Behavioral game theory
Introduction

Summer 26
Niels Boissonnet, Bielefeld

GT as a model for strategic interations

Many strategic situations in business, politics, markets, bargaining, auctions, voting, or conflict are naturally modeled as games.

GT as a model for strategic interations

Many strategic situations in business, politics, markets, bargaining, auctions, voting, or conflict are naturally modeled as games.

A *game* is a multi-person decision problem defined by its structure:

GT as a model for strategic interations

Many strategic situations in business, politics, markets, bargaining, auctions, voting, or conflict are naturally modeled as games.

A *game* is a multi-person decision problem defined by its structure:

- the players,

GT as a model for strategic interations

Many strategic situations in business, politics, markets, bargaining, auctions, voting, or conflict are naturally modeled as games.

A *game* is a multi-person decision problem defined by its structure:

- the players,
- the rules of interaction (timing, feasible actions, information),

GT as a model for strategic interations

Many strategic situations in business, politics, markets, bargaining, auctions, voting, or conflict are naturally modeled as games.

A *game* is a multi-person decision problem defined by its structure:

- the players,
- the rules of interaction (timing, feasible actions, information),
- how actions determine outcomes,

GT as a model for strategic interations

Many strategic situations in business, politics, markets, bargaining, auctions, voting, or conflict are naturally modeled as games.

A *game* is a multi-person decision problem defined by its structure:

- the players,
- the rules of interaction (timing, feasible actions, information),
- how actions determine outcomes,
- and preferences over outcomes.

Game theory (henceforth GT) provides

- a precise language for all of these elements
- predictions through equilibirum concepts
- the connection between strategic reasoning and these predictions.

Strategic reasoning and GT

Game theory studies strategic reasoning, which combines:



Rationality

+

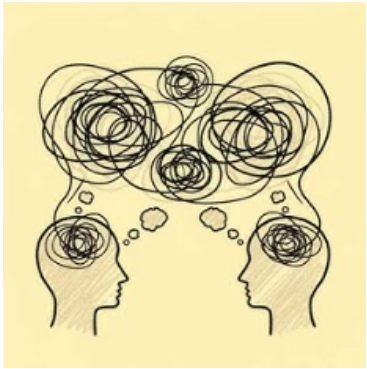
Strategic reasoning and GT

Game theory studies strategic reasoning, which combines:



Rationality

+



Common knowledge of Rationality

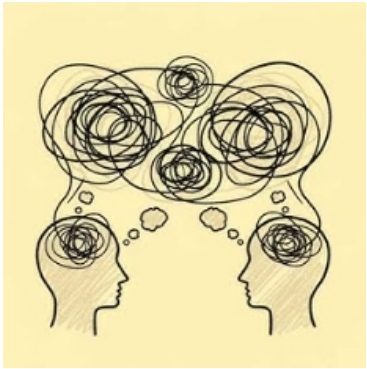
Strategic reasoning and GT

Game theory studies strategic reasoning, which combines:



Rationality

+



Common knowledge of Rationality

This results in several strategic and equilibrium concepts

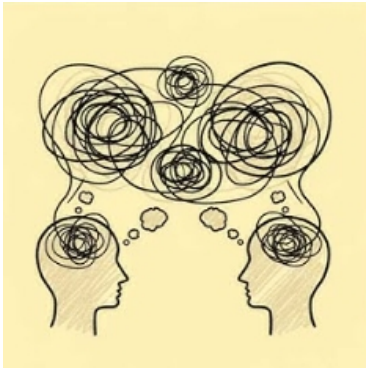
Strategic reasoning and GT

Game theory studies strategic reasoning, which combines:



Rationality

+



Common knowledge of Rationality

This results in several strategic and equilibrium concepts

- Pure and mixed strategies, extensive and repeated games
- Rationalizability, Nash equilibrium, backward induction, subgame perfection, Bayesian Nash equilibrium...

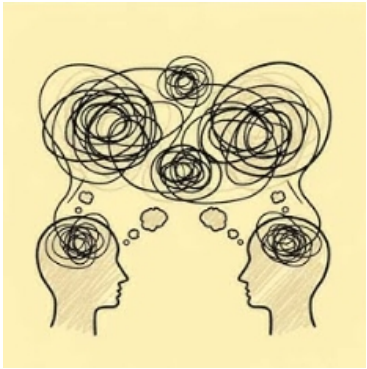
Strategic reasoning and GT

Game theory studies strategic reasoning, which combines:



Rationality

+



Common knowledge of Rationality

This results in several strategic and equilibrium concepts

- Pure and mixed strategies, extensive and repeated games
- Rationalizability, Nash equilibrium, backward induction, subgame perfection, Bayesian Nash equilibrium...

Each involves specific and complex reasoning

What standard GT gives us

The standard benchmark is equilibrium reasoning

Each player is

1. rational: chooses an action that is optimal for her,
2. believes others are rational: given beliefs about what others do,
3. assigns correct beliefs: in equilibrium these beliefs are correct.

What standard GT gives us

The standard benchmark is equilibrium reasoning

Each player is

1. rational: chooses an action that is optimal for her,
2. believes others are rational: given beliefs about what others do,
3. assigns correct beliefs: in equilibrium these beliefs are correct.

In richer environments, the benchmark also uses

1. Bayesian updating under incomplete information,
2. backward induction in sequential games,
3. and consistency of plans and beliefs across time.

The central question

When is equilibrium reasoning descriptively accurate, when does it become fragile, and what do we learn from those failures?

The central question

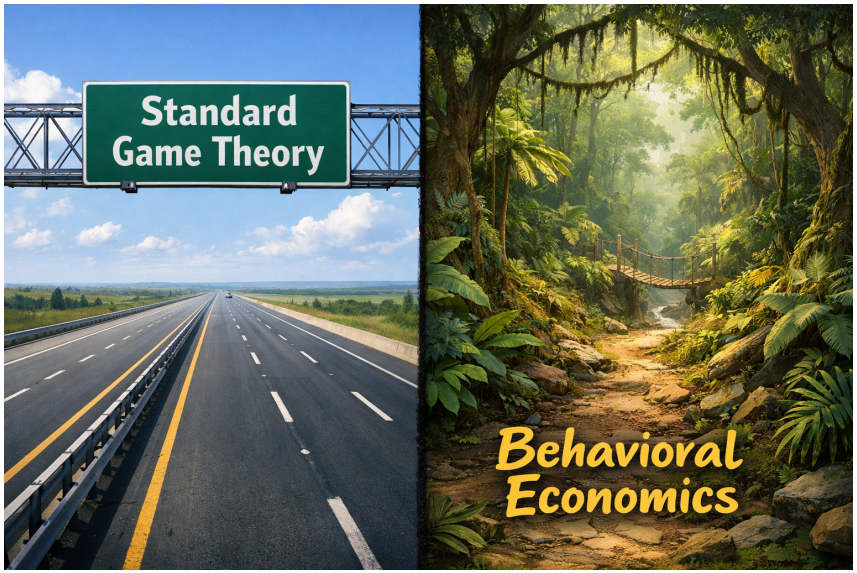
When is equilibrium reasoning descriptively accurate, when does it become fragile, and what do we learn from those failures?

This is the topic of behavioral game theory!

The central question

When is equilibrium reasoning descriptively accurate, when does it become fragile, and what do we learn from those failures?

This is the topic of behavioral game theory!



Goal of BGT

BGT retains the formal structure of GT, adding empirical evidence on

- how people actually reason,
- how people actually behave

in strategic environments.

Goal of BGT

BGT retains the formal structure of GT, adding empirical evidence on

- how people actually reason,
- how people actually behave

in strategic environments.

The goal is not

Replacing game theory with anecdotes.

The goal is

1. understanding where game theory works well,
2. where it falls short,
3. how those shortcomings can be modeled.

When benchmark equilibrium works—when it fails

Benchmark equilibrium concept perform well

- when people have experience, repeated exposure, or institutional guidance.
- they may not solve a game from first principles;
- they may learn to predict from past interaction, routines, or feedback.

When benchmark equilibrium works—when it fails

Benchmark equilibrium concept perform well

- when people have experience, repeated exposure, or institutional guidance.
- they may not solve a game from first principles;
- they may learn to predict from past interaction, routines, or feedback.

