UC Berkeley
Haas School of Business
Economic Analysis for Business Decisions
(EWMBA 201A)
Fall 2022

Block IV
Oligopoly (PR 12.2-3) and (some) Game Theory (PR 13)

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

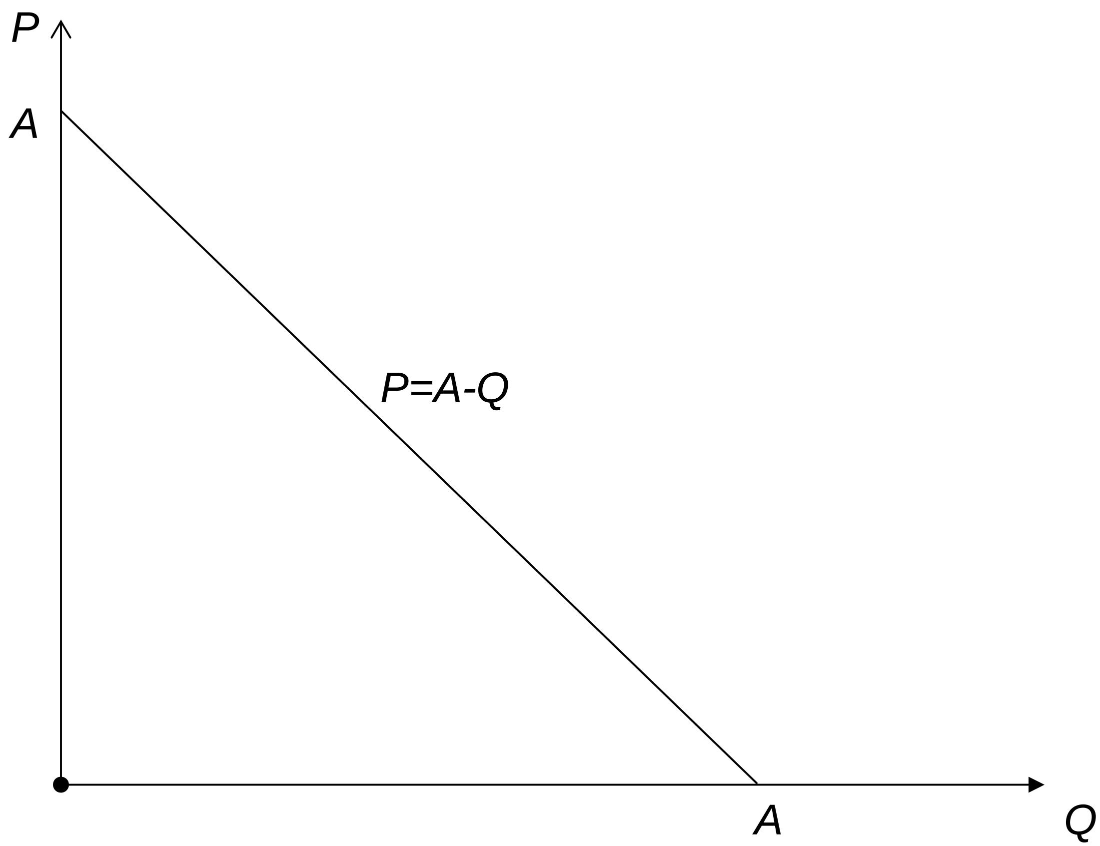

To find the Nash equilibria of the Cournot's game, we can use the procedures based on the firms' best response functions.

But first we need the firms payoffs (profits):

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

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)

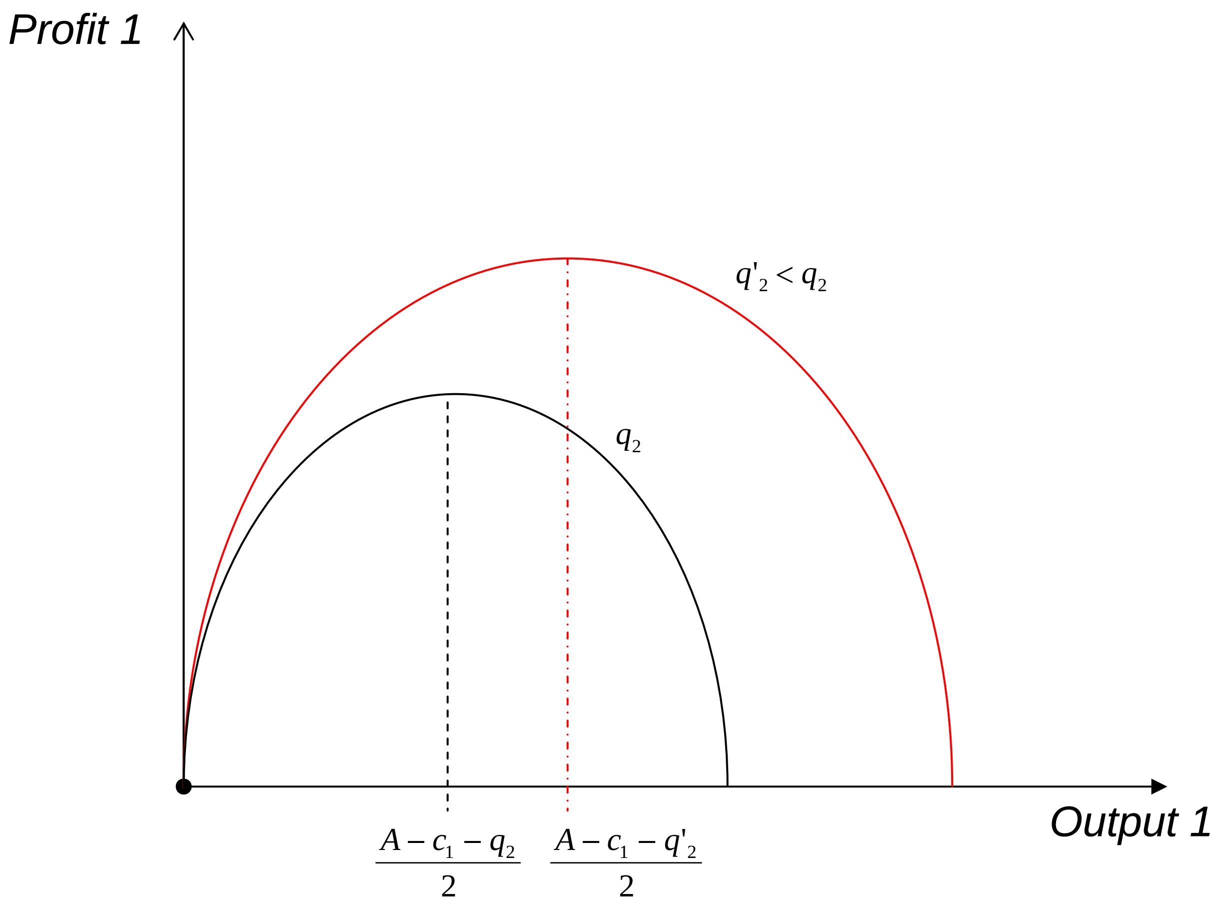

To find firm 1's best response to any given output q_2 of firm 2, we need to study firm 1's profit as a function of its output q_1 for given values of q_2.

