

Introduction to Game Theory

Irwin King

Department of Computer Science and Engineering
The Chinese University of Hong Kong, Shatin, N.T., Hong Kong



A Brief History of Game Theory

- 1713: the first known discussion of game theory in a letter written by James Waldegrave.
- 1928: became a unique field after a paper by John von Neumann was published, followed by his 1944 book *Theory of Games and Economic Behavior*.
- 1950: the first mathematical discussion of the prisoner's dilemma (囚徒困境) appeared with an experiment by Merrill M. Flood and Melvin Dresher. The concept of Nash equilibrium (纳什均衡) was developed by John Nash. Game theory experienced a flurry of activity in the 1950s and 1960s.
- 1970s: game theory was extensively applied in biology, largely as a result of the work of John Maynard Smith.



A Brief History of Game Theory

- 1994: Nash, Selten and Harsanyi became Economics Nobel Laureates for their contributions to economic game theory.
- 2005: game theorists Thomas Schelling and Robert Aumann became Nobel Laureates.
- 2007: Leonid Hurwicz, together with Eric Maskin and Roger Myerson, was awarded the Nobel Prize in Economics "for having laid the foundations of mechanism design theory."



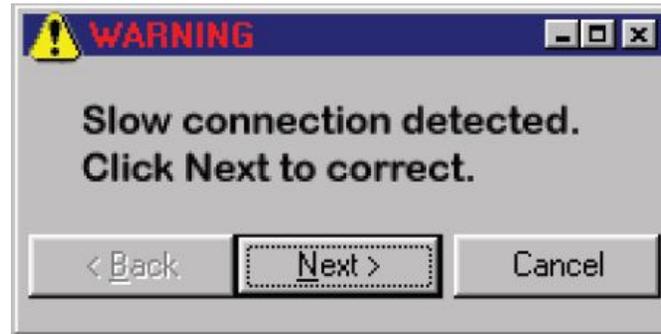
A Brief History of Game Theory

- 1994: Nash, Selten and Harsanyi became Economics Nobel Laureates for their contributions to economic game theory.
- 2005: game theorists Thomas Schelling and Robert Aumann became Nobel Laureates.
- 2007: Leonid Hurwicz, together with Eric Maskin and Roger Myerson, was awarded the Nobel Prize in Economics "for having laid the foundations of mechanism design theory."

Q: What does a game really look like?



TCP Backoff: Background



- Internet traffic is governed by the **TCP protocol**.
- A correct implementation of TCP protocol has a “**backoff mechanism**”.
- A **defective** implementation of TCP does not back off when congestion occurs.



TCP Backoff Game

- Imagine that you and a colleague are the only people using the internet. You each have two possible strategies: C (using a correct implementation) and D (using a defective one).
- The rules are:
 - **both use a correct implementation:** both get 1 ms delay
 - **one correct, one defective:** 4ms for correct, 0 ms for defective
 - **both defective:** both get a 3 ms delay.



TCP Backoff Game: Practice

- Play in pairs.
- Each player chooses one implementation (C or D) and then write down on a paper (do not let your opponent see it). After both players make the decisions, show your choice to your opponent and find your delay according to the rules.
- Play the game for 10 times and calculate your average delay.
- One with the smallest average delay is the winner.



TCP Backoff Game: Practice

- The rules are:
 - **both use a correct implementation:** both get 1 ms delay
 - **one correct, one defective:** 4ms for correct, 0 ms for defective
 - **both defective:** both get a 3 ms delay.



Discussion Time

- What strategy do you use?
- Can you make your strategy better?
- How do you think of your opponent when making decision?