But many economically relevant games are novel!

- Then, equilibrium prediction may be inaccurate depending on
  - depth of strategic reasoning,
  - belief formation,
  - ...

When benchmark equilibrium works—when it fails

Benchmark equilibrium concept perform well

- when people have experience, repeated exposure, or institutional guidance.
- they may not solve a game from first principles;
- they may learn to predict from past interaction, routines, or feedback.

But many economically relevant games are novel!

- Then, equilibrium prediction may be inaccurate depending on
  - depth of strategic reasoning,
  - belief formation,
  - ...

Question: are deviations from equilibrium prediction

- merely noise?
- or structured, systematic, and modelable?

How this course studies strategic behavior

We study behavior in

- normal-form games,
- extensive-form games,
- Bayesian games,
- repeated games.

How this course studies strategic behavior

We study behavior in

- normal-form games,
- extensive-form games,
- Bayesian games,
- repeated games.

Recurring structure in each block

1. Solve the benchmark solution to the problem

How this course studies strategic behavior

We study behavior in

- normal-form games,
- extensive-form games,
- Bayesian games,
- repeated games.

Recurring structure in each block

1. Solve the benchmark solution to the problem
2. Check experiments, historical examples, or robust intuitions.

How this course studies strategic behavior

We study behavior in

- normal-form games,
- extensive-form games,
- Bayesian games,
- repeated games.

Recurring structure in each block

1. Solve the benchmark solution to the problem
2. Check experiments, historical examples, or robust intuitions.
3. Identify systematic (rather than accidental) deviations from the benchmark.

How this course studies strategic behavior

We study behavior in

- normal-form games,
- extensive-form games,
- Bayesian games,
- repeated games.

Recurring structure in each block

1. Solve the benchmark solution to the problem
2. Check experiments, historical examples, or robust intuitions.
3. Identify systematic (rather than accidental) deviations from the benchmark.
4. Study the behavioral models proposed to explain those deviations.

Hence the recurring logic of the course

Standard GT $\longrightarrow$ evidence $\longrightarrow$ behavioral explanation $\longrightarrow$ interpretation

This is also how the modern field presents itself.

Normal-form games

Behavioral game theoryIntroduction

A puzzle: the beauty-contest game

p-Beauty contest

- Players choose a real number from $[0, 100]$.
- They win if their choice is closest to $\frac{2}{3}$ of the average of all choices
- the prize is shared in case of ties



Which one do you choose???

Beauty contest

Simultaneous-move game: players choose a number without observing others' choices.

Beauty contest

Simultaneous-move game: players choose a number without observing others' choices.

Playing 0 =NE: Let x to be the average number of an equilibrium:

Beauty contest

Simultaneous-move game: players choose a number without observing others' choices.

Playing 0 =NE: Let x to be the average number of an equilibrium:

- If $x = 0$, i.e. everybody plays 0, no player has interest to deviate

  $\Longrightarrow$ Clear Nash equilibrium

Beauty contest

Simultaneous-move game: players choose a number without observing others' choices.

Playing 0 =NE: Let x to be the average number of an equilibrium:

- If $x = 0$, i.e. everybody plays 0, no player has interest to deviate
  $\Longrightarrow$ Clear Nash equilibrium
- If $x > 0$, at least one player has interest to give a lower number.
  $\Longrightarrow$ Cannot be an equilibrium

The point of this game

This game aims at highlighting the role of higher-order beliefs in equilibrium

- reasoning about others' reasoning, iteratively.

The point of this game

This game aims at highlighting the role of higher-order beliefs in equilibrium

- reasoning about others' reasoning, iteratively.

Iterated best-response reasoning:

- Since 100 is maximum, every number above $\frac{2}{3} \cdot 100 = 67$ is dominated: no one should play it in equilibrium;

The point of this game

This game aims at highlighting the role of higher-order beliefs in equilibrium

- reasoning about others' reasoning, iteratively.

Iterated best-response reasoning:

- Since 100 is maximum, every number above $\frac{2}{3} \cdot 100 = 67$ is dominated: no one should play it in equilibrium;
- But then if everybody plays $\frac{2}{3} \cdot 67$, the best response is $\left(\frac{2}{3}\right)^2 \cdot 100$;

The point of this game

This game aims at highlighting the role of higher-order beliefs in equilibrium

- reasoning about others' reasoning, iteratively.

Iterated best-response reasoning:

- Since 100 is maximum, every number above $\frac{2}{3} \cdot 100 = 67$ is dominated: no one should play it in equilibrium;
- But then if everybody plays $\frac{2}{3} \cdot 67$, the best response is $\left(\frac{2}{3}\right)^2 \cdot 100$;
- Repeating this reasoning yields $\left(\frac{2}{3}\right)^n \cdot 100 \to 0$.

The point of this game

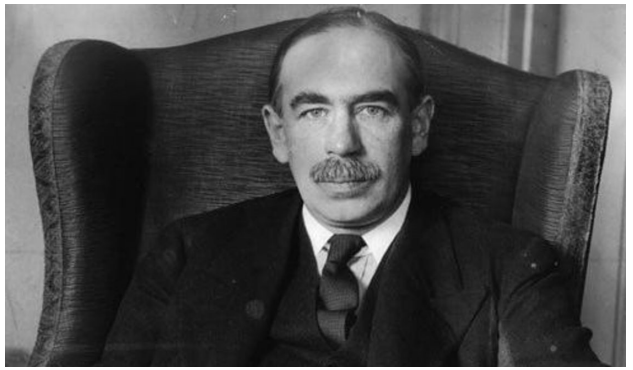
This game aims at highlighting the role of higher-order beliefs in equilibrium

- reasoning about others' reasoning, iteratively.

Iterated best-response reasoning:

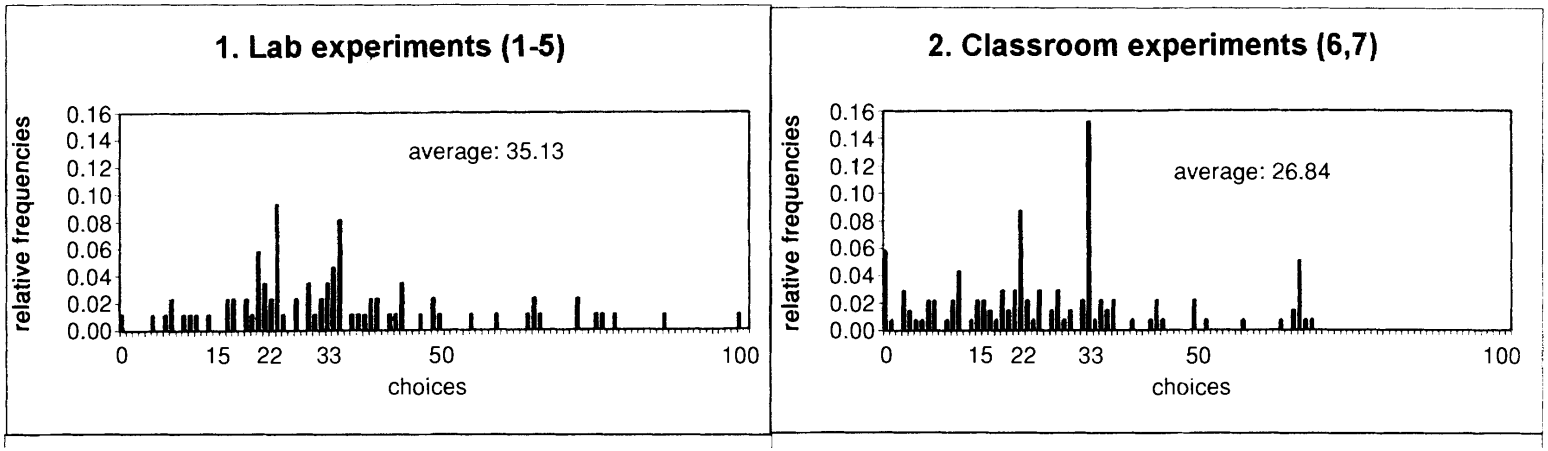
- Since 100 is maximum, every number above $\frac{2}{3} \cdot 100 = 67$ is dominated: no one should play it in equilibrium;
- But then if everybody plays $\frac{2}{3} \cdot 67$, the best response is $\left(\frac{2}{3}\right)^2 \cdot 100$;
- Repeating this reasoning yields $\left(\frac{2}{3}\right)^n \cdot 100 \to 0$.
- The unique Nash equilibrium is 0;