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

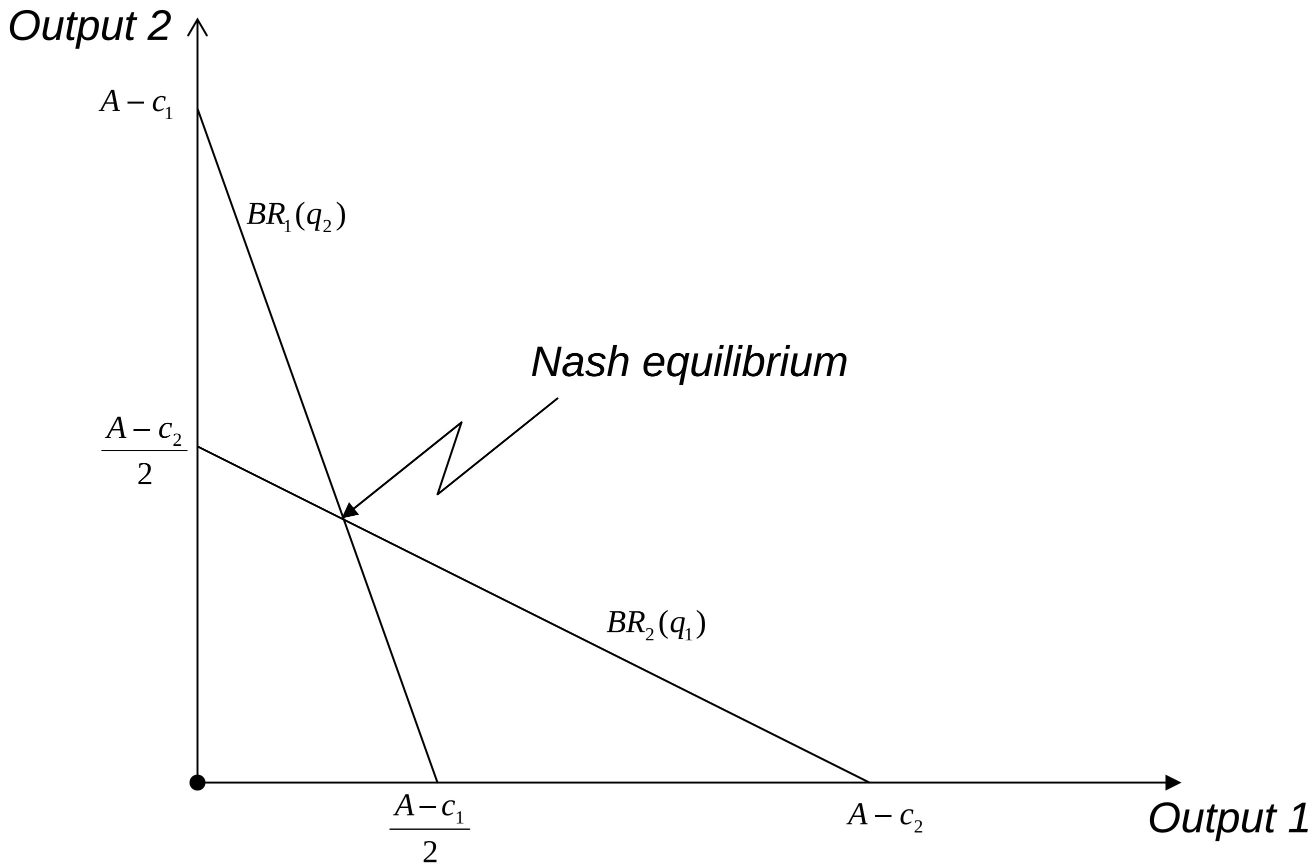

Nash equilibrium comparative statics
(a decrease in the cost of firm 2)

