Guessing games

Summer 26

Niels Boissonnet, Bielefeld

Discussion

From benchmark theory to the first behavioral puzzle

Guessing games

From last lecture to today

Last lecture established the benchmark language for simultaneous-move games:

- actions, strategies, beliefs, and payoffs,
- dominance and iterated elimination,
- best responses and Nash equilibrium.

Today we ask a first descriptive question:

> *How do people actually reason in one-shot simultaneous-move games when equilibrium reasoning is too demanding as a theory of initial play?*

We begin where standard GT is analytically strongest, but descriptively weakest:

- games that require thinking about what others think
- and what others think that others think.

Why guessing games are the right starting point

The structure of guessing games is strategically simple, but cognitively demanding.

Why guessing games are the right starting point

The structure of guessing games is strategically simple, but cognitively demanding.

- all players move simultaneously,

Why guessing games are the right starting point

The structure of guessing games is strategically simple, but cognitively demanding.

- all players move simultaneously,
- social preferences are not the main issue,

Why guessing games are the right starting point

The structure of guessing games is strategically simple, but cognitively demanding.

- all players move simultaneously,
- social preferences are not the main issue,
- the game has a sharp strategic prediction,

Why guessing games are the right starting point

The structure of guessing games is strategically simple, but cognitively demanding.

- all players move simultaneously,
- social preferences are not the main issue,
- the game has a sharp strategic prediction,
- success depends on anticipating others' reasoning.

Why guessing games are the right starting point

The structure of guessing games is strategically simple, but cognitively demanding.

- all players move simultaneously,
- social preferences are not the main issue,
- the game has a sharp strategic prediction,
- success depends on anticipating others' reasoning.

In particular, the beauty contest is a canonical testbed for strategic depth

- starting with Nagel (1995)
- later large-scale evidence such as Bosch-Domènech et al. (2002).

Why guessing games are the right starting point

The structure of guessing games is strategically simple, but cognitively demanding.

- all players move simultaneously,
- social preferences are not the main issue,
- the game has a sharp strategic prediction,
- success depends on anticipating others' reasoning.

In particular, the beauty contest is a canonical testbed for strategic depth

- starting with Nagel (1995)
- later large-scale evidence such as Bosch-Domènech et al. (2002).

More broadly, this lecture is about initial play (Fudenberg and Liang, 2019):

- before learning and repeated interaction,
- equilibrium cannot be justified as resulting from adaptation or experience

Learning goals and route through today

By the end of this lecture, you should be able to

Learning goals and route through today

By the end of this lecture, you should be able to

1. explain why the p-beauty contest is a useful test of strategic reasoning;

Learning goals and route through today

By the end of this lecture, you should be able to

1. explain why the p-beauty contest is a useful test of strategic reasoning;
2. distinguish rationality, dominance, rationalizability, and Nash equilibrium;

Learning goals and route through today

By the end of this lecture, you should be able to

1. explain why the p-beauty contest is a useful test of strategic reasoning;
2. distinguish rationality, dominance, rationalizability, and Nash equilibrium;
3. explain why non-equilibrium play need not imply irrationality;

Learning goals and route through today

By the end of this lecture, you should be able to

1. explain why the p-beauty contest is a useful test of strategic reasoning;
2. distinguish rationality, dominance, rationalizability, and Nash equilibrium;
3. explain why non-equilibrium play need not imply irrationality;
4. understand why level-k is a useful first descriptive model;

Learning goals and route through today

By the end of this lecture, you should be able to

1. explain why the p-beauty contest is a useful test of strategic reasoning;
2. distinguish rationality, dominance, rationalizability, and Nash equilibrium;
3. explain why non-equilibrium play need not imply irrationality;
4. understand why level-k is a useful first descriptive model;
5. understand why the data also leave room for richer belief-based models.

Learning goals and route through today

By the end of this lecture, you should be able to

1. explain why the p-beauty contest is a useful test of strategic reasoning;
2. distinguish rationality, dominance, rationalizability, and Nash equilibrium;
3. explain why non-equilibrium play need not imply irrationality;
4. understand why level-k is a useful first descriptive model;
5. understand why the data also leave room for richer belief-based models.

Route through today

Learning goals and route through today

By the end of this lecture, you should be able to

1. explain why the p-beauty contest is a useful test of strategic reasoning;
2. distinguish rationality, dominance, rationalizability, and Nash equilibrium;
3. explain why non-equilibrium play need not imply irrationality;
4. understand why level-k is a useful first descriptive model;
5. understand why the data also leave room for richer belief-based models.

Route through today

- introductory example and evidence,

Learning goals and route through today

By the end of this lecture, you should be able to

1. explain why the p-beauty contest is a useful test of strategic reasoning;
2. distinguish rationality, dominance, rationalizability, and Nash equilibrium;
3. explain why non-equilibrium play need not imply irrationality;
4. understand why level-k is a useful first descriptive model;
5. understand why the data also leave room for richer belief-based models.

Route through today

- introductory example and evidence,
- what benchmark theory does and does not say,

Learning goals and route through today

By the end of this lecture, you should be able to

1. explain why the p-beauty contest is a useful test of strategic reasoning;
2. distinguish rationality, dominance, rationalizability, and Nash equilibrium;
3. explain why non-equilibrium play need not imply irrationality;
4. understand why level-k is a useful first descriptive model;
5. understand why the data also leave room for richer belief-based models.

Route through today

- introductory example and evidence,
- what benchmark theory does and does not say,
- behavioral solution concepts,

Learning goals and route through today

By the end of this lecture, you should be able to

1. explain why the p-beauty contest is a useful test of strategic reasoning;
2. distinguish rationality, dominance, rationalizability, and Nash equilibrium;
3. explain why non-equilibrium play need not imply irrationality;
4. understand why level-k is a useful first descriptive model;
5. understand why the data also leave room for richer belief-based models.

Route through today

- introductory example and evidence,
- what benchmark theory does and does not say,
- behavioral solution concepts,
- level-k as the first descriptive model,

Learning goals and route through today

By the end of this lecture, you should be able to

1. explain why the p-beauty contest is a useful test of strategic reasoning;
2. distinguish rationality, dominance, rationalizability, and Nash equilibrium;
3. explain why non-equilibrium play need not imply irrationality;
4. understand why level-k is a useful first descriptive model;
5. understand why the data also leave room for richer belief-based models.

Route through today

- introductory example and evidence,
- what benchmark theory does and does not say,
- behavioral solution concepts,
- level-k as the first descriptive model,
- open question for next week.

Introductory example

Guessing games

The p-beauty contest

Basic game

Each player chooses a number from the interval $[0, 100]$. The winner is the player whose choice is closest to

$$p \times (\text{average of all choices}),$$

where typically $p = \frac{2}{3}$. Ties share the prize.

The p-beauty contest

Basic game

Each player chooses a number from the interval $[0, 100]$. The winner is the player whose choice is closest to

$$p \times (\text{average of all choices}),$$

where typically $p = \frac{2}{3}$. Ties share the prize.

Why this game is interesting

The p-beauty contest

Basic game

Each player chooses a number from the interval $[0, 100]$. The winner is the player whose choice is closest to

$$p \times (\text{average of all choices}),$$

where typically $p = \frac{2}{3}$. Ties share the prize.

Why this game is interesting

- The action space is simple: everybody chooses just one number.

The p-beauty contest

Basic game

Each player chooses a number from the interval $[0, 100]$. The winner is the player whose choice is closest to

$$p \times (\text{average of all choices}),$$

where typically $p = \frac{2}{3}$. Ties share the prize.

Why this game is interesting

- The action space is simple: everybody chooses just one number.
- The difficulty is strategic anticipation rather then heavy algebra.
  - To do well, I should not choose what *I* find reasonable.
  - I should choose what I expect *others* to choose, after taking into account what *they* expect others to choose...

The p-beauty contest

Basic game

Each player chooses a number from the interval $[0, 100]$. The winner is the player whose choice is closest to

$$p \times (\text{average of all choices}),$$

where typically $p = \frac{2}{3}$. Ties share the prize.

Why this game is interesting

- The action space is simple: everybody chooses just one number.
- The difficulty is strategic anticipation rather then heavy algebra.
  - To do well, I should not choose what *I* find reasonable.
  - I should choose what I expect *others* to choose, after taking into account what *they* expect others to choose...

Isolates a central ingredient: higher-order reasoning about others' beliefs/choices.

Why the benchmark prediction is so stark

Why the benchmark prediction is so stark

Step 1: No average can exceed 100, so no target can exceed

$$\frac{2}{3} \cdot 100 = 66.7. \implies \text{a choice above 66.7 cannot win.}$$

Why the benchmark prediction is so stark

Step 1: No average can exceed 100, so no target can exceed

$$\frac{2}{3} \cdot 100 = 66.7. \implies \text{a choice above } 66.7 \text{ cannot win.}$$

Step 2: If all players understand this, then the average cannot exceed 66.7. So the target cannot exceed

$$\frac{2}{3} \cdot 66.7 \approx 44.4. \implies \text{choice above } 44.4 \text{ is no longer plausible either.}$$

Why the benchmark prediction is so stark

Step 1: No average can exceed 100, so no target can exceed

$$\frac{2}{3} \cdot 100 = 66.7. \implies \text{a choice above } 66.7 \text{ cannot win.}$$

Step 2: If all players understand this, then the average cannot exceed 66.7. So the target cannot exceed

$$\frac{2}{3} \cdot 66.7 \approx 44.4. \implies \text{choice above } 44.4 \text{ is no longer plausible either.}$$

Repeating this reasoning: $66.7,\ 44.4,\ 29.6,\ \cdots \to 0$.

Rationality + common knowledge of rationality

the unique Nash equilibrium is that all players choose 0.

What this game tests

The beauty contest primarily tests three things:

What this game tests

The beauty contest primarily tests three things:

1. strategic anticipation: do subjects reason about others rather than only about themselves?

What this game tests

The beauty contest primarily tests three things:

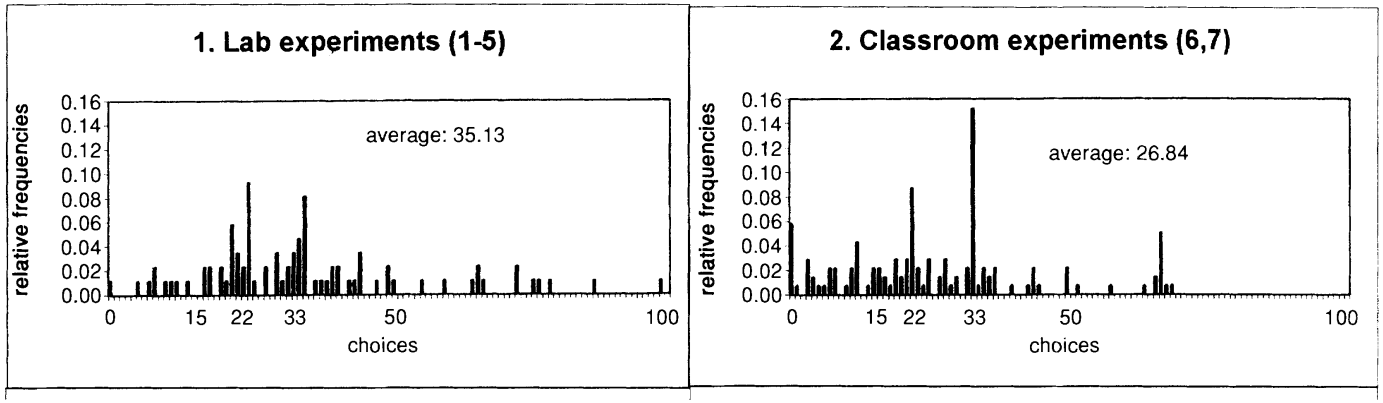
1. strategic anticipation: do subjects reason about others rather than only about themselves?
2. depth of reasoning: how many steps of "I think that you think that they think ..." are used?

