UC Berkeley
Haas School of Business
Economic Analysis for Business Decisions
(EWMBA 201A)
Fall 2018

Module III

Oligopoly

- The Cournot and Stackelberg models (PR 12.2)
- The Bertrand model (PR 12.3)

Game theory

- Gaming and strategic decisions (PR 13.1-2)
- Nash equilibrium (PR 13.3)

Oligopoly

(preface to game theory)

Oligopoly
(preface to game theory)

- Another form of market structure is **oligopoly** – a market in which only a few firms compete with one another, and entry of new firms is impeded.

- The situation is known as the Cournot model after Antoine Augustin Cournot, a French economist, philosopher and mathematician (1801-1877).

- In the basic example, a single good is produced by two firms (the industry is a "duopoly").

Cournot's oligopoly model (1838)

– A single good is produced by two firms (the industry is a "duopoly").

– The cost for firm $i = 1, 2$ for producing q_i units of the good is given by $c_i q_i$ ("unit cost" is constant equal to $c_i > 0$).

– If the firms' total output is $Q = q_1 + q_2$ then the market price is

$$P = A - Q$$

if $A \geq Q$ and zero otherwise (linear inverse demand function). We also assume that $A > c$.

The inverse demand function

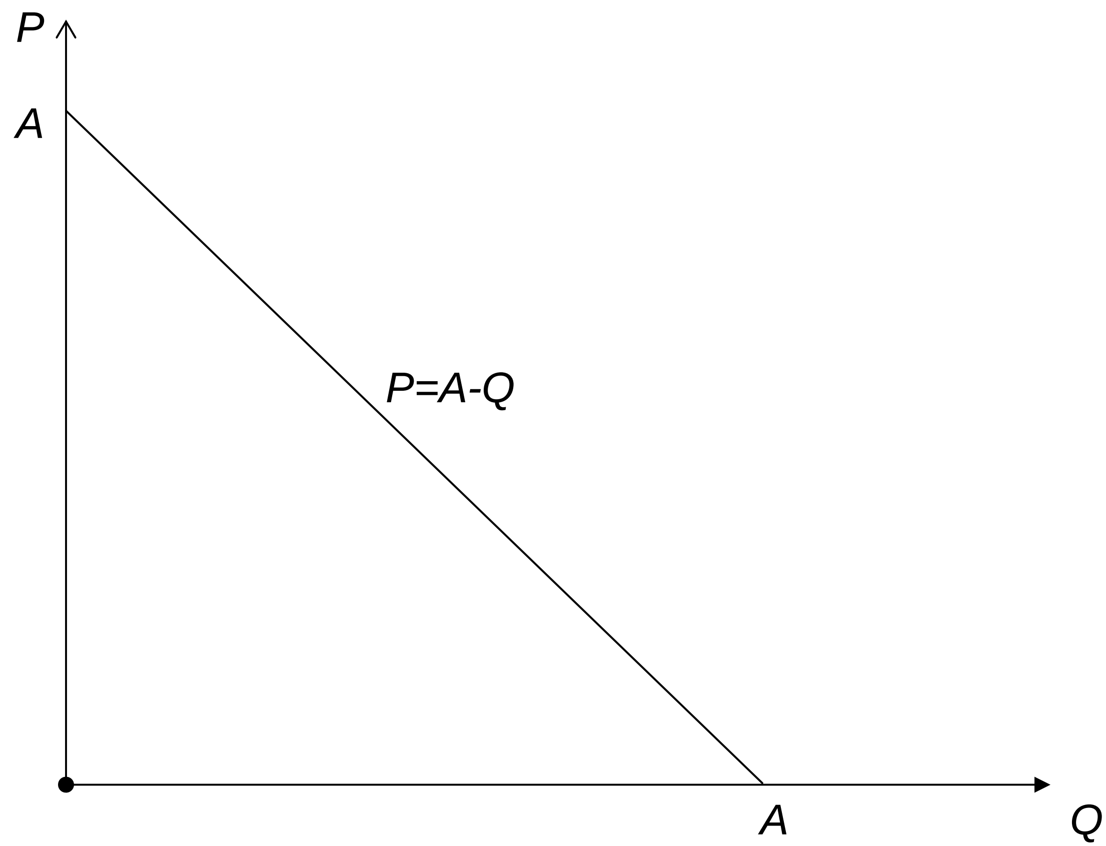

To find the Nash equilibria of the Cournot's game, we can use the procedures based on the firms' best response functions.

But first we need the firms payoffs (profits):

$$\begin{aligned}\pi_1 &= Pq_1 - c_1q_1\\ &= (A-Q)q_1 - c_1q_1\\ &= (A-q_1-q_2)q_1 - c_1q_1\\ &= (A-q_1-q_2-c_1)q_1\end{aligned}$$

and similarly,

$$\pi_2 = (A-q_1-q_2-c_2)q_2$$

Firm 1's profit as a function of its output (given firm 2's output)

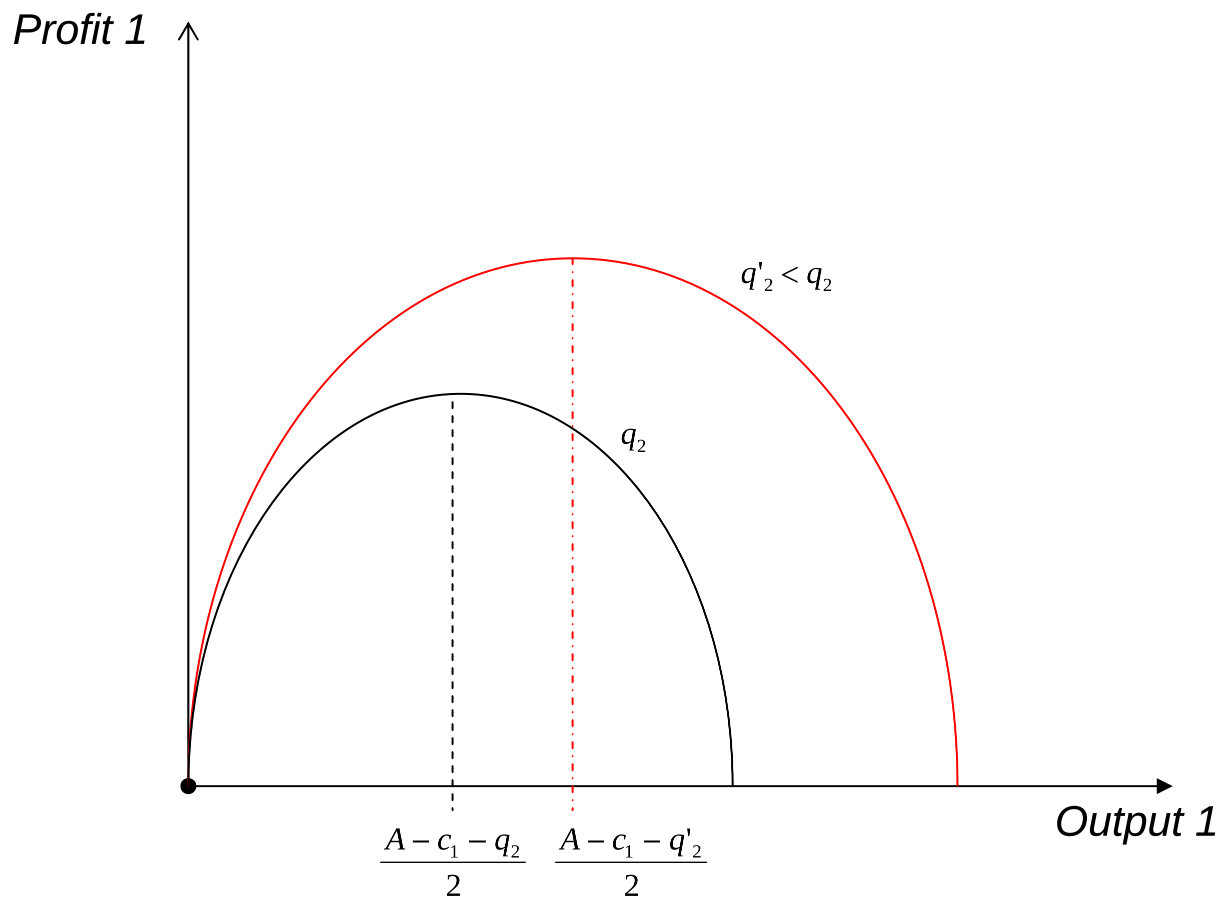

To find firm 1's best response to any given output q_2 of firm 2, we need to study firm 1's profit as a function of its output q_1 for given values of q_2.

