Using Response Times to Improve Cognitive Models

Blair RK Shevlin, PhD
Center for Computational Psychiatry
Icahn School of Medicine at Mount Sinai

October 24, 2025

Image generated with OpenAI's DALL·E via ChatGPT (October 2025)

response times are **information**

- Decision-makers can infer preferences from response times[1,2]
- Decision-makers use opponents' response times to plan and infer private information[3-5]
- eBay buyers and sellers use response times to negotiate prices[6]

PNAS RESEARCH ARTICLE | PSYCHOLOGICAL AND COGNITIVE SCIENCES OPEN ACCESS Check for updates

Deliberation during online bargaining reveals strategic information

Miruna Cotet[a], Wenjia Joyce Zhao[b], and Ian Krajbich[a,c,d,1]

Edited by Jennifer S. Trueblood, Indiana University, Bloomington, IN; received May 31, 2024; accepted January 7, 2025 by Editorial Board Member Timothy D. Wilson

1. Gates, Callaway, Ho, Griffiths (2021) *Cognition*
2. Bavard, Stuchly, Konovalov, Gluth (2024) *PLoS Biology*
3. Frydman, Krajbich (2022) *Management Science*
4. Konovalov, Krajbich (2023) *The Economic Journal*
5. Konovalov, Krajbich (2023) *Judgment & Decision Making*
6. Cotet, Zhao, Krajbich (2025) *PNAS*

Outline

- Introduction to the Diffusion Decision Model (DDM)
- Literature Incorporating RTs into Model Estimation
- Modeling Social Exchange with the Ultimatum Game
- Clinical Translation in Borderline Personality Disorder

how to understand the decision-making process?

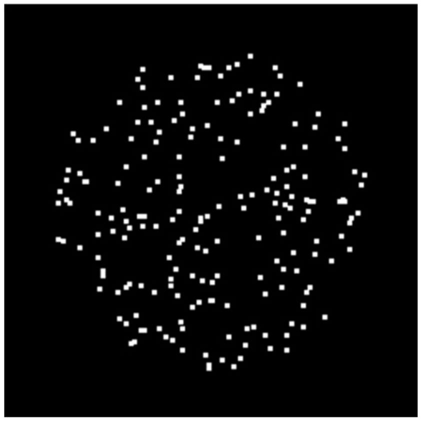

which way are the dots moving?

how to understand the decision-making process?

should I accept this offer?

what is a cognitive model

- Generative: Approximate the algorithms the brain uses to ***produce*** behavior

- Measurement: ***Quantify*** unobservable cognitive processes

- Selective influence: parameter describe ***unique*** latent constructs with ***distinct*** behavioral outputs

Example

learning rate → integration of new information → temporal evolution of behavior

the diffusion decision model (DDM)

- process terminates at fixed **thresholds**
- process **drifts** towards positive or negative values
- process can be **biased** *a priori*
- constant **non-decision time** is added to total time to threshold
- parameters can **vary across trials**

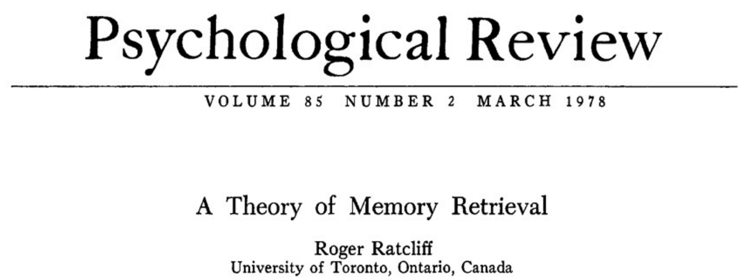
Psychological Review
VOLUME 85 NUMBER 2 MARCH 1978

A Theory of Memory Retrieval

Roger Ratcliff
University of Toronto, Ontario, Canada

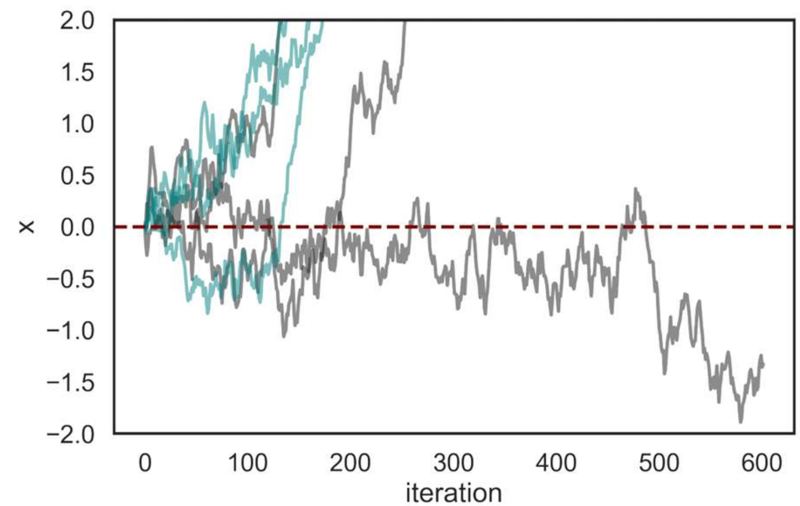

decision thresholds

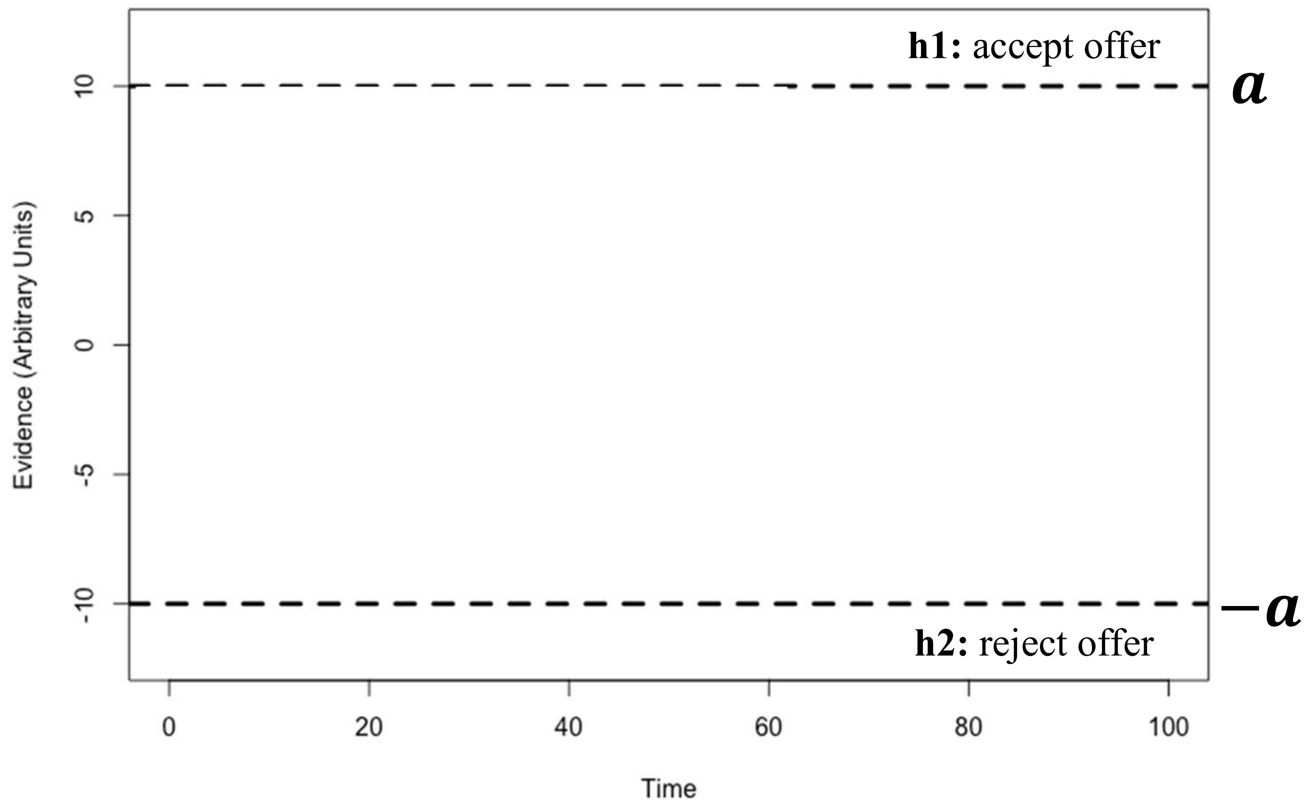

starting point

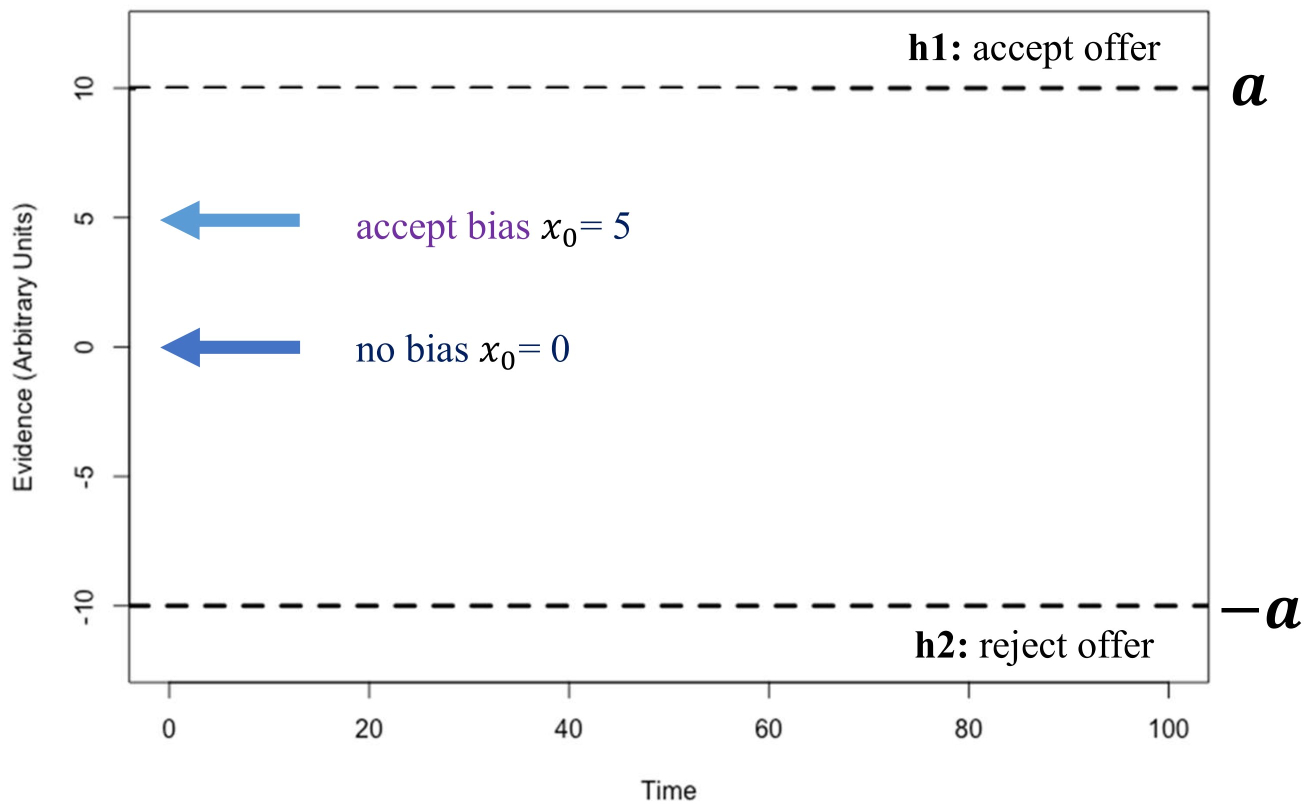

drift rate

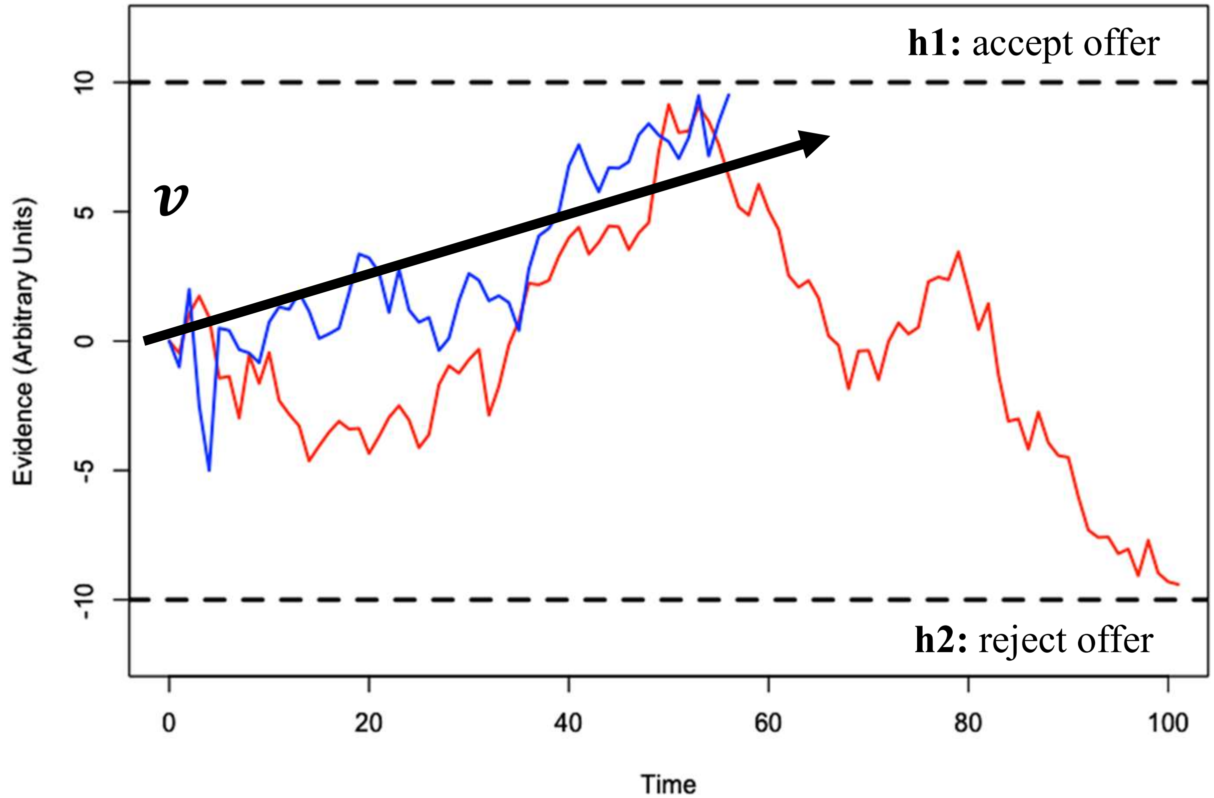

$$\frac{d}{d_i}X_i = Normal(v, s^2)$$

speed-accuracy trade-off

go faster!

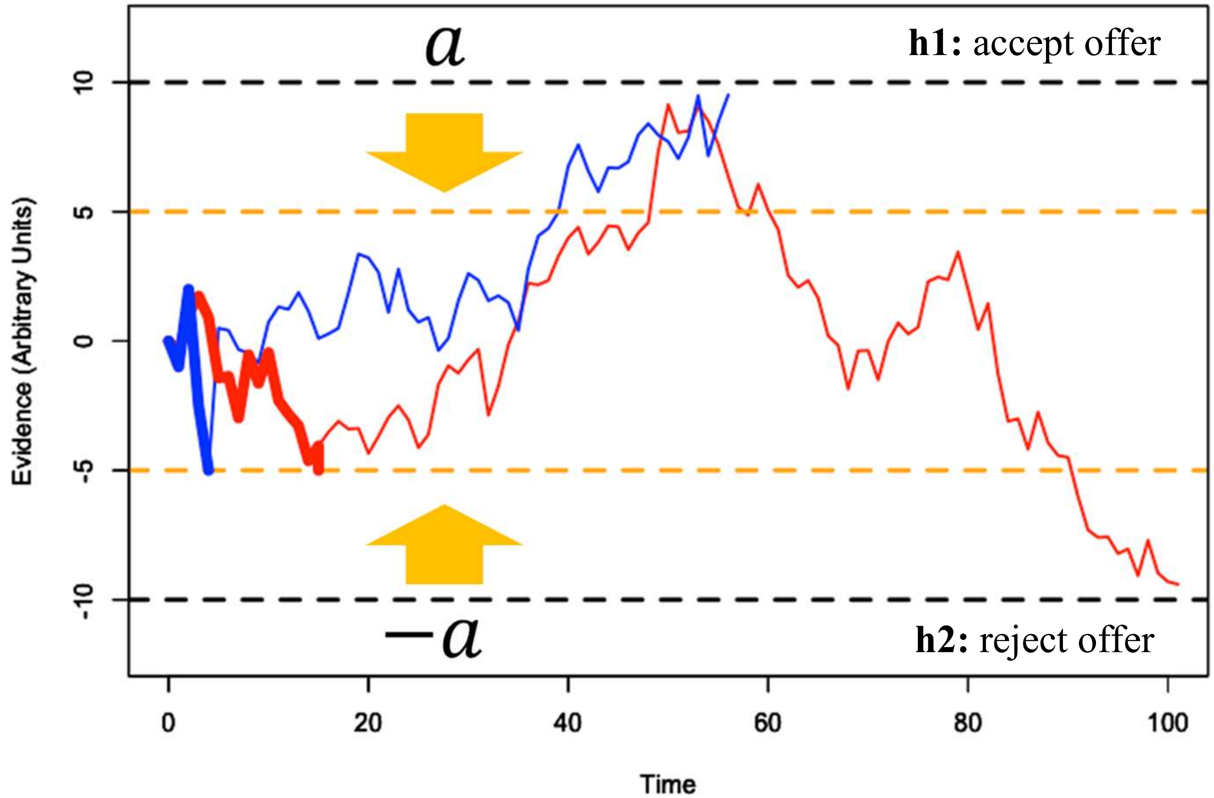

non-decision time

Decision time (DT)

Response time = Sensory delay +

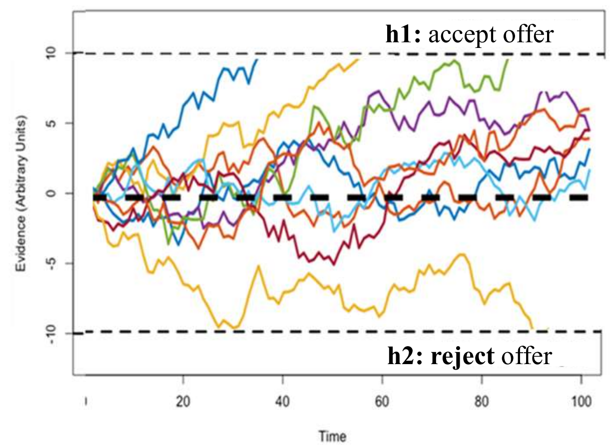

+ Motor delay

Non-decision time (T_0)

Psychological Research (2019) 83:1194–1209
https://doi.org/10.1007/s00426-017-0945-8

ORIGINAL ARTICLE

CrossMark

Experimental validation of the diffusion model based on a slow response time paradigm

Veronika Lerche[1] · Andreas Voss[1]

Received: 9 July 2017 / Accepted: 9 November 2017 / Published online: 9 December 2017

parameters

- threshold a
- starting point bias x_0
- drift rate (signal-to-noise ratio) v
- non-decision time t_0
- noise s
  - In practice, often set to 1 to .1
- variability parameters sv, sz, st
  - Across-trial fluctuations stimuli and physiological states

ELSEVIER

Journal of Mathematical Psychology

Volume 105, December 2021, 102594

Attention as a source of variability in decision-making: Accounting for overall-value effects with diffusion models

Blair R.K. Shevlin [a], Ian Krajbich [a b]

parameters

- threshold a
- starting point bias x_0
- drift rate (signal-to-noise ratio) v
- non-decision time t_0
- noise s
  - In practice, often set to 1 to .1

ELSEVIER

Journal of Mathematical Psychology

Volume 105, December 2021, 102594

Attention as a source of variability in decision-making: Accounting for overall-value effects with diffusion models

Blair R.K. Shevlin [a], Ian Krajbich [a b]

when can we use the DDM?

