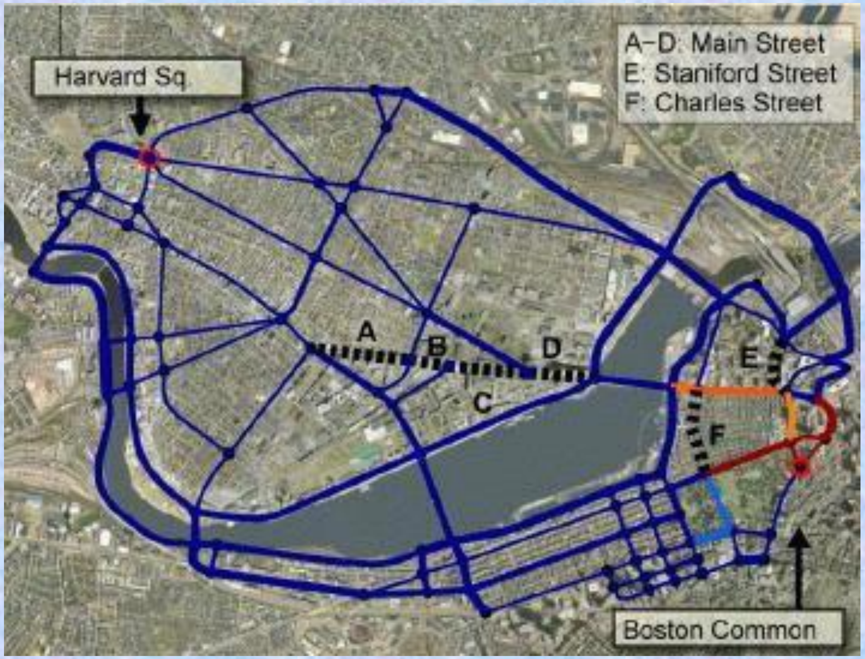

Games in Normal Form

- A (finite, n-person) **normal-form game** includes the following:
  1. An ordered set $N = (1, 2, 3, \ldots, n)$ of **agents** or **players**:
  2. Each agent i has a finite set A_i of possible actions
     - An **action profile** is an n-tuple $\mathbf{a} = (a_1, a_2, \ldots, a_n)$, where $a_1 \in A_1$, $a_2 \in A_2, \ \ldots, \ a_n \in A_n$
     - The set of all possible action profiles is $\mathbf{A} = A_1 \times \cdot \cdot \cdot \times A_n$
  3. Each agent i has a real-valued **utility** (or **payoff**) function

     $u_i\,(a_1, \ldots, a_n) = i$'s payoff if the action profile is $(a_1, \ldots, a_n)$
- Most other game representations can be reduced to normal form
- Usually represented by an n-dimensional **payoff** (or **utility**) **matrix**
  - for each action profile, shows the utilities of all the agents

	take 3	take 1
take 3	3, 3	0, 4
take 1	4, 0	1, 1

The Prisoner’s Dilemma


- Scenario: The police are holding two prisoners as suspects for committing a crime
  - For each prisoner, the police have enough evidence for a 1 year prison sentence
  - They want to get enough evidence for a 4 year prison sentence
  - They tell each prisoner,
    - “If you testify against the other prisoner, we’ll reduce your prison sentence by 1 year”
  - *C* = *Cooperate* (with the other prisoner): refuse to testify against him/her
  - *D* = *Defect*: testify against the other prisoner
  - Both prisoners cooperate => both go to prison for 1 year
  - Both prisoners defect => both go to prison for 4 – 1 = 3 years
  - One defects, other cooperates => cooperator goes to prison for 4 years; defector goes free

	C	*D*
C	–1, –1	–4, 0
D	0, –4	–3, –3

Prisoner's Dilemma

We used this:

	C	*D*
C	–1, –1	–4, 0
D	0, –4	–3, –3

Equivalent:

	take 3	take 1
take 3	3, 3	0, 4
take 1	4, 0	1, 1

Game theorists usually use this:

	C	*D*
C	3, 3	0, 5
D	5, 0	1, 1

	C	*D*
C	a, a	b, c
D	c, b	d, d

- General form:

  $c > a > d > b$

  $2a > b + c$

Utility Functions

- Idea: the preferences of a rational agent must obey some constraints
- Agent’s choices are based on rational preferences
  $\Rightarrow$ agent’s behavior is describable as maximization of expected utility
- Constraints:

  Orderability (sometimes called **Completeness**):

  $$(A \succ B) \lor (B \succ A) \lor (A \sim B)$$

  Transitivity:

  $$(A \succ B) \land (B \succ C) \Rightarrow (A \succ C)$$

- **Theorem** (Ramsey, 1931; von Neumann and Morgenstern, 1944).
- Given preferences satisfying the constraints above, there exists a real-valued function *u* such that

$$u(A) \geq u(B) \Leftrightarrow A \succeq B \qquad (*)$$

u is called a **utility function**

Utility Scales for Games

- Suppose that all the agents have rational preferences, and that this is common knowledge* to all of them
- Then games are insensitive to positive affine transformations of one or more agents' payoffs
  - Let c and d be constants, $c > 0$
  - For one or more agents i, replace every payoff x_{ij} with $cx_{ij} + d$
  - The game still models the same sets of rational preferences

	a_{21}	a_{22}
a_{11}	x_{11}, x_{21}	x_{12}, x_{22}
a_{12}	x_{13}, x_{23}	x_{14}, x_{24}

	a_{21}	a_{22}
a_{11}	$cx_{11}+d, x_{21}$	$cx_{12}+d, x_{22}$
a_{12}	$cx_{13}+d, x_{23}$	$cx_{14}+d, x_{24}$

	a_{21}	a_{22}
a_{11}	$cx_{11}+d, ex_{21}+f$	$cx_{12}+d, ex_{22}+f$
a_{12}	$cx_{13}+d, ex_{23}+f$	$cx_{14}+d, ex_{24}+f$

*Common knowledge is a complicated topic; I'll discuss it later

Common-payoff Games

- **Common-payoff game**:
  - For every action profile, all agents have the same payoff
- Also called a **pure coordination** game or a **team game**
  - Need to coordinate on an action that is maximally beneficial to all

- **Which side of the road?**
  - 2 people driving toward each other in a country with no traffic rules
  - Each driver independently decides whether to stay on the left or the right
  - Need to coordinate your action with the action of the other driver

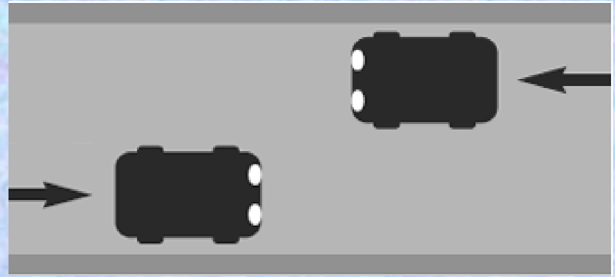

	Left	Right
Left	1, 1	0, 0
Right	0, 0	1, 1

A Brief Digression

- **Mechanism design**: set up the rules of the game, to give each agent an incentive to choose a desired outcome
- E.g., the law says what side of the road to drive on
  - Sweden on September 3, 1967:

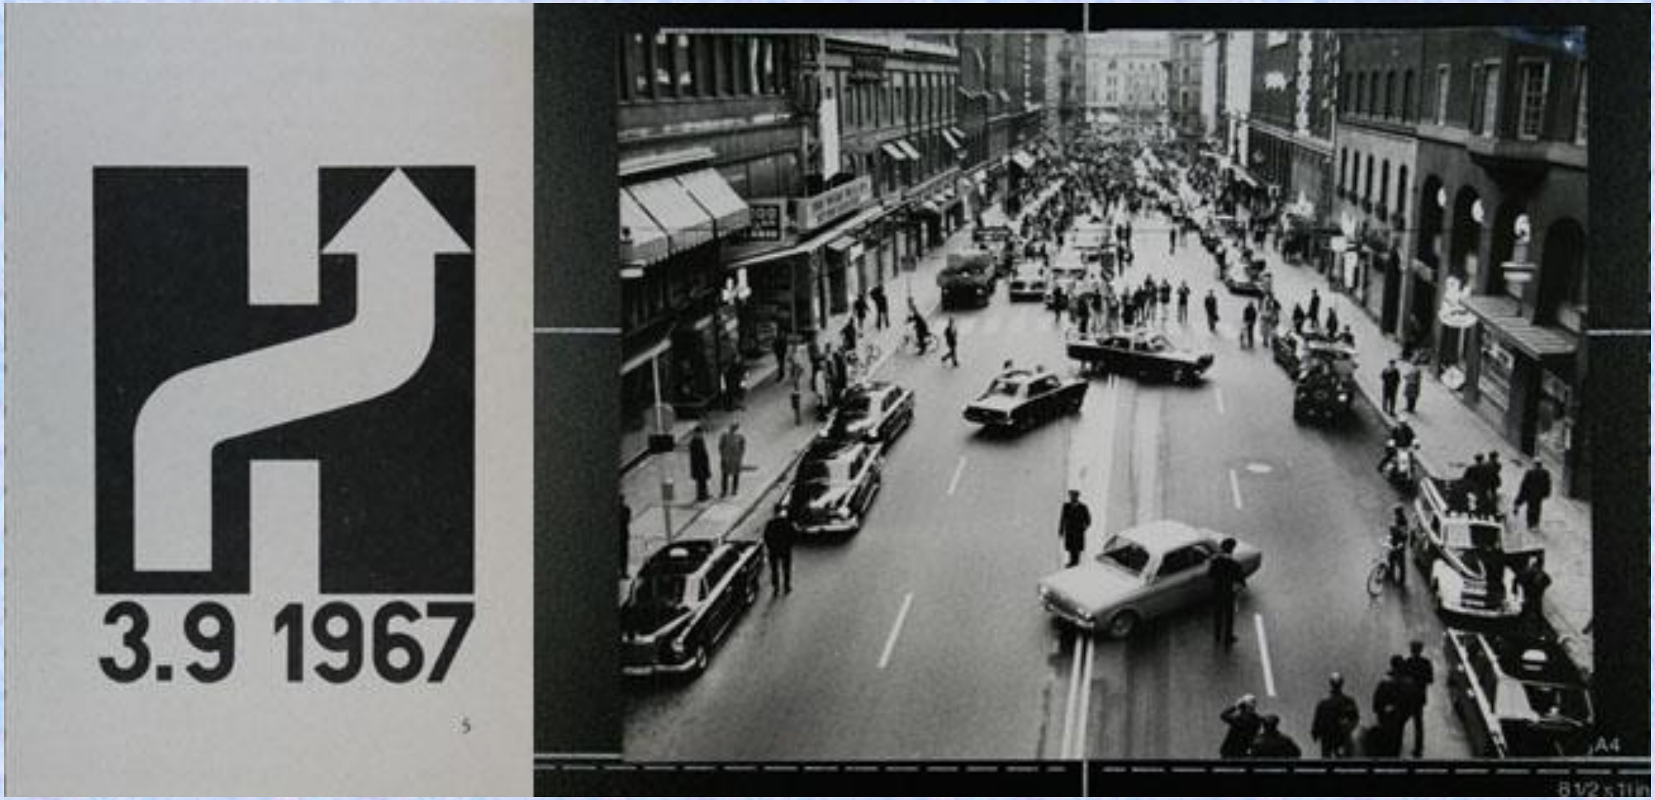

Zero-sum Games

- These games are purely competitive

- **Constant-sum** game:
  - For every action profile, the sum of the payoffs is the same, i.e.,
  - there is a constant c such for every action profile $\mathbf{a} = (a_1, \ldots, a_n)$,
    - $u_1(\mathbf{a}) + \ldots + u_n(\mathbf{a}) = c$