What this game tests

The beauty contest primarily tests three things:

1. strategic anticipation: do subjects reason about others rather than only about themselves?
2. depth of reasoning: how many steps of "I think that you think that they think ..." are used?
3. initial play: what happens before feedback repeated interaction, learning, or adaptation can discipline behavior?

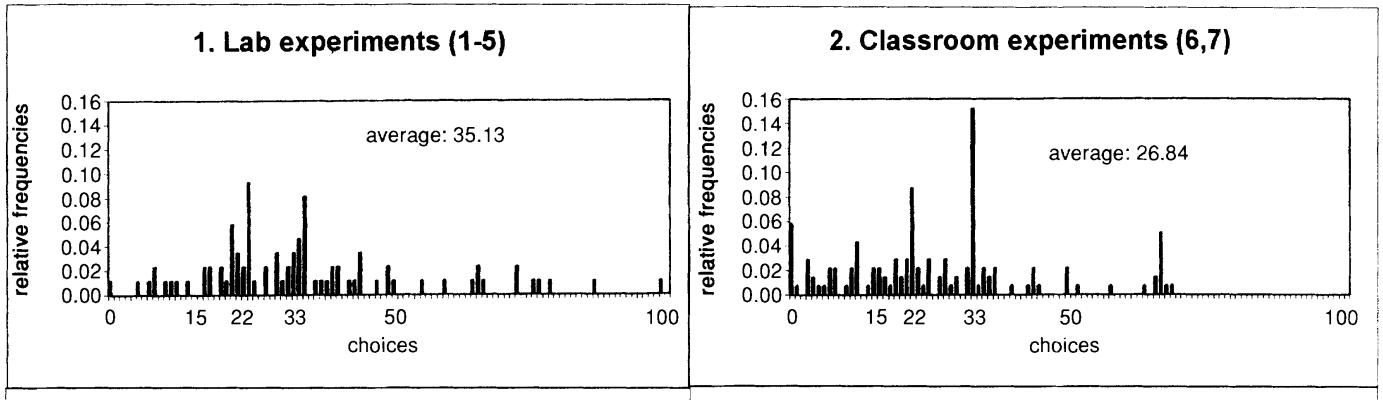
Experimental data: a first look



Source: Bosch-Domènech et al. (2002)

The broad message is immediate:

Experimental data: a first look

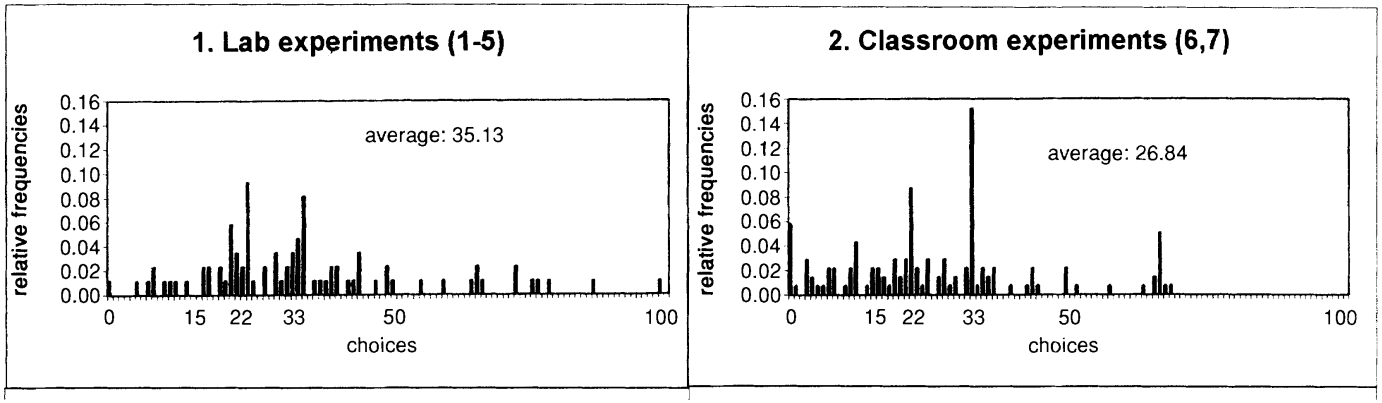


Source: Bosch-Domènech et al. (2002)

The broad message is immediate:

- the large majority of choices lie far above the Nash prediction of 0;

Experimental data: a first look

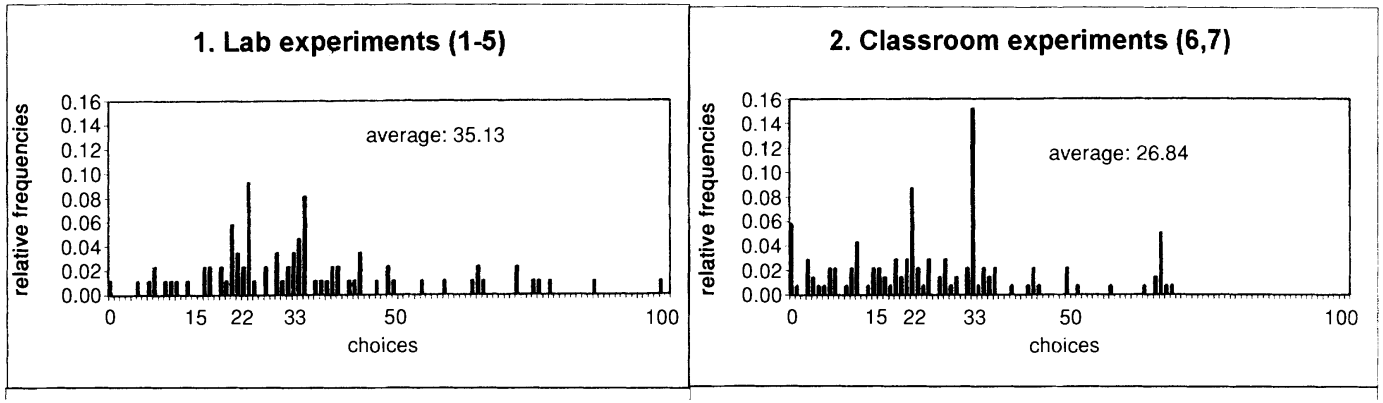


Source: Bosch-Domènech et al. (2002)

The broad message is immediate:

- the large majority of choices lie far above the Nash prediction of 0;
- observed play is nevertheless highly structured, not random;

Experimental data: a first look



Source: Bosch-Domènech et al. (2002)

The broad message is immediate:

- the large majority of choices lie far above the Nash prediction of 0;
- observed play is nevertheless highly structured, not random;
- even participants with strong formal training do not mostly choose the equilibrium.

Why the data are puzzling

The puzzle is not simply that people do not play Nash equilibrium.

The more interesting point is that choices display a clear pattern:

Why the data are puzzling

The puzzle is not simply that people do not play Nash equilibrium.

The more interesting point is that choices display a clear pattern:

- many choices are well below 50, so subjects are clearly thinking strategically;

Why the data are puzzling

The puzzle is not simply that people do not play Nash equilibrium.

The more interesting point is that choices display a clear pattern:

- many choices are well below 50, so subjects are clearly thinking strategically;
- but most choices are still far above 0, so full equilibrium reasoning is rare;

Why the data are puzzling

The puzzle is not simply that people do not play Nash equilibrium.

The more interesting point is that choices display a clear pattern:

- many choices are well below 50, so subjects are clearly thinking strategically;
- but most choices are still far above 0, so full equilibrium reasoning is rare;
- the distribution concentrates in a few regions rather than a smooth mass;

Why the data are puzzling

The puzzle is not simply that people do not play Nash equilibrium.

The more interesting point is that choices display a clear pattern:

- many choices are well below 50, so subjects are clearly thinking strategically;
- but most choices are still far above 0, so full equilibrium reasoning is rare;
- the distribution concentrates in a few regions rather than a smooth mass;
- the data therefore suggest bounded but systematic strategic reasoning.

Why the data are puzzling

The puzzle is not simply that people do not play Nash equilibrium.

The more interesting point is that choices display a clear pattern:

- many choices are well below 50, so subjects are clearly thinking strategically;
- but most choices are still far above 0, so full equilibrium reasoning is rare;
- the distribution concentrates in a few regions rather than a smooth mass;
- the data therefore suggest bounded but systematic strategic reasoning.

So the central empirical question is not

"Why are subjects irrational?"

Why the data are puzzling

The puzzle is not simply that people do not play Nash equilibrium.

The more interesting point is that choices display a clear pattern:

- many choices are well below 50, so subjects are clearly thinking strategically;
- but most choices are still far above 0, so full equilibrium reasoning is rare;
- the distribution concentrates in a few regions rather than a smooth mass;
- the data therefore suggest bounded but systematic strategic reasoning.

So the central empirical question is not

"Why are subjects irrational?"

but rather

"What kind of strategic reasoning process can generate this distribution of choices?"

A first intuitive reading of the pattern

Natural interpretation: players differ in how many reasoning steps they use.

Starting from a salient anchor such as 50:

A first intuitive reading of the pattern

Natural interpretation: players differ in how many reasoning steps they use.

Starting from a salient anchor such as 50:

- one step gives roughly $\frac{2}{3} \cdot 50 = 33.3$;

A first intuitive reading of the pattern

Natural interpretation: players differ in how many reasoning steps they use.

Starting from a salient anchor such as 50:

- one step gives roughly $\frac{2}{3} \cdot 50 = 33.3$;
- two steps give roughly $\frac{2}{3} \cdot 33.3 = 22.2$;

A first intuitive reading of the pattern

Natural interpretation: players differ in how many reasoning steps they use.

Starting from a salient anchor such as 50:

- one step gives roughly $\frac{2}{3} \cdot 50 = 33.3$;
- two steps give roughly $\frac{2}{3} \cdot 33.3 = 22.2$;
- three steps give roughly $\frac{2}{3} \cdot 22.2 = 14.8$.

This already makes the data more intelligible:

$\Longrightarrow$ choices around 33, 22, or 15 are not random;

$\Longrightarrow$ they look like the outcome of finite iterative reasoning.

A first intuitive reading of the pattern

Natural interpretation: players differ in how many reasoning steps they use.

Starting from a salient anchor such as 50:

- one step gives roughly $\frac{2}{3} \cdot 50 = 33.3$;
- two steps give roughly $\frac{2}{3} \cdot 33.3 = 22.2$;
- three steps give roughly $\frac{2}{3} \cdot 22.2 = 14.8$.

This already makes the data more intelligible:

$\Longrightarrow$ choices around 33, 22, or 15 are not random;

$\Longrightarrow$ they look like the outcome of finite iterative reasoning.

Important caution!!!

At this stage, this is just an intuition, not yet a formal model.

