

**UC Berkeley
Haas School of Business
Game Theory
(EMBA 296 & EWMBA 211)
Summer 2018**

Preliminaries

**Block 1
May 17-19, 2018**


What's 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 become a standard language in economics.
- The power of game theory is its generality and (mathematical) precision.

- 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!

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.”

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.

Using game theory to understand political decisions

Distributional preferences

Distributional preferences shape individual opinions on a range of issues related to the redistribution of income:

- unemployment benefits,
- social security,
- government-sponsored healthcare,
- and more.

These preferences are thus important inputs into any broader measure of social welfare and (should) enter every realm of voters' decision-making.

Issues around tax policy and other forms of (government) redistribution are complex and contentious:

[1] because people promote their competing private interests, and

[2] people who are motivated by morality (fairness) to promote the interests of others disagree about what constitutes an equitable or just outcome.

As a result, political and social conflicts often hover uncertainly between interest-competition [1] and moral disagreement [2].

Fair-mindedness and equality versus efficiency

Distributional preferences may naturally be divided into two qualitatively different components:

- The weight on own income versus the incomes of others (fair-mindedness)
- The weight on reducing differences in incomes versus increasing total income (equality-efficiency tradeoffs).

Fair-minded people may disagree about the extent to which efficiency should be sacrificed to combat inequality, as a comparison of Harsanyi (1955) and Rawls (1971) would suggest.

In the United States, we typically associate the Democratic party with the promotion of policies which reduce inequality, and the Republican party with the promotion of efficiency.

- However, whether Democratic voters are more willing to sacrifice efficiency – and even their own income – to reduce inequality is an open question.
- Alternatively, Democrats may be those who expect to benefit from government redistribution, as the median voter theorem would suggest.


We thus cannot understand public opinion on a number of important policy issues without understanding the individual distributional preferences of the general population:

- Distinguish fair-mindedness from preferences regarding equality-efficiency tradeoffs.
- Accurately measuring both in a large and diverse sample of American voters (and elites).

Template for analysis

- [1] A generalized dictator game where each subject faces a menu of budget sets representing the feasible monetary payoffs.
- [2] An incentivized experiment using the American Life Panel (ALP), a longitudinal survey of more than 5,000 individuals.
- [3] Combine data from the experiments with detailed individual demographic and economic information on panel members.

The experiential decision problem



A choice of the allocation (π_s, π_o) from the budget set $p_s\pi_s + p_o\pi_o = 1$ represents the payoffs to persons *self* and *other*, respectively.

The budget line configuration allows to identify the equality-efficiency tradeoffs that subjects make in their distributional preferences:

- *decreasing* $p_s\pi_s$ when p_s/p_o *increases* indicates preferences weighted towards efficiency (increasing total payoffs)
- *increasing* $p_s\pi_s$ when p_s/p_o *increases* indicates preferences weighted towards equality (reducing differences in payoffs).

An economic model of distributional preferences

We decompose distributional preferences into fair-mindedness and equality-efficiency tradeoffs by employing constant elasticity of substitution (CES) utility functions.

The CES form is commonly employed in demand analysis. In the redistribution context, the CES has the form

$$u_s(\pi_s, \pi_o) = [\alpha(\pi_s)^\rho + (1 - \alpha)(\pi_o)^\rho]^{1/\rho}$$


where α measures fair-mindedness (indexical weight on payoffs to *self*) and ρ measures the willingness to trade off equality and efficiency.

If $\rho > 0$ ($\rho < 0$) a decrease in the relative price giving p_s/p_o lowers (raises) the expenditure on tokens allocated to *self* $p_s\pi_s$:


- $\rho < 0$ indicates preferences weighted towards reducing differences in payoffs (equality).
- $\rho > 0$ indicates preferences weighted towards increasing total payoffs (efficiency).

Our experimental method generates many observations per subject, and we can therefore analyze both types of distributional preferences at the individual level.


Prototypical fair-minded distributional preferences




**The relationship between the log-price ratio and fraction kept ($\pi_s/(\pi_s + \pi_o)$)
($\alpha=0.5$ and different values of ρ)**



**The relationship between the log-price ratio and fraction kept ($\pi_s/(\pi_s + \pi_o)$)
($\alpha=0.75$ and different values of ρ)**





**The relationship between the log-price ratio and fraction kept ($\pi_s/(\pi_s + \pi_o)$)
($\alpha=0.9$ and different values of ρ)**



Foundations of Economic Analysis (1947)



Paul A. Samuelson (1915-2009) – the first American Nobel laureate in economics and the foremost (academic) economist of the 20th century (and the uncle of Larry Summers...).



