ECN 119: Psychology and Economics
Games

Jim Campbell

UC Berkeley

Summer 2021

Games

How do people behave in strategic situations? When my payoff depends not only on my choice but also on other peoples' choices, I have to consider how they think... and how they think I think, and so on. This is the object of study in the field of game theory.

Here we will look at some ways in which people systematically deviate from the 'standard' predictions of game theory.

Games

In this section:

1. Games and strategies
2. Dominance
3. Nash equilibrium
4. Mixed strategies and reaction functions
5. Level k reasoning
6. Backward induction
7. Subgame perfect Nash equilibrium
8. Repeated games in finite and infinite time
9. Community enforcement

Games

The object of interest is a **game**. The components of a game are:

1. **Players**: a list of all relevant players.
2. **Moves**: what each player can do.
3. **Information**: what each player knows when they take an action.
4. **Payoffs**: payoffs (vNM utilities) for each player as a function of every player's actions.

Games

The object of interest is a **game**. The components of a game are:

1. **Players**: a list of all relevant players.
2. **Moves**: what each player can do.
3. **Information**: what each player knows when they take an action.
4. **Payoffs**: payoffs (vNM utilities) for each player as a function of every player's actions.

The game is our model, so we would like to analyze games that capture some interesting features of the (economic) problem. Given this, our objects of interest:

- **Strategy**: a complete contingent plan of action for each player.
- **Strategy profile**: a collection of strategies for all players.
- **Solution concept**: a strategy profile that is a "likely" way to play the game.

Strategic form

Until further notice, we're working with games of **complete information**.

- The **normal form** or **strategic form** of a game is
    1. the set of players: N
    2. the pure-strategy space (the set of strategies for each player i): S_i
    3. a payoff function for each player i: $u_i : S_1 \times \ldots \times S_n \to \mathbb{R}$

Strategic form

Until further notice, we're working with games of **complete information**.

- The **normal form** or **strategic form** of a game is
  1. the set of players: N
  2. the pure-strategy space (the set of strategies for each player i): S_i
  3. a payoff function for each player i: $u_i : S_1 \times ... \times S_n \rightarrow \mathbb{R}$
- Payoffs are numbers that represent the players' preferences over the set of possible outcomes
- But it's important to remember that they do *not* have to be the same as monetary payoff

Underdetermination of models

"If there is one hypothesis that is consistent with the available evidence, there are always an infinite number that are." - Friedman

- See also Borges "The Library of Babel"

Strategic form

- For finite two-player games, we can represent this as a matrix.

	L	R
T	$2,2$	$2,0$
B	$3,0$	$0,9$

- Player 1 is the "row" player; her pure strategies are T and B.
- Player 2 is the "column" player; his pure strategies are L and R.
- Each cell is a pure strategy profile: e.g. if 1 plays T and 2 plays R, we end up in the top-right.
- The numbers are payoffs: at (T, R) player 1 gets 2 and player 2 gets 0.

Mixed strategies

A **pure strategy** is a strategy in which a player plays actions for sure. A **mixed strategy** is a probability distribution over pure strategies.

Mixed strategies

A **pure strategy** is a strategy in which a player plays actions for sure. A **mixed strategy** is a probability distribution over pure strategies.

	L	R
T	$2, 2$	$2, 0$
B	$3, 0$	$0, 9$

For example, player 1 could play T with probability $\frac{3}{4}$ and B with probability $\frac{1}{4}$.

Mixed strategies

A **pure strategy** is a strategy in which a player plays actions for sure. A **mixed strategy** is a probability distribution over pure strategies.

	L	R
T	$2,2$	$2,0$
B	$3,0$	$0,9$

For example, player 1 could play T with probability $\frac{3}{4}$ and B with probability $\frac{1}{4}$.

- Randomization is independent of opponents'.
- Payoffs are expected values of corresponding pure-strategy payoffs (invoking vNM utility).
- A pure strategy is just a degenerate mixed strategy.

Best responses

A strategy for i is a **best response** to the strategies of the other players if there exists no other strategy for i that yields a higher payoff.

Best responses

A strategy for i is a **best response** to the strategies of the other players if there exists no other strategy for i that yields a higher payoff.

	Swerve	*Don't*
Swerve	$0, 0$	$-4, 4$
Don't	$4, -2$	$-10, -10$

In Chicken, if my opponent plays “swerve”, my best response is not to swerve. If my opponent plays “don't”, my best response is to swerve.

Best responses

A strategy for i is a **best response** to the strategies of the other players if there exists no other strategy for i that yields a higher payoff.

	Swerve	*Don't*
Swerve	$0, 0$	$-4, 4$
Don't	$4, -2$	$-10, -10$

In Chicken, if my opponent plays "swerve", my best response is not to swerve. If my opponent plays "don't", my best response is to swerve.

- Best responses are to a *fixed* strategy profile for the other players.
- A strategy that is strictly dominated is never a best response.

Nash equilibrium

A foundational solution concept in game theory is **Nash equilibrium**.

A strategy profile constitutes a Nash equilibrium of a game if every player's strategy is a best response to the strategies played by the other players.

Nash equilibrium

A foundational solution concept in game theory is **Nash equilibrium**.

A strategy profile constitutes a Nash equilibrium of a game if every player's strategy is a best response to the strategies played by the other players.

	C	D
C	$-1, -1$	$-10, 0$
D	$0, -10$	$-5, -5$

In the Prisoner's Dilemma, (D, D) is a Nash equilibrium.

Nash equilibrium

A foundational solution concept in game theory is **Nash equilibrium**.

A strategy profile constitutes a Nash equilibrium of a game if every player's strategy is a best response to the strategies played by the other players.

	C	D
C	$-1, -1$	$-10, 0$
D	$0, -10$	$-5, -5$

In the Prisoner's Dilemma, (D, D) is a Nash equilibrium.

		Player 2	
		L	R
Player 1	U	$4, 5$	$1, 0$
	D	$0, 2$	$3, 1$

What about this one?

Dominated strategies

By the way:

- A strategy for player i is **strictly dominated** if there exists another strategy for i that yields a greater payoff for all strategy profiles of the other player(s).
- A strategy for player i is **weakly dominated** if there exists another strategy for I that yields a payoff at least as high for all strategy profiles of the other player(s) and a greater payoff for at least one profile.

Dominated strategies

By the way:

- A strategy for player i is **strictly dominated** if there exists another strategy for i that yields a greater payoff for all strategy profiles of the other player(s).
- A strategy for player i is **weakly dominated** if there exists another strategy for I that yields a payoff at least as high for all strategy profiles of the other player(s) and a greater payoff for at least one profile.

	C	D
C	$-1, -1$	$-10, 0$
D	$0, -10$	$-5, -5$

The Prisoner's Dilemma is a bit boring as a one-shot game because it's solvable in one step by elimination of strictly dominated strategies

Dominated strategies

By the way:

- A strategy for player i is **strictly dominated** if there exists another strategy for i that yields a greater payoff for all strategy profiles of the other player(s).
- A strategy for player i is **weakly dominated** if there exists another strategy for I that yields a payoff at least as high for all strategy profiles of the other player(s) and a greater payoff for at least one profile.

	C	D
C	$-1, -1$	$-10, 0$
D	$0, -10$	$-5, -5$

The Prisoner's Dilemma is a bit boring as a one-shot game because it's solvable in one step by elimination of strictly dominated strategies

- C is strictly dominated by D

Nash equilibrium

One more:

		Player 2		
		L	C	R
	U	$3, 0$	$1, 2$	$2, 5$
Player 1	M	$1, 1$	$0, 0$	$1, 2$
	D	$0, 3$	$2, 4$	$3, 2$

Nash equilibrium

One more:

Player 2

Player 1		L	C	R
	U	3, 0	1, 2	2, 5
	M	1, 1	0, 0	1, 2
	D	0, 3	2, 4	3, 2

- Nash equilibrium has a flavor of fulfilled expectations: best responses *given* others' strategies.
- Or no regret: if I could go back and change my strategy but everyone else had to stay the same, could I have done better?
- Again it's recursive: your action is a best response to my action which is a best response to your action which is a best response to my action...

Nash equilibrium

One more:

		Player 2 L	Player 2 C	Player 2 R
Player 1	U	$3, 0$	$1, 2$	$2, 5$
Player 1	M	$1, 1$	$0, 0$	$1, 2$
Player 1	D	$0, 3$	$2, 4$	$3, 2$

- Nash equilibrium has a flavor of fulfilled expectations: best responses *given* others' strategies.
- Or no regret: if I could go back and change my strategy but everyone else had to stay the same, could I have done better?
- Again it's recursive: your action is a best response to my action which is a best response to your action which is a best response to my action...
- What must the players be thinking?

Meeting in New York

Consider this *coordination* game:

	A	B
A	$1, 1$	$0, 0$
B	$0, 0$	$1, 1$

Meeting in New York

Consider this *coordination* game:

	A	B
A	$1, 1$	$0, 0$
B	$0, 0$	$1, 1$

There are Nash equilibria at (A, A) and (B, B).

Meeting in New York

Consider this *coordination* game:

	A	B
A	$1, 1$	$0, 0$
B	$0, 0$	$1, 1$

There are Nash equilibria at (A, A) and (B, B).

- There is another Nash equilibrium. Where?

Meeting in New York

	A	B
A	$1, 1$	$0, 0$
B	$0, 0$	$1, 1$

Say player 2 plays a mixed strategy: A with probability p and B with probability $(1 - p)$.

Meeting in New York

	A	B
A	$1, 1$	$0, 0$
B	$0, 0$	$1, 1$

Say player 2 plays a mixed strategy: A with probability p and B with probability $(1 - p)$. If player 1 plays A:

$$u_1(A) = pu_1(A, A) + (1 - p)u_1(A, B) \tag{1}$$

$$= p \tag{2}$$

Meeting in New York

	A	B
A	$1, 1$	$0, 0$
B	$0, 0$	$1, 1$

Say player 2 plays a mixed strategy: A with probability p and B with probability $(1 - p)$. If player 1 plays A:

$$u_1(A) = pu_1(A, A) + (1 - p)u_1(A, B) \quad (1)$$
$$= p \quad (2)$$

If player 2 plays B:

$$u_1(B) = pu_1(B, A) + (1 - p)u_1(B, B) \quad (3)$$
$$= (1 - p) \quad (4)$$

Meeting in New York

	A	B
A	$1, 1$	$0, 0$
B	$0, 0$	$1, 1$

Say player 2 plays a mixed strategy: A with probability p and B with probability $(1 - p)$. If player 1 plays A:

$$u_1(A) = pu_1(A, A) + (1 - p)u_1(A, B) \quad (1)$$
$$= p \quad (2)$$

If player 2 plays B:

$$u_1(B) = pu_1(B, A) + (1 - p)u_1(B, B) \quad (3)$$
$$= (1 - p) \quad (4)$$

If $p > 1 - p$, 1's best response is A; if not, B.

Meeting in New York

	A	B
A	$1, 1$	$0, 0$
B	$0, 0$	$1, 1$

But if player 2 plays precisely $p = 1 - p$ (that is, $p = \frac{1}{2}$) then player 1 is indifferent between playing A and B, or *any mix of the two*.

Meeting in New York

	A	B
A	$1, 1$	$0, 0$
B	$0, 0$	$1, 1$

But if player 2 plays precisely $p = 1 - p$ (that is, $p = \frac{1}{2}$) then player 1 is indifferent between playing A and B, or *any mix of the two*.

- Any of their strategies, pure or mixed, is a best response!
- In particular, one best response has player 1 play a mixed strategy of half probability each across A and B
- Since similar logic holds for player 2, there is a Nash equilibrium in which both players randomize over the two actions with equal probability

Meeting in New York

	A	B
A	$1,1$	$0,0$
B	$0,0$	$1,1$

But if player 2 plays precisely $p = 1 - p$ (that is, $p = \frac{1}{2}$) then player 1 is indifferent between playing A and B, or *any mix of the two*.

- Any of their strategies, pure or mixed, is a best response!
- In particular, one best response has player 1 play a mixed strategy of half probability each across A and B
- Since similar logic holds for player 2, there is a Nash equilibrium in which both players randomize over the two actions with equal probability

So this game has three Nash equilibria: (A, A), (B, B) and the mixed strategy equilibrium with each player choosing A with probability $\frac{1}{2}$.

Nash equilibria in Chicken

Chicken has Nash equilibria in pure strategies at the off-diagonals.

	Swerve	*Don′t*
Swerve	$0, 0$	$-4, 4$
Don′t	$4, -2$	$-10, -10$

Nash equilibria in Chicken

Chicken has Nash equilibria in pure strategies at the off-diagonals.

	Swerve	*Don't*
Swerve	$0, 0$	$-4, 4$
Don't	$4, -2$	$-10, -10$

Is there one in mixed strategies? Say 2 randomizes with probability p on "swerve". Player 1 is indifferent between "swerve" and "don't" if:

$$U(swerve) = U(don't) \tag{5}$$

$$0p + (-4)(1-p) = 4p + (-10)(1-p) \tag{6}$$

$$p = \frac{3}{5} \tag{7}$$

Nash equilibria in Chicken

	Swerve	*Don't*
Swerve	$0, 0$	$-4, 4$
Don't	$4, -2$	$-10, -10$

Nash equilibria in Chicken

	Swerve	*Don't*
Swerve	$0, 0$	$-4, 4$
Don't	$4, -2$	$-10, -10$

Now say 1 randomizes with probability q on "swerve". Player 2 is indifferent between "swerve" and "don't" if:

$$0q + (-2)(1-q) = 4q + (-10)(1-q) \quad (8)$$

$$q = \frac{2}{3} \quad (9)$$

Nash equilibria in Chicken

	Swerve	*Don't*
Swerve	$0, 0$	$-4, 4$
Don't	$4, -2$	$-10, -10$

Now say 1 randomizes with probability q on "swerve". Player 2 is indifferent between "swerve" and "don't" if:

$$0q + (-2)(1-q) = 4q + (-10)(1-q) \tag{8}$$

$$q = \frac{2}{3} \tag{9}$$

So if

- Player 1 is playing a mixed strategy with probability $\frac{2}{3}$ on "swerve", and
- Player 2 is playing a mixed strategy with probability $\frac{3}{5}$ on "swerve",

both players are best responding to the other: Nash equilibrium.

Nash equilibria in Chicken

	Swerve	*Don′t*
Swerve	$0, 0$	$-4, 4$
Don′t	$4, -2$	$-10, -10$

Now say 1 randomizes with probability q on "swerve". Player 2 is indifferent between "swerve" and "don't" if:

$$0q + (-2)(1 - q) = 4q + (-10)(1 - q) \quad (8)$$

$$q = \frac{2}{3} \quad (9)$$

So if

- Player 1 is playing a mixed strategy with probability $\frac{2}{3}$ on "swerve", and
- Player 2 is playing a mixed strategy with probability $\frac{3}{5}$ on "swerve",

both players are best responding to the other: Nash equilibrium.

- Note that the probability of, for example, a crash (i.e. neither swerves) is $\frac{1}{3} * \frac{2}{5} = \frac{2}{15}$.

Reaction functions

One way to see Nash equilibria in mixed strategies in 2x2 games graphically is with **reaction functions**.

