QRE and Level-k

Summer 26

Niels Boissonnet, Bielefeld

Discussion

From level-k to full-support beliefs

QRE and Level-k

From last lecture to today

Last lecture introduced the beauty contest as the first clean test of strategic reasoning in a one-shot game.

What we established

- The benchmark prediction is sharp: for $p < 1$, iterated reasoning drives play toward the unique Nash equilibrium at 0.
- The data are not close to Nash equilibrium, but they are far from random.
- Choices cluster at salient intermediate values such as 50, 33, 22, and sometimes 15.
- A natural first interpretation is that subjects differ in how many steps of reasoning they perform.

What we get from Level-k thinking

Level-k thinking is a useful first descriptive model:

$\Longrightarrow$ it organizes the most salient masses in the distribution and gives a simple language for bounded strategic reasoning.

But the same evidence already suggests a limit of literal level-k models:

$\Longrightarrow$ observed play is often smoother and more dispersed than a finite collection of pure reasoning types would predict.

- Nagel (1995), Ho et al. (1998), and Breitmoser (2012).

Purpose of this lecture

Three descriptive questions open.

1. observed behavior is often *payoff-sensitive but not literal*.
   - Many choices move in the right direction relative to incentives,
   - yet they do not sit exactly at the canonical level-1, level-2,... values.
2. empirical distribution has *full support* over a broad range of actions.
   - awkward for a model in which each type has a single best reply to a sharply specified lower type.
3. says little about beliefs rationalizing of non-canonical focal values.
   - If a subject chooses, say, 27 or 19

So the next step is not to abandon strategic reasoning, but to enrich it.

Suppose we keep one attractive idea from the previous lecture:

Subjects respond strategically, but not with exact best responses.

Then a behavioral model of one-shot play must answer more than one question.

Question for today

How should we model behavior when subjects still respond to payoff incentives, but their beliefs and responses are too smooth for literal finite level-k types?

Key modeling problem behind behavioral solution concepts

$\Longrightarrow$ not only how subjects choose,

$\Longrightarrow$ but how the hierarchy of beliefs is structured.

What has to be modeled once choices are noisy?

Own response

- How sensitive is my choice to expected-payoff differences?
- Do I almost always choose the better option, or only more often?

Beliefs about others

- How precise do I think *your* responses are?
- Do I expect you to be as precise as I am, less precise, or random?

Higher-order beliefs

- What do I think that you think about my precision?
- When does this hierarchy stop?
- Does it stop at a finite depth, converge gradually to zero precision?

Noise: Three ways to structure strategic reasoning

1. Finite reasoning depth: level-k
   - At some finite stage, beliefs reduce to "zero precision": the player treats lower-level opponents as random or exogenously simple.
   - This yields a finite chain of iterated responses.
2. Rational expectations with noisy choice: QRE
   - Everyone responds smoothly to expected payoffs.
   - Everyone is correct about everyone else's noisy behavior.
   - The model is therefore a fixed point: beliefs and behavior must coincide.
3. Infinite reasoning with fading precision: noisy introspection
   - Players think through many reasoning layers, but deeper layers become less precise.
   - Keeps strategic reasoning, but relaxes full rational-expectations assumption.

Today: focus on 2. and 3., and then return to the beauty-contest evidence.

Learning goals and route through today

Learning goals

- Explain what a quantal response equilibrium is, and why it is a fixed point rather than a finite backward-induction exercise.
- Explain the logic of noisy introspection, and how it differs from both finite level-k and QRE.
- Interpret the key parameters of these models as restrictions on *response precision* and *belief formation*.
- Use the beauty-contest evidence to compare what the different models can and cannot explain.

Route through the lecture

1. Why level-k is useful but incomplete
2. QRE: logit responses and rational expectations
3. Noisy introspection: infinite reasoning with declining precision
4. Evidence from guessing games and what it implies for model comparison

Quantal response equilibrium

QRE and Level-k

QRE is the natural next step after level-k

Level-k is often too discrete as a literal model of behavior.

- Empirical choice distributions are smoother than a finite collection of pure reasoning types.
- Many subjects react in the *right direction* to incentives.
- **Need:** a model where better actions are chosen more often, but not necessarily with probability one.

QRE keeps strategic reasoning, but smooths choice.

- Players respond to expected payoffs *probabilistically*.
- Better actions are more likely, but inferior actions may still be chosen.
- Beliefs are correct about others' stochastic behavior.

QRE is a natural setp after level-k

So QRE is best read as a compromise between two extremes:

not exact best response, but not arbitrary noise either.

The basic framework is due to McKelvey and Palfrey (1995).

Logit equilibrium: definition

Definition (Logit QRE, McKelvey and Palfrey, 1995)

A mixed-strategy profile σ is a **logit equilibrium** if there exists a precision parameter $\lambda \geq 0$ such that, for every player i and every action $a_i \in A_i$,

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

How to read the formula?

- $\pi_i(a_i, \sigma_{-i})$=exp. payoff from action a_i against opponents' mixed strategies.
- The player chooses actions with probabilities that increase in exp. payoffs.
- The same profile σ appears both in expected payoffs and in choice probabilities: QRE is therefore a fixed point.

Logit equilibrium: interpretation

Definition (Logit QRE, McKelvey and Palfrey, 1995)

A mixed-strategy profile σ is a **logit equilibrium** if there exists a precision parameter $\lambda \geq 0$ such that, for every player i and every action $a_i \in A_i$,

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

Special cases

- $\lambda = 0$: all actions are chosen uniformly.
- $\lambda > 0$: better actions are more likely, but all actions remain possible.
- $\lambda \to \infty$: choice converges toward best response and, along suitable branches, toward Nash equilibrium.

Interpreting the precision parameter

Economic meaning of the precision parameter λ

- It measures how strongly the player reacts to expected-payoff differences.
- If λ is small, payoff differences matter little and play is close to random.
- If λ is large, even small payoff diff. lead to concentrated choice probabilities.