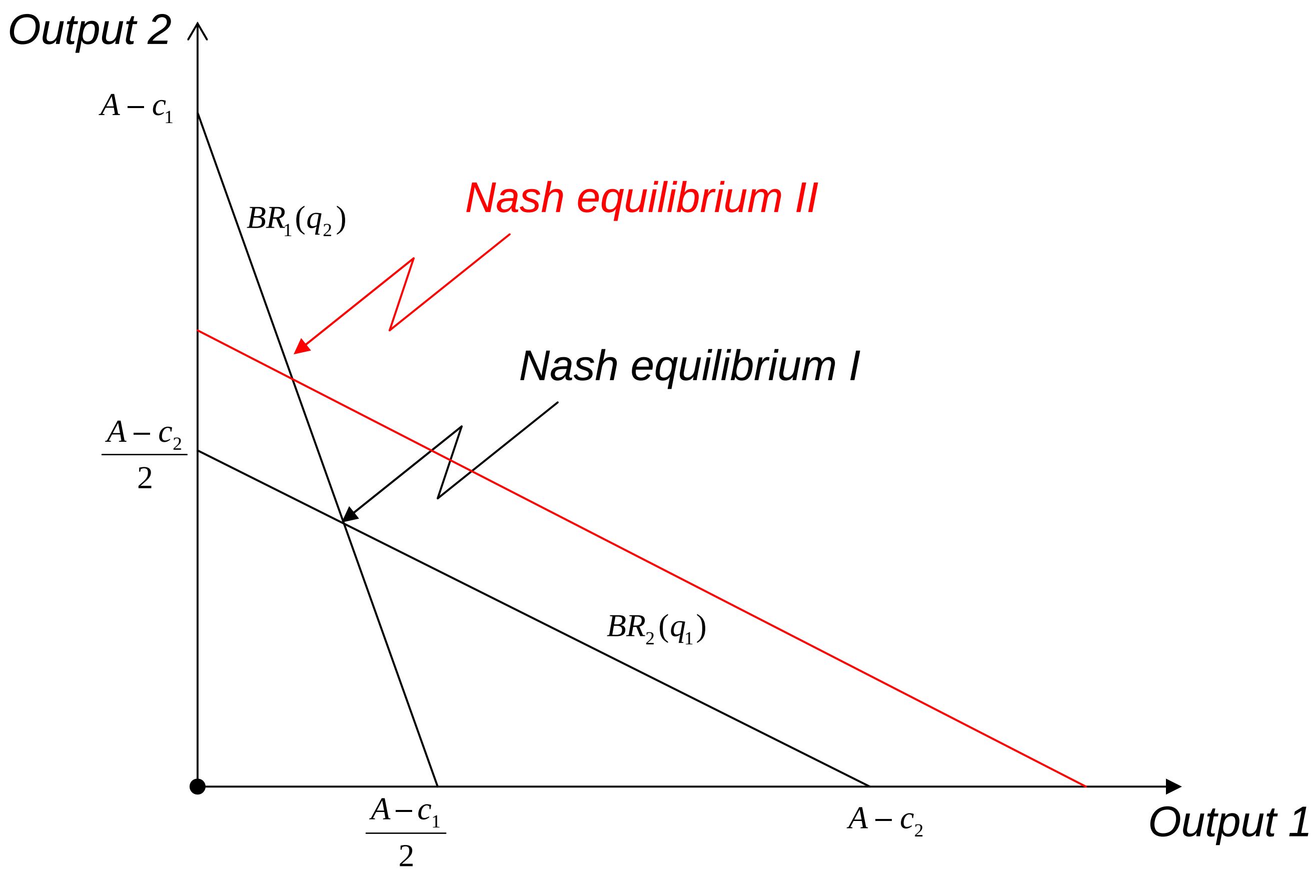

A question: what happens when consumers are willing to pay more (*A* increases)?

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

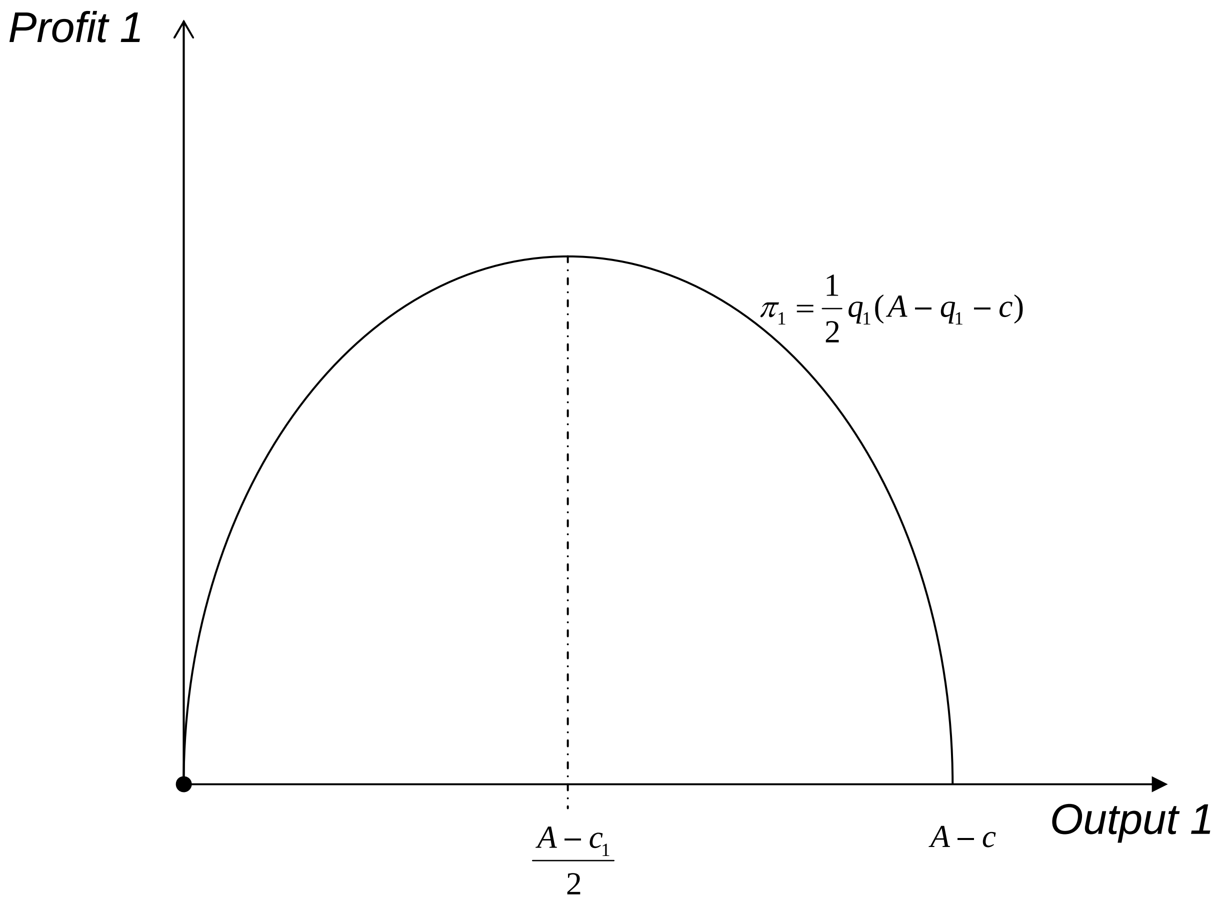

We conclude that Stackelberg's duopoly game has a unique subgame perfect equilibrium, in which firm 1's strategy is the output