What benchmark theory does and does not say

Guessing games

Non-equilibrium play does not imply irrationality

Actions and strategies

Player i chooses an action $a_i \in A_i$. A (possibly mixed) strategy is a probability distribution $\sigma_i \in \Delta(A_i)$ over actions. Let σ_{-i} denote the opponents' strategy profile.

Non-equilibrium play does not imply irrationality

Actions and strategies

Player i chooses an action $a_i \in A_i$. A (possibly mixed) strategy is a probability distribution $\sigma_i \in \Delta(A_i)$ over actions. Let σ_{-i} denote the opponents' strategy profile.

Observed play in the beauty contest differs strongly from Nash equilibrium. But this does not yet show that subjects are irrational.

Non-equilibrium play does not imply irrationality

Actions and strategies

Player i chooses an action $a_i \in A_i$. A (possibly mixed) strategy is a probability distribution $\sigma_i \in \Delta(A_i)$ over actions. Let σ_{-i} denote the opponents' strategy profile.

Observed play in the beauty contest differs strongly from Nash equilibrium. But this does not yet show that subjects are irrational.

Rationality

A player is rational if she chooses a best response to her subjective belief about others' play.

Non-equilibrium play does not imply irrationality

Actions and strategies

Player i chooses an action $a_i \in A_i$. A (possibly mixed) strategy is a probability distribution $\sigma_i \in \Delta(A_i)$ over actions. Let σ_{-i} denote the opponents' strategy profile.

Observed play in the beauty contest differs strongly from Nash equilibrium. But this does not yet show that subjects are irrational.

Rationality

A player is rational if she chooses a best response to her subjective belief about others' play.

So rationality is weaker than equilibrium:

- rationality requires only best responses to believed play;
- equilibrium requires common knowledge of rationality+correct beliefs;

Strict dominance dominance in the beauty contest

Question: can the observed choices violate very basic logic of dominance?

- In other words, can we use the dominance algorithm to interpret subjects choices?

Strict dominance

Strategy σ_i strictly dominates σ_i', if $\pi_i(\sigma_i, \sigma_{-i}) > \pi_i(\sigma_i', \sigma_{-i})$ for all σ_{-i}.

Strict dominance is too weak. Why?

- Several cases in which strategies yield the same payoff of 0
- playing 100 or 10 can yield 0 payoffs if all other players chose 0

Weak dominance dominance in the beauty contest

Question: can the observed choices violate very basic logic of dominance?

- In other words, can we use the dominance algorith to interpret subjects choices?

Weak dominance

If all choices lie in $[0, 100]$, then no target can exceed 66.7. So choices above 66.7 are weakly dominated in large groups. Repeating the logic yields thresholds such as 44.4, 29.6, 19.8, ...

So, applied literally the algorithm behind weak dominance provides a poor descriptions of observed choices.

Rationalizability

Which actions can be justified as best responses to some coherent view of others?

Rationalizability

Which actions can be justified as best responses to some coherent view of others?

Rationalizability

A strategy is rationalizable if it is a best response to opponents' strategies that are themselves rationalizable.

Rationalizability

Which actions can be justified as best responses to some coherent view of others?

Rationalizability

A strategy is rationalizable if it is a best response to opponents' strategies that are themselves rationalizable.

Operationally, one can think of rationalizability as an iterative procedure:

1. eliminate strategies that are never best responses;
2. update the set of plausible opponent strategies;
3. eliminate strategies that are no longer best responses;
4. continue until no more eliminations are possible.

Rationalizability is therefore a benchmark for best-response consistency under strategic uncertainty.

What rationalizability implies in beauty contests?

This depends on whether the action set is treated as continuous or discrete.

What rationalizability implies in beauty contests?

This depends on whether the action set is treated as continuous or discrete.

Continuous action set: $A_i = [0, 100]$

- every action below 100 is rationalizable;
- for example, 99.9 is a best response to the belief that others choose 99.99, which is a best response to 99.999, and so on.

What rationalizability implies in beauty contests?

This depends on whether the action set is treated as continuous or discrete.

Continuous action set: $A_i = [0, 100]$

- every action below 100 is rationalizable;
- for example, 99.9 is a best response to the belief that others choose 99.99, which is a best response to 99.999, and so on.

Discrete action set: $A_i = \{0, 1, \ldots, 100\}$

- the surviving rationalizable actions collapse to very low numbers;
- but the elimination process is extremely slow: eliminating action $a_i > 1$ takes roughly $101 - a_i$ rounds.

What rationalizability implies in beauty contests?

This depends on whether the action set is treated as continuous or discrete.

Continuous action set: $A_i = [0, 100]$

- every action below 100 is rationalizable;
- for example, 99.9 is a best response to the belief that others choose 99.99, which is a best response to 99.999, and so on.

Discrete action set: $A_i = \{0, 1, \ldots, 100\}$

- the surviving rationalizable actions collapse to very low numbers;
- but the elimination process is extremely slow: eliminating action $a_i > 1$ takes roughly $101 - a_i$ rounds.

So the algorithm behind rationalizability provides a poor descriptions of observed choices.

Why rationalizability does not describe the data well

Suppose the grid is fine, for instance with steps of 0.01 as in many experiments.

Why rationalizability does not describe the data well

Suppose the grid is fine, for instance with steps of 0.01 as in many experiments.

Then a choice like 50.00 is eliminated only after thousands of iterations.

$\Longrightarrow$ Taken literally, rationalizability suggests that subjects perform a huge number of elimination rounds and then suddenly stop.

Why rationalizability does not describe the data well

Suppose the grid is fine, for instance with steps of 0.01 as in many experiments.

Then a choice like 50.00 is eliminated only after thousands of iterations.

$\Longrightarrow$ Taken literally, rationalizability suggests that subjects perform a huge number of elimination rounds and then suddenly stop.

That is not a persuasive behavioral reading.

Why rationalizability does not describe the data well

Suppose the grid is fine, for instance with steps of 0.01 as in many experiments.

Then a choice like 50.00 is eliminated only after thousands of iterations.

$\Longrightarrow$ Taken literally, rationalizability suggests that subjects perform a huge number of elimination rounds and then suddenly stop.

That is not a persuasive behavioral reading.

The empirical pattern looks different:

- many subjects quickly understand that very high choices are implausible;
- many also do *not* continue all the way toward equilibrium;
- the distribution shows salient masses at intermediate values rather than thresholds generated by long elimination chains.

Summing up: rationality, rationalizability, Nash

Concept	What it requires	What it does not require
Rationality	best response to subjective beliefs	correct beliefs, mutual consistency, equilibrium
Rationalizability	best response to beliefs compatible with others' rationality	common prediction of one realized profile
Nash equilibrium	best response to correct beliefs about others' actual strategies	efficiency, uniqueness, easy discoverability

Interim takeaway

What benchmark theory does give us

- a sharp equilibrium prediction,
- clear restrictions from dominance and weak dominance,
- a precise notion of best-response consistency via rationalizability.

Interim takeaway

What benchmark theory does give us

- a sharp equilibrium prediction,
- clear restrictions from dominance and weak dominance,
- a precise notion of best-response consistency via rationalizability.

What benchmark theory does not give us

- a good description of the observed distribution of initial play,
- a plausible account of finite strategic depth,
- a model of how subjects form beliefs about others.

Interim takeaway

What benchmark theory does give us

- a sharp equilibrium prediction,
- clear restrictions from dominance and weak dominance,
- a precise notion of best-response consistency via rationalizability.

What benchmark theory does not give us

- a good description of the observed distribution of initial play,
- a plausible account of finite strategic depth,
- a model of how subjects form beliefs about others.

So the next step is natural:

Keep the idea that choices respond to beliefs, but model those beliefs explicitly.

Behavioral solution concepts

Guessing games

What a descriptive model must specify

A satisfying behavioral model of strategic choice must specify three objects:

What a descriptive model must specify

A satisfying behavioral model of strategic choice must specify three objects:

1. beliefs: what does player i expect the other players to do?

What a descriptive model must specify

A satisfying behavioral model of strategic choice must specify three objects:

1. beliefs: what does player i expect the other players to do?
2. responses: how does player i choose given those beliefs?

What a descriptive model must specify

A satisfying behavioral model of strategic choice must specify three objects:

1. beliefs: what does player i expect the other players to do?
2. responses: how does player i choose given those beliefs?
3. heterogeneity: why do different players choose differently in the same game?

In the beauty contest, two benchmark assumptions are especially natural to relax:

What a descriptive model must specify

A satisfying behavioral model of strategic choice must specify three objects:

1. beliefs: what does player i expect the other players to do?
2. responses: how does player i choose given those beliefs?
3. heterogeneity: why do different players choose differently in the same game?

In the beauty contest, two benchmark assumptions are especially natural to relax:

- exact best response: players may respond noisily to beliefs;

What a descriptive model must specify

A satisfying behavioral model of strategic choice must specify three objects:

1. beliefs: what does player i expect the other players to do?
2. responses: how does player i choose given those beliefs?
3. heterogeneity: why do different players choose differently in the same game?

In the beauty contest, two benchmark assumptions are especially natural to relax:

- exact best response: players may respond noisily to beliefs;
- beliefs consistency: players may hold simplified or incorrect views of others.

What a descriptive model must specify

A satisfying behavioral model of strategic choice must specify three objects:

1. beliefs: what does player i expect the other players to do?
2. responses: how does player i choose given those beliefs?
3. heterogeneity: why do different players choose differently in the same game?

In the beauty contest, two benchmark assumptions are especially natural to relax:

- exact best response: players may respond noisily to beliefs;
- beliefs consistency: players may hold simplified or incorrect views of others.

$\Longrightarrow$ subjects are heterogeneous in both aspects

Beliefs: what exactly is being modelled?

Let A_j denote player j's action set.

Beliefs: what exactly is being modelled?

Let A_j denote player j's action set.

Beliefs about one opponent

Player i's belief about player j is a probability distribution $\mu_{i,j} \in \Delta(A_j)$.

Expected payoffs are evaluated under these beliefs:

$$EU_i(a_i \mid \mu_{i,-i}) = \sum_{a_{-i} \in A_{-i}} \mu_{i,-i}(a_{-i})\, u_i(a_i, a_{-i}).$$

Beliefs: what exactly is being modelled?

Let A_j denote player j's action set.

Beliefs about one opponent

Player i's belief about player j is a probability distribution $\mu_{i,j} \in \Delta(A_j)$.

Expected payoffs are evaluated under these beliefs:

$$EU_i(a_i \mid \mu_{i,-i}) = \sum_{a_{-i} \in A_{-i}} \mu_{i,-i}(a_{-i})\, u_i(a_i, a_{-i}).$$

Best-responses given beliefs are

$$BR_i(\mu_{i,-i}) = \arg\max_{a_i \in A_i} EU_i(a_i \mid \mu_{i,-i})$$

Hence, once beliefs are specified, strategic choice becomes a response problem.