What λ does *not* measure

- It is *not* itself a level of strategic reasoning.
- It is *not* a direct measure of cognitive depth.
- It captures the noisiness of *response*, not the structure of the belief hierarchies.

QRE and level-k are not rival They relax different parts of the Nash benchmark.

- Level-k is mainly about the structure of beliefs;
- QRE is mainly about smooth response plus rational expectations.

QRE as a fixed-point

Level-k model is computed by backward induction over reasoning levels:

- higher levels are defined recursively.

QRE is different.

- Every player responds smoothly to expected payoffs.
- Expected payoffs depend on beliefs about others' mixed strategies.
- Beliefs are correct, they coincide with the actual mixed strategies in the population.

This has two important consequences

1. QRE preserves the equilibrium structure of mutual consistency.
2. But it can no longer be computed by finite backward induction in the way level-k can.

So QRE is conceptually attractive but computationally harder: we need a fixed point of the logit-response mapping.

Existence and basic properties of logit QRE

Theorem (McKelvey and Palfrey (1995))

For every finite normal-form game and every $\lambda \geq 0$, a logit equilibrium exists.

Proof idea

- For each player, logit response is a continuous function of opponents' mixed strategies.
- Hence the profile of logit responses defines a continuous mapping from the product of simplices into itself.
- By Brouwer's fixed-point theorem, this mapping has a fixed point.

How we compute QRE in practice

Step 1: define the logit-response mapping

$$\sigma_i^{t+1}(a_i, \lambda) = \frac{\exp\{\lambda \cdot \pi_i(a_i, \sigma_{-i}^t(\lambda)\}}{\sum_{a_i' \in A_i} \exp\{\lambda \cdot \pi_i(a_i', \sigma_{-i}^t(\lambda))\}}. \tag{1}$$

Step 2: look for a fixed point numerically

$$\sigma^{t+1}(a, \lambda) = (\sigma_1^{t+1}(a_1, \lambda), ..., \sigma_n^{t+1}(a_n, \lambda)) = (\sigma_1^t(a_1, \lambda), ..., \sigma_n^t(a_n, \lambda)) = \sigma^t(a, \lambda).$$

- start from some initial profile σ^0,
- Given $\sigma^{(t)}$ compute σ^{t+1},
- stop once $\|\sigma^{(t+1)} - \sigma^{(t)}\|$ is sufficiently small.

Step 3: Estimate λ given observed frequencies σ^* through Max log likelihood

$$\lambda^{est} = \arg\max_{\lambda} \sum_i \sum_{a_i} n_{i,a_i} \cdot \log \sigma_i^*(a_i; \lambda)$$

Pros and cons

What's nice with this method

- It makes the fixed-point nature of QRE very transparent.
- In simple games, it often works well and is easy to code.

Caveat:

- Need not converge globally: we need a more stable branch-following method.

QRE pros and cons

Why QRE is attractive empirically

- It predicts smooth deviations from NE in the direction of better responses.
- It allows substantial dispersion without abandoning strategic incentives.
- It describes one-shot laboratory data better than exact NE.

But QRE should not be oversold

- It imposes rational expectations about noisy play.
- In genuinely unfamiliar one-shot games, that may be a strong assumption.
- Limited empirical content: many distributions can be rationalized by QRE.

So for us, QRE is a benchmark — not the final word.

- It is the cleanest equilibrium-based model of smooth strategic response.
- But it may still be too demanding as a model of one-shot reasoning.
- That is exactly why the next section turns to noisy introspection.

For the empirical-content critique, see Haile et al. (2008).

Why QRE matters for beauty contests

What QRE contributes relative to level-k

- It can generate *smooth* distributions concentrated on low guesses.
- It does not require to decompose the data into a few exact reasoning types.
- It interprets intermediate guesses as payoff-sensitive resp. to smooth beliefs.

What it does *not* automatically explain

- why subjects hold rational expectations in an unfamiliar one-shot setting,
- why focal masses such as 33 or 22 remain so persistent,
- and why some data look more like finite-step reasoning.

Beauty contest Level-k vs QRE

This is the right way to think about the beauty-contest evidence:

QRE explains the smooth part of the distribution well; level-k helps explain the focal masses.

Breitmoser (2012) compares finite reasoning models with richer full-support concepts rather than treating one of them as universally correct.

Noisy introspection

QRE and Level-k

Why move beyond QRE?

But for one-shot games, QRE may still be too demanding.

- QRE assumes *rational expectations* about others' noisy play.
- In an unfamiliar one-shot game, may not correctly anticipate the exact distribution of others' choices.
- So the equilibrium restriction may be stronger than the evidence justifies.

Idea of noisy introspection

- Keep *logit responses*: better actions are chosen more often.
- Keep *iterated reasoning*: I think about your choice, your beliefs...
- But let precision *fade* as we go deeper into the hierarchy of beliefs.

Thus, noisy introspection contains two ingredients:

payoff-sensitive choice + increasingly noisy higher-order beliefs.

Noisy Introspection: Interpretation

It is not an equilibrium concept like QRE.

Instead, it is a structured model of how strategic reasoning may unravel in a novel one-shot interaction.

This is the central idea in Goeree and Holt (2004).

