Generalized concepts and evidence

Summer 26

Niels Boissonnet, Bielefeld

Discussion

From last lecture to today's question

Last lecture used Asymmetric Logit Equilibrium (ALE) to connect two recurring patterns in one-shot normal-form games:

From last lecture to today's question

Last lecture used Asymmetric Logit Equilibrium (ALE) to connect two recurring patterns in one-shot normal-form games:

- simple strategic reasoning, often close to level-1 behavior;

From last lecture to today's question

Last lecture used Asymmetric Logit Equilibrium (ALE) to connect two recurring patterns in one-shot normal-form games:

- simple strategic reasoning, often close to level-1 behavior;
- smooth payoff-sensitive response, often closer to QRE-type behavior.

This continued the strategic-reasoning line from beauty-contest evidence to richer normal-form environments (Breitmoser, 2012; Camerer et al., 2004).

Question for today

If simple level-k and standard QRE are both too restrictive, how should we model richer forms of boundedly rational one-shot play?

Why we need more than level-k or standard QRE

The benchmark models are powerful, but each leaves something out:

Why we need more than level-k or standard QRE

The benchmark models are powerful, but each leaves something out:

- Simple level-k is intuitive, but often too sharp

Why we need more than level-k or standard QRE

The benchmark models are powerful, but each leaves something out:

- Simple level-k is intuitive, but often too sharp
- Standard QRE is smooth, but often too homogeneous.

Why we need more than level-k or standard QRE

The benchmark models are powerful, but each leaves something out:

- Simple level-k is intuitive, but often too sharp
- Standard QRE is smooth, but often too homogeneous.
- Experiments suggest that bounded rationality also reflects
  - heterogeneous strategic depth
  - regret and foregone payoffs,
  - simple local procedures.

So today's lecture is not about replacing one benchmark with one new champion.

We want to identify different mechanisms that generate similar aggregate deviations from NE.

Four routes beyond the benchmark

We study four different ways to enrich one-shot behavioral game theory:

Four routes beyond the benchmark

We study four different ways to enrich one-shot behavioral game theory:

1. Cognitive Hierarchy Model (CHM): bounded rationality through *heterogeneous depths of reasoning*;

Four routes beyond the benchmark

We study four different ways to enrich one-shot behavioral game theory:

1. Cognitive Hierarchy Model (CHM): bounded rationality through *heterogeneous depths of reasoning*;
2. Impulse Balance Equilibrium (IBE): bounded rationality through *regret and foregone payoffs*;

Four routes beyond the benchmark

We study four different ways to enrich one-shot behavioral game theory:

1. Cognitive Hierarchy Model (CHM): bounded rationality through *heterogeneous depths of reasoning*;
2. Impulse Balance Equilibrium (IBE): bounded rationality through *regret and foregone payoffs*;
3. Payoff sampling / action sampling: bounded rationality through *finite procedural sampling*;

Four routes beyond the benchmark

We study four different ways to enrich one-shot behavioral game theory:

1. Cognitive Hierarchy Model (CHM): bounded rationality through *heterogeneous depths of reasoning*;
2. Impulse Balance Equilibrium (IBE): bounded rationality through *regret and foregone payoffs*;
3. Payoff sampling / action sampling: bounded rationality through *finite procedural sampling*;
4. HQRE / TQRE: bounded rationality through *heterogeneous noisy equilibrium response*.

These approaches are not isolated islands. They differ in what is noisy, what is heterogeneous, and what kind of evidence can discriminate them.

How we will compare these models

Throughout the lecture, we repeatedly ask four questions:

How we will compare these models

Throughout the lecture, we repeatedly ask four questions:

1. Are beliefs explicit or only implicit?

How we will compare these models

Throughout the lecture, we repeatedly ask four questions:

1. Are beliefs explicit or only implicit?
2. Does randomness come from response noise or from finite sampling / local procedures?

How we will compare these models

Throughout the lecture, we repeatedly ask four questions:

1. Are beliefs explicit or only implicit?
2. Does randomness come from response noise or from finite sampling / local procedures?
3. Is heterogeneity built into the model, or are all players treated alike?

How we will compare these models

Throughout the lecture, we repeatedly ask four questions:

1. Are beliefs explicit or only implicit?
2. Does randomness come from response noise or from finite sampling / local procedures?
3. Is heterogeneity built into the model, or are all players treated alike?
4. Which experimental environments are informative enough to distinguish the models?

Our recurring empirical benchmark will therefore be stationary play in simple normal-form games, especially in environments where different concepts make clearly different predictions (Crawford et al., 2013).

A useful warning from the recent literature

Before comparing the models in detail, it is useful to keep in mind:

A useful warning from the recent literature

Before comparing the models in detail, it is useful to keep in mind:

- Recent out-of-treatment comparisons do *not* identify one behavioral model that dominates all others across games.

A useful warning from the recent literature

Before comparing the models in detail, it is useful to keep in mind:

- Recent out-of-treatment comparisons do *not* identify one behavioral model that dominates all others across games.
- Different models capture different regularities in the data.

A useful warning from the recent literature

Before comparing the models in detail, it is useful to keep in mind:

- Recent out-of-treatment comparisons do *not* identify one behavioral model that dominates all others across games.
- Different models capture different regularities in the data.
- Recent synthesis work therefore suggests that model diversity is informative: combinations of behavioral models can predict better than a single universal winner.

This is exactly why we study several generalized concepts side by side (Külpmann and Kuzmics, 2022; Huang and Golman, 2025).

Learning goals and route through the lecture

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

Learning goals and route through the lecture

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

- explain what CHM adds relative to parsimonious level-k;

Learning goals and route through the lecture

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

- explain what CHM adds relative to parsimonious level-k;
- explain the intuition behind Impulse Balance Equilibrium;

Learning goals and route through the lecture

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

- explain what CHM adds relative to parsimonious level-k;
- explain the intuition behind Impulse Balance Equilibrium;
- distinguish payoff sampling from action sampling;

Learning goals and route through the lecture

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

- explain what CHM adds relative to parsimonious level-k;
- explain the intuition behind Impulse Balance Equilibrium;
- distinguish payoff sampling from action sampling;
- explain what HQRE / TQRE add to standard QRE;

Learning goals and route through the lecture

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

- explain what CHM adds relative to parsimonious level-k;
- explain the intuition behind Impulse Balance Equilibrium;
- distinguish payoff sampling from action sampling;
- explain what HQRE / TQRE add to standard QRE;
- compare these concepts in terms of beliefs, heterogeneity, noise, and evidence.

Roadmap: CHM $\rightarrow$ IBE $\rightarrow$ Sampling $\rightarrow$ HQRE/TQRE $\rightarrow$ Synthesis

We start with the benchmark model of heterogeneous strategic depth.

Cognitive Hierarchy Model

Generalized concepts and evidence

Why start with CHM?

The natural next step after ALE is to allow for heterogeneous strategic depth without committing to one literal finite level for each observation.

Why start with CHM?

The natural next step after ALE is to allow for heterogeneous strategic depth without committing to one literal finite level for each observation.

- Level-k is often too sharp: a level-k player best responds only to level $k-1$.

Why start with CHM?

The natural next step after ALE is to allow for heterogeneous strategic depth without committing to one literal finite level for each observation.

- Level-k is often too sharp: a level-k player best responds only to level $k-1$.
- CHM smooths this by letting a level-k player best respond to a *distribution* of lower levels.

Why start with CHM?

The natural next step after ALE is to allow for heterogeneous strategic depth without committing to one literal finite level for each observation.

- Level-k is often too sharp: a level-k player best responds only to level $k-1$.
- CHM smooths this by letting a level-k player best respond to a *distribution* of lower levels.
- This keeps the logic of iterated reasoning, but gives a more flexible aggregate description of play in one-shot games (Stahl and Wilson, 1995; Camerer et al., 2004; Crawford et al., 2013).

Cognitive Hierarchy Model (Camerer et al., 2004)

Cognitive Hierarchy Model (CHM)

Let K be a random variable on $\mathbb{N}_0$ and let $\lambda \in \mathbb{R}_+$. A mixed-strategy profile $(\sigma_i)_{i \in N}$ is a CH strategy profile if there exist $(\sigma_{i,0}, \sigma_{i,1}, \sigma_{i,2}, \ldots)$ for all $i \in N$ such that:

Cognitive Hierarchy Model (Camerer et al., 2004)

Cognitive Hierarchy Model (CHM)

Let K be a random variable on $\mathbb{N}_0$ and let $\lambda \in \mathbb{R}_+$. A mixed-strategy profile $(\sigma_i)_{i \in N}$ is a CH strategy profile if there exist $(\sigma_{i,0}, \sigma_{i,1}, \sigma_{i,2}, \ldots)$ for all $i \in N$ such that:

- $\sigma_i = \sum_{k \geq 0} \Pr(K = k) \cdot \sigma_{i,k}$

Cognitive Hierarchy Model (Camerer et al., 2004)

Cognitive Hierarchy Model (CHM)

Let K be a random variable on $\mathbb{N}_0$ and let $\lambda \in \mathbb{R}_+$. A mixed-strategy profile $(\sigma_i)_{i \in N}$ is a CH strategy profile if there exist $(\sigma_{i,0}, \sigma_{i,1}, \sigma_{i,2}, \ldots)$ for all $i \in N$ such that:

- $\sigma_i = \sum_{k \geq 0} \Pr(K = k) \cdot \sigma_{i,k}$
- $\sigma_{i,0}$ is the uniform distribution on A_i

Cognitive Hierarchy Model (Camerer et al., 2004)

Cognitive Hierarchy Model (CHM)

Let K be a random variable on $\mathbb{N}_0$ and let $\lambda \in \mathbb{R}_+$. A mixed-strategy profile $(\sigma_i)_{i \in N}$ is a CH strategy profile if there exist $(\sigma_{i,0}, \sigma_{i,1}, \sigma_{i,2}, \ldots)$ for all $i \in N$ such that:

- $\sigma_i = \sum_{k \geq 0} \Pr(K = k) \cdot \sigma_{i,k}$
- $\sigma_{i,0}$ is the uniform distribution on A_i
- for each $k \geq 1$, $\sigma_{i,k}$ is a logit response (precision λ) to the belief $\mu_{i,-i}(k)$

Cognitive Hierarchy Model (Camerer et al., 2004)

Cognitive Hierarchy Model (CHM)

Let K be a random variable on $\mathbb{N}_0$ and let $\lambda \in \mathbb{R}_+$. A mixed-strategy profile $(\sigma_i)_{i \in N}$ is a CH strategy profile if there exist $(\sigma_{i,0}, \sigma_{i,1}, \sigma_{i,2}, \ldots)$ for all $i \in N$ such that:

- $\sigma_i = \sum_{k \geq 0} \Pr(K = k) \cdot \sigma_{i,k}$
- $\sigma_{i,0}$ is the uniform distribution on A_i
- for each $k \geq 1$, $\sigma_{i,k}$ is a logit response (precision λ) to the belief $\mu_{i,-i}(k)$
- for each opponent $j \neq i$,

$$\mu_{i,j}(k) = \sum_{k' < k} \Pr(K = k' \mid K < k) \cdot \sigma_{j,k'}$$

Standard empirical closure: K is Poisson with mean τ.

What CHM is really assuming

CHM is best read as a parsimonious structural model of strategic heterogeneity, not as a literal psychological taxonomy.

What CHM is really assuming

CHM is best read as a parsimonious structural model of strategic heterogeneity, not as a literal psychological taxonomy.

- Level-0 is a modeling anchor: usually uniform randomization.

What CHM is really assuming

CHM is best read as a parsimonious structural model of strategic heterogeneity, not as a literal psychological taxonomy.

- Level-0 is a modeling anchor: usually uniform randomization.
- Level-k reasoning is bounded: players respond to lower levels only, not to an infinite hierarchy.

What CHM is really assuming

CHM is best read as a parsimonious structural model of strategic heterogeneity, not as a literal psychological taxonomy.

- Level-0 is a modeling anchor: usually uniform randomization.
- Level-k reasoning is bounded: players respond to lower levels only, not to an infinite hierarchy.
- Excluding own level captures the idea that each player sees herself as more sophisticated than the opponents she faces.

What CHM is really assuming

CHM is best read as a parsimonious structural model of strategic heterogeneity, not as a literal psychological taxonomy.

- Level-0 is a modeling anchor: usually uniform randomization.
- Level-k reasoning is bounded: players respond to lower levels only, not to an infinite hierarchy.
- Excluding own level captures the idea that each player sees herself as more sophisticated than the opponents she faces.
- Poisson heterogeneity is mainly a tractable way to summarize how much reasoning depth there is in the population.

Key interpretation

CHM does *not* say that every subject literally carries a fixed integer label. It says that aggregate play can often be organized well by a distribution of boundedly rational depths.

CHM versus parsimonious level-k

Parsimonious level-k

CHM versus parsimonious level-k

Parsimonious level-k

- best responds to level $k-1$

CHM versus parsimonious level-k

Parsimonious level-k

- best responds to level $k-1$
- very sharp belief restriction

CHM versus parsimonious level-k

Parsimonious level-k

- best responds to level $k-1$
- very sharp belief restriction
- easy to interpret in toy examples

CHM versus parsimonious level-k

Parsimonious level-k

- best responds to level $k-1$
- very sharp belief restriction
- easy to interpret in toy examples
- may cycle across levels if the game is not dominance solvable

CHM

CHM versus parsimonious level-k

Parsimonious level-k

- best responds to level $k-1$
- very sharp belief restriction
- easy to interpret in toy examples
- may cycle across levels if the game is not dominance solvable

CHM

- best responds to all lower levels

CHM versus parsimonious level-k

Parsimonious level-k

- best responds to level $k-1$
- very sharp belief restriction
- easy to interpret in toy examples
- may cycle across levels if the game is not dominance solvable

CHM

- best responds to all lower levels
- smoother aggregate predictions

CHM versus parsimonious level-k

Parsimonious level-k

- best responds to level $k-1$
- very sharp belief restriction
- easy to interpret in toy examples
- may cycle across levels if the game is not dominance solvable

CHM

- best responds to all lower levels
- smoother aggregate predictions
- natural population interpretation

CHM versus parsimonious level-k

Parsimonious level-k

- best responds to level $k-1$
- very sharp belief restriction
- easy to interpret in toy examples
- may cycle across levels if the game is not dominance solvable

CHM

- best responds to all lower levels
- smoother aggregate predictions
- natural population interpretation
- converges more cleanly as depth increases

Empirically, the differences are often modest when estimated depths are low (Camerer et al., 2004; Crawford et al., 2013).