- These graph each player's best response as a function of the probability the other player is putting on one of her two actions.
- Where reaction functions intersect, we have a Nash equilibrium.
- We'll draw reaction functions for Chicken

Reaction functions

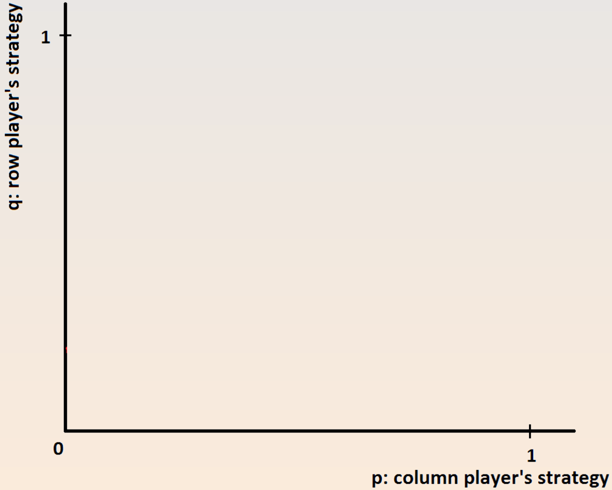

Figure: The space for drawing reaction functions

Reaction functions

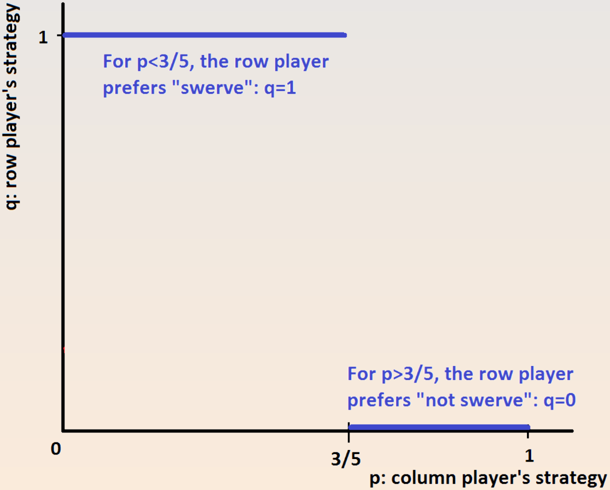

Figure: Constructing the row player's best response

Reaction functions

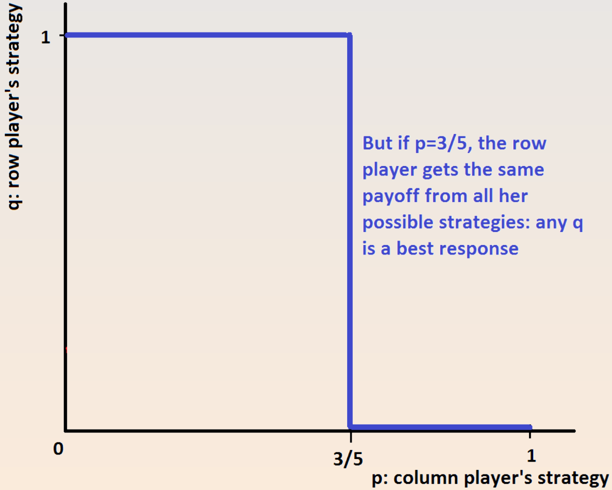

Figure: Constructing the row player's best response

Reaction functions

Figure: Constructing the column player's best response

Reaction functions

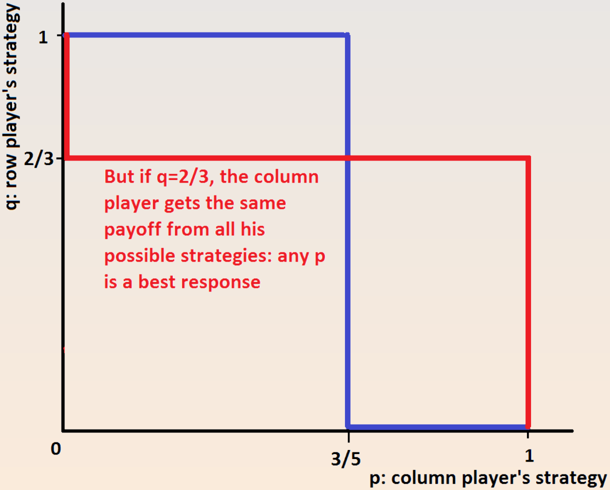

Figure: Constructing the column player's best response

Reaction functions

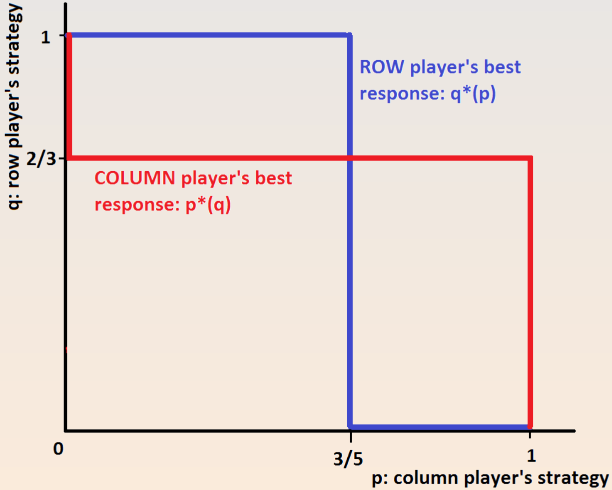

Figure: The finished product

Reaction functions

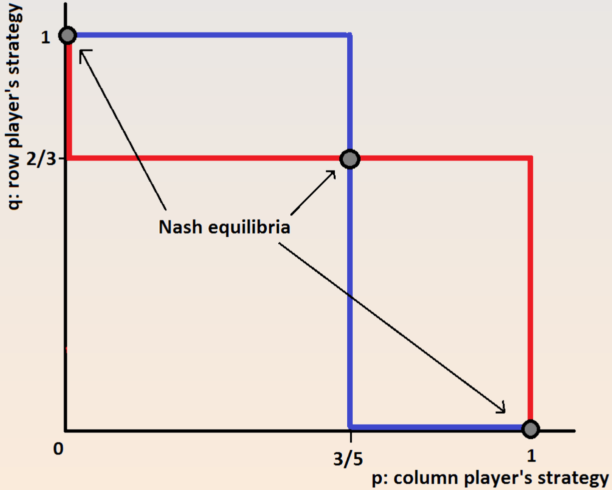

Figure: Intersections are Nash equilibria

Reaction functions

We can start to see graphically the fixed point concept of Nash equilibrium

- Very loosely speaking, this is the logic of the foundational result that Nash equilibria exist in every finite strategic-form game
- In "almost all" finite strategic forms there are a finite and odd number of Nash equilibria (Wilson 1971)
- (But there are games with an even or infinite number of Nash equilibria)

Mixed strategies

How can we interpret a mixed-strategy Nash equilibrium?

Mixed strategies

How can we interpret a mixed-strategy Nash equilibrium?

- Do people randomize?

Mixed strategies

How can we interpret a mixed-strategy Nash equilibrium?

- Do people randomize?
- Why mix? All pure strategies played with positive probability in the mixed-strategy NE yield the same payoff as the mixed strategy.

Mixed strategies

How can we interpret a mixed-strategy Nash equilibrium?

- Do people randomize?
- Why mix? All pure strategies played with positive probability in the mixed-strategy NE yield the same payoff as the mixed strategy.
- Population frequencies: evolutionary game theory.

Mixed strategies

How can we interpret a mixed-strategy Nash equilibrium?

- Do people randomize?
- Why mix? All pure strategies played with positive probability in the mixed-strategy NE yield the same payoff as the mixed strategy.
- Population frequencies: evolutionary game theory.
- 'Mixing' could be an artifact of unobservable decision process by the player. But is this a problem? We will see more along this line later in games with 'types' of players.

Mixed strategies

How can we interpret a mixed-strategy Nash equilibrium?

- Do people randomize?
- Why mix? All pure strategies played with positive probability in the mixed-strategy NE yield the same payoff as the mixed strategy.
- Population frequencies: evolutionary game theory.
- 'Mixing' could be an artifact of unobservable decision process by the player. But is this a problem? We will see more along this line later in games with 'types' of players.

Evidence from sports has been used to 'test' whether players actually mix optimally: Palacios-Huerta (1996) uses penalty kick data, Kovash & Levitt (2009) NFL run/pass data.

Matching pennies

Let's work through the steps to find the Nash equilibrium for this game:

		Player 2: *Heads*	Player 2: *Tails*
Player 1	*Heads*	$1, -1$	$-1, 1$
	Tails	$-1, 1$	$1, -1$

Matching pennies

Let's work through the steps to find the Nash equilibrium for this game:

		Player 2: *Heads*	Player 2: *Tails*
Player 1	*Heads*	$1, -1$	$-1, 1$
	Tails	$-1, 1$	$1, -1$

- In zero-sum games like this one equilibrium strategies require unpredictability
- Maximum unpredictability means making your opponent's like as difficult as possible—a *minimax* strategy means minimizing the maximum payoff your opponent can earn

Matching pennies

Let's work through the steps to find the Nash equilibrium for this game:

		Player 2	
		Heads	*Tails*
Player 1	*Heads*	$1, -1$	$-1, 1$
	Tails	$-1, 1$	$1, -1$

- In zero-sum games like this one equilibrium strategies require unpredictability
- Maximum unpredictability means making your opponent's like as difficult as possible—a *minimax* strategy means minimizing the maximum payoff your opponent can earn
- Think of how to play rock-paper-scissors optimally...
- The issue is that people are in general not that good at randomizing (we will explore this in more detail in the Beliefs topic)

Mixed strategy Nash in the wild

Palacios-Huerta (2003): test whether highly motivated professionals play according to John von Neumann's Minimax theorem and mixed strategy Nash equilibrium

- Setting: penalty kicks in soccer
  - Can you describe why this might work as an empirical setting for this question?
  - Can you think of any other settings that might work? Or why other settings might *not* work?

Mixed strategy Nash in the wild

Palacios-Huerta (2003): test whether highly motivated professionals play according to John von Neumann's Minimax theorem and mixed strategy Nash equilibrium

- Setting: penalty kicks in soccer
  - Can you describe why this might work as an empirical setting for this question?
  - Can you think of any other settings that might work? Or why other settings might *not* work?
- 1417 penalty kicks taken in professional games from 1995-2000, largely from Spain, Italy, and England
- Observe choices made by individual goalkeepers and kickers over time (i.e. across many penalties)

Palacios-Huerta (2003)

In what follows, we let the player's payoffs be the probabilities of success ("score" for the kicker and "no score" for the goalkeeper) in the penalty kick. The kicker wishes to maximize the expected probability of scoring, while the goalkeeper wishes to minimize it. Consider, for example, a simple 2×2 game-theoretical model of player's actions for the penalty kick and let π_{ij} denote the kicker's probabilities of scoring, where $i = \{L, R\}$ denotes the kicker's choice and $j = \{L, R\}$ the goalkeeper's choice, with L = left, R = right:

$i \backslash j$	L	R
L	π_{LL}	π_{LR}
R	π_{RL}	π_{RR}

This game has a unique Nash equilibrium when

$$\pi_{LR} > \pi_{LL} < \pi_{RL},$$
$$\pi_{RL} > \pi_{RR} < \pi_{LR}.$$

Palacios-Huerta (2003)

Before we begin any formal test, it is worth examining the extent to which observed behaviour is close to the Nash equilibrium predictions. For all players in the sample the empirical scoring probabilities are

	g_L	$1-g_L$
k_L	58·30	94·97
$1-k_L$	92·91	69·92

where, as indicated above, k_L and g_L denote the non-natural sides. The mixed strategy Nash equilibrium predicted frequencies for these empirical values and the actual mixing probabilities observed in the sample are

	g_L (%)	$1-g_L$ (%)	k_L (%)	$1-k_L$ (%)
Nash predicted frequencies	41·99	58·01	38·54	61·46
Actual frequencies	42·31	57·69	39·98	60·02

- Statistical tests of the data suggest (i) winning probabilities are identical across strategies, (ii) players' choices are independent draws from a random process, both consistent with minimax/MSNE
- This means choices don't depend on previous choices, opponents' previous choices, or past outcomes

Thinking about thinking about strategies

How do people think about how other people think? Why does it matter?

- Nash equilibrium and iterated dominance solution concepts are ‘complicated’ in a couple of important ways
- First, in many interesting settings it might be difficult for a player to compute which strategies are part of an equilibrium
- Second, some solution concepts require recursive reasoning on players’ rationality
- Applying game theory is abstract modeling, so we aren’t necessarily trying to capture behavior literally, but...
- Is there a way to account for players’ thought processes?
- Would including this in a solution concept improve our ability to predict the outcome of games?

Thinking deeply

Keynes, from the *General Theory*, comparing the stock market to a beauty contest prediction game:

> *"It is not the case of choosing those which, to the best of one's judgment, are really the prettiest, nor even those which average opinion genuinely thinks the prettiest. We have reached the third degree, where we devote our intelligences to anticipating what average opinion expects the average opinion to be. And there are some, I believe, who practice the fourth, fifth, and higher degrees."*

Level-k reasoning

One model of a player's thought process is **level-k reasoning**.

Level-k reasoning

One model of a player's thought process is **level-k reasoning**.

- **L0**: level 0 is an 'unsophisticated' player who chooses from his available strategies by some heuristic.

Level-k reasoning

One model of a player's thought process is **level-k reasoning**.

- **L0**: level 0 is an 'unsophisticated' player who chooses from his available strategies by some heuristic.
  - Uniform distribution over available pure strategies (picking at random)
  - Focal strategy

Level-k reasoning

One model of a player's thought process is **level-k reasoning**.

- **L0**: level 0 is an 'unsophisticated' player who chooses from his available strategies by some heuristic.
  - Uniform distribution over available pure strategies (picking at random)
  - Focal strategy
- **L1**: a level 1 strategy is a best response to the belief that opponents play level 0.

Level-k reasoning

One model of a player's thought process is **level-k reasoning**.

- **L0**: level 0 is an 'unsophisticated' player who chooses from his available strategies by some heuristic.
  - Uniform distribution over available pure strategies (picking at random)
  - Focal strategy
- **L1**: a level 1 strategy is a best response to the belief that opponents play level 0.
- **L2**: a level 2 strategy is a best response to the belief that opponents play level 1.

Level-k reasoning

One model of a player's thought process is **level-k reasoning**.

- **L0**: level 0 is an 'unsophisticated' player who chooses from his available strategies by some heuristic.
  - Uniform distribution over available pure strategies (picking at random)
  - Focal strategy
- **L1**: a level 1 strategy is a best response to the belief that opponents play level 0.
- **L2**: a level 2 strategy is a best response to the belief that opponents play level 1.
- ... **Lk**: a level k strategy is a best response to the belief that opponents play level $k-1$

Level-k reasoning

- The types $L1$ and up can be viewed as completely rational and possessing a perfect model of the game...
- But a simplified idea of other players.

Level-k reasoning

- The types $L1$ and up can be viewed as completely rational and possessing a perfect model of the game...
- But a simplified idea of other players.
- This doesn't have the fixed point characteristic of our usual equilibrium solution concepts.
- May be most useful in modeling behavior in novel situations (players have limited experience from which to predict other players' behavior) or in cases with multiple equilibria.
- Another, similar approach that we'll think about is to have Lk players best respond to some belief about the *distribution* of lower types and ignoring higher types.

Ranked coordination

Ranked coordination:

	Empire State	Grand Central
Empire State	$10, 10$	$0, 0$
Grand Central	$0, 0$	$1, 1$

Ranked coordination

Ranked coordination:

	Empire State	Grand Central
Empire State	$10, 10$	$0, 0$
Grand Central	$0, 0$	$1, 1$

- Take $L0$ to be a uniform distribution over pure strategies.
- What is the $L1$ strategy for the row player?

$$\pi(empire) = \frac{1}{2}10 + \frac{1}{2}0 = 5 \tag{10}$$

$$\pi(grand) = \frac{1}{2}0 + \frac{1}{2}1 = \frac{1}{2} \tag{11}$$

Ranked coordination

Ranked coordination:

	Empire State	Grand Central
Empire State	$10, 10$	$0, 0$
Grand Central	$0, 0$	$1, 1$

- Take $L0$ to be a uniform distribution over pure strategies.
- What is the $L1$ strategy for the row player?

$$\pi(empire) = \frac{1}{2}10 + \frac{1}{2}0 = 5 \tag{10}$$

$$\pi(grand) = \frac{1}{2}0 + \frac{1}{2}1 = \frac{1}{2} \tag{11}$$

- $L1$ row player chooses Empire State.
- And so $L2$ column player chooses Empire State...

Ranked coordination

	Empire State	Grand Central
Empire State	$10, 10$	$0, 0$
Grand Central	$0, 0$	$1, 1$

	Row player	Column player
L0	Random	Random

Ranked coordination

	Empire State	Grand Central
Empire State	$10, 10$	$0, 0$
Grand Central	$0, 0$	$1, 1$

	Row player	Column player
L0	Random	Random
L1	Empire State	Empire State

Ranked coordination

	Empire State	Grand Central
Empire State	$10, 10$	$0, 0$
Grand Central	$0, 0$	$1, 1$

	Row player	Column player
L0	Random	Random
L1	Empire State	Empire State
L2	Empire State	Empire State

Ranked coordination

	Empire State	Grand Central
Empire State	$10, 10$	$0, 0$
Grand Central	$0, 0$	$1, 1$

	Row player	Column player
L0	Random	Random
L1	Empire State	Empire State
L2	Empire State	Empire State
L3	Empire State	Empire State

Ranked coordination

	Empire State	Grand Central
Empire State	$10, 10$	$0, 0$
Grand Central	$0, 0$	$1, 1$

	Row player	Column player
L0	Random	Random
L1	Empire State	Empire State
L2	Empire State	Empire State
L3	Empire State	Empire State

Relationship to 'focal' equilibria?