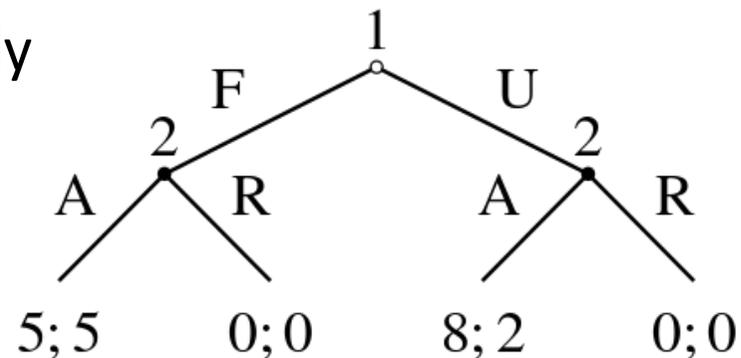
Defining Games

- **Players:** who are the decision makers?
 - People? Companies? Shooter & Goalkeeper?
- **Actions:** what can the players do?
 - Bid in an auction? Decide when to sell a stock?
- **Payoffs (Utilities):** what motivates players?
 - Can they get profit? Will they lose something?



Defining Games

- Two Standard Representations
 - Strategic Form (Normal Form 范式博弈)
 - Players move simultaneously
 - Payoffs are functions of combinations of actions
 - Extensive Form (展开形式的博弈)
 - Players move sequentially
 - Represented as a tree
 - E.g. chess



The Strategic Form

- A strategic game $\langle N, (A_i), (u_i) \rangle$ consists of:
 - Players: a finite set $N = \{1, \dots, n\}$
 - Actions: for player $i \in N$, a nonempty set A_i including all available actions for player i ;
 $A = \prod_{i \in N} A_i$ is the set of action profiles (or set of consequences)
 - Payoff functions: for player $i \in N$, $u_i: A \mapsto \mathbb{R}$ is its payoff (or utility) function



The Strategic Form: Example

- TCP Backoff Game as the strategic form
- $N = \{1, 2\}$
- $A_1 = A_2 = \{C, D\}$,
 $A = A_1 \times A_2 = \{(C, C), (C, D), (D, C), (D, D)\}$
- $u_1((C, C)) = -1, u_1((C, D)) = -4, u_1((D, C)) = 0,$
 $u_1((D, D)) = -3$
- $u_2((C, C)) = -1, u_2((C, D)) = 0, u_2((D, C)) = -4,$
 $u_2((D, D)) = -3$



The Standard Matrix Representation

- Writing a 2-player game as a matrix:
 - row player is player 1, column player is player 2
 - rows correspond to actions $a_1 \in A_1$, columns correspond to actions $a_2 \in A_2$
 - cells listing payoff values for each player: the row player first, then the column



The Standard Matrix Representation: Example

- The TCP Backoff Game written as a matrix

		2	
		C	D
1	C	-1, -1	-4, 0
	D	0, -4	-3, -3



Type of Games

- Cooperative vs. Non-cooperative (合作与非合作)
 - A game is *cooperative* if the players are able to form binding commitments. (In this tutorial, we focus on non-cooperative games.)
- Symmetric vs. Asymmetric (对称与非对称)
 - A symmetric game is a game where the payoffs for playing a particular strategy depend only on the other strategies employed, not on who is playing them.
- Zero-sum vs. Non-zero-sum (零和和非零和)
 - Zero-sum games are a special case of *constant-sum* games, in which choices by players can neither increase nor decrease the available resources.
-



		2	
		C	D
1	C	-1, -1	-4, 0
	D	0, -4	-3, -3

- **Non-Cooperative** (each player makes his/her decision independently);
- **Symmetric** (one will always get -4(0) if he/she plays C(D) and the other plays D(C));
- **Non-Zero-sum** (for (C,C), the sum of payoff is -2).



Symmetric Games

- For a 2x2 game, its payoff matrix must conform to the following.

		2	
		A	B
1	A	a, a	c, d
	B	d, c	b, b



Symmetric Games

- For a 2x2 game, its payoff matrix must conform to the following.

Q: What about a zero-sum game?