Definition of noisy introspection

Noisy introspection (Goeree and Holt, 2004)

A strategy profile σ follows from **noisy introspection** if there exist a base precision $\lambda \in \mathbb{R}_+$, a decay factor $\kappa \in (0, 1)$, and a sequence of mixed-strategy profiles $(\sigma^k)_{k \geq 0}$ such that

$$\sigma^0 = \sigma$$

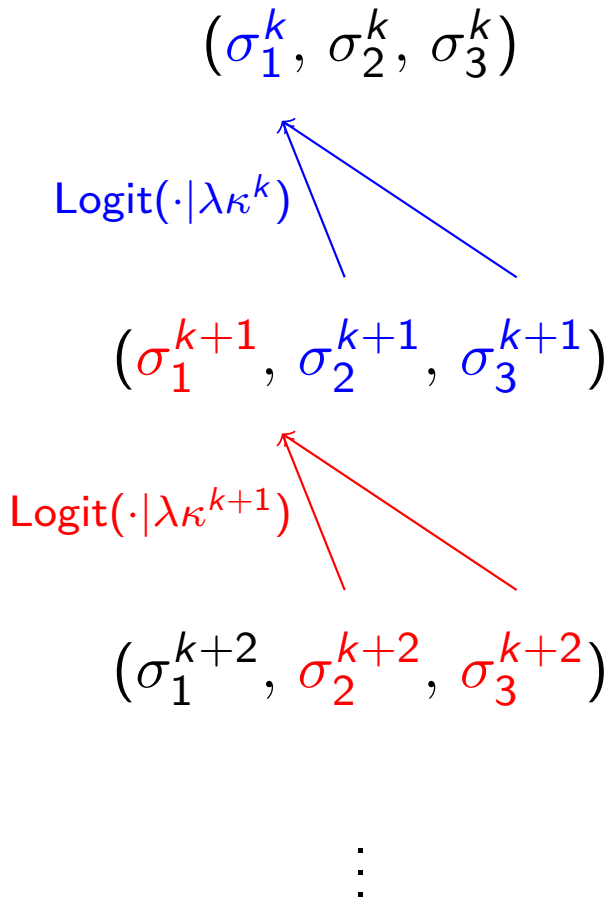
and, for every player i and every level $k \geq 0$,

$$\sigma_i^k = \text{Logit}\left(\sigma_{-i}^{k+1} \mid \lambda \kappa^k\right).$$

How to read the definition

- σ^0 is the actually observed strategy profile.
- σ^1 is the strategy profile I attribute to my opponents.
- σ^2 is what I think they attribute to their opponents, and so on.
- At level k, the relevant precision is $\lambda \kappa^k$, which becomes smaller as k grows.

Graph


$(\sigma_1^k, \sigma_2^k, \sigma_3^k)$
$\text{Logit}(\cdot|\lambda\kappa^k)$
$(\sigma_1^{k+1}, \sigma_2^{k+1}, \sigma_3^{k+1})$
$\text{Logit}(\cdot|\lambda\kappa^{k+1})$
$(\sigma_1^{k+2}, \sigma_2^{k+2}, \sigma_3^{k+2})$
$\vdots$

Economic meaning

- I am relatively precise in my own choice.
- I think others are somewhat less precise.
- I think they regard others as even less precise, and so on.

So NI keeps strategic reasoning infinite in principle, but makes distant levels of introspection increasingly noisy.

What this changes relative to level-k and QRE

Noisy introspection sits exactly between finite level-k models and QRE.

Relative to finite level-k

- NI does *not* stop reasoning after a fixed finite number of steps.
- There is no level at which beliefs suddenly become deterministic types with zero continuation.
- Instead, reasoning continues indefinitely, but its precision fades geometrically.

Relative to QRE

- NI does *not* impose rational expectations.
- My belief about your play need not coincide with your actual play.

What this changes relative to level-k and QRE

Comparison in one line

- **Level-k:** finite reasoning depth.
- **QRE:** infinite depth with common precision and correct beliefs.
- **NI:** infinite depth with declining precision and no equilibrium restriction.

For data NI can generate dispersed, payoff-sensitive behavior without requiring the strong mutual-consistency restriction built into QRE.

$\Longrightarrow$ Goeree and Holt explicitly motivate NI as a one-shot alternative to equilibrium-based smoothing (Goeree and Holt, 2004).

Interpreting the decay factor κ

$\kappa \in (0, 1)$ determines precision decline speed along the belief hierarchy.

If κ is close to 1

- deeper levels remain fairly precise,
- the model comes closer to a high-precision, equilibrium-like description,
- and NI may look empirically quite similar to QRE.

If κ is small

- precision collapses quickly at deeper levels,
- higher-order beliefs become close to randomization,
- and the model behaves more like a shallow reasoning model.

Example Suppose $\lambda = 100$ and $\kappa = \frac{1}{2}$. Then the precision sequence is

$$100,\ 50,\ 25,\ 12.5,\ 6.25,\ \ldots$$

Interpretation NI is therefore a parsimonious way to say:

I think harder about my own choice than I think you think about yours.

Computation of NI strategy profiles

Main computational advantage: Unlike QRE, NI does not require a fixed point.

- Because precision declines geometrically, for sufficiently large K the term $\lambda\kappa^K$ is arbitrarily close to 0.
- At that stage, behavior is approximately uniform randomization.

Simple computation scheme

- Choose K such that $\lambda\kappa^K$ is very small.
- Set σ^{K+1} equal to uniform randomization.
- For $k = K, K-1, \ldots, 0$, compute $\sigma^k = \text{Logit-response}(\sigma^{k+1} \mid \lambda\kappa^k)$.
- The predicted strategy profile is then σ^0.

Computation of NI strategy profiles

This makes of computing NI relatively easy

- NI can be computed by backward induction over belief levels.
- QRE generally cannot, because precision remains positive forever.

R code

```
ni <- function(lambda, kappa, sigma = c(0.5,0.5), K = 40) {
  for (k in K:0) {
    sigma <- logitResponse(lambda * kappa^k, sigma)
  }
  return(sigma)
}
```

QRE versus noisy introspection in beauty contest

QRE and NI predictions:

- *full-support*, payoff-sensitive distributions concentrated on relatively low guesses.

But the shape of the distribution differs.

- QRE produces more equilibrium-like concentration around its modal guess.
- NI produces *thinner peaks* and somewhat more dispersion around the mode.

$\Longrightarrow$ NI fits smooth distributions, yet does not require mutually correct beliefs.

Interpretation for the beauty contest

- If the data are smooth but not sharply equilibrium-like, NI is often a plausible description.
- If the data also display prominent focal masses such as 33, finite-step reasoning may still matter.

Application to guessing games

QRE and Level-k

Returning to the empirical puzzle

What we already know from the previous lecture

- The Nash benchmark is very sharp: equilibrium play is 0.
- Observed choices are far above 0, but far from random.
- Canonical masses such as 50, 33, and 22 suggest bounded steps of reasoning.