- only **two options**
- task involves **relative evidence**
- there is **perfect inhibition**
- there is no **leakage** of information
- process is **continuous** in time
- process is **single-stage**

otherwise: Leaky Competing Accumulator Model, Gaze-Weighted Accumulator Model, Linear Ballistic Accumulator Model, Piecewise Diffusion Model, Racing Diffusion Model, Circular Diffusion Model, etc.

Response Times in Model Estimation

ELSEVIER
Journal of Neuroscience Methods
Volume 317, 1 April 2019, Pages 37-44

Joint modeling of reaction times and choice improves parameter identifiability in reinforcement learning models

Ian C. Ballard [a b c], Samuel M. McClure [c]

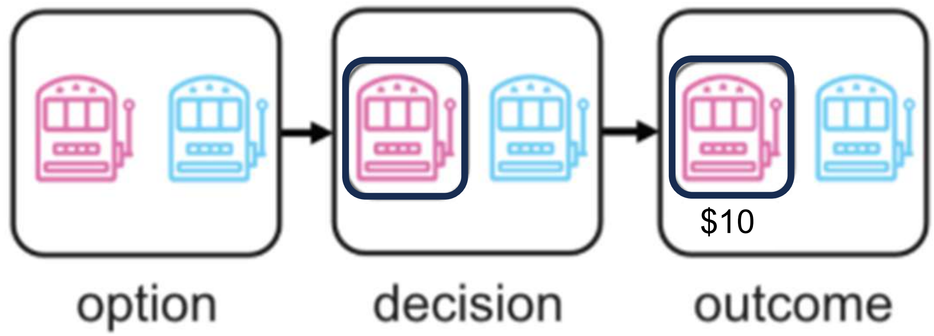

Value Function

$$Q(\text{Option}_{chosen})^{t+1} = Q(\text{Option}_{chosen})^{t} + \alpha\,(\text{Reward}^{t} - Q(\text{Option}_{chosen})^{t})$$

Softmax Decision Rule

$$p(\text{Choice}^{t} = \text{Option}_i) = \frac{e^{\beta Q(\text{Option}_i\ )^{t}}}{\sum_j e^{\beta Q(\text{Option}_j\ )^{t}}}$$

Parameters

Learning rate: α

Inverse temperature: β

Journal of Neuroscience Methods
Volume 317, 1 April 2019, Pages 37-44

Joint modeling of reaction times and choice improves parameter identifiability in reinforcement learning models

Ian C. Ballard [a b c], Samuel M. McClure [c]

Value Function

$$Q(\text{Option}_{\text{ch}})^{t+1} = Q(\text{Option}_{\text{ch}})^{t} + \alpha\,(\text{Reward}^{t} - Q(\text{Option}_{\text{ch}})^{t})$$

Softmax Decision Rule

$$p(\text{Choice}^{t} = \text{Option}_{i}) = \frac{e^{\beta Q(\text{Option}_{i})^{t}}}{\sum_{j} e^{\beta Q(\text{Option}_{j})^{t}}}$$

Parameters

Learning rate: α

Inverse temperature: β

Incorporating Response Times

$$\text{RT}^{t} = -\beta\,|Q(\text{Option}_{1})^{t} - Q(\text{Option}_{2})| + \varepsilon$$

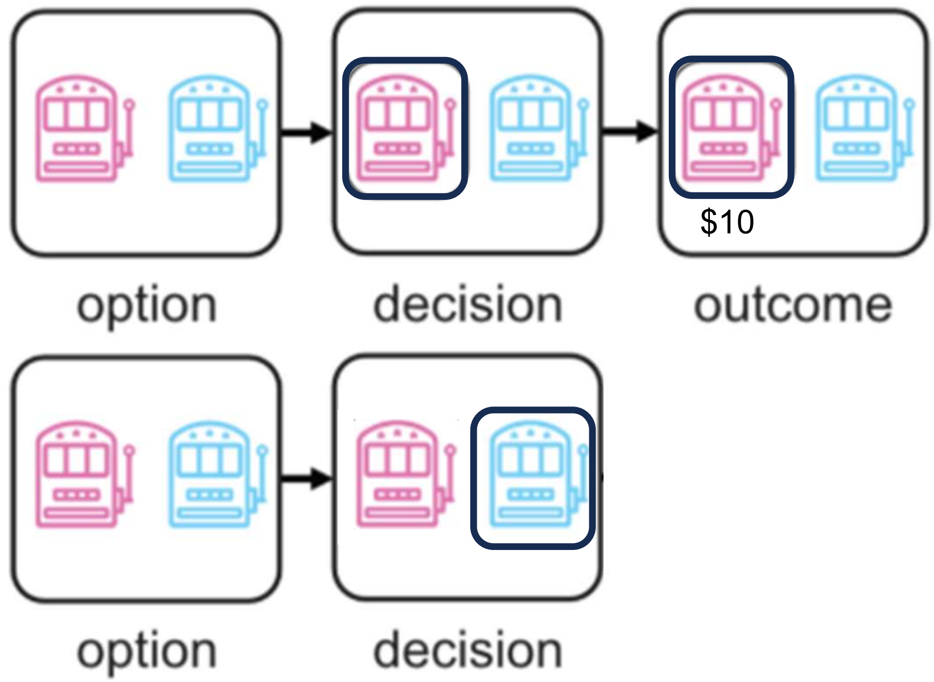

a) Did the agent fail to learn?

b) Did the agent respond noisily?

Journal of Neuroscience Methods
Volume 317, 1 April 2019, Pages 37-44

Joint modeling of reaction times and choice improves parameter identifiability in reinforcement learning models

Ian C. Ballard [a b c], Samuel M. McClure [c]

<u>Simulation Study</u>

1. Generate parameters (X_{gen}) from distributions based on prior literature;
2. Simulate choice and RT data for 10-200 subjects on 100-500 trials;
3. Extract fitted parameters (X_{fit}) from simulated data using multiple techniques (ML, MAP);
4. Assess correlations between X_{gen} and X_{fit}

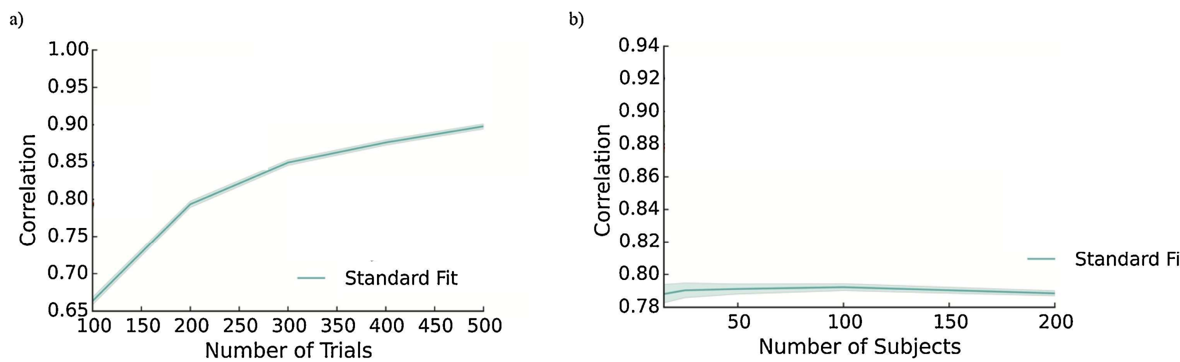

Journal of Neuroscience Methods
Volume 317, 1 April 2019, Pages 37-44

Joint modeling of reaction times and choice improves parameter identifiability in reinforcement learning models

Ian C. Ballard [a b c], Samuel M. McClure [c]

Simulation Study

1. Generate parameters (X_{gen}) from distributions based on prior literature;
2. Simulate choice and RT data for 10-200 subjects on 100-500 trials;
3. Extract fitted parameters (X_{fit}) from simulated data using multiple techniques (ML, MAP);
4. Assess correlations between X_{gen} and X_{fit}

a)

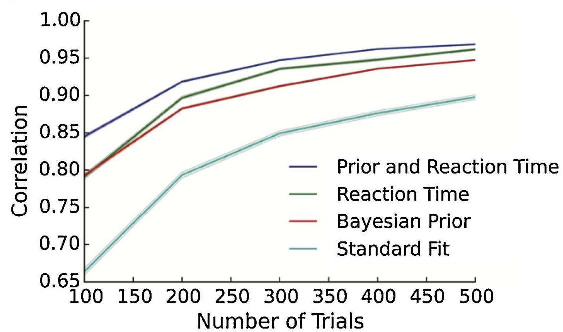

b)

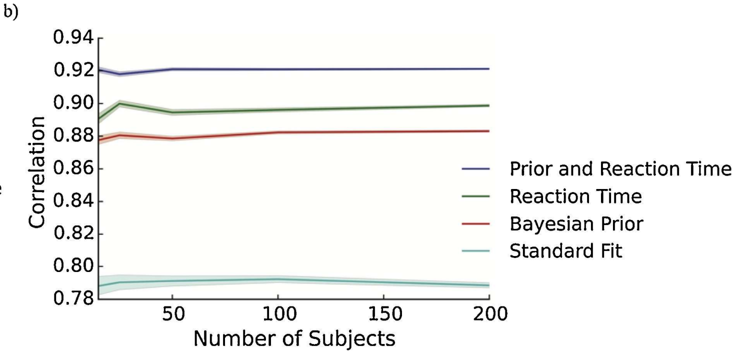

ELSEVIER

Journal of Neuroscience Methods

Volume 317, 1 April 2019, Pages 37-44

Joint modeling of reaction times and choice improves parameter identifiability in reinforcement learning models

Ian C. Ballard [a b c], Samuel M. McClure [c]

Simulation Study

1. Generate parameters (X_{gen}) from distributions based on prior literature;
2. Simulate choice and RT data for 10-200 subjects on 100-500 trials;
3. Extract fitted parameters (X_{fit}) from simulated data using multiple techniques (ML, MAP);
4. Assess correlations between X_{gen} and X_{fit}

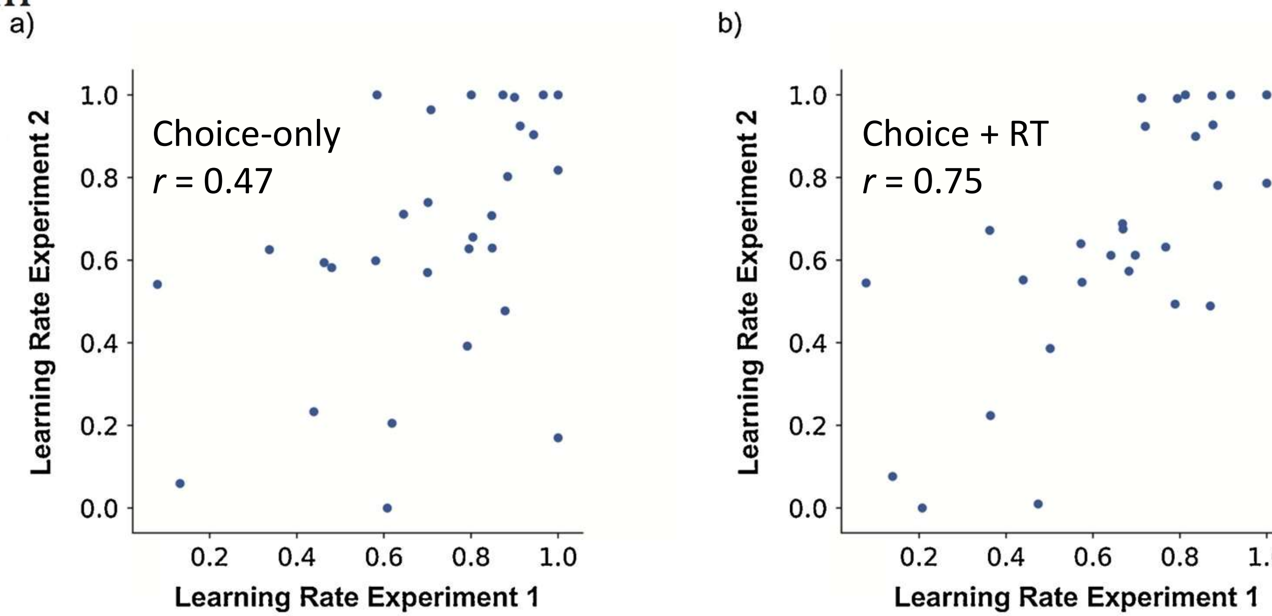

Fig. 4. Joint modeling of choice and reaction times improves parameter identifiability in real data. **A)** The correlation in estimated learning rate between two runs of a bandit task using a standard RL model. **B)** The correlation in estimated learning rate when estimated using a model of reaction times and choice.

PLOS Computational Biology

OPEN ACCESS PEER-REVIEWED

RESEARCH ARTICLE

Improving the reliability of model-based decision-making estimates in the two-stage decision task with reaction-times and drift-diffusion modeling

Nitzan Shahar, Tobias U. Hauser, Michael Moutoussis, Rani Moran, Mehdi Keramati, NSPN consortium, Raymond J. Dolan

Version 2 Published: February 13, 2019 • https://doi.org/10.1371/journal.pcbi.1006803

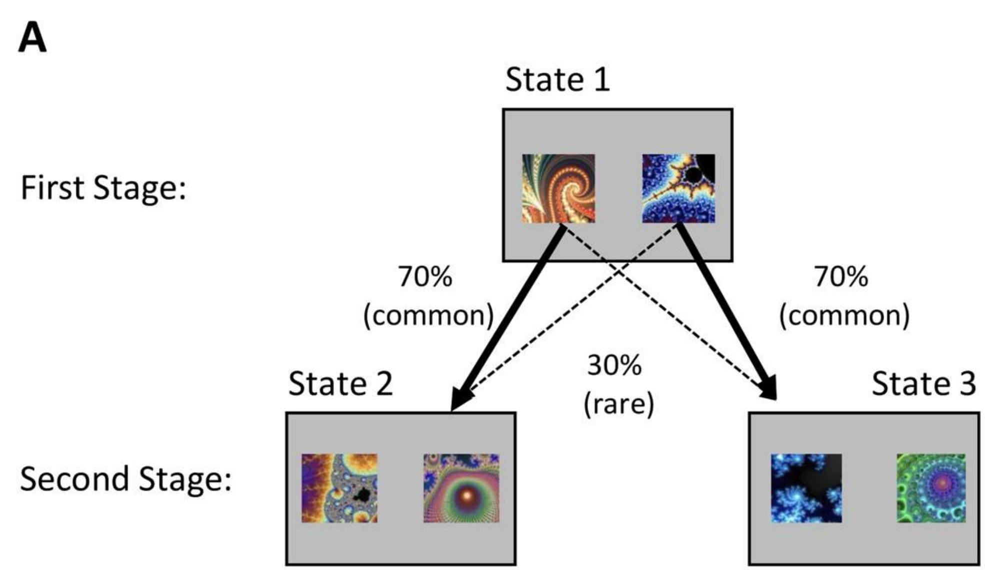

Value Functions

$$Q^{MF}(\text{Action}_2)^{t+1} = Q^{MF}(\text{Action}_2)^{t} + \alpha\,(\text{Reward}^t - Q^{MF}(\text{Action}_2)^t)$$

$$Q^{MF}(\text{Action}_1)^{t+1} = Q^{MF}(\text{Action}_1)^{t} + \alpha\left(Q^{MF}(\text{Action}_2)^t - Q^{MF}(\text{Action}_1)^t\right) + \lambda\alpha\left(\text{Reward}^t - Q^{MF}(\text{Action}_2)^t\right)$$

$$Q^{MB}(\text{Action}_1)^{t} = P(\text{State}_2 |\ \text{Action}_1) \times \max\left(Q^{MF}(\text{State}_2)^t\right) + P(\text{State}_3 |\ \text{Action}_1) \times \max\left(Q^{MF}(\text{State}_3)^t\right)$$

$$Q^{Net}(\text{Action}_1)^{t} = (1-\omega)\,Q^{MF}(\text{Action}_1)^{t} + \omega Q^{MB}(\text{Action}_2)^{t} + \gamma \text{Stay}^t$$

Softmax Decision Rule

$$p\left(\text{Choice}^t = \text{Action}_i\right) = \frac{e^{\beta Q^{Net}(\text{Action}_i)^t}}{\sum_j e^{\beta Q^{Net}(\text{Action}_j)^t}}$$