Using calculus, we set the derivative of firm 1's profit with respect to q_1 equal to zero and solve for q_1:

$$q_1 = \frac{1}{2}(A - q_2 - c_1).$$

We conclude that the best response of firm 1 to the output q_2 of firm 2 depends on the values of q_2 and c_1.

Because firm 2's cost function is $c_2 \neq c_1$, its best response function is given by

$$q_2 = \frac{1}{2}(A - q_1 - c_2).$$

A Nash equilibrium of the Cournot's game is a pair (q_1^*, q_2^*) of outputs such that q_1^* is a best response to q_2^* and q_2^* is a best response to q_1^*.

From the figure below, we see that there is exactly one such pair of outputs

$$q_1^* = \frac{A+c_2-2c_1}{3} \quad \text{and} \quad q_2^* = \frac{A+c_1-2c_2}{3}$$

which is the solution to the two equations above.

The best response functions in the Cournot's duopoly game

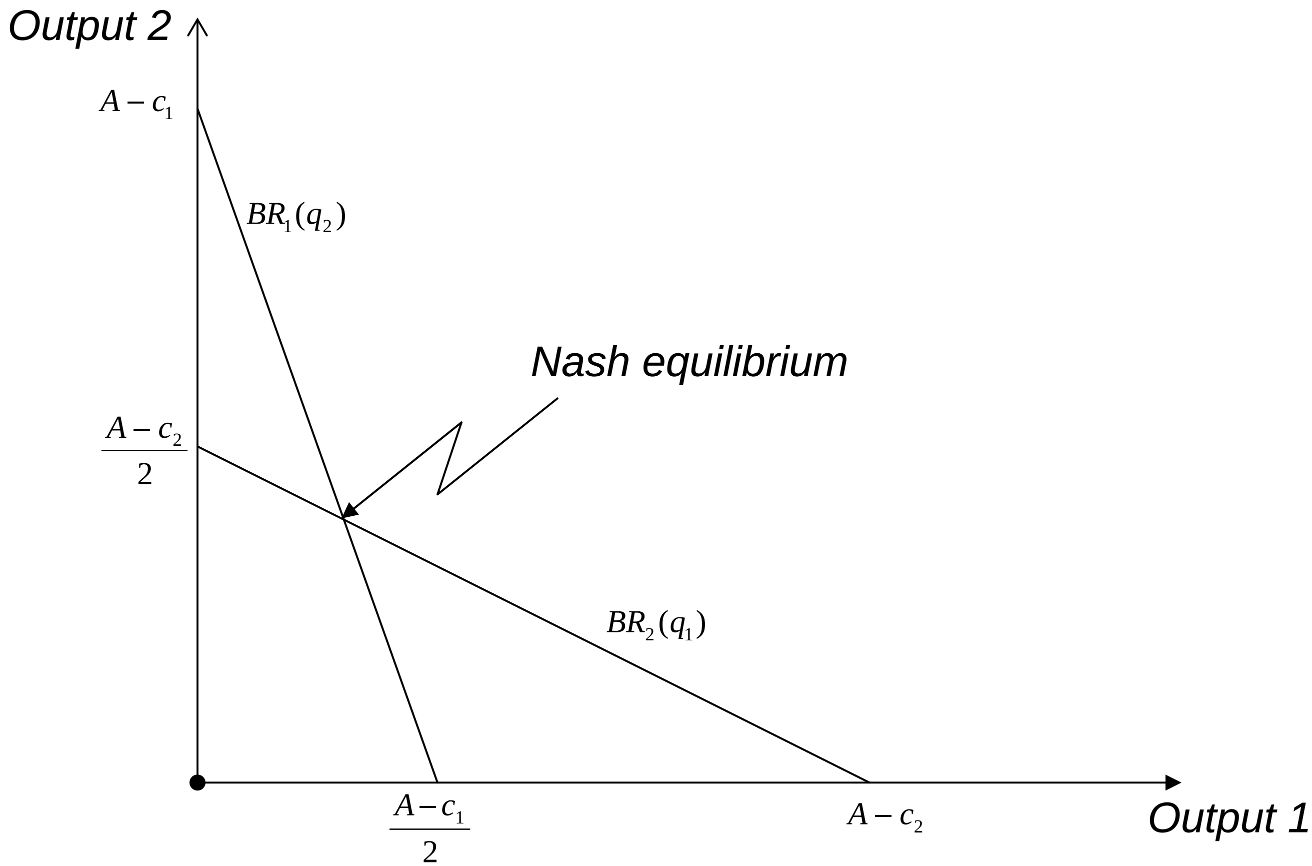

Nash equilibrium comparative statics
(a decrease in the cost of firm 2)