Responses: best response and logit response

The benchmark response concept is the best response:

$$BR_i(\mu_{i,-i}) = \arg\max_{a_i \in A_i} EU_i(a_i \mid \mu_{i,-i}).$$

A common behavioral relaxation is the logit response:

$$\sigma_i(a_i) = \frac{\exp(\lambda\, EU_i(a_i \mid \mu_{i,-i}))}{\sum_{a_i' \in A_i} \exp(\lambda\, EU_i(a_i' \mid \mu_{i,-i}))}.$$

Responses: best response and logit response

The benchmark response concept is the best response:

$$BR_i(\mu_{i,-i}) = \arg\max_{a_i \in A_i} EU_i(a_i \mid \mu_{i,-i}).$$

A common behavioral relaxation is the logit response:

$$\sigma_i(a_i) = \frac{\exp(\lambda\, EU_i(a_i \mid \mu_{i,-i}))}{\sum_{a_i' \in A_i} \exp(\lambda\, EU_i(a_i' \mid \mu_{i,-i}))}.$$

- $\lambda \geq 0$ measures payoff responsiveness;

Responses: best response and logit response

The benchmark response concept is the best response:

$$BR_i(\mu_{i,-i}) = \arg\max_{a_i \in A_i} EU_i(a_i \mid \mu_{i,-i}).$$

A common behavioral relaxation is the logit response:

$$\sigma_i(a_i) = \frac{\exp(\lambda\, EU_i(a_i \mid \mu_{i,-i}))}{\sum_{a_i' \in A_i} \exp(\lambda\, EU_i(a_i' \mid \mu_{i,-i}))}.$$

- $\lambda \geq 0$ measures payoff responsiveness;
- high-payoff actions are chosen with higher probability;

Responses: best response and logit response

The benchmark response concept is the best response:

$$BR_i(\mu_{i,-i}) = \arg\max_{a_i \in A_i} EU_i(a_i \mid \mu_{i,-i}).$$

A common behavioral relaxation is the logit response:

$$\sigma_i(a_i) = \frac{\exp(\lambda\, EU_i(a_i \mid \mu_{i,-i}))}{\sum_{a_i' \in A_i} \exp(\lambda\, EU_i(a_i' \mid \mu_{i,-i}))}.$$

- $\lambda \geq 0$ measures payoff responsiveness;
- high-payoff actions are chosen with higher probability;
- as $\lambda \to \infty$, choices approach best responses;

Responses: best response and logit response

The benchmark response concept is the best response:

$$BR_i(\mu_{i,-i}) = \arg\max_{a_i \in A_i} EU_i(a_i \mid \mu_{i,-i}).$$

A common behavioral relaxation is the logit response:

$$\sigma_i(a_i) = \frac{\exp(\lambda\, EU_i(a_i \mid \mu_{i,-i}))}{\sum_{a_i' \in A_i} \exp(\lambda\, EU_i(a_i' \mid \mu_{i,-i}))}.$$

- $\lambda \geq 0$ measures payoff responsiveness;
- high-payoff actions are chosen with higher probability;
- as $\lambda \to \infty$, choices approach best responses;
- if $\lambda = 0$, all actions are chosen equally often.

So logit response means structured, payoff-sensitive noise, not random play.

Why logit response is a useful baseline

Logit response is attractive because it relaxes exact optimization without abandoning strategic structure.

Why logit response is a useful baseline

Logit response is attractive because it relaxes exact optimization without abandoning strategic structure.

- It preserves the idea that better actions should be chosen more often.

Why logit response is a useful baseline

Logit response is attractive because it relaxes exact optimization without abandoning strategic structure.

- It preserves the idea that better actions should be chosen more often.
- It avoids knife-edge predictions when payoffs are very similar.

Why logit response is a useful baseline

Logit response is attractive because it relaxes exact optimization without abandoning strategic structure.

- It preserves the idea that better actions should be chosen more often.
- It avoids knife-edge predictions when payoffs are very similar.
- It gives a simple one-parameter way to model noisy choice.

Why logit response is a useful baseline

Logit response is attractive because it relaxes exact optimization without abandoning strategic structure.

- It preserves the idea that better actions should be chosen more often.
- It avoids knife-edge predictions when payoffs are very similar.
- It gives a simple one-parameter way to model noisy choice.
- It is analytically tractable and widely used in experimental game theory.

This logic underlies quantal response equilibrium (QRE):

$\Longrightarrow$ players choose logit responses and beliefs are correct in the sense that beliefs match the actual distribution of play (McKelvey and Palfrey, 1995).

So QRE relaxes exact best response, but keeps mutual consistency of beliefs.

QRE as a fixed point

QRE can be read as a fixed-point condition.

Idea

Each player chooses a logit response to beliefs about others, and those beliefs must coincide with the actual choice probabilities in the population.

So a QRE is a profile $(\sigma_1, \ldots, \sigma_n)$ such that, for every player i,

$$\sigma_i(a_i) = \frac{\exp(\lambda\, EU_i(a_i \mid \sigma_{-i}))}{\sum_{a_i' \in A_i} \exp(\lambda\, EU_i(a_i' \mid \sigma_{-i}))}.$$

QRE as a fixed point

QRE can be read as a fixed-point condition.

Idea

Each player chooses a logit response to beliefs about others, and those beliefs must coincide with the actual choice probabilities in the population.

So a QRE is a profile $(\sigma_1, \ldots, \sigma_n)$ such that, for every player i,

$$\sigma_i(a_i) = \frac{\exp(\lambda\, EU_i(a_i \mid \sigma_{-i}))}{\sum_{a_i' \in A_i} \exp(\lambda\, EU_i(a_i' \mid \sigma_{-i}))}.$$

Interpretation

- players may make errors,

QRE as a fixed point

QRE can be read as a fixed-point condition.

Idea

Each player chooses a logit response to beliefs about others, and those beliefs must coincide with the actual choice probabilities in the population.

So a QRE is a profile $(\sigma_1, \ldots, \sigma_n)$ such that, for every player i,

$$\sigma_i(a_i) = \frac{\exp(\lambda\, EU_i(a_i \mid \sigma_{-i}))}{\sum_{a_i' \in A_i} \exp(\lambda\, EU_i(a_i' \mid \sigma_{-i}))}.$$

Interpretation

- players may make errors,
- but errors are systematic rather than arbitrary,

QRE as a fixed point

QRE can be read as a fixed-point condition.

Idea

Each player chooses a logit response to beliefs about others, and those beliefs must coincide with the actual choice probabilities in the population.

So a QRE is a profile $(\sigma_1, \ldots, \sigma_n)$ such that, for every player i,

$$\sigma_i(a_i) = \frac{\exp(\lambda\, EU_i(a_i \mid \sigma_{-i}))}{\sum_{a_i' \in A_i} \exp(\lambda\, EU_i(a_i' \mid \sigma_{-i}))}.$$

Interpretation

- players may make errors,
- but errors are systematic rather than arbitrary,
- and beliefs are still mutually consistent.

Two different ways to relax benchmark theory

In the beauty contest, the two main behavioral directions are conceptually distinct.

Approach	**What is relaxed?**	**Leading idea**
QRE	Exact best response	Players respond noisily, but to correct beliefs
Level-k/CH	Mutual consistency of beliefs	Players best-respond to simplified models of others

Both approaches preserve strategic reasoning, but in different ways.

A simple map of the models we are about to use

The lecture now turns to models that differ along two dimensions:

A simple map of the models we are about to use

The lecture now turns to models that differ along two dimensions:

- how players form beliefs,

A simple map of the models we are about to use

The lecture now turns to models that differ along two dimensions:

- how players form beliefs,
- how sharply they respond to those beliefs.

A simple map of the models we are about to use

The lecture now turns to models that differ along two dimensions:

- how players form beliefs,
- how sharply they respond to those beliefs.

- Nash equilibrium: correct beliefs and exact best responses.

A simple map of the models we are about to use

The lecture now turns to models that differ along two dimensions:

- how players form beliefs,
- how sharply they respond to those beliefs.

- Nash equilibrium: correct beliefs and exact best responses.
- QRE: correct beliefs and noisy best responses.

A simple map of the models we are about to use

The lecture now turns to models that differ along two dimensions:

- how players form beliefs,
- how sharply they respond to those beliefs.

- Nash equilibrium: correct beliefs and exact best responses.
- QRE: correct beliefs and noisy best responses.
- Level-k: simplified, non-equilibrium beliefs and typically exact best responses.

A simple map of the models we are about to use

The lecture now turns to models that differ along two dimensions:

- how players form beliefs,
- how sharply they respond to those beliefs.

- Nash equilibrium: correct beliefs and exact best responses.
- QRE: correct beliefs and noisy best responses.
- Level-k: simplified, non-equilibrium beliefs and typically exact best responses.
- Cognitive hierarchy: heterogeneous reasoning depths with beliefs about lower-level types (Camerer et al., 2004).

Where we go next: in the remainder of this lecture, we start with level-k as the first descriptive model; next lecture returns to richer full-support belief models.

Level-k as the first descriptive model

Guessing games

Why level-k starts from a more plausible anchor

Benchmark theory asks a lot from players in a one-shot game:

Why level-k starts from a more plausible anchor

Benchmark theory asks a lot from players in a one-shot game:

- they are rational,

Why level-k starts from a more plausible anchor

Benchmark theory asks a lot from players in a one-shot game:

- they are rational,
- they know that others are rational,

Why level-k starts from a more plausible anchor

Benchmark theory asks a lot from players in a one-shot game:

- they are rational,
- they know that others are rational,
- they know that others know this,

Why level-k starts from a more plausible anchor

Benchmark theory asks a lot from players in a one-shot game:

- they are rational,
- they know that others are rational,
- they know that others know this,
- and so on.

In the beauty contest, that logic leads all the way to the equilibrium choice 0.

Level-k takes a different starting point:

Why level-k starts from a more plausible anchor

Benchmark theory asks a lot from players in a one-shot game:

- they are rational,
- they know that others are rational,
- they know that others know this,
- and so on.

In the beauty contest, that logic leads all the way to the equilibrium choice 0.

Level-k takes a different starting point:

- players begin from a simple non-equilibrium anchor,

Why level-k starts from a more plausible anchor

Benchmark theory asks a lot from players in a one-shot game:

- they are rational,
- they know that others are rational,
- they know that others know this,
- and so on.

In the beauty contest, that logic leads all the way to the equilibrium choice 0.

Level-k takes a different starting point:

- players begin from a simple non-equilibrium anchor,
- they perform only finitely many reasoning steps,

Why level-k starts from a more plausible anchor

Benchmark theory asks a lot from players in a one-shot game:

- they are rational,
- they know that others are rational,
- they know that others know this,
- and so on.

In the beauty contest, that logic leads all the way to the equilibrium choice 0.

Level-k takes a different starting point:

- players begin from a simple non-equilibrium anchor,
- they perform only finitely many reasoning steps,
- different players may stop at different depths.

So level-k is designed for initial play in one-shot games: it replaces common-knowledge reasoning by a hierarchy of finite strategic depths.

How level-k reasoning is constructed

The model is built recursively.

Level-0

Level-0 players do not reason strategically. They use a simple anchor or default rule. In beauty contests, the natural anchor is often a choice around 50.

How level-k reasoning is constructed

The model is built recursively.

Level-0

Level-0 players do not reason strategically. They use a simple anchor or default rule. In beauty contests, the natural anchor is often a choice around 50.

Higher levels

How level-k reasoning is constructed

The model is built recursively.

Level-0

Level-0 players do not reason strategically. They use a simple anchor or default rule. In beauty contests, the natural anchor is often a choice around 50.

Higher levels

- Level-1 chooses a best response to level-0.

How level-k reasoning is constructed

The model is built recursively.

Level-0

Level-0 players do not reason strategically. They use a simple anchor or default rule. In beauty contests, the natural anchor is often a choice around 50.

