Rational Heuristics for One-Shot Games

Frederick Callaway [1] Thomas L. Griffiths [1] **Gustav Karreskog Rehbinder** [2]
EEA-ESEM, August 31 2023

[1]Department of Psychology, Princeton University

[2]Department of Economics, Uppsala University

Motivation

- Perfect rationality is an insufficient model of experimental decision making.
- Biases and deviations from perfect rationality vary among settings.
- We need to understand when and why deviations from perfect rationality happen.
- Optimal use of limited cognitive resources is a potential answer.
- In particular, we think that optimal use of simple decision procedures (heuristics) can explain behavior.

Hypothesis
People make decisions using simple heuristics that are adapted to their environment in order to trade off cognitive costs and expected payoffs.

This Paper

- We present a theory that combines two perspectives:
  - Decisions by heuristics.
  - Heuristics are adapted to the environment.
- We combine these perspectives in a way that allow us to make quantitative predictions about behavior.
- We focus on one-shot games as a way of testing this hypothesis experimentally.

Introductory Example

8, 8	0, 4	0, 6	= 0.8
4, 0	4, 4	4, 6	
6, 0	6, 4	2, 2	
10%	80%	10%	

What would you do as a row player in this game?

Introductory Example

8, 8	0, 4	0, 6
4, 0	4, 4	4, 6
6, 0	6, 4	2, 2

Row 1 has an efficient outcome.

Introductory Example

8, 8	0, 4	0, 6
4, 0	4, 4	4, 6
6, 0	6, 4	2, 2

Row 1 has an efficient outcome.
But also a lot of zeros.

Introductory Example

8, 8	0, 4	0, 6
4, 0	4, 4	4, 6
6, 0	6, 4	2, 2

Row 2 has a guaranteed payoff of 4.

Introductory Example

8, 8	0, 4	0, 6	= 0.8
4, 0	4, 4	4, 6	
6, 0	6, 4	2, 2	
10%	80%	10%	

Row 3 has the highest average payoff.

Introductory Example

8, 8	0, 4	0, 6
4, 0	4, 4	4, 6
6, 0	6, 4	2, 2
10%	80%	10%

Maybe you believe that the other person is likely to be cautious.

Introductory Example

8, 8	0, 4	0, 6	= 0.8
4, 0	4, 4	4, 6	= 4
6, 0	6, 4	2, 2	= 5.6
10%	80%	10%	

Maybe you believe that the other person is likely to be cautious.
And best respond to those beliefs with row 3.

Introductory Example

$8, 8$	$0, 4$	$0, 6$
$4, 0$	$4, 4$	$4, 6$
$6, 0$	$6, 4$	$2, 2$

These are examples of possible heuristics that make sense in different settings:

- Look for the common interest.
- Maximize guaranteed payoff.
- Pick the highest row average.
- Apply an heuristic to the column player to form beliefs, and respond to those beliefs.

General Model

- A *heuristic* is a function h from a game to a probability distribution over possible actions.
- Each heuristic has an associated cognitive cost $c(h) \in \mathbb{R}_+$.
- For a given game G and opponent heuristic h_{opp}, the expected utility from using heuristic h is

$$u\left(h, h_{opp}, G\right) = \pi_G\left(h(G), h_{opp}(G^T)\right) - c(h).$$

- An *environment* is a collection $\mathcal{E} = (\mathcal{G}, \mathcal{H}, P)$, where $\mathcal{G}$ is the set of games, $\mathcal{H}$ the set of heuristics, and P the joint probability distribution.

The rational heuristic in environment $\mathcal{E}$ is given by

$$h^*(\mathcal{E}) = \arg\max_{h \in \mathcal{H}} V(h, \mathcal{E}) = \arg\max_{h \in \mathcal{H}} \mathbb{E}_{\mathcal{E}}\left[u\left(h(G), h_{opp}(G^T), G\right)\right]$$

Experiment I

- The rational heuristic $h^*(\mathcal{E})$ depends on the environment.
- With time to adapt, behavior in a given game G should thus depend on the environment.
- We use this predicted variation in behavior to test our theory in two ways.
  1. By embedding the same game in different environments.
  2. By estimating $h^*(\mathcal{E})$ in different environments and perform out of sample predictions.