- Any constant-sum game can be transformed into an equivalent game in which the sum of the payoffs is always 0
  - Positive affine transformation: subtract c/n from every payoff

- Thus constant-sum games are usually called **zero-sum** games

Examples

- **Matching Pennies**
  - Two agents, each has a penny
  - Each independently chooses to display Heads or Tails
    - If same, agent 1 gets both pennies
    - Otherwise agent 2 gets both pennies

	Heads	Tails
Heads	1, –1	–1, 1
Tails	–1, 1	1, –1

- **Penalty kicks in soccer**
  - A kicker and a goalie
  - Kicker can kick left or right
  - Goalie can jump to left or right
  - Kicker scores if he/she kicks to one side and goalie jumps to the other

Another Example:Rock-Paper-Scissors

- Two players. Each simultaneously picks an action: *Rock*, *Paper*, or *Scissors*.
- The rewards:

Rock	beats	*Scissors*
Scissors	beats	*Paper*
Paper	beats	*Rock*

- The matrices:

$$R_1 = \begin{array}{c c} & \begin{array}{ccc} \mathsf{R} & \mathsf{P} & \mathsf{S} \end{array} \\ \begin{array}{c} \mathsf{R} \\ \mathsf{P} \\ \mathsf{S} \end{array} & \begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix} \end{array} \qquad R_2 = \begin{array}{c c} & \begin{array}{ccc} \mathsf{R} & \mathsf{P} & \mathsf{S} \end{array} \\ \begin{array}{c} \mathsf{R} \\ \mathsf{P} \\ \mathsf{S} \end{array} & \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix} \end{array}$$

Nonzero-Sum Games

- A game is **nonconstant-sum** (usually called **nonzero-sum**) if there are action profiles **a** and **b** such that
  - $u_1(\mathbf{a}) + \ldots + u_n(\mathbf{a}) \neq u_1(\mathbf{b}) + \ldots + u_n(\mathbf{b})$
  - e.g., the Prisoner's Dilemma

	C	*D*
C	3, 3	0, 5
D	5, 0	1, 1

- **Battle of the Sexes**
  - Two agents need to coordinate their actions, but they have different preferences
  - Original scenario:
    - husband prefers football, wife prefers opera
  - Another scenario:
    - Two nations must act together to deal with an international crisis, and they prefer different solutions

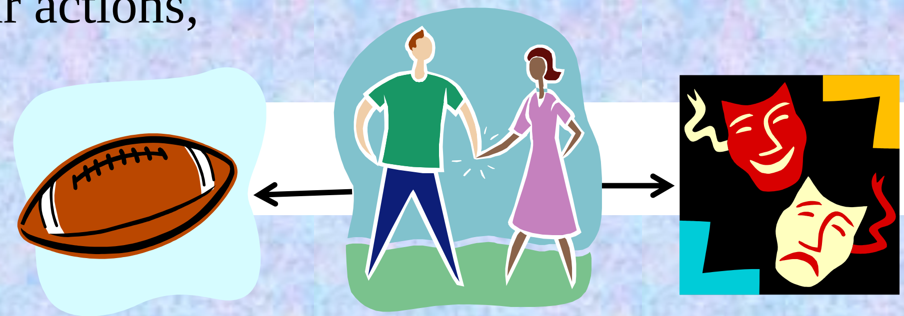

		Husband:	
		Opera	Football
Wife:	Opera	2, 1	0, 0
	Football	0, 0	1, 2

Symmetric Games

- In a **symmetric** game, every agent has the same actions and payoffs
  - If we change which agent is which, the payoff matrix will stay the same

Which side of the road?

	Left	Right
Left	1, 1	0, 0
Right	0, 0	1, 1

- For a 2x2 symmetric game, it doesn't matter whether agent 1 is the row player or the column player
  - The payoff matrix looks like this:

	a_1	a_2
a_1	w, w	x, y
a_2	y, x	z, z

- In the payoff matrix of a symmetric game, we only need to display u_1
  - If you want to know another agent's payoff, just interchange the agent with agent 1

	a_1	a_2
a_1	w	x
a_2	y	z

Strategies in Normal-Form Games

- **Pure strategy**: select a single action and play it
  - Each row or column of a payoff matrix represents both an action and a pure strategy
- **Mixed strategy**: randomize over the set of available actions according to some probability distribution
  - $s_i(a_j)$ = probability that action a_j will be played in mixed strategy s_i
- The **support** of s_i = {actions that have probability > 0 in s_i}
- A pure strategy is a special case of a mixed strategy
  - support consists of a single action
- A strategy s_i is **fully mixed** if its support is A_i
  - i.e., nonzero probability for every action available to agent *i*
- **Strategy profile**: an *n*-tuple $\mathbf{s} = (s_1, \ldots, s_n)$ of strategies, one for each agent

Expected Utility

- A payoff matrix only gives payoffs for pure-strategy profiles
- Generalization to mixed strategies uses expected utility
  - First calculate probability of each outcome, given the strategy profile (involves all agents)
  - Then calculate average payoff for agent *i*, weighted by the probabilities
  - Given strategy profile **s** = (s_1, …, s_n)
    - expected utility is the sum, over all action profiles, of the profile's utility times its probability:

$$u_i(\mathbf{s}) = \sum_{\mathbf{a} \in \mathbf{A}} u_i(\mathbf{a}) \Pr[\mathbf{a} \mid \mathbf{s}]$$

i.e.,

$$u_i(s_1,\ldots,s_n) = \sum_{(a_1,\ldots,a_n) \in \mathbf{A}} u_i(a_1,\ldots,a_n) \prod_{j=1}^{n} s_j(a_j)$$

Some Comments about Normal-Form Games

- Only two kinds of strategies in the normal-form game representation:
  - **Pure strategy**: just a single action
  - **Mixed strategy**: probability distribution over pure strategies
    - i.e., choose an action at random from the probability distribution
- The normal-form game representation may see very restricted
  - No such thing as a conditional strategy (e.g., cross the bay if the temperature is above 70)
  - No temperature or anything else to observe

	C	*D*
C	3, 3	0, 5
D	5, 0	1, 1

- However much more complicated games can be mapped into normal-form games
  - Each pure strategy is a description of what you'll do in *every* situation you might ever encounter in the game
- In later sessions, we see more examples

How to reason about games?

- In single-agent decision theory, look at an **optimal strategy**
  - Maximize the agent's expected payoff in its environment

- With multiple agents, the best strategy depends on others' choices
- Deal with this by identifying certain subsets of outcomes called **solution concepts**

- First we discuss two solution concepts:
  - Pareto optimality
  - Nash equilibrium
- Later we will discuss several others

Pareto Optimality

- A strategy profile **s** **Pareto dominates** a strategy profile **s**′ if
  - no agent gets a worse payoff with **s** than with **s**′,
    i.e., $u_i(\mathbf{s}) \geq u_i(\mathbf{s}')$ for all i ,
  - at least one agent gets a better payoff with **s** than with **s**′,
    i.e., $u_i(\mathbf{s}) > u_i(\mathbf{s}')$ for at least one i

- A strategy profile **s** is **Pareto optimal** (or **Pareto efficient**) if there's no strategy profile **s**′ that Pareto dominates **s**
  - Every game has at least one Pareto optimal profile
  - Always at least one Pareto optimal profile in which the strategies are pure

Examples

	C	*D*
C	3, 3	0, 5
D	5, 0	1, 1

The Prisoner's Dilemma

- (*D*,*C*) is Pareto optimal: no profile gives player 1 a higher payoff
- (*C*, *D*) is Pareto optimal: no profile gives player 2 a higher payoff
- (*C*,*C*) is Pareto optimal: no profile gives both players a higher payoff
- (*D*,*D*) isn't Pareto optimal: (*C*,*C*) Pareto dominates it

Which Side of the Road

- (Left,Left) and (Right,Right) are Pareto optimal
- In common-payoff games, all Pareto optimal strategy profiles have the same payoffs
  - If (Left,Left) had payoffs (2,2), then (Right,Right) wouldn't be Pareto optimal

	Left	Right
Left	1, 1	0, 0
Right	0, 0	1, 1

Best Response

- Suppose agent i knows how the others are going to play
  - Then i has an ordinary optimization problem: maximize expected utility
- We'll use $\mathbf{s}_{-i}$ to mean a strategy profile for all of the agents except i

  $$\mathbf{s}_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$$

- Let s_i be any strategy for agent i. Then

  $$(s_i, \mathbf{s}_{-i}) = (s_1, \ldots, s_{i-1}, s_i, s_{i+1}, \ldots, s_n)$$

- s_i is a **best response** to $\mathbf{s}_{-i}$ if for every strategy s_i' available to agent i,

  $$u_i(s_i, \mathbf{s}_{-i}) \geq u_i(s_i', \mathbf{s}_{-i})$$

- There is always at least one best response
- A best response s_i is **unique** if $u_i(s_i, \mathbf{s}_{-i}) > u_i(s_i', \mathbf{s}_{-i})$ for every $s_i' \neq s_i$

Best Response

- Given $\mathbf{s}_{-i}$, there are only two possibilities:
  (1) i has a pure strategy s_i that is a unique best response to $\mathbf{s}_{-i}$
  (2) i has infinitely many best responses to $\mathbf{s}_{-i}$

Proof. Suppose (1) is false. Then there are two possibilities:

- **Case 1**: s_i isn't unique, i.e., ≥ 2 strategies are best responses to $\mathbf{s}_{-i}$
  - Then they all must have the same expected utility
  - Otherwise, they aren't all "best"
  - Thus any mixture of them is also a best response
  - Thus **(2)** happens.
- **Case 2**: s_i isn't pure, i.e., it's a mixture of $k > 2$ actions
  - The actions correspond to pure strategies, so this reduces to Case 1
  - Thus **(2)** happens.
- **Theorem**: Always there exists a pure best response s_i to $\mathbf{s}_{-i}$

Proof. In both (1) and (2) above, there should be one pure best response.

Example

- Suppose we modify the Prisoner's Dilemma to give Agent 1 another possible action:
  - Suppose 2's strategy is to play action *C*
  - What are 1's best responses?
  - Suppose 2's strategy is to play action *D*
  - What are 1's best responses?

	C	*D*
C	3, 3	0, 5
D	5, 0	1, 1
E	3, 3	1, 3

Nash Equilibrium

- **Equilibrium:** it is simply a state of the world where economic forces are balanced and in the absence of external influence the equilibrium variables will not change.
  - More intuitively, a state in which no person involved in the game wants any change.
- **Famous economic equilibria**: Nash equilibrium, Correlated equilibrium, Market Clearance equilibrium
- $\mathbf{s} = (s_1, \ldots, s_n)$ is a **Nash equilibrium** if for every i, s_i is a best response to $\mathbf{s}_{-i}$
  - Every agent's strategy is a best response to the other agents' strategies
  - No agent can do better by *unilaterally* changing his strategy
- **Theorem (Nash, 1951)**: Every game with a finite number of agents and actions has at least one Nash equilibrium
- In Which Side of the Road, (Left,Left) and (Right,Right) are Nash equilibria
  In the Prisoner's Dilemma, (D,D) is a Nash equilibrium
  - Ironically, it's the only pure-strategy profile that *isn't* Pareto optimal

	Left	Right
Left	1, 1	0, 0
Right	0, 0	1, 1

	C	D
C	3, 3	0, 5
D	5, 0	1, 1

Strict Nash Equilibrium

- A Nash equilibrium $\mathbf{s} = (s_1, \ldots, s_n)$ is **strict** if for every i, s_i is the *only* best response to $\mathbf{s}_{-i}$
  - i.e., any agent who unilaterally changes strategy will do worse
- Recall that if a best response is unique, it must be pure
  - It follows that in a strict Nash equilibrium, all of the strategies are pure
- But if a Nash equilibrium is pure, it isn't necessarily strict

- Which of the following Nash equilibria are strict? Why?