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

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

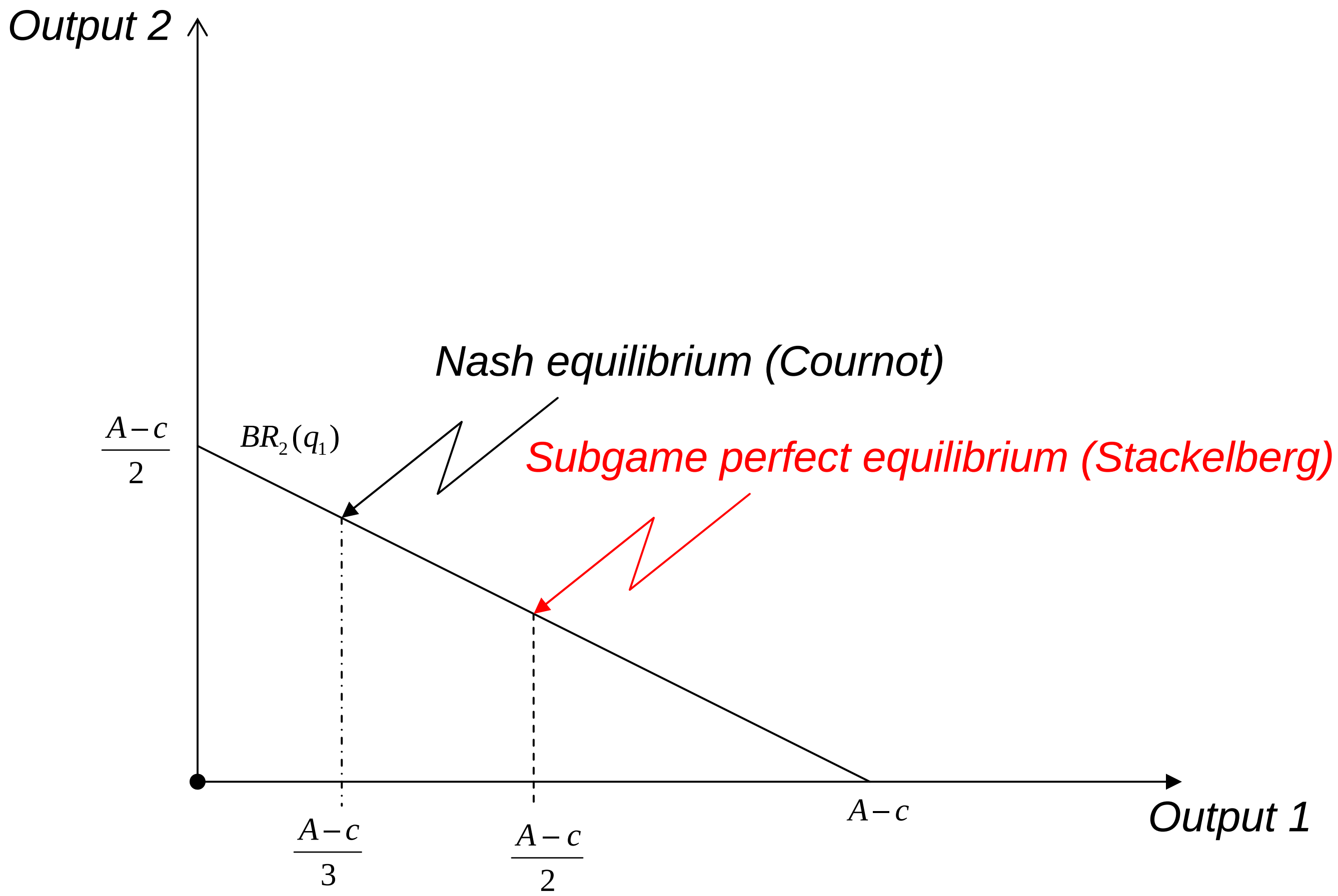

Avoiding the Bertrand trap

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

Guys, it's time for some game theory…

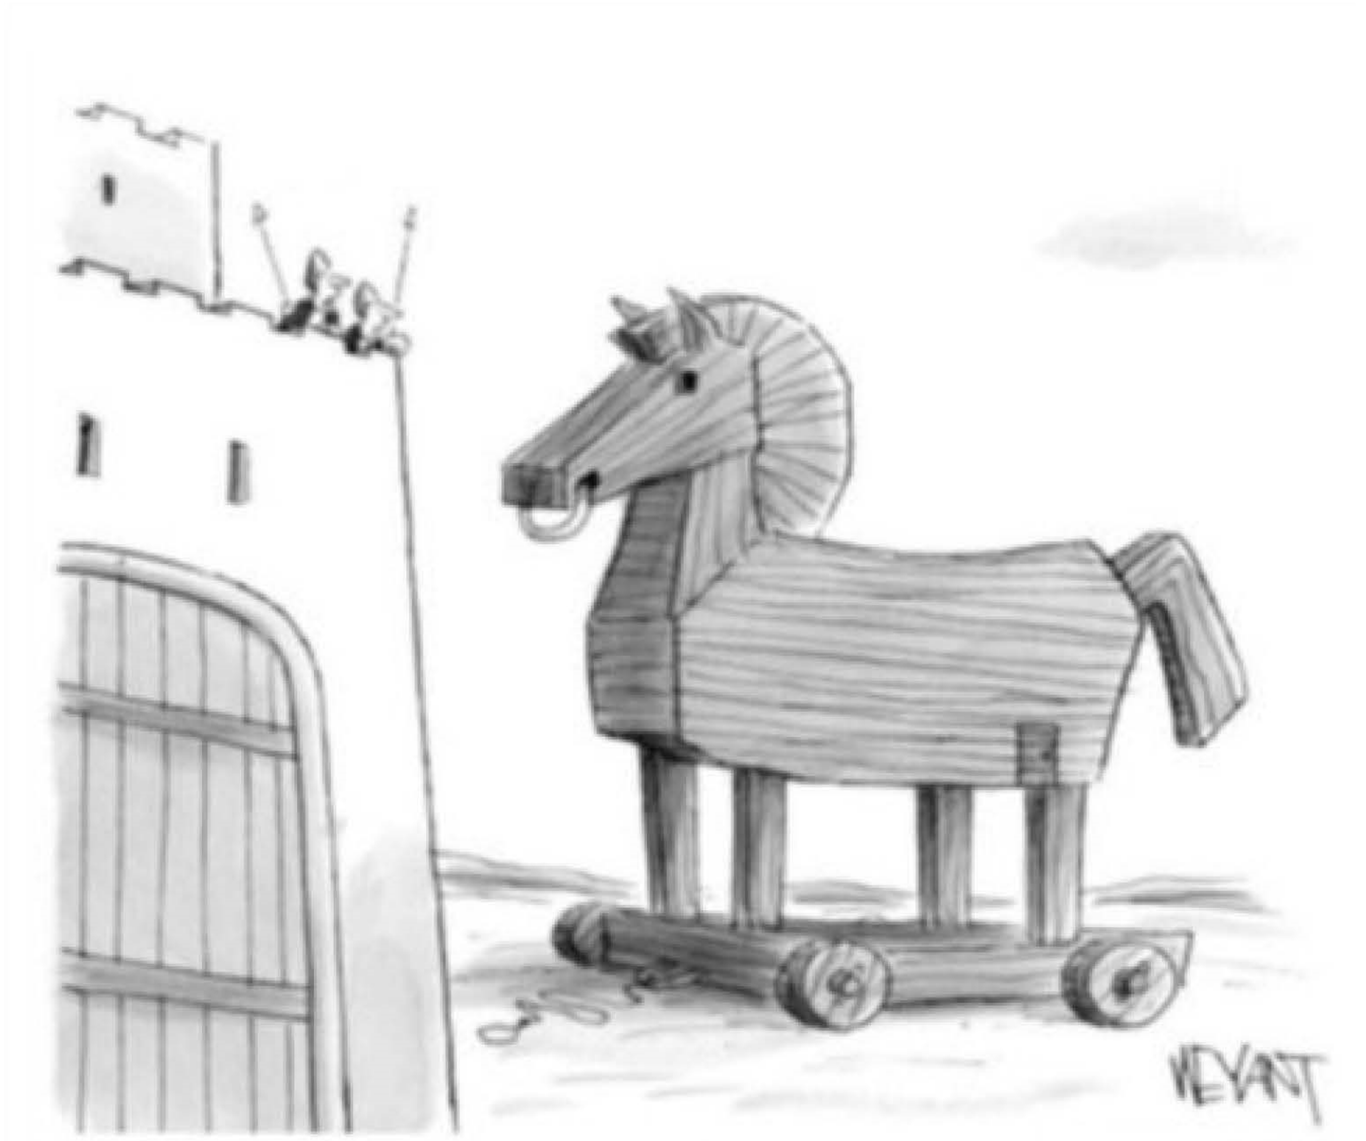

What’s game theory?