Battle of the sexes

The 'battle of the sexes' game:

	Boxing	Opera
Boxing	$2, 1$	$0, 0$
Opera	$0, 0$	$1, 2$

Battle of the sexes

The 'battle of the sexes' game:

	Boxing	Opera
Boxing	$2, 1$	$0, 0$
Opera	$0, 0$	$1, 2$

- Take $L0$ to be a uniform distribution over pure strategies.
- What is the $L1$ strategy for the row player?

$$\pi(boxing) = \frac{1}{2}2 + \frac{1}{2}0 = 1 \tag{12}$$

$$\pi(opera) = \frac{1}{2}0 + \frac{1}{2}1 = \frac{1}{2} \tag{13}$$

Battle of the sexes

The 'battle of the sexes' game:

	Boxing	Opera
Boxing	$2, 1$	$0, 0$
Opera	$0, 0$	$1, 2$

- Take $L0$ to be a uniform distribution over pure strategies.
- What is the $L1$ strategy for the row player?

$$\pi(boxing) = \frac{1}{2}2 + \frac{1}{2}0 = 1 \tag{12}$$

$$\pi(opera) = \frac{1}{2}0 + \frac{1}{2}1 = \frac{1}{2} \tag{13}$$

- $L1$ row player chooses Boxing.
- And so $L2$ column player chooses Boxing...

Battle of the sexes

	Boxing	Opera
Boxing	$2, 1$	$0, 0$
Opera	$0, 0$	$1, 2$

	Row player	Column player
L0	Random	Random

Battle of the sexes

	Boxing	Opera
Boxing	$2, 1$	$0, 0$
Opera	$0, 0$	$1, 2$

	Row player	Column player
L0	Random	Random
L1	Boxing	Opera

Battle of the sexes

	Boxing	Opera
Boxing	$2, 1$	$0, 0$
Opera	$0, 0$	$1, 2$

	Row player	Column player
L0	Random	Random
L1	Boxing	Opera
L2	Opera	Boxing

Battle of the sexes

	Boxing	Opera
Boxing	2, 1	0, 0
Opera	0, 0	1, 2

	Row player	Column player
L0	Random	Random
L1	Boxing	Opera
L2	Opera	Boxing
L3	Boxing	Opera

Rock-paper-scissors

Rock-paper-scissors:

	Rock	Paper	Scissors
Rock	0, 0	0, 1	1, 0
Paper	1, 0	0, 0	0, 1
Scissors	0, 1	1, 0	0, 0

Rock-paper-scissors

Rock-paper-scissors:

	Rock	Paper	Scissors
Rock	$0,0$	$0,1$	$1,0$
Paper	$1,0$	$0,0$	$0,1$
Scissors	$0,1$	$1,0$	$0,0$

- Take $L0$ to be a uniform distribution over pure strategies.
- What is the $L1$ strategy for the row player?

$$\pi(rock) = \pi(paper) = \pi(scissors) = \frac{1}{3}1 + \frac{1}{3}1 + \frac{1}{3}1 = \frac{1}{3} \qquad (14)$$

- $L1$ row player chooses to mix with equal probability.
- But this is not at all unique since any strategy gives expected payoff $\frac{1}{3}$.

Rock-paper-scissors

	Rock	Paper	Scissors
Rock	$0, 0$	$0, 1$	$1, 0$
Paper	$1, 0$	$0, 0$	$0, 1$
Scissors	$0, 1$	$1, 0$	$0, 0$

	Row player	Column player
L0	Random	Random

Rock-paper-scissors

	Rock	Paper	Scissors
Rock	0, 0	0, 1	1, 0
Paper	1, 0	0, 0	0, 1
Scissors	0, 1	1, 0	0, 0

	Row player	Column player
L0	Random	Random
L1	Random	Random

Rock-paper-scissors

	Rock	Paper	Scissors
Rock	0, 0	0, 1	1, 0
Paper	1, 0	0, 0	0, 1
Scissors	0, 1	1, 0	0, 0

	Row player	Column player
L0	Random	Random
L1	Random	Random
L2	Random	Random

Rock-paper-scissors

	Rock	Paper	Scissors
Rock	$0, 0$	$0, 1$	$1, 0$
Paper	$1, 0$	$0, 0$	$0, 1$
Scissors	$0, 1$	$1, 0$	$0, 0$

	Row player	Column player
L0	Random	Random
L1	Random	Random
L2	Random	Random
L3	Random	Random

Changing $L0$

Rock-paper-scissors:

	Rock	Paper	Scissors
Rock	$0, 0$	$0, 1$	$1, 0$
Paper	$1, 0$	$0, 0$	$0, 1$
Scissors	$0, 1$	$1, 0$	$0, 0$

Changing $L0$

Rock-paper-scissors:

	Rock	Paper	Scissors
Rock	$0, 0$	$0, 1$	$1, 0$
Paper	$1, 0$	$0, 0$	$0, 1$
Scissors	$0, 1$	$1, 0$	$0, 0$

- Take $L0$ to be Rock.
- What is the $L1$ strategy for the row player?

$$\pi(rock) = \pi(scissors) = 0 \tag{15}$$

$$\pi(paper) = 1 \tag{16}$$

- $L1$ row player chooses paper.

Changing $L0$

	Rock	Paper	Scissors
Rock	$0, 0$	$0, 1$	$1, 0$
Paper	$1, 0$	$0, 0$	$0, 1$
Scissors	$0, 1$	$1, 0$	$0, 0$

	Row player	Column player
L0	Rock	Rock

Changing $L0$

	Rock	Paper	Scissors
Rock	$0, 0$	$0, 1$	$1, 0$
Paper	$1, 0$	$0, 0$	$0, 1$
Scissors	$0, 1$	$1, 0$	$0, 0$

	Row player	Column player
L0	Rock	Rock
L1	Paper	Paper

Changing $L0$

	Rock	Paper	Scissors
Rock	$0, 0$	$0, 1$	$1, 0$
Paper	$1, 0$	$0, 0$	$0, 1$
Scissors	$0, 1$	$1, 0$	$0, 0$

	Row player	Column player
L0	Rock	Rock
L1	Paper	Paper
L2	Scissors	Scissors

Changing $L0$

	Rock	Paper	Scissors
Rock	$0, 0$	$0, 1$	$1, 0$
Paper	$1, 0$	$0, 0$	$0, 1$
Scissors	$0, 1$	$1, 0$	$0, 0$

	Row player	Column player
L0	Rock	Rock
L1	Paper	Paper
L2	Scissors	Scissors
L3	Rock	Rock

The beauty contest

- Everyone will write down an integer between 0 and 100 (inclusive).
- I will collect the numbers, calculate the average, and multiply the average by two-thirds. Call the result x.
- The player who wrote down the number that is closest to x will win $10.

The beauty contest

- Everyone will write down an integer between 0 and 100 (inclusive).
- I will collect the numbers, calculate the average, and multiply the average by two-thirds. Call the result x.
- The player who wrote down the number that is closest to x will win \$10.
- NE at all picking 0 or all picking 1.

Level-k in the beauty contest

- Everyone will write down an integer between 0 and 100 (inclusive).
- I will collect the numbers, calculate the average, and multiply the average by two-thirds. Call the result x.
- The player who wrote down the number that is closest to x will win $10.

Let $L0$ be random over pure strategies: that is, on average $L0$ plays 50.

Level-k in the beauty contest

- Everyone will write down an integer between 0 and 100 (inclusive).
- I will collect the numbers, calculate the average, and multiply the average by two-thirds. Call the result x.
- The player who wrote down the number that is closest to x will win $10.

Let $L0$ be random over pure strategies: that is, on average $L0$ plays 50.

- If other players are $L0$, what does an $L1$ player do?

Level-k in the beauty contest

- Everyone will write down an integer between 0 and 100 (inclusive).
- I will collect the numbers, calculate the average, and multiply the average by two-thirds. Call the result x.
- The player who wrote down the number that is closest to x will win $10.

Let $L0$ be random over pure strategies: that is, on average $L0$ plays 50.

- If other players are $L0$, what does an $L1$ player do?
- $\frac{2}{3} * 50 = 33\frac{1}{3}$: if others are $L0$, the best response—the $L1$ strategy is to play around 33.
- Note the minor point we don't have the $L1$ player account for their impact on the average, capturing either a large number of players or 'naive' play by $L1$.

Level-k in the beauty contest

- Everyone will write down an integer between 0 and 100 (inclusive).
- I will collect the numbers, calculate the average, and multiply the average by two-thirds. Call the result x.
- The player who wrote down the number that is closest to x will win $10.
- $L2$ best responds to the belief that all opponents are $L1$.

Level-k in the beauty contest

- Everyone will write down an integer between 0 and 100 (inclusive).
- I will collect the numbers, calculate the average, and multiply the average by two-thirds. Call the result x.
- The player who wrote down the number that is closest to x will win \$10.
- $L2$ best responds to the belief that all opponents are $L1$.
- $L2$ plays $\frac{2}{3} * 33 = 22$.

Level-k in the beauty contest

- Everyone will write down an integer between 0 and 100 (inclusive).
- I will collect the numbers, calculate the average, and multiply the average by two-thirds. Call the result x.
- The player who wrote down the number that is closest to x will win $10.

- $L2$ best responds to the belief that all opponents are $L1$.
- $L2$ plays $\frac{2}{3} * 33 = 22$.
- $L3$ plays $\frac{2}{3} * 22 \approx 15$.

Level-k in the beauty contest

- Everyone will write down an integer between 0 and 100 (inclusive).
- I will collect the numbers, calculate the average, and multiply the average by two-thirds. Call the result x.
- The player who wrote down the number that is closest to x will win $10.
- $L2$ best responds to the belief that all opponents are $L1$.
- $L2$ plays $\frac{2}{3} * 33 = 22$.
- $L3$ plays $\frac{2}{3} * 22 \approx 15$.
- ...
- $L\infty$ plays 0.

The beauty contest

Bosch-Domènech et al (2002) categorize some possible reasoning processes:

1. *Fixed point*: zero is the unique equilibrium; unilateral deviation from zero will not win.

The beauty contest

Bosch-Domènech et al (2002) categorize some possible reasoning processes:

1. *Fixed point*: zero is the unique equilibrium; unilateral deviation from zero will not win.
2. *Iterated deletion of weakly dominated strategies*: rational player never chooses > 66; if she believes others are rational never chooses > 44; iterate to zero.

The beauty contest

Bosch-Domènech et al (2002) categorize some possible reasoning processes:

1. *Fixed point*: zero is the unique equilibrium; unilateral deviation from zero will not win.
2. *Iterated deletion of weakly dominated strategies*: rational player never chooses > 66; if she believes others are rational never chooses > 44; iterate to zero.
3. *Iterated best reply*: level-k reasoning; L0 random so on average 50; L1 $50 * \frac{2}{3}$; Lk $50 * \frac{2}{3}^k$. Model as if all players think they're one level deeper than everyone else.

The beauty contest

Bosch-Domènech et al (2002) categorize some possible reasoning processes:

1. *Fixed point*: zero is the unique equilibrium; unilateral deviation from zero will not win.
2. *Iterated deletion of weakly dominated strategies*: rational player never chooses > 66; if she believes others are rational never chooses > 44; iterate to zero.
3. *Iterated best reply*: level-k reasoning; L0 random so on average 50; L1 $50 * \frac{2}{3}$; Lk $50 * \frac{2}{3}^{k}$. Model as if all players think they're one level deeper than everyone else.
4. *Iterated best reply II*: same as above but allow players to hold beliefs that others are at more than one level of reasoning.

The beauty contest

Bosch-Domènech et al (2002) categorize some possible reasoning processes:

1. *Fixed point*: zero is the unique equilibrium; unilateral deviation from zero will not win.
2. *Iterated deletion of weakly dominated strategies*: rational player never chooses > 66; if she believes others are rational never chooses > 44; iterate to zero.
3. *Iterated best reply*: level-k reasoning; L0 random so on average 50; L1 $50 * \frac{2}{3}$; Lk $50 * \frac{2}{3}^k$. Model as if all players think they're one level deeper than everyone else.
4. *Iterated best reply II*: same as above but allow players to hold beliefs that others are at more than one level of reasoning.
5. *Experimenters*: run an experiment!

Experimental evidence

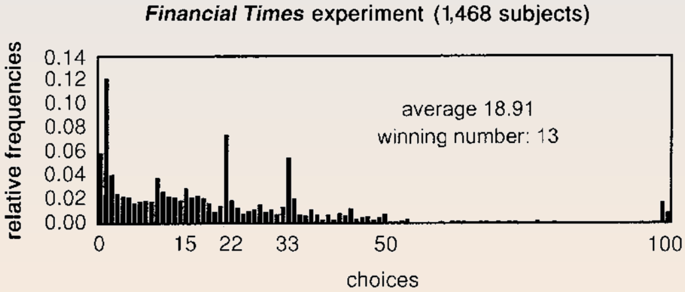

Figure: Results from the first newspaper experiment reported in Bosch-Domènech et al (2002)

Experimental evidence

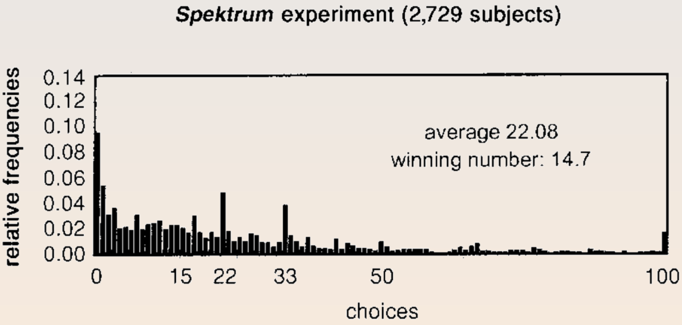

Figure: Results from the second newspaper experiment reported in Bosch-Domènech et al (2002)

Experimental evidence

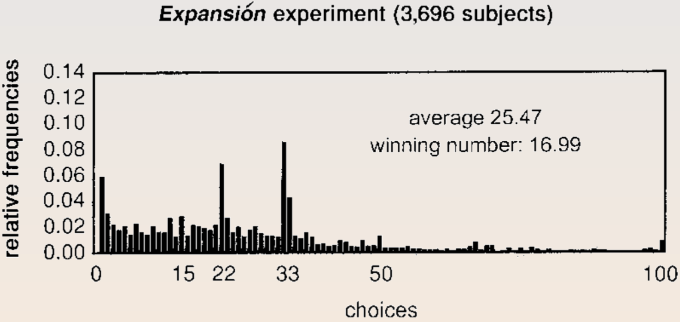

Figure: Results from the third newspaper experiment reported in Bosch-Domènech et al (2002)

Experimental evidence

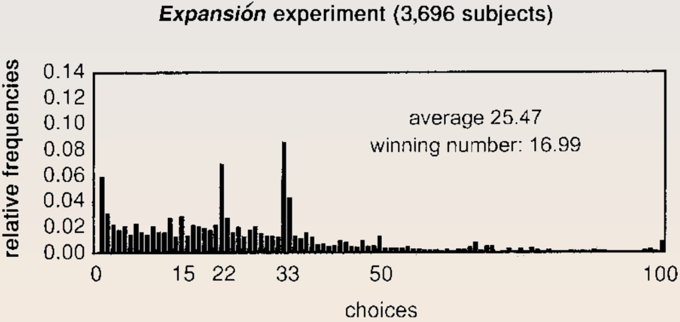

Figure: Results from the third newspaper experiment reported in Bosch-Domènech et al (2002)

- Spikes at 33 and 22... evidence of $L1$ and $L2$ types?

Experimental evidence

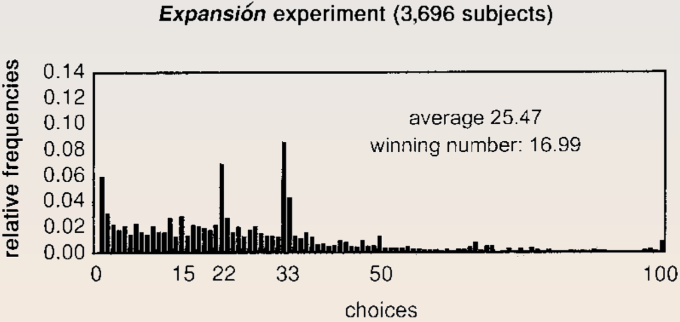

Figure: Results from the third newspaper experiment reported in Bosch-Domènech et al (2002)

- Spikes at 33 and 22... evidence of $L1$ and $L2$ types?
- Spike in the neighborhood of zero... evidence of $L\infty$? Of 'equilibrium' thinking?

Experimental evidence

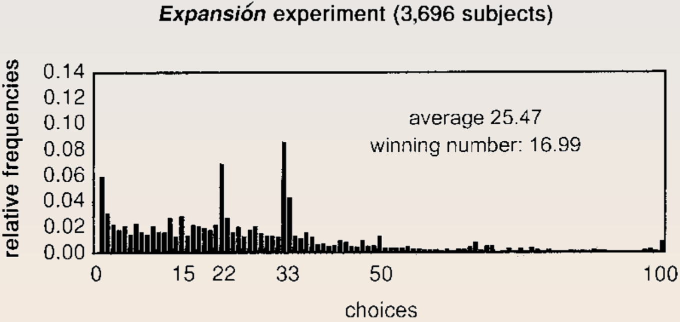

Figure: Results from the third newspaper experiment reported in Bosch-Domènech et al (2002)

- Spikes at 33 and 22... evidence of $L1$ and $L2$ types?
- Spike in the neighborhood of zero... evidence of $L\infty$? Of 'equilibrium' thinking?
- Small spike at 100... what are these people up to?