Experiment II

- Preregistered web-based experiment with 600 participants.
- 300 participants per treatment divided into 10 sessions.
- 50 randomly matched one-shot games per session.
- Four *comparison games* played by everyone in rounds 31, 38, 42, and 49.
- Two different treatment environments generated by sampling the remaining *treatment games* on the session level.
  - Common interest distribution ($\mathcal{E}^+$)
  - Competing interest distribution ($\mathcal{E}^-$)
- Model free tests using the comparison games.
- Model based tests by making predictions using $h^*(\mathcal{E}^+)$ and $h^*(\mathcal{E}^-)$.

Examples of Common Interest Games

5, 6	6, 4	5, 3
9, 4	5, 5	6, 7
2, 0	0, 1	6, 4

Common interest example 1

3, 4	5, 5	9, 7
4, 2	5, 7	5, 7
2, 4	2, 1	2, 3

Common interest example 2

9, 7	5, 9	7, 8
6, 7	9, 9	4, 6
6, 4	3, 1	6, 2

Common interest example 3

1, 4	5, 3	7, 4
3, 5	4, 2	7, 5
3, 8	3, 6	5, 3

Common interest example 4

Examples of Competing Interest Games

5, 5	6, 2	5, 3
5, 3	1, 8	8, 4
3, 6	7, 4	4, 6

Competing interest example 1

2, 4	4, 4	4, 6
1, 7	2, 6	9, 1
7, 1	4, 8	8, 6

Competing interest example 2

4, 5	1, 5	7, 1
2, 7	8, 5	5, 7
2, 6	8, 3	3, 9

Competing interest example 3

8, 0	4, 1	3, 8
4, 7	2, 7	2, 7
3, 5	3, 9	7, 5

Competing interest example 4

Model-Free Hypotheses I

- In a given game with a risky efficient outcome and safer actions we should expect:
  - In common interest environment: Focus on the efficient common interest outcome.
  - In competing interest environment: Focus on high guaranteed payoff.
- This intuition leads us to our four comparison games.

Model-Free Hypotheses II

Hypothesis (1)
The distribution of play in each of the four comparison games will be different in the two treatment populations.

Hypothesis (2)
The average payoff in the four comparison games will be higher in the common interest treatment than in the competing interests treatment.

Same Game in Different Environments I

	Common interest			Competing interest			
Strategies	**1**	**2**	**3**	**1**	**2**	**3**	
Frequencies	193	53	54	75	82	143	$p < 0.001$
Payoffs	5.09			3.64			$p < 0.001$

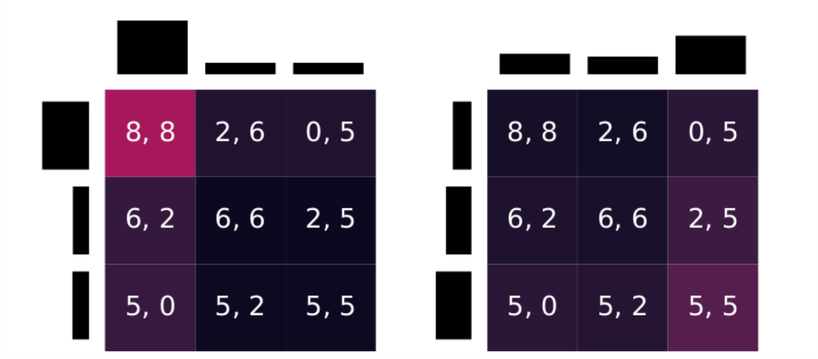

Comparison Game 1 - Weak link game

Same Game in Different Environments II

	Common interest			Competing interest			
Strategies	**1**	**2**	**3**	**1**	**2**	**3**	
Frequencies	160	139	1	103	195	2	$p < 0.001$
Payoffs		5.52			4.04		$p < 0.001$

