



Journal Club

**Capturing the Complexity of
Human Strategic Decision-Making
with Machine Learning**

Zheng Li

2024-10-01



Introduction

Paper

Zhu, J.-Q., Peterson, J. C., Enke, B., & Griffiths, T. L. (2024). *Capturing the Complexity of Human Strategic Decision-Making with Machine Learning* (No. arXiv:2408.07865). arXiv.
<https://doi.org/10.48550/arXiv.2408.07865>

1. Research Question
2. Experimental Design
3. Models
4. Results
5. Discussion



Research Question

(How) can machine learning be used to improve prediction/explanation of human strategic decision-making without the loss of interpretability



(This extra bit of explanatory power is attributed to complexity)



Can we establish complexity metrics that capture the patterns of behavioral responses?



Approach

- Conduct large scale experiment
- Develop and fit various models
- Make conjectures about differences in model performance: complexity
- Construct a bunch of complexity metrics and validate them
- Conduct another similar experiment to test out-of-sample performance

Experiment Design I: Task

		Col player (other)	
		L	R
Row player (self)	U	a x	b y
	D	c z	d w

Game Structure

- 2×2 matrix game
- Represented by an 8-item vector (a, b, c, d, x, y, z, w)
- 12 unique (orders of) payoffs for a player
 - Prisoner's Dilemma ($c > a > d > b$)
 - Stag Hunt/Trust ($a > c > d > b$)
 - Compromise, Battle of the sexes, etc.
 - *Topology of 2x2 Games* (Robinson & Goforth, 2005)
- $12 \times 12 = 144$ unique types of games

Game Generation

- Randomly generation of the 8-item vectors
 - Each item i is determined by drawing from a two-tiered uniform distribution
 - $i \sim U[1, u]$
 - $u \sim U[1, 50]$
- Excluded games with no pure-strategy NE
- Categorized games into the 144 types (only 142 after exclusion)
- Balanced the representations between double-dominance, single-dominance, and non-dominance games
 - Double-dominance: 3 instances per game type
 - Single-dominance: 8 instances per game type
 - Non-dominance: 22 instances per game type
- Ended up with 1208 instances

		L	R			L	R		
		U	17 20			30 7	U	5 14	6 19
		D	34 2	15 2			D	17 13	6 9




Experiment Design II: Procedure

Participant

- Prolific
- 4,900 participants -> 4,673 completed
 - Gender: 1942 M, 1782 F, 949 N/D
 - Age: 18 – 86, mean 37
- All from US
- At least 100 previous submissions
- minimum 95% acceptance rate

Incentive

- Base: \$2.00
- Bonus
 - One trial is randomly selected
 - Result is realized
 - Up to \$0.50 (1 pt = \$0.01)



- Basic interpretation of the payoff matrix
- No game theory (e.g. best response, NE) is taught or mentioned
- Procedure and payoff structure

- Check if participants correctly understand the payoff matrix only

- One-shot
- Different game each trial
- Matched with different opponent each trial
- Game type is randomly selected
- Normalized to row player
- Rows and columns are randomized (4 possible permutations)
- No feedback



Experiment Design III: Follow-up experiment


Col player (other)

L R

Row player (self) U D

U	a x	b y
D	c z	d w

How certain are you that choosing Row X is actually your best decision?



Game Generation

- Same process
- Ended up with 500 new games

Participant

- Prolific
- 1,013 participants -> 1,008 completed
 - Gender: 346 M, 416 F, 246 N/D
 - Age: 18 – 88, mean 35
- All from US
- At least 100 previous submissions
- minimum 95% acceptance rate

Procedure

- Same

Incentive

- Base: \$2.00
- Bonus
 - One trial is randomly selected
 - Result is realized
 - Up to \$0.50 (1 pt = \$0.01)

Models I: Context-invariant

Model

Key factors

- Limited strategic sophistication (level-k reasoning)
- Noisiness in decision-making
 - Of one's own
 - Belief of opponent's
- Risk aversion