Higher levels

- Level-1 chooses a best response to level-0.
- Level-2 chooses a best response to level-1.

How level-k reasoning is constructed

The model is built recursively.

Level-0

Level-0 players do not reason strategically. They use a simple anchor or default rule. In beauty contests, the natural anchor is often a choice around 50.

Higher levels

- Level-1 chooses a best response to level-0.
- Level-2 chooses a best response to level-1.
- More generally, level-k chooses a best response to level-$(k-1)$.

Observed play is therefore interpreted as a mixture of different reasoning depths.

Level-k predictions in the beauty contest

Take the standard game with target factor

$$p = \frac{2}{3}.$$

If level-0 is anchored at 50, then the implied predictions are:

Type	Reasoning	Predicted choice
Level-0	salient anchor	50.0
Level-1	$\frac{2}{3} \times 50$	33.3
Level-2	$\frac{2}{3} \times 33.3$	22.2
Level-3	$\frac{2}{3} \times 22.2$	14.8
Level-4	$\frac{2}{3} \times 14.8$	9.9

This makes the common empirical values 33, 22, and 15 immediately intelligible: they look like finite iterative reasoning from a salient anchor.

From individual types to a population model

Level-k is most useful when interpreted as a population model.

Different subjects may stop after different numbers of reasoning steps:

From individual types to a population model

Level-k is most useful when interpreted as a population model.

Different subjects may stop after different numbers of reasoning steps:

- some remain close to the anchor,

From individual types to a population model

Level-k is most useful when interpreted as a population model.

Different subjects may stop after different numbers of reasoning steps:

- some remain close to the anchor,
- some perform one or two iterations,

From individual types to a population model

Level-k is most useful when interpreted as a population model.

Different subjects may stop after different numbers of reasoning steps:

- some remain close to the anchor,
- some perform one or two iterations,
- only a small minority reason many steps ahead.

Then the aggregate distribution of choices is a mixture of level-0, level-1, level-2, ...

This helps explain two central empirical features:

From individual types to a population model

Level-k is most useful when interpreted as a population model.

Different subjects may stop after different numbers of reasoning steps:

- some remain close to the anchor,
- some perform one or two iterations,
- only a small minority reason many steps ahead.

Then the aggregate distribution of choices is a mixture of level-0, level-1, level-2, ...

This helps explain two central empirical features:

- why choices are far above equilibrium on average,

From individual types to a population model

Level-k is most useful when interpreted as a population model.

Different subjects may stop after different numbers of reasoning steps:

- some remain close to the anchor,
- some perform one or two iterations,
- only a small minority reason many steps ahead.

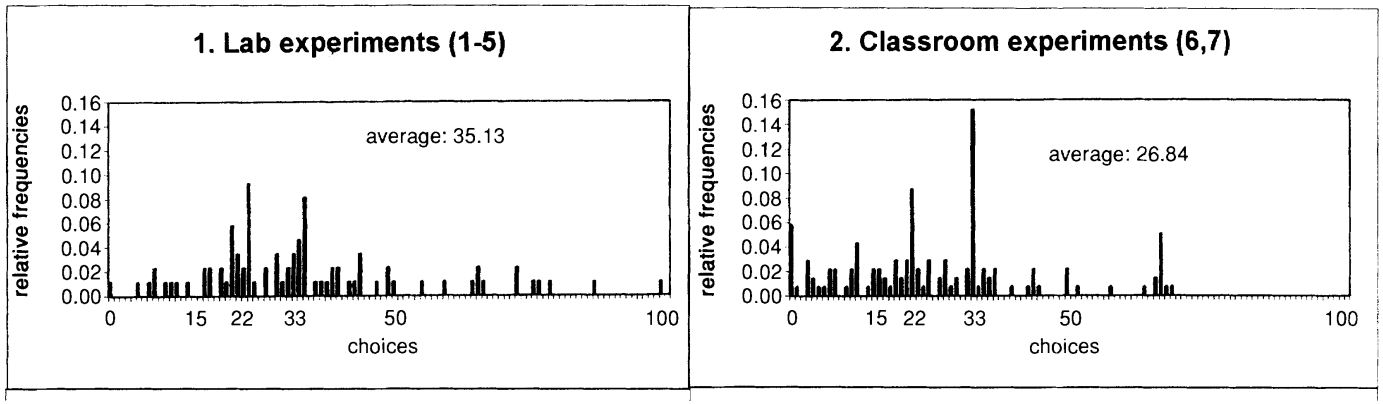
Then the aggregate distribution of choices is a mixture of level-0, level-1, level-2, ...

This helps explain two central empirical features:

- why choices are far above equilibrium on average,
- why they often cluster around a few salient values.

Cognitive hierarchy extends this idea: a level-k player best-responds to a distribution of lower-level types (Camerer et al., 2004).

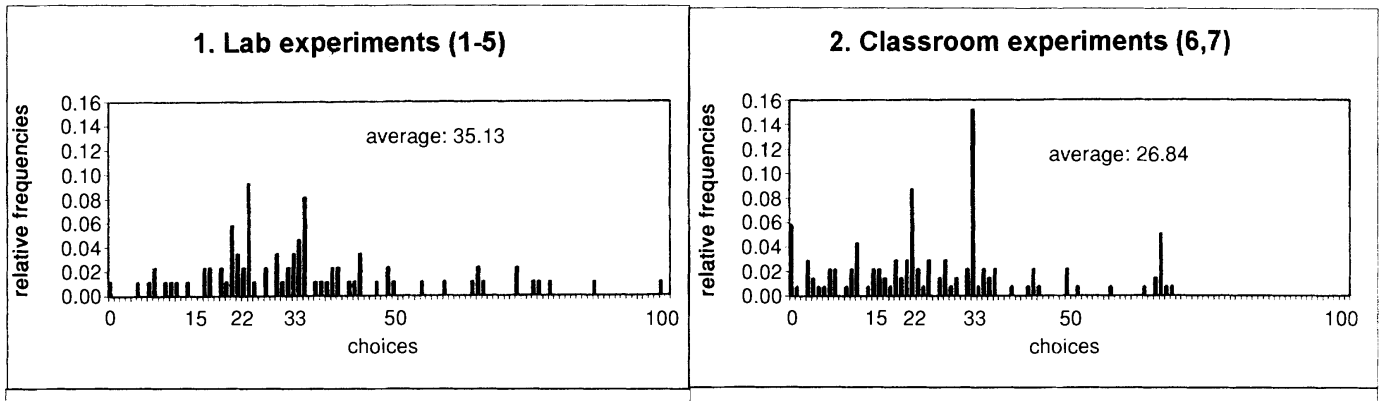
Where level-k fits the beauty-contest data well



Source: Bosch-Domènech et al. (2002)

Level-k improves sharply on both Nash and rationalizability as a descriptive account:

Where level-k fits the beauty-contest data well

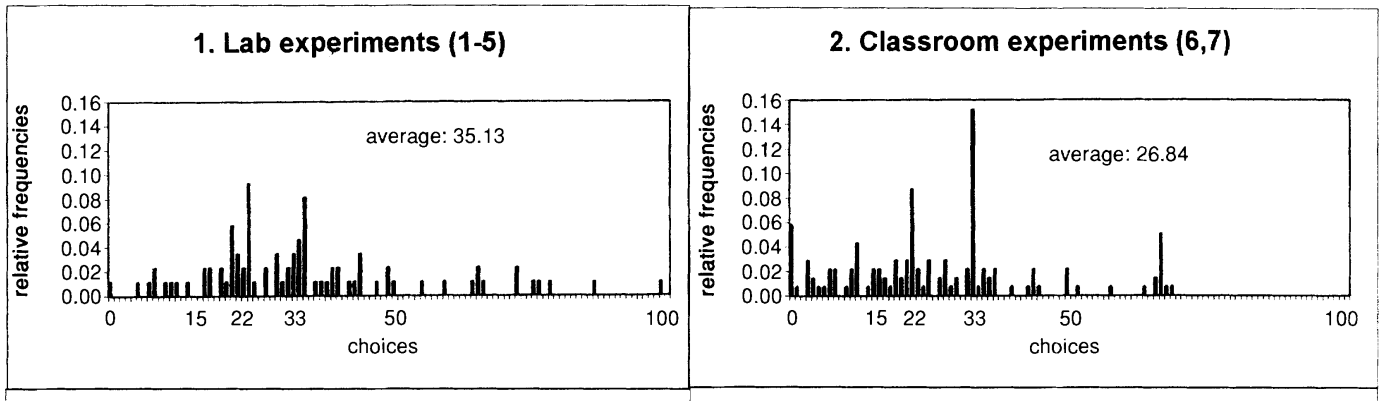


Source: Bosch-Domènech et al. (2002)

Level-k improves sharply on both Nash and rationalizability as a descriptive account:

- it explains why many subjects choose well below 50,

Where level-k fits the beauty-contest data well

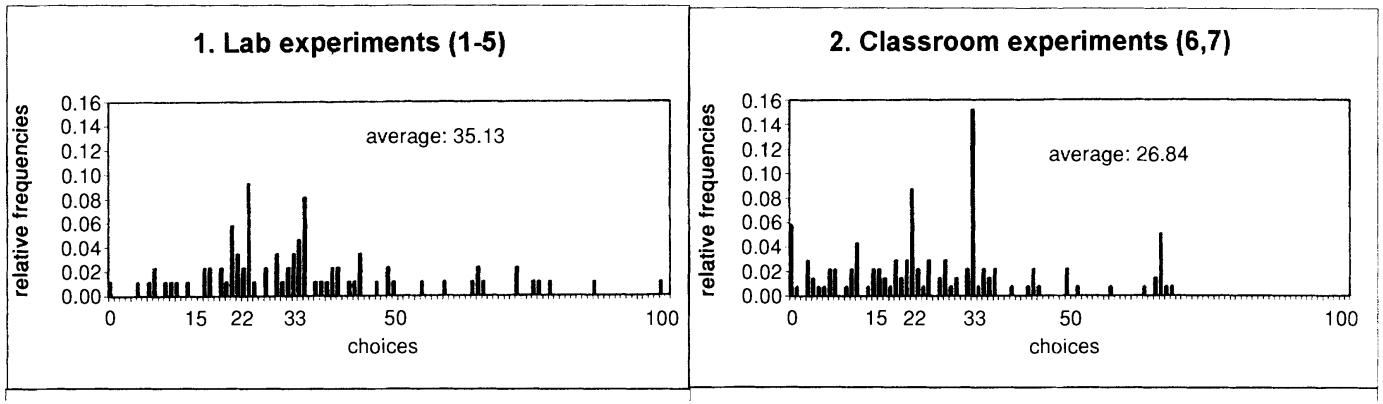


Source: Bosch-Domènech et al. (2002)

Level-k improves sharply on both Nash and rationalizability as a descriptive account:

- it explains why many subjects choose well below 50,
- it explains why equilibrium play at 0 is rare in initial responses,