Explanations from experimental subjects

From Bosch-Domènech et al:

- "I choose 1. This is what is nearest to $x = 0$, which is the only number equal to $\frac{2}{3}$ of itself. Logical answer."

Explanations from experimental subjects

From Bosch-Domènech et al:

- "I choose 1. This is what is nearest to $x = 0$, which is the only number equal to $\frac{2}{3}$ of itself. Logical answer."

Fixed point a.k.a. Nash equilibrium reasoning.

- "I choose the number 15.93. The reasoning is the following: I assume
  - 10% do not have a clue and pick the mean 50
  - 20% give a naive answer: $33 = 50*2/3$
  - 50% go a second round: $22 = 33*2/3$
  - 5% go a third round: $14 = 22*2/3$
  - 5% are really devious and choose $10 = 14*2/3$
  - 10% are crazy mathematicians who choose 1."

Explanations from experimental subjects

From Bosch-Domènech et al:

- "I choose 1. This is what is nearest to $x = 0$, which is the only number equal to $\frac{2}{3}$ of itself. Logical answer."

Fixed point a.k.a. Nash equilibrium reasoning.

- "I choose the number 15.93. The reasoning is the following: I assume
  - 10% do not have a clue and pick the mean 50
  - 20% give a naive answer: $33 = 50*2/3$
  - 50% go a second round: $22 = 33*2/3$
  - 5% go a third round: $14 = 22*2/3$
  - 5% are really devious and choose $10 = 14*2/3$
  - 10% are crazy mathematicians who choose 1."

Iterated best reply with a conjecture about the distribution of types.

Explanations from experimental subjects

- “If everybody would choose 100, the maximum number that could be chosen is 66.6. Therefore, theoretically nobody will send a number over 66.6 and, if you multiply this by 2/3 we get 44.4. Therefore, in theory, nobody should be sending either a number over 44.4. Following this reasoning process the only number that should be sent is 1. However, I understand that many different people participate in this game and not everybody will apply the reasoning process explained above. Therefore, and taking into account that the majority of people would go all the way up to 1, I choose 6.8.”

Explanations from experimental subjects

- "If everybody would choose 100, the maximum number that could be chosen is 66.6. Therefore, theoretically nobody will send a number over 66.6 and, if you multiply this by 2/3 we get 44.4. Therefore, in theory, nobody should be sending either a number over 44.4. Following this reasoning process the only number that should be sent is 1. However, I understand that many different people participate in this game and not everybody will apply the reasoning process explained above. Therefore, and taking into account that the majority of people would go all the way up to 1, I choose 6.8."

Iterated dominance plus some extra 'secret sauce' for rounding/trembling.

Explanations from experimental subjects

- "If all the numbers had the same probability of being chosen, the mean would be 50 and the choice should be $\frac{2}{3}50 = 33.33$. However, I have estimated a percentage of deviation around 33.33 of 10% and, there- fore, I choose the number 30."

Explanations from experimental subjects

- "If all the numbers had the same probability of being chosen, the mean would be 50 and the choice should be $\frac{2}{3}50 = 33.33$. However, I have estimated a percentage of deviation around 33.33 of 10% and, there- fore, I choose the number 30."

Level-1(ish) reasoning.

- "In case that all numbers are equally distributed, the average will be 50. $\frac{2}{3}$ of that is about 33. Since the readers of Spektrum are certainly not the dumbest, they will all get to 33 at the first step. However, $\frac{2}{3}$ of that is 22. Since certainly all will calculate this, one has to take $\frac{2}{3}$ of that.... The series continues ad infinitum and at the end you get 0! However, I choose, despite that logic, 2.32323."

Explanations from experimental subjects

- "If all the numbers had the same probability of being chosen, the mean would be 50 and the choice should be $\frac{2}{3}50 = 33.33$. However, I have estimated a percentage of deviation around 33.33 of 10% and, there- fore, I choose the number 30."

Level-1(ish) reasoning.

- "In case that all numbers are equally distributed, the average will be 50. $\frac{2}{3}$ of that is about 33. Since the readers of Spektrum are certainly not the dumbest, they will all get to 33 at the first step. However, $\frac{2}{3}$ of that is 22. Since certainly all will calculate this, one has to take $\frac{2}{3}$ of that.... The series continues ad infinitum and at the end you get 0! However, I choose, despite that logic, 2.32323."

Level-∞(ish) reasoning.

Explanations from experimental subjects

- "I decided to run an experiment with a group of friends. Since I believed that the sample was representative of the participants in the general experiment, I assumed the result of the experiment would be a good indicator of the solution. People used the following reasoning. One said simply the mean, 50 (!!!). Some others mul- tiplied $\frac{2}{3}$ by 50 and said 33.33. One said 25 because 'today is the 25th'. In some other cases people said 1, or a number close to 1 even though in one case the reason was 'to pick a number at random'. The mean was around 13 and, therefore, my answer is 8.66666."

Explanations from experimental subjects

- "I decided to run an experiment with a group of friends. Since I believed that the sample was representative of the participants in the general experiment, I assumed the result of the experiment would be a good indicator of the solution. People used the following reasoning. One said simply the mean, 50 (!!!). Some others mul- tiplied $\frac{2}{3}$ by 50 and said 33.33. One said 25 because 'today is the 25th'. In some other cases people said 1, or a number close to 1 even though in one case the reason was 'to pick a number at random'. The mean was around 13 and, therefore, my answer is 8.66666."

An experimenter.

Explanations from experimental subjects

- Chooses 42 with the following ex- planation: “Even though I know I won’t win, I take the answer from the question of life, universe, and the rest [see Douglas Adams, “The Hitchhiker’s Guide to the Galaxy” (1995)] and use it for everything. Maybe I will also use it for this quiz.”

Explanations from experimental subjects

- Chooses 42 with the following ex- planation: "Even though I know I won't win, I take the answer from the question of life, universe, and the rest [see Douglas Adams, "The Hitchhiker's Guide to the Galaxy" (1995)] and use it for everything. Maybe I will also use it for this quiz."

No idea. Level-0?

Experimental evidence

TABLE 2—RELATIVE FREQUENCIES OF THE DIFFERENT TYPES OF REASONING FROM THE COMMENTS OF E AND S EXPERIMENTS

Types of reasoning processes	Relative frequencies
Fixed point	2.56 percent
Equilibrium, without further explanation	14.61 percent
Iterated dominance (ID)	13.77 percent of which 11.10 percentage points are Level-∞
Iterated best-reply degenerate (IBRd)	54.71 percent of which 25.45 percentage points are Level-∞ 12.47 percentage points are Level 0
Iterated best-reply nondegenerate (IBRnd)	9.28 percent
Experimenters	5.09 percent

- Does knowing this help us to predict the outcome of other games?

Some thoughts

- Level-k is the modal reasoning process
- Relationship between level k and iterated dominance?
  - Can be a bit difficult to disentangle here since they generate similar predictions at the higher orders
- Distribution of levels as types in a Bayesian game?
  - cf. Camerer, Ho & Chong (2004) 'cognitive hierarchy'
  - Stated goal is to explain patterns of decision making across different games
  - Why are people pretty good at reaching equilibrium quickly without much learning time in some games but pretty bad in others?

Cognitive hierarchy

- Matches the ‘iterated best reply II’ categorization from the beauty contest classification
- Hierarchy as follows:
  - Step 0 players randomize
  - Step k players best respond assuming that other players are distributed over step 0 to $k-1$

Cognitive hierarchy

- Matches the 'iterated best reply II' categorization from the beauty contest classification
- Hierarchy as follows:
  - Step 0 players randomize
  - Step k players best respond assuming that other players are distributed over step 0 to $k-1$
- That is: every player assumes that they are at the highest level of reasoning relative to everyone else

Cognitive hierarchy

- Matches the 'iterated best reply II' categorization from the beauty contest classification
- Hierarchy as follows:
  - Step 0 players randomize
  - Step k players best respond assuming that other players are distributed over step 0 to $k-1$
- That is: every player assumes that they are at the highest level of reasoning relative to everyone else
- Frequency distribution $f(k)$ of step k players assumed to be Poisson
- Poisson is prob. distribution with mean and variance both equal to the same number λ
- $\lambda = 1.61$ is median estimated from 24 beauty contest data sets

Poisson distribution with $\lambda = 1.61$

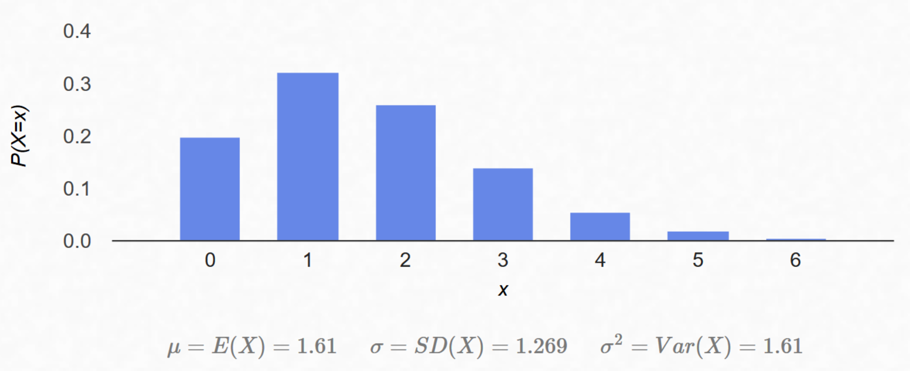

(Generated via
https://homepage.divms.uiowa.edu/ mbognar/applets/pois.html)

The 11-20 money request game

Arad & Rubinstein (2012):

- The ex post fitting of a distribution of level k types to experimental results is not ideal.
- Maybe someone using a completely different thought process is spuriously categorized as a level-k reasoner.
- Also tough to find games in which level-k is the ‘natural’ reasoning process and in which a particular $L0$ strategy is obviously the ‘right’ assumption.
- This paper is one of several recent papers to try to overcome these issues.

The 11-20 money request game

- **Players**: two individuals, 1 and 2.
- **Moves**: the individuals simultaneously choose a number of dollars between 11 and 20 (inclusive).
- **Payoffs**: each player receives the amount they chose, plus a bonus of 20 dollars if and only if she asks for exactly one dollar less than the other player.

The 11-20 money request game

- **Players**: two individuals, 1 and 2.
- **Moves**: the individuals simultaneously choose a number of dollars between 11 and 20 (inclusive).
- **Payoffs**: each player receives the amount they chose, plus a bonus of 20 dollars if and only if she asks for exactly one dollar less than the other player.

- $20 is a 'natural' $L0$; it can be justified quite easily and 'sounds good'.
- Level-k best responses are very easy to come up with.
- $L1$ is robust to quite varied $L0$ strategies.
- Lk is quite robust to beliefs about distribution of other types.
- No social aspect to payoffs.
- No pure-strategy NE exists (check) and no dominated strategies exist.

The 11-20 game: Nash and results

TABLE 1—THE 11–20 GAME ($n = 108$)

Action	11	12	13	14	15	16	17	18	19	20
Equilibrium (%)					25	25	20	15	10	5
Results (%)	4	0	3	6	1	6	32	30	12	6

The 11-20 game: Nash and results

TABLE 1—THE 11–20 GAME ($n = 108$)

Action	11	12	13	14	15	16	17	18	19	20
Equilibrium (%)					25	25	20	15	10	5
Results (%)	4	0	3	6	1	6	32	30	12	6

- We can reject the hypothesis that subjects' choices were Nash.
- Choices consistent with $L1$, $L2$ and $L3$ cover almost three quarters of the observations.
- The best statistical fit is $L0 - L3$ with noise (as opposed to Nash, L0-L4, L0-L5).

The 11-20 game: explanations

TABLE 1—THE 11–20 GAME ($n = 108$)

Action	11	12	13	14	15	16	17	18	19	20
Equilibrium (%)					25	25	20	15	10	5
Results (%)	4	0	3	6	1	6	32	30	12	6

The 11-20 game: explanations

TABLE 1—THE 11–20 GAME ($n = 108$)

Action	11	12	13	14	15	16	17	18	19	20
Equilibrium (%)					25	25	20	15	10	5
Results (%)	4	0	3	6	1	6	32	30	12	6

Subjects' explanations for their choices reveal that:

- Almost everyone in the 17-19 range used level-k reasoning.
- No-one who chose in the 11-15 range used level-k reasoning.
- Only one 16-chooser described using level-4 reasoning.
- A handful of subjects iterated back around from 11 to 20 but, despite this, ended up back within three iterations from 20.

The 11-20 game: extensions

The authors use two further experiments as robustness checks against a couple of potential competing interpretations of the data.

- The first enhances the salience of the $L0$ action:
  - Extra payoff condition: also get a \$20 bonus if she asks for \$20 and the other player asks for \$11.

The 11-20 game: extensions

The authors use two further experiments as robustness checks against a couple of potential competing interpretations of the data.

- The first enhances the salience of the $L0$ action:
  - Extra payoff condition: also get a \$20 bonus if she asks for \$20 and the other player asks for \$11.
  - Makes \$20 even more attractive than before.
  - Results show that $L0 - L3$ same as before, but of these many more are at $L1$: recognize new appeal of \$20 and so reason less deeply?

TABLE 2—THE CYCLE VERSION ($n = 72$)

Action	11	12	13	14	15	16	17	18	19	20
Equilibrium (%)					25	25	20	15	10	5
Cycle version (%)	1	1	0	1	0	4	10	22	47	13
Basic version (%)	4	0	3	6	1	6	32	30	12	6

The 11-20 game: extensions

- The first extension reduces the cost of higher levels of reasoning:
  - Change to payoffs: choose 20, get 20. Choose another number, get 17. Bonus of 20 if choose one less than the other player.

The 11-20 game: extensions

- The first extension reduces the cost of higher levels of reasoning:
  - Change to payoffs: choose 20, get 20. Choose another number, get 17. Bonus of 20 if choose one less than the other player.
  - Undercutting doesn't carry additional sacrifice of a dollar at higher levels of reasoning.
  - Again there is little evidence of orders of reasoning above $L3$, even with the indirect cost of extra reasoning removed.

TABLE 3—THE COSTLESS ITERATIONS VERSION ($n = 53$)

Action	11	12	13	14	15	16	17	18	19	20
Equilibrium (%)				10	15	15	15	15	15	15
Results (%)	0	4	0	4	4	4	9	21	40	15

Applying level-k

- The level-k model has been used to explain some empirical puzzles
- In particular it has been usefully applied to auctions and zero-sum cheap talk games
- We will look at its application to 'magical' coordination (adapted from Crawford 2007)
- And its application to hide and seek games (Crawford and Iriberri 2007 AER)

Applying level-k

- The level-k model has been used to explain some empirical puzzles
- In particular it has been usefully applied to auctions and zero-sum cheap talk games
- We will look at its application to 'magical' coordination (adapted from Crawford 2007)
- And its application to hide and seek games (Crawford and Iriberri 2007 AER)
- Another extremely important application is to auction theory (e.g. Crawford and Iriberri 2007 Econometrica) but it is a bit beyond our technical level here
  - In short, level k does a good job at explaining in a unified way a couple of famous disparate puzzles in auction theory: the winner's curse in common value auctions and overbidding in independent private value auctions
  - (Auctions, by the way, are a really important and useful application of game theory—lots and lots of important real-world transactions are auctions, and auction theory can be used as a framework to think about non-auction institutions as well)

Coordination with private preferences

Puzzle: experimental subjects are better at coordination than the mixed-strategy NE result. Kahneman 1988: "To a psychologist, it looks like magic."

Coordination with private preferences

Puzzle: experimental subjects are better at coordination than the mixed-strategy NE result. Kahneman 1988: "To a psychologist, it looks like magic."

	H	D
H	$0, 0$	$a, 1$
D	$1, a$	$0, 0$

With $a > 1$.

Coordination with private preferences

Puzzle: experimental subjects are better at coordination than the mixed-strategy NE result. Kahneman 1988: "To a psychologist, it looks like magic."

	H	D
H	$0, 0$	$a, 1$
D	$1, a$	$0, 0$

With $a > 1$.

- Unique symmetric NE is mixed strategy: $Pr(H) = \frac{a}{1+a}$.
- Thus expected coordination rate in equilibrium (i.e. how often the off-diagonals are realized) is $2p(1-p) = \frac{2a}{(1+a)^2}$.
- And the expected payoff to each player is thus $\frac{a}{1+a} < 1$.