What remains open after the level-k discussion

- The empirical distribution is mostly *smooth*.
- Many choices are incentive-sensible, even when not literal level-k choices.

Now, we revisit the beauty-contest evidence with exactly that question in mind:

How much of the data really requires finite-step thinking, and how much is better described by QRE or noisy introspection?

Main empirical finding

Core result: high-level full-support models capture choices better than by literal intermediate level-k types (Breitmoser, 2012).

What this means in practice

- Roughly, most subjects do *not* look like exact level-2 or level-3 reasoners.
- Their choices are consistent with **noisy introspection close to QRE**.
- Yet, the focal mass at 33 is hard to understand without level-1.

Interpretation

- Level-k remains useful for understanding the *focal masses*.
- QRE and NI are useful for understanding the *smooth bulk of the distribution*.
- So the evidence tells us that level-k is only part of the picture.

Bottom line The BC data are best read as a combination of *some low-level focal reasoning* and *a large amount of smooth, payoff-sensitive strategic response*.

Interpreting the main mass points

It is best avoiding reading each mass point as if it identified one model uniquely.

50: non-strategic or very shallow play

- Often interpreted as a salient anchor or a level-0 style choice.
- Useful as a starting point, but not itself a model of strategic reasoning.

33: first step of strategic thinking

- Natural level-1 choice when subjects best respond to an average of 50.
- Can also be generated by NI when deeper precision collapses quickly.

22: second-step style reasoning — but not necessarily literal level-2

- Commonly associated with a second round of reasoning.
- It can also arise from smooth high-level models, especially NI.

0: equilibrium endpoint

- Compatible with very high precision in QRE or NI.
- Observing low guesses does not by itself identify a belief model.

Model comparison

Table 1
Evaluation of the models with up to $K = 4$ components.

K	Orig LevK		Log LevK		NI		QRE
Laboratory							
1	−594.16	<<<	−565.42	<<	−562.51	>>	−566.35
2	−585.47	<<<	−560.91	<<<	−550.05	>>>	−566.34
3	−585.47	<<<	−559.26	<<<	−543	>>>	−566.33
4	−585.47	<<<	−559.18	<<<	−542.65	>>>	−566.34
Classroom							
1	−822.14	<<<	−797.19	<<<	−759.44	=	−759.43
2	−783.06	<<<	−768.49	<<<	−738.18	<	−735.84
3	−782.59	<<<	−759.52	<<<	−723.96	>>>	−735.86
4	−780.72	<<<	−759.09	<<<	−717.99	>>>	−735.82
Takehome							
1	−953.41	<<<	−901.92	<<<	−896.2	>>>	−909.92
2	−926.42	<<<	−873.73	<<<	−858.99	>>>	−887.32
3	−926.42	<<<	−873.73	<<<	−847.29	>>>	−881.21
4	−925.73	<<<	−861.59	<<<	−842.09	>>>	−879.79

Source: Breitmoser (2012)

What to take from the comparison

- Literal intermediate levels do not dominate the data once full-support alternatives are admitted.
- QRE and NI capture the smooth part of the distribution much better.
- The advantage of level-1 is mainly in organizing the focal mass at 33.

Parameter estimates in the NI model

The NI estimates are useful because they tell us *how fast precision declines* along the hierarchy of beliefs.

Table 3
Parameter estimates of the estimated NI models.

ρ_0	Component 1		Component 2		Component 3		Component 4		m	LL
	ρ	λ, μ	ρ	λ, μ	ρ	λ, μ	ρ	λ, μ		
Laboratory										
0.0338	0.11	640.27, 0	0.09	5405.76, 0	0.77	3.8, 0.86			2.4836	−543
	(0.04)	(122.2, 0)	(0)	(0.24, 0)	(0.18)	(0.01, 0)			(2e−04)	
Classroom										
0.0109	0.79	9.16, 0.93	0.06	51.35, 0.95	0.08	1883.56, 0	0.06	968.66, 0.01	3.0435	−717.99
	(0.07)	(0.03, 0)	(0)	0.44, 0	(0)	(200.56, 0)	(0)	0.83, 0	(0)	
Takehome										
0.2269	0.14	9509.16, 0	0.08	852.77, 0.01	0.05	63.6, 0.95	0.51	8.71, 0.94	3.0793	−842.09
	(0.03)	(953.17, –)	(0)	(0.57, 0)	(0)	(1.57, 0.03)	(0.08)	(0.01, 0)	(9e−04)	
Theorists										
0.1716	0.14	12.55, 0.95	0.69	24.28, 0.95					2.0202	−823.49
	(0.08)	(0.02, 0)	(0.08)	(0, 0.02)					(0)	
Newsgroup										
0.1289	0.09	77 883.29, 0	0.05	43.78, 0.84	0.17	60.48, 0.95	0.56	9.26, 0.95	3.579	−820.27
	(0)	(–, 0)	(0)	(0.04, 0)	(0.03)	(1.51, 0.02)	(0.07)	(0.01, 0)	(0.0027)	
Newspaper										
0.0019	0.32	38.99, 0.93	0.08	2118.57, 0	0.57	6.33, 0.95	0.04	2209.58, 0	2.9402	−40 052.94
	(0.01)	(1.19, 0.03)	(0)	(12.5, 0)	(0.01)	(0.01, 0)	(0)	(0.06, 0)	(0)	

Source: Breitmoser (2012)

