Game Theory

Jun Wang
Computer Science, UCL

Oct 15 2018, SJTU

Why Game Theory?

- **Game Theory** (addressing incentive issues)

-> Human computation/crowd sourcing

-> repeated games -> bittorent file sharing

-> Auction theory -> online advertising

-> ebay

Summary

- Game Theory
  - Prisoner's dilemma
  - Nash equilibrium
- ISPs Game
- Human Computation (introduction)

Game Theory

- Game Theory model situations in which multiple participants interact or affect each other's outcomes
  - In 1713, Waldegrave provides a *minimax mixed strategy solution* to a two-person version of the card game *le Her*
  - Game theory began by John von Neumann in 1928
  - More than 10 game-theorists have won the Nobel Memorial Prize in Economic Sciences

Example (1/2)

- Suppose you can either study for the *exam* or the *presentation* for a course
  - Want to maximize your average grade (equal weight)
  - Your partner has to make a similar choice
- **Exam**
  - By studying for the exam, you get 92, else 80
- **Presentation**
  - If both prepare, each gets 100; if one prepares 92, but if neither prepares you both get 84.
- Can‘t communicate, how do you decide?
  - Your partner will also need to decide...

Example (2/2)

Your partner

You

	Presentation	Exam
Presentation	90,90	86,92
Exam	92,86	88,88

Strategies

Payoff for
row player,column player

- What will happen in this game?

Each player thinks studying for exam is safer

Reasoning about behavior in a game

- How are the players likely to behave?
- Need a few assumptions
  - The payoff for a player should capture all rewards
  - Players know all possible strategies and payoffs to other players
    - Not always true!
  - Players are *rational*
    - Each players wants to maximize her own payoff
    - Each player succeeds in picking the optimal strategy

Reasoning about behavior in a game

- Let‘s look at behavior in the example
- Each player has a *strictly dominant strategy*
  - “Exam“ strategy strictly better than all other options
- No matter what your partner does, you should study for the exam
  - We can predict the outcome
- But both could be better off!
  - The 90,90 outcome won‘t happen during rational play

The prisoner‘s dilemma

- Two suspects arrested and suspected for robbery.
- Interrogated separately
  - If neither confesses, both get 1 year sentence
  - If one confesses, other person gets 10 years
  - If both confess, 4 years in prison each

	Don‘t confess	Confess
Don‘t confess	-1, -1	-10, 0
Confess	0, -10	-4, -4

Interpretations

- Confessing is a *strictly dominant strategy*
  - Not confessing is better for both of them
  - But under rational play this is not achieved
- The dilemma arises frequently
  - e.g. performance-enhancing drugs for athletes

	Don‘t use drugs	Use drugs
Don‘t use drugs	3, 3	1, 4
Use drugs	4, 1	2, 2

  - Known as “arms races“
  - The payoffs must be aligned right

Best response

- The best choice of one player, given a belief of what the other player will do
- Let‘s make this a bit more formal
  - Player 1 chooses strategy *S*
  - Player 2 chooses strategy *T*
  - Payoff to player i given S,T is $P_i(S,T)$
  - *Def*: *S* is a *best response* to *T* if $P_1(S,T) \geq P_1(S',T)$ for all other strategies S‘ of player 1.
  - *S strict best response* when $P_1(S,T) > P_1(S',T)$

Dominant strategy

- *Def:* A *dominant strategy* is a best response to *every* strategy of the other player
  - Analogous for a *strictly dominant strategy*.
  - In the prisoner's dilemma, both players had strictly dominant strategies
  - Thus easy to predict what would happen
- But what about games without dominant strategies?