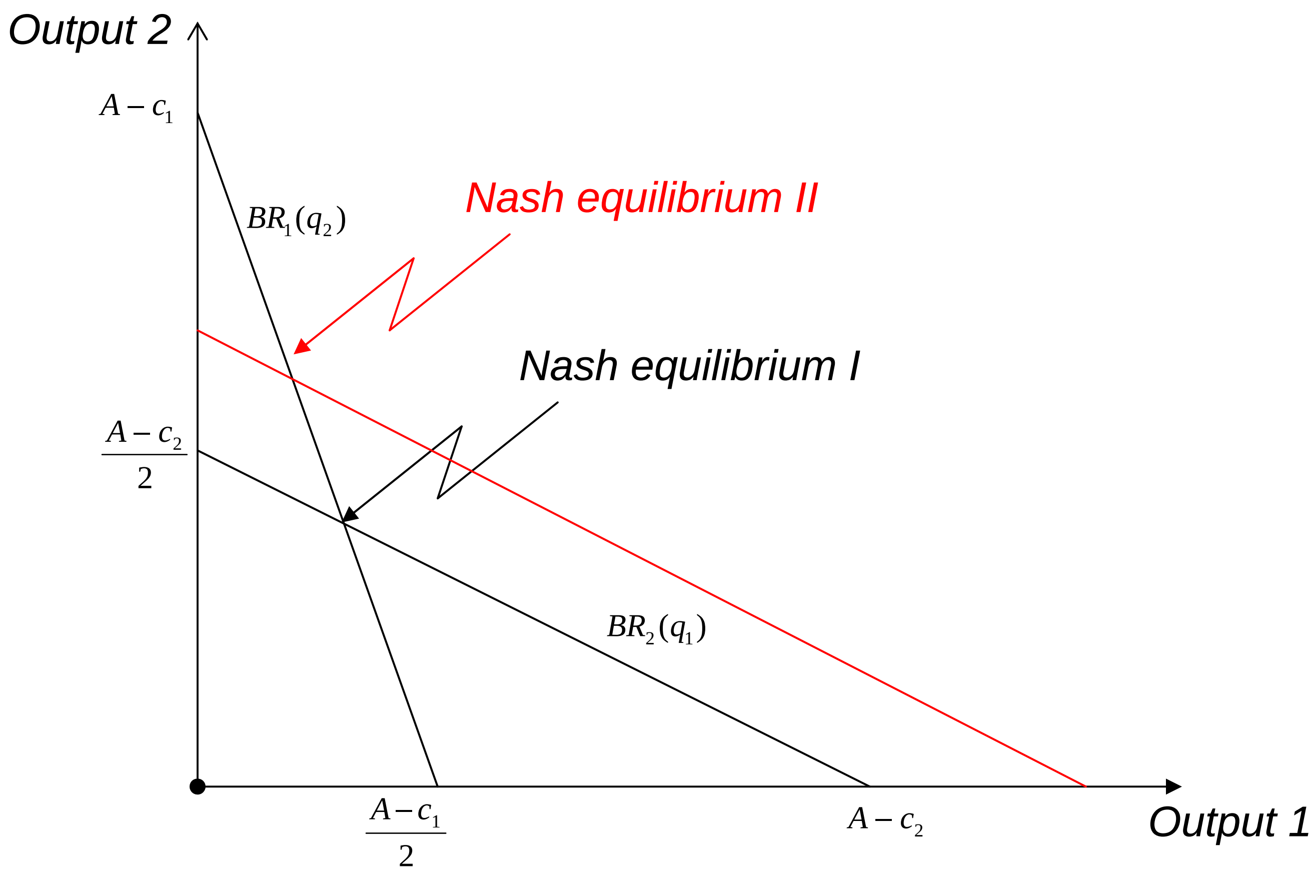

A question: what happens when consumers are willing to pay more (*A* increases)?

In summary, this simple Cournot's duopoly game has a unique Nash equilibrium.

Two economically important properties of the Nash equilibrium are (to economic regulatory agencies):

[1] The relation between the firms' equilibrium profits and the profit they could make if they act collusively.

[2] The relation between the equilibrium profits and the number of firms.

[1] Collusive outcomes: in the Cournot's duopoly game, there is a pair of outputs at which *both* firms' profits exceed their levels in a Nash equilibrium.

[2] Competition: The price at the Nash equilibrium if the two firms have the *same* unit cost $c_1 = c_2 = c$ is given by

$$
\begin{aligned}
P^* &= A - q_1^* - q_2^* \\
&= \frac{1}{3}(A + 2c)
\end{aligned}
$$

which is above the unit cost c. But as the number of firm increases, the equilibrium price deceases, approaching c (zero profits!).

Stackelberg's duopoly model (1934)

How do the conclusions of the Cournot's duopoly game change when the firms move sequentially? Is a firm better off moving before or after the other firm?

Suppose that $c_1 = c_2 = c$ and that firm 1 moves at the start of the game. We may use backward induction to find the subgame perfect equilibrium.

– First, for *any* output q_1 of firm 1, we find the output q_2 of firm 2 that maximizes its profit. Next, we find the output q_1 of firm 1 that maximizes its profit, *given the strategy* of firm 2.

Firm 2

Since firm 2 moves after firm 1, a strategy of firm 2 is a *function* that associate an output q_2 for firm 2 for each possible output q_1 of firm 1.

We found that under the assumptions of the Cournot's duopoly game Firm 2 has a unique best response to each output q_1 of firm 1, given by

$$q_2 = \frac{1}{2}(A - q_1 - c)$$

(Recall that $c_1 = c_2 = c$).

Firm 1

Firm 1's strategy is the output q_1 the maximizes

$$\pi_1 = (A - q_1 - q_2 - c)q_1 \quad \text{subject to} \quad q_2 = \tfrac{1}{2}(A - q_1 - c)$$

Thus, firm 1 maximizes

$$\pi_1 = (A - q_1 - (\frac{1}{2}(A - q_1 - c)) - c)q_1 = \frac{1}{2}q_1(A - q_1 - c).$$

This function is quadratic in q_1 that is zero when $q_1 = 0$ and when $q_1 = A - c$. Thus its maximizer is

$$q_1^* = \frac{1}{2}(A - c).$$

Firm 1's (first-mover) profit in Stackelberg's duopoly game

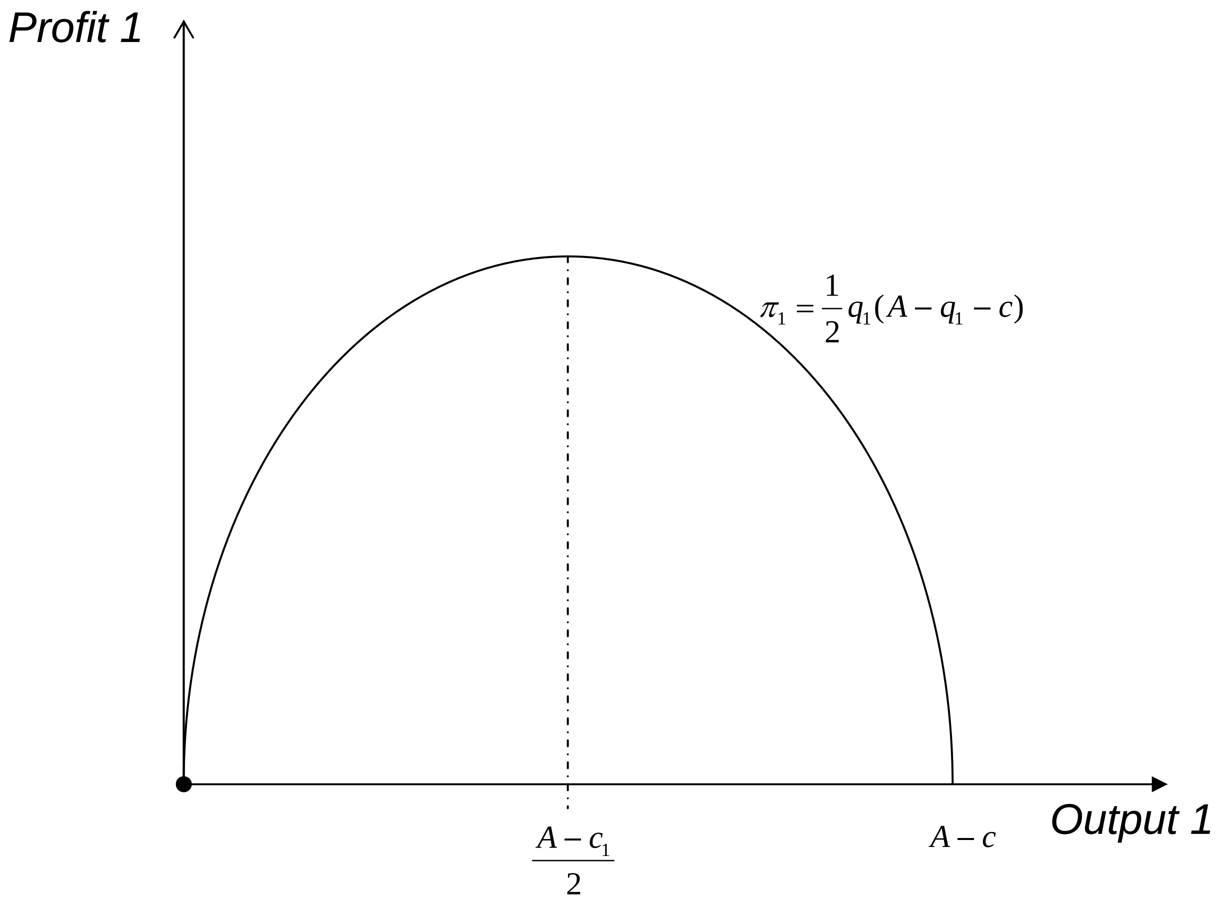

We conclude that Stackelberg's duopoly game has a unique subgame perfect equilibrium, in which firm 1's strategy is the output

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

and firm 2's output is

$$\begin{aligned} q_2^* &= \frac{1}{2}(A - q_1^* - c) \\ &= \frac{1}{2}(A - \frac{1}{2}(A - c) - c) \\ &= \frac{1}{4}(A - c). \end{aligned}$$

By contrast, in the unique Nash equilibrium of the Cournot's duopoly game under the same assumptions ($c_1 = c_2 = c$), each firm produces $\frac{1}{3}(A - c)$.