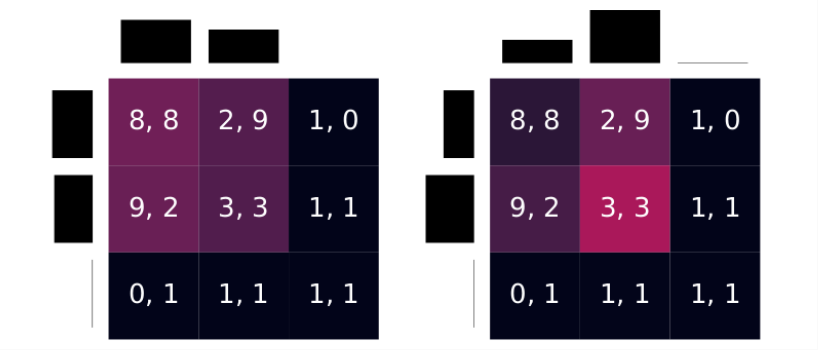

Comparison Game 2 - Prisoner's dilemma

Same Game in Different Environments III

	Common interest			Competing interest			
Strategies	**1**	**2**	**3**	**1**	**2**	**3**	
Frequencies	40	73	187	106	97	97	$p < 0.001$
Payoffs		5.00			4.31		$p = 0.004$

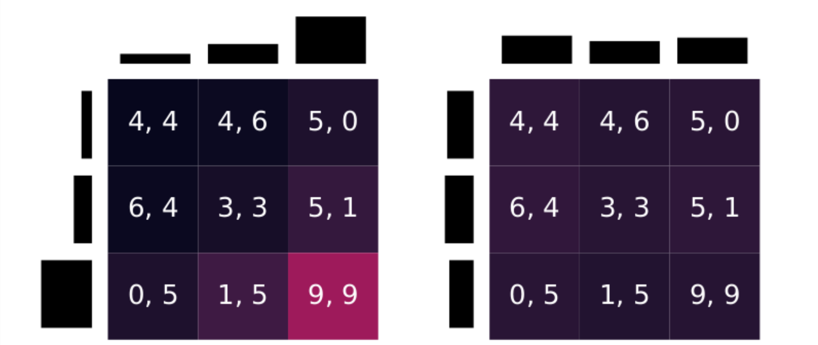

Comparison Game 3 - High risk NE efficient outcome

Same Game in Different Environments IV

	Common interest			Competing interest			
Strategies	**1**	**2**	**3**	**1**	**2**	**3**	
Frequencies	78	173	49	115	62	123	$p < 0.001$
Payoffs	5.19			3.42			$p < 0.001$

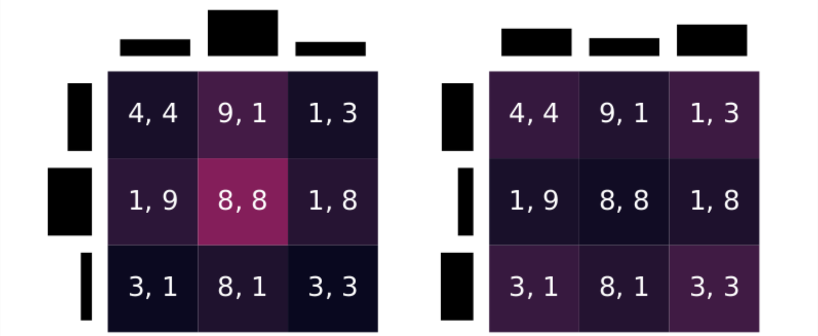

Comparison Game 4 - High risk non-NE efficient outcome

Summary of Model Free Tests

- Clear effect of the environment, different modal actions in all games.
- Differences in line with our predictions.
- Common interest environment $\implies$ coordinate on mutually beneficial outcome
- Competing interest environment $\implies$ take safe actions that are less efficient
- Game theoretical considerations not of first order importance.