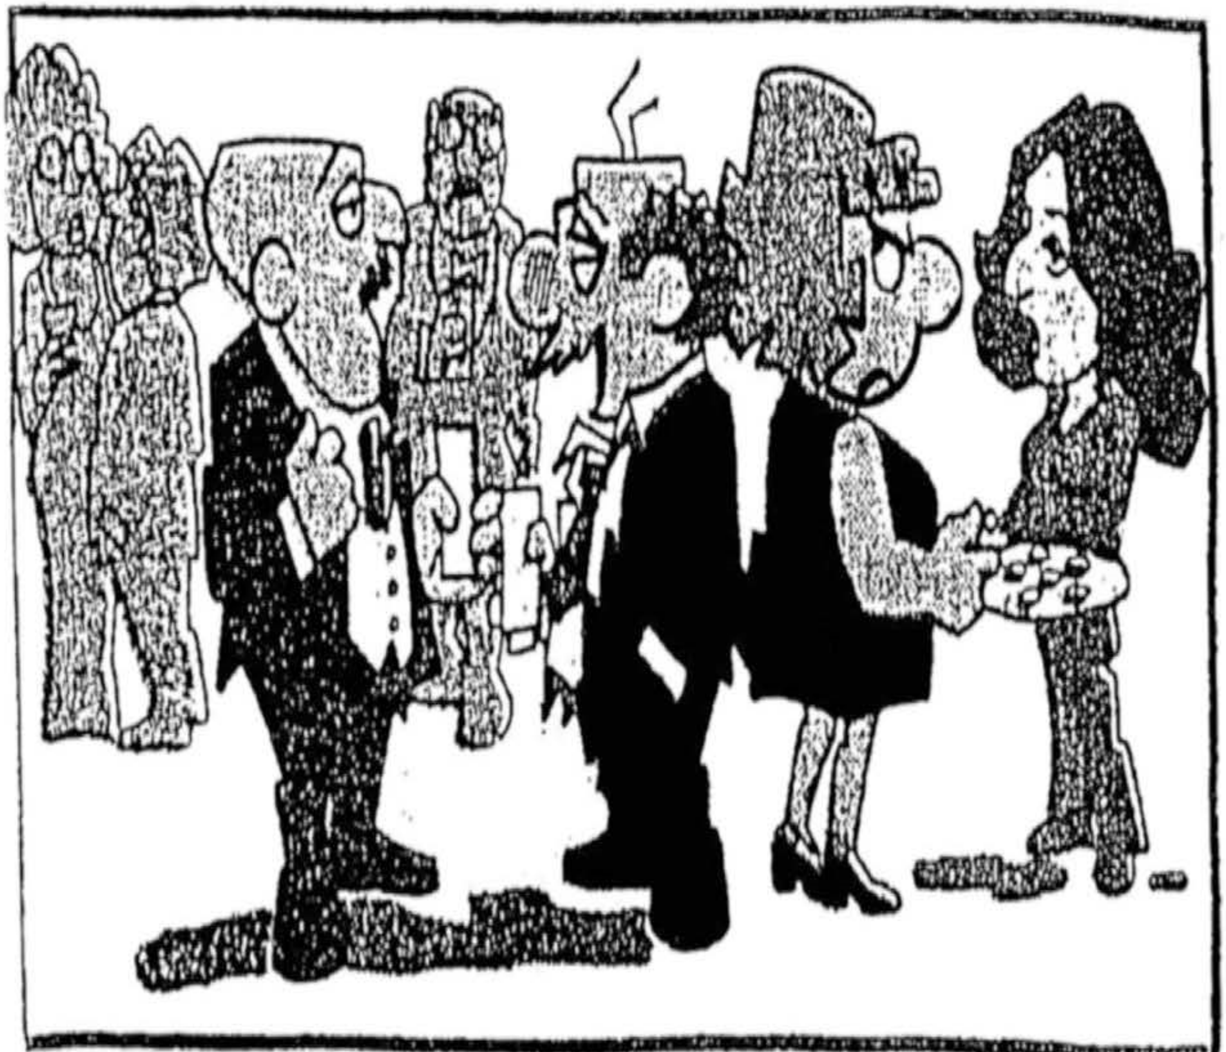

"LORETTA'S DRIVING BECAUSE I'M DRINKING, AND I'M DRINKING BECAUSE SHE'S DRIVING."

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.

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).*"

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.*

The New York Times Magazine

FEATURE

How Data (and Some Breathtaking Soccer) Brought Liverpool to the Cusp of Glory

The club is finishing a phenomenal season — thanks in part to an unrivaled reliance on analytics.

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

Four ‘simple’ games

Four 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.