Example: Marketing game (1/2)

- A game in which only one player has a dominant strategy
  - Two firms entering markets. Firm 2 established.
  - 60% of people want low-priced, 40% upscale
  - If they compete directly, Firm 1 gets 80% of sales and Firm 2 gets 20%.
  - Otherwise, each occupies own market segment

Firm 2

Firm 1	Low-Priced	Upscale
Low-Priced	.48, .12	.60, .40
Upscale	.40, .60	.32, .08

Example: Marketing game (2/2)

- What happens in this game?
  - Firm 1 has a dominant strategy: *Low-priced*, whereas Firm 2 does not
  - Firm 2 can thus assume Firm 1 will play this strategy
  - Our prediction for this game is .60, .40

Firm 2

Firm 1

	Low-Priced	Upscale
Low-Priced	.48, .12	.60, .40
Upscale	.40, .60	.32, .08

Equilibrium concepts

- What if neither player has a dominant strategy?
  - How should we reason about such games?
  - We should expect players to use strategies that are best responses to one another
- Example:
  - Three client game.
  - What is the best response to each strategy

So for player 1, there is no single strategy that is the best response to every strategy of player 2's. Player 2 does not have a dominant strategy either. Why?

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

Player 1's A is the best response to player 2's A

Player 1's B is the best response to player 2's B

Player 1's C is the best response to player 2'sC

Nash equilibrium

- Developed by John Nash in 1950.
  - Made famous in “A Beautiful Mind“
- <u>*Def*</u>:
  - For strategy *S* by player 1 and *T* by player 2, the pair (*S*,*T*) is a *Nash equilibrium* if *S* is a best response to *T*, and *T* is a best response to *S*
- More generally, at Nash equilibrium if *no player wants to unilaterally (<u>done only by one player</u>) deviate to an alternative strategy*

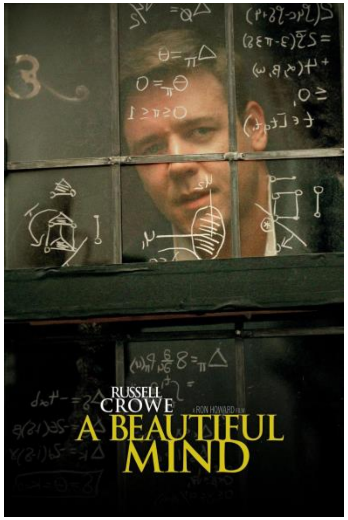

Example: Nash equilibrium

- Three client game
  - Suppose Firm 1 chooses A and Firm 2 also chooses A
  - These strategies are the best responses to each other - neither player wants to deviate
  - Thus (A,A) is a *Nash equilibrium*.
    - It is also unique – no other pair works

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

Example: Nash equilibrium

- Coordination game
  - Prepare a presentation with a partner
  - But don‘t know what software she will use
    - Incompatibility takes effort

	PowerPoint	Keynote
Powerpoint	1, 1	0, 0
Keynote	0, 0	1, 1

  - (PowerPoint,PowerPoint) is a Nash equilibrium
  - (Keynote,Keynote) is *also* a Nash equilibrium.
- Can have multiple Nash equilibria!

Formualtion: what is a game?

- Many of the motivating examples are in fact from actual games
  - Soccer penalty kick, board games
  - Model of course more widely applicable
- *Def:* A *game* consists of three things.
  - **(1)** A set of *players*
  - **(2)** Each player has set of options how to behave called *strategies*
  - **(3)** For each choice of strategies, each player receives a *payoff* from the game.

n-player game: formulation

- The players are indexed by i, where $i \in \{1, \ldots, n\}$
- Let S_i denote the set of strategies available to player *i* and let $s_i \in S_i$ denote an arbitrary member of this set
  - So, $(s_1, \ldots, s_n)$ denotes a combination of strategies
- Let $u_i(s_1, \ldots, s_n)$ be the payoff to player *i* if the players choose the strategies
- Thus, a n-player game is specified by the players' strategy space and utility functions $(s_1, \ldots, s_n)$

$$G = \{S_1, \ldots, S_n; u_1, \ldots, u_n\}$$

A dominant strategy

- When a strategy is best for a player no matter what strategy the other player uses, that strategy is said to
  - dominate all other strategies and
  - is called a dominant strategy.

$$u_i(s_1, \ldots, s_{i-1}, s'_i, s_{i+1}, \ldots, s_n) > u_i(s_1, \ldots, s_{i-1}, s''_i, s_{i+1}, \ldots, s_n)$$