	C	*D*
C	3, 3	0, 5
D	5, 0	$\boxed{1, 1}$

	C	*D*
C	3, 3	0, 4
D	4, 0	$\boxed{1, 1}$

	Left	Right
Left	$\boxed{1, 1}$	0, 0
Right	0, 0	$\boxed{1, 1}$

	Left	Right
Left	$\boxed{1, 1}$	0, 0
Right	0, 0	$\boxed{1, 1}$
Center	0, 0	$\boxed{1, ½}$

Weak Nash Equilibrium

- If a Nash equilibrium **s** isn't strict, then it is **weak**
  - At least one agent *i* has more than one best response to $\mathbf{s}_{-i}$
- If a Nash equilibrium includes a mixed strategy, then it is weak
  - If a mixture of *k* => 2 actions is a best response to $\mathbf{s}_{-i}$, then any other mixture of the actions is also a best response
- If a Nash equilibrium consists only of pure strategies, it might still be weak

	Left	Right
Left	1, 1	0, 0
Right	0, 0	1, 1
Center	0, 0	1, ½

- Weak Nash equilibria are less stable than strict Nash equilibria
  - If a Nash equilibrium is weak, then at least one agent has infinitely many best responses, and only one of them is in **s**

Finding Mixed-Strategy Nash Equilibria

- In general, it's tricky to compute mixed-strategy Nash equilibria
  - But easier if we can identify the support of the equilibrium strategies
- In 2x2 games, we can do this easily
- We especially use theorem below proved earlier

Theorem A: *Always there exists a pure best response s_i to $\mathbf{s}_{-i}$*

- **Corollary B:** If (s_1, s_2) is a pure Nash equilibrium only among pure strategies, it should be a Nash equilibrium among mixed strategies as well
- Now let (s_1, s_2) be a Nash equilibrium
- If both s_1, s_2 have supports of size one, it should be one of the cells of the normal-form matrix and we are done by Corollary B
- Thus assume at least one of s_1, s_2 has a support of size two.

Finding Mixed-Strategy Nash Equilibria

- Now if the support of one of s_1, s_2 , say s_1, is of size one, i.e., it is pure, then s_2 should be pure as well, unless both actions of player 2 have the same payoffs; in this case any mixed strategy of both actions can be Nash equilibrium.
- Thus in the rest we assume both supports have size two.
  - Thus to find s_1 assume agent 1 selects action a_1 with probability p and action a'_1 with probability 1-p.
  - Now since s_2 has a support of size two, its support must include both of agent 2's actions, and they must have the same expected utility
    - Otherwise agent 2's best response would be just one of them and its support has size one.
  - Hence find p such that $u_2(s_1, a_2) = u_2(s_1, a'_2)$, i.e., solve the equation to find p (and thus s_2)
  - Similarly, find s_2 such that $u_1(a_1, s_2) = u_1(a'_1, \mathrm{s}_2)$

Finding Mixed-Strategy Nash Equilibria

Husband / Wife	Oper a	Football
Opera	2, 1	0, 0
Football	0, 0	1, 2

Example: Battle of the Sexes

- We already saw pure Nash equilibria.
- If there's a mixed-strategy equilibrium,
  - both strategies must be mixtures of {Opera, Football}
  - each must be a best response to the other
- Suppose the husband's strategy is $s_h = \{(p, \text{Opera}), (1–p, \text{Football})\}$
- Expected utilities of the wife's actions:

$$u_w(\text{Opera}, s_h) = 2p; \qquad u_w(\text{Football}, s_h) = 1(1 - p)$$

- If the wife mixes the two actions, they must have the same expected utility
  - Otherwise the best response would be to *always* use the action whose expected utility is higher
  - Thus $2p = 1 – p$, so $p = 1/3$
- So the husband's mixed strategy is $s_h = \{(1/3, \text{Opera}), (2/3, \text{Football})\}$

Finding Mixed-Strategy Nash Equilibria

- Similarly, we can show the wife's mixed strategy is
  - s_w = {(2/3, Opera), (1/3, Football)}
- So the mixed-strategy Nash equilibrium is (s_w , s_h), where
  - s_w = {(2/3, Opera), (1/3, Football)}
  - s_h = {(1/3, Opera), (2/3, Football)}
- Questions:
  - Like all mixed-strategy Nash equilibria, (s_w , s_h) is weak
    - Both players have infinitely many other best-response strategies
    - What are they?
  - How do we know that (s_w , s_h) really is a Nash equilibrium?
    - Indeed the proof is by the way that we found Nash equilibria (s_w , s_h)

Wife \ Husband	Oper a	Football
Opera	2, 1	0, 0
Football	0, 0	1, 2

Finding Mixed-Strategy Nash Equilibria

- s_w = {(2/3, Opera), (1/3, Football)}
- s_h = {(1/3, Opera), (2/3, Football)}

- Wife's expected utility is
  - 2(2/9) + 1(2/9) + 0(5/9) = 2/3
- Husband's expected utility is also 2/3

2/3 • 1/3 = 2/9

2/3 • 2/3 = 4/9

Husband / Wife	Oper a	Football
Opera	2, 1	0, 0
Football	0, 0	1, 2

1/3 • 1/3 = 1/9

1/3 • 2/3 = 2/9

- It's "fair" in the sense that both players have the same expected payoff
- But it's Pareto-dominated by both of the pure-strategy equilibria
  - In each of them, one agent gets 1 and the other gets 2
- Can you think of a fair way of choosing actions that produces a higher expected utility?

Finding Mixed-Strategy Nash Equilibria

Matching Pennies

- Easy to see that in this game, no pure strategy could be part of a Nash equilibrium
  - ➢ For each combination of pure strategies, one of the agents can do better by changing his/her strategy

	Heads	Tails
Heads	1, –1	–1, 1
Tails	–1, 1	1, –1

- Thus there isn't a strict Nash equilibrium since it would be pure.

- But again there's a mixed-strategy equilibrium
  - ➢ Can be derived the same way as in the Battle of the Sexes
    - Result is (s,s), where s = {(½, Heads), (½, Tails)}

Another Interpretation of Mixed Strategies

- Suppose agent *i* has a deterministic method for picking a strategy, but it depends on factors that aren't part of the game itself
  - If *i* plays a game several times, *i* may pick different strategies
- If the other players don't know how *i* picks a strategy, they'll be uncertain what *i*'s strategy will be
  - Agent *i*'s mixed strategy is **everyone else's assessment** of how likely *i* is to play each pure strategy
- Example:
  - In a series of soccer penalty kicks, the kicker could kick left or right in a deterministic pattern that the goalie thinks is random

Complexity of Finding Nash Equilibria

- We've discussed how to find Nash equilibria in some special cases
  - Step 1: look for pure-strategy equilibria
    - Examine each cell of the matrix
    - If no cell in the same row is better for agent 1, and
      no cell in the same column is better for agent 2
      then the cell is a Nash equilibrium
  - Step 2: look for mixed-strategy equilibria
    - Write agent 2's strategy as $\{(q, b), (1–q, b')\}$;
      look for q such that a and a' have the same expected utility
    - Write agent 1's strategy as $\{(p, a), (1–p, a')\}$;
      look for p such that b and b' have the same expected utility

	b	b'
a	u_1, v_1	u_2, v_2
a'	u_3, v_3	u_4, v_4

2x2 games

- More generally for two-player games with any number of actions

  for each player, if we know support of each, we can find a mixed-Nash equilibrium in polynomial-time by solving linear equations (via linear program).
- What about the general case?

Complexity of Finding Nash Equilibria

- General case: *n* players, *m* actions per player, payoff matrix has m^n cells

  (not in the book)

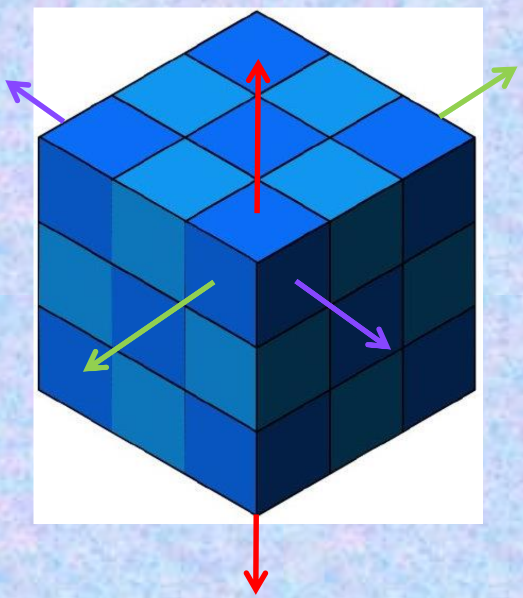

- Brute-force approach:
  - Step 1: Look for pure-strategy equilibria
    - At each cell of the matrix,
      - For each player, can that player do better by choosing a different action?
    - Polynomial time
  - Step 2: Look for mixed-strategy equilibria
    - For every possible combination of supports for $s_1, \ldots, s_n$
      - Solve sets of simultaneous equations
    - Exponentially many combinations of supports
    - Can it be done more quickly?

Complexity of Finding Nash Equilibria

- Two-player games
  - Lemke & Howson (1964): solve a set of simultaneous equations that includes all possible support sets for $s_1, \ldots, s_n$
    - Some of the equations are quadratic => worst-case exponential time
  - Porter, Nudelman, & Shoham (2004)
    - AI methods (constraint programming)
  - Sandholm, Gilpin, & Conitzer (2005)
    - Mixed Integer Programming (MIP) problem
- *n*-player games
  - van der Laan, Talma, & van der Heyden (1987)
  - Govindan, Wilson (2004)
  - Porter, Nudelman, & Shoham (2004)
- Worst-case running time still is exponential in the size of the payoff matrix

Complexity of Finding Nash Equilibria

- There are special cases that can be done in polynomial time in the size of the payoff matrix
  - Finding pure-strategy Nash equilibria
    - Check each square of the payoff matrix
  - Finding Nash equilibria in zero-sum games (see later in thi)
    - Linear programming
- For the general case,
  - It's unknown whether there are polynomial-time algorithms to do it
  - It's unknown whether there are polynomial-time algorithms to compute approximations
  - But we know both questions are PPAD-complete (but not NP-complete) even for two-player games (with some definition of PPAD

    introduced by Christos Papadimitriou in 1994)
- This is still one of the most important open problems in computational complexity theory

ε-Nash Equilibrium

- Reflects the idea that agents might not change strategies if the gain would be very small
- Let $\varepsilon > 0$. A strategy profile $\mathbf{s} = (s_1, . . . , s_n)$ is an **ε-Nash equilibrium** if for every agent i and for every strategy $s_i' \neq s_i$,

  $$u_i (s_i , \mathbf{s}_{-i}) \geq u_i (s_i', \mathbf{s}_{-i}) - \varepsilon$$

- ε-Nash equilibria exist for every $\varepsilon > 0$
  - Every Nash equilibrium is an ε-Nash equilibrium, and is surrounded by a region of ε-Nash equilibria
- This concept can be computationally useful
  - Algorithms to identify ε-Nash equilibria need consider only a finite set of mixed-strategy profiles (not the whole continuous space)
  - Because of finite precision, computers generally find only ε-Nash equilibria, where ε is roughly the machine precision
- Finding an ε-Nash equilibrium is still PPAD-complete (but not NP-complete) even for two-player games

The Price of Anarchy (PoA)

- In the Chocolate Game, recall that
  - (T3,T3) is the action profile that provides the best outcome for everyone
  - If we assume each payer acts to maximize his/her utility without regard to the other, we get (T1,T1)
  - By choosing (T3,T3), each player could have gotten 3 times as much
- Let’s generalize “best outcome for everyone”

	T3	*T1*
T3	3, 3	0, 4
T1	4, 0	1, 1

	T3	*T1*
T3	3, 3	0, 4
T1	4, 0	1, 1

The Price of Anarchy

- ***Social welfare function***: a function $w(\mathbf{s})$ that measures the players' welfare, given a strategy profile $\mathbf{s}$, e.g.,
  - Utilitarian function: $w(\mathbf{s})$ = average expected utility
  - Egalitarian function: $w(\mathbf{s})$ = minimum expected utility

- ***Social optimum***: benevolent dictator chooses $\mathbf{s}^*$ that optimizes w
  - $\mathbf{s}^* = \arg\max_{\mathbf{s}} w(\mathbf{s})$

- ***Anarchy***: no dictator; every player selfishly tries to optimize his/her own expected utility, disregarding the welfare of the other players
  - Get a strategy profile $\mathbf{s}$ (e.g., a Nash equilibrium)
  - In general, $w(\mathbf{s}) \leq w(\mathbf{s}^*)$

Price of Anarchy (PoA) $= \max_{\mathbf{s} \text{ is Nash equilibrium}} w(\mathbf{s}^*) / w(\mathbf{s})$

- PoA is the most popular measure of inefficiency of equilibria.
- We are generally interested in PoA which is closer to 1, i.e., all equilibria are good approximations of an optimal solution.