		A	B
1	A	a, a	c, d
	B	d, c	b, b



Zero-sum Games

- One instance of the zero-sum game

		2	
		A	B
1	A	$a, -a$	$b, -b$
	B	$c, -c$	$d, -d$



Example: Prisoner's Dilemma

		B	
		Cooperate	Defect
A	Cooperate (silent)	Each serves 1 year	A gets 3 years B goes free
	Defect (betray)	A goes free B gets 3 years	Each serves 2 years

Q: What's the strategic form of this game?
Is this game cooperative? symmetric? zero-sum?



Example: Prisoner's Dilemma

		B	
		Cooperate	Defect
A	Cooperate	Each serves 1 year	A gets 3 years B goes free
	Defect	A goes free B gets 3 years	Each serves 2 years

Non-cooperative, symmetric, non-zero-sum



Example: Matching Pennies

- One player wants to match; the other wants to mismatch

	Heads	Tails
Heads	1, -1	-1, 1
Tails	-1, 1	1, -1

Q: What's the strategic form of this game?
Are this game cooperative? symmetric? zero-sum?



Example: Matching Pennies

- One player wants to match; the other wants to mismatch

	Heads	Tails
Heads	1, -1	-1, 1
Tails	-1, 1	1, -1

Non-cooperative, asymmetric, zero-sum



Example: Battle of the Sexes

		Wife	
		Boxing	Opera
Husband	Boxing	2, 1	0, 0
	Opera	0, 0	1, 2

Q: What's the strategic form of this game?
Are this game cooperative? symmetric? zero-sum?



Example: Battle of the Sexes

		Wife	
		Boxing	Opera
Husband	Boxing	2, 1	0, 0
	Opera	0, 0	1, 2

Non-cooperative, asymmetric, non-zero-sum



Keynes Beauty Contest Game

- A concept developed by John Maynard Keynes to explain price fluctuations in equity markets.
- Entrants are asked to choose from a set of photographs of women that are the "most beautiful." Those who picked the most popular face will be awarded.
- The stylized version is:
 - Each player names an integer between 1 and 100.
 - The player who names the integer closest to two thirds of the **average** integer wins a prize, the other players get nothing.
 - Ties are broken uniformly at random



Keynes Beauty Contest Game: Practice

- Play in groups (5 to 7 people).
- Each player **chooses a number between 1 and 100** then write down on a paper (do not let your opponent see it).
- After all the players write down his/her number, one player collects all the number and **calculate the average**. The winners are those who name the number **closest to $2/3$ of the average**. Each winner gets a **score of $1/k$** (k is the total number of winners).
- Play the game for 10 times and the one with the highest score is the final winner.



Keynes Beauty Contest Game: Analysis

- Let the average integer be X
- X has to be less than 100, so the optimal strategy of any player has to be more than 67.
- If X is no more than 67, then the optimal strategy of any player has to be no more than $\frac{2}{3}67$.
-
- By Iterating, the unique Nash equilibrium of this game is for every player to announce 1!



Strategic Reasoning

- What will other players do?
- What should I do in response?
- Each player **best responds** to the others: Nash equilibrium



Strategic Reasoning

- For 2/3-contest
 - Level 0 player chooses randomly from the interval
 - Level 1 player assumes others are Level 0 players so it would be 50
 - Level 2 player assumes others are Level 1 players so it would be $50 \times \frac{2}{3} = 33$
 - Level 3 player assumes others are Level 2 players so it would be $33 \times \frac{2}{3} = 11$
 - Level n player assumes others are Level n-1 players



Nash Equilibrium

- A consistent list of actions in which each player's action maximizes his or her payoff given the actions of the others.
- The equilibrium action profile should be **stable**: nobody wants to change its action if the equilibrium profile is played.



Nash Equilibrium: Formal Definition

- Best Response

- Let $a_{-i} = \langle a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n \rangle$. Then $a = (a_{-i}, a_i)$.