Levels of reasoning

	H	D
H	$0, 0$	$a, 1$
D	$1, a$	$0, 0$

- Let $L0$ be random again: $Pr(H) = Pr(D) = \frac{1}{2}$.
- We will let types be $L1$ through $L4$.
- $L1$ chooses H, $L2$ chooses D, $L3$ chooses H, $L4$ chooses D.

Type frequencies

Types	*L1*	*L2*	*L3*	*L4*
L1	H, H	H, D	H, H	H, D
L2	D, H	D, D	D, H	D, D
L3	H, H	H, D	H, H	H, D
L4	D, H	D, D	D, H	D, D

Table 1. Level-*k* Outcomes

We will assume no *L*0 types and that type frequencies are independent of payoffs or player role.

Equilibrium vs. level k

- Since $L1$ and $L3$ are outcome equivalent, combine them and call their proportion in the population v.
- The level-k coordination rate is $2v(1-v)$.
- This is maximized by $v = \frac{1}{2}$ at the level $\frac{1}{2}$.

Equilibrium vs. level k

- Since $L1$ and $L3$ are outcome equivalent, combine them and call their proportion in the population v.
- The level-k coordination rate is $2v(1-v)$.
- This is maximized by $v = \frac{1}{2}$ at the level $\frac{1}{2}$.
- Compare to the rate in mixed-strategy NE: $\frac{2a}{(1+a)^2}$, which does just as well when $a = 1$ but converges to 0 as $a \to \infty$.
- Coordination rates can be dramatically higher under this model of predictable heterogeneity (population frequency) in orders of reasoning for v near $\frac{1}{2}$.

Equilibrium vs. level k

- Since $L1$ and $L3$ are outcome equivalent, combine them and call their proportion in the population v.
- The level-k coordination rate is $2v(1-v)$.
- This is maximized by $v = \frac{1}{2}$ at the level $\frac{1}{2}$.
- Compare to the rate in mixed-strategy NE: $\frac{2a}{(1+a)^2}$, which does just as well when $a = 1$ but converges to 0 as $a \to \infty$.
- Coordination rates can be dramatically higher under this model of predictable heterogeneity (population frequency) in orders of reasoning for v near $\frac{1}{2}$.
- Compare to Arad & Rubinstein... to Poisson...

Doing better by thinking less?

- The idea is that higher orders can think through the decision of lower orders and best respond, benefiting both.
- Players here are rational, yet nothing like equilibrium reasoning is being used here.
- Instead coordination is a predictable result of non-equilibrium reasoning.

Hide and seek

Hide and seek games are two player zero-sum games: one player wins by matching the opponent and the other player wins by mismatching

- Applications include roduct design and differentiation, military strategy, election campaigning, and social signaling
- Rubinstein, Tversky, and Heller (1997) is an example of a hide and seek experiment; 4 locations with non-neutral labeling
- For example: four 'boxes' labeled A-B-A-A

Hide and seek

Hide and seek games are two player zero-sum games: one player wins by matching the opponent and the other player wins by mismatching

- Applications include roduct design and differentiation, military strategy, election campaigning, and social signaling
- Rubinstein, Tversky, and Heller (1997) is an example of a hide and seek experiment; 4 locations with non-neutral labeling
- For example: four 'boxes' labeled A-B-A-A
- Zero sum games (like rock paper scissors from earlier) are clean examples in which randomization is a smart strategy
- The unique equilibrium is in mixed strategies: both players randomize uniformly across locations
- But seekers find the treasure 32% of the time

Hide and seek experiments

TABLE 1—AGGREGATE CHOICE FREQUENCIES IN RTH'S TREATMENTS

RTH-4	A	*B*	A	A
Hider (53; $p = 0.0026$)	9 percent	36 percent	40 percent	15 percent
Seeker (62; $p = 0.0003$)	13 percent	31 percent	45 percent	11 percent
RT-AABA–Treasure	A	A	*B*	A
Hider (189; $p = 0.0096$)	22 percent	35 percent	19 percent	25 percent
Seeker (85; $p = $ 9E-07)	13 percent	51 percent	21 percent	15 percent
RT-AABA–Mine	A	A	*B*	A
Hider (132; $p = 0.0012$)	24 percent	39 percent	18 percent	18 percent
Seeker (73; $p = 0.0523$)	29 percent	36 percent	14 percent	22 percent
RT-1234–Treasure	1	*2*	3	4
Hider (187; $p = 0.0036$)	25 percent	22 percent	36 percent	18 percent
Seeker (84; $p = $ 3E–05)	20 percent	18 percent	48 percent	14 percent
RT-1234–Mine	1	*2*	3	4
Hider (133; $p = $ 6E-06)	18 percent	20 percent	44 percent	17 percent
Seeker (72; $p = 0.149$)	19 percent	25 percent	36 percent	19 percent
R-ABAA	A	*B*	A	A
Hider (50; $p = 0.0186$)	16 percent	18 percent	44 percent	22 percent
Seeker (64; $p = $ 9E-07)	16 percent	19 percent	54 percent	11 percent

Notes: Sample sizes and p-values for significant differences from equilibrium in parentheses; salient labels in italics; order of presentation of locations to subjects as shown.

Rubinstein, Tversky, and Heller (1996)

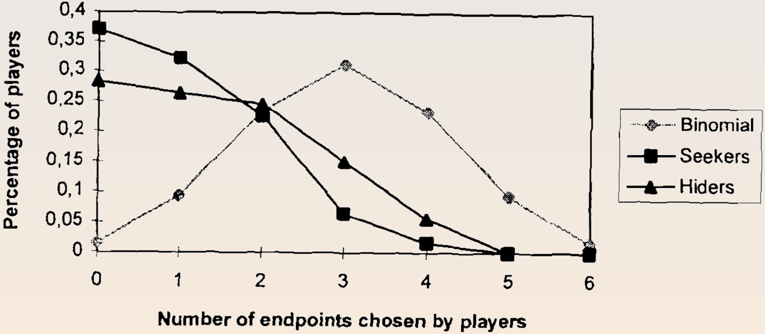

Figure 3. Distribution of patterns of endpoint choices for seekers and hiders compared to the predicted binomial distribution

Level k and hide-and-seek

Crawford and Iriberri (2007 AER) propose that a calibrated level k model can rationalize hide-and-seek behavior

- This would mean that choices are not naive but are rather sophisticated responses to a mental model of what types the other players might be
- They assign L0 probabilities that favors the supposedly 'salient' locations (most on the endpoints and next most on box B) and puts low weight on the non-salient central A
- Find 19% L1, 32% L2, 24% L3, 25% L4 explains RTH data best

Level k and hide-and-seek

Crawford and Iriberri (2007 AER) propose that a calibrated level k model can rationalize hide-and-seek behavior

- This would mean that choices are not naive but are rather sophisticated responses to a mental model of what types the other players might be
- They assign L0 probabilities that favors the supposedly 'salient' locations (most on the endpoints and next most on box B) and puts low weight on the non-salient central A
- Find 19% L1, 32% L2, 24% L3, 25% L4 explains RTH data best
- We are in a philosophical quagmire here about whether non-salience is itself salient...
- So in that spirit Wolff (2016) elicits salience from experimental subjects in hide-and-seek; finds the opposite of Crawford and Iriberri

Wolff (2016)

	$A_{(1)}$	$B_{(2)}$	$A_{(3)}$	$A_{(4)}$
Picking Task (405 participants)				
relative click frequencies (%)	21	38	35	6
mean response times (seconds)	8.8	7.7	8.5	11.9
Guessing Task (72 participants)				
average estimated relative click frequency	21	41	22	15
Beauty Contest (30 participants)				
mean rank in beauty contest	2.3	1.5	2.5	3.6
Post-H&S Guessing (156 participants)				
average estimated relative click frequency				
...by hiders (78 obs.)	19	38	24	19
...by seekers (78 obs.)	19	40	25	17
postCoord Guessing (72 participants)				
average estimated relative click frequency	19	50	18	14
postDiscoord Guessing (72 participants)				
average estimated relative click frequency	20	37	24	19
Rating Task (90 participants)				
average conspicuousness reported (scale: 0 to 10)	5.7	7.5	5.6	5.3
Post-Story Rating (90 participants)				
average conspicuousness reported (scale: 0 to 10)	3.8	7.4	4.3	4.0
Post-Story Rate-Guessing (84 participants)				
average estimated rating (scale: 0 to 10)	3.9	7.5	3.3	2.7

Level k, cognitive ability, emotional traits

David Gill and Victoria Prowse (2016): how does cognitive ability and character correlate with decisions in level k / learning settings?

- 37 sessions at U of Arizona, 780 total subjects from Experimental Science Lab subject pool (grad students in economics excluded)
- Show up fee $5 and average earnings $20
- Raven's Progressive Matrices test of analytic intelligence (nonverbal, select element to complete a visual pattern)
  - Put in 'high' or 'low' ability group if test score in the top/bottom half of the session
  - No money incentive for Raven test (avoid income effects spillover to beauty contest)
- Followed by 10 round beauty contest in groups of three $p = 0.7$
  1. Own-matched sessions: all 3 same ability
  2. Cross-matched sessions: mixed ability
  3. Known to subjects

Level k, cognitive ability, emotional traits

David Gill and Victoria Prowse (2016): how does cognitive ability and character correlate with decisions in level k / learning settings?

- 37 sessions at U of Arizona, 780 total subjects from Experimental Science Lab subject pool (grad students in economics excluded)
- Show up fee \$5 and average earnings \$20
- Raven's Progressive Matrices test of analytic intelligence (nonverbal, select element to complete a visual pattern)
  - Put in 'high' or 'low' ability group if test score in the top/bottom half of the session
  - No money incentive for Raven test (avoid income effects spillover to beauty contest)
- Followed by 10 round beauty contest in groups of three $p = 0.7$
  1. Own-matched sessions: all 3 same ability
  2. Cross-matched sessions: mixed ability
  3. Known to subjects
- Some sessions preceded by questionnaire about character skills

Gill & Prowse (2016)

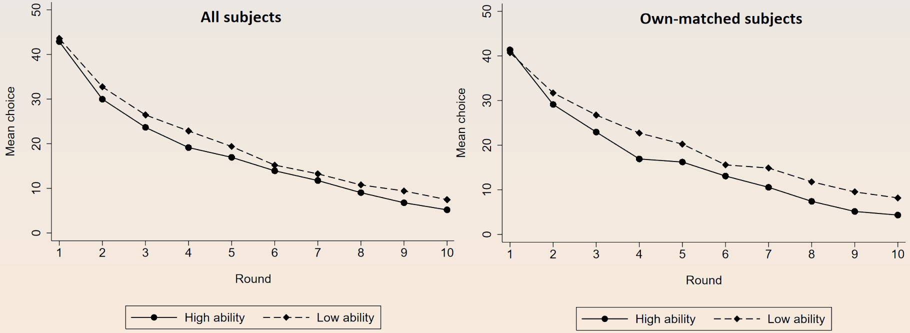

Gill & Prowse (2016)

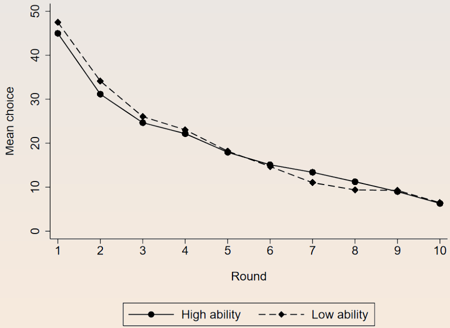

(a) Round-by-round means of choices.

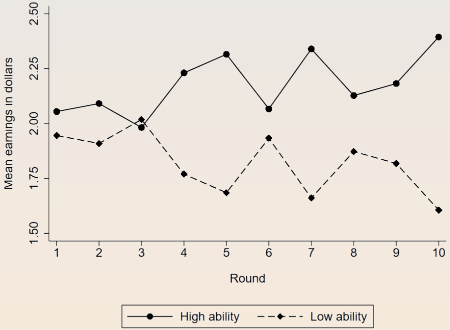

(b) Round-by-round means of earnings.

Figure 4: p-beauty contest choices and earnings of cross-matched subjects.

Level k

Authors estimate structural model of level k reasoning that fits the behavioral differences in the experimental data

- Level 0: ‘follow the crowd’ and copy group behavior from previous round
- Rule learning: include types who switch up a level during the experiment

Level k

Authors estimate structural model of level k reasoning that fits the behavioral differences in the experimental data

- Level 0: 'follow the crowd' and copy group behavior from previous round
- Rule learning: include types who switch up a level during the experiment
- Systematic positive relationship between cognitive ability and levels
- Average level of high ability subjects responds positively to opponents' ability; not true for low ability subjects

Level k

Authors estimate structural model of level k reasoning that fits the behavioral differences in the experimental data

- Level 0: 'follow the crowd' and copy group behavior from previous round
- Rule learning: include types who switch up a level during the experiment
- Systematic positive relationship between cognitive ability and levels
- Average level of high ability subjects responds positively to opponents' ability; not true for low ability subjects
- More agreeable and emotionally stable subjects chose lower numbers and had higher levels; effect less than cognitive ability
- Effects of cognition and character persist when controlling for the other

Some robustness checks from the paper

A couple of examples of other things analyzed:

- Authors present evidence that the effect of being allocated to the 'top half' versus 'bottom half' of Raven's results didn't itself affect behavior
  - To check this, can exploit that the variation in session composition meant that some 'bottom half' subjects had higher scores than some 'top half' subjects from other sessions
- Check to see if composition of cross-matched groups changes choices
  - i.e. does it matter if you're the only 'top half' person in a mixed group or if there is another?
  - Differences in earnings persist

Poe's riddle

Eliaz and Rubinstein (2011) study a repeated matching pennies game motivated by Poe's "The Purloined Letter"

- Similar kind of idea to Crawford and Iriberri: use level k to explain the deviations of experimental data from randomization
- The game:

	L	R
T	$0, 1$	$1, 0$
B	$1, 0$	$0, 1$

- Nash equilibrium is unique in mixed strategies; each player randomizes half-half over their pure strategies and so each 'wins' with probability $\frac{1}{2}$
- In an n-round repeated version, each player's win rate 'should' be distributed $Binomial(n, 0.5)$

Eliaz & Rubinstein (2011)

- Baseline treatment:
  - Players move sequentially (2nd player does not observe 1st player's move)
  - 1st player is labeled 'misleader' and 2nd player 'guesser'
  - Actions are labeled 0 and 1 so that the guesser wins when choosing the same action as the misleader

Eliaz & Rubinstein (2011)

- Baseline treatment:
  - Players move sequentially (2nd player does not observe 1st player's move)
  - 1st player is labeled 'misleader' and 2nd player 'guesser'
  - Actions are labeled 0 and 1 so that the guesser wins when choosing the same action as the misleader
- Guessers won more points than the misleader in around 51% of pairs; the misleader won more points in around 29% of pairs; 20% of pairs ended in a draw
- About 53% of rounds are won by the guesser—a seemingly small edge, but one that is statistically significant and requires large deviations from equilibrium strategies to generate
- There is a tendency in the data for guessers to repeat their action after success slightly more frequently than misleaders, which partly accounts for the slight edge for guessers

Eliaz & Rubinstein (2011)

The actual distribution vs. the binomial prediction

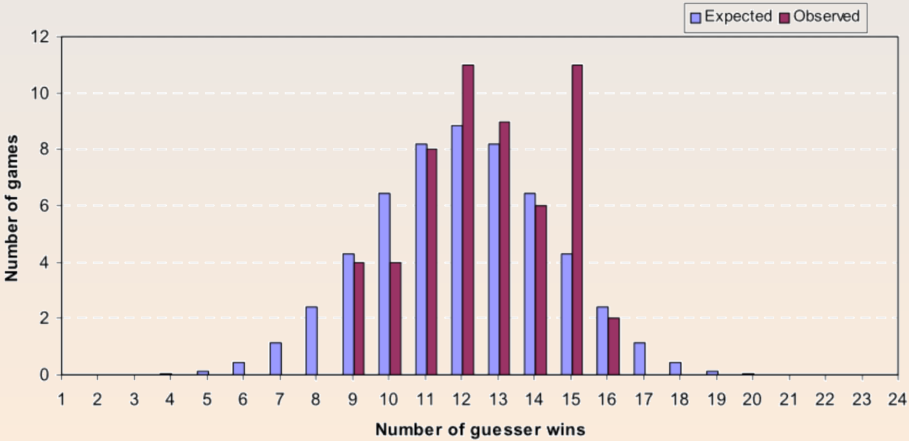

Eliaz & Rubinstein (2011)

Some variants to explore the mechanism:

1. Drop the guesser-misleader frame; players labeled 'odd' (P1) and 'even' (P2) and get a point depending on whether the sum of the numbers chosen is odd or even
    - Even player won 54% of rounds: 2nd mover and having to match the opponent's action seem enough to generate the same result

Eliaz & Rubinstein (2011)

Some variants to explore the mechanism:

1. Drop the guesser-misleader frame; players labeled 'odd' (P1) and 'even' (P2) and get a point depending on whether the sum of the numbers chosen is odd or even
   - Even player won 54% of rounds: 2nd mover and having to match the opponent's action seem enough to generate the same result
2. Check if effect is a second-mover advantage: have P1 guess P2's future choice, and repeat variant 1 with P1 the 'even' player
   - Misleader/odd-player does not get a significant advantage when moving second

Eliaz & Rubinstein (2011)

Some variants to explore the mechanism:

1. Drop the guesser-misleader frame; players labeled ‘odd’ (P1) and ‘even’ (P2) and get a point depending on whether the sum of the numbers chosen is odd or even
   - Even player won 54% of rounds: 2nd mover and having to match the opponent’s action seem enough to generate the same result
2. Check if effect is a second-mover advantage: have P1 guess P2’s future choice, and repeat variant 1 with P1 the ‘even’ player
   - Misleader/odd-player does not get a significant advantage when moving second
3. Drop all non-neutral labels: P1 chooses ‘a’ or ‘i’; P2 chooses ‘s’ or ‘t’. P1 wins with ‘at’ or ‘is’ and P2 wins with ‘is’ or ‘at’
   - Cannot reject hypothesis that each player is equally likely to win: timing-without-observability does not seem to be enough to generate the second-mover advantage

Backward induction

A core concept in game theory is the idea of **backward induction**

- What do I think you will do tomorrow?
- Given that, what should I do today?
- It is another classic case in which the predictions of mutual rationality do not match the evidence very well
- ... as we go along: do you feel like the demands of mutual rationality in a one-shot game is like or unlike the demands of mutual rationality in backward induction

Ferocious pirates

Five pirates, A, B, C, D, E. 100 gold pieces are up for grabs.