How Camerer et al. (2004) estimate CHM

Their baseline implementation is deliberately parsimonious:

How Camerer et al. (2004) estimate CHM

Their baseline implementation is deliberately parsimonious:

- Level distribution: Poisson, with mean τ

$$\Pr(K = k) = e^{-\tau} \frac{\tau^k}{k!}$$

How Camerer et al. (2004) estimate CHM

Their baseline implementation is deliberately parsimonious:

- Level distribution: Poisson, with mean τ

$$\Pr(K = k) = e^{-\tau} \frac{\tau^k}{k!}$$

- Choice rule: in the benchmark specification, players best respond to their level-specific beliefs, i.e. effectively $\lambda = \infty$.

How Camerer et al. (2004) estimate CHM

Their baseline implementation is deliberately parsimonious:

- Level distribution: Poisson, with mean τ

$$\Pr(K = k) = e^{-\tau} \frac{\tau^k}{k!}$$

- Choice rule: in the benchmark specification, players best respond to their level-specific beliefs, i.e. effectively $\lambda = \infty$.
- Estimation: maximum likelihood and least squares, plus cross-validation.

How Camerer et al. (2004) estimate CHM

Their baseline implementation is deliberately parsimonious:

- Level distribution: Poisson, with mean τ

$$\Pr(K = k) = e^{-\tau} \frac{\tau^k}{k!}$$

- Choice rule: in the benchmark specification, players best respond to their level-specific beliefs, i.e. effectively $\lambda = \infty$.
- Estimation: maximum likelihood and least squares, plus cross-validation.
- Why estimation is feasible: since level-0 has positive probability, every action obtains positive overall probability.

What to look for

Does one common distribution of depths organize behavior well across games, and how much does the implied mean depth vary across environments?

Evidence: estimated depth varies across games

TABLE III
PARAMETER ESTIMATE τ FOR COGNITIVE HIERARCHY MODELS

Data set	Stahl and Wilson	Cooper and Van Huyck	Costa-Gomes et al.	Mixed	Entry
Game-specific τ					
Game 1	2.93	15.90	2.28	0.98	0.70
Game 2	0.00	1.07	2.27	1.71	0.85
Game 3	1.40	0.18	2.29	0.86	—
Game 4	2.34	1.28	1.26	3.85	0.73
Game 5	2.01	0.52	1.80	1.08	0.70
Game 6	0.00	0.82	1.67	1.13	
Game 7	5.37	0.96	0.88	3.29	
Game 8	0.00	1.54	2.18	1.84	
Game 9	1.35		1.89	1.06	
Game 10	11.33		2.26	2.26	
Game 11	6.48		1.23	0.87	
Game 12	1.71		1.03	2.06	
Game 13			2.28	1.88	
Game 14				9.07	
Game 15				3.49	
Game 16				2.07	
Game 17				1.14	
Game 18				1.14	
Game 19				1.55	
Game 20				1.95	
Game 21				1.68	
Game 22				3.06	
Median τ	1.86	1.01	1.89	1.77	0.71
Common τ	1.54	0.82	1.73	1.48	0.73

Takeaway

The estimated mean depth τ varies substantially across games. CHM organizes behavior, but it does not suggest one universal sophistication level (Camerer et al., 2004).

Economic value of CHM

Economic value: difference between the payoff from the strategy a theory recommends and the average payoff subjects actually achieve.

Economic value of CHM

Economic value: difference between the payoff from the strategy a theory recommends and the average payoff subjects actually achieve.

- The model is fitted to $n - 1$ data sets and evaluated on the remaining one.

Economic value of CHM

Economic value: difference between the payoff from the strategy a theory recommends and the average payoff subjects actually achieve.

- The model is fitted to $n - 1$ data sets and evaluated on the remaining one.
- This is stricter than in-sample fit: it asks whether the model travels.

Economic value of CHM

Economic value: difference between the payoff from the strategy a theory recommends and the average payoff subjects actually achieve.

- The model is fitted to $n-1$ data sets and evaluated on the remaining one.
- This is stricter than in-sample fit: it asks whether the model travels.
- Clairvoyance is the benchmark value of knowing the data itself.

TABLE VIII
ECONOMIC VALUE OF VARIOUS THEORIES

Data set	Stahl and Wilson	Cooper and Van Huyck	Costa-Gomes et al.	Mixed	Entry
Observed payoff	195	586	264	328	118
Clairvoyance payoff	243	664	306	708	176
Economic value					
Clairvoyance	48	78	42	380	58
Cognitive hierarchy (Common τ)	13	55	22	132	10
Nash equilibrium	5	30	15	−17	2
% Maximum economic value achieved					
Cognitive hierarchy (Common τ)	26%	71%	52%	35%	17%
Nash equilibrium	10%	39%	35%	−4%	3%

The economic value is the total value (in experimental payoffs) of all rounds that a "hypothetical" subject will earn using the respective model to predict other's behavior and best responds with the strategy that yields the highest expected payoff in each round.

Modern question: is reasoning depth a stable type?

A natural temptation is to read CHM literally: estimate a subject's "depth" once and carry it across games. More recent evidence suggests caution.

Modern question: is reasoning depth a stable type?

A natural temptation is to read CHM literally: estimate a subject's "depth" once and carry it across games. More recent evidence suggests caution.

- Hyndman et al. (2022) find fairly high stability of beliefs and actions within the same strategically equivalent game, but much lower stability across different games.

Modern question: is reasoning depth a stable type?

A natural temptation is to read CHM literally: estimate a subject's "depth" once and carry it across games. More recent evidence suggests caution.

- Hyndman et al. (2022) find fairly high stability of beliefs and actions within the same strategically equivalent game, but much lower stability across different games.
- Cooper et al. (2024) show that a fixed depth of reasoning is often rejected even across closely related level-k environments.

Implication

CHM is very useful as a model of population heterogeneity in a given environment, but less compelling if interpreted as an immutable subject trait across all games.

Section takeaway: what CHM gives us

Section takeaway: what CHM gives us

- CHM is the canonical model of heterogeneous depth of reasoning.

Section takeaway: what CHM gives us

- CHM is the canonical model of heterogeneous depth of reasoning.
- Relative to parsimonious level-k, it smooths the belief side by letting players respond to a distribution of lower levels.

Section takeaway: what CHM gives us

- CHM is the canonical model of heterogeneous depth of reasoning.
- Relative to parsimonious level-k, it smooths the belief side by letting players respond to a distribution of lower levels.
- Empirically, it often captures one-shot play well and improves on Nash in many environments (Camerer et al., 2004; Crawford et al., 2013).

Section takeaway: what CHM gives us

- CHM is the canonical model of heterogeneous depth of reasoning.
- Relative to parsimonious level-k, it smooths the belief side by letting players respond to a distribution of lower levels.
- Empirically, it often captures one-shot play well and improves on Nash in many environments (Camerer et al., 2004; Crawford et al., 2013).
- But estimated depth varies across games, and recent evidence warns against taking "depth" too literally as a stable characteristic (Hyndman et al., 2022; Cooper et al., 2024).

Transition

CHM is still a belief-based account of bounded rationality. We now turn to a concept built around regret and foregone payoffs.

Impulse Balance Equilibrium

Generalized concepts and evidence

Why move beyond belief-based models?

So far, most of our models of one-shot play have been belief-based: level-k, CHM, ALE, and QRE describe behavior as responses to beliefs about others.

But bounded rationality may also work differently:

Why move beyond belief-based models?

So far, most of our models of one-shot play have been belief-based: level-k, CHM, ALE, and QRE describe behavior as responses to beliefs about others.

But bounded rationality may also work differently:

- subjects may react to foregone payoffs,

Why move beyond belief-based models?

So far, most of our models of one-shot play have been belief-based: level-k, CHM, ALE, and QRE describe behavior as responses to beliefs about others.

But bounded rationality may also work differently:

- subjects may react to foregone payoffs,
- they may learn from regret rather than from articulated beliefs,

Why move beyond belief-based models?

So far, most of our models of one-shot play have been belief-based: level-k, CHM, ALE, and QRE describe behavior as responses to beliefs about others.

But bounded rationality may also work differently:

- subjects may react to foregone payoffs,
- they may learn from regret rather than from articulated beliefs,
- and long-run frequencies may then reflect balanced impulses rather than one-shot best responses.

Alternative foundation

Impulse Balance Equilibrium (IBE) models bounded rationality via regret, loss aversion, and stationary adjustment (Selten and Chmura, 2008).

Impulse balance: intuition and notation

We focus on a 2×2 game. For each player $i \in \{1, 2\}$, let a_i and a_i' denote the two actions.

Impulse balance: intuition and notation

We focus on a 2×2 game. For each player $i \in \{1, 2\}$, let a_i and a_i' denote the two actions.

- The starting point is a reference point, taken to be the pure-strategy maximin payoff

$$s_i = \max_{a_i} \min_{a_j} \pi_i(a_i, a_j).$$

Impulse balance: intuition and notation

We focus on a 2×2 game. For each player $i \in \{1, 2\}$, let a_i and a_i' denote the two actions.

- The starting point is a reference point, taken to be the pure-strategy maximin payoff

$$s_i = \max_{a_i} \min_{a_j} \pi_i(a_i, a_j).$$

- Outcomes above the reference point are treated more mildly than outcomes below it: gains are worth only half as much as losses.

Impulse balance: intuition and notation

We focus on a 2×2 game. For each player $i \in \{1, 2\}$, let a_i and a_i' denote the two actions.

- The starting point is a reference point, taken to be the pure-strategy maximin payoff

$$s_i = \max_{a_i} \min_{a_j} \pi_i(a_i, a_j).$$

- Outcomes above the reference point are treated more mildly than outcomes below it: gains are worth only half as much as losses.
- If the unchosen action would have yielded a better transformed payoff, the player experiences regret.

Core idea

Regret creates an impulse to switch. In the long run, play settles where the average impulses toward the two actions are balanced.

From transformed payoffs to regret

Relative to the maximin reference point s_i, define

$$u_i(a_i, a_j) = \begin{cases} s_i + \big(\pi_i(a_i, a_j) - s_i\big)/2, & \text{if } \pi_i(a_i, a_j) \geq s_i, \\ \pi_i(a_i, a_j), & \text{if } \pi_i(a_i, a_j) < s_i. \end{cases}$$

Regret after choosing a_i against a_j is

$$r_i(a_i, a_j) = \max_{a_i'} \big(u_i(a_i', a_j) - u_i(a_i, a_j)\big)$$

From transformed payoffs to regret

Relative to the maximin reference point s_i, define

$$u_i(a_i, a_j) = \begin{cases} s_i + (\pi_i(a_i, a_j) - s_i)/2, & \text{if } \pi_i(a_i, a_j) \geq s_i, \\ \pi_i(a_i, a_j), & \text{if } \pi_i(a_i, a_j) < s_i. \end{cases}$$

Regret after choosing a_i against a_j is

$$r_i(a_i, a_j) = \max_{a_i'} \left(u_i(a_i', a_j) - u_i(a_i, a_j) \right)$$

- Regret is directional: it points from the chosen action to the better forgone action.

From transformed payoffs to regret

Relative to the maximin reference point s_i, define

$$u_i(a_i, a_j) = \begin{cases} s_i + (\pi_i(a_i, a_j) - s_i)/2, & \text{if } \pi_i(a_i, a_j) \geq s_i, \\ \pi_i(a_i, a_j), & \text{if } \pi_i(a_i, a_j) < s_i. \end{cases}$$

Regret after choosing a_i against a_j is

$$r_i(a_i, a_j) = \max_{a_i'} \left(u_i(a_i', a_j) - u_i(a_i, a_j) \right)$$

- Regret is directional: it points from the chosen action to the better forgone action.

$\Longrightarrow$ IBE is closer to direction learning than to best-response reasoning.

Impulse Balance Equilibrium

For a mixed strategy σ_j of the opponent, anticipated regret from action a_i is

$$r_i(a_i, \sigma_j) = \sigma_j(a_j)\, r_i(a_i, a_j) + \sigma_j(a_j')\, r_i(a_i, a_j').$$

Impulse Balance Equilibrium (IBE)

A mixed-strategy profile σ is an impulse balance equilibrium if, for each player i,

$$\sigma_i(a_i)\, r_i(a_i, \sigma_j) = \sigma_i(a_i')\, r_i(a_i', \sigma_j).$$

The more often an action is played, the more often its associated regret is experienced. In equilibrium, the experienced impulse mass toward the two actions is the same (Selten and Chmura, 2008).

Why IBE is not Nash in transformed utilities

Although IBE uses transformed payoffs, it is not just Nash equilibrium of the transformed game.

Why IBE is not Nash in transformed utilities

Although IBE uses transformed payoffs, it is not just Nash equilibrium of the transformed game.

- If players maximized expected transformed utility directly

  $\Longrightarrow$ i.e. minimizing regret r_i

  indifference would require $r_i(a_i, \sigma_j) = r_i(a_i', \sigma_j)$.