- a_i is the best response of a_{-i} if and only if $\forall a_i \in A_i, u_i(a_i^*, a_{-i}) \geq u_i(a_i, a_{-i})$.



Nash Equilibrium: Formal Definition

- Nash Equilibrium
 - $a = \langle a_1, \dots, a_n \rangle$ is a pure strategy Nash equilibrium if and only if $\forall i, a_i$ is the best response of a_{-i} .
 - Another kind of Nash equilibrium is mixed strategy Nash equilibrium



Example: TCP Backoff Game

	C	D
C	-1, -1	-4, 0
D	0, -4	-3, -3



Practice: Find the pure strategy Nash equilibrium of previous examples



Example: Prisoner's Dilemma

	Cooperate	Defect
Cooperate	freedom, freedom	freedom & award, imprisonment
Defect	imprisonment, freedom & award	remission, remission



Example: Matching Pennies

- One player wants to match; the other wants to mismatch

	Heads	Tails
Heads	1, -1	-1, 1
Tails	-1, 1	1, -1



Example: Battle of the Sexes

		Husband	
		Boxing	Opera
Wife	Boxing	2, 1	0, 0
	Opera	0, 0	1, 2



Challenges

- Which kind of game has a pure strategy Nash equilibrium?
- Generally, how to find a pure strategy Nash equilibrium if it exists?



Dominant Strategy (支配性策略)

- A player i 's strategy a_i'' **strictly dominates** his strategy a_i' if $(a_{-i}, a_i'') >_i (a_{-i}, a_i')$ for every a_{-i} .
- A player i 's strategy a_i'' **weakly dominates** his strategy a_i' if $(a_{-i}, a_i'') \geq_i (a_{-i}, a_i')$ for every a_{-i} , and $(a_{-i}, a_i'') >_i (a_{-i}, a_i')$ for some a_{-i} .



Dominant Strategy

- A **dominant strategy** for player i is a strategy that is the best no matter what other players do.
- A **strictly dominant strategy** for player i is a strategy that strictly dominates all of his other strategies.
- A **weakly dominant strategy** for player i is a strategy that weakly dominates all of his other strategies.



Example: TCP Backoff Game

	C	D
C	-1, -1	-4, 0
D	0, -4	-3, -3



Practice: Find the Dominant Strategy
of each player in previous examples



Example: Prisoner's Dilemma

	Cooperate	Defect
Cooperate	freedom, freedom	freedom & award, imprisonment
Defect	imprisonment, freedom & award	remission, remission



Example: Battle of the Sexes

		Wife	
		Boxing	Opera
Husband	Boxing	2, 1	0, 0
	Opera	0, 0	1, 2



Challenges

- Does a game with a pure strategy Nash equilibrium always have dominant strategy for each player?
- Generally, how to find a dominant strategy if it exists?
- If every player in a game has its dominant strategy, does this game guarantee to have pure strategy Nash equilibrium?



TCP Backoff Game: Revisit

	C	D
C	-1, -1	-4, 0
D	0, -4	-3, -3

Q: If you are a web administrator (i.e. an outsider of this game), which strategies do you prefer them to choose? Why?



Q: What if we change the payoffs like this?

	C	D
C	-1, -1	-4, 0
D	0, -4	-100, -100



Q: What if we change the payoffs like this?

	C	D
C	-1, -1	-2, 0
D	0, -2	-3, -3



The Concept of Incentives (激勵)

- Definition: something that motivates an individual to perform an action.
- One common taxonomy for incentives:
 - Remunerative incentives (financial incentives)
 - Moral incentives
 - Coercive incentives
 - Natural incentives
- **Main target:** all economic activity (both in terms of individual decision-making and in terms of cooperation and competition within a larger institutional structure).
- **Aim:** provide value for money and contribute to organizational success.



Example: The Auction

- First-Price Sealed-Bid Auctions
 - bidders submit written bids without knowing the bid of the other people in the auction, and in which the highest bidder wins the auction.