Side note I: individual preferences

Consider some (finite) set of alternatives (x, y, z, \dots) .

- Formally, we represent the decision-maker's preferences by a binary relation \succsim defined on the set of consumption bundles.
- For any pair of bundles x and y , if the decision-maker says that x is at least as good as y , we write

$$x \succsim y$$

and say that x is *weakly preferred* to y .

Bear in mind: economic theory often seeks to convince you with simple examples and then gets you to extrapolate. This simple construction works in wider (and wilder circumstances).

From the weak preference relation \succsim we derive two other relations on the set of alternatives:

- Strict performance relation

$$x \succ y \text{ if and only if } x \succsim y \text{ and not } y \succsim x.$$

The phrase $x \succ y$ is read x is *strictly preferred* to y .

- Indifference relation

$$x \sim y \text{ if and only if } x \succsim y \text{ and } y \succsim x.$$

The phrase $x \sim y$ is read x is *indifferent* to y .

Side note II: individual rationality

Economic theory begins with two assumptions about preferences. These assumptions are so fundamental that we can refer to them as “axioms” of decision theory.

[1] Completeness

$$x \succsim y \text{ or } y \succsim x$$

for any pair of bundles x and y .

[2] Transitivity

$$\text{if } x \succsim y \text{ and } y \succsim z \text{ then } x \succsim z$$

for any three bundles x , y and z .

Together, completeness and transitivity constitute the formal definition of *rationality* as the term is used in economics. Rational economic agents are ones who

have the ability to make choices [1], and whose choices display a logical consistency [2].

(Only) the preferences of a rational agent can be represented, or summarized, by a *utility function*.


The Generalized Axiom of Revealed Preference (GARP)

The most basic question to ask about choice data is whether it is consistent with individual utility maximization, and classical revealed preference theory provides a direct test:

- choices are consistent with maximizing a well-behaved (piecewise linear, continuous, increasing, and concave) utility function if and only if they satisfy GARP.

The obvious difficulty: GARP provides an exact test of utility maximization – either the data satisfy GARP or they do not – but individual choices frequently involve at least some errors.

Testing for GARP




The critical cost efficiency index (CCEI)

The CCEI measures the fraction by which each budget constraint must be shifted in order to remove all violations of GARP.


- The CCEI is between 0 and 1 – indices closer to 1 mean the data are closer to perfect consistency with GARP and hence with utility maximization.

Because our subjects make choices in a wide range of budget sets, our data provides a stringent test of utility maximization.