Nash Equilibrium

In the n-player game $G = \{S_1, \ldots, S_n; u_1, \ldots, u_n\}$ the strategies $(s^*_1, \ldots, s^*_n)$ are a **Nash equilibrium** if, for each *i*, s^*_i is player *i*'s best response to the strategies by others, $(s^*_1, \ldots, s^*_{i-1}, s^*_{i+1}, \ldots, s^*_n)$

$$u_i(s^*_1, \ldots, s^*_{i-1}, s^*_i, s^*_{i+1}, \ldots, s^*_n)$$

$$> u_i(s^*_1, \ldots, s^*_{i-1}, s_i, s^*_{i+1}, \ldots, s^*_n)$$

for every feasible strategy $s_i \in S_i$; that is

$$s^*_i = \arg\max_{s_i \in S_i} u_i(s^*_1, \ldots, s^*_{i-1}, s_i, s^*_{i+1}, \ldots, s^*_n)$$

Important games

- Battle of the sexes
  - Which kind of movie to rent?

	Romance	Thriller
Romance	1, 2	0, 0
Thriller	0, 0	2, 1

  - Two equilibria, but which one will be played?
- Hard to predict the outcome
  - Depends on social conventions

Important games

- Stag Hunt
  - If hunters work together, they can catch a stag
  - On their own they can each catch a hare (rabbit)
  - If one hunter tries for a stag, he gets nothing

	Hunt Stag	Hunt Hare
Hunt Stag	4, 4	0, 3
Hunt Hare	3, 0	3, 3

- Two equilibria, but “riskier“ to hunt stag
  - What if other player hunts hare? Get nothing
  - Similar to prisoner‘s dilemma
    - Must trust other person to get best outcome!

Important games

- Hawk-Dove (or Game of Chicken) refers to a situation where there is a competition for a shared resource
  - Each player either aggressive (H) or passive (D)
  - If both passive, divide resources evenly
  - If both aggressive – war! Disastrous for both
  - Otherwise aggressor wins

	Dove	Hawk
Dove	3, 3	1, 5
Hawk	5, 1	0, 0

- Can model the foreign policy of countries
  - each player, in attempting to secure her best outcome, risks the worst.

Mixed strategies

- Do Nash equilibria always exist?
  - Matching Pennies game
  - Player 1 wins on same outcome, 2 on different

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

- Example of a *zero-sum* game
  - What one player gains, the other loses
  - E.g. Allied landing in Europe on June 6, 1944
- How would you play this game?

Mixed strategies

- You are *randomizing* your strategy
  - Instead of choosing H/T directly, choose a *probability* you will choose H.
- Player 1 commits to play H with some probability p
  - Similarly, player 2 plays H with probability q
- This is called a *mixed strategy*
  - As opposed to a *pure* strategy (e.g. p=0)
- What about the payoffs?

Mixed strategies

- Suppose player 1 evaluates pure strategies
  - Player 2 meanwhile chooses strategy q
  - If Player 1 chooses H, he gets a payoff of -1 with probability q and +1 with probability $1-q$
  - If Player 1 chooses T, he gets -1 with probability $1-q$ and +1 with probability q
- Is H or T more appealing to player 1?
  - Rank the *expected values*
  - Pick H: expect $(-1)(q) + (+1)(1-q) = 1-2q$
  - Pick T: expect $(+1)(q) + (-1)(1-q) = 2q - 1$

Mixed strategies

- *Def:* Nash equilibrium for mixed strategies
  - A pair of strategies (now probabilities) such that each is a best response to the other.
  - Thm: Nash proved that this *always* exists.
- In Matching Pennies, no Nash equilibrium can use a pure strategy (by examining the table):
  - Player 2 would have a unique best response which is a pure strategy
  - But this is not the best response for player 1...
- What is Player 1's best response to strategy *q*?
  - If $1-2q \neq 2q-1$ (play 1 chooses H or T), then a pure strategy (either H or T) is a unique best response to player 1.
  - This can't be part of a Nash equilibrium by the above
  - So must have $1-2q=2q-1$ in any Nash equilibrium
    - Which gives $q=1/2$. Similarly $p=1/2$ for Player 1.
    - This is a unique Nash equilibrium (check!)