General Theory of Keynes

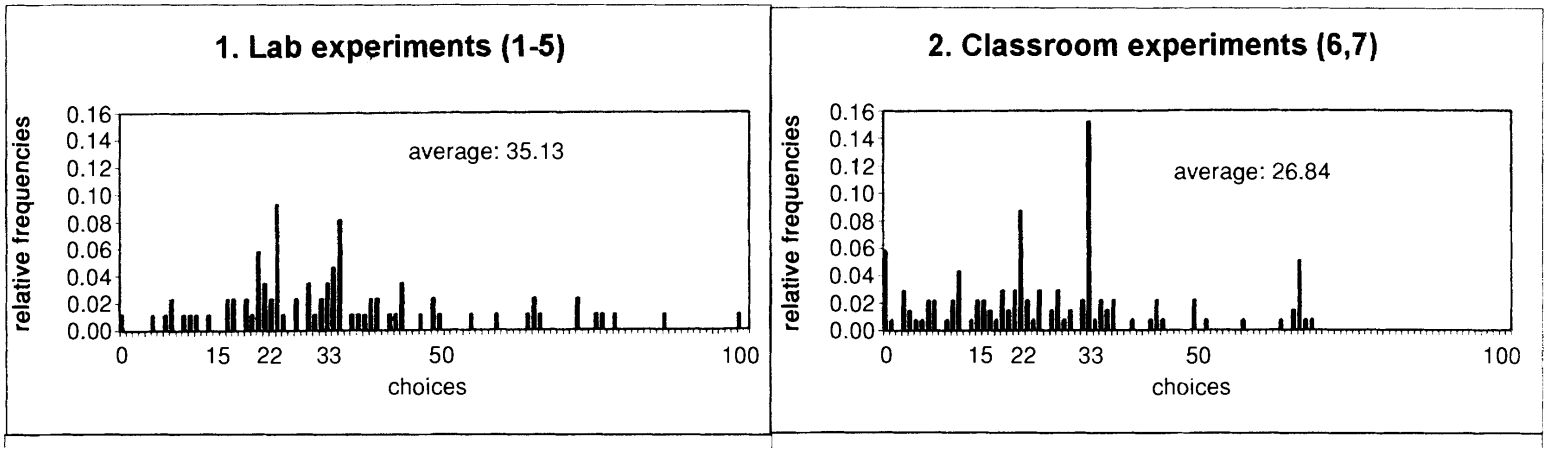


[P]rofessional investment may be likened to those newspaper competitions in which the competitors have to pick out the six prettiest faces from a hundred photographs, the winner being the one who comes nearest to the average preferences of the competitors as a whole; so that each competitor has to pick, not those faces which he himself finds prettiest, but those which he thinks likeliest to catch the fancy of the other competitors, all of whom are looking at the problem from the same point of view.
It is not a case of choosing those which, to the best of one's judgment, are really the prettiest, nor even those which 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 practise the fourth, fifth and higher degrees.

Lessons from the data

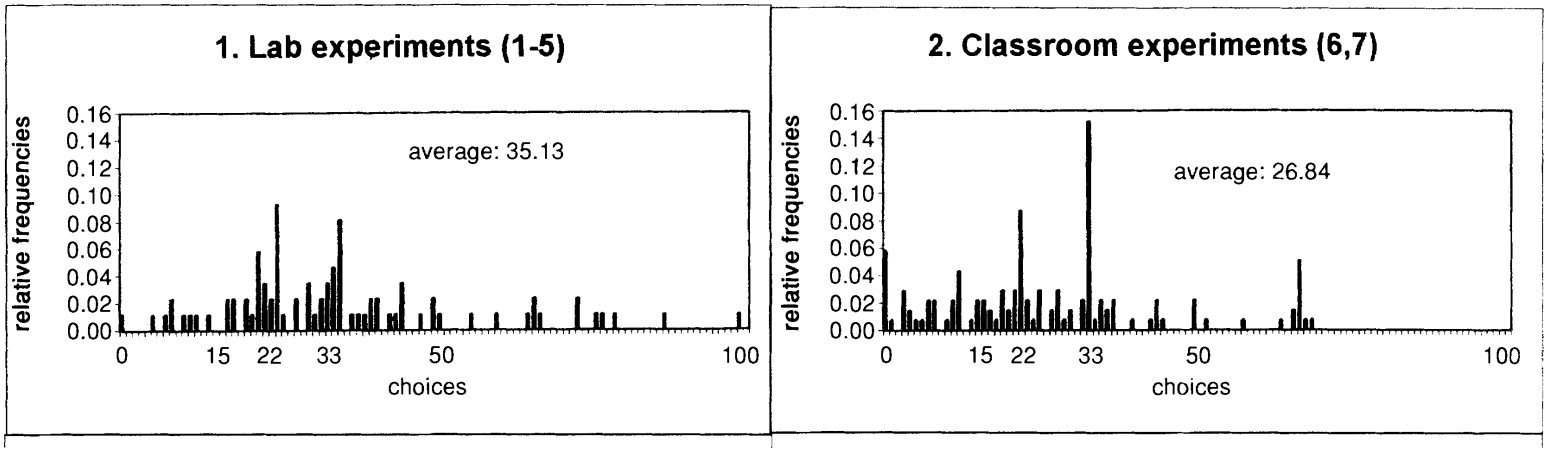


Lessons from the data



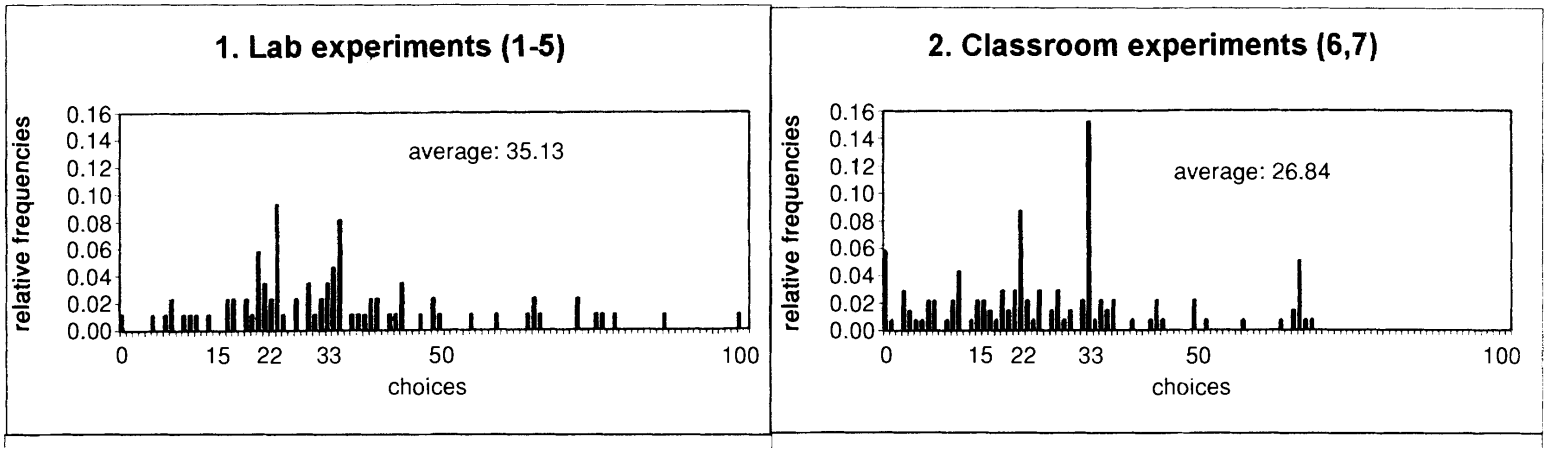
- Data typically show substantial mass well above the equilibrium prediction.