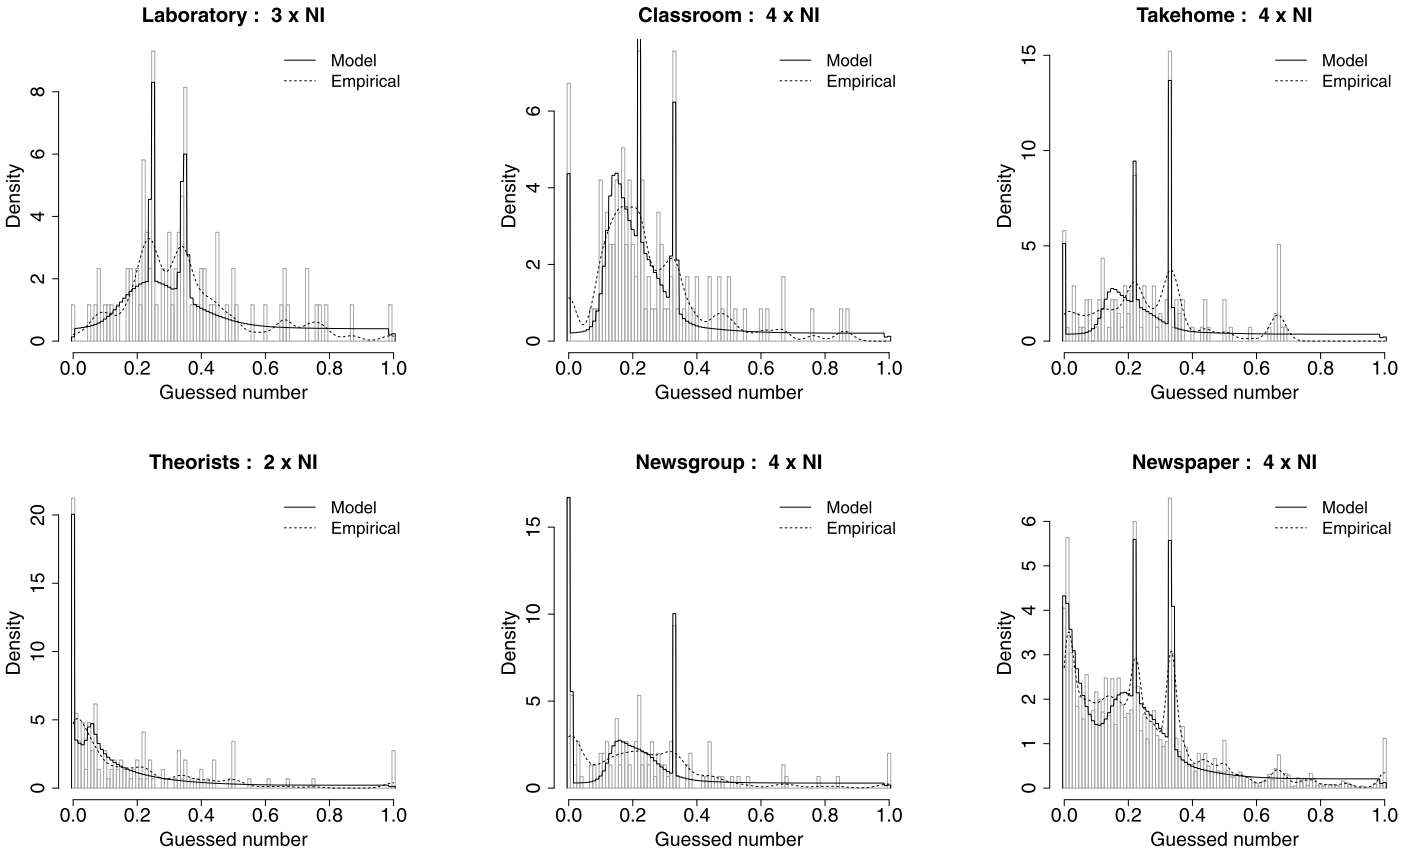
Parameter estimates in the NI model

Interpretation of the NI parameter

- In the estimation, the decay parameter μ captures how persistent precision remains at deeper levels.
- $\mu \approx 0$ corresponds to very rapid collapse of deeper reasoning and hence a level-1-like description.
- μ close to its upper bound corresponds to very slow decay and hence behavior close to logit equilibrium.
- Intermediate values indicate genuine noisy introspection: subjects reason strategically, but increasingly imprecisely at deeper levels.

Main lesson The data do not mostly point to a population of literal level-2 and level-3 types. They point more strongly to *smooth strategic reasoning with heterogeneous precision*.

Overall fit and what we should conclude



Note: The predictions of the models with the "estimated numbers of components" (highest number of components that is not significantly improved upon at $\alpha = 0.1$) are plotted on top of the histograms of the respective data set, and the predictions are aggregated to match the bins of the histograms. The broken line is the kernel estimate labeled "Kernel-0.5" in Fig. 1.

Fig. 7. Predictions of the Noisy Introspection models with the estimated numbers of components.

Source: Breitmoser (2012)

What the evidence supports

The right question is which model organizes the data best without over-interpreting it.

- The smooth bulk of beauty-contest play is better captured by QRE and NI than by literal finite-step types.
- Level-1 reasoning remains important, especially for the persistent mass at 33.
- There is little direct evidence that many subjects are exact intermediate level-k types.

Take-away

Moving from level-k to QRE and NI does not replace the earlier evidence. It explains more of it, with a richer model of beliefs and response.

Where the literature has moved

QRE and Level-k

What changed after recently in the literature?

The early beauty-contest literature established two robust facts:

- observed play is far from the Nash equilibrium at 0,
- but it is also far from random and exhibits clear strategic structure.

What the earlier debate mainly asked

- Are subjects best described by finite steps of reasoning?
- Or is behavior better captured by full-support models such as QRE or NI?

What the more recent literature asks instead

- Is depth of reasoning a stable characteristic across games?
- How do cognitive ability and judgment relate to strategic sophistication?
- Which theories predict behavior *out of treatment*: beyond the specific game?

Toward infinity and beyond

So the literature has moved from:

"Which model fits this game?"

to:

"Which restrictions on beliefs and responses travel across games and subjects?"

This is not the end of the story: it is the starting point for a broader model-comparison perspective.

Recent challenge: depth of reasoning $\neq$ fixed trait

Cooper et al. (2024) study choices across five classes of games designed to identify reasoning depth within and across environments.