Why IBE is not Nash in transformed utilities

Although IBE uses transformed payoffs, it is not just Nash equilibrium of the transformed game.

- If players maximized expected transformed utility directly

  $\Longrightarrow$ i.e. minimizing regret r_i

  indifference would require $r_i(a_i, \sigma_j) = r_i(a_i', \sigma_j)$.

- IBE instead requires

$$\sigma_i(a_i)\, r_i(a_i, \sigma_j) = \sigma_i(a_i')\, r_i(a_i', \sigma_j).$$

Why IBE is not Nash in transformed utilities

Although IBE uses transformed payoffs, it is not just Nash equilibrium of the transformed game.

- If players maximized expected transformed utility directly

  $\Longrightarrow$ i.e. minimizing regret r_i

  indifference would require $r_i(a_i, \sigma_j) = r_i(a_i', \sigma_j)$.

- IBE instead requires

$$\sigma_i(a_i)\, r_i(a_i, \sigma_j) = \sigma_i(a_i')\, r_i(a_i', \sigma_j).$$

- Thus, regret terms are weighted by the frequencies with which they are experienced.

Interpretation

IBE is a stationary consequence of ex post adjustment, not a one-shot best-response condition.

Selten and Chmura (2008): experimental design

Game 1

10 8	0 18
9 9	10 8

Game 7

10 12	4 22
9 9	14 8

Selten and Chmura (2008): experimental design

Game 1

10 8	0 18
9 9	10 8

Game 7

10 12	4 22
9 9	14 8

- Twelve completely mixed 2×2 games: 6 constant-sum and 6 non-constant-sum games.

Selten and Chmura (2008): experimental design

Game 1

10 8	0 18
9 9	10 8

Game 7

10 12	4 22
9 9	14 8

- Twelve completely mixed 2×2 games: 6 constant-sum and 6 non-constant-sum games.
- Fixed roles, anonymous random matching, 200 periods.

Selten and Chmura (2008): experimental design

Game 1

10, 8	0, 18
9, 9	10, 8

Game 7

10, 12	4, 22
9, 9	14, 8

- Twelve completely mixed 2×2 games: 6 constant-sum and 6 non-constant-sum games.
- Fixed roles, anonymous random matching, 200 periods.
- The design is tailored to compare stationary concepts.

In the figure, Player 1 is the row player and the row player's payoff is shown top-left.

Why this design is informative

The key variation in Selten and Chmura (2008) is that some games differ only by additive payoff shifts.

Why this design is informative

The key variation in Selten and Chmura (2008) is that some games differ only by additive payoff shifts.

- For Nash and logit / QRE, adding a constant in this way leaves the relevant payoff *differences* unchanged.

Why this design is informative

The key variation in Selten and Chmura (2008) is that some games differ only by additive payoff shifts.

- For Nash and logit / QRE, adding a constant in this way leaves the relevant payoff *differences* unchanged.
- Hence these concepts predict the same mixed behavior in the translated games.

Why this design is informative

The key variation in Selten and Chmura (2008) is that some games differ only by additive payoff shifts.

- For Nash and logit / QRE, adding a constant in this way leaves the relevant payoff *differences* unchanged.
- Hence these concepts predict the same mixed behavior in the translated games.
- For IBE, the maximin reference point and the regret terms can change, so the prediction may move.

Key point

The experiment is a clean test of the translation invariance underlying logit-style choice (Breitmoser, 2021).

Headline result in Selten and Chmura (2008)

Selten and Chmura compare five stationary concepts:

Headline result in Selten and Chmura (2008)

Selten and Chmura compare five stationary concepts:

- Nash equilibrium

Headline result in Selten and Chmura (2008)

Selten and Chmura compare five stationary concepts:

- Nash equilibrium
- quantal response equilibrium

Headline result in Selten and Chmura (2008)

Selten and Chmura compare five stationary concepts:

- Nash equilibrium
- quantal response equilibrium
- action-sampling equilibrium

Headline result in Selten and Chmura (2008)

Selten and Chmura compare five stationary concepts:

- Nash equilibrium
- quantal response equilibrium
- action-sampling equilibrium
- payoff-sampling equilibrium

Headline result in Selten and Chmura (2008)

Selten and Chmura compare five stationary concepts:

- Nash equilibrium
- quantal response equilibrium
- action-sampling equilibrium
- payoff-sampling equilibrium
- impulse balance equilibrium

Their overall ranking is

$$\text{IBE} > \text{payoff sampling} > \text{action sampling} > \text{QRE} > \text{Nash}.$$

Immediate lesson

Nash fits worst. Among non-Nash concepts, models that incorporate regret or local procedures perform substantially better.

Econometric reassessment (Brunner et al., 2011)

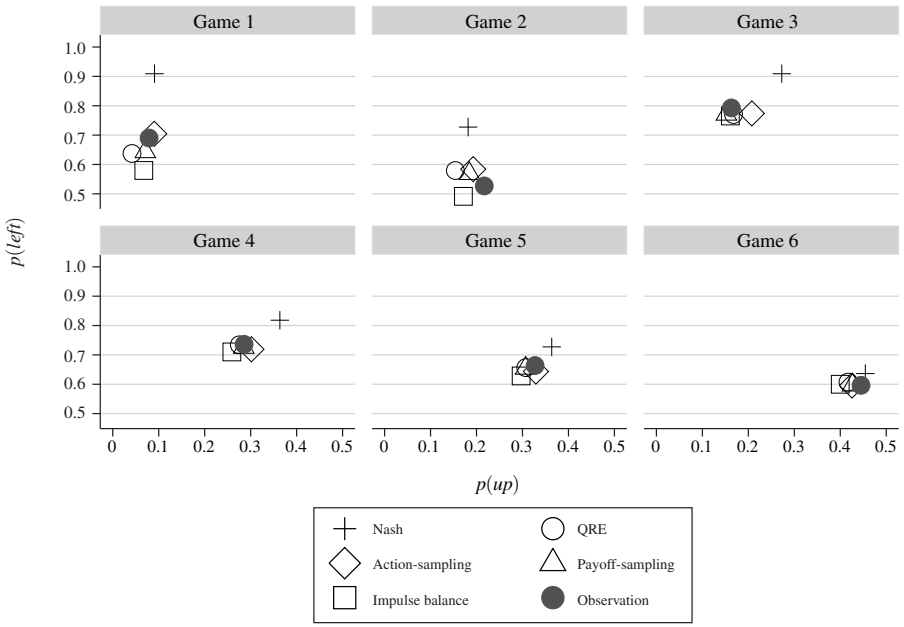


FIGURE 2. VISUALIZATION OF THE THEORETICAL EQUILIBRIA AND THE OBSERVED AVERAGE IN THE CONSTANT SUM GAMES (*cf. SC 2008, Figure 8*)

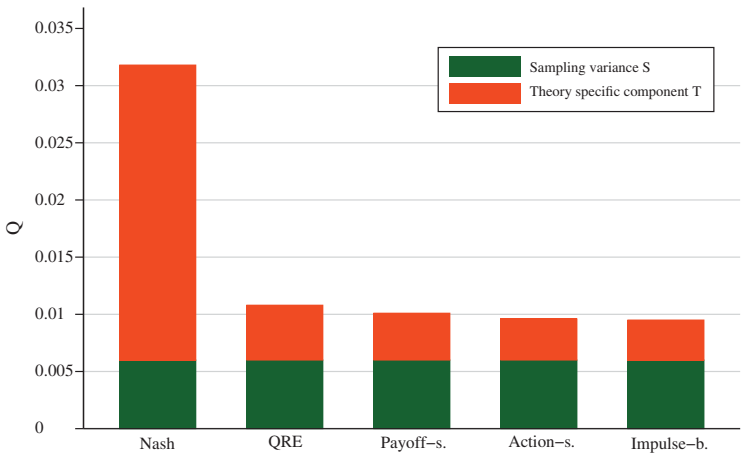


FIGURE 3. OVERALL MEAN SQUARED DISTANCES OF THE FIVE STATIONARY CONCEPTS COMPARED TO THE OBSERVED AVERAGE (*cf. SC 2008, Figure 12*)

Econometric reassessment (Brunner et al., 2011)

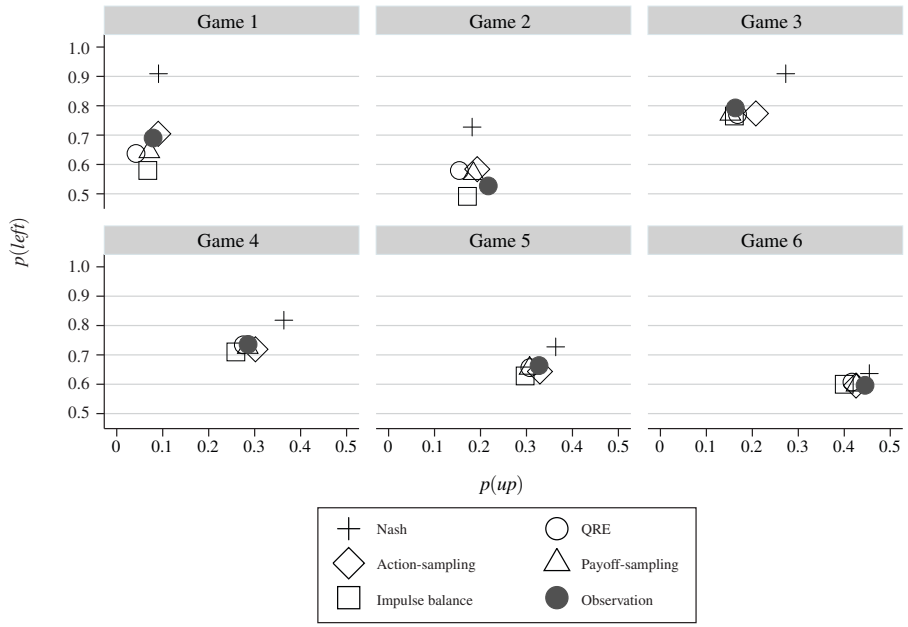


FIGURE 2. VISUALIZATION OF THE THEORETICAL EQUILIBRIA AND THE OBSERVED AVERAGE IN THE CONSTANT SUM GAMES (*cf. SC 2008, Figure 8*)

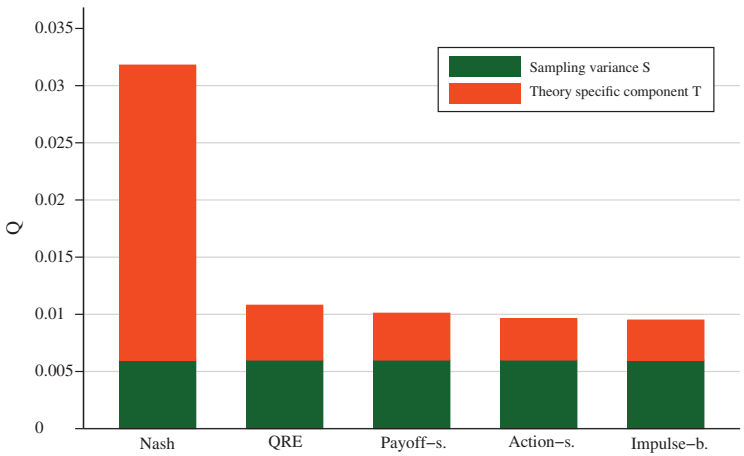


FIGURE 3. OVERALL MEAN SQUARED DISTANCES OF THE FIVE STATIONARY CONCEPTS COMPARED TO THE OBSERVED AVERAGE (*cf. SC 2008, Figure 12*)

- Brunner, Camerer, and Goeree reanalyze the Selten–Chmura data and correct errors.

Econometric reassessment (Brunner et al., 2011)

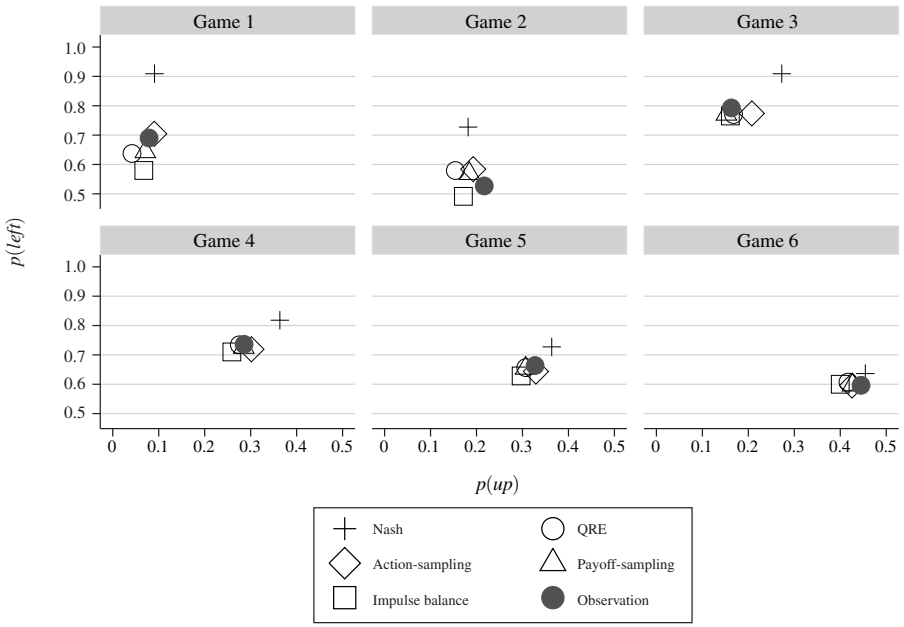


FIGURE 2. VISUALIZATION OF THE THEORETICAL EQUILIBRIA AND THE OBSERVED AVERAGE IN THE CONSTANT SUM GAMES (*cf. SC 2008, Figure 8*)

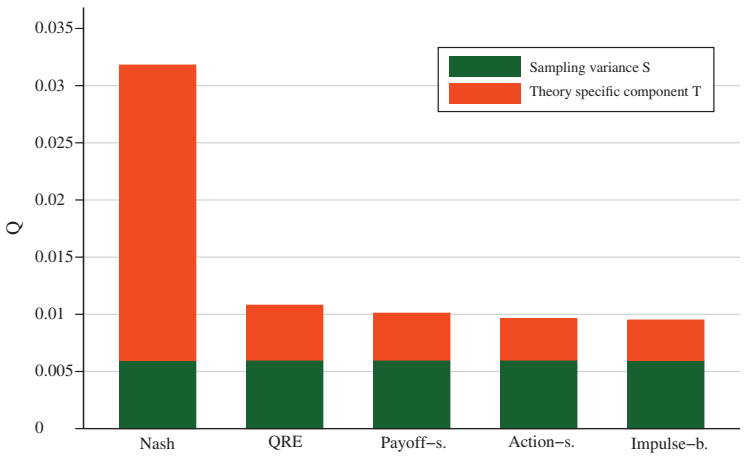


FIGURE 3. OVERALL MEAN SQUARED DISTANCES OF THE FIVE STATIONARY CONCEPTS COMPARED TO THE OBSERVED AVERAGE (*cf. SC 2008, Figure 12*)

- Brunner, Camerer, and Goeree reanalyze the Selten–Chmura data and correct errors.
- Nash still performs clearly worst.