Mixed strategies

- Intuitively, mixed strategies are used to make it harder for the opponent to predict what will be played
  - By setting q=1/2, Player 2 makes Player 1 *indifferent* between playing H or T.
- How do we interpret mixed equilibria?
  - In sports (or real games)
    - Players are indeed randomizing their actions
  - Competition for food among species
    - Individuals are hardwired play certain strategies
    - Mixed strategies are *proportions* within populations
    - Population as a *whole* is a mixed equilibrium
  - Nash equilibrium is an equilibrium in beliefs
    - If you believe other person will play a Nash equilibrium strategy, so will you.
    - It is self-reinforcing – an equilibrium

Mixed strategies: Examples

- American football
  - Offense can run with the ball, or pass forward

	Defend pass	Defend run
Pass	0, 0	10, -10
Run	5, -5	0, 0

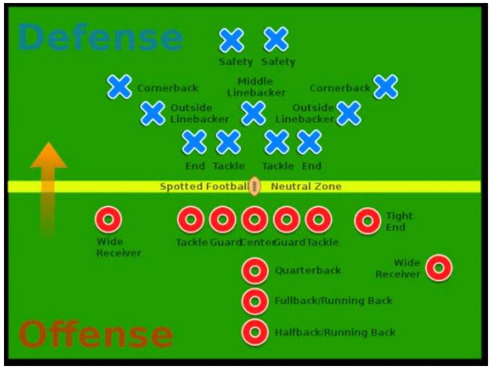

- What happens?
  - Suppose the defense defends against a pass with probability q
  - P: expect $(0)(q) + (10)(1\text{-}q) = 10\text{-}10q$
  - R: expect $(5)(q) + (0)(1\text{-}q) = 5q$
  - Offense is indifferent when $q=2/3$

Mixed strategies: Examples

- American football
  - Offense can run with the ball, or pass forward

	Defend pass	Defend run
Pass	0, 0	10, -10
Run	5, -5	0, 0

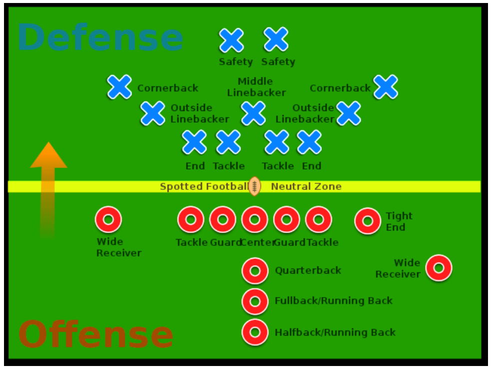

- What happens?
  - Suppose offense passes with probability p
  - Similarly, defense is indifferent when p=1/3
  - (1/3,2/3) is a Nash equilibrium
    - Expected payoff to offense: 10/3 (yard gain)

Mixed strategies: Examples

- Penalty-kick game
  - An economist analyzed 1,417 penalty kicks from five years of professional soccer matches among European clubs
  - The success rates of penalty kickers given the decision by both the goalkeeper and the kicker to kick or dive to the left or the right are as follows:

goalkeeper

kicker

	Defend left	Defend right
Left	58%	95%
Right	93%	70%

Left and right are from kicker's perspective

Mixed strategies: Examples

- Penalty-kick game
  - Soccer penalties have been studied extensively

	Defend left	Defend right
Left	0.58, -0.58	0.95, -0.95
Right	0.93, -0.93	0.70, -0.70

  - Suppose goalkeeper defends left with probability q
  - Kicker indifferent when
    - $(0.58)(q) + (0.95)(1-q) = (0.93)(q) + (0.70)(1-q)$
  - Get q =0.42. Similarly p=0.39
  - True values from data? q=0.42 , p=0.40 !!
    - The theory predicts reality very well

Palacios-Huerta, Ignacio. "Professionals play minimax." *The Review of Economic Studies* 70.2 (2003): 395-415.
http://pricetheory.uchicago.edu/levitt/Papers/ChiapporiGrosecloseLevitt2002.pdf
http://www.mikeshor.com/courses/gametheory/docs/topic4/mixedsoccer.html