PLOS Computational Biology

OPEN ACCESS PEER-REVIEWED

RESEARCH ARTICLE

Improving the reliability of model-based decision-making estimates in the two-stage decision task with reaction-times and drift-diffusion modeling

Nitzan Shahar, Tobias U. Hauser, Michael Moutoussis, Rani Moran, Mehdi Keramati, NSPN consortium, Raymond J. Dolan

Version 2 Published: February 13, 2019 • https://doi.org/10.1371/journal.pcbi.1006803

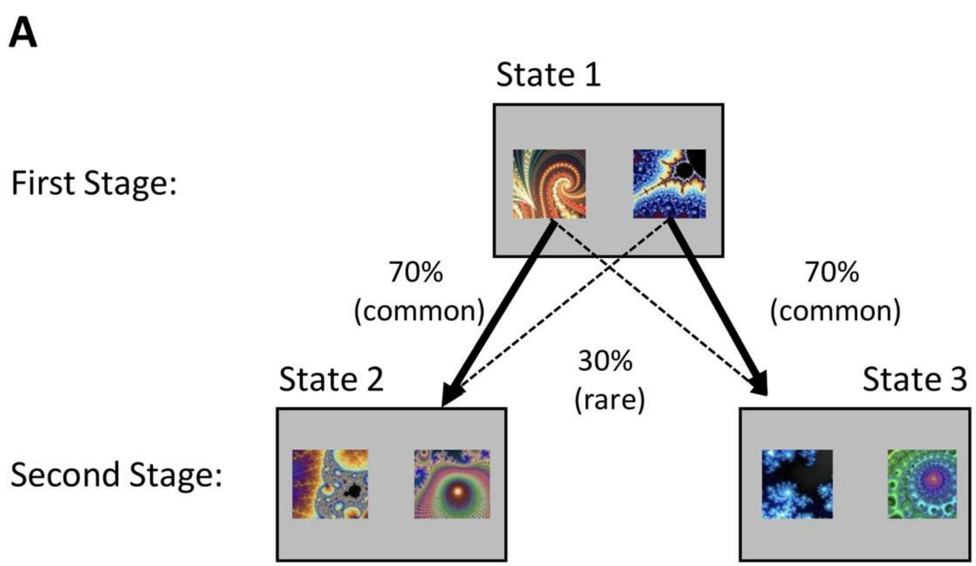

Value Functions

$$Q^{MF}(\text{Action}_2)^{t+1} = Q^{MF}(\text{Action}_2)^{t} + \alpha\,(\text{Reward}^t - Q^{MF}(\text{Action}_2)^t)$$

$$Q^{MF}(\text{Action}_1)^{t+1} = Q^{MF}(\text{Action}_1)^{t} + \alpha\left(Q^{MF}(\text{Action}_2)^t - Q^{MF}(\text{Action}_1)^t\right) + \lambda\alpha\left(\text{Reward}^t - Q^{MF}(\text{Action}_2)^t\right)$$

$$Q^{MB}(\text{Action}_1)^{t} = P(\text{State}_2 \mid \text{Action}_1) \times \max\left(Q^{MF}(\text{State}_2)^t\right) + P(\text{State}_3 \mid \text{Action}_1) \times \max\left(Q^{MF}(\text{State}_3)^t\right)$$

$$Q^{Net}(\text{Action}_1)^{t} = (\mathbf{1} - \boldsymbol{\omega})\, Q^{MF}(\text{Action}_1)^t + \boldsymbol{\omega} Q^{MB}(\text{Action}_2)^t + \gamma \text{Stay}^t$$

Softmax Decision Rule

$$p\left(\text{Choice}^t = \text{Action}_i\right) = \frac{e^{\beta Q^{Net}(\text{Action}_i)^t}}{\sum_j e^{\beta Q^{Net}(\text{Action}_j)^t}}$$

PLOS Computational Biology

OPEN ACCESS PEER-REVIEWED

RESEARCH ARTICLE

Improving the reliability of model-based decision-making estimates in the two-stage decision task with reaction-times and drift-diffusion modeling

Nitzan Shahar, Tobias U. Hauser, Michael Moutoussis, Rani Moran, Mehdi Keramati, NSPN consortium, Raymond J. Dolan

Version 2 Published: February 13, 2019 • https://doi.org/10.1371/journal.pcbi.1006803

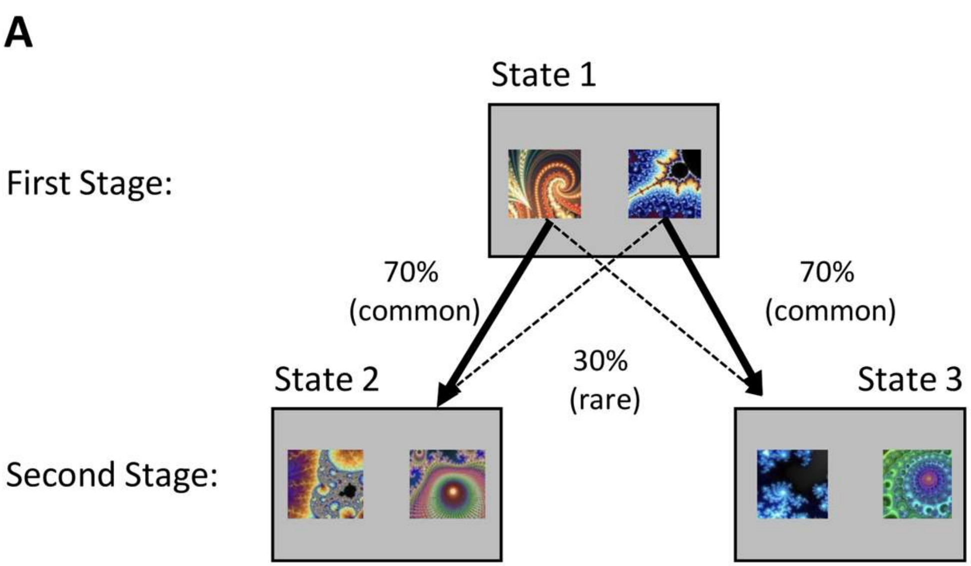

Value Functions

$$Q^{MF}(\text{Action}_2)^{t+1} = Q^{MF}(\text{Action}_2)^{t} + \alpha\left(\text{Reward}^t - Q^{MF}(\text{Action}_2)^t\right)$$

$$Q^{MF}(\text{Action}_1)^{t+1} = Q^{MF}(\text{Action}_1)^{t} + \alpha\left(Q^{MF}(\text{Action}_2)^t - Q^{MF}(\text{Action}_1)^t\right) + \lambda\alpha\left(\text{Reward}^t - Q^{MF}(\text{Action}_2)^t\right)$$

$$Q^{MB}(\text{Action}_1)^{t} = P(\text{State}_2 |\ \text{Action}_1) \times \max\left(Q^{MF}(\text{State}_2)^t\right) + P(\text{State}_3 |\ \text{Action}_1) \times \max\left(Q^{MF}(\text{State}_3)^t\right)$$

$$Q^{Net}(\text{Action}_1)^{t} = (\mathbf{1}-\boldsymbol{\omega})\, Q^{MF}(\text{Action}_1)^t + \boldsymbol{\omega} Q^{MB}(\text{Action}_2)^t + \gamma \text{Stay}^t$$

DDM Decision Rule

$$\frac{d}{di}X_i \sim N\left(v, s^2\right)$$

$$v = d\left(Q^{Net}(\text{Action}_{chosen})^t - Q^{Net}(\text{Action}_{unchosen})^t\right)$$

PLOS Computational Biology

OPEN ACCESS PEER-REVIEWED

RESEARCH ARTICLE

Improving the reliability of model-based decision-making estimates in the two-stage decision task with reaction-times and drift-diffusion modeling

Nitzan Shahar, Tobias U. Hauser, Michael Moutoussis, Rani Moran, Mehdi Keramati, NSPN consortium, Raymond J. Dolan

Version 2 Published: February 13, 2019 • https://doi.org/10.1371/journal.pcbi.1006803

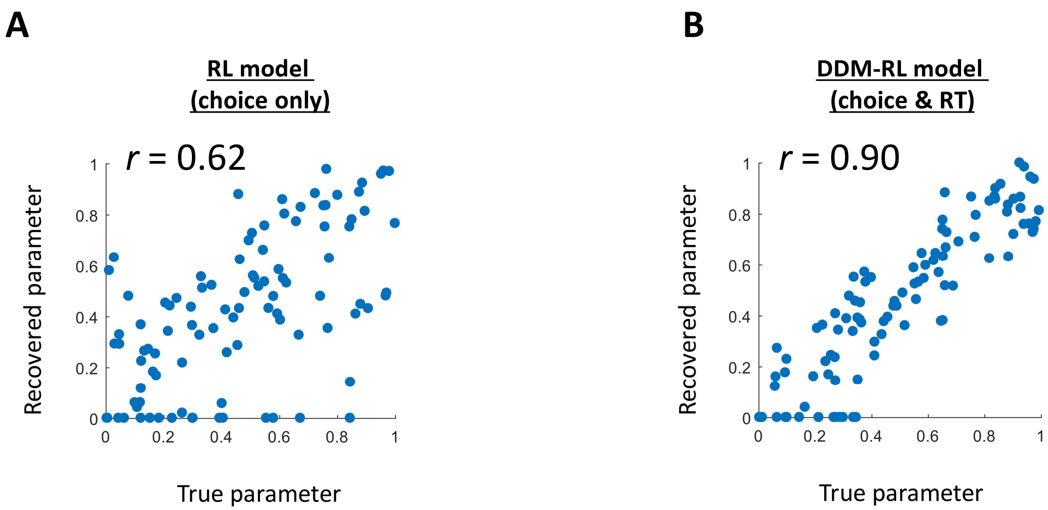

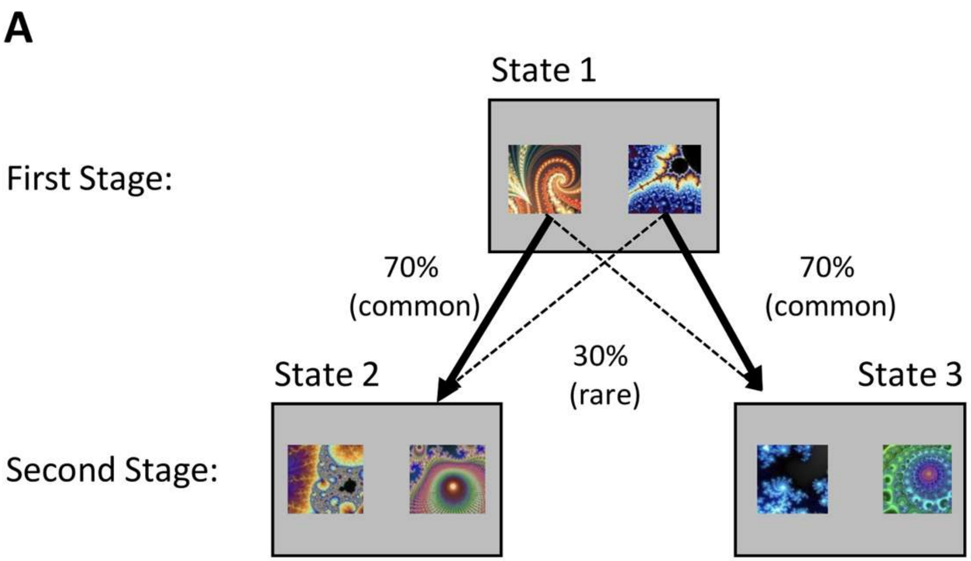

		Trials in the analysis			
		200	500	1000	5000
RL model (choice)	α_1	.54	.62	.68	.92
	α_2	.95	.98	.99	.99
	λ	.53	.71	.71	.88
	w	**.61**	**.69**	**.82**	**.97**
	p	.82	.90	.91	.97
	β_1	.82	.90	.93	.98
	β_2	.89	.96	.98	.99
DDM-RL model (choice & RT)	α_1	.68	.72	.84	.94
	α_2	.99	.99	.99	.99
	λ	.58	.75	.83	.92
	w	**.90**	**.95**	**.96**	**.99**
	p	.91	.94	.97	.99
	b_1	.93	.93	.99	.99
	a_1	.93	.98	.99	.99
	τ_1	.99	.99	.99	.99
	b_2	.99	.99	.99	.99
	a_2	.97	.99	.99	.99
	τ_2	.99	.99	.99	.99

Other relevant literature

Behavior Research Methods (2020) 52:2142–2155
https://doi.org/10.3758/s13428-020-01372-w

Quantifying the benefits of using decision models with response time and accuracy data

Tom Stafford[1] · Angelo Pirrone[2] · Mike Croucher[3] · Anna Krystalli[4]

Received: 26 August 2019 / Revised: 25 January 2020 / Accepted: 28 January 2020 / Published online: 30 March 2020

American Psychological Association
Psychology and Aging

0882-7974/18/$12.00
2018, Vol. 33, No. 7, 1093–1104
http://dx.doi.org/10.1037/pag0000298

The Mapping Between Transformed Reaction Time Costs and Models of Processing in Aging and Cognition

Craig Hedge, Georgina Powell, and Petroc Sumner
Cardiff University

The Ultimatum Game

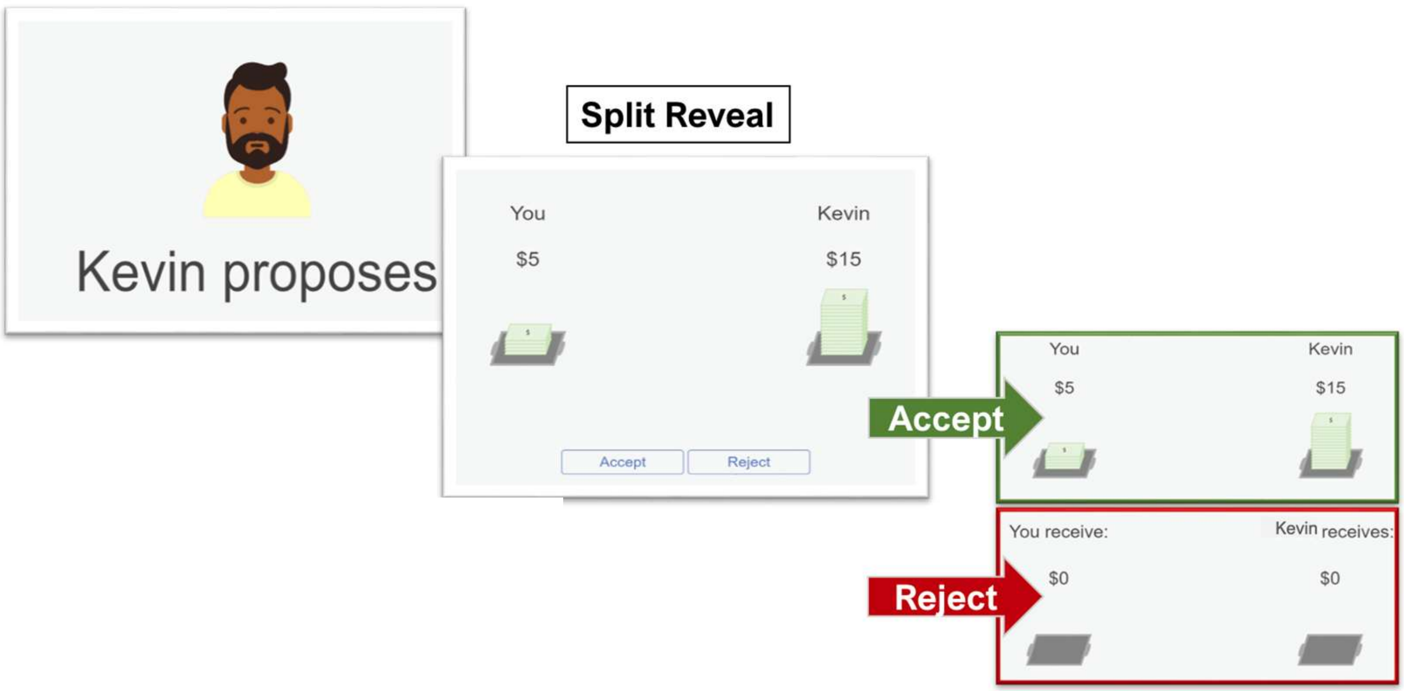

The Ultimatum Game

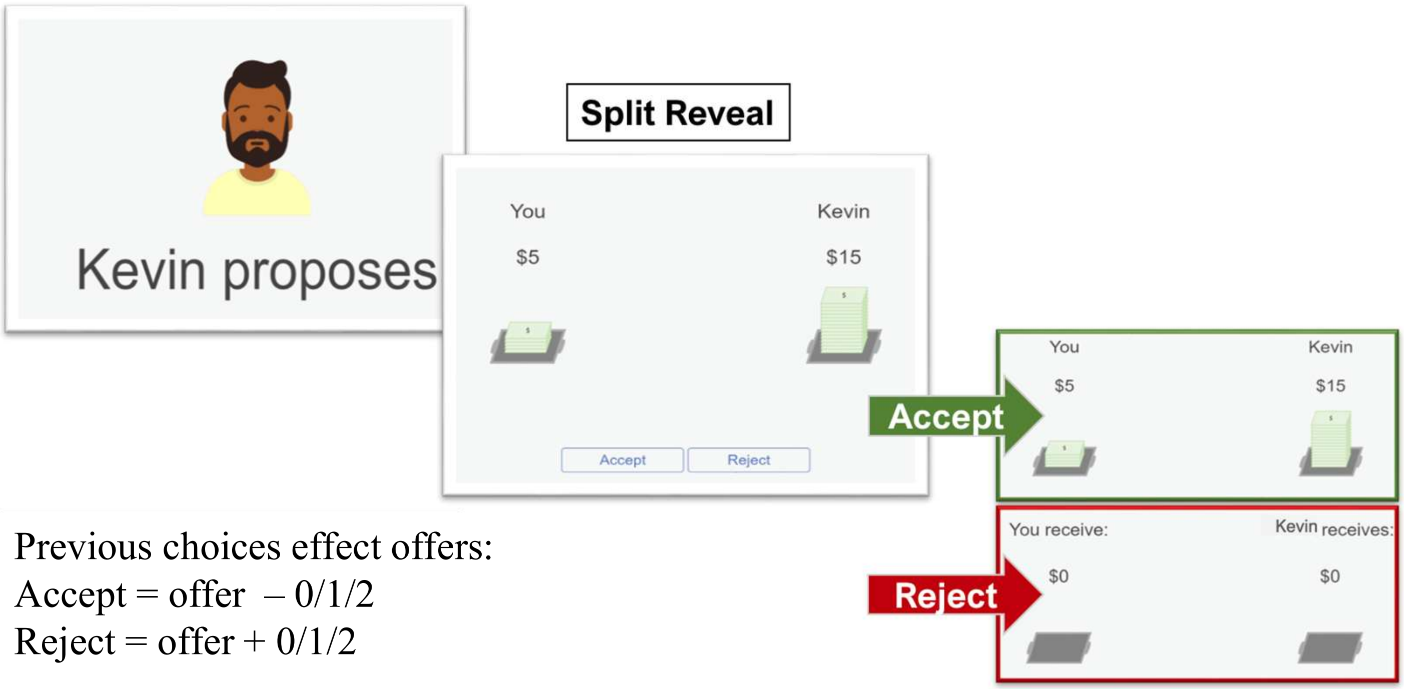

Previous choices effect offers:
Accept = offer – 0/1/2
Reject = offer + 0/1/2

Research Article
Neuroscience

Humans use forward thinking to exploit social controllability

Soojung Na, Dongil Chung, Andreas Hula, Ofer Perl, Jennifer Jung, Matthew Heflin, Sylvia Blackmore, Vincenzo G Fiore, Peter Dayan, Xiaosi Gu