Econometric reassessment (Brunner et al., 2011)

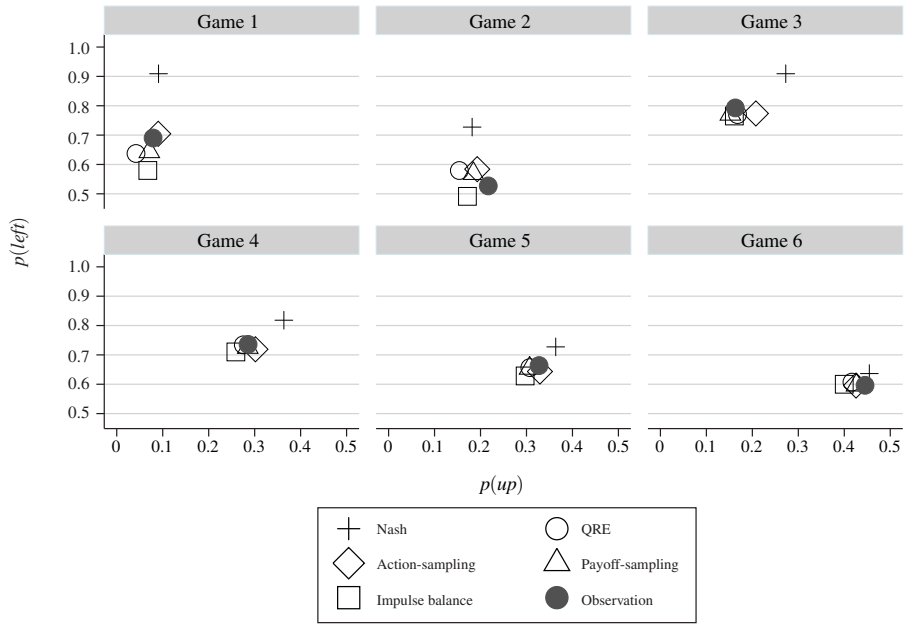


FIGURE 2. VISUALIZATION OF THE THEORETICAL EQUILIBRIA AND THE OBSERVED AVERAGE IN THE CONSTANT SUM GAMES (*cf. SC 2008, Figure 8*)

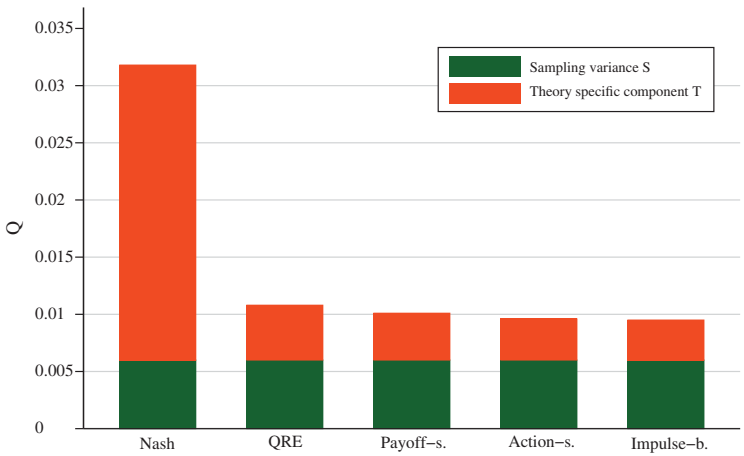


FIGURE 3. OVERALL MEAN SQUARED DISTANCES OF THE FIVE STATIONARY CONCEPTS COMPARED TO THE OBSERVED AVERAGE (*cf. SC 2008, Figure 12*)

- Brunner, Camerer, and Goeree reanalyze the Selten–Chmura data and correct errors.
- Nash still performs clearly worst.
- Among the four non-Nash concepts, differences are much smaller than the original ranking suggests.

Econometric reassessment (Brunner et al., 2011)

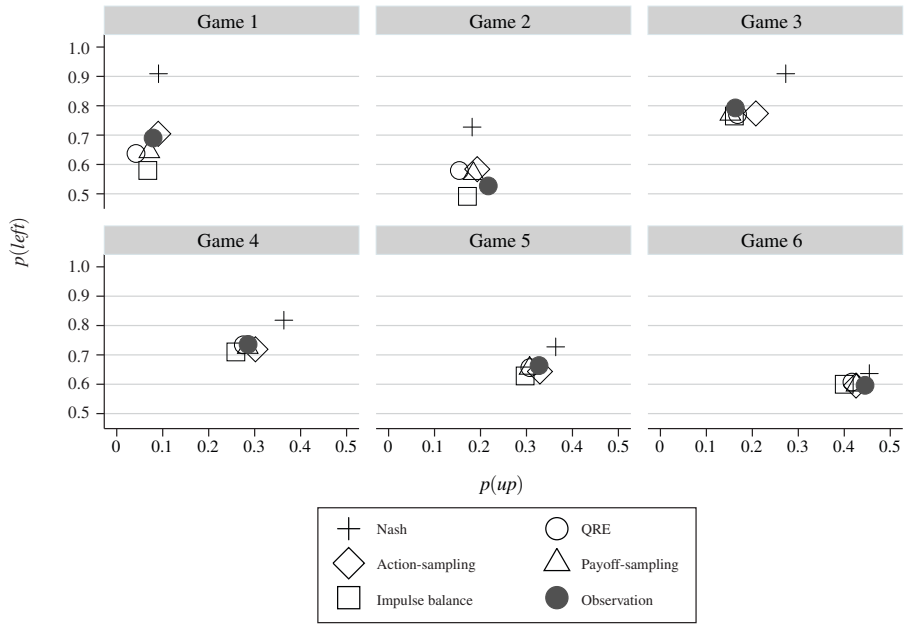


FIGURE 2. VISUALIZATION OF THE THEORETICAL EQUILIBRIA AND THE OBSERVED AVERAGE IN THE CONSTANT SUM GAMES (*cf. SC 2008, Figure 8*)

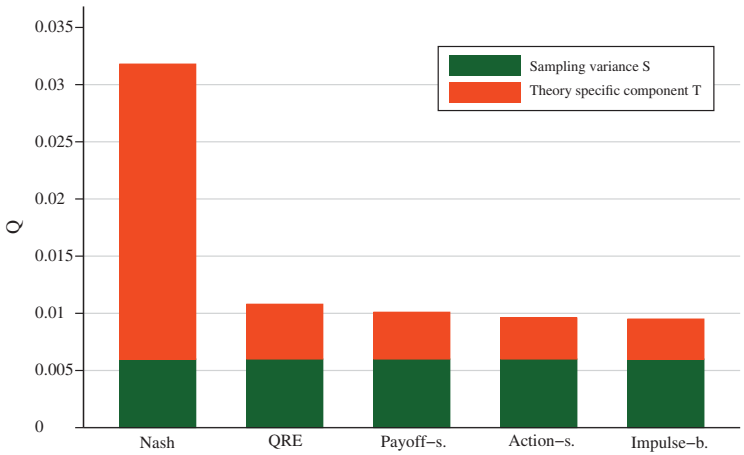


FIGURE 3. OVERALL MEAN SQUARED DISTANCES OF THE FIVE STATIONARY CONCEPTS COMPARED TO THE OBSERVED AVERAGE (*cf. SC 2008, Figure 12*)

- Brunner, Camerer, and Goeree reanalyze the Selten–Chmura data and correct errors.
- Nash still performs clearly worst.
- Among the four non-Nash concepts, differences are much smaller than the original ranking suggests.
- In additional data, concepts augmented with loss aversion can outperform plain IBE.

Implication

IBE is important, but its empirical success may partly come from the loss-aversion ingredient hardwired into the concept.

Section takeaway: what IBE contributes

Section takeaway: what IBE contributes

- IBE offers a genuinely different route beyond Nash: bounded rationality via regret and directional impulses.

Section takeaway: what IBE contributes

- IBE offers a genuinely different route beyond Nash: bounded rationality via regret and directional impulses.
- It is a stationary concept, not a one-shot best-response concept.

Section takeaway: what IBE contributes

- IBE offers a genuinely different route beyond Nash: bounded rationality via regret and directional impulses.
- It is a stationary concept, not a one-shot best-response concept.
- In the Selten–Chmura horse race, IBE performs very well, but later reassessment shows that it should not be treated as a universal winner.

Section takeaway: what IBE contributes

- IBE offers a genuinely different route beyond Nash: bounded rationality via regret and directional impulses.
- It is a stationary concept, not a one-shot best-response concept.
- In the Selten–Chmura horse race, IBE performs very well, but later reassessment shows that it should not be treated as a universal winner.
- The main conceptual gain is that deviations from Nash need not come only from noisy beliefs; they may also reflect how foregone payoffs are processed.

Transition

We now move from regret-based adjustment to finite local procedures: payoff sampling and action sampling.

Osborne–Rubinstein Sampling Concepts

Generalized concepts and evidence

Why add sampling concepts here?

CHM and QRE-type models describe bounded rationality mainly through beliefs and responses to beliefs.

Osborne and Rubinstein propose a different route:

Why add sampling concepts here?

CHM and QRE-type models describe bounded rationality mainly through beliefs and responses to beliefs.

Osborne and Rubinstein propose a different route:

- players need not solve the full strategic problem directly;

Why add sampling concepts here?

CHM and QRE-type models describe bounded rationality mainly through beliefs and responses to beliefs.

Osborne and Rubinstein propose a different route:

- players need not solve the full strategic problem directly;
- instead, they use a simple decision procedure;

Why add sampling concepts here?

CHM and QRE-type models describe bounded rationality mainly through beliefs and responses to beliefs.

Osborne and Rubinstein propose a different route:

- players need not solve the full strategic problem directly;
- instead, they use a simple decision procedure;
- randomness comes from finite sampling, not from an exogenous error term;

Why add sampling concepts here?

CHM and QRE-type models describe bounded rationality mainly through beliefs and responses to beliefs.

Osborne and Rubinstein propose a different route:

- players need not solve the full strategic problem directly;
- instead, they use a simple decision procedure;
- randomness comes from finite sampling, not from an exogenous error term;
- equilibrium is a steady state in which the procedure reproduces the observed population behavior.

Sampling concepts therefore sit naturally between IBE and HQRE.

Procedural rationality

The core idea is to replace full strategic reasoning by a simple choice procedure.

Procedural rationality

The core idea is to replace full strategic reasoning by a simple choice procedure.

- In GT, a player forms a belief about opponents and evaluates all relevant contingencies.

Procedural rationality

The core idea is to replace full strategic reasoning by a simple choice procedure.

- In GT, a player forms a belief about opponents and evaluates all relevant contingencies.
- In procedural rationality, the player does not solve the full model.

Procedural rationality

The core idea is to replace full strategic reasoning by a simple choice procedure.

- In GT, a player forms a belief about opponents and evaluates all relevant contingencies.
- In procedural rationality, the player does not solve the full model.
- Instead, she constructs an action–consequence correspondence from a small amount of sampled information and then chooses the action that looks best.

Bounded rationality here

The player ignores the fact that the sample itself is random and treats it as sufficiently informative for choice.

Payoff sampling: the Osborne–Rubinstein idea

Consider a player i with two actions a_i and a_i' facing opponent distribution σ_j.

Payoff sampling: the Osborne–Rubinstein idea

Consider a player i with two actions a_i and a_i' facing opponent distribution σ_j.

- The player samples one consequence for each of her own actions.

Payoff sampling: the Osborne–Rubinstein idea

Consider a player i with two actions a_i and a_i' facing opponent distribution σ_j.

- The player samples one consequence for each of her own actions.
- Formally, draw $\tilde{a}_j, \tilde{a}_j'$ independently from σ_j.

Payoff sampling: the Osborne–Rubinstein idea

Consider a player i with two actions a_i and a_i' facing opponent distribution σ_j.

- The player samples one consequence for each of her own actions.
- Formally, draw $\tilde{a}_j, \tilde{a}_j'$ independently from σ_j.
- Compare the sampled payoffs

$$\pi_i(a_i, \tilde{a}_j) \qquad \text{and} \qquad \pi_i(a_i', \tilde{a}_j').$$

Payoff sampling: the Osborne–Rubinstein idea

Consider a player i with two actions a_i and a_i' facing opponent distribution σ_j.