Model-based Testing I

- We have seen that there are strong treatment effects.
- If we specify $\mathcal{H}$, $c(h)$ and $\mathcal{E}$, we can take things further and make quantitative predictions.
- We take an out of sample prediction approach for evaluating the models.
- We do this via a preregistered train/test split. The early rounds (first 30) are used as training data, and the later (last 16) are used as testing data.
- Denote an empirical environment as for example $\mathcal{E}^{+}_{\text{train}}$.

Model-based Testing II

- If different heuristics are used in the two environments, we should be able to predict these differences.

- In particular, heuristics trained on $\mathcal{E}^{+}_{\text{train}}$ should predict $\mathcal{E}^{+}_{\text{test}}$ better than if trained on $\mathcal{E}^{-}_{\text{train}}$.

- Two ways of estimating:
  - **Fitting:** Find the heuristics and joint cognitive costs that best match the training data.
  - **Optimizing:** Find the joint cognitive costs such that the rational heuristics best match the training data.

Model-based Testing III

- We try two completely different explicit models of $\mathcal{H}$ and $c(h)$.
- *Metaheuristics* is our primary model. We assume a set of parameterized primitive heuristics, that combine into a larger heuristic. The primitive heuristics are: *row*, *cell*, and *simulation* heuristics.
- *Deep heuristics* are based on a special neural network design. Allows for a much broader set of heuristics, at the cost of interpretability.
- Both models confirm our hypotheses.

Out of Sample Performance - Metaheuristic

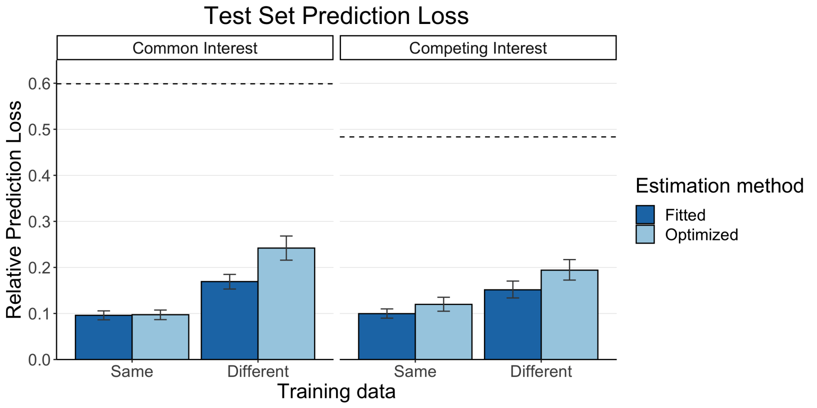

Deep Heuristic

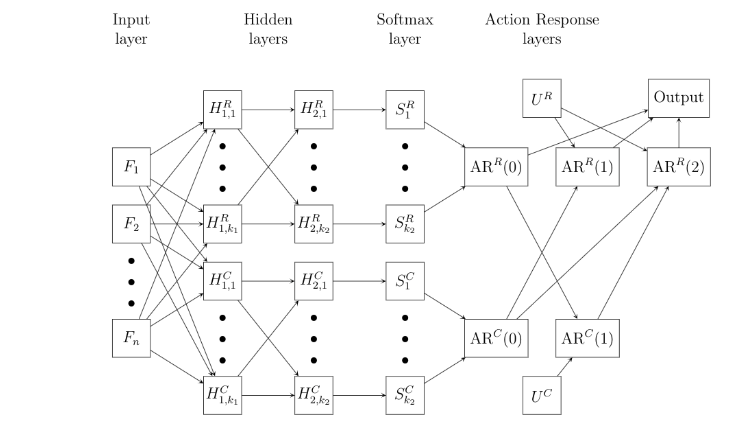

Inspired from Hartford, Wright and Leyton-Brown (2016). Cognitive cost inversely proportional to entropy of prediction plus extra cost for simulation.