Forward-Thinking Model

- Initial norms (*f0*)
- Envy (α)
- Adaptation rate (ε)
- Influence (δ)
- Inverse temperature (β)

Norm Learning

$$\text{Norm}^{t=1} = f0$$
$$\text{Norm}^{t+1} = \text{Norm}^{t} - \varepsilon(\text{Norm}^{t} - \text{Offer}^{t})$$

Value Function

$$U(\text{Offer})^{t} = \text{Offer}^{t} - \alpha(\text{Norm}^{t} - \text{Offer}^{t})$$
$$\text{Offer}^{t+i} = \begin{matrix} \text{Offer}^{t} + \delta \text{ if Choice} = \text{Reject} \\ \text{Offer}^{t} - \delta \text{ if Choice} = \text{Accept} \end{matrix}$$

$$Q(\text{Accept})^{t} = U(\text{Offer})^{t} + U(\text{Offer}|\text{Accept})^{t+1} + U(\text{Offer}|\text{Accept})^{t+2}$$
$$Q(\text{Reject})^{t} = U(\text{Offer}|\text{Reject})^{t+1} + U(\text{Offer}|\text{Reject})^{t+2}$$
$$Q(\text{Action})^{t} = Q(\text{Accept})^{t} - Q(\text{Reject})^{t}$$

Softmax Decision Rule

$$p\left(\text{Choice}^{t} = \text{Action}_{i}\right) = \frac{e^{\beta Q(\text{Action}_{i})^{t}}}{\sum_{j} e^{\beta Q(\text{Action}_{j})^{t}}}$$

Recoverability using best fitting parameter values

Adaptation rate

$r = 0.39$

Expected influence

$r = 0.88$

Sensitivity to norm PE

$r = 0.57$

Inv temp

$r = 0.77$

Initial norm

$r = 0.66$

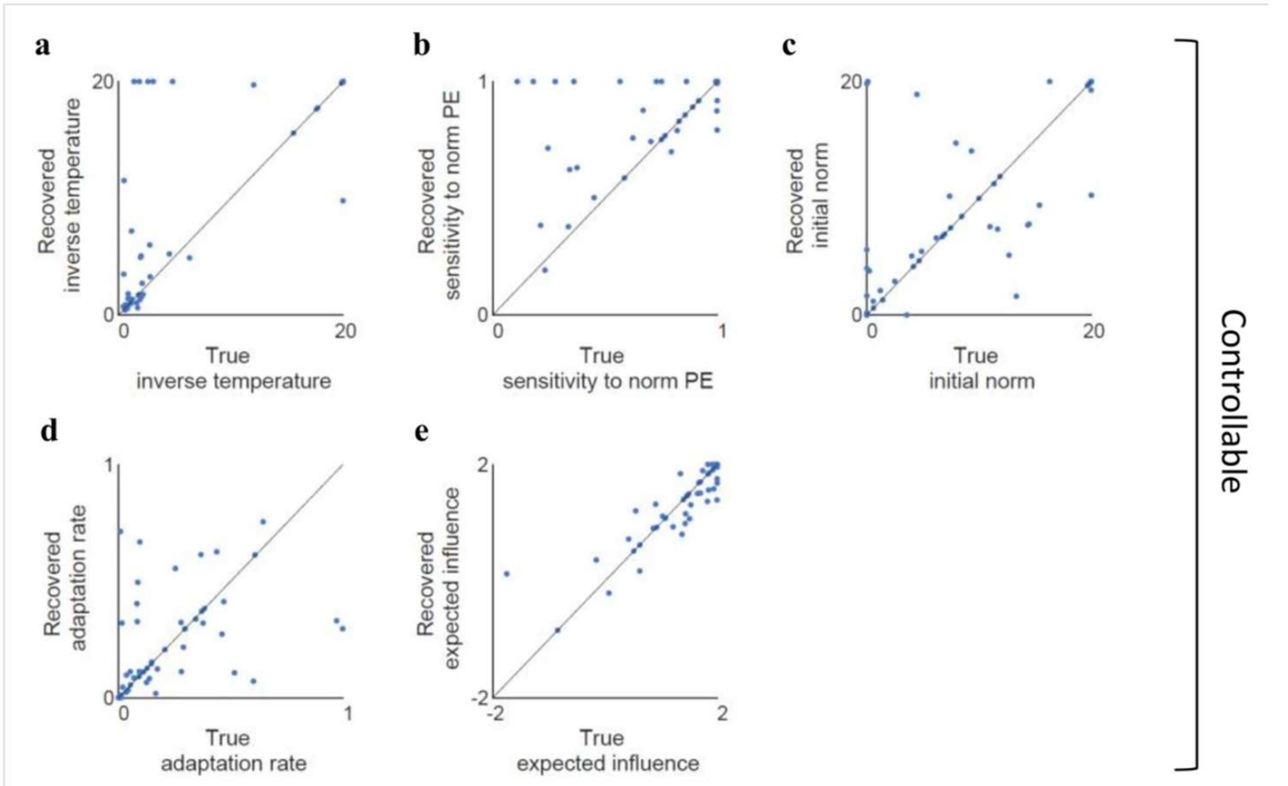

Recoverability using best fitting parameter values

Adaptation rate

$r = 0.39$

Expected influence

$r = 0.88$

Sensitivity to norm PE

$r = 0.57$

Inv temp

$r = 0.77$

Initial norm

$r = 0.66$

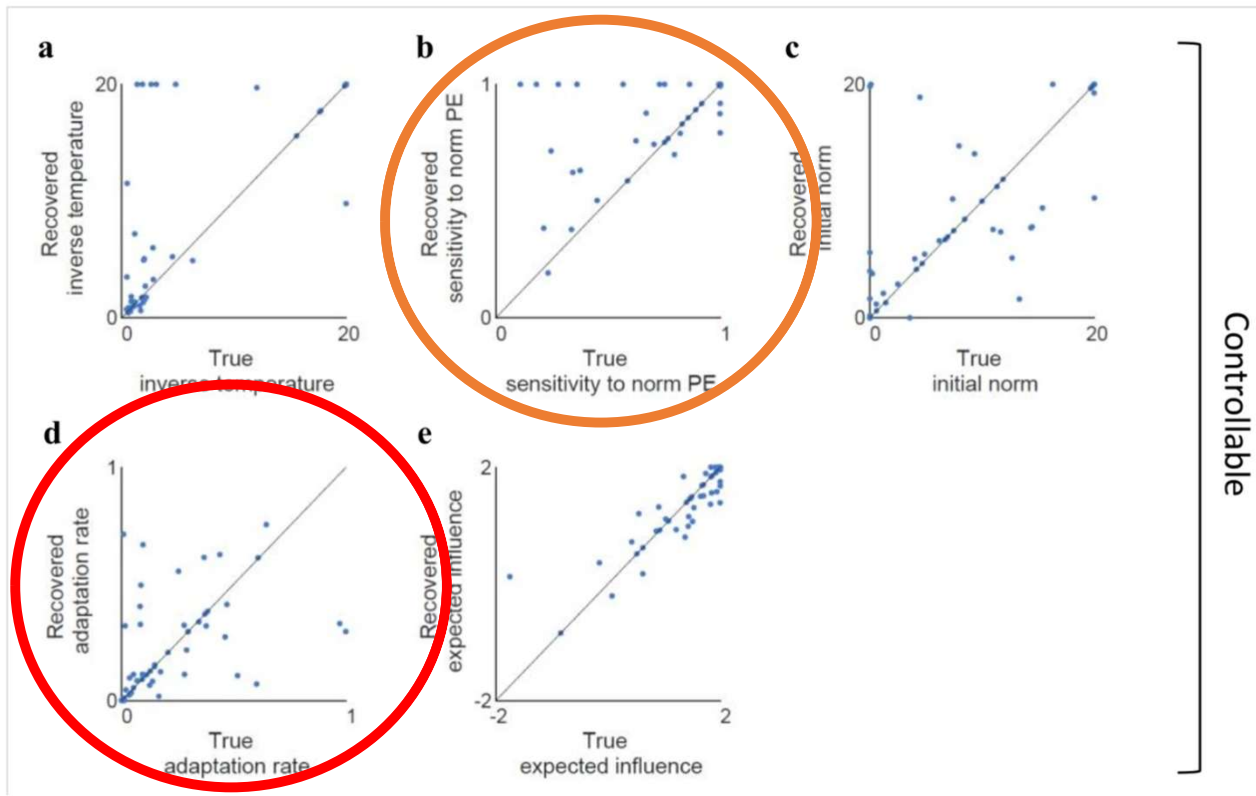

Replication of parameter recoverability

Recoverability changed

- Adaptation rate
  - r = 0.39-> 0.61
- Expected influence
  - r= 0.88 -> 0.91
- Sensitivity to norm PE
  - r = 0.57 -> 0.35
- Inv temp
  - r = 0.77 -> 0.82
- Initial norm
  - r = 0.66 -> 0.72

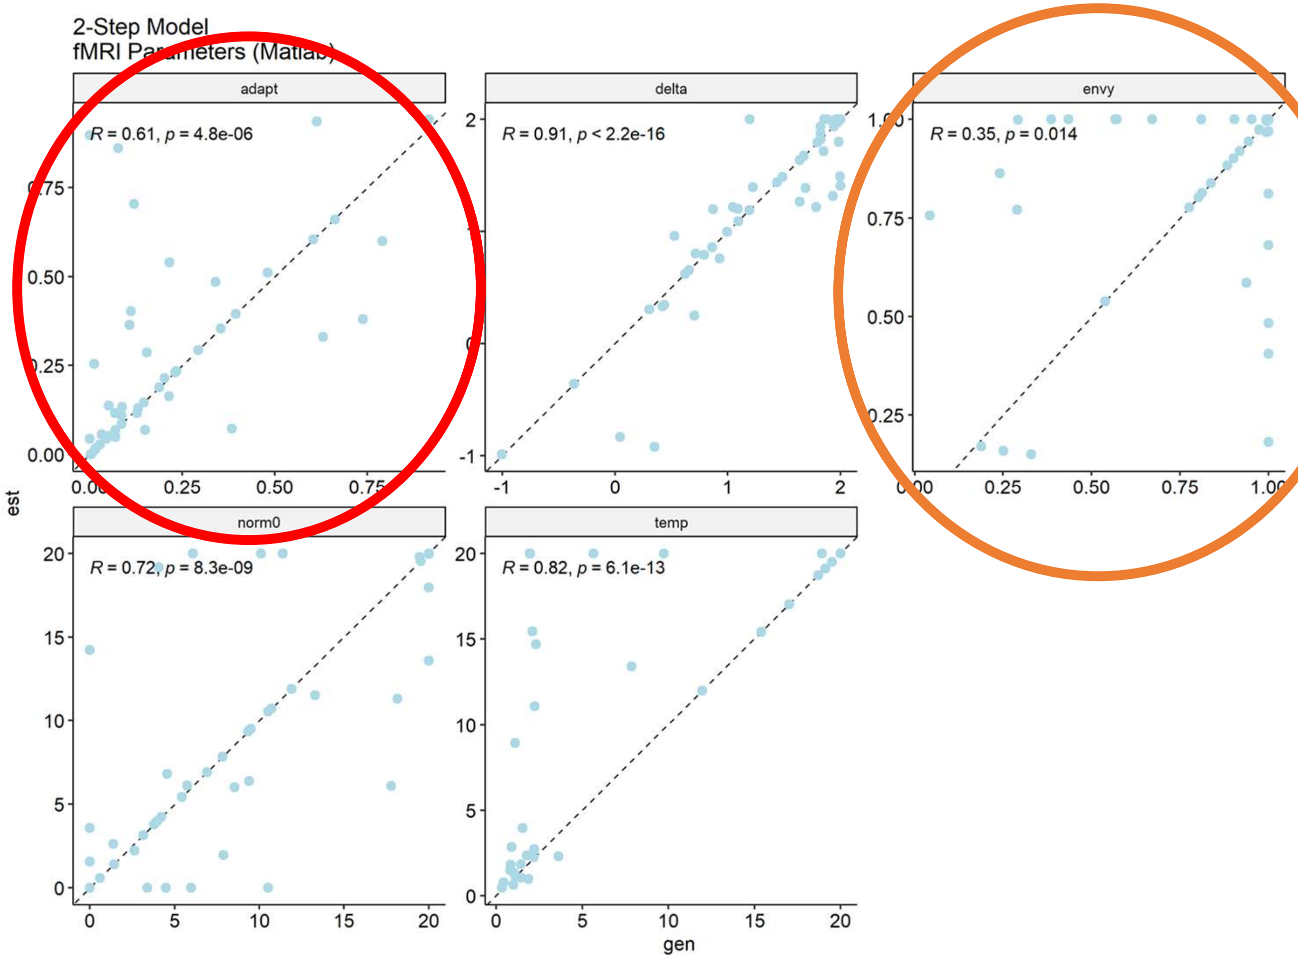

Recoverability with **random** parameters

Recoverability gets worse!

- Adaptation rate
  - 0.39-> 0.61 -> 0.25
- Expected influence
  - 0.88 -> 0.91 -> 0.90
- Sensitivity to norm PE
  - 0.57 -> 0.35 -> 0.31
- Inv temp
  - 0.77 -> 0.82 -> 0.38
- Initial norm
  - 0.66 -> 0.72 -> -0.009

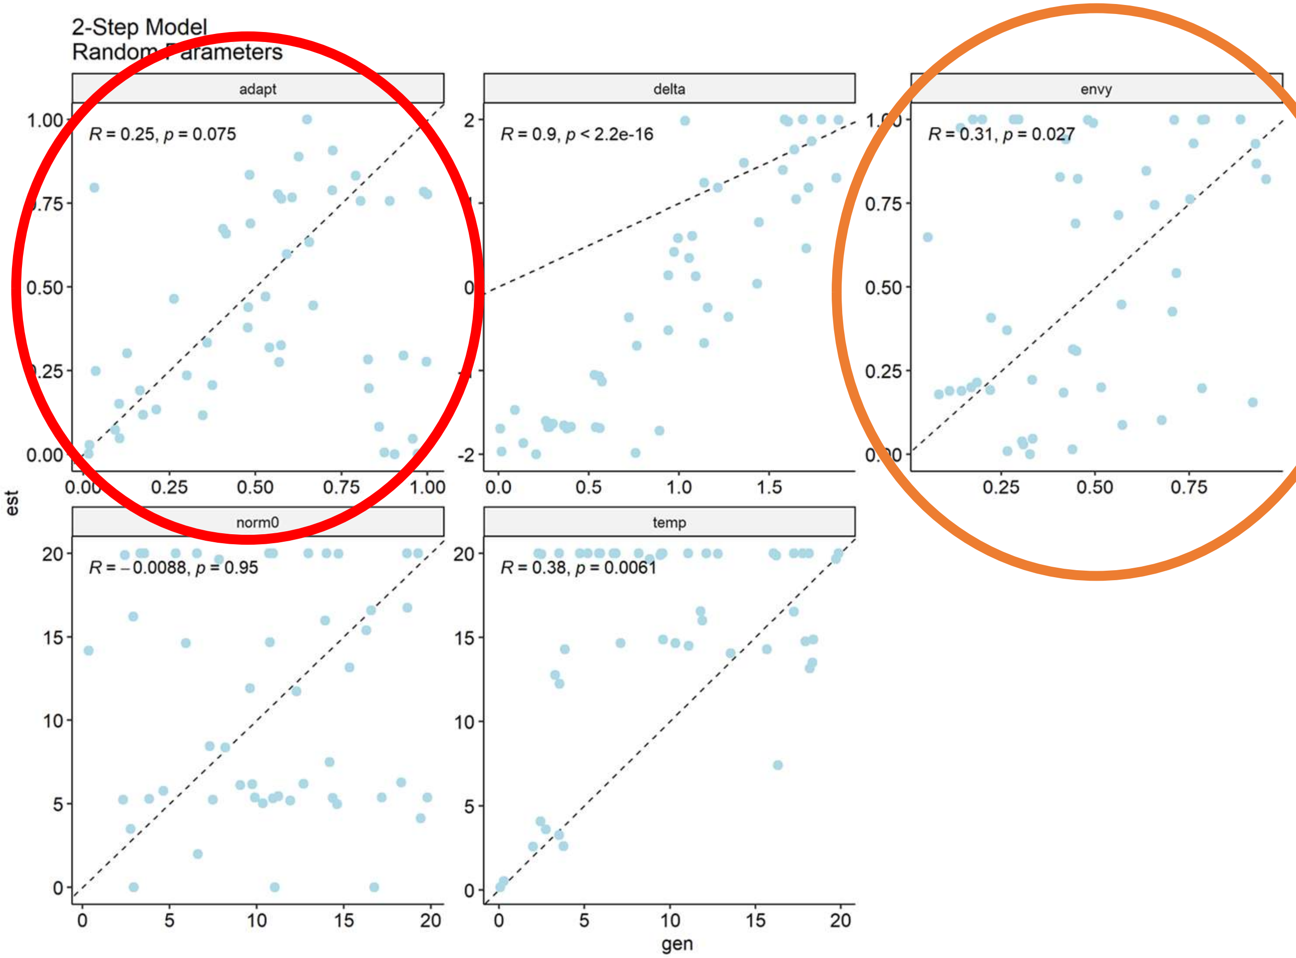

Recoverability with **random** parameters

Recoverability gets worse!