- The player samples one consequence for each of her own actions.
- Formally, draw $\tilde{a}_j, \tilde{a}_j'$ independently from σ_j.
- Compare the sampled payoffs

$$\pi_i(a_i, \tilde{a}_j) \qquad \text{and} \qquad \pi_i(a_i', \tilde{a}_j').$$

- Choose the action with the larger sampled payoff; if the sampled payoffs are equal, break the tie uniformly.

Interpretation

The player does *not* compute expected payoffs. She chooses the action that looks best in a small local experiment.

Payoff-sampling equilibrium in a 2×2 game

Let $\tilde{a}_j, \tilde{a}'_j \sim \sigma_j$ independently. The probability that player i chooses a_i is

$$\sigma_i(a_i) = \Pr\left[\pi_i(a_i, \tilde{a}_j) > \pi_i(a'_i, \tilde{a}'_j)\right] + \frac{1}{2}\Pr\left[\pi_i(a_i, \tilde{a}_j) = \pi_i(a'_i, \tilde{a}'_j)\right].$$

A payoff-sampling equilibrium is a mixed profile for which these choice probabilities reproduce the same population frequencies

Payoff-sampling equilibrium in a 2×2 game

Let $\tilde{a}_j, \tilde{a}'_j \sim \sigma_j$ independently. The probability that player i chooses a_i is

$$\sigma_i(a_i) = \Pr\left[\pi_i(a_i, \tilde{a}_j) > \pi_i(a'_i, \tilde{a}'_j)\right] + \frac{1}{2}\Pr\left[\pi_i(a_i, \tilde{a}_j) = \pi_i(a'_i, \tilde{a}'_j)\right].$$

A payoff-sampling equilibrium is a mixed profile for which these choice probabilities reproduce the same population frequencies

- In the original paper, each action is sampled once.

Payoff-sampling equilibrium in a 2×2 game

Let $\tilde{a}_j, \tilde{a}'_j \sim \sigma_j$ independently. The probability that player i chooses a_i is

$$\sigma_i(a_i) = \Pr\left[\pi_i(a_i, \tilde{a}_j) > \pi_i(a'_i, \tilde{a}'_j)\right] + \frac{1}{2}\Pr\left[\pi_i(a_i, \tilde{a}_j) = \pi_i(a'_i, \tilde{a}'_j)\right].$$

A payoff-sampling equilibrium is a mixed profile for which these choice probabilities reproduce the same population frequencies

- In the original paper, each action is sampled once.
- In the experimental implementation of Selten and Chmura (2008), larger equal samples per action are used to obtain sharper stationary predictions.

Example: Payoff-sampling equilibrium in Hawk–Dove

Game.

	H	D
H	$-1, -1$	$4, 0$
D	$0, 4$	$2, 2$

Nash benchmark.

$-q + 4(1-q) = 2(1-q)$

$$\Rightarrow \quad q^*_{\text{Nash}} = \tfrac{2}{3}.$$

Sampling choice. Opponent plays H w.p. q.

$(\tilde{a}_j, \tilde{a}'_j)$	$\pi(H, \tilde{a}_j)$	$\pi(D, \tilde{a}'_j)$	pick
(H, H)	-1	0	D
(H, D)	-1	2	D
(D, H)	4	0	H
(D, D)	4	2	H

Player picks H iff $\tilde{a}_j = D$: $\sigma_i(H) = 1 - q$.

Equilibrium. Reproduce the population: $q = 1 - q$, so

$$\boxed{q^{\text{PSE}} = \tfrac{1}{2}.}$$

Action sampling and sampling equilibrium

The later Osborne–Rubinstein idea shifts the sampling from payoffs of own actions to opponents' actions.

Action sampling and sampling equilibrium

The later Osborne–Rubinstein idea shifts the sampling from payoffs of own actions to opponents' actions.

- A player observes a sample of k actions drawn from the opponent population.

Action sampling and sampling equilibrium

The later Osborne–Rubinstein idea shifts the sampling from payoffs of own actions to opponents' actions.

- A player observes a sample of k actions drawn from the opponent population.
- She treats the empirical frequency in this sample as a prediction of the population distribution.

Action sampling and sampling equilibrium

The later Osborne–Rubinstein idea shifts the sampling from payoffs of own actions to opponents' actions.

- A player observes a sample of k actions drawn from the opponent population.
- She treats the empirical frequency in this sample as a prediction of the population distribution.
- She then plays a best response to that empirical distribution.

Let $\hat{\sigma}_j^k$ be the empirical distribution from k independent draws from σ_j. Then

$$\sigma_i(a_i) = \Pr\left[a_i \text{ is the unique best response to } \hat{\sigma}_j^k\right] + \frac{1}{2}\Pr[a_i \text{ ties for best response}].$$

Payoff sampling versus action sampling

Payoff sampling

Payoff sampling versus action sampling

Payoff sampling

- samples *consequences*

Payoff sampling versus action sampling

Payoff sampling

- samples *consequences*
- each own action is locally "tested"

Payoff sampling versus action sampling

Payoff sampling

- samples *consequences*
- each own action is locally "tested"
- randomness comes from payoff comparisons

Payoff sampling versus action sampling

Payoff sampling

- samples *consequences*
- each own action is locally "tested"
- randomness comes from payoff comparisons
- closer to a primitive choice experiment

Action sampling

Payoff sampling versus action sampling

Payoff sampling

- samples *consequences*
- each own action is locally "tested"
- randomness comes from payoff comparisons
- closer to a primitive choice experiment

Action sampling

- samples *opponents' actions*

Payoff sampling versus action sampling

Payoff sampling

- samples *consequences*
- each own action is locally "tested"
- randomness comes from payoff comparisons
- closer to a primitive choice experiment

Action sampling

- samples *opponents' actions*
- constructs a local empirical belief

Payoff sampling versus action sampling

Payoff sampling

- samples *consequences*
- each own action is locally "tested"
- randomness comes from payoff comparisons
- closer to a primitive choice experiment

Action sampling

- samples *opponents' actions*
- constructs a local empirical belief
- randomness comes from sampling error in observed frequencies

Payoff sampling versus action sampling

Payoff sampling

- samples *consequences*
- each own action is locally "tested"
- randomness comes from payoff comparisons
- closer to a primitive choice experiment

Action sampling

- samples *opponents' actions*
- constructs a local empirical belief
- randomness comes from sampling error in observed frequencies
- closer to bounded strategic inference

Common feature

In both models, noise is endogenous: it is generated by finite samples rather than imposed through a payoff-sensitivity parameter as in logit/QRE.

Why these concepts matter empirically

Sampling concepts are attractive because they are behaviorally concrete.

Why these concepts matter empirically

Sampling concepts are attractive because they are behaviorally concrete.

- They make bounded rationality procedural rather than purely reduced-form.

Why these concepts matter empirically

Sampling concepts are attractive because they are behaviorally concrete.

- They make bounded rationality procedural rather than purely reduced-form.
- They generate stochastic play without an ad hoc error term.

Why these concepts matter empirically

Sampling concepts are attractive because they are behaviorally concrete.

- They make bounded rationality procedural rather than purely reduced-form.
- They generate stochastic play without an ad hoc error term.
- The amount of noise is linked to sample size and payoff structure.

Why these concepts matter empirically

Sampling concepts are attractive because they are behaviorally concrete.

- They make bounded rationality procedural rather than purely reduced-form.
- They generate stochastic play without an ad hoc error term.
- The amount of noise is linked to sample size and payoff structure.
- In Selten and Chmura (2008), both sampling models outperform Nash, and payoff sampling outperforms action sampling in the overall ranking.

Why these concepts matter empirically

Sampling concepts are attractive because they are behaviorally concrete.

- They make bounded rationality procedural rather than purely reduced-form.
- They generate stochastic play without an ad hoc error term.
- The amount of noise is linked to sample size and payoff structure.
- In Selten and Chmura (2008), both sampling models outperform Nash, and payoff sampling outperforms action sampling in the overall ranking.
- In Brunner et al. (2011), the differences among the non-Nash models become less sharp.

Main lesson

Noisy aggregate behavior can come from simple local procedures, not only from noisy optimization.

Section takeaway: what Osborne–Rubinstein add

Section takeaway: what Osborne–Rubinstein add

- Osborne–Rubinstein add a distinct behavioral foundation: procedural rationality.

Section takeaway: what Osborne–Rubinstein add

- Osborne–Rubinstein add a distinct behavioral foundation: procedural rationality.
- Payoff sampling models players as testing their own actions through sampled consequences.

Section takeaway: what Osborne–Rubinstein add

- Osborne–Rubinstein add a distinct behavioral foundation: procedural rationality.
- Payoff sampling models players as testing their own actions through sampled consequences.
- Action sampling models players as observing a small sample of others' actions and best-responding to the sample.

Section takeaway: what Osborne–Rubinstein add

- Osborne–Rubinstein add a distinct behavioral foundation: procedural rationality.
- Payoff sampling models players as testing their own actions through sampled consequences.
- Action sampling models players as observing a small sample of others' actions and best-responding to the sample.
- Both concepts are equilibrium-based, but the randomness in behavior is generated by finite sampling.

Transition

We now return to equilibrium language and ask how far we can get by allowing heterogeneous noisy response directly.

Heterogeneous QRE

Generalized concepts and evidence

Why heterogeneous QRE?

Standard QRE assumes that all players share the same payoff-sensitivity parameter. This is often too restrictive in experimental data.

Why heterogeneous QRE?

Standard QRE assumes that all players share the same payoff-sensitivity parameter. This is often too restrictive in experimental data.

- Some subjects respond very sharply to payoff differences.

Why heterogeneous QRE?

Standard QRE assumes that all players share the same payoff-sensitivity parameter. This is often too restrictive in experimental data.

- Some subjects respond very sharply to payoff differences.
- Others choose more noisily, even in the same game.

Why heterogeneous QRE?

Standard QRE assumes that all players share the same payoff-sensitivity parameter. This is often too restrictive in experimental data.

- Some subjects respond very sharply to payoff differences.
- Others choose more noisily, even in the same game.
- More sophisticated players may also hold more accurate beliefs about the responsiveness of others.

Question

Can we generalize QRE so that precision is heterogeneous, while preserving the equilibrium logic of mutually consistent stochastic choice?

Heterogeneous logit equilibrium (HLE / HQRE)

Heterogeneous logit equilibrium

Player i has payoff sensitivity $\lambda_i \in \mathbb{R}_+$. The type λ_i is private information, but its density $f_i(\lambda_i)$ is common knowledge.

Given a mixed strategy profile σ, type λ_i chooses action $a_i \in A_i$ with probability

$$p_i(a_i \mid \lambda_i) = \frac{\exp\{\lambda_i U_i(a_i \mid \sigma)\}}{\sum_{a_i' \in A_i} \exp\{\lambda_i U_i(a_i' \mid \sigma)\}}.$$

The induced ex-ante mixed strategy of player i is

$$\sigma_i(a_i) = \int_0^\infty p_i(a_i \mid \lambda)\, f_i(\lambda)\, d\lambda.$$

Expected utility and equilibrium consistency

In HQRE, each type of player i evaluates actions against the induced mixed strategies of the opponents:

$$U_i(a_i \mid \sigma) = \sum_{a_{-i} \in A_{-i}} \left(\prod_{j \neq i} \sigma_j(a_j) \right) u_i(a_i, a_{-i}).$$

Expected utility and equilibrium consistency

In HQRE, each type of player i evaluates actions against the induced mixed strategies of the opponents:

$$U_i(a_i \mid \sigma) = \sum_{a_{-i} \in A_{-i}} \left(\prod_{j \neq i} \sigma_j(a_j) \right) u_i(a_i, a_{-i}).$$

- The response rule is stochastic, as in standard QRE.

Expected utility and equilibrium consistency

In HQRE, each type of player i evaluates actions against the induced mixed strategies of the opponents:

$$U_i(a_i \mid \sigma) = \sum_{a_{-i} \in A_{-i}} \left(\prod_{j \neq i} \sigma_j(a_j) \right) u_i(a_i, a_{-i}).$$

- The response rule is stochastic, as in standard QRE.
- Heterogeneity enters through the distribution of λ_i.

Expected utility and equilibrium consistency

In HQRE, each type of player i evaluates actions against the induced mixed strategies of the opponents:

$$U_i(a_i \mid \sigma) = \sum_{a_{-i} \in A_{-i}} \left(\prod_{j \neq i} \sigma_j(a_j) \right) u_i(a_i, a_{-i}).$$

- The response rule is stochastic, as in standard QRE.
- Heterogeneity enters through the distribution of λ_i.
- Beliefs are still correct on average: all players use the same induced profile σ.

Interpretation

HQRE is the cleanest extension of QRE when we want heterogeneous noisy response without abandoning equilibrium consistency (McKelvey and Palfrey, 1995; Rogers et al., 2009).

Truncated QRE (TLE / TQRE): heterogeneity plus limited insight

HQRE still assumes that all players correctly understand the full distribution of others' precisions. TQRE relaxes this.

Truncated QRE

A type λ_i player believes that each opponent's precision satisfies $\lambda_j \leq \theta_i(\lambda_i)$, where $\theta_i(\lambda_i)$ is increasing and common knowledge.

Thus player i's belief about opponent j is characterized by the truncated density

$$f_j(\lambda_j \mid \lambda_i) = \frac{f_j(\lambda_j)}{F_j(\theta_i(\lambda_i))} \qquad \text{for } \lambda_j \leq \theta_i(\lambda_i).$$

More precise players are assumed to imagine a wider range of opponent types; self-limited players satisfy $\theta_i(\lambda_i) \leq \lambda_i$ for all λ_i.

How TQRE connects QRE and CHM

TQRE is important because it provides a formal bridge between stochastic QRE models and level-based models of strategic thinking.