The Price of Anarchy

- Example: the Chocolate Game
  - Utilitarian welfare function:
    $w(\mathbf{s})$ = average expected utility

	T 3	*T1*
T3	3, 3	0, 4
T1	4, 0	1, 1

- Social optimum: $\mathbf{s}^* = (T3,T3)$
  - $w(\mathbf{s}^*) = 3$

	T3	*T1*
T3	3, 3	0, 4
T1	4, 0	1, 1

- Anarchy: $\mathbf{s} = (T1,T1)$
  - $w(\mathbf{s}) = 1$

- Price of anarchy

  $= w(\mathbf{s}^*) / w(\mathbf{s}) = 3/1 = 3$

- What would the answer be if we used the egalitarian welfare function?

The Price of Anarchy

- Sometimes instead of ***maximizing*** a welfare function w, we want to ***minimize*** a cost function c (e.g. in Prisoner's Dilemma)
  - Utilitarian function: $c(\mathbf{s})$ = avg. expected cost
  - Egalitarian function: $c(\mathbf{s})$ = max. expected cost

	C	*D*
C	3, 3	0, 5
D	5, 0	$\boxed{1, 1}$

- Need to adjust the definitions
  - ***Social optimum***: $\mathbf{s}^* = \arg\min_{\mathbf{s}} c(\mathbf{s})$
  - ***Anarchy***: every player selfishly tries to minimize his/her own cost, disregarding the costs of the other players
    - Get a strategy profile $\mathbf{s}$ (e.g., a Nash equilibrium)
    - In general, $c(\mathbf{s}) \geq c(\mathbf{s}^*)$
  - ***Price of Anarchy (PoA)*** = $\max_{\mathbf{s}\text{ is Nash equilibrium}} c(\mathbf{s}) / c(\mathbf{s}^*)$
    - i.e., the reciprocal of what we had before
    - E.g. in Prisoner's dilemma PoA= 3

Rationalizability

- A strategy is **rationalizable** if a *perfectly rational agent* could justifiably play it against *perfectly rational opponents*
  - The formal definition complicated
- Informally:
  - A strategy for agent *i* is rationalizable if it's a best response to strategies that *i* could *reasonably* believe the other agents have
  - To be reasonable, *i*'s beliefs must take into account
    - the other agents' knowledge of *i*'s rationality,
    - their knowledge of *i*'s knowledge of *their* rationality,
    - and so on so forth recursively
- A **rationalizable strategy profile** is a strategy profile that consists only of rationalizable strategies

Rationalizability

	Left	Right
Left	1, 1	0, 0
Right	0, 0	1, 1

- Every Nash equilibrium is composed of rationalizable strategies
- Thus the set of rationalizable strategies (and strategy profiles) is always nonempty

Example: Which Side of the Road

- For Agent 1, the pure strategy s_1 = *Left* is rationalizable because
  - s_1 = *Left* is 1's best response if 2 uses s_2 = *Left*,
  - and 1 can reasonably believe 2 would rationally use s_2 = *Left*, because
    - s_2 = *Left* is 2's best response if 1 uses s_1 = *Left*,
    - and 2 can reasonably believe 1 would rationally use s_1 = *Left*, because
      - s_1 = *Left* is 1's best response if 2 uses s_2 = *Left*,
      - and 1 can reasonably believe 2 would rationally use s_2 = *Left*, because
        - … and so on so forth…

Rationalizability

	Heads	Tails
Heads	1, –1	–1, 1
Tails	–1, 1	1, –1

- Some rationalizable strategies are not part of any Nash equilibrium

Example: Matching Pennies

- For Agent 1, the pure strategy s_1 = *Heads* is rationalizable because
  - s_1 = *Heads* is 1's best response if 2 uses s_2 = *Heads*,
  - and 1 can reasonably believe 2 would rationally use s_2 = *Heads*, because
    - s_2 = *Heads* is 2's best response if 1 uses s_1 = *Tails*,
    - and 2 can reasonably believe 1 would rationally use s_1 = *Tails*, because
      - s_1 = *Tails* is 1's best response if 2 uses s_2 = *Tails*,
      - and 1 can reasonably believe 2 would rationally use s_2 = *Tails*, because
        - … and so on so forth…

Common Knowledge

- The definition of common knowledge is recursive analogous to the definition of rationalizability
- A property p is *common knowledge* if
  - Everyone knows p
  - Everyone knows that everyone knows p
  - Everyone knows that everyone knows that everyone knows p
  - …

We Aren't Rational

- More evidence that we aren't game-theoretically rational agents

- Why choose an "irrational" strategy?
  - Several possible reasons …

Reasons for Choosing “Irrational” Strategies

(1) Limitations in reasoning ability

- Didn’t calculate the Nash equilibrium correctly
- Don’t know how to calculate it
- Don’t even know the concept

(2) Wrong payoff matrix - doesn’t encode agent’s actual preferences

- It’s a common error to take an external measure (money, points, etc.) and assume it’s all that an agent cares about
- Other things may be more important than winning
  - Being helpful
  - Curiosity
  - Creating mischief
  - Venting frustration

(3) Beliefs about the other agents’ likely actions (next slide)

Beliefs about Other Agents' Actions

- A Nash equilibrium strategy is best for you if the other agents also use their Nash equilibrium strategies
- In many cases, the other agents won't use Nash equilibrium strategies
  - If you can guess what actions they'll choose, then
    - You can compute your best response to those actions
      - maximize your expected payoff, given their actions
    - Good guess => you may do much better than the Nash equilibrium
    - Bad guess => you may do much worse

Worst-Case Expected Utility

- For agent i, the **worst-case** expected utility of a strategy s_i is the minimum over all possible combinations of strategies for the other agents:

$$\min_{\mathbf{s}_{-i}} \; u_i\left(s_i, \mathbf{s}_{-i}\right)$$

Wife \ Husband	Opera	Football
Opera	2, 1	0, 0
Football	0, 0	1, 2

- **Example: Battle of the Sexes**
  - Wife's strategy $s_w = \{(p, \text{Opera}), (1 - p, \text{Football})\}$
  - Husband's strategy $s_h = \{(q, \text{Opera}), (1 - q, \text{Football})\}$
  - $u_w(p,q) = 2pq + (1 - p)(1 - q) = 3pq - p - q + 1$
  - For any fixed p, $u_w(p,q)$ is linear in q
    - e.g., if $p = ½$, then $u_w(½,q) = ½\, q + ½$
  - $0 \le q \le 1$, so the min must be at $q = 0$ or $q = 1$
    - e.g., $\min_q$ $(½\, q + ½)$ is at $q = 0$
  - $\min_q u_w(p,q) = \min\,(u_w(p,0), u_w(p,1)) = \min\,(1 - p, 2p)$

We can write $u_w(p,q)$ instead of $u_w(s_w, s_h)$

Maxmin Strategies

Also called **maximin**

- A **maxmin strategy** for agent *i*
  - A strategy s_1 that makes *i*'s worst-case expected utility as high as possible:

$$\arg\max_{s_i} \min_{\mathbf{s}_{-i}} u_i\left(s_i, \mathbf{s}_{-i}\right)$$

  - This isn't necessarily unique
  - Often it is mixed

- Agent *i*'s **maxmin value**, or **security level**, is the maxmin strategy's worst-case expected utility:

$$\max_{s_i} \min_{\mathbf{s}_{-i}} u_i\left(s_i, \mathbf{s}_{-i}\right)$$

- For 2 players it simplifies to

$$\max_{s_1} \min_{s_2} u_1\left(s_1, s_2\right)$$

Example

- Wife's and husband's strategies
  - $s_w = \{(p, \text{Opera}), (1 - p, \text{Football})\}$
  - $s_h = \{(q, \text{Opera}), (1 - q, \text{Football})\}$
- Recall that wife's worst-case expected utility is

$$\min_q u_w(p,q) = \min(1 - p, 2p)$$

  - Find p that maximizes it
- Max is at $1 - p = 2p$, i.e., $p = 1/3$
  - Wife's maxmin value is $1 - p = 2/3$
  - Wife's maxmin strategy is $\{(1/3, \text{Opera}), (2/3, \text{Football})\}$
- Similarly,
  - Husband's maxmin value is 2/3
  - Husband's maxmin strategy is $\{(2/3, \text{Opera}), (1/3, \text{Football})\}$

Wife \ Husband	Opera	Football
Opera	2, 1	0, 0
Football	0, 0	1, 2

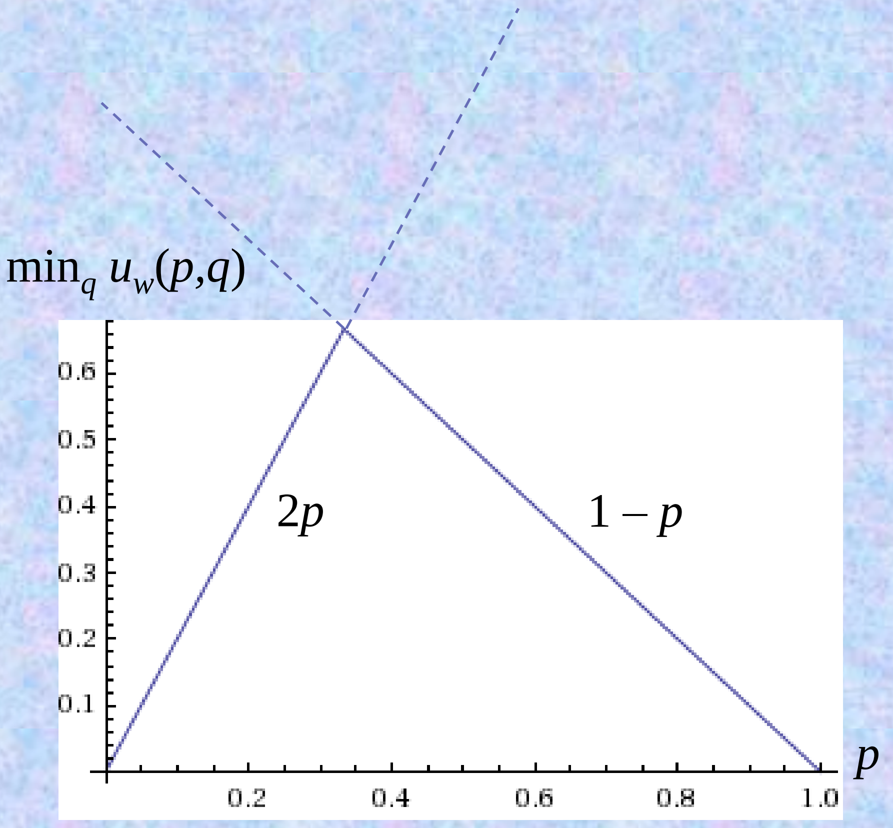

Minmax Strategies (in 2-Player Games)