Player					
					
Submitted Bid	b_1	b_2	b_3	b_4	b_5
Valuation	v_1	v_2	v_3	v_4	v_5



Example: The Auction

- $A_i = [0, +\infty)$.
- $u_i(b_i) = \begin{cases} v_i - b_i & \text{Player } i \text{ wins} \\ 0 & \text{Otherwise} \end{cases}$
- Assumptions:
 - $v_1 \geq v_2 \geq v_3 \geq v_4 \geq v_5$
 - Player with the lowest index wins if more than one player submits the highest bid.



Example: The Auction

- $A_i = [0, +\infty)$.
- $u_i(b_i) = \begin{cases} v_i - b_i & \text{Player } i \text{ wins} \\ 0 & \text{Otherwise} \end{cases}$
- Assumptions:
 - $v_1 \geq v_2 \geq v_3 \geq v_4 \geq v_5$
 - Player with the lowest index wins if more than one player submits the highest bid.

Q: What are the Nash equilibria?



Example: The Vickrey Auction

- Vickrey Auctions
 - bidders submit written bids without knowing the bid of the other people in the auction, and in which the highest bidder wins, but the price paid is the second-highest bid.



Example: The Vickrey Auction

- Vickrey Auctions
 - bidders submit written bids without knowing the bid of the other people in the auction, and in which the highest bidder wins, but the price paid is the second-highest bid.

Q: If we follow the same assumption, what are the Nash equilibria now?

Q: If you are the owner of an auction company, which type of auction do you prefer? Why?



When game theory meets crowdsourcing and human computation



Crowdsourcing Sites

- The typical structure of crowdsourcing sites:
 - Task
 - Reward
 - Time period
 - Users
- When the time period ends, a subset of submissions are selected, and the corresponding users are granted the reward.
- E.g. In TopCoder.com, users select among several tasks asking for a Quality Assurance plan for a software, each offering different rewards.



Crowdsourcing Modeled as Games

- Crowdsourcing can be modeled as a **two-stage** game [DiPalantino2009]:
 - Users select among tasks offering different rewards;
 - Upon joining a task, users those who selected it compete amongst themselves for the reward.
- More assumptions:
 - Users are endowed with a private skill;
 - Skills for different users are drawn independently at random;
 - The second stage of the game can be modeled as an **all-pay auction** (the bidders who do not win the auction should also pay their bids).



Challenges

- What are the differences between all-pay auction and previous types of auctions we introduced?
- Why all-pay auction is more suitable for this scenario?
- What is the incentive in this scenario?



The ESP Game

- A two-player game for labeling images on the web



red, car,
Ferrari,
sportscar,
exhibition,
...

- Observation: players tend to coordinate on “easy” words rather than “hard” words



The ESP Game: Model

- A two-player game of incomplete information (players may or may not know some information about the other players)
- A player makes two decisions:
 - Picks an “effort level” (the difficulty of the word domain)
 - Samples a sequence of words to report
- Each player first prefers to match rather than not, and then prefers to match earlier rather than later
- **Conclusion:** The equilibrium is to choose the low effort level [Weber2008]



The ESP Game: Incentives

- An important question in the area of incentive design for the ESP game is how to design incentives to elicit high effort from players and as a result, richer, more descriptive labels.
- One solution: alternative scoring mechanism
 - rare-word first preference;
 - Taboo words;
 - ?



References

- Books

- Fudenberg, Drew, and Jean Tirole. *Game Theory*. Cambridge, MA: MIT Press, 1991.
- Osborne, Martin, and Ariel Rubinstein. *A Course in Game Theory*. Cambridge, MA: MIT Press, 1994.