Lessons from the data



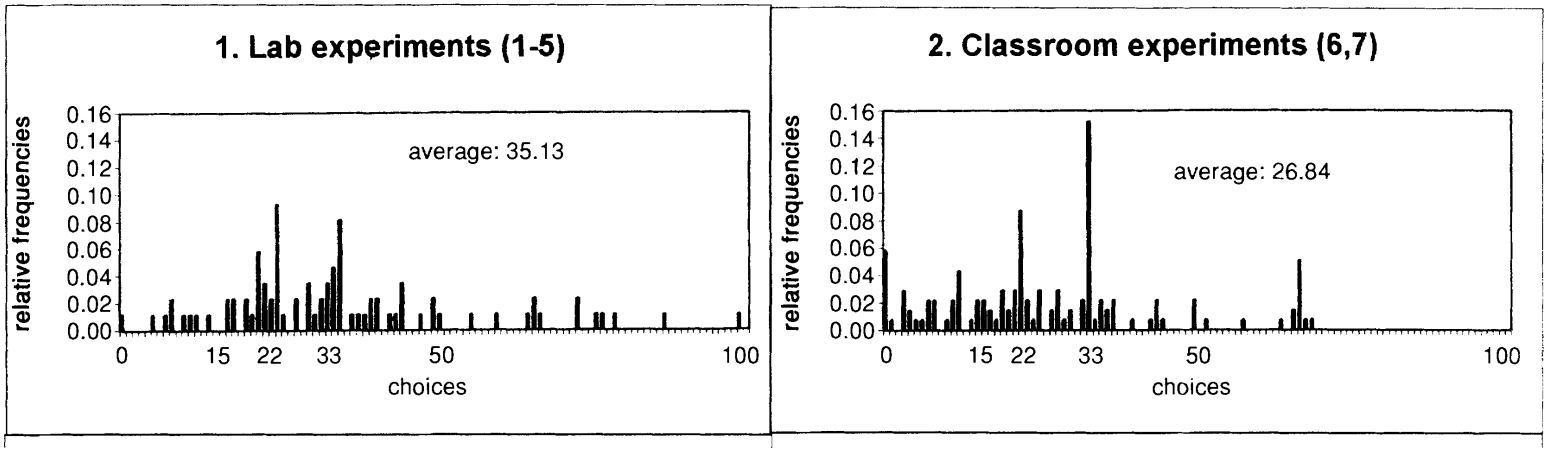
- Data typically show substantial mass well above the equilibrium prediction.
- Yet, choices are not random: they concentrate around a few salient levels.

Lessons from the data



- Data typically show substantial mass well above the equilibrium prediction.
- Yet, choices are not random: they concentrate around a few salient levels.
- Natural reading: as limited numbers of reasoning steps:
  - some respond to a simple average, and others best respond to that.
  - Yet there also seems to be noise or error (Quantil response)

Lessons from the data



- Data typically show substantial mass well above the equilibrium prediction.
- Yet, choices are not random: they concentrate around a few salient levels.
- Natural reading: as limited numbers of reasoning steps:
  - some respond to a simple average, and others best respond to that.
  - Yet there also seems to be noise or error (Quantil response)
- Hence the empirical problem is not simply "failure of rationality," but heterogeneous and finite strategic reasoning.

Equilibrium as a weak guide for coordination

- The issues with coordination games differs from the beauty contest:
  - not too little sophistication, but the existence of several equilibria.
- The main question becomes: how do players form expectations about *which* equilibrium others will aim for?

We will see that

- groups often coordinate, but not necessarily on the efficient outcome.
- history, labels, salience, and expectations affect which equilibrium is selected.

What normal-form games will teach us in this course

Strategic issues

- dominance and iterated reasoning
- Nash equilibrium/mixed strategies
- equilibrium selection
- the role of beliefs in one-shot play
- ambiguity/strategic uncertainty

Canonical evidence

- guessing games
- small matrix games
- coordination games

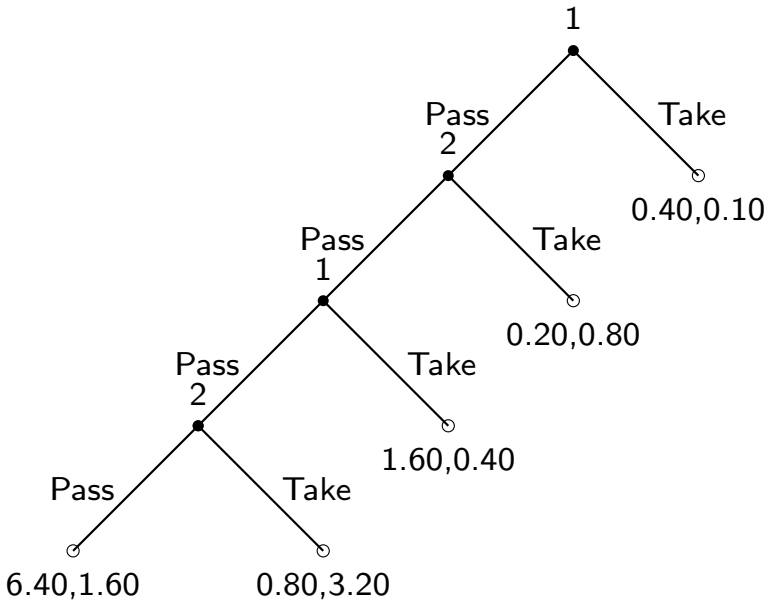
Behavioral models

- QRE (McKelvey and Palfrey, 1995)
- AQRE, truncated QRE (Weizsäcker, 2003; Rogers et al., 2009)
- level-k, CH (Stahl and Wilson, 1995; Camerer, 2004)
- noisy introspection (Goeree and Holt, 2004)
- impulse balance, sampling equilibrium (Ockenfels and Selten, 2005; Osborne and Rubinstein, 1998; Osborne, 2003)

Extensive-form games

Behavioral game theoryIntroduction

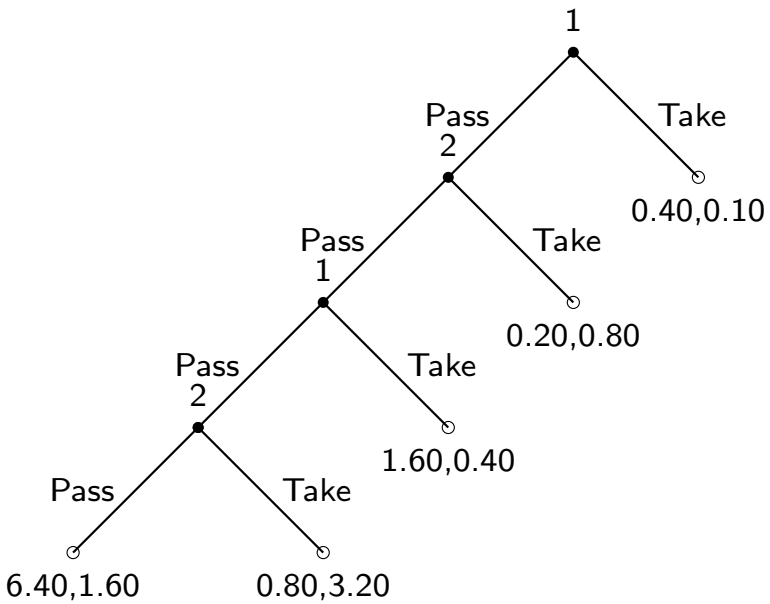
Centipede game



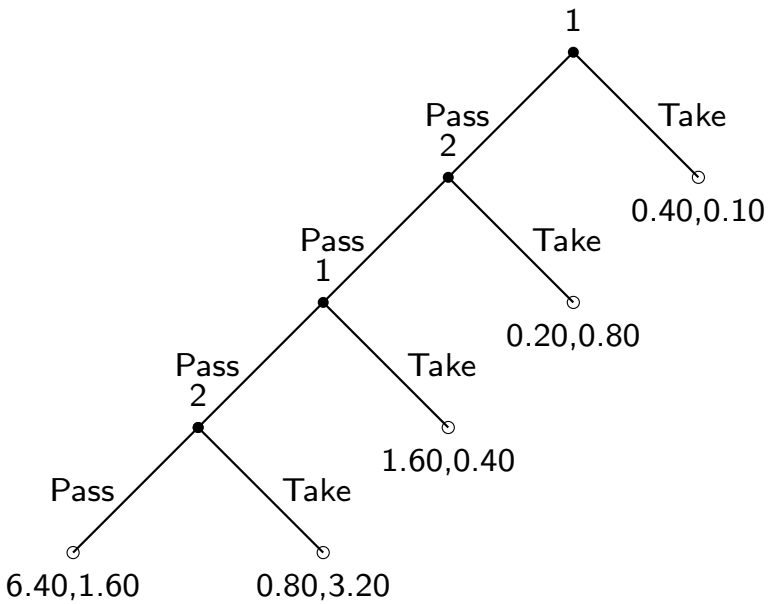
1. P1 chooses Take (and the payoffs are $0.40, 0.10$) or Pass...
2. If P1 chose Pass, then P2 can choose Take (and the payoffs are $0.20, 0.80$) or Pass...
3. If P2 chose Pass, then P1 can choose Take (and the payoffs are $1.60, 0.40$) or Pass...
4. If P1 chose Pass, then P2 can choose Take (and the payoffs are $0.80, 3.20$) or Pass (and the payoffs are $6.40, 1.60$)

This is a complete information extensive game so we can apply backward induction.