Mixed strategies: Examples

- Penalty-kick game
  - More impressive is that the players' ability to randomize without patterns
    - Neither the player's past kicks nor his opponent's past behavior is useful in predicting the next kick
  - The overall chance of scoring on a penalty kick, in equilibrium is 80%
    - Why?

http://www2.owen.vanderbilt.edu/mike.shor/courses/game-theory/docs/lecture05/MixedSoccer.html

Pareto optimality

- Even playing best responses does not always reach a good outcome as a group
  - E.g. prisoner's dilemma
- Want to define a *socially good outcome*
- *Def*:
  - A choice of strategies is *Pareto optimal* if no other choice of strategies gives all players a payoff at least as high, and at least one player gets a strictly higher payoff
  - In other words, no player can gain higher payoff without sacrificing others' payoffs
- Note: *Everyone* must do at least as well

Social optimality

- *Def:*
  - A choice of strategies is a *social welfare maximizer* (or *socially optimal*) if it maximizes the sum of the players‘ payoffs.
- Example:
  - The unique Nash equilibrium in this game is socially optimal

	Presentation	Exam
Presentation	98,98	94,96
Exam	96,94	92,92

ISP COMPETITION

Competition between ISPs

- How does the outcome of competition among two Internet Service Providers (ISPs) in a competitive market?

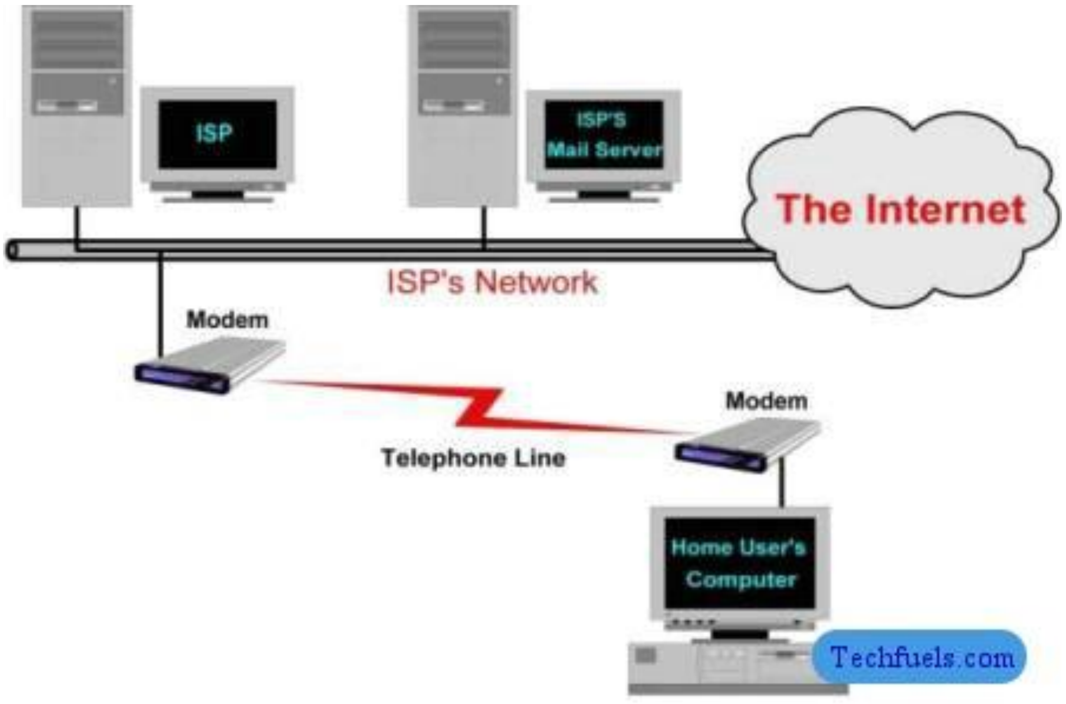

Cournot Model of Duopoly

- *n* number of ISPs offer their service (Bandwidth) to the customers
- The ***cost*** to ISP *i* of producing q_i units of the Bandwidth is $C_i(q_i)$, where C_i is an increasing function
- The Bandwidth is sold at a single ***price*** $P(Q)$, where Q is the total output and P is a decreasing function unless it is already zero.
- If the output of each ISP *i* is q_i, then the price is $P(q_1 + \dots + q_n)$, so ***revenue*** of *i* is $q_i P(q_1 + \dots + q_n)$.