Main result

- Depth of reasoning is *pervasively inconsistent* across games.
- The same individual looks deeper in some environments, shallower in others.

Implication

- Level-k remains useful as a *local descriptive approximation*.
- But we should interpret estimated levels as stable personal types.
- This strengthens the case for richer models that separate *belief structure*, *response precision*, and *game-specific complexity*.

Cognitive skill matters — but not the whole story

Gill et al. (2025)

- Higher cognitive ability predicts higher estimated strategic sophistication.
- But choosing a very low number is *not* a pure measure of cognitive skill.
- Strategic behavior reflects a combination of reasoning ability, judgment, and beliefs about others.

Gill and Prowse (2023)

- Strategic complexity varies strongly across situations.
- Response times reveal substantial heterogeneity in how subjects react to that complexity.
- Thinking longer than usual need not improve performance.

Broader lesson: no one-shot theory wins everywhere

Natural question: can one model predict behavior well across many games.

Külpmann and Kuzmics (2022)

- compare several theories using an out-of-treatment likelihood approach,
- calibrate theories on some games and evaluate them on different games,
- and find that many theories have predictive power, but no theory is uniformly best.

What this changes in our interpretation of beauty contests

- The right question is not whether level-k, QRE, or NI is *the* true model.
- Rather, which restrictions are empirically helpful in a given class of games.

Our question is thus

Do the same modeling lessons survive outside the beauty contest?

Summary

QRE and Level-k

What we learned from QRE and noisy introspection

How should we model strategic behavior when observed choices are smoother and more dispersed than literal level-k types suggest?

1. QRE: smooth strategic response with rational expectations
   - Generalizes Nash by replacing mutual best responses with logit responses.
   - Preserves equilibrium: beliefs about others' noisy behavior must be correct.
   - Attractive when subjects react to incentives, but with inaccurate best resp.
2. Noisy introspection: smooth strategic response not full equilibrium
   - Keeps iterated strategic reasoning, but precision decline along beliefs hierar.
   - It is therefore a natural one-shot alternative to QRE in unfamiliar games.
   - Computationally, NI is easier to handle.
3. What the beauty-contest evidence suggests
   - Canonical level-1 reasoning remains visible, especially in the mass at 33.
   - But the smooth bulk distribution better captured by QRE and NI.

What this means for the rest of the course

The broader lesson of this lecture is not that one model has "won."

What we should retain

- **Level-k** is useful for organizing salient focal masses and finite-step reasoning.
- **QRE** is useful when equilibrium discipline and smooth response seem empirically plausible.
- **NI** is useful when subjects appear to reason strategically, but without fully correct beliefs about others' behavior.

Next lecture

From guessing games to broader evidence on strategic behavior: beliefs, actions, and heterogeneity across normal-form games.

This broader comparative perspective is also in line with the recent literature emphasizing cross-game instability of fixed depth and the value of out-of-treatment model comparison (Cooper et al., 2024; Külpmann and Kuzmics, 2022).

References

QRE and Level-k

Guide to the literature

This lecture builds on four layers of literature.

1. Beauty-contest evidence and strategic reasoning
   - Nagel (1995): the classic experimental beauty-contest paper
   - Bosch-Domènech et al. (2002): field and classroom evidence on guessing games
   - Breitmoser (2012): structural comparison of level-k, QRE, and noisy introspection
2. Smooth strategic response and equilibrium
   - McKelvey and Palfrey (1995): definition and existence of logit QRE
   - Turocy (2005): principal branch and homotopy interpretation
   - Haile et al. (2008): limits of the empirical content of QRE
3. One-shot reasoning beyond rational expectations
   - Goeree and Holt (2004): noisy introspection as a one-shot alternative to QRE
   - Goeree et al. (2018): later application of NI to another reasoning game

References I

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.

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.

Gill, D., Knepper, Z., Prowse, V. L., and Zhou, J. (2025). How cognitive skills affect strategic behavior: Cognitive ability, fluid intelligence and judgment. *Games and Economic Behavior*, 149:82–95.

Gill, D. and Prowse, V. L. (2023). Strategic complexity and the value of thinking. *The Economic Journal*, 133(650):761–786.

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

Goeree, J. K., Louis, P., and Zhang, J. (2018). Noisy introspection in the 11–20 game. *The Economic Journal*, 128(611):1509–1530.

References II

Haile, P. A., Hortaçsu, A., and Kosenok, G. (2008). On the empirical content of quantal response equilibrium. *American Economic Review*, 98(1):180–200.

Ho, T.-H., Camerer, C., and Weigelt, K. (1998). Iterated dominance and iterated best response in experimental "p-beauty contests.". *American Economic Review*, 88(4):947–969.

Külpmann, P. and Kuzmics, C. (2022). Comparing theories of one-shot play out of treatment. *Journal of Economic Theory*, 205:105554.

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.

Turocy, T. L. (2005). A dynamic homotopy interpretation of the logistic quantal response equilibrium correspondence. *Games and Economic Behavior*, 51(2):243–263.

Repetition

QRE and Level-k

Repetition I: Why do we need models beyond level-k?

1. Why is level-k a useful starting point in beauty-contest games?
Answer sketch: It explains why many subjects choose salient intermediate values such as 33 or 22 instead of the Nash equilibrium 0.

2. Why is literal finite level-k not enough as a complete description of the data?
Answer sketch: Because the empirical distribution is often smoother and more dispersed than a small set of exact reasoning types would predict.

3. Once choices are noisy, what is the real modeling problem?
Answer sketch: Not only own response noise, but also how beliefs about others' precision and higher-order beliefs are disciplined.

4. What are the three main ways discussed in this lecture to discipline strategic reasoning?
Answer sketch: Finite depth (level-k), rational-expectations smoothing (QRE), and fading-precision hierarchies (noisy introspection).

5. What is the main difference between lecture 03 and lecture 04?
Answer sketch: Lecture 03 introduces finite-step reasoning; lecture 04 asks how to model the smooth bulk of behavior beyond literal level-k types.