The subgame perfect equilibrium of Stackelberg's duopoly game

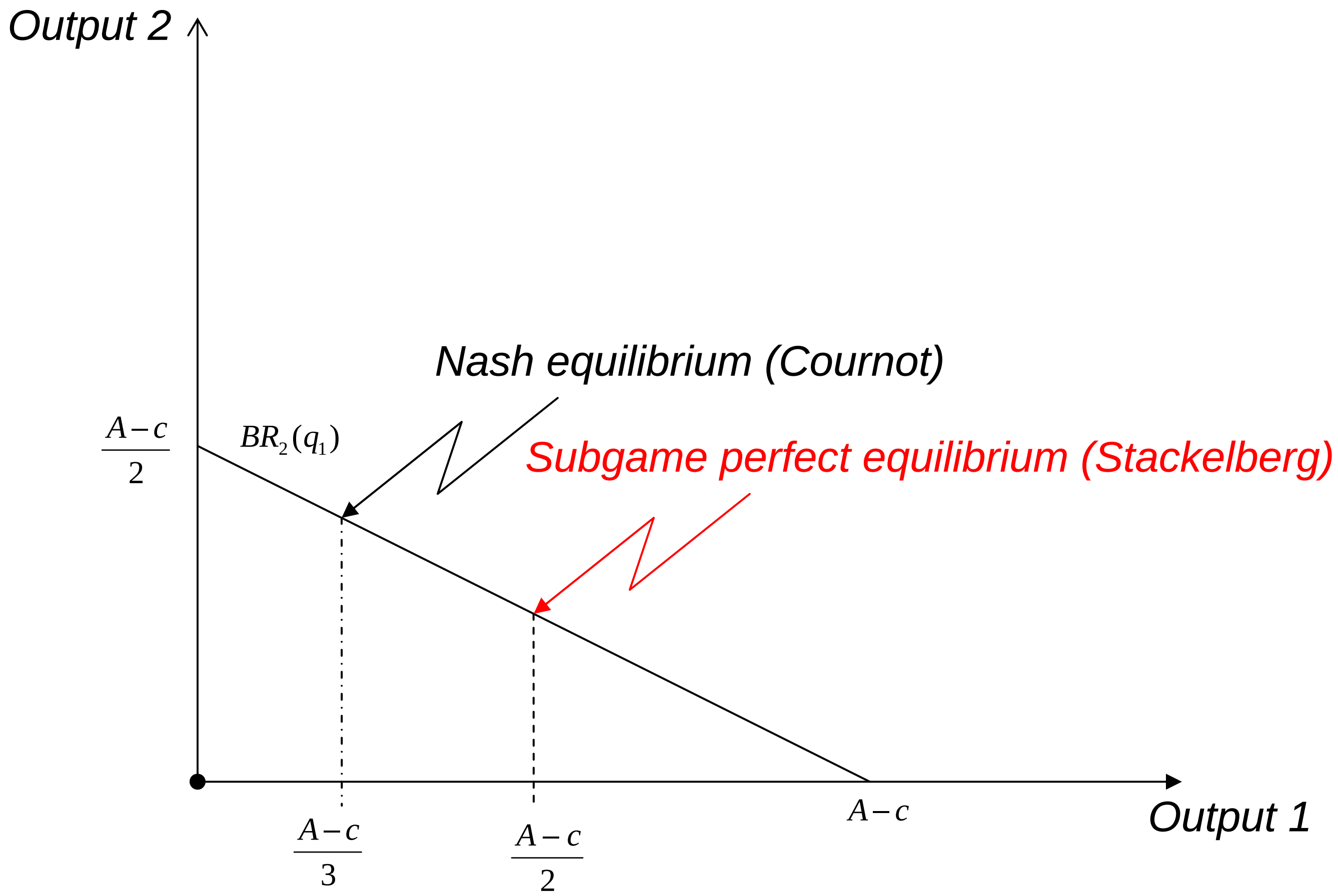

Bertrand's oligopoly model (1883)

In Cournot's game, each firm chooses an output, and the price is determined by the market demand in relation to the total output produced.

An alternative model, suggested by Bertrand, assumes that each firm chooses a price, and produces enough output to meet the demand it faces, given the prices chosen by *all* the firms.

$\Longrightarrow$ As we shell see, some of the answers it gives are different from the answers of Cournot.

Suppose again that there are two firms (the industry is a "duopoly") and that the cost for firm $i = 1, 2$ for producing q_i units of the good is given by cq_i (equal constant "unit cost").

Assume that the demand function (rather than the inverse demand function as we did for the Cournot's game) is

$$D(p) = A - p$$

for $A \geq p$ and zero otherwise, and that $A > c$ (the demand function in PR 12.3 is different).

Because the cost of producing each until is the same, equal to c, firm i makes the profit of $p_i - c$ on every unit it sells. Thus its profit is

$$\pi_i = \begin{cases} (p_i - c)(A - p_i) & \text{if } \ p_i < p_j \\ \frac{1}{2}(p_i - c)(A - p_i) & \text{if } \ p_i = p_j \\ 0 & \text{if } \ p_i > p_j \end{cases}$$

where j is the other firm.

In Bertrand's game we can easily argue as follows: $(p_1, p_2) = (c, c)$ is the unique Nash equilibrium.

Using intuition,

- If one firm charges the price c, then the other firm can do no better than charge the price c.

- If $p_1 > c$ and $p_2 > c$, then each firm i can increase its profit by lowering its price p_i slightly below p_j.

$\Longrightarrow$ In Cournot's game, the market price decreases toward c as the number of firms increases, whereas in Bertrand's game it is c (so profits are zero) even if there are only two firms (but the price remains c when the number of firm increases).

Avoiding the Bertrand trap

If you are in a situation satisfying the following assumptions, then you will end up in a Bertrand trap (zero profits):

[1] Homogenous products

[2] Consumers know all firm prices

[3] No switching costs

[4] No cost advantages

[5] No capacity constraints

[6] No future considerations

Game theory

Game theory

- Game theory is about what happens when decision makers (spouses, workers, managers, presidents) interact.

- In the past fifty years, game theory has gradually became a standard language in economics.

- The power of game theory is its <u>generality</u> and (mathematical) <u>precision</u>.

- Because game theory is rich and crisp, it could unify many parts of social science.

- The spread of game theory outside of economics has suffered because of the misconception that it requires a lot of fancy math.

- Game theory is also a natural tool for understanding complex social and economic phenomena in the real world.