Ferocious pirates

Five pirates, A, B, C, D, E. 100 gold pieces are up for grabs.

- Pirate A proposes a division of the gold among the five groups: the division must give between 0 and 100 to each pirate, using whole pieces only, and must sum to 100.

Ferocious pirates

Five pirates, A, B, C, D, E. 100 gold pieces are up for grabs.

- Pirate A proposes a division of the gold among the five groups: the division must give between 0 and 100 to each pirate, using whole pieces only, and must sum to 100.
- All pirates (including the proposer) vote on the proposal. If a majority of pirates vote in favour, the proposal passes, the division is implemented, and the game ends. Ties are broken in favour of the proposer. If a majority of groups vote against, the motion fails and pirate A walks the plank.

Ferocious pirates

Five pirates, A, B, C, D, E. 100 gold pieces are up for grabs.

- Pirate A proposes a division of the gold among the five groups: the division must give between 0 and 100 to each pirate, using whole pieces only, and must sum to 100.
- All pirates (including the proposer) vote on the proposal. If a majority of pirates vote in favour, the proposal passes, the division is implemented, and the game ends. Ties are broken in favour of the proposer. If a majority of groups vote against, the motion fails and pirate A walks the plank.
- If A's proposal fails, B proposes a division of the gold among the remaining four pirates, which again must sum to 100. The voting stage is repeated.

Ferocious pirates

Five pirates, A, B, C, D, E. 100 gold pieces are up for grabs.

- Pirate A proposes a division of the gold among the five groups: the division must give between 0 and 100 to each pirate, using whole pieces only, and must sum to 100.
- All pirates (including the proposer) vote on the proposal. If a majority of pirates vote in favour, the proposal passes, the division is implemented, and the game ends. Ties are broken in favour of the proposer. If a majority of groups vote against, the motion fails and pirate A walks the plank.
- If A's proposal fails, B proposes a division of the gold among the remaining four pirates, which again must sum to 100. The voting stage is repeated.
- And so on until a proposal is accepted or there are no more pirates.

Ferocious pirates

Five pirates, A, B, C, D, E. 100 gold pieces are up for grabs.

- Pirate A proposes a division of the gold among the five groups: the division must give between 0 and 100 to each pirate, using whole pieces only, and must sum to 100.
- All pirates (including the proposer) vote on the proposal. If a majority of pirates vote in favour, the proposal passes, the division is implemented, and the game ends. Ties are broken in favour of the proposer. If a majority of groups vote against, the motion fails and pirate A walks the plank.
- If A's proposal fails, B proposes a division of the gold among the remaining four pirates, which again must sum to 100. The voting stage is repeated.
- And so on until a proposal is accepted or there are no more pirates.

What should pirate A propose and why?

Predation

Consider the following game:

1. **Players**: two firms, entrant E and incumbent I.
2. **Moves**: E chooses whether to 'enter' or 'stay out', then I chooses whether to 'fight' or 'accommodate'.
3. **Information**: the firms know everything about the game and payoffs, and the incumbent observes the entrant's decision before making its own choice.
4. **Payoffs**: specified according to the following matrix representation:

	Fight	*Accommodate*
Stay out	$0, 2$	$0, 2$
Enter	$-3, -1$	$2, 1$

Where are the Nash equilibria in pure strategies?

Predation

	Fight	*Accommodate*
Stay out	$0, 2$	$0, 2$
Enter	$-3, -1$	$2, 1$

There are Nash equilibria at the strategy pairs (*out*, *fight*) and (*enter*, *accommodate*).

Predation

	Fight	*Accommodate*
Stay out	$0, 2$	$0, 2$
Enter	$-3, -1$	$2, 1$

There are Nash equilibria at the strategy pairs (*out*, *fight*) and (*enter*, *accommodate*). But since the moves are sequential, the first of these relies on a non-credible threat. Once the entrant has entered, the incumbent does better by accommodating.

Predation

	Fight	*Accommodate*
Stay out	$0, 2$	$0, 2$
Enter	$-3, -1$	$2, 1$

There are Nash equilibria at the strategy pairs (*out*, *fight*) and (*enter*, *accommodate*). But since the moves are sequential, the first of these relies on a non-credible threat. Once the entrant has entered, the incumbent does better by accommodating.

- The Nash equilibrium solution concept admits such insensible predictions in sequential games.
- Can we find a solution concept that is 'better' in these settings?

Extensive form

First let's look at a natural way to represent sequential-move games.

- The **extensive form** of a game uses a **game tree** to represent the order of the game.
- Like the matrix representation of normal form, it captures the players, their possible moves, and the payoffs to each strategy profile.
- But it also captures who moves when and what they will have already seen before they are called on to make some move.

Extensive form of the predation game

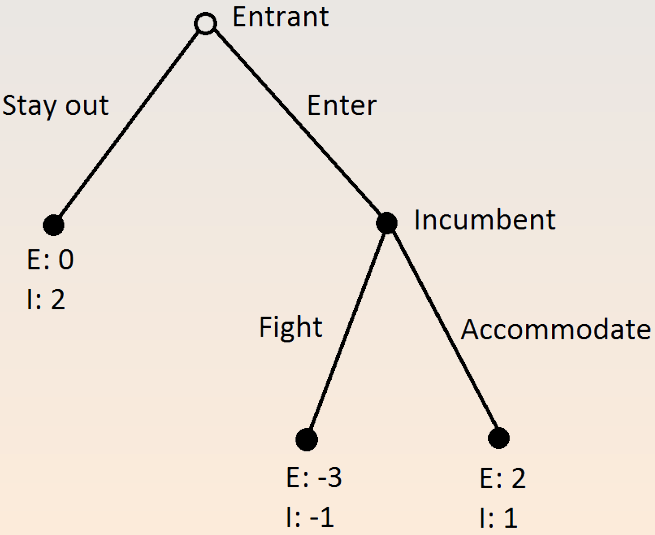

Extensive form terminology

Some terminology for extensive form:

- A **decision node** is a point at which a player is called upon to move.
- The **initial decision node** is the first such point.
- A **branch** represents a choice from a decision node.
- A **terminal node** is a point at which there are no further moves; the game ends with payoffs to players specified at the terminal node.
- An **information set** links decision nodes for player i such that when play reaches one decision node in the information set, player i does not not know which decision node in the information set has been reached.

Extensive form terminology

Some terminology for extensive form:

- A **decision node** is a point at which a player is called upon to move.
- The **initial decision node** is the first such point.
- A **branch** represents a choice from a decision node.
- A **terminal node** is a point at which there are no further moves; the game ends with payoffs to players specified at the terminal node.
- An **information set** links decision nodes for player i such that when play reaches one decision node in the information set, player i does not not know which decision node in the information set has been reached.

Today we are looking only at games of perfect information, so that we don't have to worry about information sets: each information set contains only one decision node.

Sequential rationality

A player's strategy is **sequentially rational** if it is optimal at every point in the game tree.

Sequential rationality

A player's strategy is **sequentially rational** if it is optimal at every point in the game tree.

Given that the game has arrived at some point in the game tree, sequentially rational strategies are those that prescribe optimal play from that point on, given opponents' strategies.

Sequential rationality

A player's strategy is **sequentially rational** if it is optimal at every point in the game tree.

Given that the game has arrived at some point in the game tree, sequentially rational strategies are those that prescribe optimal play from that point on, given opponents' strategies.

- For the incumbent to play 'fight' at the point after the incumbent has chosen to enter is *not* sequentially rational.
- Sequential rationality rules out non-credible threats.

Backward induction

A procedure that repeatedly applies sequential rationality is **backward induction**.

- Starting at the bottom of the game tree, we determine optimal actions at the final decision nodes.
- Non-optimal actions are ruled out if we require that strategies are sequentially rational, so we can ‘prune’ them from the game tree.
- We repeat this logic back up the layers of the tree.

Backward induction

A procedure that repeatedly applies sequential rationality is **backward induction**.

- Starting at the bottom of the game tree, we determine optimal actions at the final decision nodes.
- Non-optimal actions are ruled out if we require that strategies are sequentially rational, so we can 'prune' them from the game tree.
- We repeat this logic back up the layers of the tree.
- In the predation game there is only one layer to unpick in this way. Consider the more complicated game on the next slide.

Backward induction

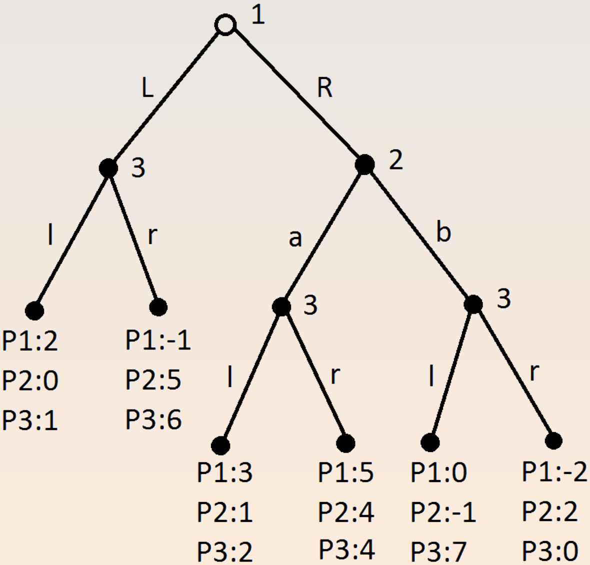

Backward induction

- At the nodes immediately before the terminal nodes, we find the optimal actions for player 3 given that we have reached each such node.
- After pruning non-optimal actions, we move back to the preceding layer and find the optimal action for player 2 given that we have reached the node at which she moves, *given* player 3's optimal actions.
- After pruning player 2's non-optimal actions, we move back again to the preceding layer and find the optimal action for player 1...

Backward induction

- At the nodes immediately before the terminal nodes, we find the optimal actions for player 3 given that we have reached each such node.
- After pruning non-optimal actions, we move back to the preceding layer and find the optimal action for player 2 given that we have reached the node at which she moves, *given* player 3's optimal actions.
- After pruning player 2's non-optimal actions, we move back again to the preceding layer and find the optimal action for player 1...

The unique strategy profile satisfying sequential rationality features (R, a, r) on the equilibrium path. Note that there are certainly other Nash equilibria in this game!

Backward induction

But remember: a strategy is a complete contingent plan of action. The unique strategy profile satisfying sequential rationality has

$$\sigma_1 = R \tag{17}$$

$$\sigma_2 = a \text{ if player 1 plays R} \tag{18}$$

$$\sigma_3 = \begin{cases} r \text{ if player 1 plays L} \\ r \text{ if player 1 plays R and player 2 plays a} \\ l \text{ if player 1 plays R and player 2 plays b} \end{cases} \tag{19}$$

It is very important to remember to completely specify strategies!

Backward induction

Applying sequential rationality can thus rule out 'unreasonable' Nash equilibria in sequential games.

Every finite game of perfect information has a pure strategy Nash equilibrium that can be derived through backward induction. If no player has the same payoffs at any two terminal nodes, there is a unique Nash equilibrium that can be derived through backward induction.

Backward induction

Applying sequential rationality can thus rule out 'unreasonable' Nash equilibria in sequential games.

Every finite game of perfect information has a pure strategy Nash equilibrium that can be derived through backward induction. If no player has the same payoffs at any two terminal nodes, there is a unique Nash equilibrium that can be derived through backward induction.

For finite games of perfect information, the set of Nash equilibria that survive backward induction coincide with the set of **Subgame Perfect Nash Equilibria** (SPNE).

Subgames

A **subgame** of an extensive form game is a subset of the game which

1. begins at an information set containing a single decision node, contains all the decision nodes that are its successors, and contains *only* these nodes, and
2. if it includes some node x, also includes all nodes in the information set containing node x.

Subgames

A **subgame** of an extensive form game is a subset of the game which

1. begins at an information set containing a single decision node, contains all the decision nodes that are its successors, and contains *only* these nodes, and
2. if it includes some node x, also includes all nodes in the information set containing node x.

Since we are still looking only at perfect information games, all information sets have only a single decision node. Then the part of the game from any one decision node to the end of the tree is a subgame.

Subgame Perfect Nash Equilibrium

Our next solution concept is **subgame perfect Nash equilibrium** (SPNE).

A strategy profile constitutes a SPNE of an extensive form game G if it induces a Nash equilibrium in every subgame of G.

Subgame Perfect Nash Equilibrium

Our next solution concept is **subgame perfect Nash equilibrium** (SPNE).

A strategy profile constitutes a SPNE of an extensive form game G if it induces a Nash equilibrium in every subgame of G.

- The backward induction procedure we identified earlier will find SPNE in finite games of perfect information.
- As we move backward through the tree, we are implicitly identifying subgames, and when we prune the tree we are eliminating non-NE play in those subgames.
- Next section we will analyze SPNE in games of imperfect information.

Interpreting SPNE

In our perfect information games, is SPNE a ‘good’ solution concept?

Interpreting SPNE

In our perfect information games, is SPNE a 'good' solution concept?

- Are players *able* to perform backward induction? Chess is a sequential-move, perfect information game...
- Does it yield plausible predictions?

Centipede

There are **two players**, 1 and 2. A referee puts down \$2 on the table. Player 1 can either 'take' and pocket the \$2 or 'split' and give each player \$1. If she plays 'take', the game ends. If not, the referee puts down another \$2. Player 2 can either take or split this new \$2. If he takes, the game ends. If not, the referee puts down another \$2 and player 1 chooses again. This continues until either someone plays 'take' or the referee has put down \$20, at which point the referee runs out of money.

Centipede

There are **two players**, 1 and 2. A referee puts down \$2 on the table. Player 1 can either 'take' and pocket the \$2 or 'split' and give each player \$1. If she plays 'take', the game ends. If not, the referee puts down another \$2. Player 2 can either take or split this new \$2. If he takes, the game ends. If not, the referee puts down another \$2 and player 1 chooses again. This continues until either someone plays 'take' or the referee has put down \$20, at which point the referee runs out of money.

i. Should player 1 play 'take' or 'split'?

Centipede

There are **two players**, 1 and 2. A referee puts down \$2 on the table. Player 1 can either 'take' and pocket the \$2 or 'split' and give each player \$1. If she plays 'take', the game ends. If not, the referee puts down another \$2. Player 2 can either take or split this new \$2. If he takes, the game ends. If not, the referee puts down another \$2 and player 1 chooses again. This continues until either someone plays 'take' or the referee has put down \$20, at which point the referee runs out of money.

i. Should player 1 play 'take' or 'split'?

ii. Can you suggest a real-world situation that could be modeled by this game?

Centipede

There are **two players**, 1 and 2. A referee puts down \$2 on the table. Player 1 can either 'take' and pocket the \$2 or 'split' and give each player \$1. If she plays 'take', the game ends. If not, the referee puts down another \$2. Player 2 can either take or split this new \$2. If he takes, the game ends. If not, the referee puts down another \$2 and player 1 chooses again. This continues until either someone plays 'take' or the referee has put down \$20, at which point the referee runs out of money.

i. Should player 1 play 'take' or 'split'?
ii. Can you suggest a real-world situation that could be modeled by this game?
iii. Find the unique outcome that survives backward induction in this game.