- Online Courses

- Open Yale courses - ECON 159 Game Theory
<http://oyc.yale.edu/economics/econ-159>
- Game Theory by Yoav Shoham, Matthew O. Jackson, Kevin Leyton-Brown
<https://www.coursera.org/course/gametheory>

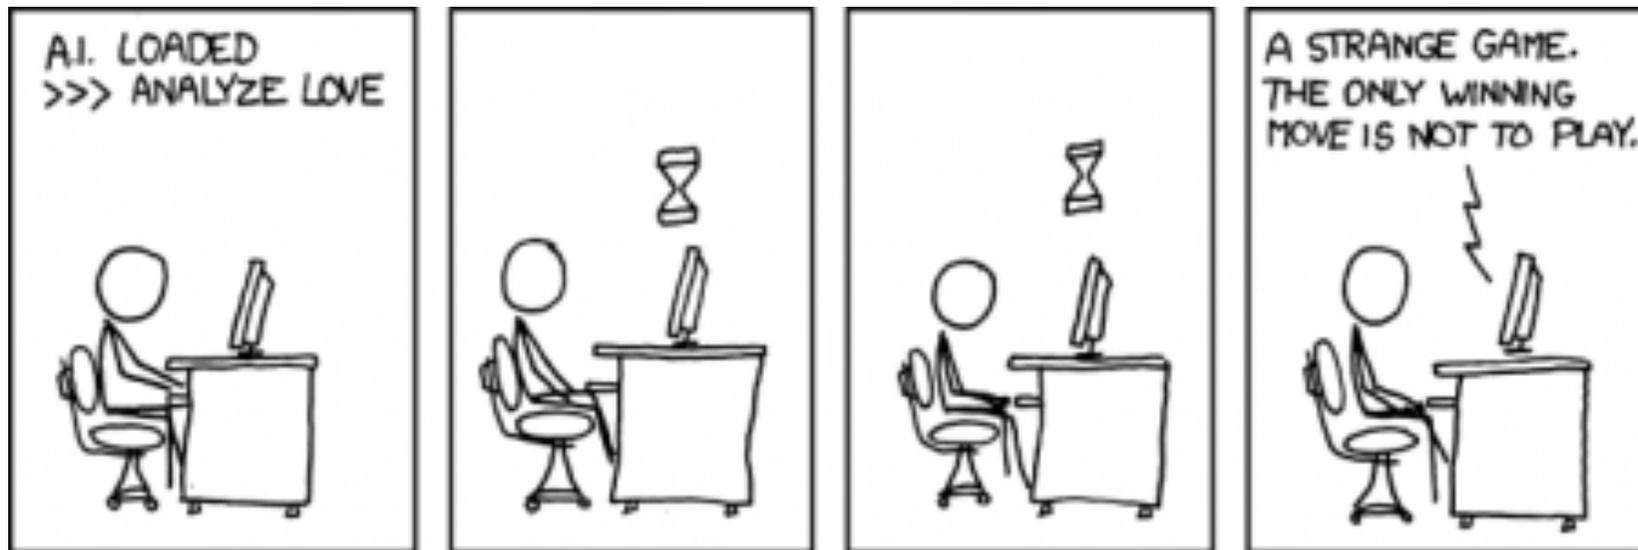


References

- Journal and Conference Papers
 - S. Jain, and D. C. Parkes. "The role of game theory in human computation systems." *Proceedings of the ACM SIGKDD Workshop on Human Computation*. ACM, 2009.
 - D. DiPalantino, and M. Vojnovic. "Crowdsourcing and all-pay auctions." *Proceedings of the 10th ACM conference on Electronic commerce*. ACM, 2009.
 - I. Weber, S. Robertson, and M. Vojnovic. Rethinking the ESP game. Technical report, Microsoft Research, 2008.
 - D. Yang, et al. "Crowdsourcing to smartphones: incentive mechanism design for mobile phone sensing." *Proceedings of the 18th annual international conference on Mobile computing and networking*. ACM, 2012.



Q&A



- Game Theory (博弈论)
- Incentives (奖励)
- Nash equilibrium (纳什均衡)
- Strategic Form (战略形式)
- Extended Form (扩展形式)