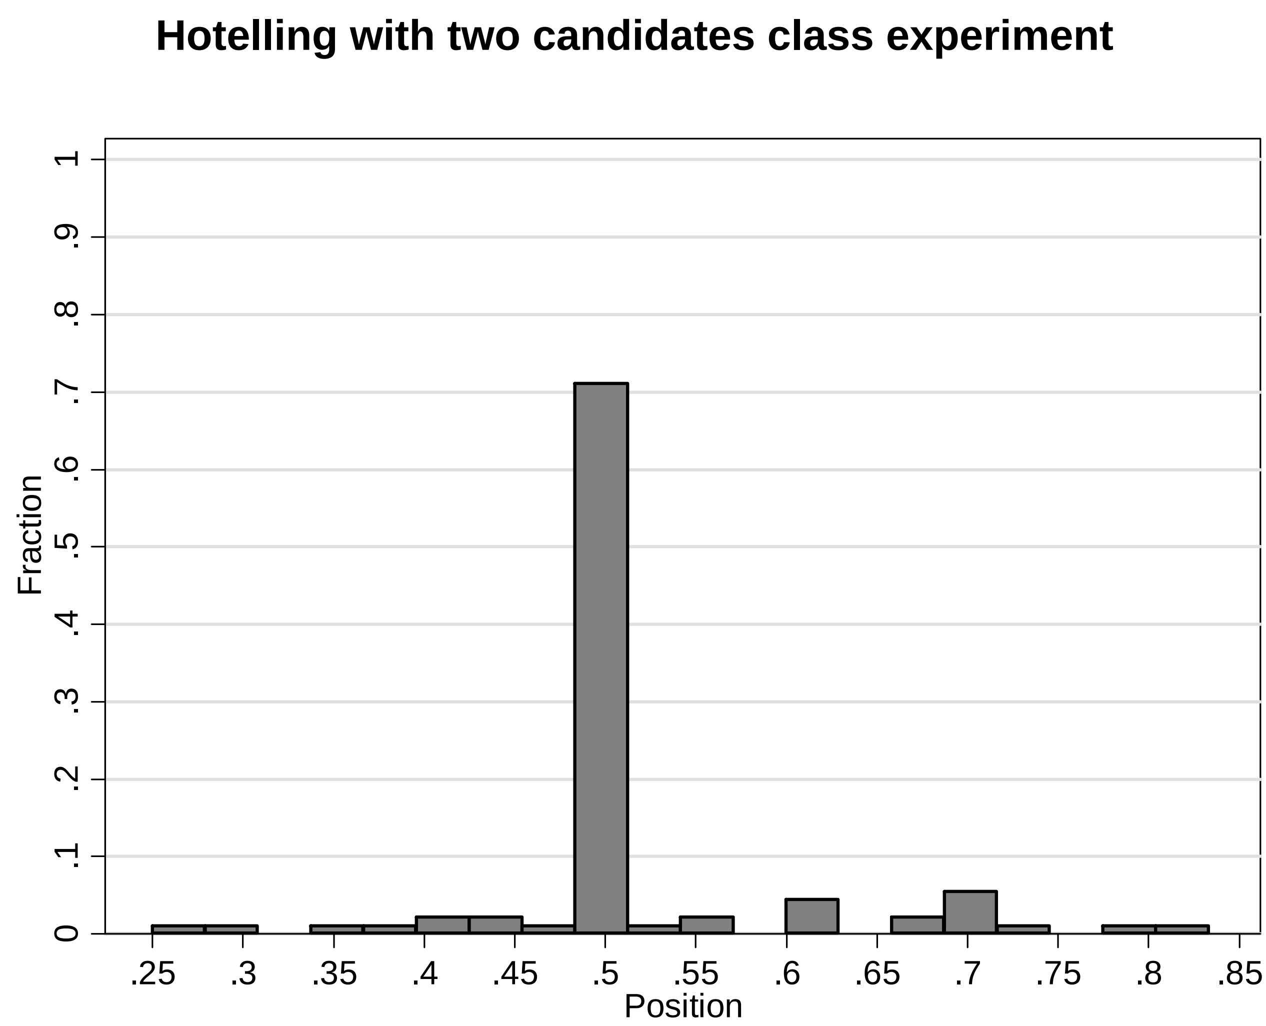
Hotelling with two candidates class experiment
Fraction
0
.1
.2
.3
.4
.5
.6
.7
.8
.9
1
.25
.3
.35
.4
.45
.5
.55
.6
.65
.7
.75
.8
.85
Position

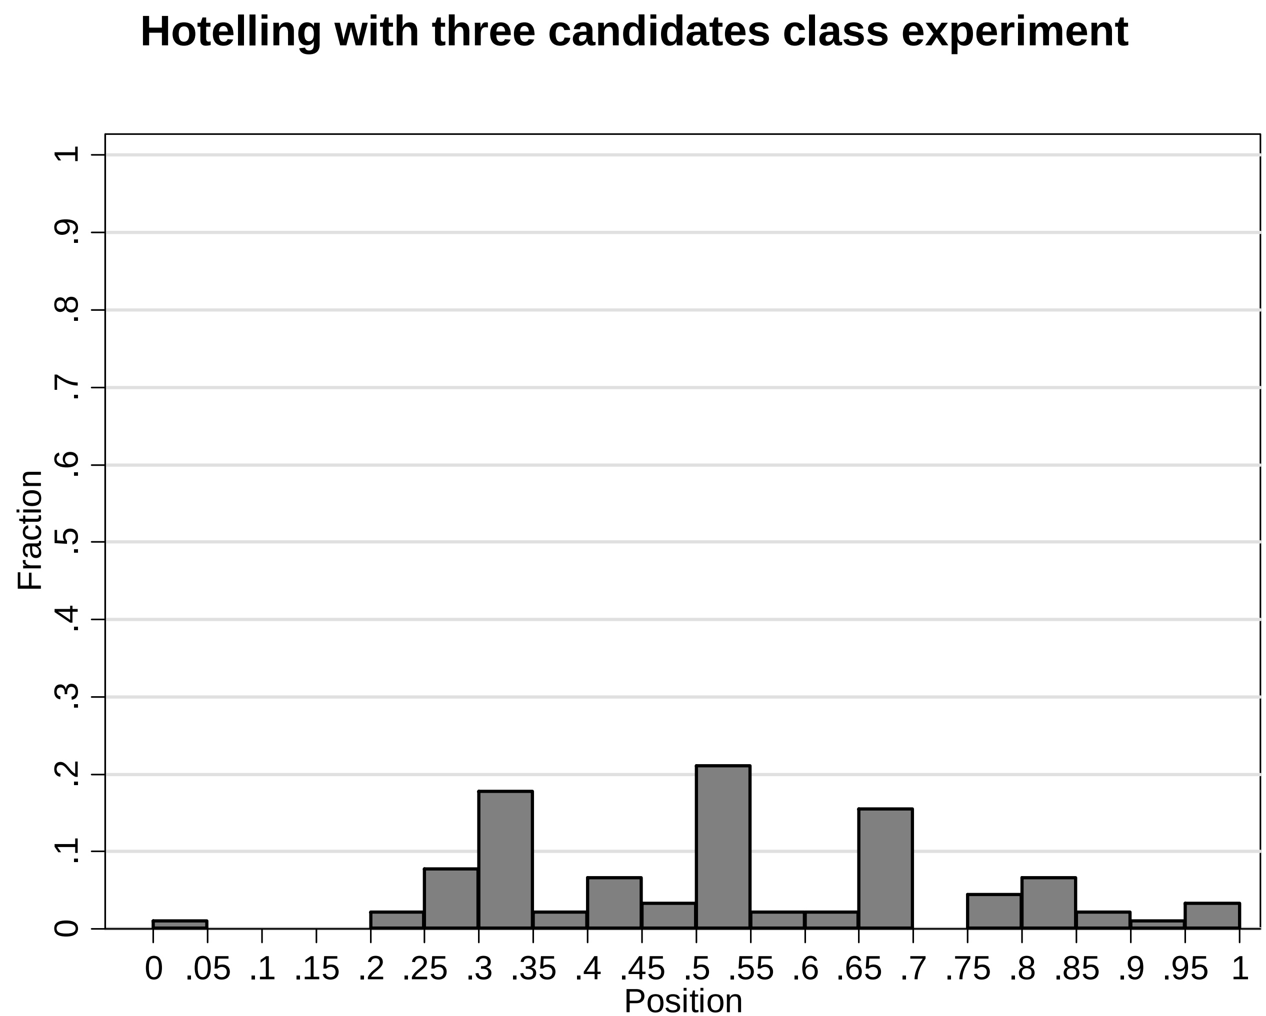
Hotelling with three candidates class experiment
Fraction
0
.1
.2
.3
.4
.5
.6
.7
.8
.9
1
0 .05 .1 .15 .2 .25 .3 .35 .4 .45 .5 .55 .6 .65 .7 .75 .8 .85 .9 .95 1
Position