Adaptation rate

0.39-> 0.61 -> 0.25

Expected influence

0.88 -> 0.91 -> 0.90

Sensitivity to norm PE

0.57 -> 0.35 -> 0.31

Inv temp

0.77 -> 0.82 -> 0.38

Initial norm

0.66 -> 0.72 -> -0.009

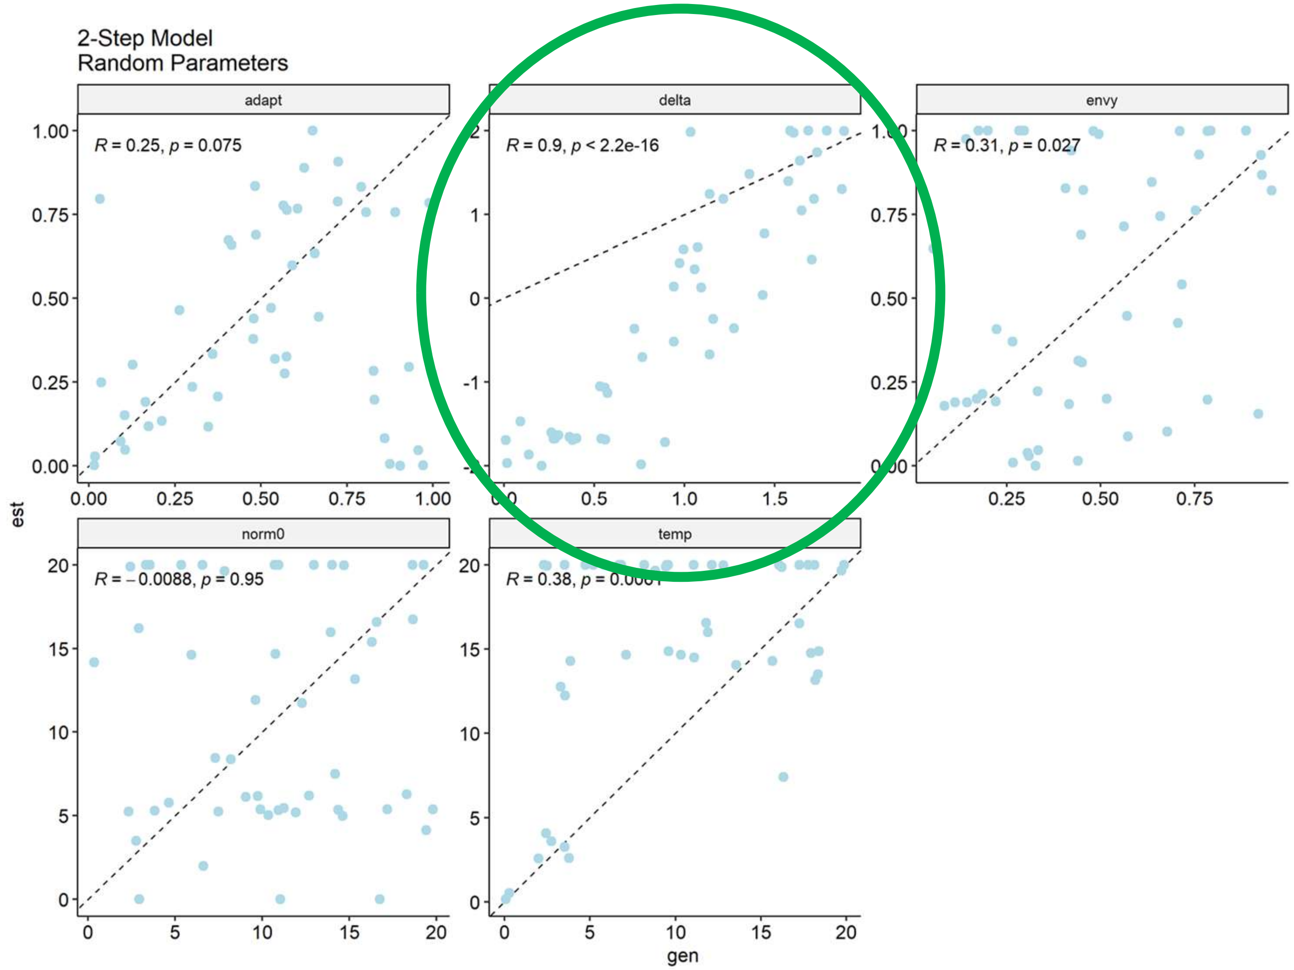

recoverability with **DDM** decision rule

Key changes:
- Removed two parameters
  - Initial norm = max(Offer) ÷ 2
  - No controllability
- Added four parameters
  - Boundary separation
    $r = 0.94$
  - Drift rate
    $r = 0.98$
  - Starting point
    $r = 0.71$
  - Non-decision time
    $r = 0.96$

Using **random** parameter values

Adaptation rate

$r = 0.25$ -> 0.6

Sensitivity to norm PE

$r = 0.31$ - > 0.55

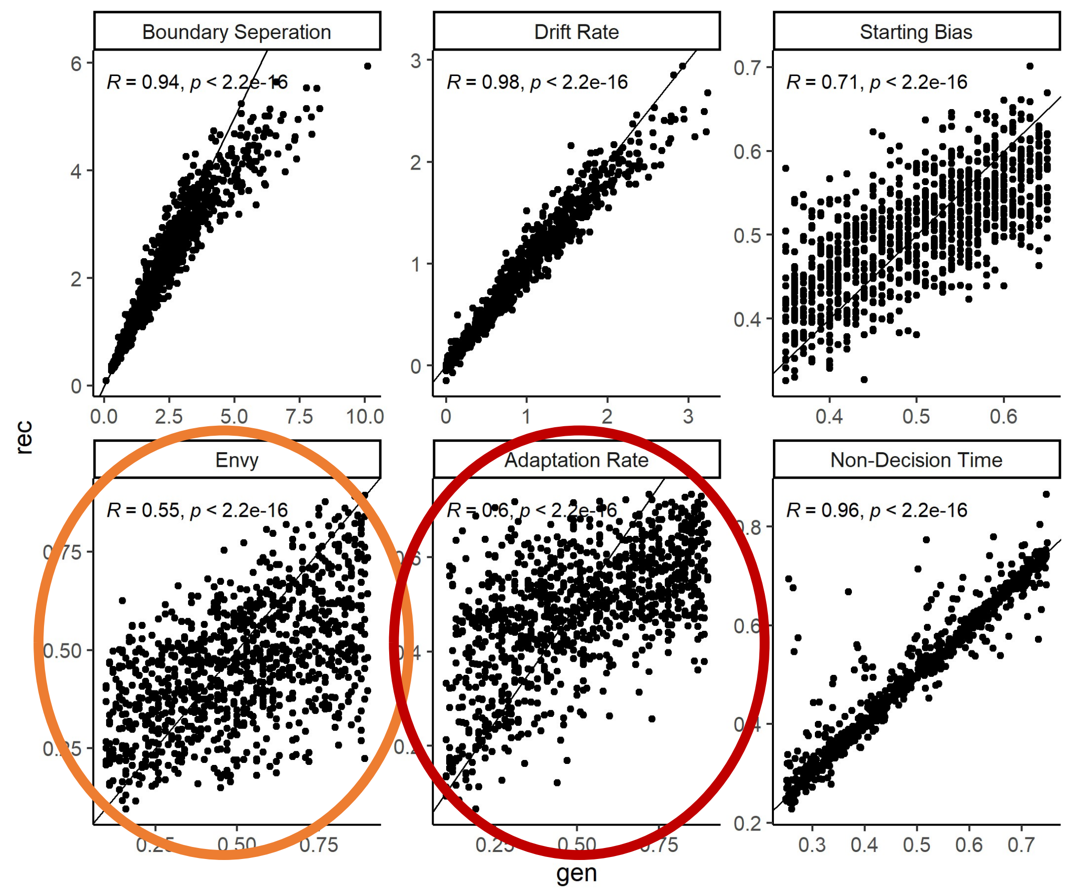

recoverability with **DDM** decision rule

Key changes:

- Removed two parameters
  - Initial norm = max(Offer) ÷ 2
  - No norm learning
- Added four parameters
  - Boundary separation
    $r = 0.94$
  - Drift rate
    $r = 0.98$
  - Starting point
    $r = 0.71$
  - Non-decision time
    $r = 0.96$

Using **random** parameter values

Controllability

$r = 0.55$

Sensitivity to unfairness

$r = 0.79$

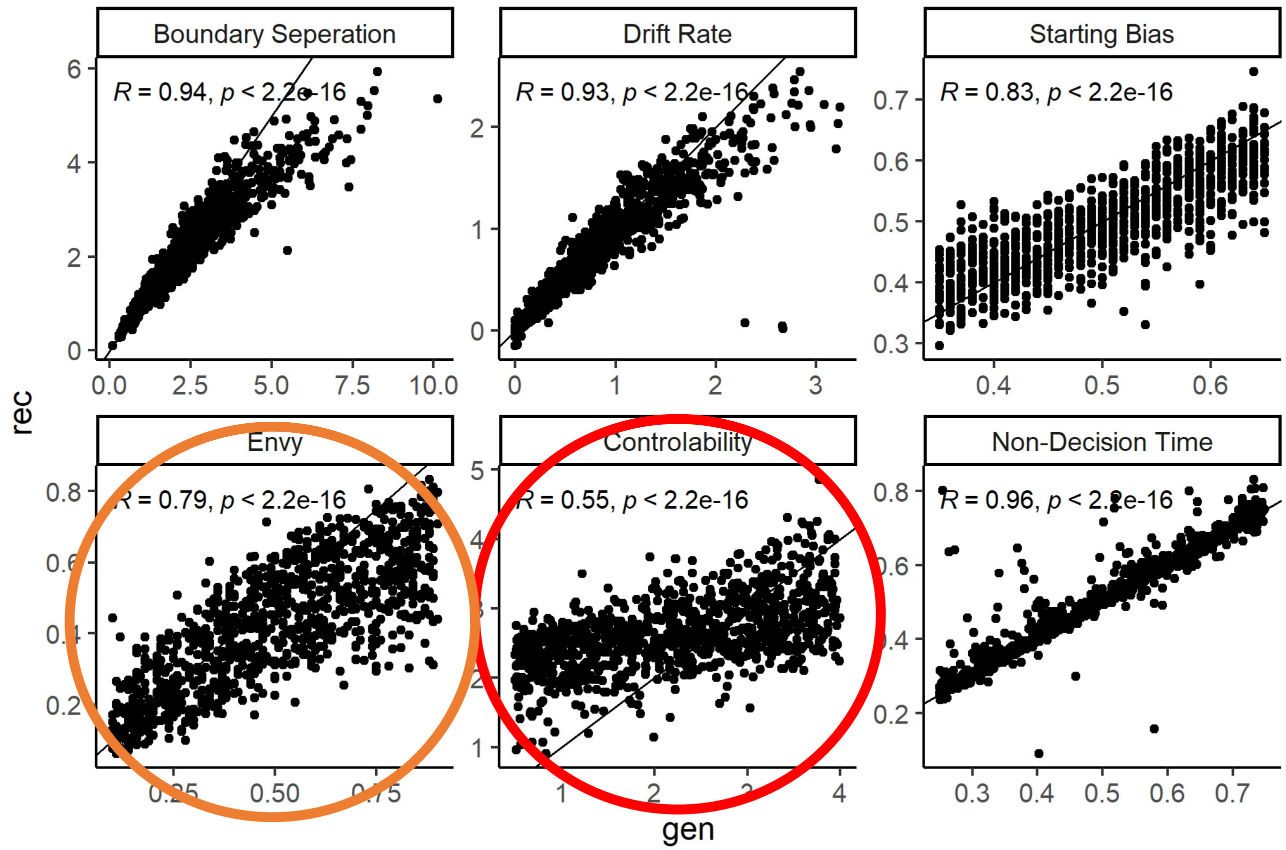

Application to Empirical Data

- Social Interactions & Borderline Personality Disorder (BPD)
  - BPD characterized by inability to maintain relationships, inflexibility, & need to control others
  - Hyp: BPD patients should show deficits in capacity to negotiate in social exchange task
- Data
  - BPD (n=29), Anxiety (ANX, n=33), healthy controls (HC, n=42)
  - UG task with and without controllability

behavioral analyses

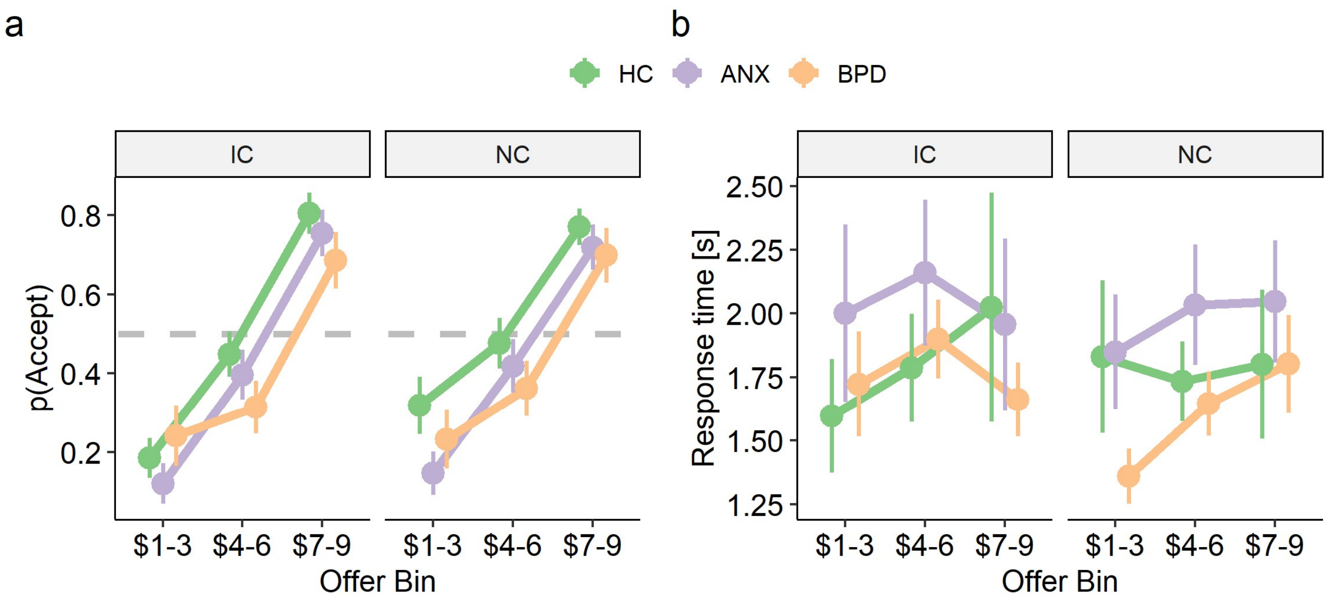

<u>Choice</u>

Main effect of group

- F(2, 531) = 3.47, P = 0.032

Post-hoc analysis

- $BPD - HC = -0.08,$ $P_{adj} = 0.065$
- $ANX - HC = -0.08,$ $P_{adj} = 0.070$

Note:
IC = In-control Condition
NC = No-control Condition

behavioral analyses

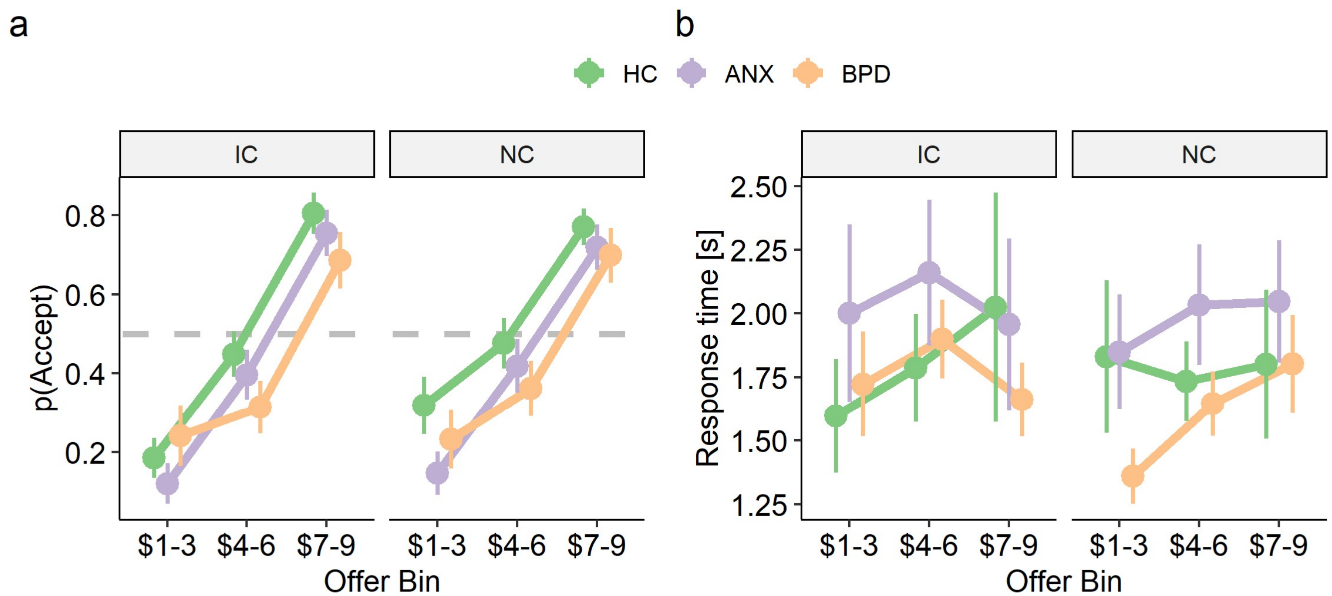

<u>Response time</u>

Main effect of group

- F(2, 531) = 4.31, P = 0.014

Post-hoc analysis

- $ANX - HC = 0.13,$ $P_{adj} = 0.018$
- $ANX - BPD = 0.12,$ $P_{adj} = 0.057$

Note:
IC = In-control Condition
NC = No-control Condition

preliminary modeling results

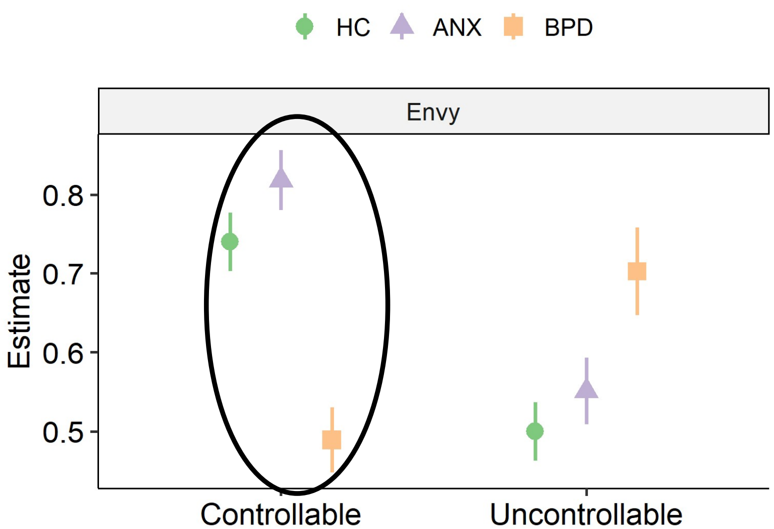

Envy

- *How strongly do decision-makers weigh norm violations?*
- Group x Block Interaction
  - $F(2,202) = 19.07$, $P < 0.001$
- Post-hoc analyses
  - $BPD_{\text{IC}} - HC_{\text{IC}} = -0.25, P_{adj} < 0.001$
  - $BPD_{\text{IC}} - ANX_{\text{IC}} = -0.33, P_{adj} < 0.001$