Repetition II: Quantal response equilibrium

1. What is a logit QRE?
Answer sketch: A mixed-strategy profile in which each player chooses actions with probabilities increasing in expected payoffs, and beliefs about others' noisy play are correct (McKelvey and Palfrey, 1995).

2. What does the precision parameter λ measure?
Answer sketch: How strongly choice probabilities react to expected-payoff differences.

3. Why is QRE a fixed-point concept?
Answer sketch: Because expected payoffs depend on beliefs about others' behavior, while those beliefs must coincide with the actual stochastic behavior in equilibrium.

4. What happens at $\lambda = 0$ and as $\lambda \to \infty$?
Answer sketch: At $\lambda = 0$, play is uniform randomization; as λ becomes large, behavior approaches best response and may converge toward Nash along suitable branches.

5. Why is QRE computationally harder than finite level-k?
Answer sketch: Because QRE requires solving a fixed point, whereas finite level-k can be computed recursively from a lowest level.

Repetition III: Noisy introspection

1. What is the core idea of noisy introspection?
Answer sketch: Players reason strategically through many layers, but precision declines with depth in the hierarchy of beliefs (Goeree and Holt, 2004).

2. What is the role of the decay parameter $\kappa \in (0, 1)$?
Answer sketch: It determines how quickly precision fades at deeper levels of introspection.

3. How does NI differ from level-k?
Answer sketch: Level-k stops at finite depth; NI continues indefinitely, but with increasingly noisy deeper reasoning.

4. How does NI differ from QRE?
Answer sketch: NI does not require rational expectations; my belief about your play need not equal your actual play.

5. Why is NI computationally easier than QRE?
Answer sketch: Because once deeper precision becomes negligible, the hierarchy can be truncated and solved by backward induction from near-randomization.

6. When should NI be especially attractive as a model?
Answer sketch: In unfamiliar one-shot games where subjects seem to reason strategically, but fully correct beliefs about others are implausible.

Repetition IV: What do the beauty-contest data tell us?

1. Why is the beauty contest a useful environment for comparing models of strategic reasoning?
Answer sketch: Because the Nash prediction is sharp, the salient level-k benchmarks are clear, and deviations are highly informative (Nagel, 1995; Bosch-Domènech et al., 2002).

2. What does the mass at 33 usually suggest?
Answer sketch: A first step of strategic reasoning, often interpreted as level-1 best response to an average of 50.

3. Why should we not read a choice like 22 or 0 as uniquely identifying one model?
Answer sketch: Because the same observed choice can arise from different belief structures, including finite-step reasoning and full-support models.

4. What is the main empirical message of Breitmoser (2012)?
Answer sketch: Much of the smooth bulk of beauty-contest play is better captured by high-level full-support models such as QRE and NI, while level-1 reasoning remains important for focal masses such as 33.

5. Does lecture 04 overturn the message of lecture 03?

Repetition V: Broader interpretation and recent literature

1. What does recent evidence imply about interpreting level-k estimates as stable personal traits?
Answer sketch: That is too strong: reasoning depth often varies across games and environments (Cooper et al., 2024).

2. Why is choosing a very low number not a pure measure of strategic sophistication?
Answer sketch: Because behavior depends not only on cognitive ability, but also on judgment and beliefs about others (Gill et al., 2025).

3. What broader methodological lesson should we take from this lecture?
Answer sketch: We should compare models by the restrictions they impose on beliefs and responses, not just ask which named model "wins."

4. Why is out-of-treatment comparison valuable?
Answer sketch: Because a model that fits one game well may not predict well in other games; predictive performance should travel across environments (Külpmann and Kuzmics, 2022).

5. What is the main take-away from lecture 04 in one sentence?
Answer sketch: Beauty contest data combine focal low-level reasoning with