From backward induction

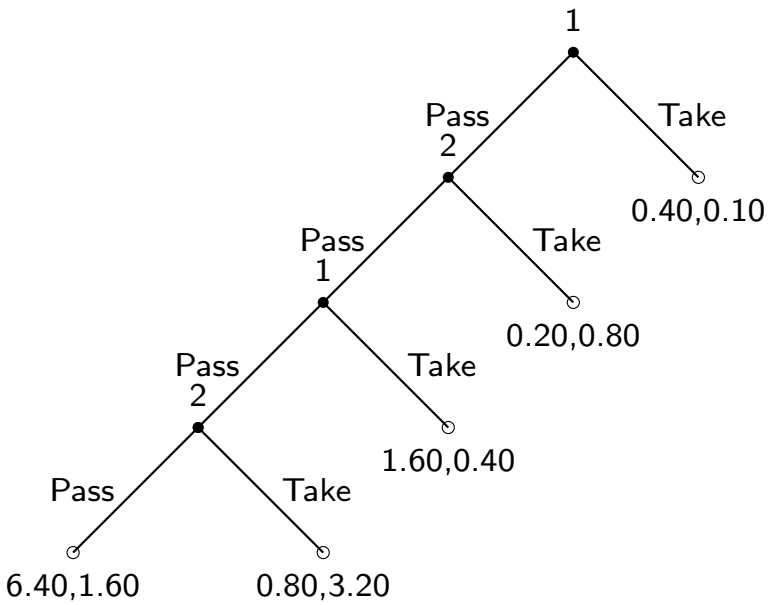


From backward induction



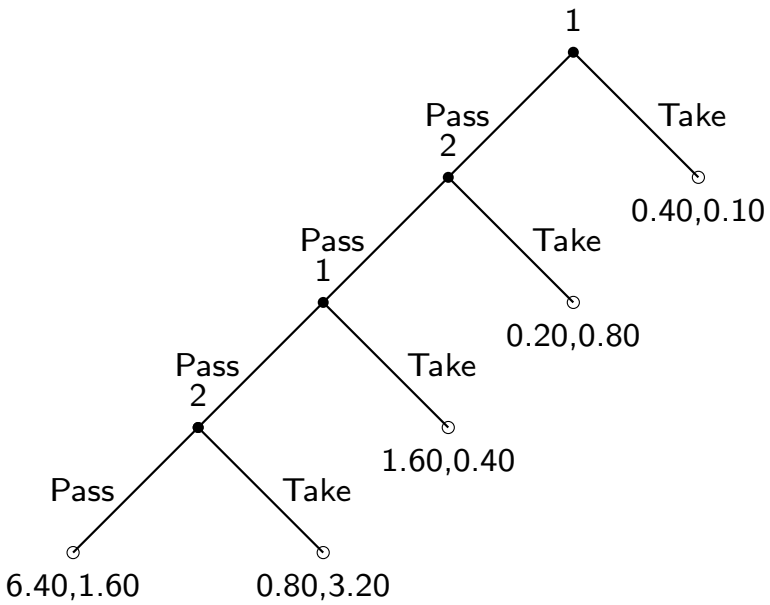
1. Starting at the last node, player 2 should take ($3.2 > 1.6$)

From backward induction



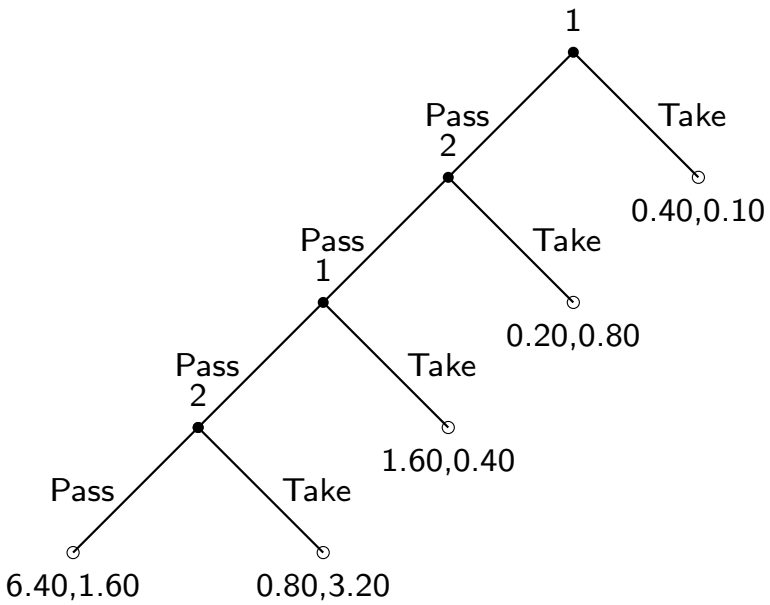
1. Starting at the last node, player 2 should take ($3.2 > 1.6$)
2. Player 1 knows this, so before this node, he should take ($1.6 > 0.8$)

From backward induction



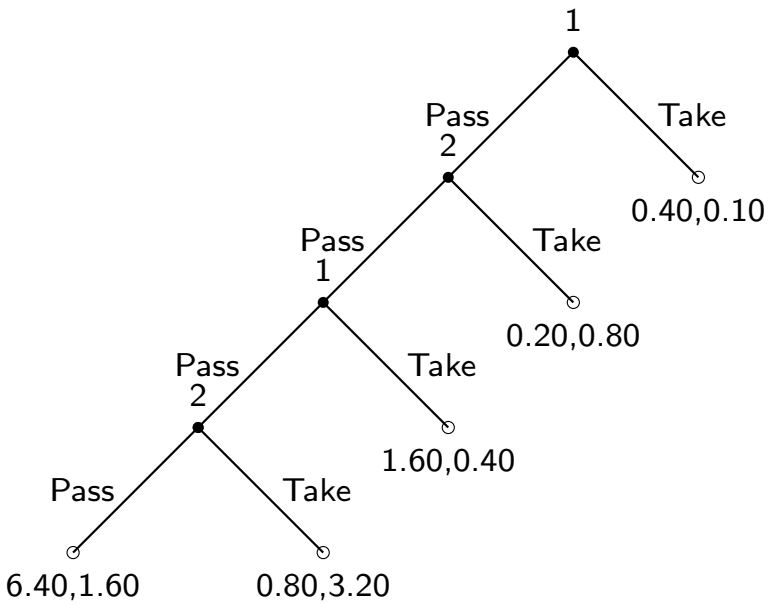
1. Starting at the last node, player 2 should take ($3.2 > 1.6$)
2. Player 1 knows this, so before this node, he should take ($1.6 > 0.8$)
3. Player 2 knows this, so before this node, he should take ($0.8 > 0.4$)

From backward induction



1. Starting at the last node, player 2 should take ($3.2 > 1.6$)
2. Player 1 knows this, so before this node, he should take ($1.6 > 0.8$)
3. Player 2 knows this, so before this node, he should take ($0.8 > 0.4$)
4. Player 2 knows this, so before this node, he should take ($0.4 > 0.2$)

From backward induction



1. Starting at the last node, player 2 should take ($3.2 > 1.6$)
2. Player 1 knows this, so before this node, he should take ($1.6 > 0.8$)
3. Player 2 knows this, so before this node, he should take ($0.8 > 0.4$)
4. Player 2 knows this, so before this node, he should take ($0.4 > 0.2$)

$\Longrightarrow$ In equilibrium player 1 takes in the first round

Empirical findings

In the lab:

- Subjects often pass for several rounds before taking.
- Are these deviations from the benchmark predictions systematic or do players play randomly?