The paternity of game theory

 

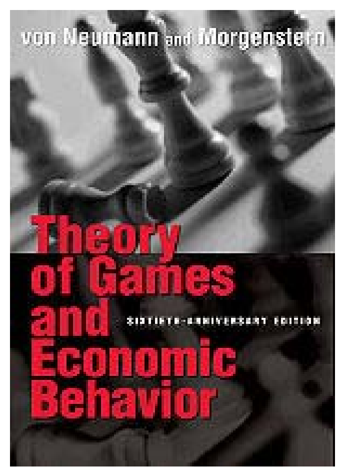

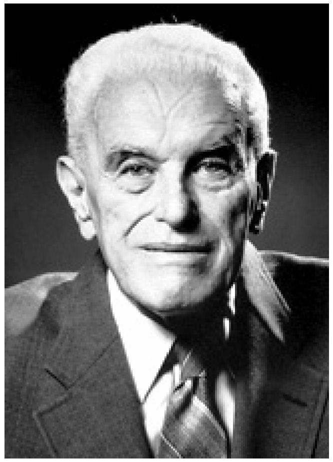
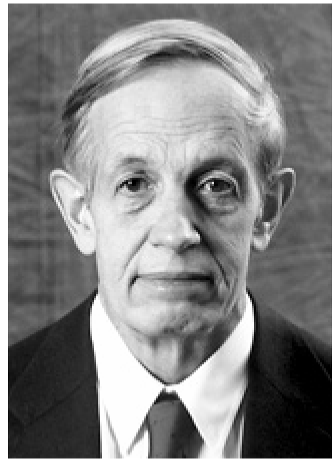

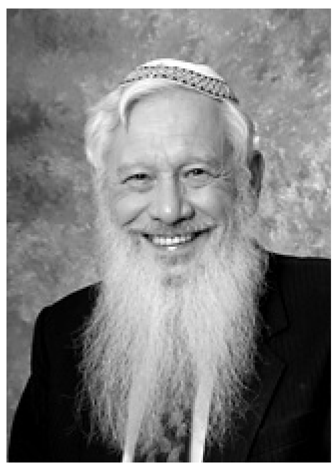
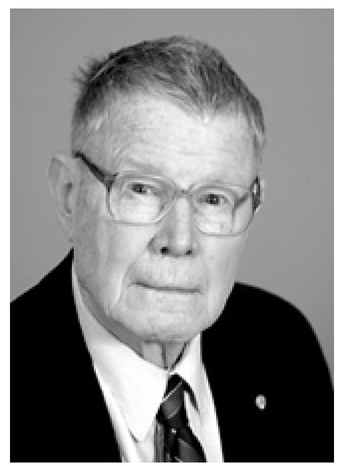

What is game theory good for?

Q Is game theory meant to predict what decision makers do, to give them advice, or what?

A The tools of analytical game theory are used to predict, postdict (explain), and prescribe.

Remember: even if game theory is not always accurate, descriptive failure is prescriptive opportunity!

What is game theory good for?

Q Is game theory meant to predict what decision makers do, to give them advice, or what?

A The tools of analytical game theory are used to predict, postdict (explain), and prescribe.

Remember: even if game theory is not always accurate, descriptive failure is prescriptive opportunity!

As Milton Friedman said famously observed “theories do not have to be realistic to be useful.” A theory can be *useful* in three ways:

$A.$ descriptive (how people actually choose)

$B.$ prescriptive (as a practical aid to choice)

$C.$ normative (how people ought to choose)

Aumann (1987):

> “*Game theory is a sort of umbrella or ‘unified field’ theory for the rational side of social science, where ‘social’ is interpreted broadly, to include human as well as non-human players (computers, animals, plants).*”

Game theory in practice

Farhan Zaidi, the General Manager of the LA Dodgers (PHD in economics from UC Berkeley), and the person Billy Beane called “absolutely brilliant.”

Three examples

Example I: Hotelling's electoral competition game

- There are two candidates and a continuum of voters, each with a favorite position on the interval $[0, 1]$.

- Each voter's distaste for any position is given by the distance between the position and her favorite position.

- A candidate attracts the votes off all citizens whose favorite positions are closer to her position.

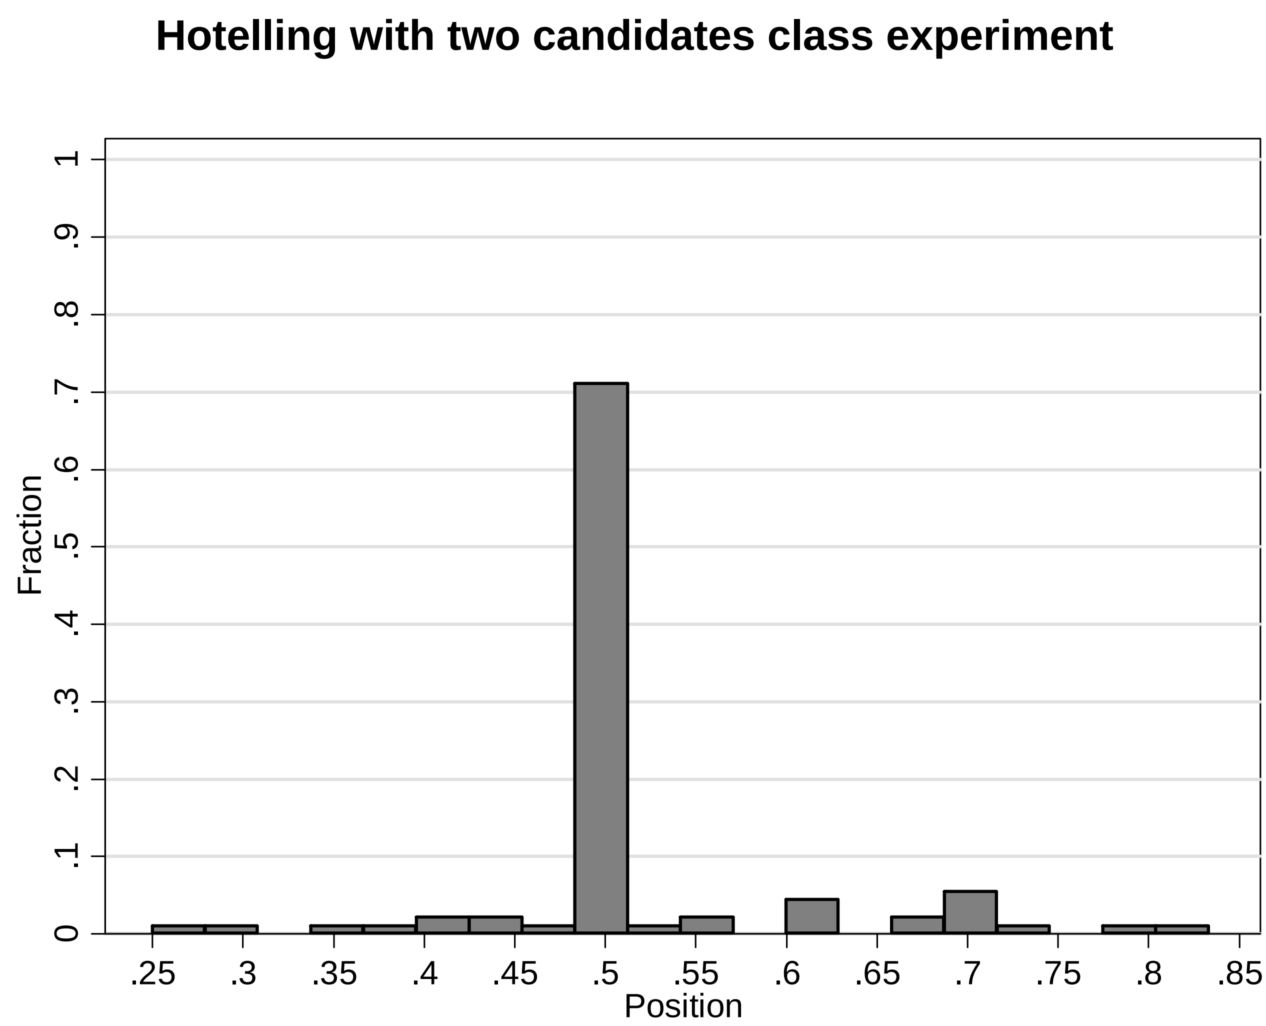
Hotelling with two candidates class experiment
Fraction
0
.1
.2
.3
.4
.5
.6
.7
.8
.9
1
.25
.3
.35
.4
.45
.5
.55
.6
.65
.7
.75
.8
.85
Position