The construction of the CCEI




Economic rationality – CCEI scores



Fair-mindedness – estimated α parameters

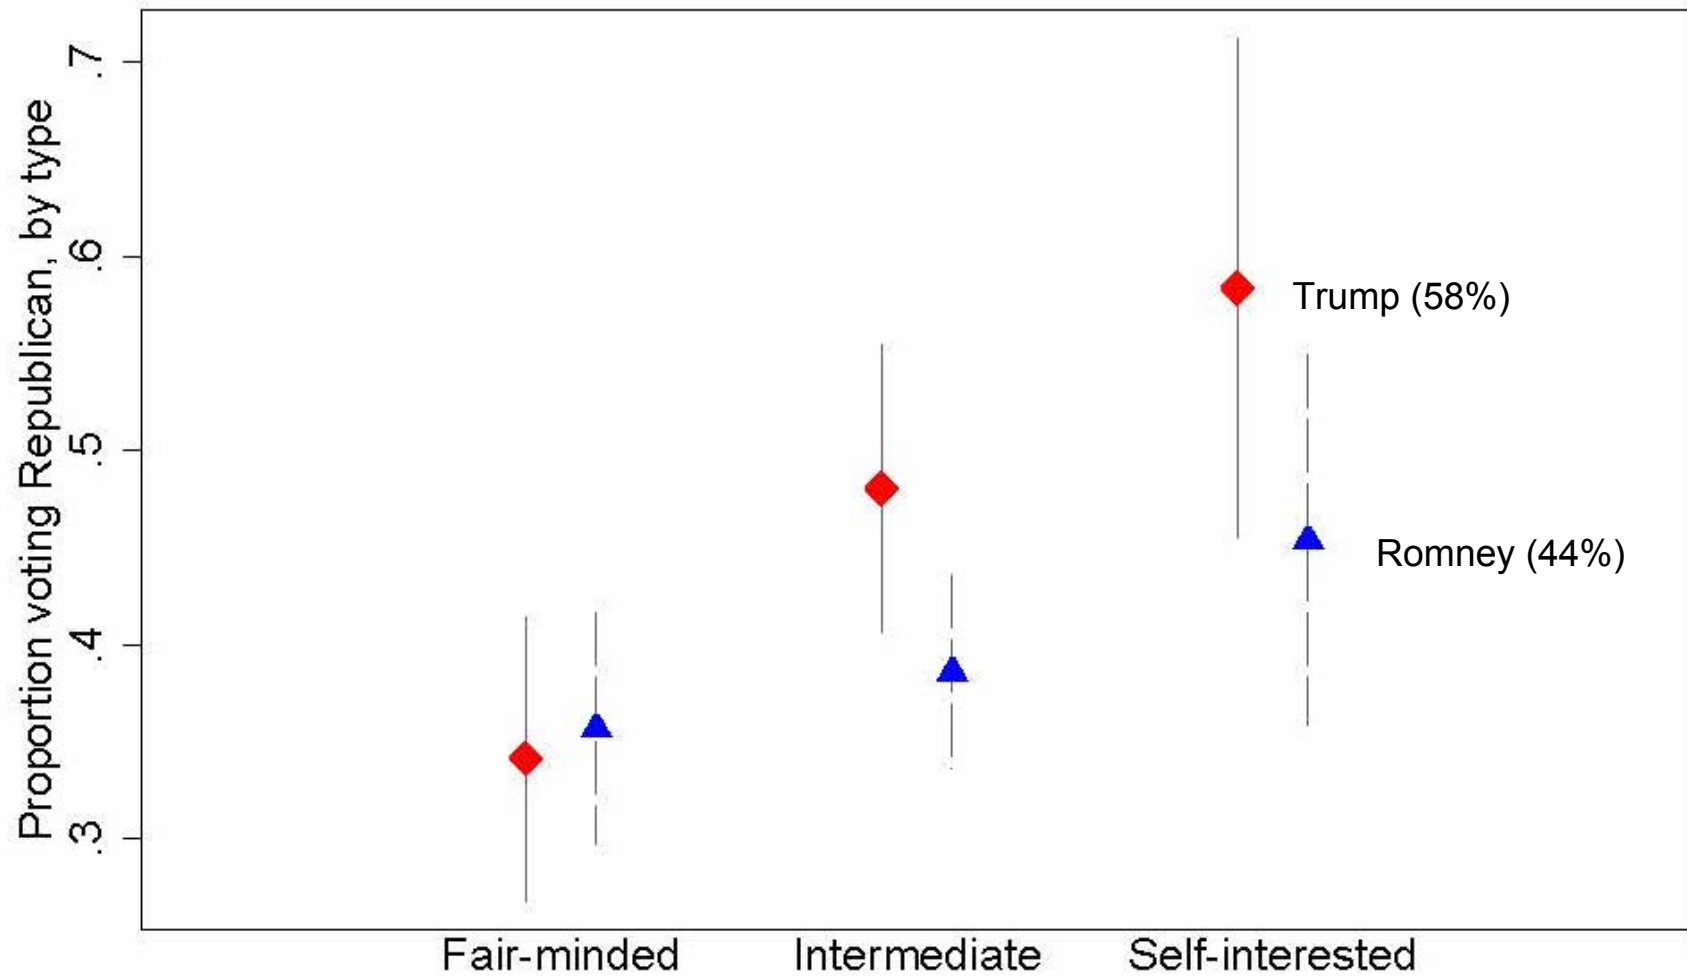


Equality versus efficiency – estimated ρ parameters






Distributional Preferences and Voting Behavior




Fair-minded, intermediate and self-interested voters in the 2012 versus 2016 presidential elections



Romney versus non-Romney voters



Trump versus non-Trump voters



Voting for Trump

	Dept. var: Voted for Trump								Romney
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$\hat{\alpha}_{n2}$	0.377*** [0.120]	0.356*** [0.123]	0.265** [0.129]	0.209 [0.138]					
$\hat{\rho}_{n2}$		0.00306 [0.00364]		0.00523 [0.00384]					
Decile of $\hat{\alpha}_{n2}$					0.0256*** [0.00853]	0.0256*** [0.00855]	0.0179** [0.00895]	0.0173* [0.00899]	
Decile of $\hat{\rho}_{n2}$						-0.00344 [0.00836]		0.00685 [0.00826]	
Decile of $\hat{\alpha}_{n1}$									-0.0114 [0.00748]
Decile of $\hat{\rho}_{n1}$									0.0141* [0.00728]
CCEI	No	No	Yes	Yes	No	No	Yes	Yes	Yes
Demographics	No	No	Yes	Yes	No	No	Yes	Yes	Yes
State FE	No	No	Yes	Yes	No	No	Yes	Yes	Yes
Observations	403	403	403	403	403	403	403	403	578
R2	0.0231	0.0247	0.357	0.361	0.0220	0.0224	0.357	0.358	0.247

The Distributional Preferences of Elites

The distributional preferences of law students

Elite law students hold especial interest because they assume positions of substantial power in national and indeed global social, economic and political affairs:

- All eight sitting Supreme Court Justices (as well as Garland and Gorsuch nominated to succeed Scalia) are graduates of either Yale or Harvard Law Schools.
- Over the past century more than half of the presidents attended Yale, Harvard or Princeton, and the last four before Donald Trump are graduates of Yale or Harvard.