Empirical findings

In the lab:

- Subjects often pass for several rounds before taking.
- Are these deviations from the benchmark predictions systematic or do players play randomly?

The point of this game: Sequential play makes reasoning demands visible.

- Does the failure comes from limited look-ahead, uncertainty about others, noisy choice, or a different interpretation of the game?

Incomplete Information: Cascade games

- One the following Urns is selected with probability $\frac{1}{2}$:

$$A = (\bullet, \bullet, \bullet) \qquad B = (\bullet, \bullet, \bullet)$$

Incomplete Information: Cascade games

- One the following Urns is selected with probability $\frac{1}{2}$:

$$A = (\bullet, \bullet, \bullet) \qquad B = (\bullet, \bullet, \bullet)$$

- 3 players (P_1, P_2, P_3) have to correctly guess which urn is selected they only know the probabilities that one urn be selected

Play is sequential and information partially disclosed

Incomplete Information: Cascade games

- One the following Urns is selected with probability $\frac{1}{2}$:

$$A = (\bullet, \bullet, \bullet) \qquad B = (\bullet, \bullet, \bullet)$$

- 3 players (P_1, P_2, P_3) have to correctly guess which urn is selected they only know the probabilities that one urn be selected

Play is sequential and information partially disclosed

- P_1 draws a ball from the chosen urn and selects an urn;

Incomplete Information: Cascade games

- One the following Urns is selected with probability $\frac{1}{2}$:

$$A = (\bullet, \bullet, \bullet) \qquad B = (\bullet, \bullet, \bullet)$$

- 3 players (P_1, P_2, P_3) have to correctly guess which urn is selected they only know the probabilities that one urn be selected

Play is sequential and information partially disclosed

- P_1 draws a ball from the chosen urn and selects an urn;
- P_2 observes P_1's choice, draws a ball, and selects an urn;

Incomplete Information: Cascade games

- One the following Urns is selected with probability $\frac{1}{2}$:

$$A = (\bullet, \bullet, \bullet) \qquad B = (\bullet, \bullet, \bullet)$$

- 3 players (P_1, P_2, P_3) have to correctly guess which urn is selected they only know the probabilities that one urn be selected

Play is sequential and information partially disclosed

- P_1 draws a ball from the chosen urn and selects an urn;
- P_2 observes P_1's choice, draws a ball, and selects an urn;
- P_3 observes previous choices, draws a ball, and selects an urn;

In equilibrium

A feature of perfect Bayesian equilibrium is:

- If P_3 observes P_1 and P_2 choosing urn A, she must imitate their choice, ignoring the result of her own drawing.

Intuition:

- he infers that player 1 has drawn a red ball
- he infers that player 2 has drawn *probably* drawn a red ball

$\Longrightarrow$ more red balls $\Longrightarrow$ choosing A is optimal

Empirical findings

The point of this game

- Testing whether and how people use information they can infer from others' choice

Empirical findings

The point of this game

- Testing whether and how people use information they can infer from others' choice

In the lab:

- Individuals generally use information efficiently
- Follow the decisions of others when rational.

Empirical findings

The point of this game

- Testing whether and how people use information they can infer from others' choice

In the lab:

- Individuals generally use information efficiently
- Follow the decisions of others when rational.
- However, some errors suggest that they rely more on their own private information.

Topics in extensive-form games

Strategic issues

- timing and information sets
- strategies in trees
- backward induction and subgame perfection
- inference from observed moves
- strategic interpretation of off-path behavior

Canonical evidence

- centipede games
- backward-induction experiments
- forward-induction tests
- information cascades

Behavioral themes

- limited look-ahead and noisy sequential choice
- AQRE and related agent-based error models
- belief-based interpretations of early actions
- social learning and cascade behavior

Preview of later lectures

- BI: (Binmore et al., 2002; Johnson et al., 2002)
- centipede, AQRE: (McKelvey and Palfrey, 1998)
- cascades: (Kübler and Weizsäcker, 2004; Goeree et al., 2007; Weizsäcker, 2010)

Bayesian games

Behavioral game theoryIntroduction

A simple auction

A first-price auction

Two bidders which value for object is an independent random draw from the uniform distribution on $[0, 100]$. You simultaneously submit bids without observing the other bid. The highest bidder wins and pays her bid.

A simple auction

A first-price auction

Two bidders which value for object is an independent random draw from the uniform distribution on $[0, 100]$. You simultaneously submit bids without observing the other bid. The highest bidder wins and pays her bid.

- Suppose your value is 50. How much do you bid?

A simple auction

A first-price auction

Two bidders which value for object is an independent random draw from the uniform distribution on $[0, 100]$. You simultaneously submit bids without observing the other bid. The highest bidder wins and pays her bid.

- Suppose your value is 50. How much do you bid?
- Suppose your value is 80. How much do you bid now?

A simple auction

A first-price auction

Two bidders which value for object is an independent random draw from the uniform distribution on $[0, 100]$. You simultaneously submit bids without observing the other bid. The highest bidder wins and pays her bid.

- Suppose your value is 50. How much do you bid?
- Suppose your value is 80. How much do you bid now?
- Suppose your value is only 20. How much do you bid then?

- The basic intuition is clear: you should usually bid below your value.
- But how far below value should you bid?

A simple auction

A first-price auction

Two bidders which value for object is an independent random draw from the uniform distribution on $[0, 100]$. You simultaneously submit bids without observing the other bid. The highest bidder wins and pays her bid.

- Suppose your value is 50. How much do you bid?
- Suppose your value is 80. How much do you bid now?
- Suppose your value is only 20. How much do you bid then?

- The basic intuition is clear: you should usually bid below your value.
- But how far below value should you bid?

  $\Longrightarrow$ depends on your beliefs about the other value and bidding behavior.
- This is a Bayesian game:

  $\Longrightarrow$ The strategic problem is more complex than in a complete-information.

Why this is a Bayesian game

- In a Bayesian game, players do not know the type of the others.

Why this is a Bayesian game

- In a Bayesian game, players do not know the type of the others.
- Here, each bidder knows her own value v_i, but not the opponent's value v_j.
- a strategy is no longer just an action:

  $\Longrightarrow$ it is a rule that maps types into actions, here a bidding function $b_i(v_i)$.

Why this is a Bayesian game

- In a Bayesian game, players do not know the type of the others.
- Here, each bidder knows her own value v_i, but not the opponent's value v_j.
- a strategy is no longer just an action:

  $\Longrightarrow$ it is a rule that maps types into actions, here a bidding function $b_i(v_i)$.

In Bayesian games, players must reason about others' actions, *but also* about hidden others' payoffs, making **belief formation central**.

Why this is a Bayesian game

- In a Bayesian game, players do not know the type of the others.
- Here, each bidder knows her own value v_i, but not the opponent's value v_j.
- a strategy is no longer just an action:

  $\Longrightarrow$ it is a rule that maps types into actions, here a bidding function $b_i(v_i)$.

In Bayesian games, players must reason about others' actions, *but also* about hidden others' payoffs, making **belief formation central**.

Empirical puzzle: In many auction environments, bidders bid more aggressively than this benchmark predicts. Why?

What Bayesian games will teach us in this course

Strategic issues

- incomplete information
- Bayesian Nash equilibrium
- auctions and mechanism design
- private values, common values
- purification/revelation principle

Canonical evidence

- first-price auctions
- second-price auctions
- overbidding and underbidding
- winner's-curse effects

Behavioral models

- risk aversion (Goeree et al., 2002; Bajari and Hortacsu, 2005)
- joy of winning
- cursed equilibrium (Eyster and Rabin, 2005)
- level-k in auctions (Crawford and Iriberri, 2007)
- type projection (Breitmoser, 2019)

Preview of later lectures

- classic auction evidence: (Kagel et al., 1987; Dyer et al., 1989; Goeree and Offerman, 2002)
- modern incomplete-information puzzles

Repeated games

Behavioral game theoryIntroduction

A first repeated-game puzzle: cooperation over time

	c	d
c	$3, 3$	$0, 4$
d	$4, 0$	$1, 1$

One-shot PD:

defection strictly dominates cooperation.