Hotelling with three candidates class experiment

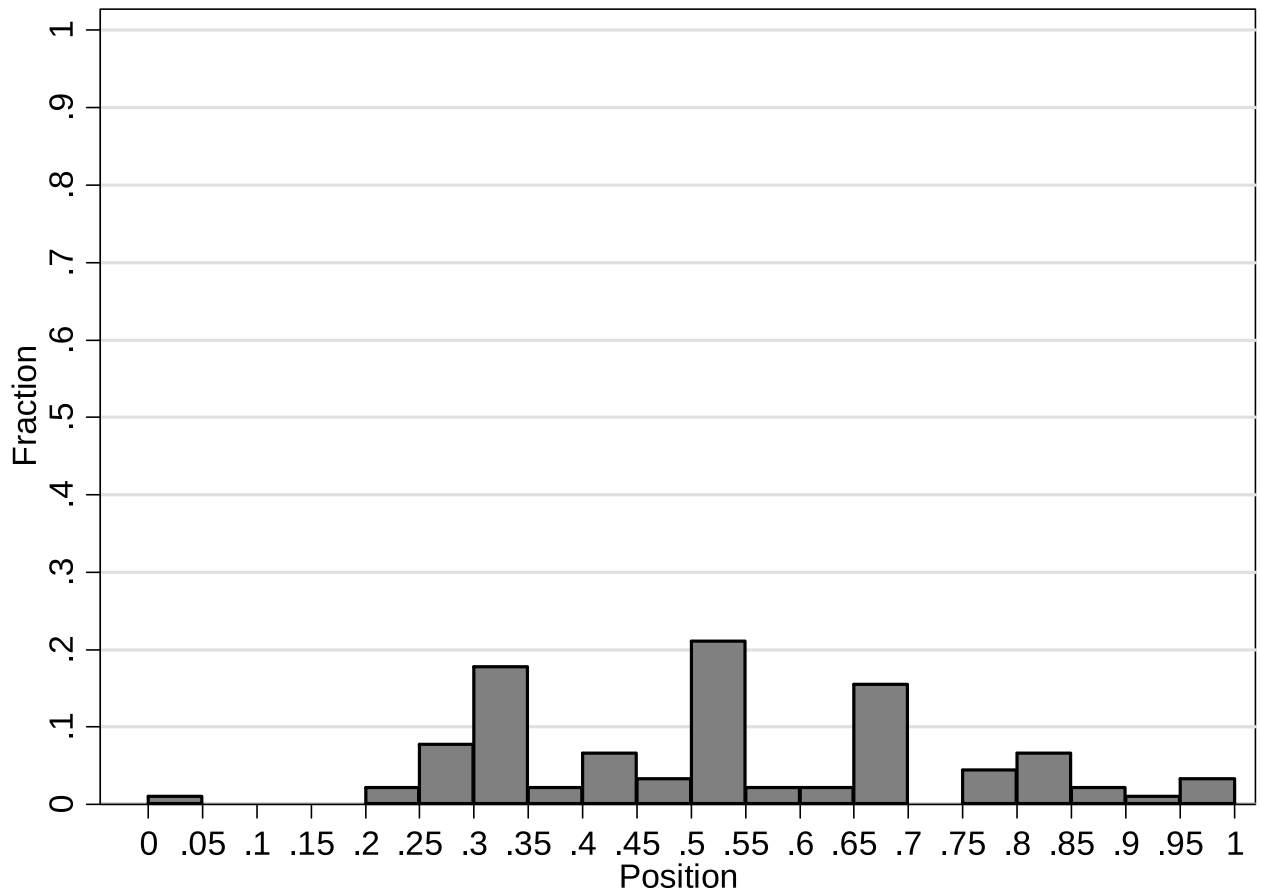

Example II: Keynes's beauty contest game

– Simultaneously, everyone choose a number (integer) in the interval $[0, 100]$.

– The person whose number is closest to $2/3$ of the average number wins a fixed prize.

John Maynard Keynes (1936):

> "*It is not a case of choosing those [faces] that, to the best of one's judgment, are really the prettiest, nor even those that average opinion genuinely thinks the prettiest. We have reached the third degree where we devote our intelligences to anticipating what average opinion expects the average opinion to be. And there are some, I believe, who practice the fourth, fifth and higher degrees.*"

$\Longrightarrow$ self-fulfilling price bubbles!

Beauty contest results

	Portfolio Managers	Economics PhDs	CEOs	Caltech students	Caltech trustees
Mean	24.3	27.4	37.8	21.9	42.6
Median	24.4	30.0	36.5	23.0	40.0
Fraction choosing zero	7.7%	12.5%	10.0%	7.4%	2.7%

	Germany	Singapore	UCLA	Wharton	High school (US)
Mean	36.7	46.1	42.3	37.9	32.4
Median	33.0	50.0	40.5	35.0	28.0
Fraction choosing zero	3.0%	2.0%	0.0%	0.0%	3.8%

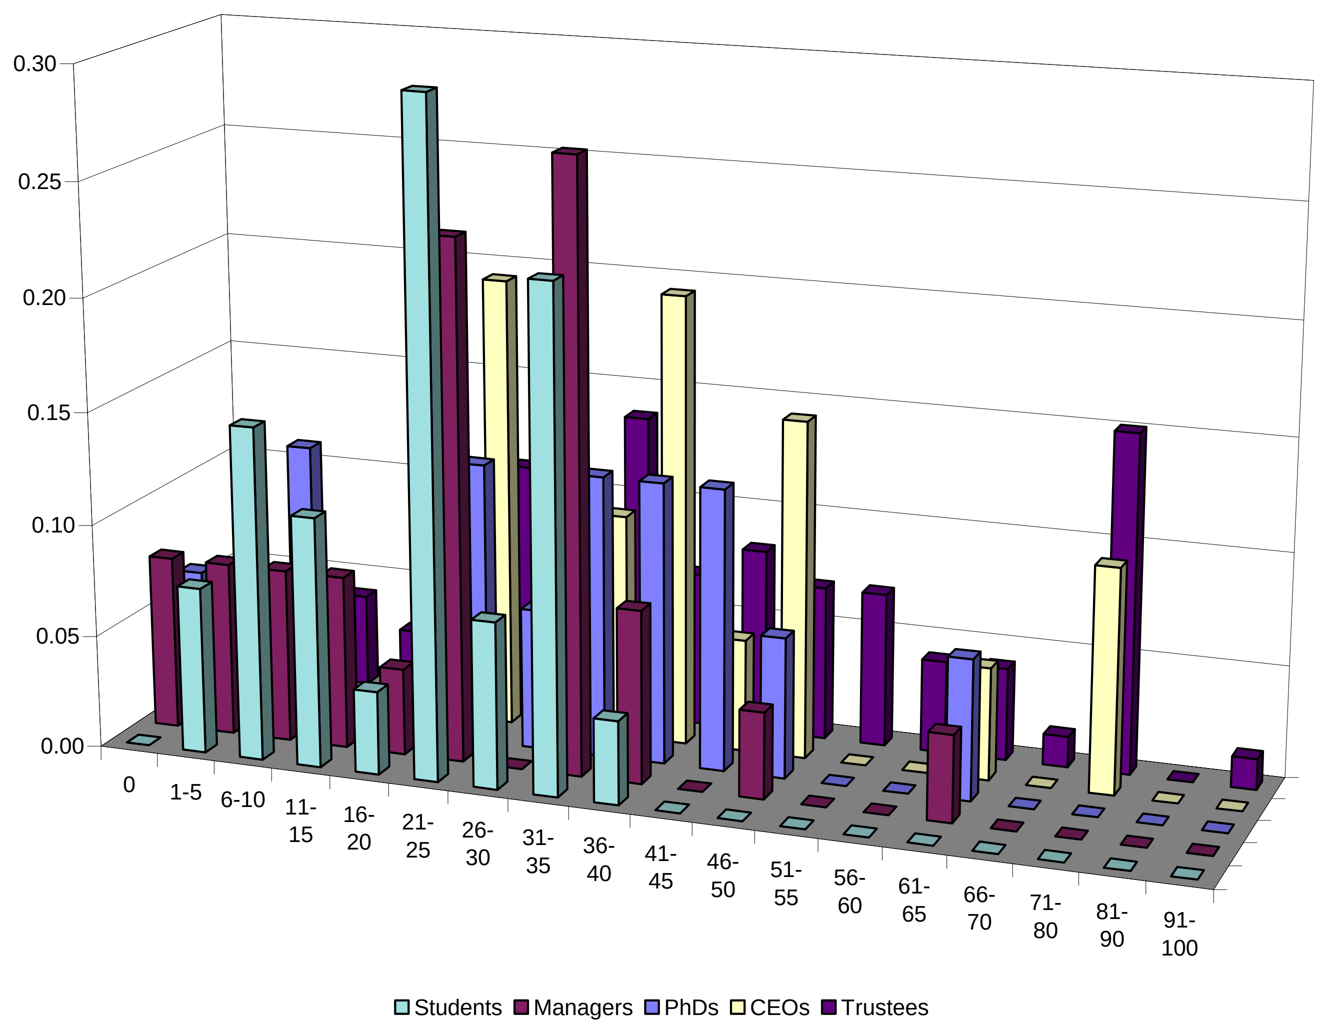
0.30
0.25
0.20
0.15
0.10
0.05
0.00
0
1-5
6-10
11-15
16-20
21-25
26-30
31-35
36-40
41-45
46-50
51-55
56-60
61-65
66-70
71-80
81-90
91-100
Students
Managers
PhDs
CEOs
Trustees