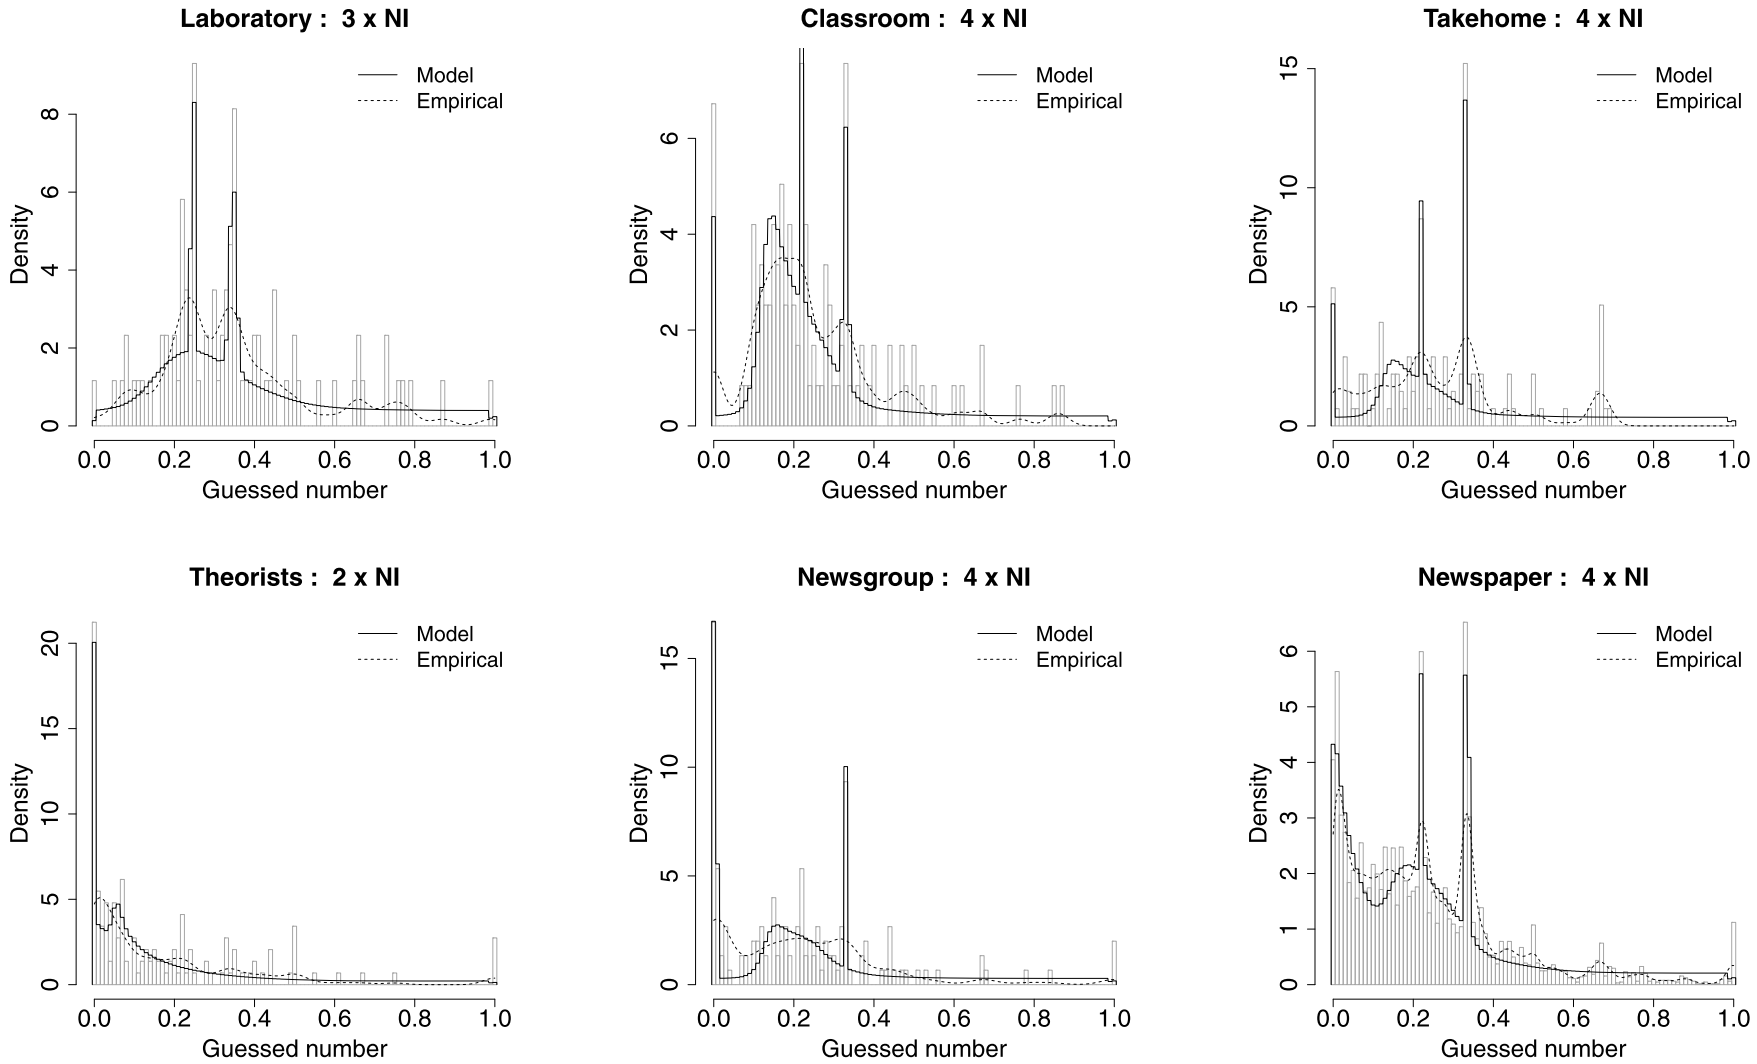
Definition (Logit Equilibrium (McKelvey and Palfrey, 1995))

Strategy profile σ is a **logit equilibrium** (special case of QRE) if players play mutual logit responses for some $\lambda \in_+$, i.e. there exists $\lambda \in_+$ such that for all $i \in N$,

$$\sigma_i(a_i) = \frac{\exp\{\lambda \cdot \pi_i(a_i, \sigma_{-i})\}}{\sum_{a_i'} \exp\{\lambda \cdot \pi_i(a_i', \sigma_{-i})\}}.$$

- Logit equilibrium = mutual logit response
- $\lambda \to \infty$ implies convergence to mutual best response (Nash)
- For all $\lambda < \infty$: All actions have positive probability
- principal branch of logit equilibria is a closed path connecting $\lambda = 0$ and $\lambda \to \infty$, converging to a NE
- In 2×2 games, the principal branch converges to the risk dominant Nash equilibrium (Turocy, 2005)
- Existence | Computation | Correspondence | Multiple | Not Perfect
- Guessing games Majority of participants (70%) fits logit equilibrium

Overall fit is okay



Note: The predictions of the models with the "estimated numbers of components" (highest number of components that is not significantly improved upon at $\alpha = 0.1$) are plotted on top of the histograms of the respective data set, and the predictions are aggregated to match the bins of the histograms. The broken line is the kernel estimate labeled "Kernel-0.5" in Fig. 1.

Fig. 7. Predictions of the Noisy Introspection models with the estimated numbers of components.

Noisy Introspection (Goeree and Holt, 2004)

- Agents play logit response σ with precision λ to belief σ^1
- Belief σ^1 is a logit response with precision $\kappa\lambda$ to belief σ^2
- Belief σ^2 is a logit response with precision $\kappa^2\lambda$ to belief σ^3
- Belief σ^3 is a logit response with precision $\kappa^3\lambda$ to belief σ^4
- ...

- Logit response to logit response to logit response ...
- With λ decreasing by constant factor κ
- Useful, as it nests many other concepts as special cases
  - Logit Eq $\lambda < \infty$, $\kappa \to 1$
  - Nash Eq $\lambda \to \infty$, $\kappa \to 1$
  - Level 0 $\lambda = 0$
  - Level 1 $\lambda > 0$, $\kappa = 0$
- Computation