- **Minmax strategy** and **minmax value**
  - Duals of their maxmin counterparts
- Suppose agent 1 wants to punish agent 2, regardless of how it affects agent 1's own payoff
- Agent 1's **minmax strategy** against agent 2
  - A strategy s_1 that minimizes the expected utility of 2's best response to s_1

$$\arg\min_{s_1}\max_{s_2} u_2(s_1, s_2)$$

- **Agent 2's minmax value** is 2's maximum expected utility if agent 1 plays his/her minmax strategy:

$$\min_{s_1}\max_{s_2} u_2(s_1, s_2)$$

- **Minmax strategy profile**: both players use their minmax strategies

Also called **minimax**

Example

Wife \ Husband	Opera	Football
Opera	2, 1	0, 0
Football	0, 0	1, 2

- Wife's and husband's strategies
  - $s_w = \{(p, \text{Opera}), (1 - p, \text{Football})\}$
  - $s_h = \{(q, \text{Opera}), (1 - q, \text{Football})\}$
- $u_h(p,q) = pq + 2(1 - p)(1 - q) = 3pq - 2p - 2q + 2$
- Given wife's strategy *p*, husband's expected utility is linear in *q*
  - e.g., if $p = ½$, then $u_h(½,q) = -½\, q + 1$
- Max is at $q = 0$ or $q = 1$

$$\max_q u_h(p,q) = (2-2p, p)$$

- Find *p* that minimizes this
- Min is at $-2p + 2 = p \rightarrow p = 2/3$
- Husband/s minmax value is 2/3
- Wife's minmax strategy is $\{(2/3, \text{Opera}), (1/3, \text{Football})\}$

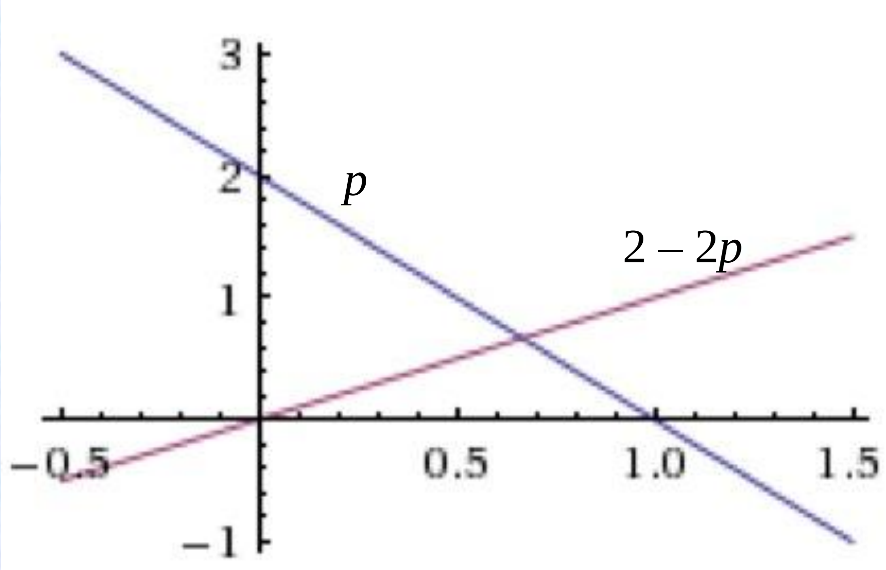

Minmax Strategies in *n*-Agent Games

- In n-agent games ($n > 2$), agent i usually can't minimize agent j's payoff by acting unilaterally
- But suppose all the agents "gang up" on agent j
  - Let $\mathbf{s}^*_{-j}$ be a mixed-strategy profile that minimizes j's maximum payoff, i.e.,

$$\mathbf{s}^*_{-j} = \arg\min_{\mathbf{s}_{-j}} \left(\max_{s_j} u_j\left(s_j, \mathbf{s}_{-j}\right) \right)$$

  - For every agent $i \neq j$, a **minmax strategy for *i*** is i's component of $\mathbf{s}_{-j}^*$

- **Agent *j*'s minmax value** is j's maximum payoff against $\mathbf{s}_{-j}^*$

$$\max_{s_j} u_j\left(s_j, \mathbf{s}^*_{-j}\right) = \min_{\mathbf{s}_{-j}} \max_{s_j} u_j\left(s_j, \mathbf{s}_{-j}\right)$$

- We have equality since we just replaced $\mathbf{s}_{-j}^*$ by its value above

Minimax Theorem (von Neumann, 1928)

- **Theorem.** Let G be any finite two-player zero-sum game. For each player i,
  - i's expected utility in any Nash equilibrium
    = i's maxmin value
    = i's minmax value
  - In other words, for every Nash equilibrium (s_1*, s_2*),

$$u_1(s_1^*, s_2^*) = \min_{s_1} \max_{s_2} u_1(s_1, s_2) = \max_{s_1} \min_{s_2} u_1(s_1, s_2) = -u_2(s_1^*, s_2^*)$$

  - Note that since $-u_{2=} u_1$ the third term does not mention u_2

- **Corollary.** For two-player zer-sum games:{Nash equilibria} = {maxmin strategy profiles}= {minmax strategy profiles}
- Note that this is **not necessary true** for **non-zero-sum** games as we saw for Battle of Sexes in previous slides
- Terminology: the **value** (or **minmax value**) of G is agent 1's minmax value

Proof of Minimax Theorem

- Let-u_2=u_1 = u and let mixed strategies $s_1 = \mathrm{x} = (\mathrm{x}_1, \dots, \mathrm{x}_\mathrm{k})$ and $\mathrm{s}_2 = \mathrm{y} = (\mathrm{y}_1, \dots, \mathrm{y}_\mathrm{r})$.
- Then $\mathrm{u}(x, y) = \sum_i \sum_j x_i y_j u_{i,j}$=$\sum_j y_j \sum_i x_i u_{i,j}$
- We want to find x^* which optimizes $v^1 = \max_x \min_y u(\mathrm{x,y})$
- Since player 2 is doing his best response (in $\min_y u(\mathrm{x,y})$) he sets $y_j > 0$ only if $\sum_i x_i u_{i,j}$ is minimized.

Thus v^1= $\sum_j \sum_i x_i y_j u_{i,j}$ = $(\sum_j y_j)\min_j \sum_i x_i u_{i,j} = \min_\mathrm{j} \sum_i x_i u_{i,j} \leq \sum_i x_i u_{i,j}$

for any j

Thus we have the following LP1 to find x^*

$$\max v^1$$

$$\text{such that } v^1 \leq \sum_i x_i u_{i,j} \text{ for all j}$$

$$\sum_i x_i = 1$$

$$x_i \geq 0$$

Proof of Minimax Theorem (continued)

- Similarly for $v^2 = \min_y \max_x u$(x,y) we have LP2

$$\min v^2$$

$$\text{such that } v^2 \geq \textstyle\sum_j y_j u_{i,j} \text{ for all i}$$

$$\textstyle\sum_j y_j = 1$$

$$y_j \geq 0$$

- But LP1 And LP2 are duals of each other and by the (strong) duality theorem $v^1 = v^2$
- Also note that if (x,y) is a Nash equilibrium, x should satisfy LP1 (since we used only the fact that y is a best response to x in the proof) and y should satisfy LP2 (since we used only the fact that x is a best response to y in the proof) and thus $u_1(x, y) = v^1 = v^2$

Dominant Strategies

- Let s_i and s_i' be two strategies for agent i
  - Intuitively, s_i dominates s_i' if agent i does better with s_i than with s_i' for *every* strategy profile $\mathbf{s}_{-i}$ of the remaining agents
- Mathematically, there are three gradations of dominance:
  - s_i **strictly dominates** s_i' if for every $\mathbf{s}_{-i}$,

    $$u_i(s_i, \mathbf{s}_{-i}) > u_i(s_i', \mathbf{s}_{-i})$$

  - s_i **weakly dominates** s_i' if for every $\mathbf{s}_{-i}$,

    $$u_i(s_i, \mathbf{s}_{-i}) \geq u_i(s_i', \mathbf{s}_{-i})$$

    and for at least one $\mathbf{s}_{-i}$,

    $$u_i(s_i, \mathbf{s}_{-i}) > u_i(s_i', \mathbf{s}_{-i})$$

  - s_i **very weakly dominates** s_i' if for every $\mathbf{s}_{-i}$,

    $$u_i(s_i, \mathbf{s}_{-i}) \geq u_i(s_i', \mathbf{s}_{-i})$$

Dominant Strategy Equilibria

- A strategy is **strictly** (resp., **weakly**, **very weakly**) **dominant** for an agent if it strictly (weakly, very weakly) dominates any other strategy for that agent
- A strategy profile $(s_1, . . . , s_n)$ in which every s_i is dominant for agent i (strictly, weakly, or very weakly) is a Nash equilibrium
  - Why?
  - Such a strategy profile forms an **equilibrium in strictly** (**weakly**, **very weakly**) **dominant strategies**

Examples

- Example: **the Prisoner's Dilemma**
  - http://www.youtube.com/watch?v=ED9gaAb2BEw

	C	*D*
C	3, 3	0, 5
D	**5**, 0	**1**, 1

- For agent 1, *D* is strictly dominant
  - If agent 2 uses *C, then*
    - Agent 1's payoff is higher with *D* than with *C*
  - If agent 2 uses *D, then*
    - Agent 1's payoff is higher with *D* than with *C*
- Similarly, *D* is strictly dominant for agent 2
- So (*D*,*D*) is a Nash equilibrium in strictly dominant strategies

	C	*D*
C	3, 3	0, 5
D	5, 0	1, 1

- How do strictly dominant strategies relate to strict Nash equilibria?

Example: Matching Pennies

- **Matching Pennies**
  - If agent 2 uses Heads, then
    - For agent 1, Heads is better than Tails
  - If agent 2 uses Tails, then
    - For agent 1, Tails is better than Heads
  - Agent 1 doesn't have a dominant strategy

    => no Nash equilibrium in dominant strategies

	Heads	Tails
Heads	**1**, –1	–1, 1
Tails	–1, 1	**1**, –1

- **Which Side of the Road**
  - Same kind of argument as above
  - No Nash equilibrium in dominant strategies

	Left	Right
Left	**1**, 1	0, 0
Right	0, 0	**1**, 1

Elimination of Strictly Dominated Strategies

- A strategy s_i is **strictly** (**weakly**, **very weakly**) **dominated** for an agent i if some other strategy s_i' strictly (weakly, very weakly) dominates s_i

	L	R
~~U~~	~~3, 3~~	~~0, 5~~
D	5, 1	1, 0

- A strictly dominated strategy can't be a best response to any move, so we can eliminate it (remove it from the payoff matrix)
  - This gives a **reduced** game
  - Other strategies may now be strictly dominated, even if they weren't dominated before

	L	~~R~~
D	5, 1	~~1, 0~~

- **IESDS** (Iterated Elimination of Strictly Dominated Strategies):
  - Do elimination repeatedly until no more eliminations are possible
  - When no more eliminations are possible, we have the **maximal reduction** of the original game

	L
D	5, 1

IESDS

- If you eliminate a strictly dominated strategy, the reduced game has the same Nash equilibria as the original one
- Thus

  {Nash equilibria of the original game}

  = {Nash equilibria of the maximally reduced game}

	L	*R*
~~*U*~~	~~3, 3~~	~~0, 5~~
D	5, 1	1, 0

- Use this technique to simplify finding Nash equilibria
  - Look for Nash equilibria on the maximally reduced game

	L	~~*R*~~
D	5, 1	~~1, 0~~

- In the example, we ended up with a single cell
  - The single cell *must* be a unique Nash equilibrium in all three of the games

	L
D	5, 1

IESDS

- Even if s_i isn't strictly dominated by a pure strategy, it may be strictly dominated by a mixed strategy
- **Example**: the three games shown at right
  - 1st game:
    - R is strictly dominated by L (and by C)
    - Eliminate it, get 2nd game
  - 2nd game:
    - Neither *U* nor *D* dominates *M*
    - But {(½, *U*), (½, *D*)} strictly dominates *M*
      - This wasn't true before we removed *R*
    - Eliminate it, get 3rd game
  - 3rd game is maximally reduced

	L	*C*	*R*
U	3, 1	0, 1	0, 0
M	1, 1	1, 1	5, 0
D	0, 1	4, 1	0, 0

	L	*C*
U	3, 1	0, 1
M	1, 1	1, 1
D	0, 1	4, 1

	L	*C*
U	3, 1	0, 1
D	0, 1	4, 1

Correlated Equilibrium: Pithy Quote

If there is intelligent life on other planets, in a majority of them, they would have discovered correlated equilibrium before Nash equilibrium.