Backward induction in Centipede

- At the final decision node (after 'continue' all the way but before one more 'continue' gives both \$10) player 2 gets \$11 by stopping and \$10 by continuing
- Thus it isn't sequentially rational for player 2 to continue
- And so at the previous turn player 1 gets \$9 by stopping and \$8 by continuing...
- The unique backward induction 'solution' has both players saying 'stop' whenever it is their turn
- But this is terrible for the players... they could have had \$10 each!

Backward induction in Centipede

Is this a good predictor of play?

- If player 1 plays ‘stop’ in the very first round, it must be because she expects player 2 to play stop
- What if player 1 plays ‘continue’?
- If player 2 expected ‘stop’, how should he interpret this thing that SPNE predicted would never happen?
- Perhaps player 2 believes that this shows that there’s a chance that player 1 is crazy or stupid and will always say ‘continue’; then perhaps he should continue too...
- To be an SPNE, a strategy profile must prescribe SPNE in any subgame, even ones whose arrival has already contradicted SPNE!

Backward induction in Centipede

Is this a good predictor of play?

- If player 1 plays ‘stop’ in the very first round, it must be because she expects player 2 to play stop
- What if player 1 plays ‘continue’?
- If player 2 expected ‘stop’, how should he interpret this thing that SPNE predicted would never happen?
- Perhaps player 2 believes that this shows that there’s a chance that player 1 is crazy or stupid and will always say ‘continue’; then perhaps he should continue too...
- To be an SPNE, a strategy profile must prescribe SPNE in any subgame, even ones whose arrival has already contradicted SPNE!

Some variations on the centipede:

- What if the referee never runs out of money?
- What if the players don’t know how much money the referee has?
- What if some players are ‘crazy’ and always play ‘continue’?

Centipede

Centipede is a multistage game in which players alternately choose whether to 'take' the pie, ending the game, or 'pass', increasing the size of the pie

- First studied in Rosenthal (1981) and named 'centipede' by Ken Binmore because its extensive form looks like a bug with many legs
- These are multistage trust games that model situations in which gains grow over time but players face constant temptation to end the relationship by grabbing more

Centipede

Centipede is a multistage game in which players alternately choose whether to ‘take’ the pie, ending the game, or ‘pass’, increasing the size of the pie

- First studied in Rosenthal (1981) and named ‘centipede’ by Ken Binmore because its extensive form looks like a bug with many legs
- These are multistage trust games that model situations in which gains grow over time but players face constant temptation to end the relationship by grabbing more
- Different than a repeated prisoners’ dilemma because of the alternating nature of the choices and the increasing pie—there is a similar flavor but centipede vs PD highlights some of the key distinctions between sequential and repeated games

Lab evidence

McKelvey and Pelfrey (1992) studied the centipede experimentally

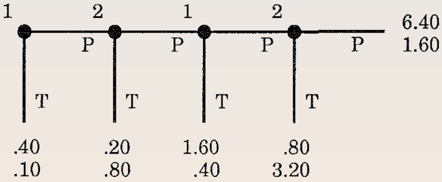

Figure 1: The Four Move Centipede Game.

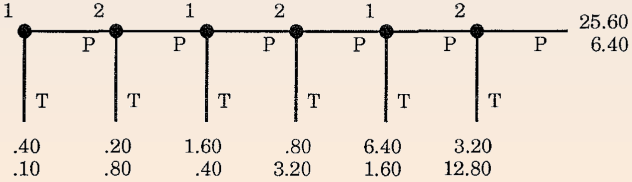

Figure 2: The Six Move Centipede Game.

Number of experiments ending at each node

	Session		n_1	n_2	n_3	n_4	n_5	n_6	n_7
Four Moves	1	(PCC)	6	26	44	20	4		
	2	(PCC)	8	31	32	9	1		
	3	(CIT)	6	43	28	14	9		
	Total		20	100	104	43	14		
High Payoff	4	(CIT)	15	37	32	11	5		
Six Moves	5	(CIT)	2	9	39	28	20	1	1
	6	(PCC)	0	2	3	37	28	9	2
	7	(PCC)	0	7	14	43	23	12	1
	Total		2	18	56	108	71	22	4

- Only 1 out of 138 participants chose ‘take’ every chance they got (i.e. the backward induction strategy)
- 9 out of 138 always chose ‘pass’ at every possibility
- If 5% were altruistic in the sense of always choosing pass, the experimental data can be well explained

Lab evidence

Fey, McKelvey, and Pelfrey (1996) study a constant-sum centipede, since fairness considerations could have been at work in the original experiments

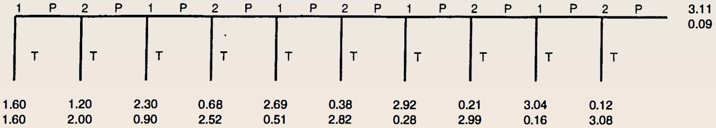

Figure 1: A Ten-Move Constant-Sum Centipede Game.

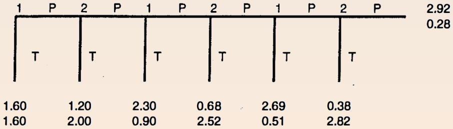

Figure 2: A Six-Move Constant-Sum Centipede Game.

Lab evidence

In the constant-sum centipede, results looked quite different

- 45 of 176 subjects chose ‘take’ at every opportunity
- 2 chose ‘pass’ at every opportunity, but none were in the first-mover role
- The altruism model proposed in the previous paper doesn’t do well here to explain the data

Making mistakes or being different

There are a couple of theoretical frameworks that are consistent with the centipede evidence

- The main idea is that if there is a sufficient chance that the other player will pass, your best response may be to pass too
- As the length of the centipede increases, the chance that the other player will pass doesn't have to be so large in order for you to want to pass
- That is: the longer the game, the bigger the potential reward if the other player keeps passing

Making mistakes or being different

There are a couple of theoretical frameworks that are consistent with the centipede evidence

- The main idea is that if there is a sufficient chance that the other player will pass, your best response may be to pass too
- As the length of the centipede increases, the chance that the other player will pass doesn't have to be so large in order for you to want to pass
- That is: the longer the game, the bigger the potential reward if the other player keeps passing
- Two possible reasons why:
  1. Mistakes: the player tries to play 'rationally' but just messes up sometimes
  2. Types: some players might have a different goal in mind than maximizing their own cash payoff

'Lab' evidence with sophisticated players

Two papers that studied the centipede with high-level chess players are Palacios-Huerta and Volij (2009) and Levitt, List, and Sadoff (2011); we will look at these in turn

- The idea is to find people who are experts in backward induction
- ... and *known* to be experts: we need common knowledge of rationality for backward induction to be a valid model (Aumann 1995)

'Lab' evidence with sophisticated players

Two papers that studied the centipede with high-level chess players are Palacios-Huerta and Volij (2009) and Levitt, List, and Sadoff (2011); we will look at these in turn

- The idea is to find people who are experts in backward induction
- ... and *known* to be experts: we need common knowledge of rationality for backward induction to be a valid model (Aumann 1995)
- Goal: vary 'distance' from common knowledge of rationality to help to distinguish this explanation for failure of SPNE in the centipede from social preferences
- Field experiment: chess players at tournaments matched up for one play of the centipede game; comparison group of college students playing a lab version
- Lab experiment: students and chess players matched up for ten rounds against different opponents each time; order of play of the two types of subjects varied across treatments

Palacios-Huerta and Volij (2009) field experiment

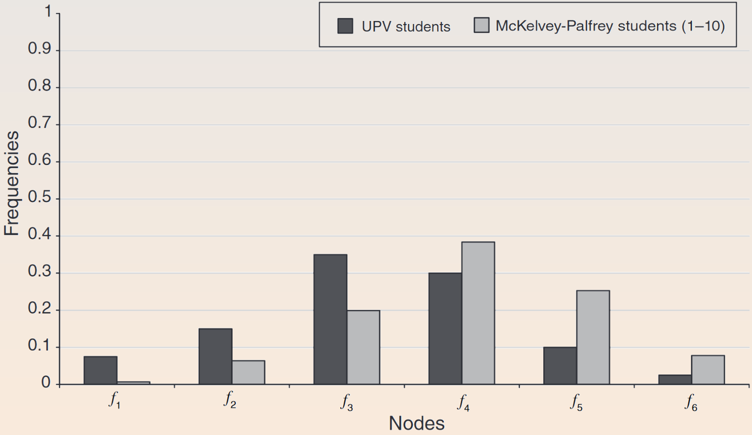

FIGURE 2. COLLEGE STUDENTS: PROPORTION OF OBSERVATIONS AT EACH TERMINAL NODE

Palacios-Huerta and Volij (2009) field experiment

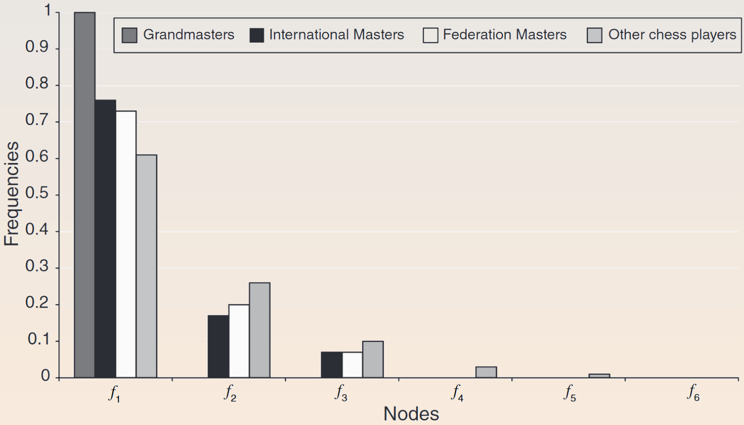

FIGURE 3. CHESS PLAYERS: PROPORTION OF OBSERVATIONS AT EACH TERMINAL NODE BY TYPE OF PLAYER 1 IN THE PAIR

Palacios-Huerta and Volij (2009) lab experiment

Treatment 1: student vs. student; treatment 2: student vs. chess player

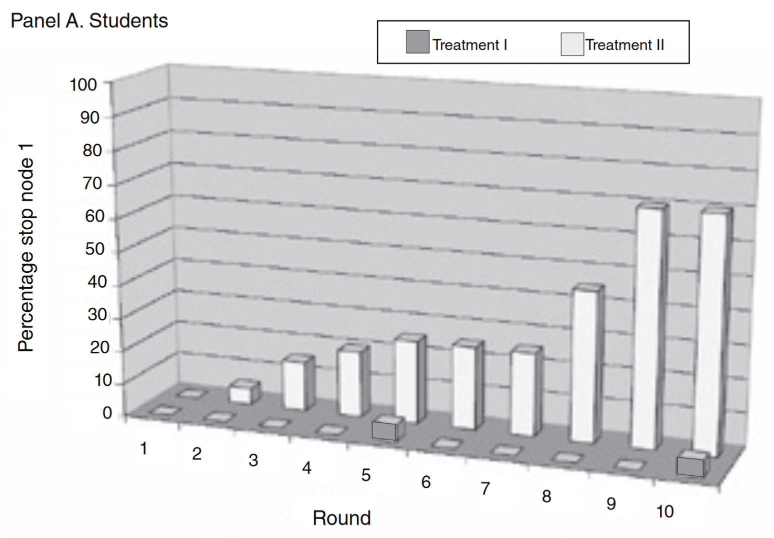

Palacios-Huerta and Volij (2009) lab experiment

Treatment 3: chess player vs. student; treatment 4: chess player vs. chess player

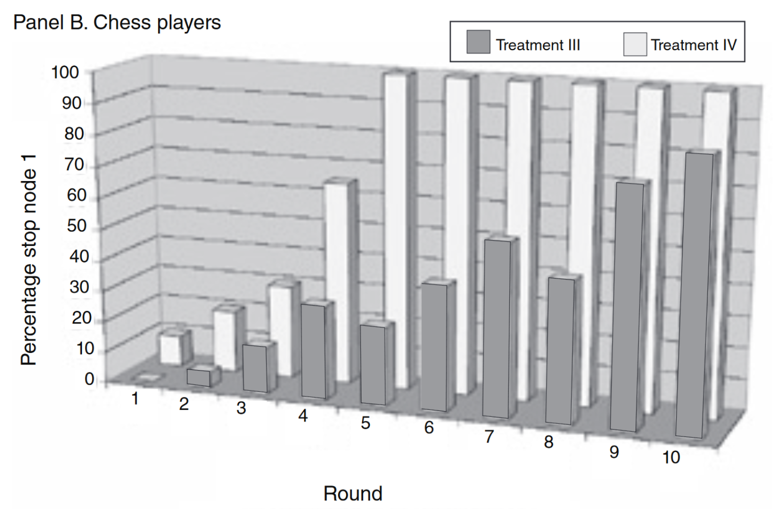

Levitt, List, and Sadoff (2011)

Again recruited subjects from chess tournaments, focusing on highly ranked players

- Two conference rooms, subjects in different rooms; didn't know who they were paired with
- Each pair played one centipede and two versions of the race to 100 in a random order
- Designed to compare the centipede to the 'race to 100 game' which is a constant sum winner-take-all game with tough backward induction logic
- Decisions communicated by instant message on experimenter-controlled computers
- Players recorded decisions of themselves and their opponent along the way; correctness confirmed by the experimenter

The race to 100 game

- Two players
- Alternate choosing numbers in a given range
- This paper: ranges 1-10 and 1-9 (to vary the difficulty of the backward induction steps)
- Running total is kept
- Winner is the one who chooses a number that makes the total add up to exactly 100

The race to 100 game

- Two players
- Alternate choosing numbers in a given range
- This paper: ranges 1-10 and 1-9 (to vary the difficulty of the backward induction steps)
- Running total is kept
- Winner is the one who chooses a number that makes the total add up to exactly 100
- Helpful because optimal strategy is robust to all but the most extreme assumptions on players' preferences
- Requires 10 steps of backward induction to 'solve'

Levitt, List, and Sadoff (2011) centipede

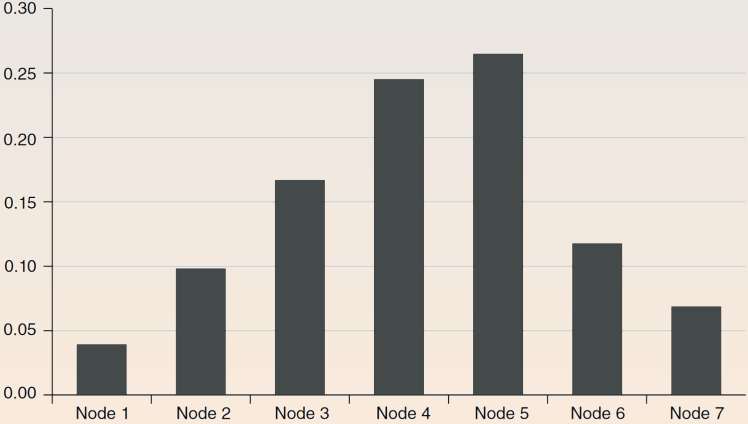

Figure 2. Distribution of Centipede Game Stopping Nodes

Levitt, List, and Sadoff (2011) race to 100

TABLE 3—SUMMARY OF RACE TO 100 RESULTS

Node at which game solved	1	2	3	4	5	6	7	8	9
Number at which game solved (1–9)	10	20	30	40	50	60	70	80	90
Percentage of time solved (1–9)	0.573	0.087	0.029	0.039	0.019	0.01	0.029	0.078	0.136
Number at which game solved (1–10)	1	12	23	34	45	56	67	78	89
Percentage of time solved (1–10)	0.126	0.087	0.019	0.01	0.01	0.029	0.049	0.214	0.447

Notes: Table 3 reports the distribution of nodes at which a race to 100 game was solved. Rows 1 and 3 report the "key number" from which a win may be forced. Rows 2 and 4 report the percent of the time a corresponding game was solved at that node; a solution is a choice of number that summed to a "key number," in conjunction with never deviating from subsequent "key numbers" afterward.

TABLE 6—PERCENTAGE OF GAMES SOLVED AT EACH NODE

	1	2	3	4	5	6	7	8	9
Race to 100 1–9	10	20	30	40	50	60	70	80	90
1–9 Played first ($N = 52$)	0.615	0.115	0.038	0.019	0.019	0	0.019	0.058	0.115
1–9 Played second ($N = 51$)	0.529	0.059	0.02	0.059	0.02	0.02	0.039	0.098	0.157
Race to 100 1–10	1	12	23	34	45	56	67	78	89
1–10 Played first ($N = 51$)	0.078	0.059	0	0.02	0	0.059	0.039	0.176	0.569
1–10 Played second ($N = 51$)	0.176	0.118	0.039	0	0.02	0	0.059	0.255	0.333

Notes: Table 6 reports the distribution of when a race was solved, conditional on the order races were played in. Columns correspond to the "key number" at which a game was solved. One pair of players dropped out after playing the race (1–9) first, which accounts for the differing number of observations.