Cournot Model of Duopoly

- Thus ISP i's profit, equal to its revenue minus its cost, is

  $$F_i(q_1, \ldots, q_n) = q_i P(q_1 + \ldots + q_n) - C_i(q_i)$$

- The ISP Games:
  - Players: the ISPs.
  - Actions: the set of its possible outputs $q_1, \ldots, q_n$
  - Utilities: ISPs preferences are represented by their profits $F_1, \ldots, F_n$

Cournot Model of Duopoly

- With certain forms of C_i and P a Nash equilibrium can be computed
  - Suppose there are two ISPs, each ISP's cost function is the same, given by

$C_i(q_i) = cq_i$ for all q_i, where c (>=0) is a constant unit cost

  - The inverse demand function is linear where it is positive, given by

$$P(Q) = a - Q \quad \text{if } Q <= a$$

$$P(Q) = 0 \quad \text{if } Q > a, \quad \text{where } a > 0$$

Cournot Model of Duopoly

- Assume that $c < a$, so that there is some value of total output Q for which the market price P(Q) is greater than the ISPs' common unit cost c
  - If c were to exceed a, there would be no output for the ISPs at which they could make any profit
- ISP *1* s' profit (the same case for ISP *2*):

$$F_1(q_1, q_2) = q_1(P(q_1 + q_2) - c)$$

$$= q_1(a - c - (q_1 + q_2))$$

Cournot Model of Duopoly

To find the Nash equilibrium, recall for every feasible strategy $s_i \in S_i$ there is

$$s^*_i = \arg\max_{s_i \in S_i} u_i(s^*_1, \ldots, s^*_{i-1}, s_i, s^*_{i+1}, \ldots, s^*_n)$$

Thus $$q_1^* = \arg\max_{0 \le q_1 < \infty} F_1(q_1,\ q_2^*)$$

$$= \arg\max_{0 \le q_1 < \infty}\ q_1(a - c - (q_1 + q_2^*))$$

The first-order condition for the optimization problem yields (assume $q_1 < a - c$):

$$q_1^* = \frac{1}{2}(a - q_2^* - c)$$

The same case for q_2: $$q_2^* = \frac{1}{2}(a - q_1^* - c)$$

Cournot Model of Duopoly

Solving this pair of equations yields:

$$q_1^* = q_2^* = \frac{a - c}{3},$$

which is indeed smaller than *a-c and the profit in the equilibrium is*

$$F_1^* = F_2^* = \frac{a - c}{3}\left(a - \frac{a - c}{3} - c\right) = \frac{2(a - c)^2}{9},$$