preliminary modeling results

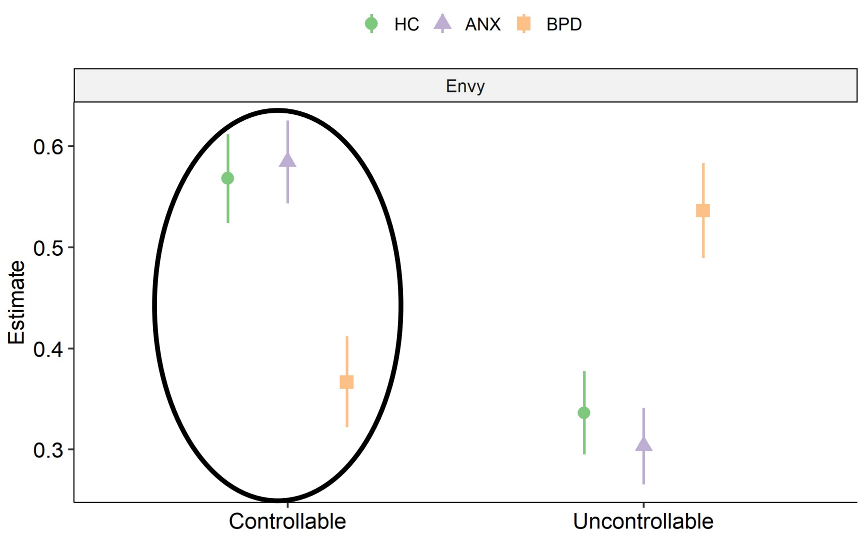

Envy

- *How strongly do decision-makers weigh norm violations?*
- Group x Block Interaction
  - $F(2,202) = 19.07$, $P < 0.001$
- Post-hoc analyses
  - $BPD_{\mathrm{IC}} - HC_{\mathrm{IC}} = -0.25, P_{adj} < 0.001$
  - $BPD_{\mathrm{IC}} - ANX_{\mathrm{IC}} = -0.33, P_{adj} < 0.001$

preliminary modeling results

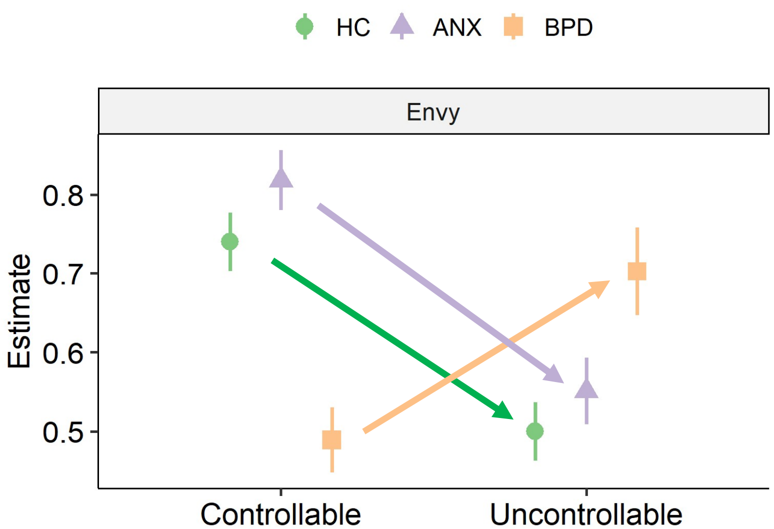

Aversion to Norm Violations

- *How strongly do decision-makers weigh norm violations?*
- Group x Block Interaction
  - $F(2,202) = 19.07$, $P < 0.001$
- Post-hoc analyses
  - $BPD_{IC} - HC_{IC} = -0.25, P_{adj} < 0.001$
  - $BPD_{IC} - ANX_{IC} = -0.33, P_{adj} < 0.001$
- Interpretation
  - When offers were controllable, BPD group weighed norm violations less strongly than both controls and ANX group
  - When offers were *not* controllable, BPD group weighed norm violations more strongly than control group

preliminary modeling results

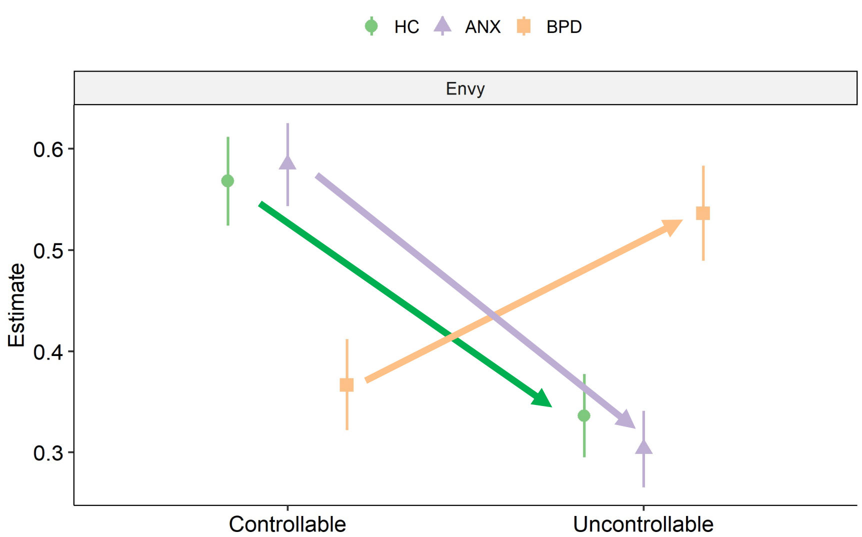

Envy

- *How strongly do decision-makers weigh norm violations?*
- Group x Block Interaction
  - $F(2,202) = 19.07$, $P < 0.001$
- Post-hoc analyses
  - $BPD_{\text{IC}} - HC_{\text{IC}} = -0.25, P_{adj} < 0.001$
  - $BPD_{\text{IC}} - ANX_{\text{IC}} = -0.33, P_{adj} < 0.001$
- Interpretation
  - When offers were controllable, BPD group weighed norm violations less strongly than both controls and ANX group
  - When offers were *not* controllable, BPD group weighed norm violations more strongly than control group

preliminary modeling results

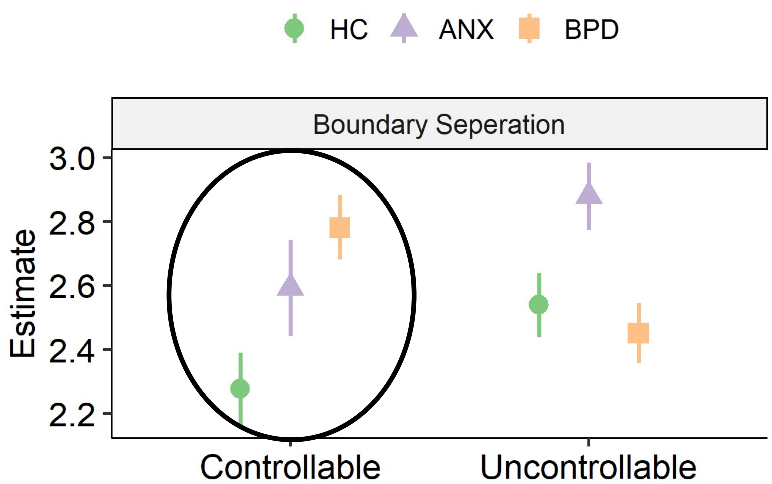

Boundary Separation

- *How cautious are decision-makers?*
- Group x Block Interaction
    - $F(2,202) = 4.27$, $P = 0.015$
- Post-hoc analyses
    - $BPD_{\mathrm{IC}} - HC_{\mathrm{IC}} = 0.50, P_{adj} = 0.024$
- Interpretation
    - When offers are controllable, BPD patients are more cautious than healthy controls

preliminary modeling results

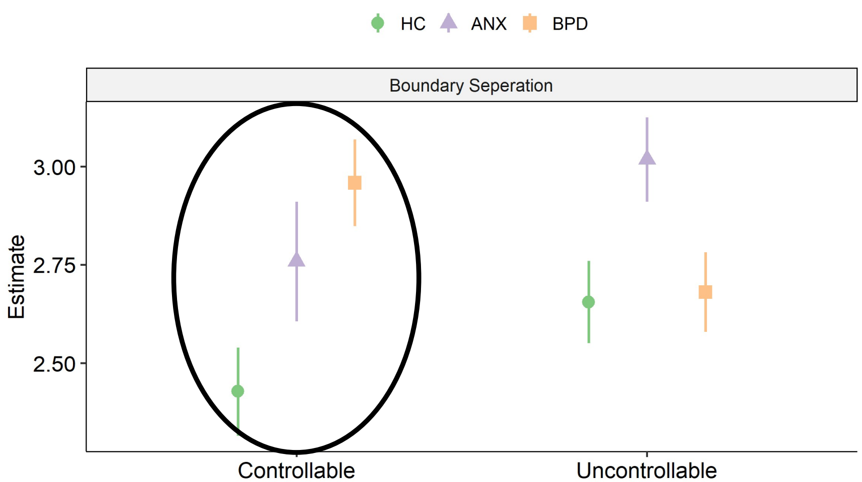

<u>Boundary Separation</u>

- *How cautious are decision-makers?*
- Group x Block Interaction
  - $F(2{,}202) = 4.27$, $P = 0.015$
- Post-hoc analyses
  - $BPD_{\mathrm{IC}} - HC_{\mathrm{IC}} = 0.50, P_{adj} = 0.024$
- Interpretation
  - When offers are controllable, BPD patients are more cautious than healthy controls

preliminary modeling results

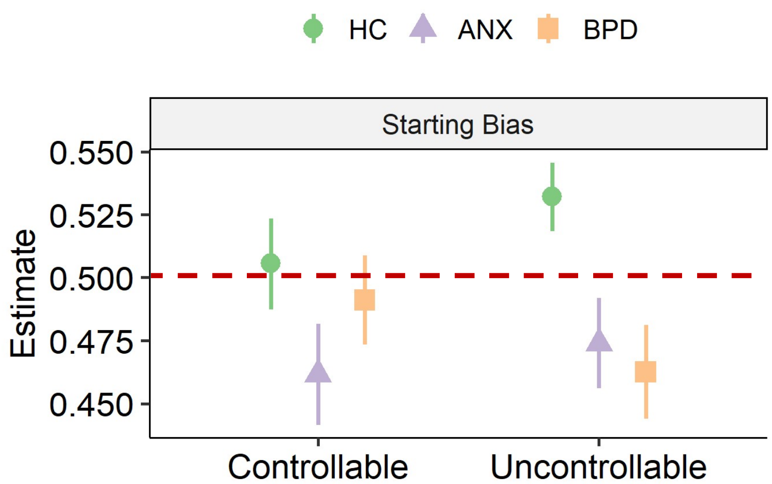

<u>Starting Bias</u>

- *How biased are decision-makers towards one type of response?*
- Main effect of Group
    - $F(2,202) = 5.20$, P = 0.006
- Post-hoc analyses
    - $ANX - HC = -0.05, P_{adj} = 0.009$
    - $BPD - HC = -0.04, P_{adj} = 0.049$
- Interpretation
    - Both clinical groups are more likely to punish opponents, regardless of amount

preliminary modeling results

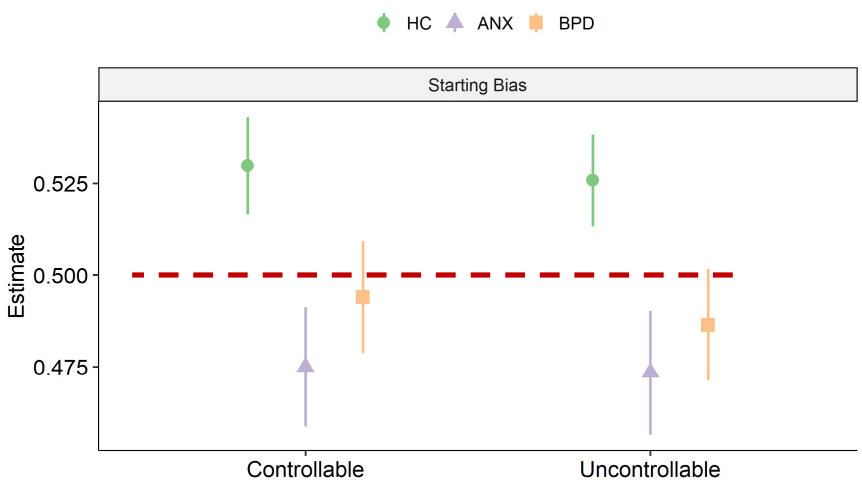

<u>Starting Bias</u>

- *How biased are decision-makers towards one type of response?*
- Main effect of Group
  - $F(2{,}202) = 5.20$, $P = 0.006$
- Post-hoc analyses
  - $ANX - HC = -0.05, P_{adj} = 0.009$
  - $BPD - HC = -0.04, P_{adj} = 0.049$
- Interpretation
  - Both clinical groups are more likely to punish opponents, regardless of amount

preliminary modeling results

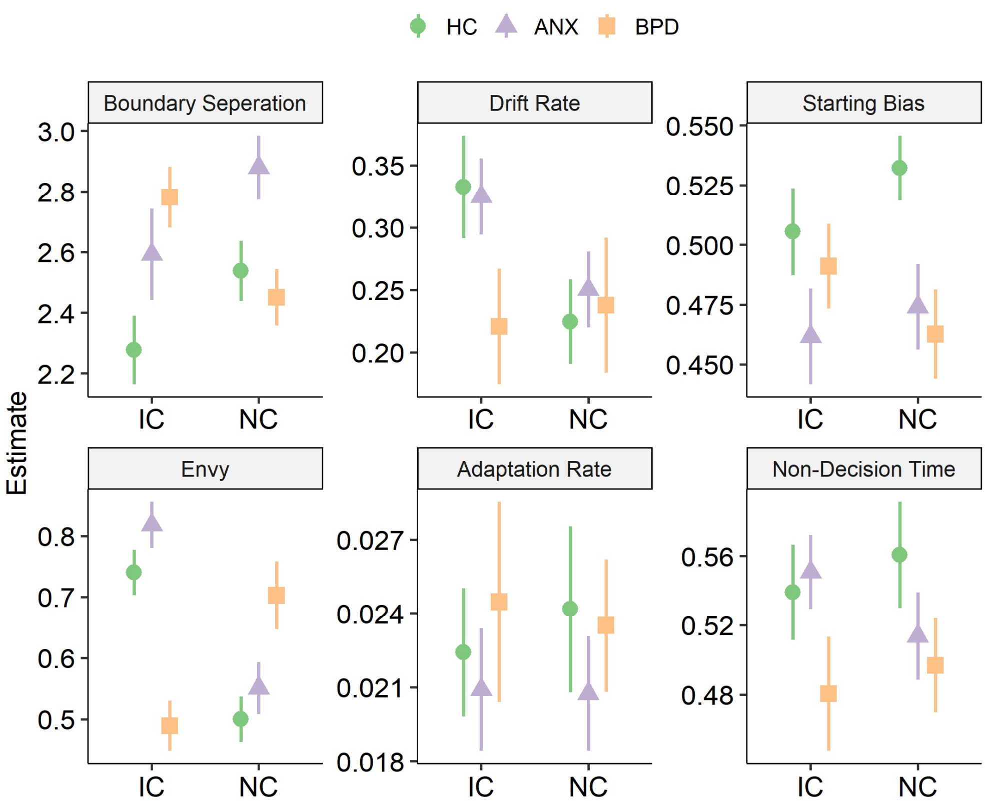

No significant group differences or group-condition interactions for *Drift Rate*, *Adaptation Rate*, or *Non-Decision time*