Where level-k fits the beauty-contest data well



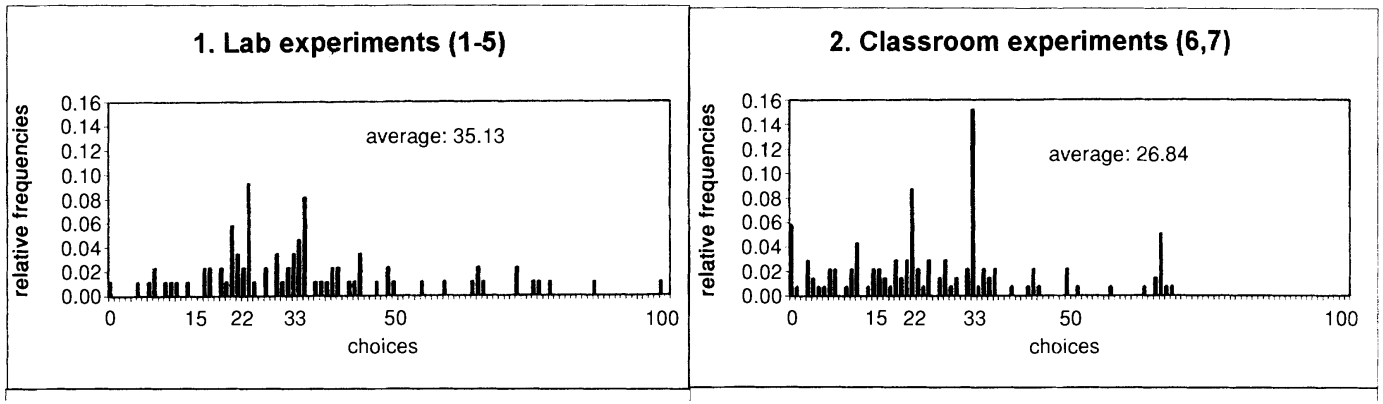
Source: Bosch-Domènech et al. (2002)

Level-k improves sharply on both Nash and rationalizability as a descriptive account:

- it explains why many subjects choose well below 50,
- it explains why equilibrium play at 0 is rare in initial responses,
- it explains why values near 33, 22, and 15 often attract substantial mass.

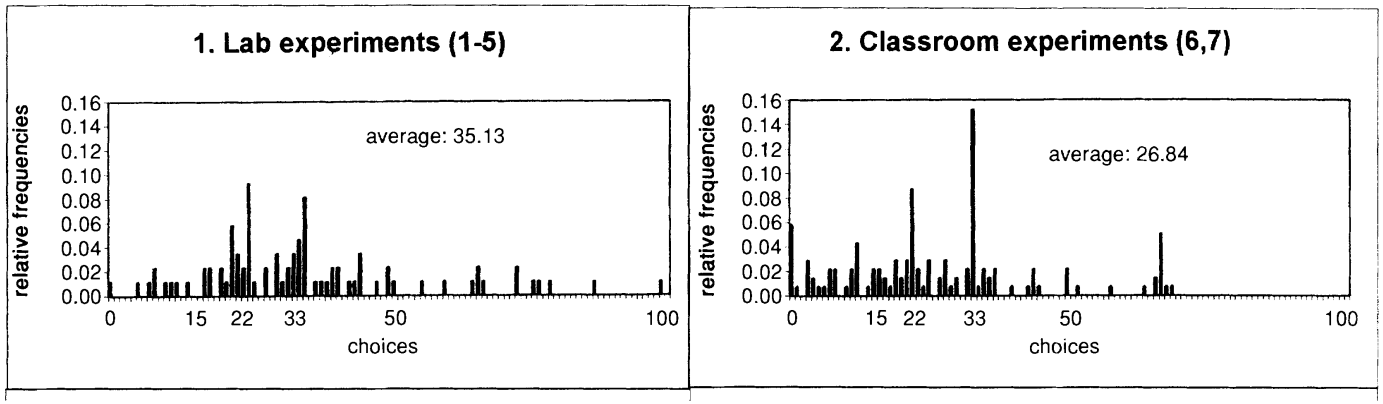
Reading: subjects do reason strategically, but only to finite depth.

Where level-k misses the data



Level-k does not explain everything we see in beauty-contest data.

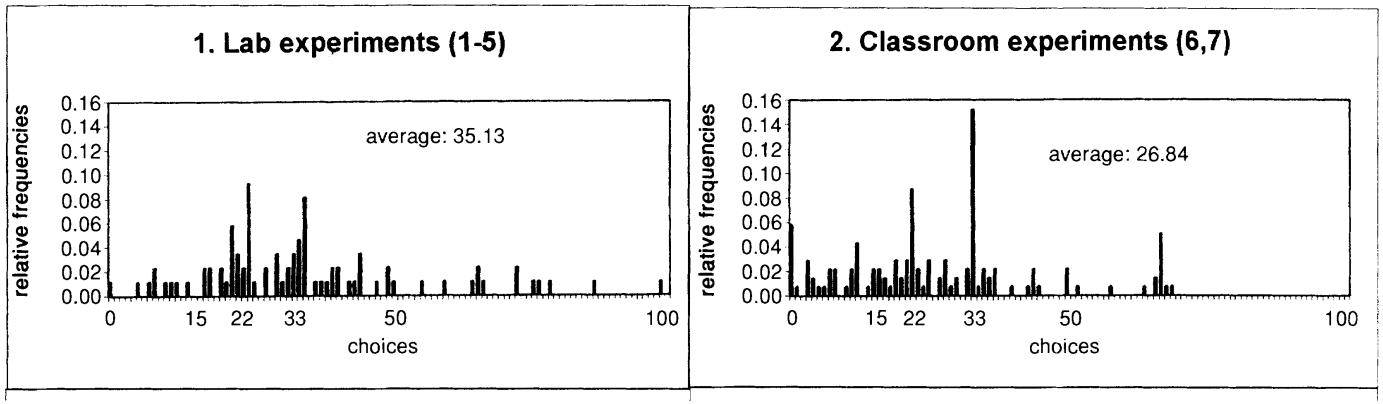
Where level-k misses the data



Level-k does not explain everything we see in beauty-contest data.

- Observed mass is smooth rather than concentrated at canonical levels.

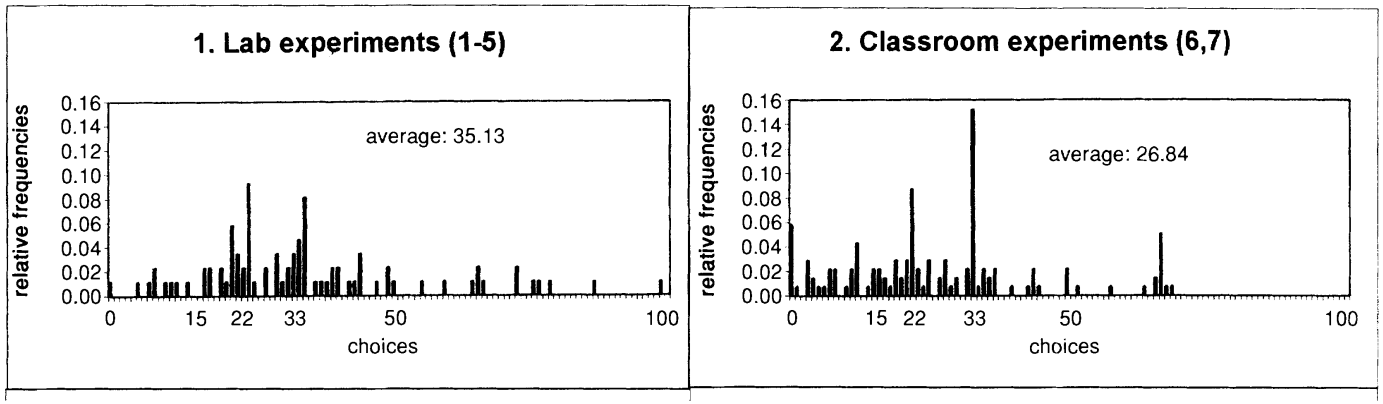
Where level-k misses the data



Level-k does not explain everything we see in beauty-contest data.

- Observed mass is smooth rather than concentrated at canonical levels.
- The fit depends sensitively on how level-0 is specified.

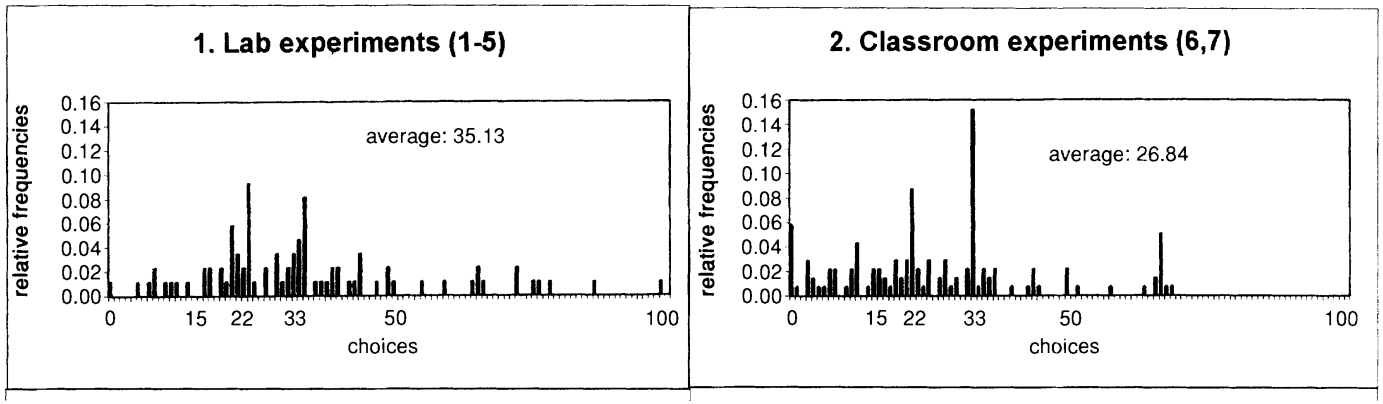
Where level-k misses the data



Level-k does not explain everything we see in beauty-contest data.

- Observed mass is smooth rather than concentrated at canonical levels.
- The fit depends sensitively on how level-0 is specified.
- Similar observed choice can be rationalized by different underlying beliefs.

Where level-k misses the data



Level-k does not explain everything we see in beauty-contest data.

- Observed mass is smooth rather than concentrated at canonical levels.
- The fit depends sensitively on how level-0 is specified.
- Similar observed choice can be rationalized by different underlying beliefs.
- Some of the distribution looks more like responses to full-support beliefs than to one-step type beliefs.

So level-k is a major improvement, but it should not be read too literally as a complete map from choices to cognitive types.

What can be inferred from 33, 22, and 15

The values 33, 22, and 15 are highly informative, but they do not justify overly strong conclusions.

What we may conclude

What can be inferred from 33, 22, and 15

The values 33, 22, and 15 are highly informative, but they do not justify overly strong conclusions.

What we may conclude

- choices are not random,

What can be inferred from 33, 22, and 15

The values 33, 22, and 15 are highly informative, but they do not justify overly strong conclusions.

What we may conclude

- choices are not random,
- finite strategic reasoning is a real feature of the data,

What can be inferred from 33, 22, and 15

The values 33, 22, and 15 are highly informative, but they do not justify overly strong conclusions.

What we may conclude

- choices are not random,
- finite strategic reasoning is a real feature of the data,
- level-k captures an important part of the empirical pattern.

What we should not conclude

What can be inferred from 33, 22, and 15

The values 33, 22, and 15 are highly informative, but they do not justify overly strong conclusions.

What we may conclude

- choices are not random,
- finite strategic reasoning is a real feature of the data,
- level-k captures an important part of the empirical pattern.