Solution by Best Response

- ISP 1's best response
- ISP 2's best response

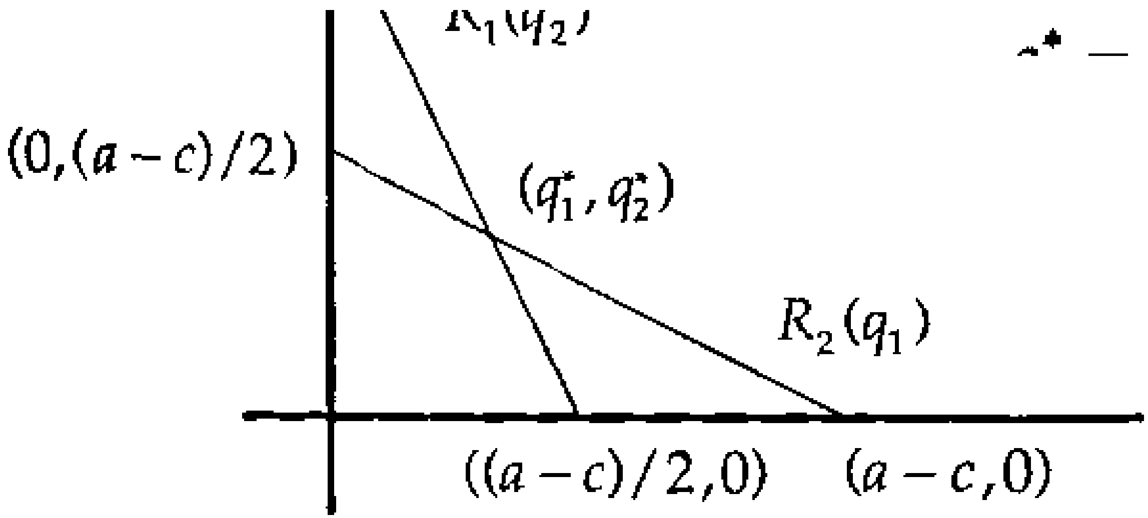

Figure 1.2.1.

These two best response functions intersect only once at the equilibrium quantity pair (p_1^*, p_2^*)

wn in Figure 1.2.1, these two best-response functions
ly once, at the equilibrium quantity pair (q_1^*, q_2^*).