John Maynard Keynes
1883-1946

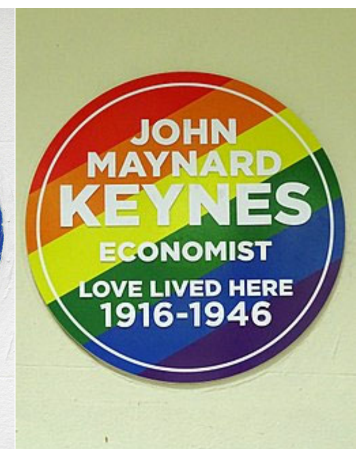

Example II: Keynes's beauty contest game

- 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%

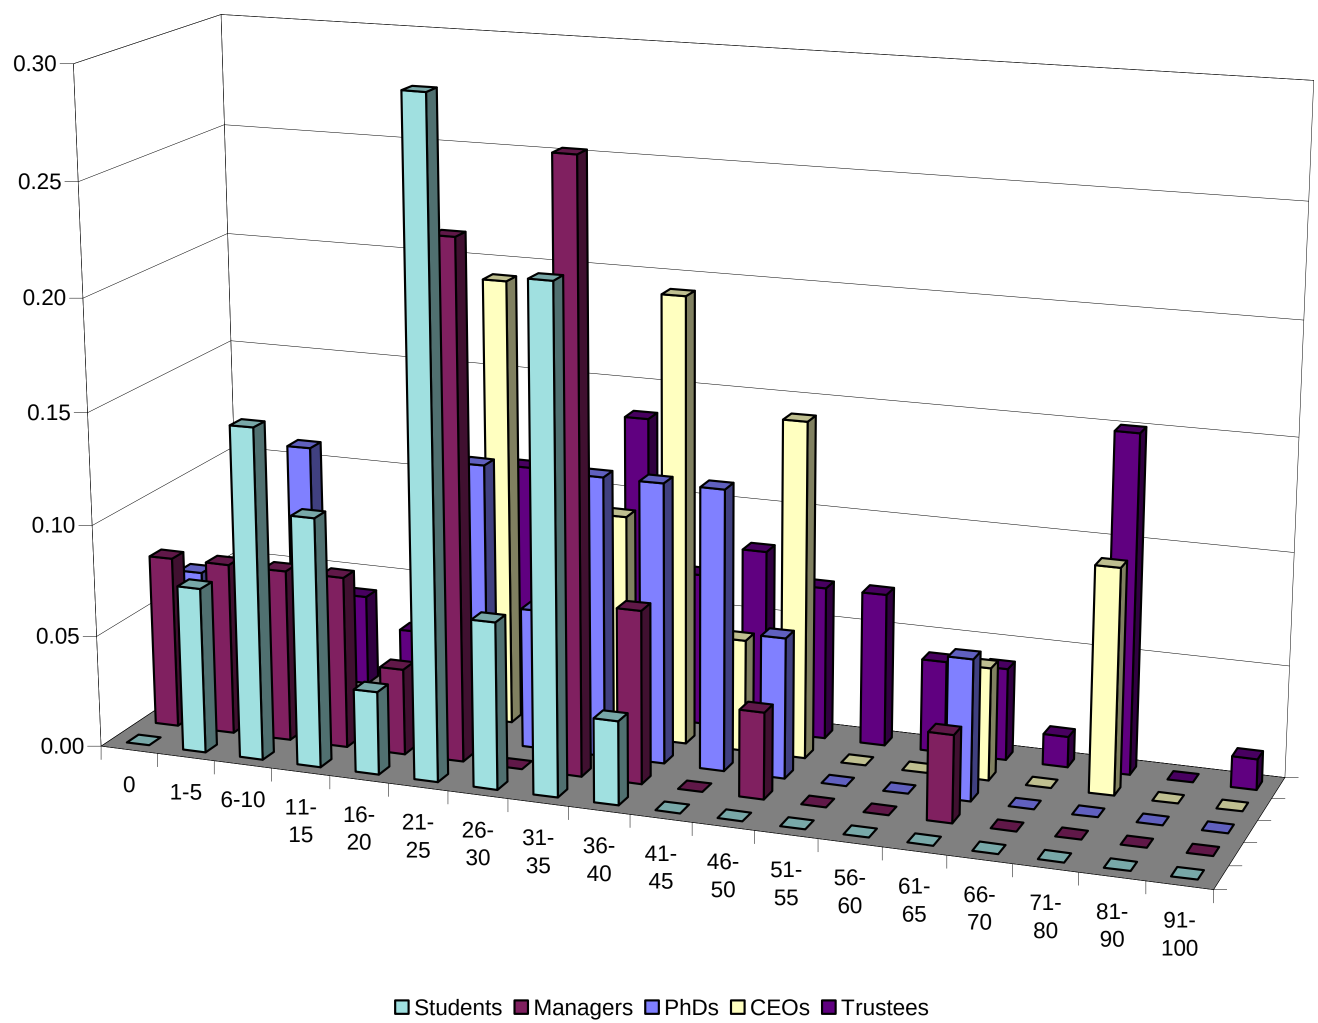

0.30
0.25
0.20
0.15
0.10
0.05
0.00
0
1-5
6-10
11-15
16-20
21-25
26-30
31-35
36-40
41-45
46-50
51-55
56-60
61-65
66-70
71-80
81-90
91-100
Students
Managers
PhDs
CEOs
Trustees

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

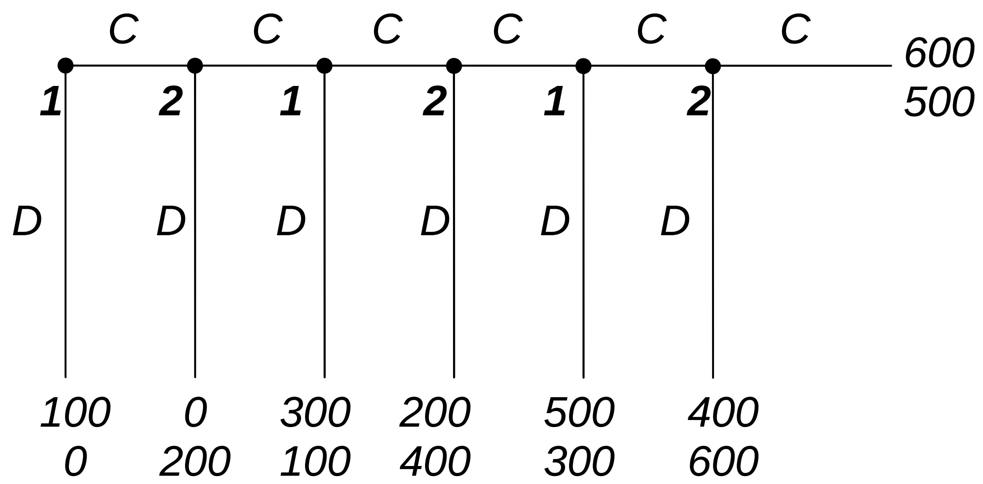

The centipede game class experiment

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).