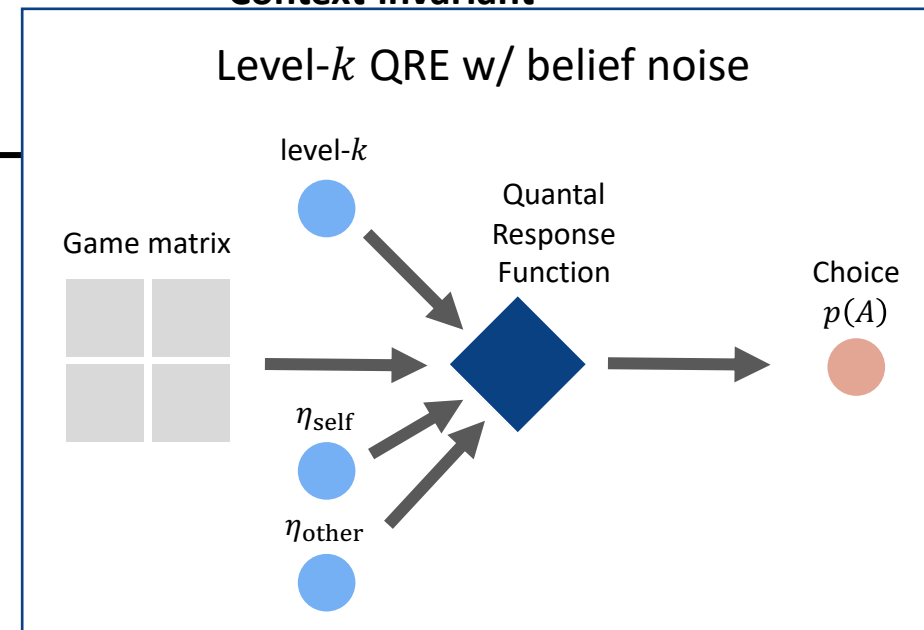
Mathematical form

- Noisy logit choice with expected CARA utility based on belief
- $p(A) = 1 / (1 + e^{-\eta_{\text{self}}(EU(A) - EU(B))})$
- $EU(A) = p^s(C|k, \eta_{\text{other}}^s)U(x_{A,C}) + p^s(D|k, \eta_{\text{other}}^s)U(x_{A,D})$
- $U(x) = (1 - e^{-\alpha x}) / \alpha$
- Parameters: $k, \eta_{\text{self}}, \eta_{\text{other}}^s$
 - **Context-invariant**

Variations

Description

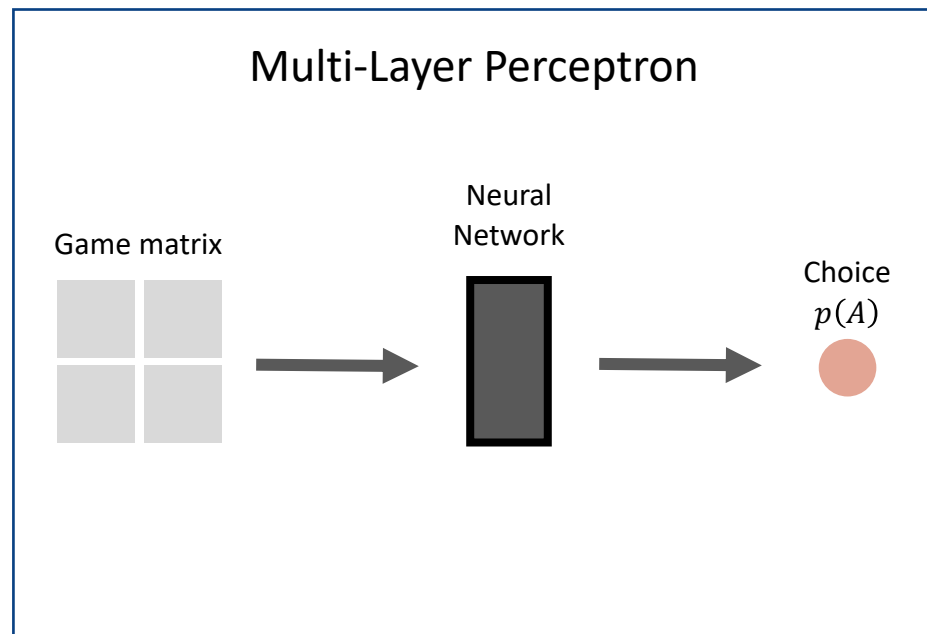
Nash Equilibrium (PSNE, MSNE)	PSNE: play NE choice with 100% MSNE: play choices with MS probability
Level-k	Level-0: play randomly Level-k: play BR assuming opponent is at $k - 1$
Level-k QRE	Play QR according to $p(A; k, \eta_{\text{self}})$ assuming opponent is playing $p(A; k - 1, \eta_{\text{self}})$
Level-k QRE w/ belief noise	Play QR according to $p(A; k, \eta_{\text{self}})$ assuming opponent is playing $p(A; k - 1, \eta_{\text{other}})$