What we should not conclude

- every choice near 33 identifies a literal level-1 subject,

What can be inferred from 33, 22, and 15

The values 33, 22, and 15 are highly informative, but they do not justify overly strong conclusions.

What we may conclude

- choices are not random,
- finite strategic reasoning is a real feature of the data,
- level-k captures an important part of the empirical pattern.

What we should not conclude

- every choice near 33 identifies a literal level-1 subject,
- every cluster corresponds one-for-one to a cognitive type,

What can be inferred from 33, 22, and 15

The values 33, 22, and 15 are highly informative, but they do not justify overly strong conclusions.

What we may conclude

- choices are not random,
- finite strategic reasoning is a real feature of the data,
- level-k captures an important part of the empirical pattern.

What we should not conclude

- every choice near 33 identifies a literal level-1 subject,
- every cluster corresponds one-for-one to a cognitive type,
- level-k is the only plausible explanation of the data.

This is one of the main lessons from Breitmoser (2012).

A useful model, but not the final word

Level-k is an excellent first descriptive model of beauty-contest play.

It succeeds because it captures:

A useful model, but not the final word

Level-k is an excellent first descriptive model of beauty-contest play.

It succeeds because it captures:

- finite strategic depth,

A useful model, but not the final word

Level-k is an excellent first descriptive model of beauty-contest play.

It succeeds because it captures:

- finite strategic depth,
- heterogeneity across subjects,

A useful model, but not the final word

Level-k is an excellent first descriptive model of beauty-contest play.

It succeeds because it captures:

- finite strategic depth,
- heterogeneity across subjects,
- structured non-equilibrium play.

But more recent work also gives two reasons for caution:

A useful model, but not the final word

Level-k is an excellent first descriptive model of beauty-contest play.

It succeeds because it captures:

- finite strategic depth,
- heterogeneity across subjects,
- structured non-equilibrium play.

But more recent work also gives two reasons for caution:

- reasoning depth may be endogenous, depending on incentives and cognitive costs (Alaoui and Penta, 2016);

A useful model, but not the final word

Level-k is an excellent first descriptive model of beauty-contest play.

It succeeds because it captures:

- finite strategic depth,
- heterogeneity across subjects,
- structured non-equilibrium play.

But more recent work also gives two reasons for caution:

- reasoning depth may be endogenous, depending on incentives and cognitive costs (Alaoui and Penta, 2016);
- measured depth is often not stable across games, so level-k should not be treated as a fixed personality label (Cooper et al., 2024).

Bottom line: level-k is the right first model for this lecture, but not the end of the story.

Question for next week

Guessing games

The open question after level-k

Level-k is a major improvement over Nash equilibrium and dominance iteration.

It explains well why:

The open question after level-k

Level-k is a major improvement over Nash equilibrium and dominance iteration.

It explains well why:

- equilibrium play at 0 is rare in initial responses,

The open question after level-k

Level-k is a major improvement over Nash equilibrium and dominance iteration.

It explains well why:

- equilibrium play at 0 is rare in initial responses,
- many subjects choose values such as 33, 22, or 15,

The open question after level-k

Level-k is a major improvement over Nash equilibrium and dominance iteration.

It explains well why:

- equilibrium play at 0 is rare in initial responses,
- many subjects choose values such as 33, 22, or 15,
- observed play reflects finite strategic depth rather than no reasoning.

The open question after level-k

Level-k is a major improvement over Nash equilibrium and dominance iteration.

It explains well why:

- equilibrium play at 0 is rare in initial responses,
- many subjects choose values such as 33, 22, or 15,
- observed play reflects finite strategic depth rather than no reasoning.

But an important part of the beauty-contest distribution remains insufficiently explained:

The open question after level-k

Level-k is a major improvement over Nash equilibrium and dominance iteration.

It explains well why:

- equilibrium play at 0 is rare in initial responses,
- many subjects choose values such as 33, 22, or 15,
- observed play reflects finite strategic depth rather than no reasoning.

But an important part of the beauty-contest distribution remains insufficiently explained:

- choice frequencies are often smoother than a pure type-by-type story suggests;

The open question after level-k

Level-k is a major improvement over Nash equilibrium and dominance iteration.

It explains well why:

- equilibrium play at 0 is rare in initial responses,
- many subjects choose values such as 33, 22, or 15,
- observed play reflects finite strategic depth rather than no reasoning.

But an important part of the beauty-contest distribution remains insufficiently explained:

- choice frequencies are often smoother than a pure type-by-type story suggests;
- many observations do not sit naturally at canonical level-k values;

The open question after level-k

Level-k is a major improvement over Nash equilibrium and dominance iteration.

It explains well why:

- equilibrium play at 0 is rare in initial responses,
- many subjects choose values such as 33, 22, or 15,
- observed play reflects finite strategic depth rather than no reasoning.

But an important part of the beauty-contest distribution remains insufficiently explained:

- choice frequencies are often smoother than a pure type-by-type story suggests;
- many observations do not sit naturally at canonical level-k values;
- the same observed choice can be consistent with several different belief structures.

The open question after level-k

Level-k is a major improvement over Nash equilibrium and dominance iteration.

It explains well why:

- equilibrium play at 0 is rare in initial responses,
- many subjects choose values such as 33, 22, or 15,
- observed play reflects finite strategic depth rather than no reasoning.

But an important part of the beauty-contest distribution remains insufficiently explained:

- choice frequencies are often smoother than a pure type-by-type story suggests;
- many observations do not sit naturally at canonical level-k values;
- the same observed choice can be consistent with several different belief structures.

So the next question is:

What if subjects respond to diffuse, full-support beliefs rather than to one specific lower type?

Main takeaways

Main takeaways

1. The beauty contest is a clean test of strategic depth in one-shot games.

Main takeaways

1. The beauty contest is a clean test of strategic depth in one-shot games.
2. Nash equilibrium gives a sharp benchmark, but it does not organize observed initial play well.

Main takeaways

1. The beauty contest is a clean test of strategic depth in one-shot games.
2. Nash equilibrium gives a sharp benchmark, but it does not organize observed initial play well.
3. Non-equilibrium play does not imply irrationality: subjects may best-respond to mistaken or simplified beliefs.

Main takeaways

1. The beauty contest is a clean test of strategic depth in one-shot games.
2. Nash equilibrium gives a sharp benchmark, but it does not organize observed initial play well.
3. Non-equilibrium play does not imply irrationality: subjects may best-respond to mistaken or simplified beliefs.
4. Rationalizability and dominance provide useful restrictions, but not a convincing empirical description.

Main takeaways

1. The beauty contest is a clean test of strategic depth in one-shot games.
2. Nash equilibrium gives a sharp benchmark, but it does not organize observed initial play well.
3. Non-equilibrium play does not imply irrationality: subjects may best-respond to mistaken or simplified beliefs.
4. Rationalizability and dominance provide useful restrictions, but not a convincing empirical description.
5. Level-k is the first genuinely successful descriptive model in this setting.

Main takeaways

1. The beauty contest is a clean test of strategic depth in one-shot games.
2. Nash equilibrium gives a sharp benchmark, but it does not organize observed initial play well.
3. Non-equilibrium play does not imply irrationality: subjects may best-respond to mistaken or simplified beliefs.
4. Rationalizability and dominance provide useful restrictions, but not a convincing empirical description.
5. Level-k is the first genuinely successful descriptive model in this setting.
6. Level-k is still incomplete, which motivates richer models of beliefs and responses.

Next lecture

When the data are too smooth for a literal level-k interpretation, is the main problem bounded depth, payoff-sensitive noise, or both?

References

Guessing games

References I

Alaoui, L. and Penta, A. (2016). Endogenous depth of reasoning. *The Review of Economic Studies*, 83(4):1297–1333.

Bosch-Domènech, A., García-Montalvo, J. A., Nagel, R., and Satorra, A. (2002). One, two, (three), infinity: Newspaper and lab beauty-contest experiments. *American Economic Review*, 92(5):1687–1701.

Breitmoser, Y. (2012). Strategic reasoning in p-beauty contests. *Games and Economic Behavior*, 75(2):555–569.

Camerer, C., Ho, T.-H., and Chong, J. K. (2004). A cognitive hierarchy model of games. *Quarterly Journal of Economics*, 119(3):861–898.

Cooper, D. J., Fatas, E., Morales, A. J., and Qi, S. (2024). Consistent depth of reasoning in level-k models. *American Economic Journal: Microeconomics*, 16(4):40–76.

Fudenberg, D. and Liang, A. (2019). Predicting and understanding initial play. *American Economic Review*, 109(12):4112–4141.

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

Nagel, R. (1995). Unraveling in guessing games: An experimental study. *American Economic Review*, 85(5):1313–1326.

Repetition

Guessing games

Repetition I: benchmark and evidence

1. Why is 0 the Nash equilibrium in the standard beauty contest?
Answer sketch: Iterated reasoning drives all admissible choices downward because with $p = \frac{2}{3} < 1$, the target is always below the average. Repeating this logic yields the unique equilibrium at 0.

2. Why is the beauty contest a useful test of strategic reasoning?
Answer sketch: The action space is simple, but success depends on anticipating others' beliefs and choices. So the game isolates higher-order strategic reasoning in a transparent way.

3. What is the main empirical puzzle in the data?
Answer sketch: Choices are typically far above the Nash equilibrium, but they are not random. They show structured non-equilibrium play with substantial mass at salient intermediate values.

Repetition II: benchmark concepts

4. Why does non-equilibrium play not imply irrationality?
Answer sketch: A player can be rational if she best-responds to her subjective beliefs, even when those beliefs are incorrect. Equilibrium additionally requires correct beliefs.

5. What is the role of weak dominance in the beauty contest?
Answer sketch: Weak dominance generates downward thresholds such as 66.7, 44.4, and 29.6, showing that very high choices are implausible. But those thresholds do not explain the observed clusters well.

6. Why is rationalizability not a convincing description of the data?
Answer sketch: In continuous action spaces it is too permissive; in fine discrete spaces it may require implausibly many elimination rounds. So it gives useful restrictions, but not a persuasive account of actual initial play.

Repetition III: behavioral models

7. What does logit response add relative to exact best response?
Answer sketch: It allows payoff-sensitive noise: better actions are chosen more often, but not with certainty. So choice is structured, not perfectly deterministic.

8. What is the central idea of level-k reasoning?
Answer sketch: Players begin from a simple non-equilibrium anchor and perform only finitely many reasoning steps. Higher levels best-respond to lower-level behavior.

9. Why are values such as 33, 22, and 15 important?
Answer sketch: With level-0 anchored around 50, these values arise from one, two, and three rounds of multiplying by $\frac{2}{3}$. They are therefore natural markers of finite strategic depth.

Repetition IV: interpretation and transition

10. Why should we not interpret every choice near 33 as literal level-1 play?
Answer sketch: The same observed choice may be generated by different belief structures or noisy responses. So clusters are suggestive, but they do not identify cognitive types one-for-one.

11. In what sense is level-k successful, and in what sense is it incomplete?
Answer sketch: It successfully explains finite-depth reasoning and the main clusters in the data. But it leaves part of the distribution too smooth or too diffuse to be explained convincingly by pure discrete types alone.

12. What is the main question left for the next lecture?
Answer sketch: Whether models with richer, full-support beliefs and noisy responses can explain the remaining distribution better than a literal level-k account.