How TQRE connects QRE and CHM

TQRE is important because it provides a formal bridge between stochastic QRE models and level-based models of strategic thinking.

- If $\theta_i(\lambda_i) \to \infty$, TQRE approaches HQRE.

How TQRE connects QRE and CHM

TQRE is important because it provides a formal bridge between stochastic QRE models and level-based models of strategic thinking.

- If $\theta_i(\lambda_i) \to \infty$, TQRE approaches HQRE.
- If $\theta_i(\lambda_i) \to 0$, behavior approaches a level-1-type limit.

How TQRE connects QRE and CHM

TQRE is important because it provides a formal bridge between stochastic QRE models and level-based models of strategic thinking.

- If $\theta_i(\lambda_i) \to \infty$, TQRE approaches HQRE.
- If $\theta_i(\lambda_i) \to 0$, behavior approaches a level-1-type limit.
- If the precision distribution is discretized, TQRE can approximate CHM arbitrarily closely (Rogers et al., 2009; Camerer et al., 2004).

Big picture

CHM can therefore be interpreted as a limiting case of a model with heterogeneous stochastic choice and downward-looking beliefs.

Original evidence: HQRE, TQRE, and CHM

Matrix game estimates												
Game	Description	Random neg Log L	Empirical $-\log L$	CH		QRE		TQRE			HQRE		
				τ	$-\log L$	λ	$-\log L$	Y	τ	$-\log L$	Λ	ε	$-\log L$
1	Unprofitable	81.30	73.12	10.19	73.65	0.06	75.97	0.35	3.73	73.12	0.07	0.13	75.9
2	Cloned MP	81.30	52.34	0.90	58.39	0.12	71.72	∞	0.90	58.39	0.12	0	71.7
3	Cloned SH (Low)	51.96	46.50	0.44	46.72	0.15	46.74	0.50	0.62	46.67	0.15	0.3	46.7
4	Cloned SH (Hi)	51.96	40.18	1.00	41.91	0.18	41.88	0.89	1.40	40.65	0.19	0.29	41.8
5	SW 1	81.30	39.55	3.10	40.92	0.13	41.51	0.17	2.66	39.93	0.15	0.23	40.5
6	SW 2	81.30	46.00	5.59	46.03	0.13	46.75	∞	5.59	46.03	0.13	0	46.7
7	SW 3	81.30	53.26	0.72	63.26	0.07	73.94	∞	0.72	63.26	0.08	0.16	73.6
8	SW 4	81.30	55.02	1.96	56.11	∞	56.40	∞	1.96	56.11	∞	–	56.4
9	SW 5	81.30	79.27	0.00	81.30	0.14	80.50	0.00	0.00	81.30	0.15	0.3	80.4
10	SW 6	81.30	79.54	1.27	79.69	0.05	80.86	0.11	2.81	79.54	0.05	0.1	80.8
11	SW 7	81.30	73.69	0.67	73.73	∞	73.73	∞	0.67	73.73	∞	–	73.7
12	SW 8	81.30	70.76	0.00	81.30	0.98	73.25	0.02	51.30	78.30	1	1.5	73.2
13	SW 9	81.30	54.30	1.01	60.94	0.08	66.56	∞	1.01	60.94	0.08	0	66.5
14	SW 10	81.30	73.12	13.90	77.35	0.42	73.17	0.15	7.13	73.31	0.44	0.88	73.1
15	SW 11	81.30	73.81	0.83	73.82	0.17	74.55	0.17	0.93	73.81	0.18	0.36	74.5
16	SW 12	81.30	66.81	0.94	67.02	0.03	72.94	0.16	0.87	66.81	0.04	0.07	72.7
17	Cloned Joker	86.88	77.28	0.83	81.08	0.16	83.62	∞	0.82	81.07	0.17	0	83.6
	Sum	1329.00	1054.55	–	1103.21	–	1134.08	–	–	1092.96	–	–	1132.3
	Pooled	–	–	0.53	1239.26	0.10	1192.00	0.11	1.57	1164.41	0.12	0.24	1182.9

Original evidence: HQRE, TQRE, and CHM

Matrix game estimates												
Game	Description	Random neg Log L	Empirical $-\log L$	CH		QRE		TQRE			HQRE		
				τ	$-\log L$	λ	$-\log L$	Y	τ	$-\log L$	Λ	ε	$-\log L$
1	Unprofitable	81.30	73.12	10.19	73.65	0.06	75.97	0.35	3.73	73.12	0.07	0.13	75.9
2	Cloned MP	81.30	52.34	0.90	58.39	0.12	71.72	∞	0.90	58.39	0.12	0	71.7
3	Cloned SH (Low)	51.96	46.50	0.44	46.72	0.15	46.74	0.50	0.62	46.67	0.15	0.3	46.7
4	Cloned SH (Hi)	51.96	40.18	1.00	41.91	0.18	41.88	0.89	1.40	40.65	0.19	0.29	41.8
5	SW 1	81.30	39.55	3.10	40.92	0.13	41.51	0.17	2.66	39.93	0.15	0.23	40.5
6	SW 2	81.30	46.00	5.59	46.03	0.13	46.75	∞	5.59	46.03	0.13	0	46.7
7	SW 3	81.30	53.26	0.72	63.26	0.07	73.94	∞	0.72	63.26	0.08	0.16	73.6
8	SW 4	81.30	55.02	1.96	56.11	∞	56.40	∞	1.96	56.11	∞	–	56.4
9	SW 5	81.30	79.27	0.00	81.30	0.14	80.50	0.00	0.00	81.30	0.15	0.3	80.4
10	SW 6	81.30	79.54	1.27	79.69	0.05	80.86	0.11	2.81	79.54	0.05	0.1	80.8
11	SW 7	81.30	73.69	0.67	73.73	∞	73.73	∞	0.67	73.73	∞	–	73.7
12	SW 8	81.30	70.76	0.00	81.30	0.98	73.25	0.02	51.30	78.30	1	1.5	73.2
13	SW 9	81.30	54.30	1.01	60.94	0.08	66.56	∞	1.01	60.94	0.08	0	66.5
14	SW 10	81.30	73.12	13.90	77.35	0.42	73.17	0.15	7.13	73.31	0.44	0.88	73.1
15	SW 11	81.30	73.81	0.83	73.82	0.17	74.55	0.17	0.93	73.81	0.18	0.36	74.5
16	SW 12	81.30	66.81	0.94	67.02	0.03	72.94	0.16	0.87	66.81	0.04	0.07	72.7
17	Cloned Joker	86.88	77.28	0.83	81.08	0.16	83.62	∞	0.82	81.07	0.17	0	83.6
	Sum	1329.00	1054.55	–	1103.21	–	1134.08	–	–	1092.96	–	–	1132.3
	Pooled	–	–	0.53	1239.26	0.10	1192.00	0.11	1.57	1164.41	0.12	0.24	1182.9

HQRE:

$$\lambda \sim \text{Unif}[\Lambda - \varepsilon/2, \Lambda + \varepsilon/2]$$

Original evidence: HQRE, TQRE, and CHM

Matrix game estimates												
Game	Description	Random neg Log L	Empirical $-\log L$	CH		QRE		TQRE			HQRE		
				τ	$-\log L$	λ	$-\log L$	Y	τ	$-\log L$	Λ	ε	$-\log L$
1	Unprofitable	81.30	73.12	10.19	73.65	0.06	75.97	0.35	3.73	73.12	0.07	0.13	75.9
2	Cloned MP	81.30	52.34	0.90	58.39	0.12	71.72	∞	0.90	58.39	0.12	0	71.7
3	Cloned SH (Low)	51.96	46.50	0.44	46.72	0.15	46.74	0.50	0.62	46.67	0.15	0.3	46.7
4	Cloned SH (Hi)	51.96	40.18	1.00	41.91	0.18	41.88	0.89	1.40	40.65	0.19	0.29	41.8
5	SW 1	81.30	39.55	3.10	40.92	0.13	41.51	0.17	2.66	39.93	0.15	0.23	40.5
6	SW 2	81.30	46.00	5.59	46.03	0.13	46.75	∞	5.59	46.03	0.13	0	46.7
7	SW 3	81.30	53.26	0.72	63.26	0.07	73.94	∞	0.72	63.26	0.08	0.16	73.6
8	SW 4	81.30	55.02	1.96	56.11	∞	56.40	∞	1.96	56.11	∞	–	56.4
9	SW 5	81.30	79.27	0.00	81.30	0.14	80.50	0.00	0.00	81.30	0.15	0.3	80.4
10	SW 6	81.30	79.54	1.27	79.69	0.05	80.86	0.11	2.81	79.54	0.05	0.1	80.8
11	SW 7	81.30	73.69	0.67	73.73	∞	73.73	∞	0.67	73.73	∞	–	73.7
12	SW 8	81.30	70.76	0.00	81.30	0.98	73.25	0.02	51.30	78.30	1	1.5	73.2
13	SW 9	81.30	54.30	1.01	60.94	0.08	66.56	∞	1.01	60.94	0.08	0	66.5
14	SW 10	81.30	73.12	13.90	77.35	0.42	73.17	0.15	7.13	73.31	0.44	0.88	73.1
15	SW 11	81.30	73.81	0.83	73.82	0.17	74.55	0.17	0.93	73.81	0.18	0.36	74.5
16	SW 12	81.30	66.81	0.94	67.02	0.03	72.94	0.16	0.87	66.81	0.04	0.07	72.7
17	Cloned Joker	86.88	77.28	0.83	81.08	0.16	83.62	∞	0.82	81.07	0.17	0	83.6
	Sum	1329.00	1054.55	–	1103.21	–	1134.08	–	–	1092.96	–	–	1132.3
	Pooled	–	–	0.53	1239.26	0.10	1192.00	0.11	1.57	1164.41	0.12	0.24	1182.9

HQRE:

$\lambda \sim \text{Unif}[\Lambda - \varepsilon/2, \Lambda + \varepsilon/2]$

TQRE:

$\lambda = k\gamma$ with Poisson-distributed k and self-limited types

Takeaway: TQRE often fits pooled data better than CHM, but overall fit differences across HQRE, TQRE, and CHM remain fairly small (Rogers et al., 2009).

What HQRE and TQRE add

What HQRE and TQRE add

- HQRE adds heterogeneity to standard QRE while preserving common beliefs about the distribution of precision.

What HQRE and TQRE add

- HQRE adds heterogeneity to standard QRE while preserving common beliefs about the distribution of precision.
- TQRE adds a second ingredient: more precise players may still underestimate how sophisticated others are.

What HQRE and TQRE add

- HQRE adds heterogeneity to standard QRE while preserving common beliefs about the distribution of precision.
- TQRE adds a second ingredient: more precise players may still underestimate how sophisticated others are.
- These models provide a unified language for comparing QRE-style noise and CHM-style downward-looking reasoning.

What HQRE and TQRE add

- HQRE adds heterogeneity to standard QRE while preserving common beliefs about the distribution of precision.
- TQRE adds a second ingredient: more precise players may still underestimate how sophisticated others are.
- These models provide a unified language for comparing QRE-style noise and CHM-style downward-looking reasoning.
- Empirically, they often improve flexibility, but they do not create a large clean winner across one-shot games (Rogers et al., 2009; Crawford et al., 2013).

Transition

We now step back and compare all four routes studied today.

Synthesis and Conclusion

Generalized concepts and evidence

Four routes beyond simple Nash

Today's lecture developed four distinct ways to enrich one-shot behavioral game theory:

Four routes beyond simple Nash

Today's lecture developed four distinct ways to enrich one-shot behavioral game theory:

- CHM: bounded rationality through *heterogeneous depths of reasoning*

Four routes beyond simple Nash

Today's lecture developed four distinct ways to enrich one-shot behavioral game theory:

- CHM: bounded rationality through *heterogeneous depths of reasoning*
- IBE: bounded rationality through *regret and foregone payoffs*

Four routes beyond simple Nash

Today's lecture developed four distinct ways to enrich one-shot behavioral game theory:

- CHM: bounded rationality through *heterogeneous depths of reasoning*
- IBE: bounded rationality through *regret and foregone payoffs*
- Sampling: bounded rationality through *finite procedural sampling*

Four routes beyond simple Nash

Today's lecture developed four distinct ways to enrich one-shot behavioral game theory:

- CHM: bounded rationality through *heterogeneous depths of reasoning*
- IBE: bounded rationality through *regret and foregone payoffs*
- Sampling: bounded rationality through *finite procedural sampling*
- HQRE / TQRE: bounded rationality through *heterogeneous noisy equilibrium*

Common theme

All four approaches explain systematic deviations from Nash, but they locate bounded rationality in different places.

What differs across the four approaches?

Beliefs and reasoning

What differs across the four approaches?

Beliefs and reasoning

- CHM: explicit downward-looking beliefs over lower levels

What differs across the four approaches?

Beliefs and reasoning

- CHM: explicit downward-looking beliefs over lower levels
- HQRE / TQRE: explicit stochastic response to beliefs

What differs across the four approaches?

Beliefs and reasoning

- CHM: explicit downward-looking beliefs over lower levels
- HQRE / TQRE: explicit stochastic response to beliefs
- Sampling: beliefs are local and implicit

What differs across the four approaches?

Beliefs and reasoning

- CHM: explicit downward-looking beliefs over lower levels
- HQRE / TQRE: explicit stochastic response to beliefs
- Sampling: beliefs are local and implicit
- IBE: no central role for fully articulated beliefs

Noise and heterogeneity

What differs across the four approaches?

Beliefs and reasoning

- CHM: explicit downward-looking beliefs over lower levels
- HQRE / TQRE: explicit stochastic response to beliefs
- Sampling: beliefs are local and implicit
- IBE: no central role for fully articulated beliefs

Noise and heterogeneity

- HQRE / TQRE: noise is payoff-sensitive and parameterized

What differs across the four approaches?