- In a one-shot interaction, the benchmark prediction is clear: both players defect.

A first repeated-game puzzle: cooperation over time

	c	d
c	$3, 3$	$0, 4$
d	$4, 0$	$1, 1$

One-shot PD:

defection strictly dominates cooperation.

- In a one-shot interaction, the benchmark prediction is clear: both players defect.
- In a repeated interaction, current behavior can affect future behavior.

A first repeated-game puzzle: cooperation over time

	c	d
c	$3, 3$	$0, 4$
d	$4, 0$	$1, 1$

One-shot PD:

defection strictly dominates cooperation.

- In a one-shot interaction, the benchmark prediction is clear: both players defect.
- In a repeated interaction, current behavior can affect future behavior.
- This creates room for **rewards, punishments, reputation, and cooperation**.

A first repeated-game puzzle: cooperation over time

	c	d
c	$3, 3$	$0, 4$
d	$4, 0$	$1, 1$

One-shot PD:

defection strictly dominates cooperation.

- In a one-shot interaction, the benchmark prediction is clear: both players defect.
- In a repeated interaction, current behavior can affect future behavior.
- This creates room for **rewards, punishments, reputation, and cooperation**.

Repetition shifts the strategic question:

A first repeated-game puzzle: cooperation over time

	c	d
c	$3, 3$	$0, 4$
d	$4, 0$	$1, 1$

One-shot PD:

defection strictly dominates cooperation.

- In a one-shot interaction, the benchmark prediction is clear: both players defect.
- In a repeated interaction, current behavior can affect future behavior.
- This creates room for **rewards, punishments, reputation, and cooperation**.

Repetition shifts the strategic question:

- From: *"What is optimal in this round?"*

A first repeated-game puzzle: cooperation over time

	c	d
c	$3, 3$	$0, 4$
d	$4, 0$	$1, 1$

One-shot PD:

defection strictly dominates cooperation.

- In a one-shot interaction, the benchmark prediction is clear: both players defect.
- In a repeated interaction, current behavior can affect future behavior.
- This creates room for **rewards, punishments, reputation, and cooperation**.

Repetition shifts the strategic question:

- From: *"What is optimal in this round?"*
- To: *"What pattern of play can be sustained over time?"*

Why repeated interaction changes the problem

- In repeated games, strategies can be conditioned on past play.

Why repeated interaction changes the problem

- In repeated games, strategies can be conditioned on past play.
- A player's action today can therefore influence the opponent's future actions.

Why repeated interaction changes the problem

- In repeated games, strategies can be conditioned on past play.
- A player's action today can therefore influence the opponent's future actions.
- This creates a new strategic logic:
  - cooperation may be sustained by the **threat of future punishment**,

Why repeated interaction changes the problem

- In repeated games, strategies can be conditioned on past play.
- A player's action today can therefore influence the opponent's future actions.
- This creates a new strategic logic:
  - cooperation may be sustained by the **threat of future punishment**,
  - deviations may be deterred if **continuation is sufficiently likely**,

Why repeated interaction changes the problem

- In repeated games, strategies can be conditioned on past play.
- A player's action today can therefore influence the opponent's future actions.
- This creates a new strategic logic:
  - cooperation may be sustained by the **threat of future punishment**,
  - deviations may be deterred if **continuation is sufficiently likely**,
  - but coordinating on a cooperative path is itself non-trivial.
- Example of sustainable strategy: Tit for tat

Repeated games: empirics and the behavioral puzzle

- Empirically, repeated PD experiments show **substantial cooperation**, especially when continuation is likely and incentives align.

Repeated games: empirics and the behavioral puzzle

- Empirically, repeated PD experiments show **substantial cooperation**, especially when continuation is likely and incentives align.
- Yet cooperation is far from guaranteed: outcomes vary widely across groups and settings.

Repeated games: empirics and the behavioral puzzle

- Empirically, repeated PD experiments show **substantial cooperation**, especially when continuation is likely and incentives align.
- Yet cooperation is far from guaranteed: outcomes vary widely across groups and settings.

Behavioral puzzle

Even when cooperation is theoretically possible, it does not arise automatically. Why do some groups sustain it, while others unravel quickly?

Repeated games: empirics and the behavioral puzzle

- Empirically, repeated PD experiments show **substantial cooperation**, especially when continuation is likely and incentives align.
- Yet cooperation is far from guaranteed: outcomes vary widely across groups and settings.

Behavioral puzzle

Even when cooperation is theoretically possible, it does not arise automatically. Why do some groups sustain it, while others unravel quickly?

Key factors we will examine:

- Role of the **discount factor** and horizon length

Repeated games: empirics and the behavioral puzzle

- Empirically, repeated PD experiments show **substantial cooperation**, especially when continuation is likely and incentives align.
- Yet cooperation is far from guaranteed: outcomes vary widely across groups and settings.

Behavioral puzzle

Even when cooperation is theoretically possible, it does not arise automatically. Why do some groups sustain it, while others unravel quickly?

Key factors we will examine:

- Role of the **discount factor** and horizon length
- Sensitivity to **strategies** (tit-for-tat, grim trigger, . . .)

Repeated games: empirics and the behavioral puzzle

- Empirically, repeated PD experiments show **substantial cooperation**, especially when continuation is likely and incentives align.
- Yet cooperation is far from guaranteed: outcomes vary widely across groups and settings.

Behavioral puzzle

Even when cooperation is theoretically possible, it does not arise automatically. Why do some groups sustain it, while others unravel quickly?

Key factors we will examine:

- Role of the **discount factor** and horizon length
- Sensitivity to **strategies** (tit-for-tat, grim trigger, . . .)
- Effects of **communication, norms, and experience**

What theory leaves open: behavioral questions

These tools characterise *what can* be sustained — but not *what will* be. Several behavioral questions remain:

- Do players **believe** others will actually carry out threatened punishments?

What theory leaves open: behavioral questions

These tools characterise *what can* be sustained — but not *what will* be. Several behavioral questions remain:

- Do players **believe** others will actually carry out threatened punishments?
- Can they **coordinate** on a cooperative strategy in the first place?

What theory leaves open: behavioral questions

These tools characterise *what can* be sustained — but not *what will* be. Several behavioral questions remain:

- Do players **believe** others will actually carry out threatened punishments?
- Can they **coordinate** on a cooperative strategy in the first place?
- How important are **errors, limited foresight, and heterogeneity** in strategies?

What theory leaves open: behavioral questions

These tools characterise *what can* be sustained — but not *what will* be. Several behavioral questions remain:

- Do players **believe** others will actually carry out threatened punishments?
- Can they **coordinate** on a cooperative strategy in the first place?
- How important are **errors, limited foresight, and heterogeneity** in strategies?

Broader takeaway

Repeated games are not just about whether cooperation *is* an equilibrium outcome, but whether players can build, interpret, and maintain cooperation dynamically.

What repeated games will teach us in this course

Strategic issues

- one-shot deviation principle
- finite versus indefinite repetition
- punishment, credibility, incentives
- cooperation and collusion
- equilibrium multiplicity

Canonical evidence

- repeated prisoner's dilemma
- cooperation rates over time
- breakdown and recovery
- elicited beliefs/continuation play

Behavioral themes

- heterogeneity in repeated-game
- belief formation/belief updating
- implementation errors/noisy responses
- learning and experience effects

Preview of later lectures

- repeated PD evidence: (Dal Bó, 2005; Dal Bó and Fréchette, 2011)
- cooperation and repeated-game theory: (Breitmoser, 2015)

Overview

Normal-form games

- simultaneous choice
- strategic depth and iterated reasoning
- beliefs about what others will do
- multiplicity and equilibrium selection

Extensive-form games

- timing and information sets
- backward induction and subgame perfection
- interpretation of observed moves

Bayesian games

- incomplete information and private types
- beliefs over hidden characteristics
- bidding and strategic uncertainty

Repeated games

- incentives over time
- cooperation, punishment, credibility
- beliefs and learning across rounds
- breakdown and recovery of cooperation

What comes next?