Example III: the centipede game (graphically resembles a centipede insect)

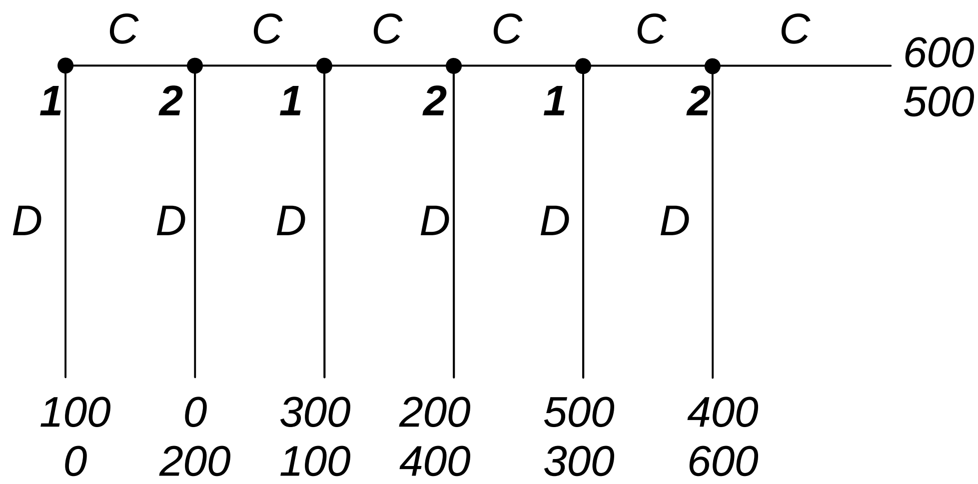

The centipede game class experiment

Down	*0.311*
Continue, Down	*0.311*
Continue, Continue, Down	*0.267*
Continue, Continue, Continue	*0.111*

Eye movements can tell us a lot about how people play this game (and others).

Food for thought

LUPI

Many players simultaneously chose an integer between 1 and 99,999. Whoever chooses the lowest unique positive integer (LUPI) wins.

Question What does an equilibrium model of behavior predict in this game?

The field version of LUPI, called Limbo, was introduced by the government-owned Swedish gambling monopoly Svenska Spel. Despite its complexity, there is a surprising degree of convergence toward equilibrium.

Morra

A two-player game in which each player simultaneously hold either one or two fingers and each guesses the total number of fingers held up.

If exactly one player guesses correctly, then the other player pays her the amount of her guess.

Question Model the situation as a strategic game and describe the equilibrium model of behavior predict in this game.

The game was played in ancient Rome, where it was known as "micatio."

Maximal game
(sealed-bid second-price auction)

Two bidders, each of whom privately observes a signal X_i that is independent and identically distributed (i.i.d.) from a uniform distribution on $[0, 10]$.

Let $X^{\mathsf{max}} = \mathsf{max}\{X_1, X_2\}$ and assume the ex-post common value to the bidders is X^{max}.

Bidders bid in a sealed-bid second-price auction where the highest bidder wins, earns the common value X^{max} and pays the second highest bid.

Types of games

We study four groups of game theoretic models:

I strategic games

II extensive games (with perfect and imperfect information)

III repeated games

IV coalitional games

Strategic games

A strategic game consists of

– a set of players (decision makers)

– for each player, a set of possible actions

– for each player, preferences over the set of action profiles (outcomes).

In strategic games, players move simultaneously. A wide range of situations may be modeled as strategic games.

A two-player (finite) strategic game can be described conveniently in a so-called bi-matrix.

For example, a generic 2×2 (two players and two possible actions for each player) game

	L	R
T	A_1, A_2	B_1, B_2
B	C_1, C_2	D_1, D_2

where the two rows (resp. columns) correspond to the possible actions of player 1 (resp. 2).

For example, rock-paper-scissors (over a dollar):

	R	P	S
R	$0, 0$	$-1, 1$	$1, -1$
P	$1, -1$	$0, 0$	$-1, 1$
S	$-1, 1$	$1, -1$	$0, 0$

Each player's set of actions is $\{Rock, Papar, Scissors\}$ and the set of action profiles is

$$\{RR, RP, RS, PR, PP, PS, SR.SP, SS\}.$$

In rock-paper-scissors

$$PR \sim_1 SP \sim_1 RS \succ_1 PP \sim_1 RR \sim_1 SS \succ_1 PS \sim_1 SR \sim_1 PS$$

and

$$PR \sim_2 SP \sim_2 RS \prec_2 PP \sim_2 RR \sim_2 SS \prec_2 PS \sim_2 SR \sim_2 PS$$

This is a zero-sum or a strictly competitive game.

Classical 2×2 **games**

- The following simple 2×2 games represent a variety of strategic situations.

- Despite their simplicity, each game captures the essence of a type of strategic interaction that is present in more complex situations.

- These classical games "span" the set of almost *all* games (strategic equivalence).

Game I: Prisoner's Dilemma

	$Work$	$Goof$
$Work$	$3,3$	$0,4$
$Goof$	$4,0$	$1,1$

A situation where there are gains from cooperation but each player has an incentive to "free ride."

Examples: team work, duopoly, arm/advertisement/R&D race, public goods, and more.

Game II: Battle of the Sexes (BoS)

	$Ball$	$Show$
$Ball$	$2, 1$	$0, 0$
$Show$	$0, 0$	$1, 2$

Like the Prisoner's Dilemma, Battle of the Sexes models a wide variety of situations.

Examples: political stands, mergers, among others.

Game III-V: Coordination, Hawk-Dove, and Matching Pennies

	Ball	*Show*
Ball	2, 2	0, 0
Show	0, 0	1, 1

	Dove	*Hawk*
Dove	3, 3	1, 4
Hawk	4, 1	0, 0

	Head	*Tail*
Head	1, −1	−1, 1
Tail	−1, 1	1, −1

Best response and dominated actions

Action a is player 1's *best response* to action b player 1 if it is the optimal choice when 1 *conjectures* that 2 will play b.

In any game, player 1's action a' is *strictly* dominated if it is never a best response (inferior no matter what the other players do).

In the Prisoner's Dilemma, for example, action $Work$ is strictly dominated by action $Goof$. As we will see, a strictly dominated action is not used in any Nash equilibrium.

Nash equilibrium

Nash equilibrium (NE) is a steady state of the play of a strategic game – no player has a profitable deviation given the actions of the other players.