Beliefs and reasoning

- CHM: explicit downward-looking beliefs over lower levels
- HQRE / TQRE: explicit stochastic response to beliefs
- Sampling: beliefs are local and implicit
- IBE: no central role for fully articulated beliefs

Noise and heterogeneity

- HQRE / TQRE: noise is payoff-sensitive and parameterized
- Sampling: noise is endogenous to finite samples

What differs across the four approaches?

Beliefs and reasoning

- CHM: explicit downward-looking beliefs over lower levels
- HQRE / TQRE: explicit stochastic response to beliefs
- Sampling: beliefs are local and implicit
- IBE: no central role for fully articulated beliefs

Noise and heterogeneity

- HQRE / TQRE: noise is payoff-sensitive and parameterized
- Sampling: noise is endogenous to finite samples
- IBE: behavior is driven by regret and impulse balance

What differs across the four approaches?

Beliefs and reasoning

- CHM: explicit downward-looking beliefs over lower levels
- HQRE / TQRE: explicit stochastic response to beliefs
- Sampling: beliefs are local and implicit
- IBE: no central role for fully articulated beliefs

Noise and heterogeneity

- HQRE / TQRE: noise is payoff-sensitive and parameterized
- Sampling: noise is endogenous to finite samples
- IBE: behavior is driven by regret and impulse balance
- CHM / HQRE / TQRE: heterogeneity is explicit

Similar aggregate play can arise from very different behavioral mechanisms.

What recent literature adds

More recent work sharpens the message of this lecture:

What recent literature adds

More recent work sharpens the message of this lecture:

- Külpmann and Kuzmics (2022) compare theories of one-shot play *out of treatment*: most theories have predictive power, but no theory is uniformly best.

What recent literature adds

More recent work sharpens the message of this lecture:

- Külpmann and Kuzmics (2022) compare theories of one-shot play *out of treatment*: most theories have predictive power, but no theory is uniformly best.
- Huang and Golman (2025) show that weighted combinations of behavioral models can predict strategic choice better than any individual model.

What recent literature adds

More recent work sharpens the message of this lecture:

- Külpmann and Kuzmics (2022) compare theories of one-shot play *out of treatment*: most theories have predictive power, but no theory is uniformly best.
- Huang and Golman (2025) show that weighted combinations of behavioral models can predict strategic choice better than any individual model.
- Hyndman et al. (2022) and Cooper et al. (2024) caution against reading reasoning depth too literally as a fixed subject trait across all games.

Interpretation

CHM, IBE, sampling, and HQRE are best viewed as complementary lenses on boundedly rational strategic behavior.

What you should take away from this lecture

What you should take away from this lecture

- CHM is the natural benchmark for *heterogeneous reasoning depth*.

What you should take away from this lecture

- CHM is the natural benchmark for *heterogeneous reasoning depth*.
- IBE shows that bounded rationality can be modeled through *regret and impulse balancing*, not only through beliefs.

What you should take away from this lecture

- CHM is the natural benchmark for *heterogeneous reasoning depth*.
- IBE shows that bounded rationality can be modeled through *regret and impulse balancing*, not only through beliefs.
- Sampling concepts show that stochastic play can arise from *finite procedures*, not only from ad hoc error terms.

What you should take away from this lecture

- CHM is the natural benchmark for *heterogeneous reasoning depth*.
- IBE shows that bounded rationality can be modeled through *regret and impulse balancing*, not only through beliefs.
- Sampling concepts show that stochastic play can arise from *finite procedures*, not only from ad hoc error terms.
- HQRE and TQRE reconnect these ideas to the QRE tradition by allowing *heterogeneous payoff-sensitive response*.

What you should take away from this lecture

- CHM is the natural benchmark for *heterogeneous reasoning depth*.
- IBE shows that bounded rationality can be modeled through *regret and impulse balancing*, not only through beliefs.
- Sampling concepts show that stochastic play can arise from *finite procedures*, not only from ad hoc error terms.
- HQRE and TQRE reconnect these ideas to the QRE tradition by allowing *heterogeneous payoff-sensitive response*.
- Together with earlier strategic-reasoning evidence (Breitmoser, 2012), today's lecture broadens the toolkit for analyzing one-shot play beyond simple Nash.

Looking ahead: from general models to specific game classes

Today we studied general models of one-shot play. Next, we move to environments in which the differences between these models become especially visible.

Looking ahead: from general models to specific game classes

Today we studied general models of one-shot play. Next, we move to environments in which the differences between these models become especially visible.

- In mixed-equilibrium games, the issue is not equilibrium multiplicity, but whether players randomize in the benchmark way at all.

Looking ahead: from general models to specific game classes

Today we studied general models of one-shot play. Next, we move to environments in which the differences between these models become especially visible.

- In mixed-equilibrium games, the issue is not equilibrium multiplicity, but whether players randomize in the benchmark way at all.
- In coordination games, the issue is not randomization, but which equilibrium is selected and how behavior adjusts to strategic uncertainty.

Looking ahead: from general models to specific game classes

Today we studied general models of one-shot play. Next, we move to environments in which the differences between these models become especially visible.

- In mixed-equilibrium games, the issue is not equilibrium multiplicity, but whether players randomize in the benchmark way at all.
- In coordination games, the issue is not randomization, but which equilibrium is selected and how behavior adjusts to strategic uncertainty.
- This is where the distinction between noisy response, strategic depth, procedural sampling, and regret-based adjustment becomes especially concrete.

Next lecture

We therefore turn from general one-shot models to mixed equilibrium and coordination games.

References

Generalized concepts and evidence

Core references for this lecture

The main references from today's lecture fall into four groups.

Core references for this lecture

The main references from today's lecture fall into four groups.

- Heterogeneous strategic reasoning: Camerer et al. (2004), Crawford et al. (2013)

Core references for this lecture

The main references from today's lecture fall into four groups.

- Heterogeneous strategic reasoning: Camerer et al. (2004), Crawford et al. (2013)
- Regret and stationary play: Selten and Chmura (2008), Brunner et al. (2011)

Core references for this lecture

The main references from today's lecture fall into four groups.

- Heterogeneous strategic reasoning: Camerer et al. (2004), Crawford et al. (2013)
- Regret and stationary play: Selten and Chmura (2008), Brunner et al. (2011)
- Procedural sampling: Osborne and Rubinstein (1998), Osborne and Rubinstein (2003)

Core references for this lecture

The main references from today's lecture fall into four groups.

- Heterogeneous strategic reasoning: Camerer et al. (2004), Crawford et al. (2013)
- Regret and stationary play: Selten and Chmura (2008), Brunner et al. (2011)
- Procedural sampling: Osborne and Rubinstein (1998), Osborne and Rubinstein (2003)
- Heterogeneous noisy equilibrium: McKelvey and Palfrey (1995), Rogers et al. (2009)

For the broader behavioral-game-theory perspective behind this lecture, see Crawford et al. (2013).

Further reading and how the papers fit together

Further reading and how the papers fit together

- For strategic reasoning in guessing games and the earlier course line, see Breitmoser (2012).

Further reading and how the papers fit together

- For strategic reasoning in guessing games and the earlier course line, see Breitmoser (2012).
- For recent evidence on the limits of stable reasoning levels, see Hyndman et al. (2022) and Cooper et al. (2024).

Further reading and how the papers fit together

- For strategic reasoning in guessing games and the earlier course line, see Breitmoser (2012).
- For recent evidence on the limits of stable reasoning levels, see Hyndman et al. (2022) and Cooper et al. (2024).
- For recent synthesis across one-shot models, see Külpmann and Kuzmics (2022) and Huang and Golman (2025).

Further reading and how the papers fit together

- For strategic reasoning in guessing games and the earlier course line, see Breitmoser (2012).
- For recent evidence on the limits of stable reasoning levels, see Hyndman et al. (2022) and Cooper et al. (2024).
- For recent synthesis across one-shot models, see Külpmann and Kuzmics (2022) and Huang and Golman (2025).
- For a deeper foundation of translation-invariant logit-style response, useful when comparing logit/QRE to payoff-sensitive alternatives, see Breitmoser (2021).

How to use the literature

The slides contain the main line of argument. The papers are most useful for seeing the original experimental designs, econometric implementations, and historical development of the models.

References I

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

Breitmoser, Y. (2021). An axiomatic foundation of conditional logit. *Economic Theory*, 72(1):245–261.

Brunner, C., Camerer, C. F., and Goeree, J. K. (2011). Stationary concepts for experimental 2 x 2 games: Comment. *American Economic Review*, 101(2):1029–1040.

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.

Crawford, V. P., Costa-Gomes, M. A., and Iriberri, N. (2013). Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications. *Journal of Economic Literature*, 51(1):5–62.

Huang, S. and Golman, R. (2025). The collective wisdom of behavioral game theory. *Economic Theory*, 79:341–356.

References II

Hyndman, K., Terracol, A., and Vaksmann, J. (2022). Beliefs and (in)stability in normal-form games. *Experimental Economics*, 25(4):1146–1172.

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.

Osborne, M. J. and Rubinstein, A. (1998). Games with procedurally rational players. *American Economic Review*, 88(4):834–847.

Osborne, M. J. and Rubinstein, A. (2003). Sampling equilibrium, with an application to strategic voting. *Games and Economic Behavior*, 45(2):434–441.

Rogers, B. W., Palfrey, T. R., and Camerer, C. F. (2009). Heterogeneous quantal response equilibrium and cognitive hierarchies. *Journal of Economic Theory*, 144(4):1440–1467.

Selten, R. and Chmura, T. (2008). Stationary concepts for experimental 2x2-games. *American Economic Review*, 98(3):938–966.

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

Repetition

Generalized concepts and evidence

Repetition I: Cognitive Hierarchy Model

Repetition I: Cognitive Hierarchy Model

1. **What does CHM add relative to simple level-k?**
   Answer sketch: A level-k player best responds to a distribution of lower levels, not only to level $k-1$.

Repetition I: Cognitive Hierarchy Model

1. **What does CHM add relative to simple level-k?**
   Answer sketch: A level-k player best responds to a distribution of lower levels, not only to level $k-1$.
2. **Why is CHM a model of heterogeneous strategic depth?**
   Answer sketch: The population is described by a distribution of reasoning levels, usually Poisson with mean τ (Camerer et al., 2004).

Repetition I: Cognitive Hierarchy Model

1. **What does CHM add relative to simple level-k?**
   Answer sketch: A level-k player best responds to a distribution of lower levels, not only to level $k-1$.
2. **Why is CHM a model of heterogeneous strategic depth?**
   Answer sketch: The population is described by a distribution of reasoning levels, usually Poisson with mean τ (Camerer et al., 2004).
3. **What is the role of level-0 in CHM?**
   Answer sketch: Level-0 is a modeling anchor, usually uniform randomization; it closes the hierarchy and guarantees positive probability for all actions.

Repetition I: Cognitive Hierarchy Model

1. **What does CHM add relative to simple level-k?**
   Answer sketch: A level-k player best responds to a distribution of lower levels, not only to level $k-1$.
2. **Why is CHM a model of heterogeneous strategic depth?**
   Answer sketch: The population is described by a distribution of reasoning levels, usually Poisson with mean τ (Camerer et al., 2004).
3. **What is the role of level-0 in CHM?**
   Answer sketch: Level-0 is a modeling anchor, usually uniform randomization; it closes the hierarchy and guarantees positive probability for all actions.
4. **Why should CHM not be interpreted too literally as a fixed subject type?**
   Answer sketch: Recent evidence suggests that reasoning depth is only partly stable across games (Hyndman et al., 2022; Cooper et al., 2024).

Repetition II: Impulse Balance Equilibrium

Repetition II: Impulse Balance Equilibrium

5. **What is the basic idea of IBE?**
   Answer sketch: Behavior is shaped by regret and directional impulses; in equilibrium, the impulse mass toward the two actions is balanced (Selten and Chmura, 2008).

Repetition II: Impulse Balance Equilibrium

5. **What is the basic idea of IBE?**
   Answer sketch: Behavior is shaped by regret and directional impulses; in equilibrium, the impulse mass toward the two actions is balanced (Selten and Chmura, 2008).
6. **Why does IBE use a reference point?**
   Answer sketch: The maximin payoff provides a benchmark relative to which outcomes are coded as gains and losses.

Repetition II: Impulse Balance Equilibrium

5. **What is the basic idea of IBE?**
   Answer sketch: Behavior is shaped by regret and directional impulses; in equilibrium, the impulse mass toward the two actions is balanced (Selten and Chmura, 2008).
6. **Why does IBE use a reference point?**
   Answer sketch: The maximin payoff provides a benchmark relative to which outcomes are coded as gains and losses.
7. **Why is IBE not simply Nash equilibrium in transformed utilities?**
   Answer sketch: IBE balances experienced regret flows, not expected transformed utilities directly.

Repetition II: Impulse Balance Equilibrium

5. **What is the basic idea of IBE?**
   Answer sketch: Behavior is shaped by regret and directional impulses; in equilibrium, the impulse mass toward the two actions is balanced (Selten and Chmura, 2008).
6. **Why does IBE use a reference point?**
   Answer sketch: The maximin payoff provides a benchmark relative to which outcomes are coded as gains and losses.
7. **Why is IBE not simply Nash equilibrium in transformed utilities?**
   Answer sketch: IBE balances experienced regret flows, not expected transformed utilities directly.
8. **What is the empirical message of Selten–Chmura and its reassessment?**
   Answer sketch: Nash fits worst; IBE performs well, but later reassessment shows that differences among non-Nash models are smaller than first claimed (Selten and Chmura, 2008; Brunner et al., 2011).