Out of Sample Performance - Deep Heuristic

Concluding Remarks

- Our model also outperforms the Quantal Cognitive Hierarchy model and a model with noisy best reply and pro-social preferences. Details

- By assuming optimal use of simple heuristics we:
  - Get accurate predictions of behavior in one-shot games
  - Capture the influence of the environment

- Key insights:
  - The environment shapes how people reason and what they pay attention to in interactions. A novel channel for how "culture" matters
  - As researchers, we should not look for THE heuristic or bias. Behavior adapts to the environment.
  - The simple heuristics people use can appear irrational in specific cases, but might work well with respect to an environement.

Thank you!

Comparison with Alternative Models

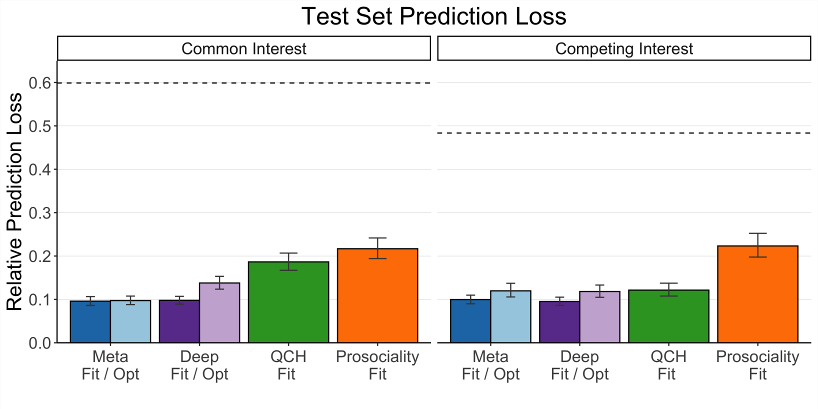

Go back

Metaheuristics

- **Row heuristic:** Evaluates each row via a weighted average. Goes from maximin to maximax via a single parameter γ.

- **Cell heuristic (common interest heuristic):** Assigns a value to each possible outcome (cell), choose an outcome and play the part of that outcome. We assume that the "best" outcome is the one that maximize the lowest of the two payoffs (the common interest).

- **Simulation heuristic:** Use a row or cell heuristic h^{belief} to form beliefs and best reply to those beliefs.

- All heuristics play a noisy best reply with sensitivity φ.

Metaheuristics

- **Cognitive costs:**
  - Row: $c(h^{row}) = C_{row} \cdot \varphi_{row}$
  - Cell: $c(h^{cell}) = C_{cell} \cdot \varphi_{cell}$
  - Simulation: $c(h^{sim}) = c(h^{belief}) + C_{mul} + C_{row} \cdot \varphi_{sim}$
- **Metaheuristic:** We borrow the functional form of rational inattention from Matějka and McKay (2015) which has three properties we want.
  - A primitive heuristic that is often used is more likely to be used in a given game.
  - A primitive heuristic that performs well in a game is more likely be used in for that game.
  - This is nicely captured by a prior probability and an adjustment cost λ.

$$\mathbb{P}\left[\{i \text{ use } h_i^n \text{ in } G\}\right] = \frac{p_n \exp\left[V_i(h_i^n, \mathcal{E}|G)/\lambda\right]}{\sum_{j=1}^{N} p_j \exp\left[(V_i(h_i^j, \mathcal{E}|G))/\lambda\right]}$$

Go back

References i

Hartford, Jason S., James R. Wright, and Kevin Leyton-Brown. 2016. "Deep Learning for Predicting Human Strategic Behavior." *Advances in Neural Information Processing Systems*, , (Nips): 2424–2432.

Matějka, Filip, and Alisdair McKay. 2015. "Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model." *American Economic Review*, 105(1): 272–298.