Key topics

Key topics

- Standard assumption Subjects play logit responses to their beliefs

Key topics

- Standard assumption Subjects play logit responses to their beliefs
- Main question How can we model belief formation?

p-Beauty contest

Each player picks a real number from $[0, 100]$. The player closest to $2/3$ of the mean of all guesses wins a prize (sharing the prize in case of ties).

Key topics

- Standard assumption Subjects play logit responses to their beliefs
- Main question How can we model belief formation?

p-Beauty contest

Each player picks a real number from $[0, 100]$. The player closest to $2/3$ of the mean of all guesses wins a prize (sharing the prize in case of ties).

Key topics

- Standard assumption Subjects play logit responses to their beliefs
- Main question How can we model belief formation?

p-Beauty contest

Each player picks a real number from $[0, 100]$. The player closest to $2/3$ of the mean of all guesses wins a prize (sharing the prize in case of ties).

- The only dominated/not rationalizable choice is 100

Key topics

- Standard assumption Subjects play logit responses to their beliefs
- Main question How can we model belief formation?

p-Beauty contest

Each player picks a real number from $[0, 100]$. The player closest to $2/3$ of the mean of all guesses wins a prize (sharing the prize in case of ties).

- The only dominated/not rationalizable choice is 100
  - observed behavior does not violate rationality or common knowledge of rationality

Key topics

- Standard assumption Subjects play logit responses to their beliefs
- Main question How can we model belief formation?

p-Beauty contest

Each player picks a real number from $[0, 100]$. The player closest to $2/3$ of the mean of all guesses wins a prize (sharing the prize in case of ties).

- The only dominated/not rationalizable choice is 100
  - observed behavior does not violate rationality or common knowledge of rationality
- All $\sigma_i \geq 2/3 \cdot 100$ are weakly dominated

Key topics

- Standard assumption Subjects play logit responses to their beliefs
- Main question How can we model belief formation?

p-Beauty contest

Each player picks a real number from $[0, 100]$. The player closest to $2/3$ of the mean of all guesses wins a prize (sharing the prize in case of ties).

- The only dominated/not rationalizable choice is 100
  - observed behavior does not violate rationality or common knowledge of rationality
- All $\sigma_i \geq 2/3 \cdot 100$ are weakly dominated
  - not rationalizable in presence of noise

Key topics

- Standard assumption Subjects play logit responses to their beliefs
- Main question How can we model belief formation?

p-Beauty contest

Each player picks a real number from $[0, 100]$. The player closest to $2/3$ of the mean of all guesses wins a prize (sharing the prize in case of ties).

- The only dominated/not rationalizable choice is 100
  - observed behavior does not violate rationality or common knowledge of rationality
- All $\sigma_i \geq 2/3 \cdot 100$ are weakly dominated
  - not rationalizable in presence of noise
  - same holds true for all actions above $100 \cdot (2/3)^2 \approx 44.4$, $100 \cdot (2/3)^3 \approx 29.6$, ...

Notation: Actions and strategies

Player i has to choose an action $a_i \in A_i$. Player i's strategy $\sigma_i \in \Delta(A_i)$ is a probability distribution on A_i. $\sigma_{-i} = \cup_{j \neq i} \sigma_j$

Belief

Player i's belief about j's strategy is a probability distribution $\mu_{i,j} \in \Delta(A_j)$. E.g. $\mu_{i,j}(a_j)$ is the prob that i assigns to j playing a_j. Define $\mu_{i,-i} = \cup_{j \neq i} \mu_{i,j}$

Notation: Actions and strategies

Player i has to choose an action $a_i \in A_i$. Player i's strategy $\sigma_i \in \Delta(A_i)$ is a probability distribution on A_i. $\sigma_{-i} = \cup_{j \neq i} \sigma_j$

Belief

Player i's belief about j's strategy is a probability distribution $\mu_{i,j} \in \Delta(A_j)$. E.g. $\mu_{i,j}(a_j)$ is the prob that i assigns to j playing a_j. Define $\mu_{i,-i} = \cup_{j \neq i} \mu_{i,j}$

Logit Response

A strategy σ_i is a logit response to belief $\mu_{i,-i}$ if there exists $\lambda \geq 0$ such that

$$\text{for all } a_i: \quad \sigma_i(a_i|\mu_{i,-i}) = \frac{\exp\{\lambda \cdot \pi_i(a_i, \mu_{i,-i})\}}{\sum_{a_i'} \exp\{\lambda \cdot \pi_i(a_i', \mu_{i,-i})\}}.$$

Notation: Actions and strategies

Player i has to choose an action $a_i \in A_i$. Player i's strategy $\sigma_i \in \Delta(A_i)$ is a probability distribution on A_i. $\sigma_{-i} = \cup_{j \neq i} \sigma_j$

Belief

Player i's belief about j's strategy is a probability distribution $\mu_{i,j} \in \Delta(A_j)$. E.g. $\mu_{i,j}(a_j)$ is the prob that i assigns to j playing a_j. Define $\mu_{i,-i} = \cup_{j \neq i} \mu_{i,j}$

Logit Response

A strategy σ_i is a logit response to belief $\mu_{i,-i}$ if there exists $\lambda \geq 0$ such that

$$\text{for all } a_i: \quad \sigma_i(a_i|\mu_{i,-i}) = \frac{\exp\{\lambda \cdot \pi_i(a_i, \mu_{i,-i})\}}{\sum_{a_i'} \exp\{\lambda \cdot \pi_i(a_i', \mu_{i,-i})\}}.$$

Notation: Actions and strategies

Player i has to choose an action $a_i \in A_i$. Player i's strategy $\sigma_i \in \Delta(A_i)$ is a probability distribution on A_i. $\sigma_{-i} = \cup_{j \neq i} \sigma_j$

Belief

Player i's belief about j's strategy is a probability distribution $\mu_{i,j} \in \Delta(A_j)$. E.g. $\mu_{i,j}(a_j)$ is the prob that i assigns to j playing a_j. Define $\mu_{i,-i} = \cup_{j \neq i} \mu_{i,j}$

Logit Response

A strategy σ_i is a logit response to belief $\mu_{i,-i}$ if there exists $\lambda \geq 0$ such that

$$\text{for all } a_i: \quad \sigma_i(a_i|\mu_{i,-i}) = \frac{\exp\{\lambda \cdot \pi_i(a_i, \mu_{i,-i})\}}{\sum_{a_i'} \exp\{\lambda \cdot \pi_i(a_i', \mu_{i,-i})\}}.$$

- $\lambda \in_+$ is called "precision": the higher λ, the more likely the best decision

Notation: Actions and strategies

Player i has to choose an action $a_i \in A_i$. Player i's strategy $\sigma_i \in \Delta(A_i)$ is a probability distribution on A_i. $\sigma_{-i} = \cup_{j \neq i} \sigma_j$

Belief

Player i's belief about j's strategy is a probability distribution $\mu_{i,j} \in \Delta(A_j)$. E.g. $\mu_{i,j}(a_j)$ is the prob that i assigns to j playing a_j. Define $\mu_{i,-i} = \cup_{j \neq i} \mu_{i,j}$

Logit Response

A strategy σ_i is a logit response to belief $\mu_{i,-i}$ if there exists $\lambda \geq 0$ such that

$$\text{for all } a_i: \quad \sigma_i(a_i|\mu_{i,-i}) = \frac{\exp\{\lambda \cdot \pi_i(a_i, \mu_{i,-i})\}}{\sum_{a_i'} \exp\{\lambda \cdot \pi_i(a_i', \mu_{i,-i})\}}.$$

- $\lambda \in_+$ is called "precision": the higher λ, the more likely the best decision
- $\lambda = 0$: Uniform randomization

Notation: Actions and strategies

Player i has to choose an action $a_i \in A_i$. Player i's strategy $\sigma_i \in \Delta(A_i)$ is a probability distribution on A_i. $\sigma_{-i} = \cup_{j \neq i} \sigma_j$

Belief

Player i's belief about j's strategy is a probability distribution $\mu_{i,j} \in \Delta(A_j)$. E.g. $\mu_{i,j}(a_j)$ is the prob that i assigns to j playing a_j. Define $\mu_{i,-i} = \cup_{j \neq i} \mu_{i,j}$

Logit Response

A strategy σ_i is a logit response to belief $\mu_{i,-i}$ if there exists $\lambda \geq 0$ such that

$$\text{for all } a_i: \quad \sigma_i(a_i|\mu_{i,-i}) = \frac{\exp\{\lambda \cdot \pi_i(a_i, \mu_{i,-i})\}}{\sum_{a_i'} \exp\{\lambda \cdot \pi_i(a_i', \mu_{i,-i})\}}.$$

- $\lambda \in_+$ is called "precision": the higher λ, the more likely the best decision
- $\lambda = 0$: Uniform randomization
- $\lambda \to \infty$ implies convergence to a rationality

Notation: Actions and strategies

Player i has to choose an action $a_i \in A_i$. Player i's strategy $\sigma_i \in \Delta(A_i)$ is a probability distribution on A_i. $\sigma_{-i} = \cup_{j \neq i} \sigma_j$

Belief

Player i's belief about j's strategy is a probability distribution $\mu_{i,j} \in \Delta(A_j)$. E.g. $\mu_{i,j}(a_j)$ is the prob that i assigns to j playing a_j. Define $\mu_{i,-i} = \cup_{j \neq i} \mu_{i,j}$

Logit Response

A strategy σ_i is a logit response to belief $\mu_{i,-i}$ if there exists $\lambda \geq 0$ such that

$$\text{for all } a_i: \quad \sigma_i(a_i|\mu_{i,-i}) = \frac{\exp\{\lambda \cdot \pi_i(a_i, \mu_{i,-i})\}}{\sum_{a_i'} \exp\{\lambda \cdot \pi_i(a_i', \mu_{i,-i})\}}.$$

- $\lambda \in_+$ is called "precision": the higher λ, the more likely the best decision
- $\lambda = 0$: Uniform randomization
- $\lambda \to \infty$ implies convergence to a rationality
- For all $\lambda < \infty$: All actions have positive probability

Notation: Actions and strategies

Player i has to choose an action $a_i \in A_i$. Player i's strategy $\sigma_i \in \Delta(A_i)$ is a probability distribution on A_i. $\sigma_{-i} = \cup_{j \neq i} \sigma_j$

Belief

Player i's belief about j's strategy is a probability distribution $\mu_{i,j} \in \Delta(A_j)$. E.g. $\mu_{i,j}(a_j)$ is the prob that i assigns to j playing a_j. Define $\mu_{i,-i} = \cup_{j \neq i} \mu_{i,j}$

Logit Response

A strategy σ_i is a logit response to belief $\mu_{i,-i}$ if there exists $\lambda \geq 0$ such that

$$\text{for all } a_i: \quad \sigma_i(a_i|\mu_{i,-i}) = \frac{\exp\{\lambda \cdot \pi_i(a_i, \mu_{i,-i})\}}{\sum_{a_i'} \exp\{\lambda \cdot \pi_i(a_i', \mu_{i,-i})\}}.$$

- $\lambda \in_+$ is called "precision": the higher λ, the more likely the best decision
- $\lambda = 0$: Uniform randomization
- $\lambda \to \infty$ implies convergence to a rationality
- For all $\lambda < \infty$: All actions have positive probability
- See "Example to illustrate the computation"