Models II: Context-dependent

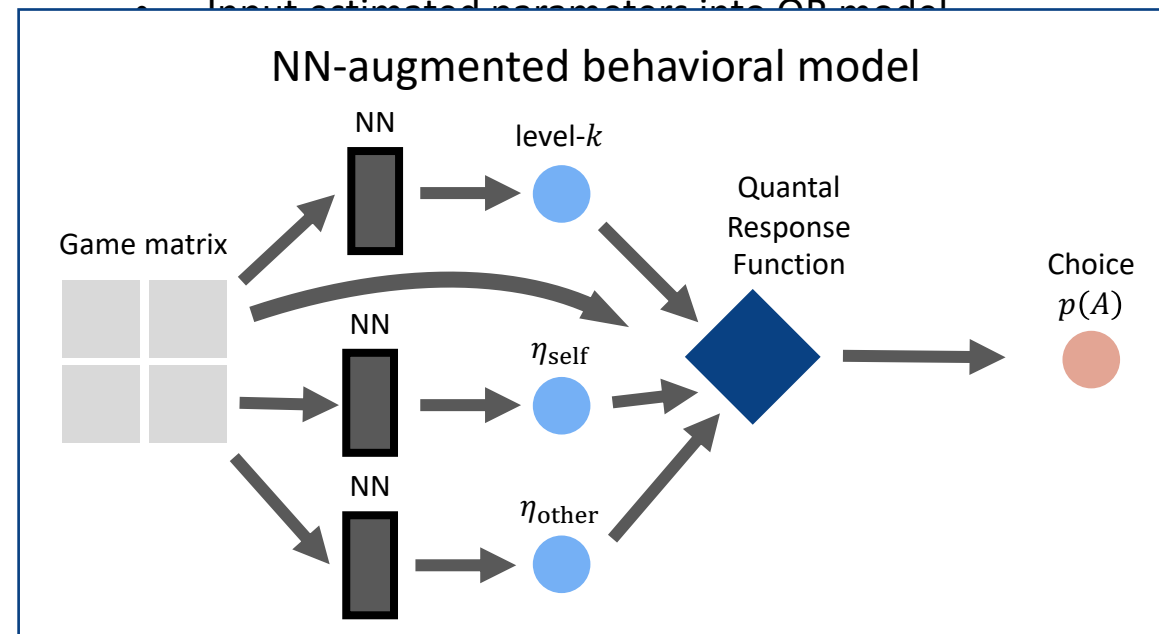
Multi-Layer Perceptron

- 3 hidden layers each with 300 neurons and sigmoid activation
- Softmax activation in output layer
- $p(A) = f_{MLP}(\text{game matrix})$




Neural Network Augmented Behavioral Model

- Connection b/w context-invariant and context-dependent
 - Estimate 3 parameters using separate MLP
 - $k = f_{MLP_0}(\text{game matrix})$
 - $\eta_{\text{self}} = f_{MLP_1}(\text{game matrix})$
 - $\eta_{\text{other}} = f_{MLP_2}(\text{game matrix})$