An analysis of the result

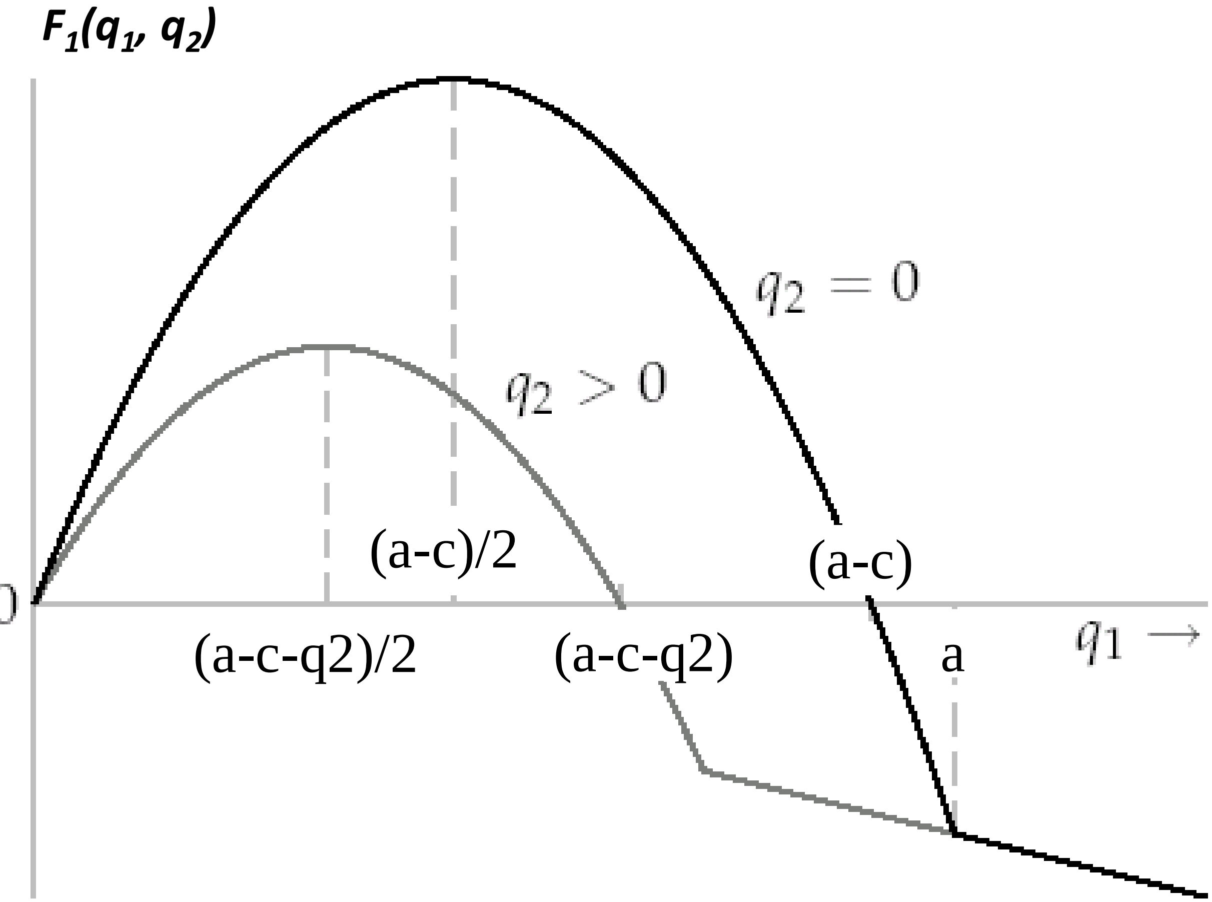

- Each ISP would like to be a monopolist, (the right cure)
- In equilibrium, the output of ISP 1 reduces to $q_1=(a-c)/3$ (the left curve)

An analysis of the result

- Given two ISPs, the aggregate profits for the *duopoly* would be maximized by setting $q_1 = q_2 = q_m/2 = (a-c)/4$ the half of the monopoly quantity
- However, because q_1 or q_2 is low, each ISP has an incentive to deviate
  - in order to increase their individual profits
- As the aggregate quantity goes higher, the price is lower; the temptation to increase output is reduced
- In the equilibrium, $q_1 = q_2 = (a-c)/3$

HUMAN COMPUTATION

Human Computation

- There is a lot of things that **human can easy do** that **computers can not yet do**
  - Speech recognition
  - Natural language understanding
  - Computer graphics
  - …

Games with a Purpose

- It combine **computation** with **game**
- People spend a lot of time playing games
- It makes Human Computation more efficient
- There are a lot of GWAP systems has been created
  - e.g. ESP (Extra Sensory Perception) game and Google Image Labeler)

What is the ESP game?

Alice

agreement on “flower”

shoe

rocks

flower

Bob

flower

What is the ESP game?

- it is efficient
  - 200,000+ players have contributed 50+ million labels
  - each player plays for a total of 91 minutes
  - 233 labels/human/hour (i.e. one label every 15 seconds)
  - Google bought a license to create its own version of the game in 2006 to solve their *Image Retrieval* problem

Analysis of the ESP game

- Can you model/analyze the EPS game using what we have learned?
- Suggest to read:
  - Von Ahn, Luis. "Games with a purpose." Computer 39.6 (2006): 92-94.
  - Weber, Ingmar, Stephen Robertson, and Milan Vojnovic. "Rethinking the ESP game." Proc. of 27th intl. conf. on Human factors in Computing Systems, ser. CHI. Vol. 9. 2008.

References and Further Readings

- I. E. Sutherland. 1968. A futures market in computer time. Commun. ACM 11, 6 (June 1968), 449-451. DOI=10.1145/363347.363396 http://doi.acm.org/10.1145/363347.363396
- Orna Agmon Ben-Yehuda, Muli Ben-Yehuda, Assaf Schuster, and Dan Tsafrir. 2011. Deconstructing Amazon EC2 Spot Instance Pricing. In Proceedings of the 2011 IEEE Third International Conference on Cloud Computing Technology and Science (CLOUDCOM '11)
- Martin J. Osborne, An Introduction to Game Theory

References and Further Readings

- Networks, Crowds, and Markets. Reasoning about a Highly Connected World. by David Easley and Jon Kleinberg. Cambridge University Press 2010
- D. Braess, Über ein Paradoxon aus der Verkehrsplanung. Unternehmensforschung 12, 258–268 (1969)

English Translation: http://homepage.rub.de/Dietrich.Braess/Paradox-BNW.pdf

- Hardin, G. (1968). "The Tragedy of the Commons". Science 162 (3859): 1243–1248.
- Ignacio Palacios-Huerta, "Professionals Play Minimax," Review of Economic Studies, vol. 70(2), April 2003, pp. 395-415.