----Roger Myerson

Correlated Equilibrium: Intuition

- Not every correlated equilibrium is a Nash equilibrium but every Nash equilibrium is a correlated equilibrium
- We have a **traffic light**: a fair randomizing device that tells one of the agents to go and the other to wait.
- Benefits:
  - easier to compute than Nash, e.g., it is polynomial-time computable
  - fairness is achieved
  - the sum of social welfare exceeds that of any Nash equilibrium

Correlated Equilibrium

Wife \ Husband	Oper a	Football
Opera	2, 1	0, 0
Football	0, 0	1, 2

- Recall the mixed-strategy equilibrium for the Battle of the Sexes
  - s_w = {(2/3, Opera), (1/3, Football)}
  - s_h = {(1/3, Opera), (2/3, Football)}
- This is "fair": each agent is equally likely to get his/her preferred activity
- But 5/9 of the time, they'll choose different activities => utility 0 for both
  - Thus each agent's expected utility is only 2/3
  - We've required them to make their choices independently
- Coordinate their choices (e.g., flip a coin) => eliminate cases where they choose different activities
  - Each agent's payoff will always be 1 or 2; expected utility 1.5
- Solution concept: **correlated** equilibrium
  - Generalization of a Nash equilibrium

Correlated Equilibrium Definition

- Let G be an *2*-agent game (for now).
- Recall that in a (mixed) Nash Equilibrium at the end we compute a probability matrix (also known as joint probability distribution) $P = [p_{i,j}]$ where $\Sigma_{i,j} p_{i,j} = 1$ and in addition $p_{i,j} = q_i . q'_j$ where $\Sigma_i q_i = 1$ and $\Sigma_j q'_j = 1$ (here q and q' are the mixed strategies of the first agent and the second agent).
- Now if we remove the constraint $p_{i,j} = q_i . q'_j$ (and thus $\Sigma_i q_i = 1$ and $\Sigma_j q'_j = 1$) but still keep all other properties of Nash Equilibrium then we have a ***Correlated Equilibrium***.
- Surely it is clear that by this definition of Correlated Equilibrium, every Nash Equilibrium is a Correlated Equilibrium as well but note vice versa.
- Even for a more general n-player game, we can compute a Correlated Equilibrium in polynomial time by a linear program (as we see in the next slide).
- Indeed the constraint $p_{i,j} = q_i . q'_j$ is the one that makes computing Nash Equilibrium harder.

Computing CE

$$\sum_{a \in A | a_i \in a} p(a) u_i(a) \geq \sum_{a \in A | a'_i \in a} p(a) u_i(a'_i, a_{-i}) \quad \forall i \in N, \forall a_i, a'_i \in A_i$$

$$p(a) \geq 0 \qquad \forall a \in A$$

$$\sum_{a \in A} p(a) = 1$$

- variables: $p(a)$; constants: $u_i(a)$
- we could find the social-welfare maximizing CE by adding an objective function

$$\text{maximize:} \quad \sum_{a \in A} p(a) \sum_{i \in N} u_i(a).$$

Motivation of Correlated Equilibrium

- Let G be an n-agent game
- Let “Nature”(e.g., a ***traffic light***) choose action profile $\mathbf{a} = (a_1, \ldots, a_n)$ randomly according to our computed joint probability distribution (Correlated Equilibirum) p.
- Then “Nature” tells each agent i the value of a_i (privately)
  - An agent can condition his/her action based on (private) value a_i
- However by the definition of best response in Nash Equilibrium (which also exists in Correlated Equilibrium), agent i will not deviate from suggested action a_i
  - Note that here we implicitly assume because other agents are rational as well, they choose the suggested actions by the “Nature” which are given to them privately.
- Since there is no randomization in the actions, the correlated equilibrium might seem more natural.

Auctions

- An auction is a way (other than bargaining) to sell a fixed supply of a ***commodity*** (an item to be sold) for which there is no well-established ongoing market
- ***Bidders*** make ***bids***
  - proposals to pay various amounts of money for the commodity
- Often the commodity is sold to the bidder who makes the largest bid
- Example applications
  - Real estate, art, oil leases, electromagnetic spectrum, electricity, eBay, google ads
- **Private-value auctions**
  - Each bidder may have a different ***bidder value or bidder valuation (BV)***, i.e., how much the commodity is worth to that bidder
  - A bidder's BV is his/her private information, not known to others
  - E.g., flowers, art, antiques

Types of Auctions

- Classification according to the rules for bidding
    - English
    - Dutch
    - First price sealed bid
    - Vickrey
    - many others
  - On the following pages, I'll describe several of these and will analyze their equilibria

- A possible problem is ***collusion*** (secret agreements for fraudulent purposes)
  - Groups of bidders who won't bid against each other, to keep the price low
  - Bidders who place phony (phantom) bids to raise the price (hence the auctioneer's profit)
- If there's collusion, the equilibrium analysis is no longer valid

English Auction

- The name comes from oral auctions in English-speaking countries, but I think this kind of auction was also used in ancient Rome
- Commodities:
  - antiques, artworks, cattle, horses, wholesale fruits and vegetables, old books, etc.
- Typical rules:
  - Auctioneer solicits an opening bid from the group
  - Anyone who wants to bid should call out a new price at least c higher than the previous high bid (e.g., c = 1 dollar)
  - The bidding continues until all bidders but one have dropped out
  - The highest bidder gets the object being sold, for a price equal to his/her final bid
- For each bidder i, let
  - v_i = i's valuation of the commodity (private information)
  - B_i = i's final bid
- If i wins, then i's profit is $\pi_i = v_i - B_i$ and everyone else's profit = 0

English Auction (continued)

- Nash equilibrium:
  - Each bidder i participates until the bidding reaches v_i , then drops out
  - The highest bidder, i, gets the object, at price $B_i < v_i$, so $\pi_i = B_i - v_i > 0$
    - B_i is close to the second highest bidder's valuation
  - For every bidder $j \neq i$, $\pi_j = 0$
- Why is this an equilibrium?
- Suppose bidder j deviates and none of the other bidders deviate
  - If j deviates by dropping out earlier,
    - Then j's profit will be 0, no better than before
  - If u deviates by bidding $B_i > v_j$, then
    - j win's the auction but j's profit is $v_j - B_j < 0$, worse than before

English Auction (continued)

- If there is a large range of bidder valuations, then the difference between the highest and 2nd-highest valuations may be large
  - Thus if there's wide disagreement about the item's value, the winner might be able to get it for much less than his/her valuation
- Let *n* be the number of bidders
  - The higher *n* is, the more likely it is that the highest and 2nd-highest valuations are close
    - Thus, the more likely it is that the winner pays close to his/her valuation

First-Price Sealed-Bid Auctions

- Examples:
  - construction contracts (lowest bidder)
  - real estate
  - art treasures
- Typical rules
  - Bidders write their bids for the object and their names on slips of paper and deliver them to the auctioneer
  - The auctioneer opens the bid and finds the highest bidder
  - The highest bidder gets the object being sold, for a price equal to his/her own bid
  - Winner's profit = BV– price paid
  - Everyone else's profit = 0

First-Price Sealed-Bid (continued)

- Suppose that
  - There are n bidders
  - Each bidder has a private valuation, v_i, which is private information
  - But a probability distribution for v_i is common knowledge
    - Let's say v_i is uniformly distributed over [0, 100]
  - Let B_i denote the bid of player i
  - Let π_i denote the profit of player i

- What is the Nash equilibrium bidding strategy for the players?
  - Need to find the optimal bidding strategies

- First we'll look at the case where $n = 2$

First-Price Sealed-Bid (continued)

- Finding the optimal bidding strategies
  - Let B_i be agent *i*'s bid, and π_i be agent *i*'s profit
  - If $B_i \geq v_i$, then $\pi_i \leq 0$
    - So, assuming rationality, $B_i < v_i$
  - Thus
    - $\pi_i = 0 \quad\quad \text{if } B_i \neq \max_j \{B_j\}$
    - $\pi_i = v_i - B_i \quad \text{if } B_i = \max_j \{B_j\}$
  - How much below v_i should your bid be?
  - The smaller B_i is,
    - the less likely that *i* will win the object
    - the more profit *i* will make if *i* wins the object

First-Price Sealed-Bid (continued)

- Case $n = 2$
  - Suppose your BV is v and your bid is B
  - Let x be the other bidder's BV
    and αx be his/her bid, where $0 < \alpha < 1$
    - You don't know the values of x and α
  - Your expected profit is
    - $E(\pi) = P(\text{your bid is higher}) \cdot (v - B) + P(\text{your bid is lower}) \cdot 0$
- If x is uniformly distributed over [0, 100], then the pdf is $f(x) = 1/100$, $0 \leq x \leq 100$
  - $P(\text{your bid is higher}) = P(\alpha x < B) = P(x < B/\alpha) = \int_0^{B/\alpha} (1/100)\, dx = B/100\alpha$
  - so $E(\pi) = B(v - B)/100\alpha$
- If you want to maximize your expected profit (hence your valuation of money is risk-neutral), then your maximum bid is
    - $\max_B B(v-B)/100\alpha = \max_B B(v-B) = \max_B Bv - B^2$
    - maximum occurs when $v - 2B = 0 \Rightarrow B = v/2$
- So, bid ½ of what the item is worth to you!