Models III: Complexity


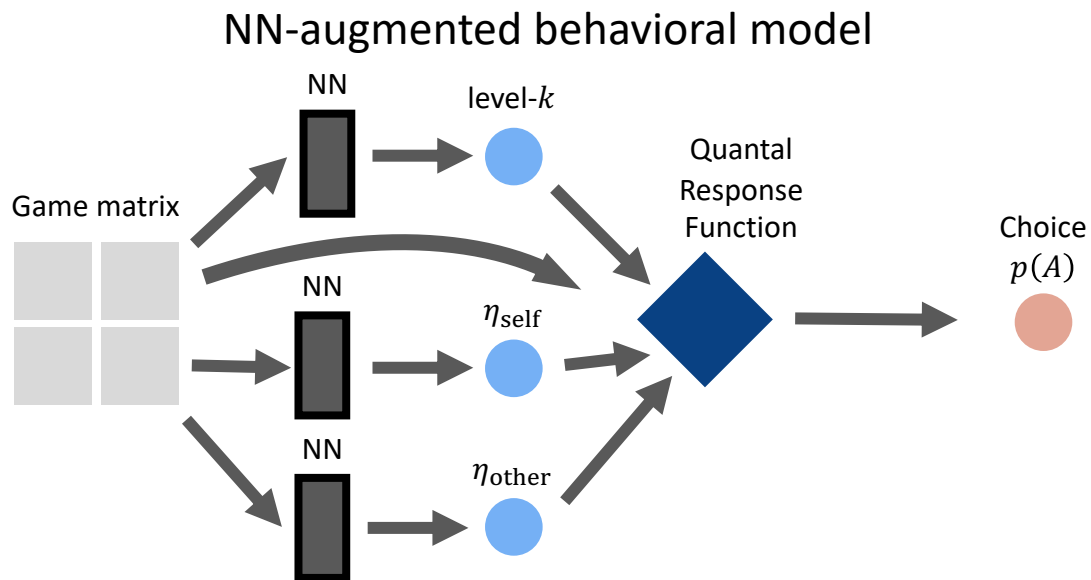

Intermediate result



- Context matters – but what exactly properties of the context is making a difference?
- Note that player's own noise and belief of opponent's noisiness have significant impact in prediction
- **Conjecture:** *noises are largely due to the complexity of games*
- **Question:** *can we establish quantifiable measure of complexity?*

Models III: Complexity

Intermediate result





Models III: Complexity

Col player (other)

L

R

Row player (self)

U	a x	b y
D	c z	d w

Complexity index


- **Dominant Solvability:** whether the game is solvable by iterated deletion of dominated action
- **Excess Dissimilarity:** $\text{Dissimilarity}_{\text{self}} = \frac{|a-c|}{2} + \frac{|b-d|}{2} - |\mu_U - \mu_D|$, $\text{Dissimilarity}_{\text{other}} = \frac{|x-y|}{2} + \frac{|z-w|}{2} - |\mu_L - \mu_R|$
- **Levels of Iterative Rationality:** minimum level k at which the best-response dynamic converges to a fixed strategy (up to 3)
- **Number of Nash Equilibria:** ...
- **Nash Equilibrium Payoff Dominance:** whether at least one PSNE in a game yields the maximum payoff
- **Nash Equilibrium Pareto Dominance:** whether a PSNE offers superior payouts compared to all other PSNE
- **Pure Motives:** whether player's choice preferences are expected to exhibit perfect rank correlation
- **Max Payouts:** $\text{Max}_{\text{self}} = \max\{a, b, c, d\}$, $\text{Max}_{\text{other}} = \max\{x, y, z, w\}$
- **Payoff Variances:** $\text{Var}_{\text{self}} = \frac{1}{4}((a - \mu_U)^2 + (b - \mu_U)^2 + (c - \mu_D)^2 + (d - \mu_D)^2)$
- **Deviations from Zero-Sum Games:** $\text{NonZeroSum} = |a - c + x - z| + |a - b + x - y| + |c - d - x + z| + |c - d - y + w|$
- **Inequality in Payouts:** $\text{Inequality} = \max\{a, b, c, d\} - \max\{x, y, z, w\}$
- **Asymmetry in Payouts:** $\text{Asymmetry} = \frac{1}{4}(|a - x| + |b - y| + |c - z| + |d - w|)$

The results of LASSO regressions.

Game feature	$-\eta_{\text{self}}$	$-\eta_{\text{other}}^s$
DominantSolvable _{self}	.	.
DominantSolvable _{other}	.	.
Dissimilarity _{self}	0.28	.
Dissimilarity _{other}	.	0.08
LevelIterRational	0.38	.
NumPSNE	.	.
NumMSNE	.	.
PayoffDomEquilibrium	-0.80	.
PayoffDomNonEquilibrium	.	.
ParetoDomEquilibrium	.	.
PureMotives	.	.
Max _{self}	0.30	.
Max _{other}	.	0.01
PayoffVar _{self}	0.40	.
PayoffVar _{other}	.	.
NonZeroSum	.	.
Inequality	0.85	.
Asymmetry	-0.09	.
Intercept	-9.28	-1.09
R ²	0.33	0.02

Further selection

- *Level-2 Neural QR and Neural Belief Noise*
- **Decision tree**
 - Max depth = 3
- **LASSO regressions**
 - normalized the game features
 - L1 term at 0.2



Results I: Predictions

Comparisons Between Content-independent and Neurally-Augmented Content-Independent