Note:
IC = In-control Condition
NC = No-control Condition

preliminary modeling results

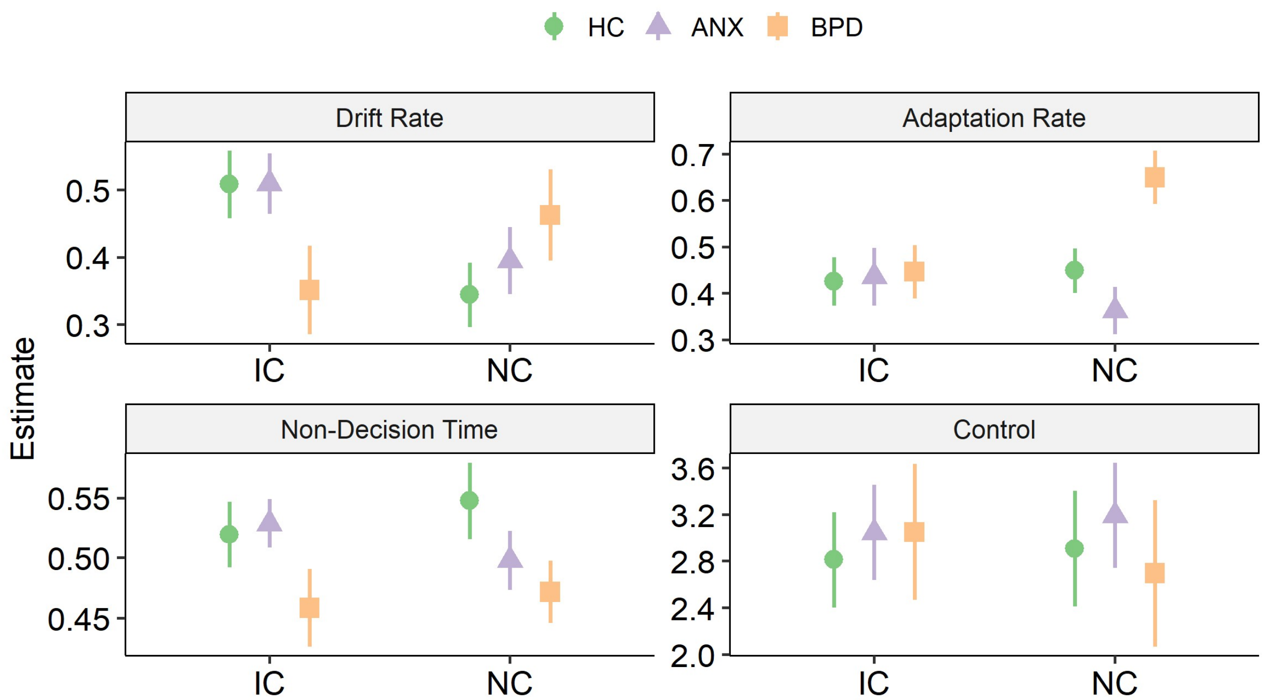

No significant group differences or group-condition interactions for *Drift Rate*, *Adaptation Rate*, *Non-Decision time*, or *Control*

Note:
IC = In-control Condition
NC = No-control Condition

preliminary modeling results

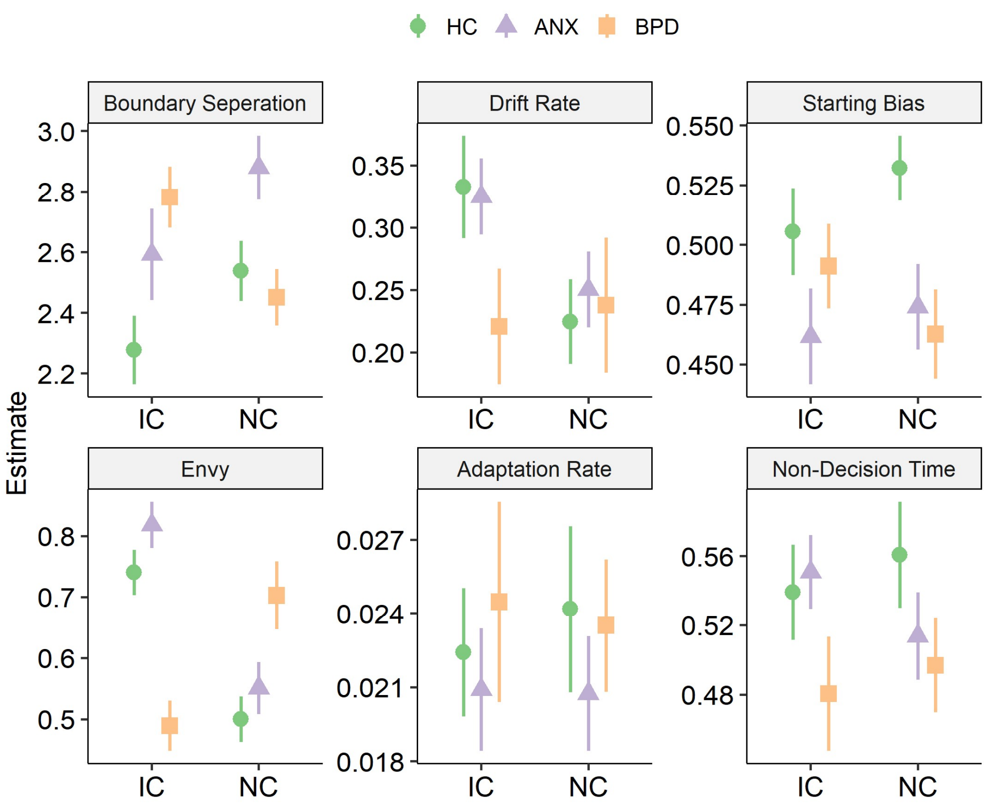

Summary

- Individuals with BPD or Anxiety are more punitive than controls
- Individuals with BPD are generally more cautious than controls when making decisions *when they have control*
- Individuals with BPD use of norm violations differs greatly from both controls and those with anxiety

Note:
IC = In-control Condition
NC = No-control Condition

take-aways

- Response times are information
- Response times improve reliability and portability of cognitive models

best practices

- Parameter recovery and model identifiability
- Build your models with your tasks

NYCPW

HOME ABOUT TEAM SCHEDULE FAQ APPLY

NYComputational Psychiatry Workshop

November 10-12, 2025

A specialized workshop bringing together computational methods and clinical expertise to advance understanding of mental health and psychiatric disorders.

What is NYCPW?

The **New York Computational Psychiatry Workshop** (NYCPW) is an immersive, hands-on three-day workshop designed to provide trainees with the latest tools and knowledge in computational psychiatry. Our aim is to bridge the gap between computational methods and clinical applications.

The workshop will take place at the Mount Sinai **Center for Computational Psychiatry** (**55 W 125th St, New York, NY 10027**) on **November 10-12, 2025**.

Thank you

Acknowledgements
Xiaosi Gu (Yale)
Laura Berner (ISMMS)
Kianté Fernandez (UCLA)
Ian Krajbich (UCLA)
Roger Ratcliff (OSU)
Laura Fontanesi (Basel)
Robert (Bob) Wilson (Arizona, Georgia Tech)

The Forward Thinking Model

- Five parameter model:
    - Initial norms (*f0*)
    - Envy (α)
    - Adaptation rate (ε)
    - Influence (δ)
    - Inverse temperature (β)

The Forward Thinking Model

- Five parameter model:
  - Initial norms (*f0*)
  - Envy (α)
  - Adaptation rate (ε)
  - Influence (δ)
  - Inverse temperature (β)

What is fair?

Expectation at start of trial

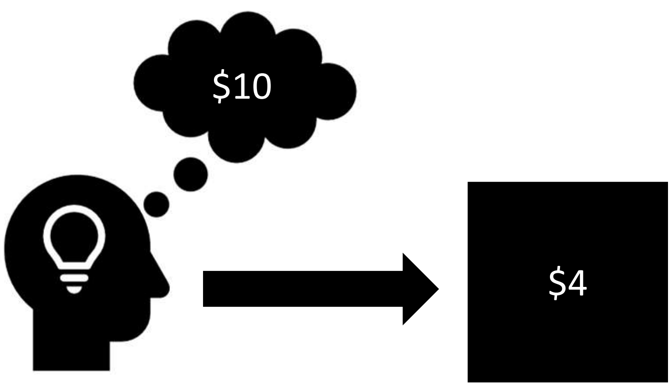

The Forward Thinking Model

- Five parameter model:
  - Initial norms (*f0*)
  - Envy (α)
  - Adaptation rate (ε)
  - Influence (δ)
  - Inverse temperature (β)

Desire for fairness

Aversion for inequality

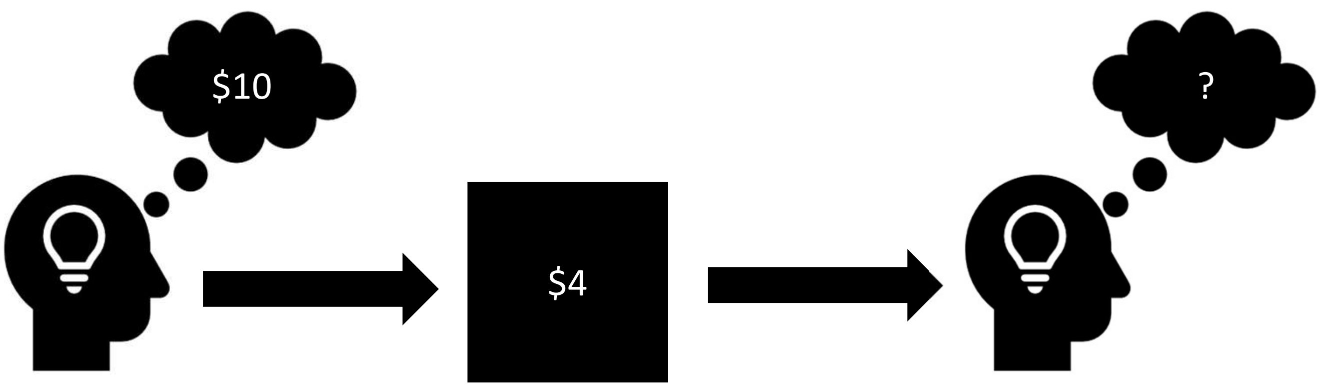

The Forward Thinking Model

- Five parameter model:
  - Initial norms (*f0*)
  - Envy (α)
  - Adaptation rate (ε)
  - Influence (δ)
  - Inverse temperature (β)

Desire for fairness

Aversion for inequality

$$u(\text{Offer}|a) = \text{Offer} - a(\text{Norm} - \text{Offer})$$

$$u(\$4|a) = \$4 - \alpha(\$10 - \$4) = \$4 - \alpha(\$6)$$

$$u(\$4|0.4) = \$4 - (0.4)(\$6) = \$4 - \$2.4 = \$1.60$$

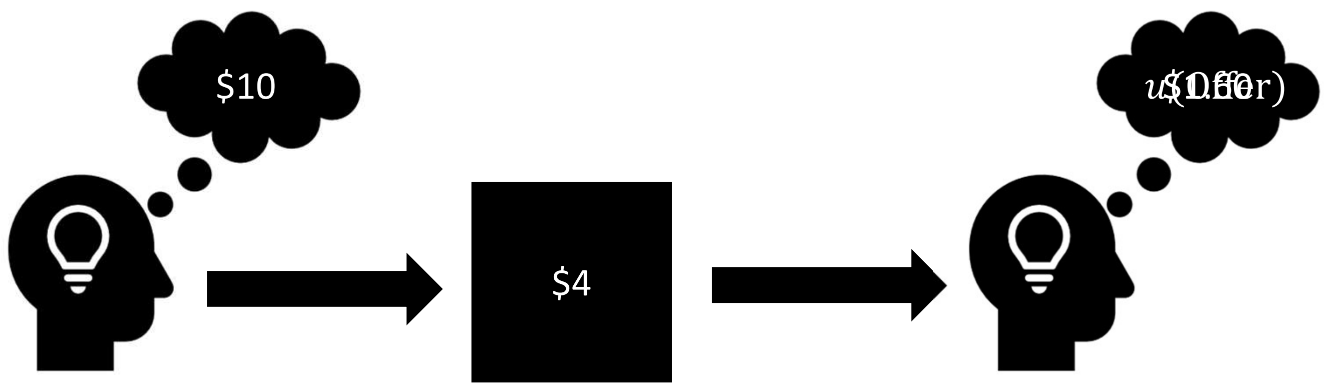

The Forward Thinking Model

- Five parameter model:
  - Initial norms (*f0*)
  - Envy (α)
  - Adaptation rate (ε)
  - Influence (δ)
  - Inverse temperature (β)

Rate of updating norms

Influence of new information

$$\text{pr}(\text{Norm}_{t+1}|\varepsilon) = \text{Norm}_t + \varepsilon(\text{Offer} - \text{Norm}_t)$$

$$\text{pr}(\text{Norm}_{t+1}|\varepsilon) = 10 + \varepsilon(4 - 10) = 10 + \varepsilon(-6)$$

$$\text{pr}(\text{Norm}_{t+1}|.65) = 10 + (0.65)(-6) = 6.10$$

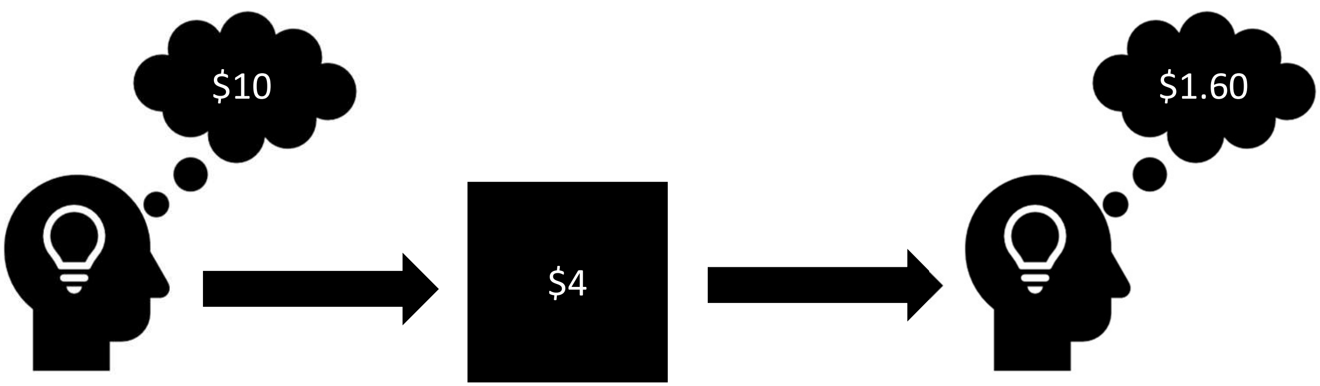

The Forward Thinking Model

- Five parameter model:
  - Initial norms (*f0*)
  - Envy (α)
  - Adaptation rate (ε)
  - Influence (δ)
  - Inverse temperature (β)

Influence on new offers

How much change is possible

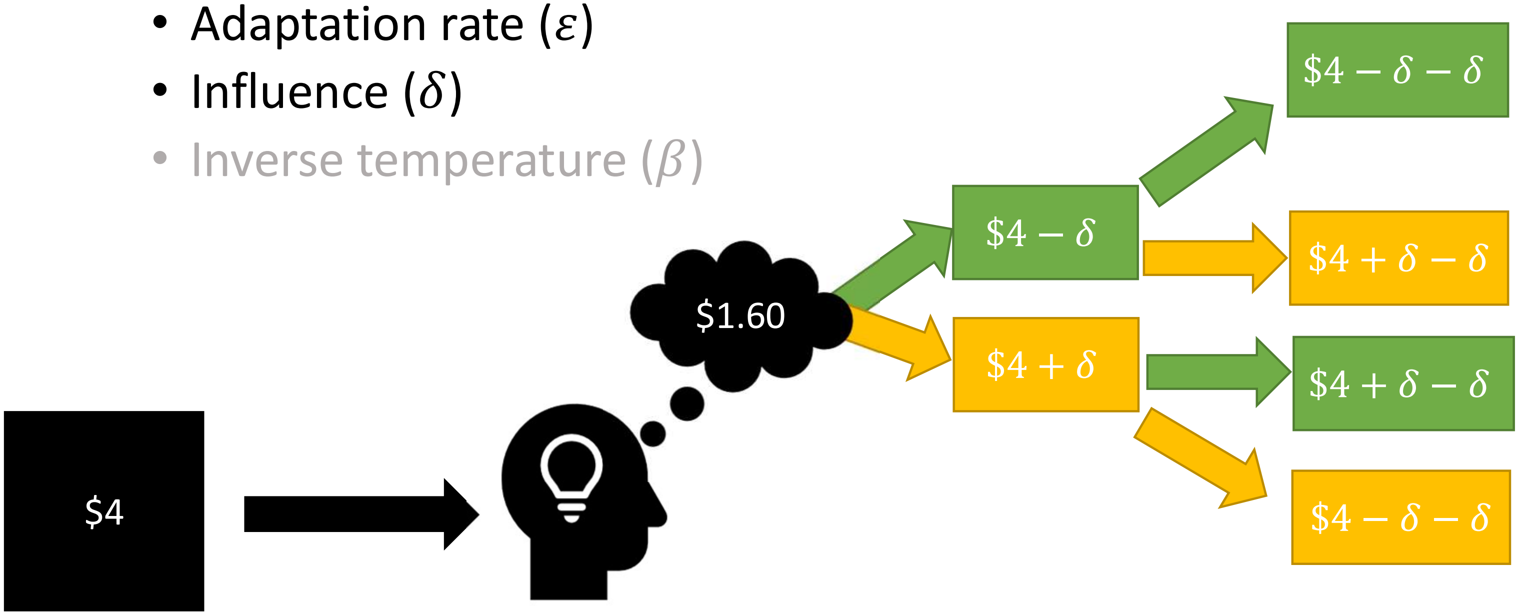

The Forward Thinking Model

- Five parameter model:
  - Initial norms (*f0*)
  - Envy (α)
  - Adaptation rate (ε)
  - Influence (δ)
  - Inverse temperature (β)