First-Price Sealed-Bid (continued)

- With n bidders, if your bid is B, then
  - P(your bid is the highest) = $(B/100\alpha)^{n-1}$

- Assuming risk neutrality, you choose your bid to be
  - $\max_B B^{n-1}(v-B) = v(n-1)/n$
- As n increases, $B \to v$
  - I.e., increased competition drives bids close to the valuations

Dutch Auctions

- Examples
  - flowers in the Netherlands, fish market in England and Israel, tobacco market in Canada
- Typical rules
  - Auctioneer starts with a high price
  - Auctioneer lowers the price gradually, until some buyer shouts “Mine!”
  - The first buyer to shout “Mine!” gets the object at the price the auctioneer just called
  - Winner’s profit = BV – price
  - Everyone else’s profit = 0
- Dutch auctions are game-theoretically equivalent to first-price, sealed-bid auctions
  - The object goes to the highest bidder at the highest price
  - A bidder must choose a bid without knowing the bids of any other bidders
  - The optimal bidding strategies are the same

Sealed-Bid, Second-Price Auctions

- Background: Vickrey (1961)
- Used for
  - stamp collectors' auctions
  - US Treasury's long-term bonds
  - Airwaves auction in New Zealand
  - eBay and Amazon
- Typical rules
  - Bidders write their bids for the object and their names on slips of paper and deliver them to the auctioneer
  - The auctioneer opens the bid and finds the highest bidder
  - The highest bidder gets the object being sold, for a price equal to the *second highest* bid
- Winner's profit = BV – price
- Everyone else's profit = 0

Sealed-Bid, Second-Price (continued)

- Equilibrium bidding strategy:
  - It is a weakly dominant strategy to bid your true value: This property is also called **truthfulness** or **strategyproofness** of an auction.
- To show this, need to show that overbidding or underbidding cannot increase your profit and might decrease it.
- Let V be your valuation of the object, and X be the highest bid made by anyone else
- Let s_V be the strategy of bidding V, and π_V be your profit when using s_V
- Let s_B be a strategy that bids some $B \neq V$, and π_B be your profit when using s_B
- There are 3! = 6 possible numeric orderings of B, V, and X:
  - Case 1, $X > B > V$: You don't get the commodity either way, so $\pi_B = \pi_V = 0$.
  - Case 2, $B > X > V$: $\pi_B = V - X < 0$, but $\pi_V = 0$
  - Case 3, $B > V > X$: you pay X rather than your bid, so $\pi_B = \pi_V = V - X > 0$
  - Case 4, $X < B < V$: you pay X rather than your bid, so $\pi_B = \pi_V = V - X > 0$
  - Case 5, $B < X < V$: $\pi_B = 0$, but $\pi_V = V - X > 0$
  - Case 6, $B < V < X$: You don't get the commodity either way, so $\pi_B = \pi_V = 0$

Sealed-Bid, Second-Price (continued)

- Sealed-bid, 2nd-price auctions are nearly equivalent to English auctions
  - The object goes to the highest bidder
  - Price is close to the second highest BV (close since the second highest bids just a bit below his actual BV)

Coalitional Games with Transferable Utility

- Given a set of agents, a coalitional game defines how well each group (or **coalition**) of agents can do for itself—its payoff
  - Not concerned with
    - how the agents make individual choices within a coalition,
    - how they coordinate, or
    - any other such detail
- **Transferable utility** assumption: the payoffs to a coalition may be freely redistributed among its members
  - Satisfied whenever there is a universal **currency** that is used for exchange in the system
  - Implies that each coalition can be assigned a single value as its payoff

Coalitional Games with Transferable Utility

- A **coalitional game with transferable utility** is a pair $G = (N,v)$, where
  - N = {1, 2, …, n} is a finite set of players
  - **(nu)** $v : 2^N \rightarrow \Re$ associates with each coalition $S \subseteq N$ a real-valued payoff $v(S)$, that the coalition members can distribute among themselves
- v is the **characteristic function**
  - We assume $v(\emptyset) = 0$ and that v is non-negative.
- A coalition's payoff is also called its **worth**
- Coalitional game theory is normally used to answer two questions:

  (1) Which coalition will form?

  (2) How should that coalition divide its payoff among its members?
- The answer to (1) is often "the grand coalition" (all of the agents)
  - But this answer can depend on making the right choice about (2)

Example: A Voting Game

- Consider a parliament that contains 100 representatives from four political parties:
  - A (45 reps.), B (25 reps.), C (15 reps.), D (15 reps.)
- They're going to vote on whether to pass a $100 million spending bill| (and how much of it should be controlled by each party)
- Need a majority ($\geq$ 51 votes) to pass legislation
  - If the bill doesn't pass, then every party gets 0
- More generally, a voting game would include
  - a set of agents N
  - a set of *winning* coalitions $W \subseteq 2^N$
    - In the example, all coalitions that have enough votes to pass the bill
  - $v(S) = 1$ for each coalition $S \in W$
    - Or equivalently, we could use $v(S) =$ $100 million
  - $v(S) = 0$ for each coalition $S \notin W$

Superadditive Games

- A coalitional game $G = (N,v)$ is **superadditive** if the union of two disjoint coalitions is worth at least the sum of its members' worths
  - for all S, $T \subseteq N$, if $S \cap T = \varnothing$, then $v(S \cup T) \geq v(S) + v(T)$
- The voting-game example is superadditive
  - If $S \cap T = \varnothing$, $v(S) = 0$, and $v(T) = 0$, then $v(S \cup T) \geq 0$
  - If $S \cap T = \varnothing$ and $v(S) = 1$, then $v(T) = 0$ and $v(S \cup T) = 1$
  - Hence $v(S \cup T) \geq v(S) + v(T)$
- If G is superadditive, the grand coalition always has the highest possible payoff
  - For any $S \neq N$, $v(N) \geq v(S) + v(N–S) \geq v(S)$
- $G = (N,v)$ is **additive** (or **inessential**) if
  - For S, $T \subseteq N$ and $S \cap T = \varnothing$, then $v(S \cup T) = v(S) + v(T)$

Constant-Sum Games

- G is **constant-sum** if the worth of the grand coalition equals the sum of the worths of any two coalitions that partition N
  - $v(S) + v(N - S) = v(N)$, for every $S \subseteq N$

- Every additive game is constant-sum
  - additive => $v(S) + v(N - S) = v(S \cup (N - S)) = v(N)$
- But not every constant-sum game is additive
  - Example is a good exercise

Convex Games

- G is **convex (supermodular)** if for all $S,T \subseteq N$,
  - $v(S \cup T) + v(S \cap T) \geq v(S) + v(T)$
- It can be shown the above definition is equivalent to for all i in N and for all $S \subseteq T \subseteq N$-{$i$},
  - $v(T \cup \{i\})- v(T) \geq v(S \cup \{i\}) - v(S)$
  - Prove it as an exercise
- Recall the definition of a superadditive game:
  - for all $S,T \subseteq N$, if $S \cap T = \emptyset$, then $v(S \cup T) \geq v(S) + v(T)$
- It follows immediately that every super-additive game is a convex game

Simple Coalitional Games

- A game $G = (N, v)$ is **simple** for every coalition S,
    - either $v(S) = 1$ (i.e., S **wins**) or $v(S) = 0$ (i.e., S **loses**)
  - Used to model voting situations (e.g., the example earlier)
- Often add a requirement that if S wins, all supersets of S would also win:
    - if $v(S) = 1$, then for all $T \supseteq S$, $v(T) = 1$

- This doesn't quite imply superadditivity
  - Consider a voting game G in which 50% of the votes is sufficient to pass a bill
  - Two coalitions S and T, each is exactly 50% N
  - $v(S) = 1$ and $v(T) = 1$
  - But $v(S \cup T) \neq 2$

Proper-Simple Games

- G is a **proper simple game** if it is both simple and constant-sum
  - If S is a winning coalition, then $N - S$ is a losing coalition
    - $v(S) + v(N - S) = 1$, so if $v(S) = 1$ then $v(N - S) = 0$

- Relations among the classes of games:

{Additive games} $\subseteq$ {Super-additive games} $\subseteq$ {Convex games}

{Additive games} $\subseteq$ {Constant-sum game}

{Proper-simple games} $\subseteq$ {Constant-sum games}

{Proper-simple games} $\subseteq$ {Simple game}

Analyzing Coalitional Games

- Main question in coalitional game theory
  - How to divide the payoff to the grand coalition?
- Why focus on the grand coalition?
  - Many widely studied games are super-additive
    - Expect the grand coalition to form because it has the highest payoff
  - Agents may be required to join
    - E.g., public projects often legally bound to include all participants
- Given a coalitional game $G = (N, v)$, where $N = \{1, \ldots, n\}$
  - We'll want to look at the agents' shares in the grand coalition's payoff
    - The book writes this as **(Psi)** $\psi(N,v) = \mathbf{x} = (x_1, \ldots, x_n)$, where $\psi_i(N,v) = x_i$ is the agent's payoff
  - We won't use the ψ notation much
    - Can be useful for talking about several different coalitional games at once, but we usually won't be doing that

Terminology

- **Feasible payoff set**

  = {all payoff profiles that don't distribute more than the worth of the grand coalition}

  $= \{(x_1, \ldots, x_n) \mid x_1 + x_2 + \ldots + x_n\} \leq v(N)$

- **Pre-imputation set**

  $\mathcal{P}$ = {feasible payoff profiles that are **efficient**, i.e., distribute the entire worth of the grand coalition}

  $= \{(x_1, \ldots, x_n) \mid x_1 + x_2 + \ldots + x_n\} = v(N)$

- **Imputation set**

  $\mathcal{C}$ = {payoffs in $\mathcal{P}$ in which each agent gets at least what he/she would get by going alone (i.e., forming a singleton coalition)}

  $= \{(x_1, \ldots, x_n) \in \mathcal{P} : \forall i \in N, x_i \geq v(\{i\})\}$

im•pute: verb [trans.] represent as being done, caused, or possessed by someone; attribute *: the crimes **imputed to** Richard.*

Fairness, Symmetry

- What is a **fair** division of the payoffs?
  - Three axioms describing fairness
    - *Symmetry*, *dummy player*, and *additivity* axioms

- Definition: agents i and j are **interchangeable** if they always contribute the same amount to every coalition of the other agents
  - i.e., for every S that contains neither i nor j , $v(S \cup \{i\}) = v(S \cup \{j\})$

- **Symmetry axiom**: in a fair division of the payoffs, interchangeable agents should receive the same payments, i.e.,
  - if i and j are interchangeable and $(x_1, \ldots, x_n)$ is the payoff profile, then $x_i = x_j$

Dummy Players

- Agent i is a **dummy player** if i's contributes to any coalition is exactly the amount i can achieve alone
  - i.e., for all S s.t. $i \notin S$, $v(S \cup \{i\}) = v(S) + v(\{i\})$

- **Dummy player axiom**: in a fair distribution of payoffs, dummy players should receive payment equal to the amount they achieve on their own
  - i.e., if i is a dummy player and $(x_1, \ldots, x_n)$ is the payoff profile, then $x_i = v(\{i\})$

Additivity

- Let $G_1 = (N, v_1)$ and $G_2 = (N, v_2)$ be two coalitional games with the same agents
- Consider the combined game $G = (N, v_1 + v_2)$, where
  - $(v_1 + v_2)(S) = v_1(S) + v_2(S)$
- **Additivity axiom**: in a fair distribution of payoffs for G, the agents should get the sum of what they would get in the two separate games
  - i.e., for each player i, $\psi_i(N, v_1 + v_2) = \psi_i(N, v_1) + \psi_i(N, v_2)$

Shapley Values

- Recall that a pre-imputation is a payoff division that is both feasible and efficient

- **Theorem.** Given a coalitional game (N,v), there's a unique pre-imputation $\varphi(N,v)$ that satisfies the Symmetry, Dummy player, and Additivity axioms. For each player i, i's share of $\varphi(N,v)$ is

$$\varphi_i(N,v) = \frac{1}{|N|!} \sum_{S \subseteq N - \{i\}} |S|!\ (|N| - |S| - 1)!\ (v(S \cup \{i\}) - v(S))$$

- $\varphi_i(N,v)$ is called i's **Shapley value**
  - Lloyd Shapley introduced it in 1953

- It captures agent i's **average marginal contribution**
  - The average contribution that i makes to the coalition, averaged over every possible sequence in which the grand coalition can be built up from the empty coalition

Shapley Values

- Suppose agents join the grand coalition one by one, all sequences equally likely
- Let S = {agents that joined before i} and T = {agents that joined after i}
  - i's marginal contribution is $v(S \cup \{i\}) - v(S)$
    - independent of how S is ordered, independent of how T is ordered
  - Pr[S, then i, then T]

    = (# of sequences that include S then i then T) / (total # of sequences)