Model	MSE (SE) ↓	R^2 (SE) ↑
Nash	.1680	.2561
L1+QR+Risk	.0203	.758
L2+QR+Belief+Risk	.016	.8113
L3+QR+Belief+Risk	.0167	.8026
QRE+Belief+Risk	.0166	.8039
L1+QR+Risk	.0128 (.0001)	.8547 (.0032)
L2+QR+Belief+Risk	.0109 (.0002)	.875 (.0027)
L2+QR+Belief+Risk	.0102 (.0001)	.8868 (.0018)
L3+QR+Belief+Risk	.0108 (.0002)	.8798 (.0022)
L3+QR+Belief+Risk	.0104 (.0001)	.8855 (.0018)
QRE+Belief+Risk	.0111 (.0002)	.877 (.0029)
QRE+Belief+Risk	.0105 (.0001)	.8835 (.0018)
L+QR+Belief+Risk	.0081 (.0002)	.9065 (.0023)
MLP	.0073 (.0001)	.9194 (.0022)


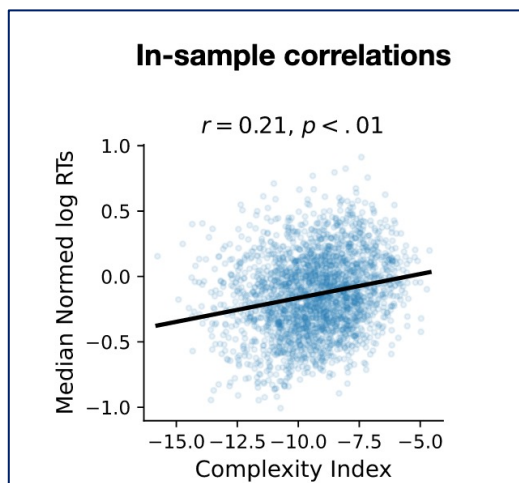

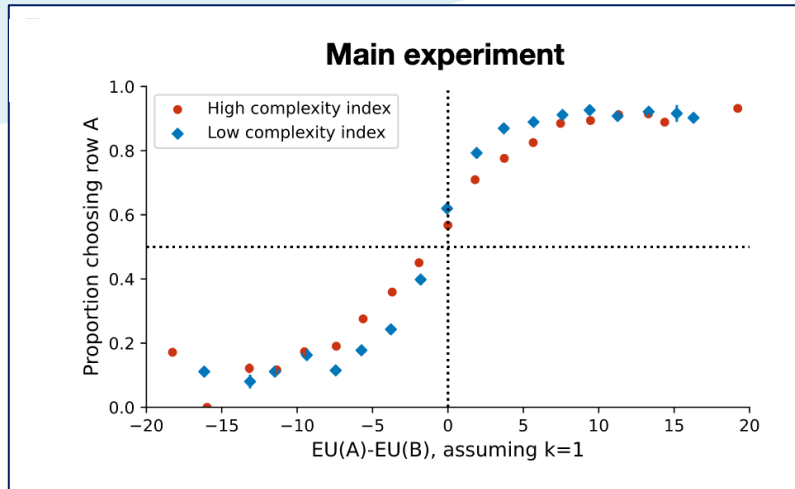
Note. Colorized texts denote components of the model that are implemented as neural networks. Specifically, **QR** indicates that η_{self} is a neural network, **Belief** signifies that η_{other}^s is a neural network, and **L** means that a neural network predicts a mixture of level- k players. The numbers in parentheses represent the standard errors from the 10-fold cross-validation.

QRE+Belief+Risk	.0116 (.0004)	.8732 (.0036)	95
L+QR+Risk	.0194 (.0003)	.7834 (.0052)	85
L+QR+Belief+Risk	.0096 (.0003)	.8931 (.0039)	97
MLP	.0073 (.0001)	.9194 (.0022)	100 (upper bound)

“Completeness”

- Scaled performance metric
- 0%: performance of a 50-50 random choice model
- 100%: performance of the MLP

Results II: Effect of Complexity Metrics





Discussion

1. Complexity is definitely getting more and more attention even in economics
 - Attribute behavioral patterns/irregularities to complexity
 - “on the surface”, not into the neural/cognitive mechanism though
 - Derive measures of complexity from features of the problem/task/game/choice itself
 - “social science approach”, not computational complexity
2. Use of machine learning algorithms in studying human behaviors
 - Algorithms explain/predict better than traditional theory -> incorporate them together
 - Interpretability
3. Extensions...



Journal Club

**Capturing the Complexity of
Human Strategic Decision-Making
with Machine Learning**

Zheng Li

2024-10-01