Notice the difference between the 'easier' and more difficult versions

Levitt, List, and Sadoff (2011) comparison

TABLE 7—CENTIPEDE BEHAVIOR BY INDUCTION ABILITY: IMPLIED STOP PROBABILITIES

	N	F1	F2	F3	F4	F5	F6
Best	15	0	0.2	0.125	0.25	0.667	0
		(10)	(5)	(8)	(4)	(3)	(2)
Second best	66	0	0.106	0.375	0.412	0.571	0.889
		(17)	(47)	(16)	(34)	(7)	(9)
Second worst	36	0.1	0.133	0.278	0.417	0.833	1
		(20)	(15)	(18)	(12)	(12)	(1)
Bad	87	0.036	0.065	0.109	0.238	0.458	0.429
		(55)	(31)	(46)	(21)	(24)	(7)

Notes: Table 7 displays implied stop probability by inductor ability, rather than title for players. Odd numbered columns refer to player 1's decisions; even numbered columns refer to player 2's decisions.

The Chainstore Paradox

We can also apply SPNE to repetitions of games.

- Consider a game that consists of the Predation game from earlier repeated 10 times.
- For example, a chainstore is trying to deter entry by separate local competitors into 10 markets where it has outlets.
- The unique SPNE in Predation was for the entrant to enter and the incumbent to accommodate.
- But perhaps with 10 markets the chainstore will fight the first entrant to try to deter the other 9 from trying?

This is the famous **Chainstore Paradox** (Selten 1978).

The Chainstore Paradox

- Consider the 10th market. By backward induction, the last game is identical to the one-shot version.
- The unique SPNE in the subgame consisting of the last repetition is at $(enter, accommodate)$.
- In the 9th repetition, the incumbent has no incentive to fight entry to try to deter the 10th entrant, since the 10th entrant will be accommodated anyway.
- We can apply this logic right back to the first market.
- The unique SPNE features $(enter, accommodate)$ in every market.

Repeated games

This equally applies to repeated simultaneous-move games.

Repeated games

This equally applies to repeated simultaneous-move games.

- Consider a repeated Prisoner's Dilemma.

	C	D
C	$-1, -1$	$-10, 0$
D	$0, -10$	$-5, -5$

Repeated games

This equally applies to repeated simultaneous-move games.

- Consider a repeated Prisoner's Dilemma.

	C	D
C	$-1, -1$	$-10, 0$
D	$0, -10$	$-5, -5$

- The unique Nash equilibrium has both players choose D.
- Notice that strategies can be conditional on history and so can become *very* complicated. To see why, imagine what the extensive form looks like...
- The five-round repeated Prisoner's Dilemma has a strategy set for each player with over two billion strategies!

Finitely repeated games

What if the game is repeated T times?

Finitely repeated games

What if the game is repeated T times?

- As in the chainstore game, we can apply backward induction and SPNE: each repetition is the beginning of a subgame that consists of all remaining repetitions.
- Consider the last period: players have no incentive to cooperate since there is no tomorrow.
- Thus in SPNE (D, D) must be played in the last period.
- In the second-to-last period, players have no incentive to cooperate since they will not cooperate tomorrow...

In both cases, this is not perhaps what we'd expect to observe with real players. Why does SPNE not seem entirely satisfying?

The Flood-Dresher experiment

Flood (1952) is an early reported repeated prisoner's dilemma

- 100 rounds (known in advance); payoffs asymmetric (relabeled here for clarity)

	L	R
U	$-1, 2$	$0.5, 1$
D	$0, 0.5$	$1, -1$

- It is notable because the players were two friends, an economist and a mathematician, and they were asked to keep a running log of their personal comments on what was happening
- Players were familiar with zero-sum game theory but not with the full extent of what were then the new findings from Nash
- Payoffs concocted deliberately to include 'split-the-difference' solution

Flood-Dresher

Strategy Frequencies

AA	JW 1	JW 2	Total
1	8	60	68
2	14	18	32
Total.	22	78	100

Start of comments (table from De Herdt 2003)

Table 1: The Flood–Dresher experiment: moves and comments°

Game	*Moves*		*Armen Alchian's comments*	*John Williams's comments*
	AA	*JW*		
1	D	C	JW will play D—sure win. Hence if I play C—I lose.	Hope he's bright.
2	D	C	What is he doing ?!!	He isn't but maybe he'll wise up
3	D	D	Trying mixed ?	Okay, dope.
4	D	D	Has he settled on D ?	Okay, dope.
5	C	D	Perverse!	It isn't the best of all possible worlds.
6	D	C	I'm sticking to D since he will mix for at least 4 more times.	Oh ho! Guess I'll have to give him another chance.
7	D	C		Cagey, ain't he ? Well...
8	D	D		In time he could learn, but not in ten moves so:
9	D	D	If I mix occasionally, he will switch—but why will he ever switch from D ?	

Response from Nash

[11] Dr. Nash makes the following comment (private communication) on this experiment:

"The flaw in this experiment as a test of equilibrium point theory is that the experiment really amounts to having the players play one large multimove game. One cannot just as well think of the thing as a sequence of independent games as one can in zero-sum cases. There is much too much interaction, which is obvious in the results of the experiment.

"Viewing it as a multimove game a strategy is a complete program of action, including reactions to what the other player has done. In this view it is still true the only real absolute equilibrium point is for A (AA) always to play 2, B (JW) always 1.

"However, the strategies:

A plays 1 'til B plays 1, then 2 ever after,

B plays 2 'til A plays 2, then 1 ever after,

are very nearly at equilibrium and in a game with an indeterminate stop point or an infinite game with interest on utility it *is* an equilibrium point.

Response from Nash

"Since 100 trials are so long that the Hangman's Paradox cannot possibly be well reasoned through on it, it's fairly clear that one should expect an approximation to this behavior which is most appropriate for indeterminate end games with a little flurry of aggressiveness at the end and perhaps a few sallies, to test the opponent's mettle during the game.

"It is really striking, however, how inefficient AA and JW were in obtaining the rewards. One would have thought them more rational.

"If this experiment were conducted with various different players rotating the competition and with *no information given to a player of what choices the others have been making until the end* of all the trials, then the experimental results would have been quite different, for this modification of procedure would remove the interaction between the trials."

Dr. Dresher and I were glad to receive these comments, and to include them here, even though we would not change our interpretation of the experiment along the lines indicated by Dr. Nash.

Flood-Dresher

- Interesting since there is not yet a totally ingrained received wisdom
- Game theory seems at that point ‘up for grabs’ relative to how the Nash program came to dominate
- Responses from Nash include several chapters worth of material for a game theory textbook...
- Eliciting the commentary makes this a really fun artifact—but is the game with prompt for introspective commentary the same game? How can we elicit comments in our experiments?

SPNE in finitely repeated games

- If the repeated game has more than one Nash equilibrium, SPNE is less restrictive.
- This is because more than one payoff can be realized in the final subgame...
- Strategies can then promise rewards and punishments without violating that each subgame must feature a Nash equilibrium.

SPNE in finitely repeated games

- If the repeated game has more than one Nash equilibrium, SPNE is less restrictive.
- This is because more than one payoff can be realized in the final subgame...
- Strategies can then promise rewards and punishments without violating that each subgame must feature a Nash equilibrium.

What if a game is repeated indefinitely?

- There is no 'end' from which to apply backward induction and pin down SPNE.

Infinitely repeated Prisoner's Dilemma

Consider the infinitely repeated Prisoner's Dilemma.

	C	D
C	$1, 1$	$-1, 3$
D	$3, -1$	$0, 0$

Infinitely repeated Prisoner's Dilemma

Consider the infinitely repeated Prisoner's Dilemma.

	C	D
C	$1, 1$	$-1, 3$
D	$3, -1$	$0, 0$

What if player 2 plays a strategy X that says "Play C unless my opponent has ever played D, then play D forever" ("Grim" strategy)?

Infinitely repeated Prisoner's Dilemma

Consider the infinitely repeated Prisoner's Dilemma.

	C	D
C	$1, 1$	$-1, 3$
D	$3, -1$	$0, 0$

What if player 2 plays a strategy X that says "Play C unless my opponent has ever played D, then play D forever" ("Grim" strategy)? Consider some period before which both players have always played C. Player 1's payoffs to also playing X and to playing D forever:

$$\pi_1(X) = 1 + 1 + 1 + 1 + ... \tag{20}$$

$$\pi_2(D) = 3 + 0 + 0 + 0 + ... \tag{21}$$

For some discount rate δ, X dominates D if

$$\frac{1}{1-\delta} > 3 \tag{22}$$

$$\delta > \frac{2}{3}. \tag{23}$$

Collusion

We can apply this logic to analyze collusion: consider an infinitely repeated version of the Bertrand game we saw last time.

- The unique one-shot Nash equilibrium featured both firms pricing at cost and making zero profits.
- Can there be an equilibrium in which both set some high price p^* and earn positive profits $\pi^* > 0$?

Collusion

Say firm i is playing a strategy that says "set a high price, unless my opponent has ever undercut, then play the Bertrand price forever."

- If firm j plays the same strategy it earns π^* every period, forever.
- If firm j undercuts, it earns a higher profit today (π_H) and at most zero for each subsequent period.

The 'collude' strategy beats undercutting if:

$$\frac{\pi^*}{1-\delta} > \pi_H \tag{24}$$

$$\delta > 1 - \frac{\pi^*}{\pi_H} \tag{25}$$

A *sufficiently patient* firm prefers to collude than undercut. Note that this 'collusion' requires no talking!

Collusion

Say firm *i* is playing a strategy that says "set a high price, unless my opponent has ever undercut, then play the Bertrand price forever."

- If firm *j* plays the same strategy it earns π^* every period, forever.
- If firm *j* undercuts, it earns a higher profit today (π_H) and at most zero for each subsequent period.

The 'collude' strategy beats undercutting if:

$$\frac{\pi^*}{1-\delta} > \pi_H \tag{24}$$

$$\delta > 1 - \frac{\pi^*}{\pi_H} \tag{25}$$

A *sufficiently patient* firm prefers to collude than undercut. Note that this 'collusion' requires no talking!

- We can also interpret δ as the probability that the game will end today...

Punishment and enforcement

Variations on the same framework can be applied to *community enforcement*: the idea that our friends have our backs

- e.g. repeated interaction between a business and a population of potential customers
- Temptation to rip off the customer, but long term gains from playing nice

Punishment and enforcement

Variations on the same framework can be applied to *community enforcement*: the idea that our friends have our backs

- e.g. repeated interaction between a business and a population of potential customers
- Temptation to rip off the customer, but long term gains from playing nice
- Threat of lost business from the customer and her friends (Ali and Miller 2015)
- If the firm cares enough about the future, this can enforce good behavior (Kandori 1992)
- But fly-by-night firms and snake oil salesmen may strike...
- One of my papers is on how the availability of social network data affects the incentives to rip off customers
- What are the costs and benefits of swindling a consumer? It depends on where they are positioned in the network
- Applying standard repeated game concepts of reputation and trigger strategies to the networked model

The Folk Theorem

These ideas in infinitely repeated games are captured in the **Folk Theorem**. For two player games:

For any feasible pair of individually rational payoffs $(\pi_1, \pi_2) \gg (\underline{\pi_1}, \underline{\pi_2})$, there exists a $\underline{\delta}$ such that, for all $\delta > \underline{\delta}$, (π_1, π_2) are the average payoffs arising in an SPNE.

The Folk Theorem

These ideas in infinitely repeated games are captured in the **Folk Theorem**. For two player games:

For any feasible pair of individually rational payoffs $(\pi_1, \pi_2) \gg (\underline{\pi_1}, \underline{\pi_2})$, there exists a $\underline{\delta}$ such that, for all $\delta > \underline{\delta}$, (π_1, π_2) are the average payoffs arising in an SPNE.

Any average payoff bigger than what I can guarantee alone can be supported in SPNE for sufficiently patient players.

- Intuitively: the more patient the player, the harder punishment bites, and so the smaller is the payoff in excess of the punishment payoff that the threat of punishment can sustain.
- This also implies that punishment strategies need not be so brutal as the grim strategy above.

The shadow of the future

Dal Bó (2005) studies an ‘infinitely repeated’ prisoner’s dilemma in the lab

- Random continuation rule stands in for discounting
  - In the ‘infinite’ treatments, $\delta = 0, \frac{1}{2}, \frac{3}{4}$ is a commonly known probability that the game will continue after each round
- Finite treatments with the same expected length as each of the ‘infinite’ treatments admits a clean comparison: no ‘length of game’ effect
- Two different payoff matrices used; which strategies are equilibrium strategies for $\delta = \frac{1}{2}$ is different in the two games
- Subjects quite good at understanding the expected length of ‘infinite’ games

The effect of continuation probability

TABLE 5—PERCENTAGE OF COOPERATION BY MATCH AND TREATMENT*

		Match									
		1	2	3	4	5	6	7	8	9	10
Dice	$\delta = 0$	26.26	18.18	10.61	11.62	12.63	12.63	5.56	5.26	5.26	5
	$\delta = 1/2$	28.36	27.12	34.58	35.53	21.60	19.08	29.84	35.96	28.16	50
	$\delta = 3/4$	40.44	28.57	27.78	32.92	46.51	33.09	44.05	53.51	42.26	45.83
Finite	$H = 1$	26.56	18.23	16.67	17.19	11.98	8.02	6.79	10.49	6.14	6.67
	$H = 2$	19.79	15.89	14.84	9.64	11.46	10.80	12.04	10.19	6.58	6.67
	$H = 4$	31.64	30.34	30.47	25.52	25.13	23.77	16.36	19.75	14.91	20.83

* *All rounds and sessions.*

TABLE 6—PERCENTAGE OF COOPERATION BY ROUND AND TREATMENT*

		Round											
		1	2	3	4	5	6	7	8	9	10	11	12
Dice	$\delta = 0$	9.17											
	$\delta = 1/2$	30.93	26.10	19.87	12.50	12.96							
	$\delta = 3/4$	46.20	40.76	38.76	34.58	33.04	27.27	24.75	26.28	29.17	26.04	32.29	31.25
Finite	$H = 1$	10.34											
	$H = 2$	13.31	6.90										
	$H = 4$	34.58	21.55	18.97	10.63								

* *All sessions, matches four through ten.*

Cooperation in finite time

Consider again the finitely repeated Prisoner's Dilemma, played between 2 players $T+1$ times, from $t=0$ to $t=T$. Payoffs are the sum of the stage payoffs.

	C	D
C	$1,1$	$-1,3$
D	$3,-1$	$0,0$

We know that the unique SPNE strategy profile is (D,D) forever. This is the chain store result.

Cooperation in finite time

	C	D
C	$1, 1$	$-1, 3$
D	$3, -1$	$0, 0$

What if player 1 believes that there is some probability α that 2 will *not* play this SPNE strategy but instead will play the grim profile?

- Player 1 considers two options: "play SPNE" (A) or "play Grim until the last period, then defect" (B).

Cooperation in finite time

	C	D
C	$1, 1$	$-1, 3$
D	$3, -1$	$0, 0$

What if player 1 believes that there is some probability α that 2 will *not* play this SPNE strategy but instead will play the grim profile?

- Player 1 considers two options: "play SPNE" (A) or "play Grim until the last period, then defect" (B).

$$\pi_A = \underbrace{\alpha[3 + 0.T]}_{\text{meet grim player}} + \underbrace{(1-\alpha)[0.(T+1)]}_{\text{meet SPNE player}} \tag{26}$$

$$\pi_B = \underbrace{\alpha[1.(T) + 3]}_{\text{meet grim player}} + \underbrace{(1-\alpha)[-1 + 0.T]}_{\text{meet SPNE player}} \tag{27}$$

Cooperation in finite time

The payoff to B beats the payoff to A if

$$\alpha[1.(T) + 3] + (1-\alpha)[-1 + 0.T] > \alpha[3 + 0.T] + (1-\alpha)[0.(T+1)] \quad (28)$$

$$(T+1)\alpha > 1 \quad (29)$$

Cooperation in finite time

The payoff to B beats the payoff to A if

$$\alpha[1.(T) + 3] + (1-\alpha)[-1 + 0.T] > \alpha[3 + 0.T] + (1-\alpha)[0.(T+1)] \quad (28)$$

$$(T+1)\alpha > 1 \quad (29)$$

- If the game goes on sufficiently long, or there is a sufficiently high chance of meeting the ‘grim’ player, the expected payoff to trying the ‘play nice then punish’ strategy is higher than the expected payoff to always playing SPNE
- Willing to try being nice in the hope of meeting a nice opponent
- Shades of level k?