The distributional preferences of elite law students will likely exercise a major influence over public and private orderings in the United States.

The distributional preferences of medical students

Patients rely on physicians to act in their best interest, healthcare systems rely on physicians to efficiently ration limited care, and physicians must balance these often conflicting imperatives against their own self-interest.




The distributional preferences of physicians thus have profound implications for patient outcomes and wellbeing, as well as the success of reforms attempting to provide more equitable, higher quality and more efficient healthcare.

Physicians' fair-mindedness – the concern for patient health and wellbeing beyond own self-interest – has been reinforced by ethical guidelines such as in the Hippocratic Oath.




“...the behavior expected of sellers of medical care is different from that of business men in general...His behavior is supposed to be governed by a concern for the customer’s welfare which would not be expected of a salesman.” (Kenneth Arrow, 1963)

“...medicine is one of the few spheres of human activity in which the purposes are unambiguously altruistic.” (Editors, *New England Journal of Medicine*, 2000)




Law students (YLS), medical students (MS) and the general population (ALP)



medical students attending tier 1 versus tier 2 medical schools



Low-income (<\$300K) versus high-income (>\$300K) medical specialties



Takeaways

1. We characterize, via experiments, the distributional preferences of the general population of the United States.
2. Overall, the data indicate a high degree of heterogeneity within each demographic or economic category.
3. Provide links from underlying distributional preferences to voter preferences over policy outcomes.
4. The distributional preferences of those (who will be) in power differ from the preferences of voters.


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 of all citizens whose favorite positions are closer to her position.

Hotelling with two candidates class experiment



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.”

⇒ 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%



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



The centipede game class experiment

<i>Down</i>	<i>0.311</i>
<i>Continue, Down</i>	<i>0.311</i>
<i>Continue, Continue, Down</i>	<i>0.267</i>
<i>Continue, Continue, Continue</i>	<i>0.111</i>

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 economic 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 sold to the highest bidder.
- Dutch – the seller begins by offering units at a “high” price and reduces it until all units are sold.
- 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.

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

	<i>L</i>	<i>R</i>
<i>T</i>	A_1, A_2	B_1, B_2
<i>B</i>	C_1, C_2	D_1, D_2

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

Applying the definition of a strategic game to the 2×2 game above yields:

- Players: $\{1, 2\}$
- Action sets: $A_1 = \{T, B\}$ and $A_2 = \{L, R\}$
- Action profiles (outcomes):

$$A = A_1 \times A_2 = \{(T, L), (T, R), (B, L), (B, R)\}$$

- Preferences: \succsim_1 and \succsim_2 are given by the bi-matrix.

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 $\{Rock, Paper, Scissors\}$ 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

	<i>Work</i>	<i>Goof</i>
<i>Work</i>	3, 3	0, 4
<i>Goof</i>	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)

	<i>Ball</i>	<i>Show</i>
<i>Ball</i>	2, 1	0, 0
<i>Show</i>	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

	<i>Ball</i>	<i>Show</i>
<i>Ball</i>	2, 2	0, 0
<i>Show</i>	0, 0	1, 1

	<i>Dove</i>	<i>Hawk</i>
<i>Dove</i>	3, 3	1, 4
<i>Hawk</i>	4, 1	0, 0

	<i>Head</i>	<i>Tail</i>
<i>Head</i>	1, -1	-1, 1
<i>Tail</i>	-1, 1	1, -1

Best response and dominated actions

Action T is player 1's *best response* to action L player 2 if T is the optimal choice when 1 *conjectures* that 2 will play L .

Player 1's action T is *strictly* dominated if it is never a best response (inferior to B 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	$1 - q$
		L	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,

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

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,

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

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:

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

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

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).

Oligopoly

(only if time permits)

- Oligopoly is 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) (Antoine Augustin Cournot, an economist, philosopher and mathematician, 1801-1877).

- 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 $A > c$.

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 decreases, approaching c (zero profits).

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 $X^{\max} = \max\{X_1, X_2\}$ 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.