Put differently, a NE is a set of actions such that all players are doing their best given the actions of the other players.

Mixed strategy Nash equilibrium in the BoS

Suppose that, each player can randomize among all her strategies so choices are not deterministic:

		q L	$1-q$ R
p	T	pq	$p(1-q)$
$1-p$	B	$(1-p)q$	$(1-p)(1-q)$

Let p and q be the probabilities that player 1 and 2 respectively assign to the strategy *Ball*.

Player 2 will be indifferent between using her strategy B and S when player 1 assigns a probability p such that her expected payoffs from playing B and S are the same. That is,

$$\begin{gathered} 1p + 0(1-p) = 0p + 2(1-p) \\ p = 2 - 2p \\ p^* = 2/3 \end{gathered}$$

Hence, when player 1 assigns probability $p^* = 2/3$ to her strategy B and probability $1 - p^* = 1/3$ to her strategy S, player 2 is indifferent between playing B or S any mixture of them.

Similarly, player 1 will be indifferent between using her strategy B and S when player 2 assigns a probability q such that her expected payoffs from playing B and S are the same. That is,

$$\begin{gathered} 2q + 0(1-q) = 0q + 1(1-q) \\ 2q = 1 - q \\ q^* = 1/3 \end{gathered}$$

Hence, when player 2 assigns probability $q^* = 1/3$ to her strategy B and probability $1 - q^* = 2/3$ to her strategy S, player 2 is indifferent between playing B or S any mixture of them.

In terms of best responses:

$$B_1(q) = \begin{cases} p = 1 & if \quad p > 1/3 \\ p \in [0,1] & if \quad p = 1/3 \\ p = 0 & if \quad p < 1/3 \end{cases}$$

$$B_2(p) = \begin{cases} q = 1 & if \quad p > 2/3 \\ q \in [0,1] & if \quad p = 2/3 \\ q = 0 & if \quad p < 2/3 \end{cases}$$

The BoS has two Nash equilibria in pure strategies $\{(B, B), (S, S)\}$ and one in mixed strategies $\{(2/3, 1/3)\}$. In fact, any game with a finite number of players and a finite number of strategies for each player has Nash equilibrium (Nash, 1950).

Auctions
(if time permits...)

From Babylonia to eBay, auctioning has a very long history.

Babylon:

- women at marriageable age.

Athens, Rome, and medieval Europe:

- rights to collect taxes, dispose of confiscated property, lease of land and mines,

and many more...

The word “auction” comes from the Latin *augere*, meaning “to increase.”

The earliest use of the English word “auction” given by the *Oxford English Dictionary* dates from 1595 and concerns an auction “when will be sold Slaves, household goods, etc.”

In this era, the auctioneer lit a short candle and bids were valid only if made before the flame went out – Samuel Pepys (1633-1703) –

- Auctions, broadly defined, are used to allocate significant economics resources.

  Examples: works of art, government bonds, offshore tracts for oil exploration, radio spectrum, and more.

- Auctions take many forms. A game-theoretic framework enables to understand the consequences of various auction designs.

- Game theory can suggest the design likely to be most effective, and the one likely to raise the most revenues.

Types of auctions

Sequential / simultaneous

Bids may be called out sequentially or may be submitted simultaneously in sealed envelopes:

– English (or oral) – the seller actively solicits progressively higher bids and the item is soled to the highest bidder.

– Dutch – the seller begins by offering units at a "high" price and reduces it until all units are soled.

– Sealed-bid – all bids are made simultaneously, and the item is sold to the highest bidder.

First-price / second-price

The price paid may be the highest bid or some other price:

- First-price – the bidder who submits the highest bid wins and pay a price equal to her bid.

- Second-prices – the bidder who submits the highest bid wins and pay a price equal to the second highest bid.

Variants: all-pay (lobbying), discriminatory, uniform, Vickrey (William Vickrey, Nobel Laureate 1996), and more.

Private-value / common-value

Bidders can be certain or uncertain about each other's valuation:

- In private-value auctions, valuations differ among bidders, and each bidder is certain of her own valuation and can be certain or uncertain of every other bidder's valuation.

- In common-value auctions, all bidders have the same valuation, but bidders do not know this value precisely and their estimates of it vary.

First-price auction (with perfect information)

To define the game precisely, denote by v_i the value that bidder i attaches to the object. If she obtains the object at price p then her payoff is $v_i - p$.

Assume that bidders' valuations are all different and all positive. Number the bidders 1 through n in such a way that

$$v_1 > v_2 > \cdots > v_n > 0.$$

Each bidder i submits a (sealed) bid b_i. If bidder i obtains the object, she receives a payoff $v_i - b_i$. Otherwise, her payoff is zero.

Tie-breaking – if two or more bidders are in a tie for the highest bid, the winner is the bidder with the highest valuation.

In summary, a first-price sealed-bid auction with perfect information is the following strategic game:

– Players: the n bidders.

– Actions: the set of possible bids b_i of each player i (nonnegative numbers).

– Payoffs: the preferences of player i are given by

$$u_i = \begin{cases} v_i - \overline{b} & \text{if} \quad b_i = \overline{b} \text{ and } v_i > v_j \text{ if } b_j = \overline{b} \\ 0 & \text{if} \quad b_i < \overline{b} \end{cases}$$

where $\overline{b}$ is the highest bid.

The set of Nash equilibria is the set of profiles $(b_1, ..., b_n)$ of bids with the following properties:

$$\begin{array}{ll} [1] & v_2 \leq b_1 \leq v_1 \\ [2] & b_j \leq b_1 \ \underline{\text{for all}}\ j \neq 1 \\ [3] & b_j = b_1 \ \underline{\text{for some}}\ j \neq 1 \end{array}$$

It is easy to verify that all these profiles are Nash equilibria. It is harder to show that there are no other equilibria. We can easily argue, however, that there is no equilibrium in which player 1 does not obtain the object.

$\Longrightarrow$ The first-price sealed-bid auction is socially efficient, but does not necessarily raise the most revenues.

Second-price auction (with perfect information)

A second-price sealed-bid auction with perfect information is the following strategic game:

- Players: the n bidders.

- Actions: the set of possible bids b_i of each player i (nonnegative numbers).

- Payoffs: the preferences of player i are given by

$$u_i = \begin{cases} v_i - \overline{b} & \text{if} \quad b_i > \overline{b} \text{ or } b_i = \overline{b} \text{ and } v_i > v_j \text{ if } b_j = \overline{b} \\ 0 & \text{if} \quad b_i < \overline{b} \end{cases}$$

  where $\overline{b}$ is the highest bid submitted by a player other than i.

First note that for any player i the bid $b_i = v_i$ is a (weakly) dominant action (a "truthful" bid), in contrast to the first-price auction.

The second-price auction has many equilibria, but the equilibrium $b_i = v_i$ for all i is distinguished by the fact that every player's action dominates all other actions.

Another equilibrium in which player $j \neq 1$ obtains the good is that in which

[1] $b_1 < v_j$ and $b_j > v_1$
[2] $b_i = 0$ <u>for all</u> $i \neq \{1, j\}$

Common-value auctions and the winner's curse

Suppose we all participate in a sealed-bid auction for a jar of coins. Once you have estimated the amount of money in the jar, what are your bidding strategies in first- and second-price auctions?

The winning bidder is likely to be the bidder with the largest positive error (the largest overestimate).

In this case, the winner has fallen prey to the so-called the winner's curse. Auctions where the winner's curse is significant are oil fields, spectrum auctions, pay per click, and more.

Conclusions

Adam Brandenburger:

> *There is nothing so practical as a good [game] theory. A good theory confirms the conventional wisdom that "less is more." A good theory does less because it does not give answers. At the same time, it does a lot more because it helps people organize what they know and uncover what they do not know. A good theory gives people the tools to discover what is best for them.*