Next lecture: normal-form benchmark and first behavioral models

- dominance and iterated dominance,
- best responses and Nash equilibrium,
- mixed strategies as a benchmark,
- first systematic deviations in one-shot simultaneous-move play.

You should review beforehand

- strategic-form representation of games,
- pure and mixed Nash equilibrium,
- dominance and rationalizability,
- the logic of best responses.

Guiding question

How much strategic sophistication can we infer from initial play in simple one-shot games?

Literature

Behavioral game theoryIntroduction

References for Lecture 01 I

Bajari, P. and Hortacsu, A. (2005). Are structural estimates of auction models reasonable? evidence from experimental data. *Journal of Political Economy*, 113(4):703–741.

Binmore, K., McCarthy, J., Ponti, G., Samuelson, L., and Shaked, A. (2002). A backward induction experiment. *Journal of Economic theory*, 104(1):48–88.

Breitmoser, Y. (2015). Cooperation, but no reciprocity: Individual strategies in the repeated prisoner's dilemma. *American Economic Review*, 105(9):2882–2910.

Breitmoser, Y. (2019). Knowing me, imagining you: Projection and overbidding in auctions. *Games and Economic Behavior*, 113:423–447.

Camerer, C. F. (2004). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press.

Crawford, V. P. and Iriberri, N. (2007). Level-k auctions: Can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions? *Econometrica*, 75(6):1721–1770.

Dal Bó, P. (2005). Cooperation under the shadow of the future: Experimental evidence from infinitely repeated games. *American Economic Review*, 95(5):1591–1604.

Dal Bó, P. and Fréchette, G. R. (2011). The evolution of cooperation in infinitely repeated games: Experimental evidence. *American Economic Review*, 101(1):411–429.

Dyer, D., Kagel, J. H., and Levin, D. (1989). A comparison of naive and experienced bidders in common value offer auctions: A laboratory analysis. *Economic Journal*, 99(394):108–115.

Eyster, E. and Rabin, M. (2005). Cursed equilibrium. *Quarterly Journal of Economics*, 120(3):1643–1672.

References for Lecture 01 II

Goeree, J. K. and Holt, C. A. (2004). A model of noisy introspection. *Games and Economic Behavior*, 46(2):365–382.

Goeree, J. K., Holt, C. A., and Palfrey, T. R. (2002). Quantal response equilibrium and overbidding in private-value auctions. *Journal of Economic Theory*, 104(1):247–272.

Goeree, J. K. and Offerman, T. (2002). Efficiency in auctions with private and common values: An experimental study. *American Economic Review*, 92(3):625–643.

Goeree, J. K., Palfrey, T. R., Rogers, B. W., and McKelvey, R. D. (2007). Self-correcting information cascades. *The Review of Economic Studies*, 74(3):733–762.

Johnson, E. J., Camerer, C., Sen, S., and Rymon, T. (2002). Detecting failures of backward induction: Monitoring information search in sequential bargaining. *Journal of Economic Theory*, 104(1):16–47.

Kagel, J. H., Harstad, R. M., and Levin, D. (1987). Information impact and allocation rules in auctions with affiliated private values: A laboratory study. *Econometrica*, 55(6):1275–1304.

Kübler, D. and Weizsäcker, G. (2004). Limited depth of reasoning and failure of cascade formation in the laboratory. *The Review of Economic Studies*, 71(2):425–441.

McKelvey, R. D. and Palfrey, T. R. (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1):6–38.

McKelvey, R. D. and Palfrey, T. R. (1998). Quantal response equilibria for extensive form games. *Experimental Economics*, 1(1):9–41.

References for Lecture 01 III

Ockenfels, A. and Selten, R. (2005). Impulse balance equilibrium and feedback in first price auctions. *Games and Economic Behavior*, 51(1):155–170.

Osborne, M. J. (2003). *An Introduction to Game Theory*. Oxford University Press.

Osborne, M. J. and Rubinstein, A. (1998). Games with procedurally rational players. *American Economic Review Papers and Proceedings*, 88(2):834–847.

Rogers, B. W., Gale, A. M. G., and Morgan, J. P. (2009). Heterogeneous quantal response equilibrium. *Journal of Economic Theory*, 144(4):1440–1467.

Stahl, D. O. and Wilson, P. W. (1995). On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 10(1):218–254.

Weizsäcker, G. (2003). Ignoring the rationality of others: Evidence from experimental normal-form games. *Games and Economic Behavior*, 44(1):145–171.

Weizsäcker, G. (2010). Do we follow others when we should? a simple test of rational expectations. *American Economic Review*, 100(5):2340–60.

Repetition questions

Behavioral game theoryIntroduction

Repetition questions I: benchmark and motivation

1. **What is the basic object studied in game theory?**
   Sketch: A structured strategic situation with players, feasible actions, rules/timing/information, an outcome mapping, and preferences over outcomes.
2. **Why does this course start from standard game theory rather than from behavioral critiques?**
   Sketch: Because equilibrium and related concepts provide the benchmark that organizes strategic reasoning and makes deviations interpretable.
3. **What is the central benchmark used throughout the course?**
   Sketch: Equilibrium reasoning: players choose optimally given beliefs, and in equilibrium those beliefs are correct.
4. **Why can equilibrium predictions work well in familiar environments?**
   Sketch: Experience, routines, feedback, and institutional guidance can substitute for explicit deep reasoning.
5. **Why can equilibrium predictions become fragile in novel environments?**
   Sketch: Because prediction then depends more heavily on players' beliefs, strategic depth, and ability to work through the logic of the game.

Repetition questions II: normal-form and extensive-form games

1. **What does the beauty-contest game illustrate?**
   Sketch: The equilibrium prediction requires iterated reasoning; observed play typically shows finite and heterogeneous strategic depth.
2. **Why are deviations in beauty-contest games informative rather than just "mistakes"?**
   Sketch: Because choices are often clustered at interpretable levels corresponding to limited reasoning steps.
3. **What is the main behavioral issue in coordination games?**
   Sketch: Not just whether players are rational, but how they coordinate on one equilibrium among several possibilities.
4. **What does the centipede game illustrate?**
   Sketch: Backward induction is often descriptively weak in early play; subjects typically continue longer than the benchmark predicts.
5. **What is the basic idea of forward induction?**
   Sketch: Later players may interpret an earlier move in light of why a rational player would have chosen it.

Repetition questions III: Bayesian and repeated games

1. **What makes a game a Bayesian game?**
   Sketch: Some payoff-relevant characteristics are private information, so players must form beliefs about others' types.
2. **What changes when we move from complete information to incomplete information?**
   Sketch: Strategies become type-contingent rules, and equilibrium requires correct beliefs about hidden characteristics and actions.
3. **Why are auctions a natural laboratory for Bayesian games?**
   Sketch: Bidders must choose strategically while reasoning about opponents' unknown values and bids.
4. **What is the winner's curse?**
   Sketch: In common-value settings, winning can reveal that one was too optimistic, so winning itself is bad news.
5. **Why do repeated games differ fundamentally from one-shot games?**
   Sketch: Actions today affect future incentives, so cooperation and punishment can become strategically relevant.

Repetition questions IV: synthesis across the lecture

1. **What is the recurring four-step logic of this course?**
   Sketch: Benchmark theory → evidence → behavioral model → interpretation.
2. **What is the difference between saying that behavior deviates from equilibrium and saying that behavior is systematically structured?**
   Sketch: A deviation from equilibrium need not be random; if it follows stable patterns, it can be modeled and compared across environments.
3. **How do the four classes of games differ in the main difficulty they pose to players?**
   Sketch: Normal form: strategic depth and beliefs; extensive form: look-ahead and interpretation of moves; Bayesian games: beliefs about types; repeated games: maintaining cooperation over time.
4. **What role do the papers shown in the lecture play?**
   Sketch: They provide the original evidence and more detail, while the slides present the core line of argument needed to follow the course.
5. **What is the central message of Lecture 01 in one sentence?**
   Sketch: Standard game theory is the benchmark, and behavioral game theory studies the systematic ways in which actual strategic behavior departs from that benchmark.