Example IV: auctions

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.

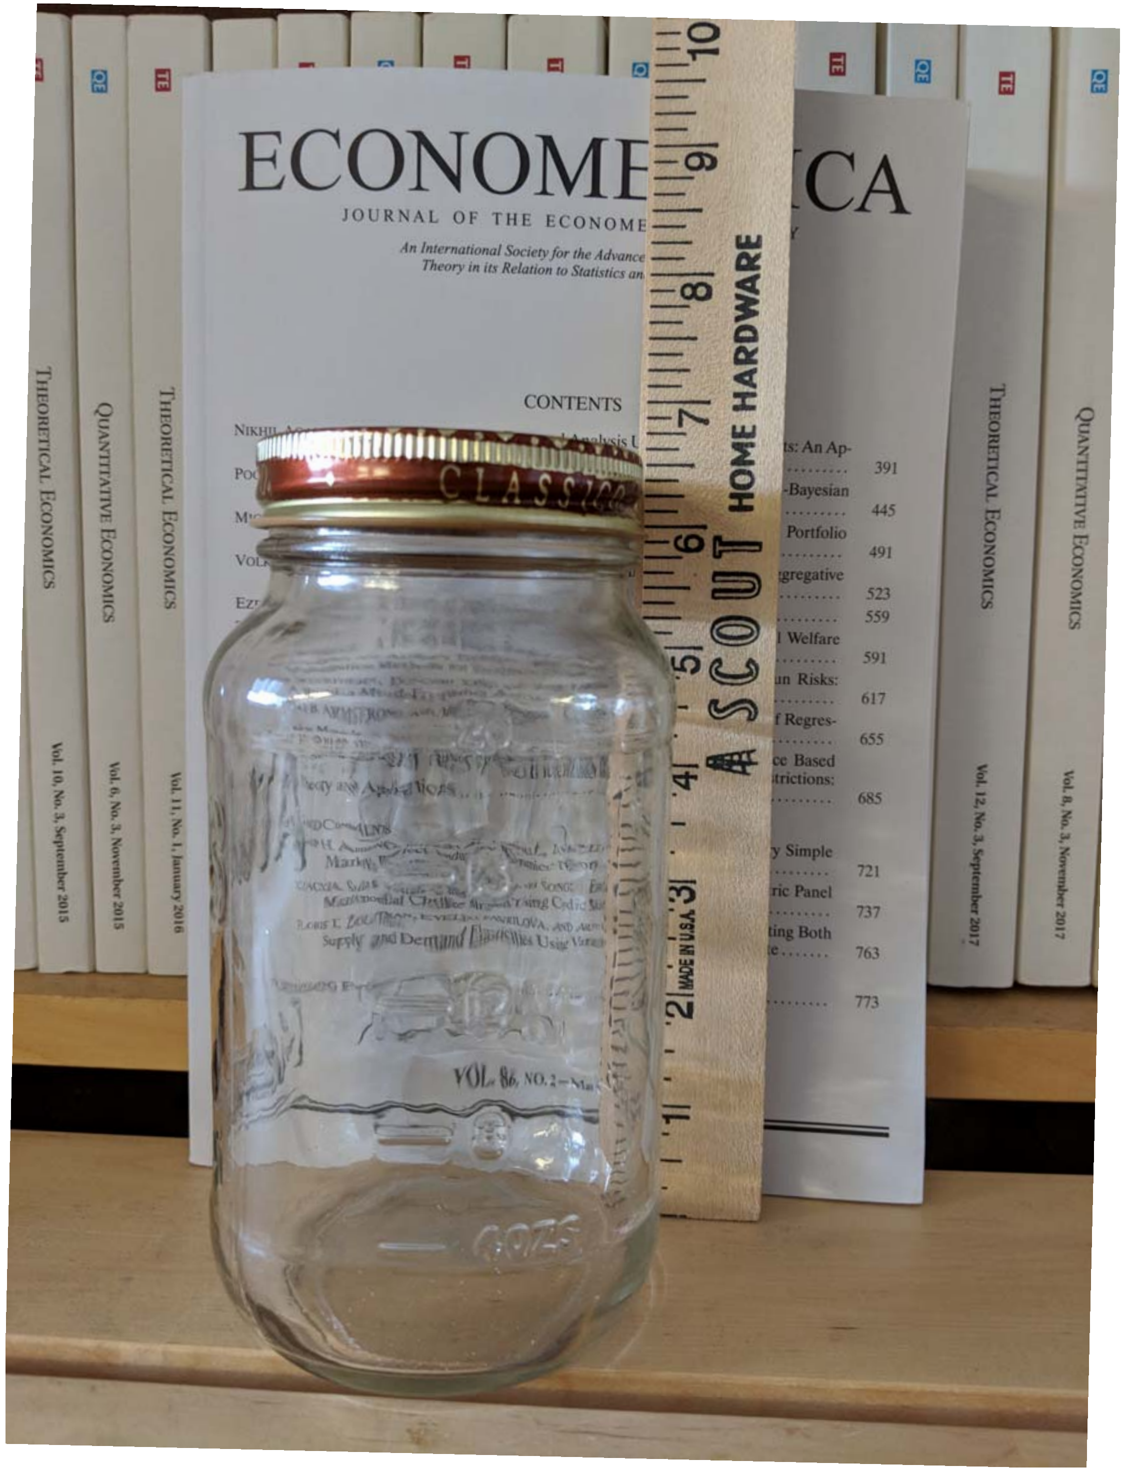
THEORETICAL ECONOMICS
Vol. 10, No. 3, September 2015
QUANTITATIVE ECONOMICS
Vol. 6, No. 3, November 2015
THEORETICAL ECONOMICS
Vol. 11, No. 1, January 2016
CONTENTS
HOME HARDWARE
A SCOUT
MADE IN U.S.A.
THEORETICAL ECONOMICS
Vol. 12, No. 3, September 2017
QUANTITATIVE ECONOMICS
Vol. 8, No. 3, November 2017

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.