    = $|S|!\ |T|!\ /\ |N|!$
- Let $\varphi_{i,S}$ = Pr[S, then i, then T] × i's marginal contribution when it joins
- Then $\varphi_{i,S} = \dfrac{|S|!\,(|N| - |S| - 1)!}{|N|!}\,(v(S \cup \{i\}) - v(S))$
- Let $\varphi_i(N,v)$ = expected contribution over all possible sequences
- Then $\varphi_i(N,v) = \sum_{S \subseteq N - \{i\}} \varphi_{i,S} = \dfrac{1}{|N|!} \sum_{S \subseteq N - \{i\}} |S|!\,(|N| - |S| - 1)!\,(v(S \cup \{i\}) - v(S))$

Example

- The voting game again
  - Parties *A*, *B*, *C*, and *D* have 45, 25, 15, and 15 representatives
  - A simple majority (51 votes) is required to pass the $100M bill
- How much money is it fair for each party to demand?
  - Calculate the Shapley values of the game
- Every coalition with ≥ 51 members has value 1; other coalitions have value 0

- Recall what it means for two agents *i* and *j* to be interchangeable:
  - for every *S* that contains neither *i* nor *j* , $v\,(S \cup \{i\}) = v\,(S \cup \{j\})$
- *B* and *C* are interchangeable
  - Each adds 0 to $\varnothing$, 1 to $\{A\}$, 0 to $\{D\}$, and 0 to $\{A,D\}$
- Similarly, *B* and *D* are interchangeable, and so are *C* and *D*
- So the fairness axiom says that *B*, *C*, and *D* should each get the same amount

- Recall that

$$\varphi_{i,S} = \frac{|S|!\,(|N| - |S| - 1)!\,(v(S \cup \{i\}) - v(S))}{|N|!}$$

$$\varphi_i(N, v) = \sum_{S \subseteq N - \{i\}} \varphi_{i,S} = \frac{1}{|N|!} \sum_{S \subseteq N - \{i\}} \underbrace{|S|!\,(|N| - |S| - 1)!\,(v(S \cup \{i\}) - v(S))}$$

- In the example, it will be useful to let $\varphi'_{i,S}$ be the term inside the summation
  - Hence $\varphi'_{i,S} = |N|!\ \varphi_{i,S}$

- Let's compute $\varphi_A(N, v)$
- $N = |\{A,B,C,D\}| = 4$, so $\varphi'_{A,S} = |S|!\,(3 - |S|)!\,(v(S \cup A) - v(S))$
- S may be any of the following:
  - $\varnothing$, {B}, {C}, {D}, {B,C}, {B,D}, {C,D}
- We need to sum over all of them:

$$\varphi_A(N, v) = \frac{1}{4!}\left(\varphi'_{A,\varnothing} + \varphi'_{A,\{B\}} + \varphi'_{A,\{C\}} + \varphi'_{A,\{D\}} + \varphi'_{A,\{B,C\}} + \varphi'_{A,\{B,D\}} + \varphi'_{A,\{C,D\}} + \varphi'_{A,\{B,C,D\}}\right)$$

A has 45 members
B has 25 members
C has 15 members
D has 15 members

$$\varphi'_{A,S} = |S|!\,(3-|S|)!\,(v(S \cup A) - v(S))$$

$S = \varnothing$	→	$v(\{A\}) - v(\varnothing) = 0 - 0 = 0$	→	$\varphi'_{A,\varnothing} = 0!\ 3!\ 0 = 0$
$S = \{B\}$	→	$v(\{A,B\}) - v(\{B\}) = 1 - 0 = 1$	→	$\varphi'_{A,\{B\}} = 1!\ 2!\ 1 = 2$
$S = \{C\}$	→	same		
$S = \{D\}$	→	same		
$S = \{B,C\}$	→	$v(\{A,B,C\}) - v(\{B,C\}) = 1 - 0 = 1$	→	$\varphi'_{A,\{B,C\}} = 2!\ 1!\ 1 = 2$
$S = \{B,D\}$	→	same		
$S = \{C,D\}$	→	same		
$S = \{B,C,D\}$	→	$v(\{A,B,C,D\}) - v(\{B,C,D\}) = 1 - 1 = 0$	→	$\varphi'_{A,\{B,C,D\}} = 3!\ 0!\ 0 = 0$

$$\varphi_A(N,v) = \frac{1}{4!}\left(\varphi'_{A,\varnothing} + \varphi'_{A,\{B\}} + \varphi'_{A,\{C\}} + \varphi'_{A,\{D\}} + \varphi'_{A,\{B,C\}} + \varphi'_{A,\{B,D\}} + \varphi'_{A,\{C,D\}} + \varphi'_{A,\{B,C,D\}}\right)$$

$$= \frac{1}{24}(0+2+2+2+2+2+2+0) = 12/24 = 1/2$$

$$\varphi_i(N,v) = \frac{1}{|N|!} \sum_{S \subseteq N - \{i\}} |S|!\ (|N| - |S| - 1)!\ (v(S \cup \{i\}) - v(S))$$

- Similarly, $\varphi_B = \varphi_C = \varphi_D = 1/6$
  - The text calculates it using Shapley's formula
- Here's another way to get it:
  - If *A* gets ½, then the other ½ will be divided among *B*, *C*, and *D*
  - They are interchangeable, so a fair division will give them equal amounts: 1/6 each

- So distribute the money as follows:
  - *A* gets (1/2) \$100M = \$50M
  - *B*, *C*, *D* each get (1/6) \$100M = \$16⅔M

Stability of the Grand Coalition

- Agents have incentive to form the grand coalition iff there aren't any smaller coalitions in which they could get higher payoffs
- Sometimes a subset of the agents may prefer a smaller coalition
- Recall the Shapley values for our voting example:
    - *A* gets \$50M; *B, C, D* each get \$ $16\frac{2}{3}$M
  - *A* on its own can't do better
  - But {*A, B*} have incentive to defect and divide the \$100M
    - e.g., \$75M for *A* and \$25M for *B*

- What payment divisions would make the agents want to join the grand coalition?

The Core

- The **core** of a coalitional game includes every payoff vector **x** that gives every sub-coalition S at least as much in the grand coalition as S could get by itself
  - All feasible payoff vectors $\mathbf{x} = (x_1, \ldots, x_n)$ such that for every $S \subseteq N$,

$$\sum_{i \in S} x_i \geq v(S)$$

- For every payoff vector **x** in the core, no S has any incentive to **deviate** from the grand coalition
  - i.e., form their own coalition, excluding the others
- It follows immediately that if **x** is in the core then **x** is efficient
  - Why?

Analogy to Nash Equilibria

- The core is an analog of the set of all Nash equilibria in a noncooperative game
  - There, no agent can do better by deviating from the equilibrium
- But the core is stricter
  - No set of agents can do better by deviating from the grand coalition
- Analogous to the set of **strong** Nash equilibria
  - Equilibria in which no coalition of agents can do better by deviating

- Unlike the set of Nash equilibria, the core may sometimes be empty
  - In some cases, no matter what the payoff vector is, some agent or group of agents has incentive to deviate

Example of an Empty Core

- Consider the voting example again:
  - Shapley values are \$50M to *A*, and \$16.33M each to *B*, *C*, *D*
- The minimal coalitions that achieve 51 votes are
  - {A,B}, {A,C}, {A,D}, {B,C,D}
- If the sum of the payoffs to B, C, and D is < \$100M, this set of agents has incentive to deviate from the grand coalition
  - Thus if **x** is in the core, **x** must allocate \$100M to {B, C, D}
  - But if B, C, and D get the entire \$100M, then A (getting \$0) has incentive to join with whichever of B, C, and D got the least
    - e.g., form a coalition {A,B} without the others
  - So if **x** allocates the entire \$100M to {B,C,D} then **x** cannot be in the core
- So the core is empty

Simple Games

- There are several situations in which the core is either guaranteed to exist, or guaranteed not to exist
  - The first one involves simple games
- Recall: G is **simple** for every coalition S, either $v(S) = 1$ or $v(S) = 0$
- Player i is a **veto player** if $v(S) = 0$ for any $S \subseteq N - \{i\}$
- **Theorem**. In a simple game, the core is empty iff there is no veto player
- Example: previous slide

Simple Games

- **Theorem**. In a simple game in which there are veto players, the core is {all payoff vectors in which non-veto players get 0}

- **Example**: consider a modified version of the voting game
  - An 80% majority is required to pass the bill
- Recall that *A*, *B*, *C*, and *D* have 45, 25, 15, and 15 representatives
  - The minimal winning coalitions are {A, B, C} and {A, B, D}
  - All winning coalitions must include both A and B
  - So A and B are veto players
    - The core includes all distributions of the $100M among A and B
    - Neither A nor B can do better by deviating

Non-Additive Constant-Sum Games

- Recall:
  - G is constant-sum if for all S, $v(S) + v(N - S) = v(N)$
  - G is additive if $v(S \cup T) = v(S) + v(T)$ whenever S and T are disjoint
- **Theorem**. Every non-additive constant-sum game has an empty core
- **Example:** consider a constant-sum game G with 3 players a, b, c
  - Suppose v(a) = 1, v(b) = 1, v(c) = 1, v({a,b,c})=4
  - Then v(a) + v({b,c}) = v({a,b})+v(c) = v({a,c}) + v(b) = 4
  - Thus v({b,c}) = 4 – 1 = 3 ≠ v(b) + v(c)
  - So G is not additive
- Consider **x** = (1.333, 1.333, 1.333)
  - v({a,b}) = 3, so if {a,b} deviate, they can allocate (1.5,1.5)
- To keep {a,b} from deviating, suppose we use **x** = (1.5, 1.5, 1)
  - v({a,c}) = 3, so if {a,c} deviate, they can allocate (1.667, 1.333)

Convex Games

- Recall:
  - G is **convex** if for all $S,T \subseteq N$, $v(S \cup T) \geq v(S) + v(T) - v(S \cap T)$
- **Theorem**. Every convex game has a nonempty core
- **Theorem**. In every convex game, the Shapley value is in the core

Modified Parliament Example

- 100 representatives from four political parties:
  - *A* (45 reps.), *B* (25 reps.), *C* (15 reps.), *D* (15 reps.)
- Any coalition of parties can approve a spending bill worth \$1K times the number of representatives in the coalition:

$$v(S) = \sum_{i \in S} \$1000 \times \text{size}(i)$$

$v(A)$ = \$45K, $v(B)$ = \$25K, $v(C)$ = \$15K, $v(D)$ = \$15K,

$v(\{A,B\})$ = \$70K, $v(\{A,C\})$ = \$60K, $v(\{A,D\})$ = \$60K,

$v(\{B,C\})$ = \$40K, $v(\{B,D\})$ = \$40K, $v(\{C,D\})$ = \$30K, …

$v(\{A,B,C,D\})$ = \$100K

- Is the game convex?

Modified Parliament Example

- Let S be the grand coalition
  - What is each party's Shapley value in S?

- Each party's Shapley value is the average value it adds to S, averaged over all 24 of the possible sequences in which S might be formed:

A, B, C, D;   *A, B, D, C;*   *A, C, B, D;*   *A, C, D, B;*   *etc*

- In every sequence, every party adds exactly \$1K times its size
- Thus every party's Shapley value is \$1K times its size:
  - φ_A = \$45K, φ_B = \$25K, φ_C = \$15K, φ_D = \$15K

Modified Parliament Example

- Suppose we distribute $v(S)$ by giving each party its Shapley value
- Does any party or group of parties have an incentive to leave and form a smaller coalition T?
  - $v(T)$ = \$1K times the number of representatives in T
    = the sum of the Shapley values of the parties in T
  - If each party in T gets its Shapley value, it does no better in T than in S
  - If some party in T gets more than its Shapley value, then another party in T will get less than its Shapley value
- No case in which every party in T does better in T than in S
- No case in which all of the parties in T will have an incentive to leave S and join T

- Thus the Shapley value is in the core