Repetition III: Sampling concepts and heterogeneous QRE

Repetition III: Sampling concepts and heterogeneous QRE

9. **What is payoff sampling?**
   Answer sketch: A player samples payoff consequences for her own actions and chooses the action that performs better in the sample (Osborne and Rubinstein, 1998).

Repetition III: Sampling concepts and heterogeneous QRE

9. **What is payoff sampling?**
   Answer sketch: A player samples payoff consequences for her own actions and chooses the action that performs better in the sample (Osborne and Rubinstein, 1998).

10. **What is action sampling?**
    Answer sketch: A player samples a finite set of opponents' actions, treats the empirical frequency as informative, and best responds to that sample (Osborne and Rubinstein, 2003).

Repetition III: Sampling concepts and heterogeneous QRE

9. **What is payoff sampling?**
   Answer sketch: A player samples payoff consequences for her own actions and chooses the action that performs better in the sample (Osborne and Rubinstein, 1998).

10. **What is action sampling?**
    Answer sketch: A player samples a finite set of opponents' actions, treats the empirical frequency as informative, and best responds to that sample (Osborne and Rubinstein, 2003).

11. **How do sampling concepts differ from QRE?**
    Answer sketch: In sampling models, randomness comes from finite samples; in QRE, it comes from a parametric stochastic response rule.

Repetition III: Sampling concepts and heterogeneous QRE

9. **What is payoff sampling?**
   Answer sketch: A player samples payoff consequences for her own actions and chooses the action that performs better in the sample (Osborne and Rubinstein, 1998).

10. **What is action sampling?**
    Answer sketch: A player samples a finite set of opponents' actions, treats the empirical frequency as informative, and best responds to that sample (Osborne and Rubinstein, 2003).

11. **How do sampling concepts differ from QRE?**
    Answer sketch: In sampling models, randomness comes from finite samples; in QRE, it comes from a parametric stochastic response rule.

12. **What does HQRE add relative to standard QRE?**
    Answer sketch: Players can differ in payoff sensitivity, so precision is heterogeneous (Rogers et al., 2009).

Repetition III: Sampling concepts and heterogeneous QRE

9. **What is payoff sampling?**
   Answer sketch: A player samples payoff consequences for her own actions and chooses the action that performs better in the sample (Osborne and Rubinstein, 1998).

10. **What is action sampling?**
    Answer sketch: A player samples a finite set of opponents' actions, treats the empirical frequency as informative, and best responds to that sample (Osborne and Rubinstein, 2003).

11. **How do sampling concepts differ from QRE?**
    Answer sketch: In sampling models, randomness comes from finite samples; in QRE, it comes from a parametric stochastic response rule.

12. **What does HQRE add relative to standard QRE?**
    Answer sketch: Players can differ in payoff sensitivity, so precision is heterogeneous (Rogers et al., 2009).

13. **What does TQRE add relative to HQRE?**
    Answer sketch: More sophisticated players may still have downward-looking beliefs about others' precision; TQRE therefore connects HQRE back to CHM (Rogers et al., 2009).

Repetition IV: Synthesis

Repetition IV: Synthesis

14. **Which four routes beyond simple Nash did we study today?**
Answer sketch: CHM (heterogeneous depth), IBE (regret / impulses), sampling concepts (finite procedures), HQRE / TQRE (heterogeneous noisy equilibrium).

Repetition IV: Synthesis

14. **Which four routes beyond simple Nash did we study today?**
Answer sketch: CHM (heterogeneous depth), IBE (regret / impulses), sampling concepts (finite procedures), HQRE / TQRE (heterogeneous noisy equilibrium).

15. **Which of today's models are most belief-based?**
Answer sketch: CHM and HQRE / TQRE are most explicitly belief-based; sampling concepts use only local implicit beliefs, and IBE is less belief-centered.

Repetition IV: Synthesis

14. **Which four routes beyond simple Nash did we study today?**
Answer sketch: CHM (heterogeneous depth), IBE (regret / impulses), sampling concepts (finite procedures), HQRE / TQRE (heterogeneous noisy equilibrium).

15. **Which of today's models are most belief-based?**
Answer sketch: CHM and HQRE / TQRE are most explicitly belief-based; sampling concepts use only local implicit beliefs, and IBE is less belief-centered.

16. **Which models generate noise endogenously?**
Answer sketch: Sampling concepts generate noise from finite samples; IBE from regret balancing; HQRE / TQRE instead use explicit stochastic choice.

Repetition IV: Synthesis

14. **Which four routes beyond simple Nash did we study today?**
Answer sketch: CHM (heterogeneous depth), IBE (regret / impulses), sampling concepts (finite procedures), HQRE / TQRE (heterogeneous noisy equilibrium).

15. **Which of today's models are most belief-based?**
Answer sketch: CHM and HQRE / TQRE are most explicitly belief-based; sampling concepts use only local implicit beliefs, and IBE is less belief-centered.

16. **Which models generate noise endogenously?**
Answer sketch: Sampling concepts generate noise from finite samples; IBE from regret balancing; HQRE / TQRE instead use explicit stochastic choice.

17. **What is the main empirical lesson from the recent literature?**
Answer sketch: No one-shot model is uniformly best across games, and combinations of models can outperform single-model champions (Külpmann and Kuzmics, 2022; Huang and Golman, 2025).

Repetition IV: Synthesis

14. **Which four routes beyond simple Nash did we study today?**
Answer sketch: CHM (heterogeneous depth), IBE (regret / impulses), sampling concepts (finite procedures), HQRE / TQRE (heterogeneous noisy equilibrium).

15. **Which of today's models are most belief-based?**
Answer sketch: CHM and HQRE / TQRE are most explicitly belief-based; sampling concepts use only local implicit beliefs, and IBE is less belief-centered.

16. **Which models generate noise endogenously?**
Answer sketch: Sampling concepts generate noise from finite samples; IBE from regret balancing; HQRE / TQRE instead use explicit stochastic choice.

17. **What is the main empirical lesson from the recent literature?**
Answer sketch: No one-shot model is uniformly best across games, and combinations of models can outperform single-model champions (Külpmann and Kuzmics, 2022; Huang and Golman, 2025).

18. **What should you now be able to explain after this lecture?**
Answer sketch: Why observed deviations from Nash can reflect different behavioral mechanisms, and why model comparison matters.

Heterogeneous QRE (Rogers et al., 2009)

Heterogeneous logit equilibrium (HLE, HQRE)

Player $i \in N$ has precision λ_i, which is private information but its density $f_i(\lambda_i)$ is common knowledge. An HLE σ is a logit equilibrium of this Bayesian game. The probability that type λ_i of player $i \in N$ plays action $a_i \in A_i$ is

$$p_i(a_i|\lambda_i) = \frac{\exp\{\lambda_i\, U_i(a_i|\sigma)\}}{\sum_{a_i'} \exp\{\lambda_i\, U_i(a_i'|\sigma)\}}.$$

The *induced mixed strategy* of $i \in N$, i.e. i's ex-ante strategy, is

$$\sigma_i(a_i) = \int_0^\infty p_i(a_i|\lambda)\, f_i(\lambda)\, d\lambda.$$

Given the induced mixed strategy profile $\sigma = (\sigma_i)$, i's expected utility is

$$U_i(a_i|\sigma) = \sum_{a_{-i} \in A_{-i}} \left(\prod_{j \neq i} \sigma_j(a_j) \right) u_i(a_i, a_{-i})$$

Note HLE contains ALE and LE as (limiting) special cases

Matrix game estimates

Game	Description	Random neg Log L	Empirical $-\log L$	CH		QRE		TQRE			HQRE		
				τ	$-\log L$	λ	$-\log L$	Y	τ	$-\log L$	Λ	ε	$-\log L$
1	Unprofitable	81.30	73.12	10.19	73.65	0.06	75.97	0.35	3.73	73.12	0.07	0.13	75.96
2	Cloned MP	81.30	52.34	0.90	58.39	0.12	71.72	∞	0.90	58.39	0.12	0	71.72
3	Cloned SH (Low)	51.96	46.50	0.44	46.72	0.15	46.74	0.50	0.62	46.67	0.15	0.3	46.73
4	Cloned SH (Hi)	51.96	40.18	1.00	41.91	0.18	41.88	0.89	1.40	40.65	0.19	0.29	41.88
5	SW 1	81.30	39.55	3.10	40.92	0.13	41.51	0.17	2.66	39.93	0.15	0.23	40.51
6	SW 2	81.30	46.00	5.59	46.03	0.13	46.75	∞	5.59	46.03	0.13	0	46.75
7	SW 3	81.30	53.26	0.72	63.26	0.07	73.94	∞	0.72	63.26	0.08	0.16	73.63
8	SW 4	81.30	55.02	1.96	56.11	∞	56.40	∞	1.96	56.11	∞	–	56.40
9	SW 5	81.30	79.27	0.00	81.30	0.14	80.50	0.00	0.00	81.30	0.15	0.3	80.40
10	SW 6	81.30	79.54	1.27	79.69	0.05	80.86	0.11	2.81	79.54	0.05	0.1	80.85
11	SW 7	81.30	73.69	0.67	73.73	∞	73.73	∞	0.67	73.73	∞	–	73.73
12	SW 8	81.30	70.76	0.00	81.30	0.98	73.25	0.02	51.30	78.30	1	1.5	73.23
13	SW 9	81.30	54.30	1.01	60.94	0.08	66.56	∞	1.01	60.94	0.08	0	66.56
14	SW 10	81.30	73.12	13.90	77.35	0.42	73.17	0.15	7.13	73.31	0.44	0.88	73.15
15	SW 11	81.30	73.81	0.83	73.82	0.17	74.55	0.17	0.93	73.81	0.18	0.36	74.52
16	SW 12	81.30	66.81	0.94	67.02	0.03	72.94	0.16	0.87	66.81	0.04	0.07	72.70
17	Cloned Joker	86.88	77.28	0.83	81.08	0.16	83.62	∞	0.82	81.07	0.17	0	83.62
	Sum	1329.00	1054.55	–	1103.21	–	1134.08	–	–	1092.96	–	–	1132.34
	Pooled	–	–	0.53	1239.26	0.10	1192.00	0.11	1.57	1164.41	0.12	0.24	1182.95

HQRE $\lambda \sim \text{Unif}[\Lambda - \varepsilon/2, \Lambda + \varepsilon/2]$. TQRE $\lambda = k\,\gamma$ where k has Poisson distribution $f(k) = \tau^k e^{-k}/k!$ and all players are self-limited $\theta_i(\lambda_i) = \lambda_i$

Matrix game estimates													
Game	Description	Random neg Log L	Empirical $-\log L$	CH		QRE		TQRE			HQRE		
				τ	$-\log L$	λ	$-\log L$	Y	τ	$-\log L$	Λ	ε	$-\log L$
1	Unprofitable	81.30	73.12	10.19	73.65	0.06	75.97	0.35	3.73	73.12	0.07	0.13	75.96
2	Cloned MP	81.30	52.34	0.90	58.39	0.12	71.72	∞	0.90	58.39	0.12	0	71.72
3	Cloned SH (Low)	51.96	46.50	0.44	46.72	0.15	46.74	0.50	0.62	46.67	0.15	0.3	46.73
4	Cloned SH (Hi)	51.96	40.18	1.00	41.91	0.18	41.88	0.89	1.40	40.65	0.19	0.29	41.88
5	SW 1	81.30	39.55	3.10	40.92	0.13	41.51	0.17	2.66	39.93	0.15	0.23	40.51
6	SW 2	81.30	46.00	5.59	46.03	0.13	46.75	∞	5.59	46.03	0.13	0	46.75
7	SW 3	81.30	53.26	0.72	63.26	0.07	73.94	∞	0.72	63.26	0.08	0.16	73.63
8	SW 4	81.30	55.02	1.96	56.11	∞	56.40	∞	1.96	56.11	∞	–	56.40
9	SW 5	81.30	79.27	0.00	81.30	0.14	80.50	0.00	0.00	81.30	0.15	0.3	80.40
10	SW 6	81.30	79.54	1.27	79.69	0.05	80.86	0.11	2.81	79.54	0.05	0.1	80.85
11	SW 7	81.30	73.69	0.67	73.73	∞	73.73	∞	0.67	73.73	∞	–	73.73
12	SW 8	81.30	70.76	0.00	81.30	0.98	73.25	0.02	51.30	78.30	1	1.5	73.23
13	SW 9	81.30	54.30	1.01	60.94	0.08	66.56	∞	1.01	60.94	0.08	0	66.56
14	SW 10	81.30	73.12	13.90	77.35	0.42	73.17	0.15	7.13	73.31	0.44	0.88	73.15
15	SW 11	81.30	73.81	0.83	73.82	0.17	74.55	0.17	0.93	73.81	0.18	0.36	74.52
16	SW 12	81.30	66.81	0.94	67.02	0.03	72.94	0.16	0.87	66.81	0.04	0.07	72.70
17	Cloned Joker	86.88	77.28	0.83	81.08	0.16	83.62	∞	0.82	81.07	0.17	0	83.62
	Sum	1329.00	1054.55	–	1103.21	–	1134.08	–	–	1092.96	–	–	1132.34
	Pooled	–	–	0.53	1239.26	0.10	1192.00	0.11	1.57	1164.41	0.12	0.24	1182.95

HQRE $\lambda \sim \text{Unif}[\Lambda - \varepsilon/2, \Lambda + \varepsilon/2]$. TQRE $\lambda = k\,\gamma$ where k has Poisson distribution $f(k) = \tau^k e^{-k}/k!$ and all players are self-limited $\theta_i(\lambda_i) = \lambda_i$

- TQRE fits well "Sum" and "Pooled"; CHM fits well "Sum", not "Pooled" (Goodness of fit of CHM seems to stem from flexibility, it is not robust)