Influence on new offers

How much change is possible

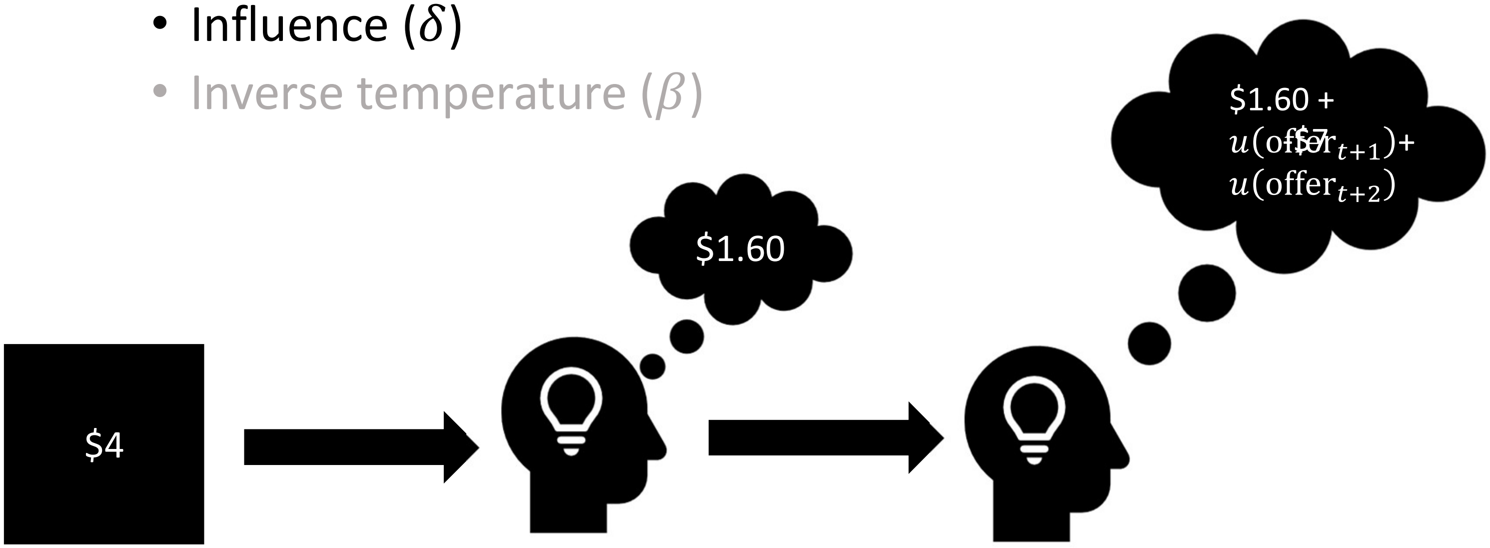

The Forward Thinking Model

- Five parameter model:
  - Initial norms (*f0*)
  - Envy (α)
  - Adaptation rate (ε)
  - Influence (δ)
  - Inverse temperature (β)

How much value influences choice

Choice consistency

$$P(\text{Accept}) = \frac{1}{(1 + e^{-\beta \cdot U})}$$

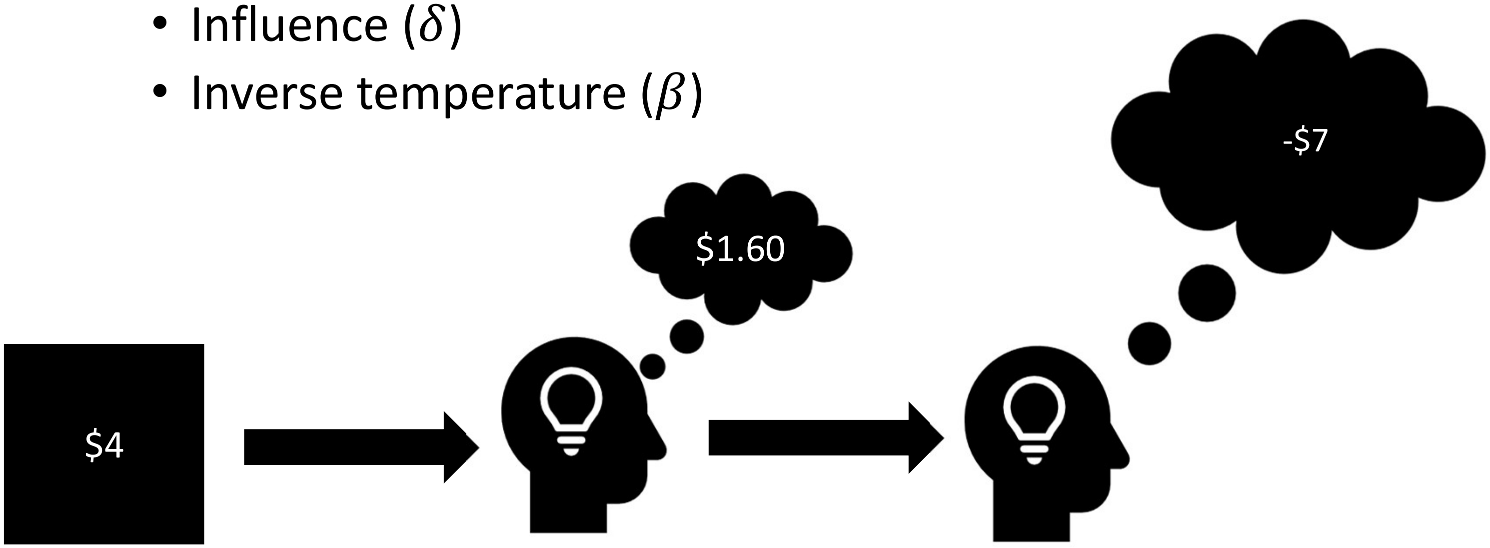

The Forward Thinking Model

- Five parameter model:
  - Initial norms (*f0*)
  - Envy (α)
  - Adaptation rate (ε)
  - Influence (δ)
  - Inverse temperature (β)

How much value influences choice

Choice consistency

$$P(\text{Accept}) = \frac{1}{(1+e^{-\beta \cdot U})} = \frac{1}{\left(1+e^{-\beta \cdot -7}\right)} = \frac{1}{(1+e^{-(1.5) \cdot -7})}$$

$$P(\text{Accept}) = 2.8e^{-5}$$

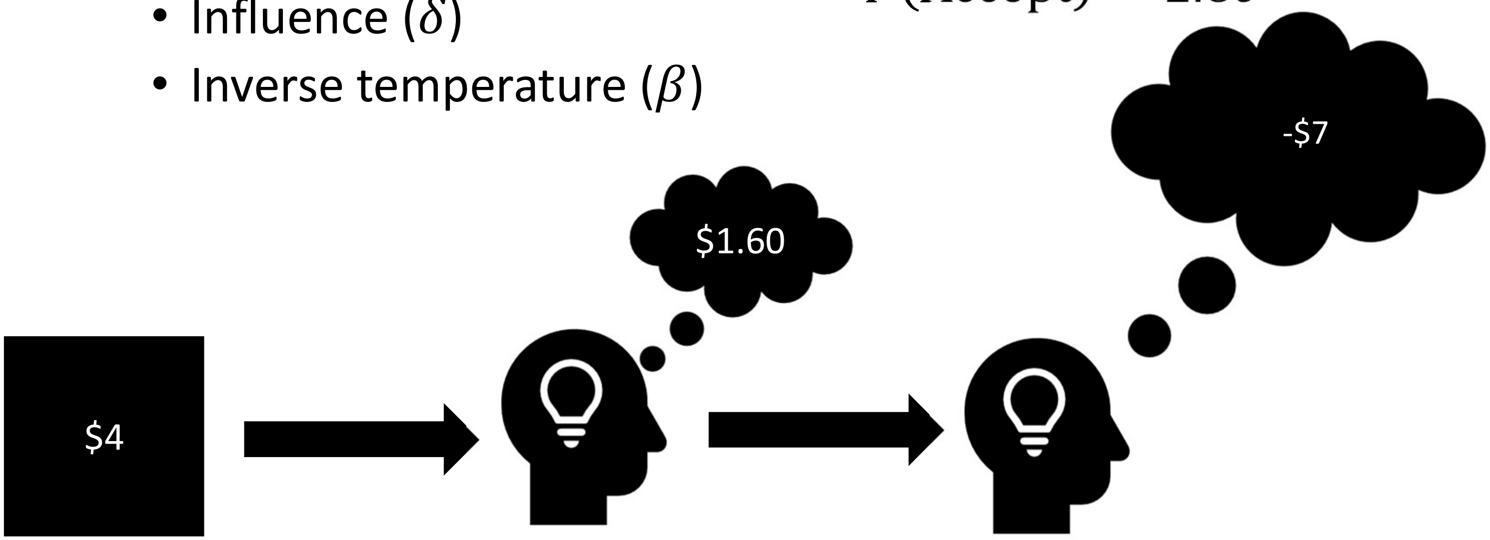

Estimation Procedure

- Quasi-Grid Search
  - Grid of 256 value combinations
  - Initial values [-1,0,1]
- Optimizer
  - BFGS Quasi-Newtonian method
  - Cubic line search

5, 4, 3,…,9

1, 1, 0, …, 1

$\{f0, \alpha, \varepsilon, \delta, \beta\}$

$\{0,1,1,1,1\} \rightarrow \{9,0.65,0.13,1.5,7.2\} \rightarrow$ LL

Estimation Procedure

- Quasi-Grid Search
  - Grid of 256 value combinations
  - Initial values [-1,0,1]
- Optimizer
  - BFGS Quasi-Newtonian method
  - Cubic line search

5, 4, 3,…,9

1, 1, 0, …, 1

$\{f0, \alpha, \varepsilon, \delta, \beta\}$

$\{0,1,1,1,1\} \rightarrow \{9,0.65,0.13,1.5,7.2\} \rightarrow -13$

$\{0,0,1,1,1\} \rightarrow \{8,0.99,0.21,1.7,7.3\} \rightarrow -11$

$\{0,0,0,1,1\} \rightarrow \{8,0.01,0.11,1.3,2.2\} \rightarrow -19$

…

$\{1,1,1,1,1\} \rightarrow \{10,0.75,0.45,0.6,1.1\} \rightarrow -14$

Parameter Recoverability

- Procedure
  - Take a set of parameters X_{gen}
  - Generate simulated data using X_{gen}
  - Fit model to simulated data to estimate X_{est}
  - Assess correlations between $X_{gen}$and X_{est}
- Parameter sets
  - Study 1: estimates from fMRI data (*N = 48; t*=30*;* Na et al., 2020, *eLife*)
  - Study 2: estimates from Prolific data (*N* = 711; t=20 *;* Na et al., 2020, *eLife*)
  - Study 3: random parameters (*N = 50; t*=100)

Baseline Recoverability

eLife | Home Magazine Community About

Research Article
Neuroscience

Humans use forward thinking to exploit social controllability

Soojung Na, Dongil Chung, Andreas Hula, Ofer Perl, Jennifer Jung, Matthew Heflin, Sylvia Blackmore, Vincenzo G Fiore, Peter Dayan, Xiaosi Gu

Recoverability with 30 trial (out of 40, fMRI sample)

- Adaptation rate
  - 0.39
- Delta
  - 0.88
- Envy
  - 0.57
- Temp
  - 0.77
- Initial norm
  - 0.66

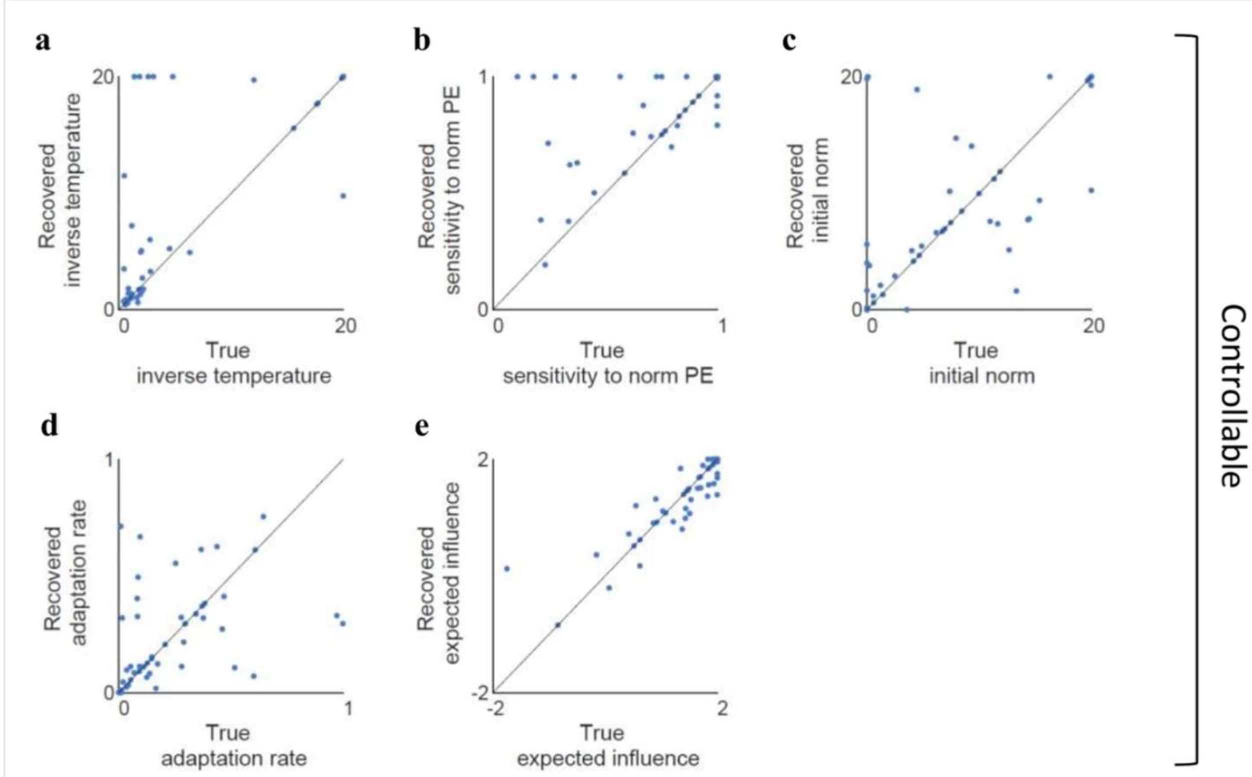

Baseline Recoverability

Recoverability with 30 trial (out of 40, fMRI sample)

- Adaptation rate
  - 0.39
- Delta
  - 0.88
- Envy
  - 0.57
- Temp
  - 0.77
- Initial norm
  - 0.66

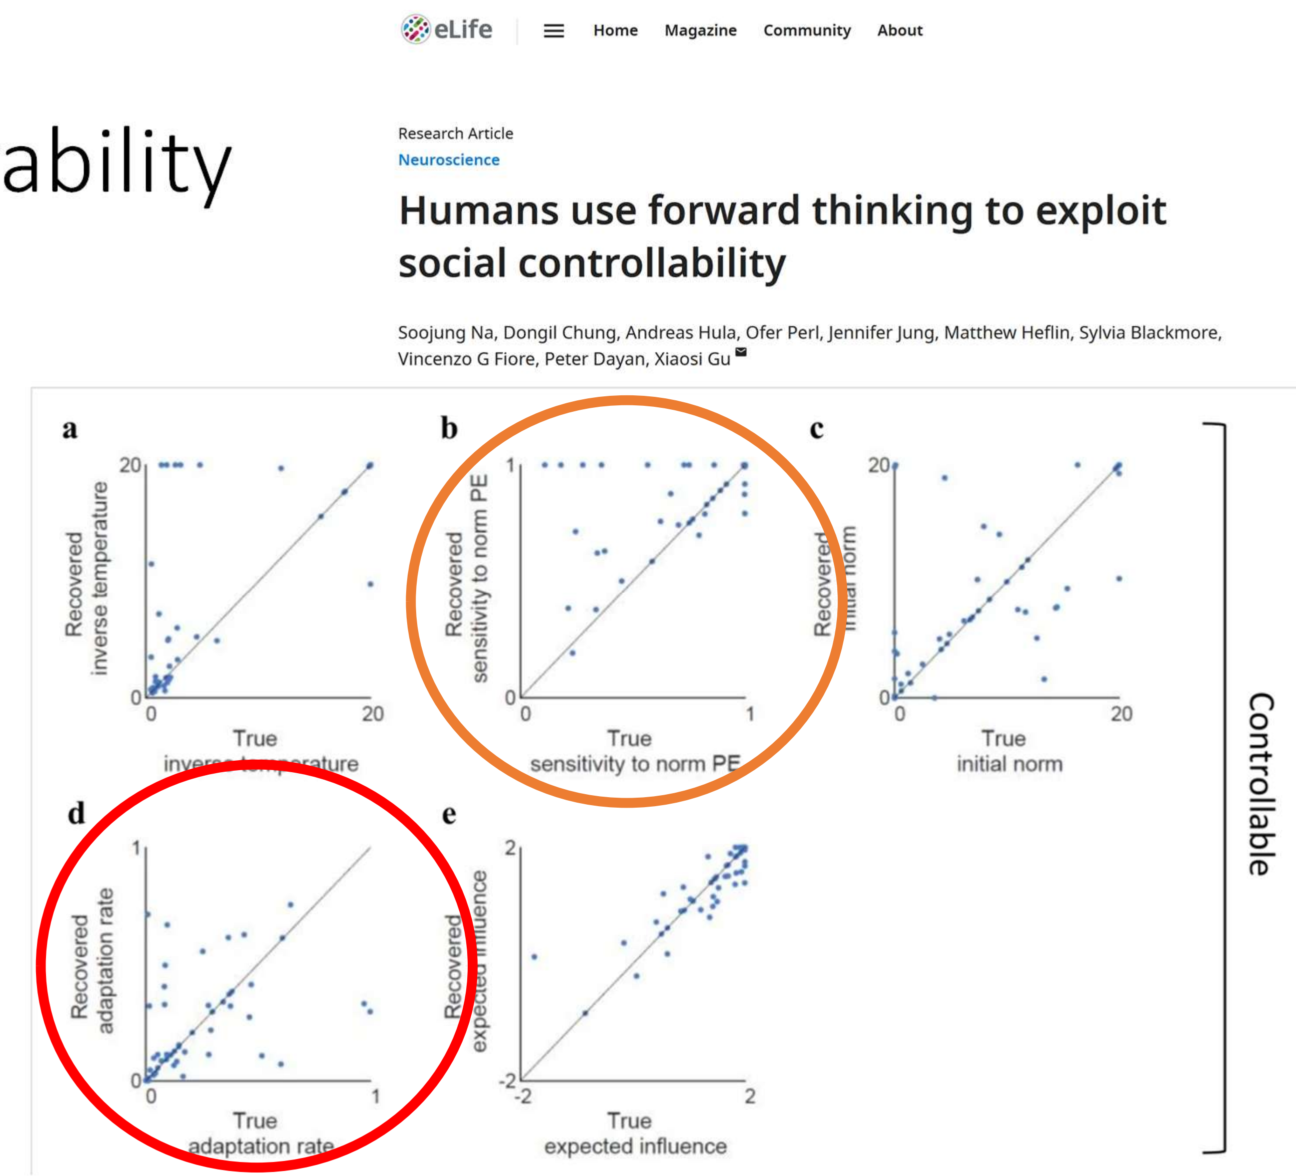

Replication of parameter recoverability

Recoverability changed

- Adaptation rate
  - 0.39-> 0.61
- Delta
  - 0.88 -> 0.91
- Envy
  - 0.57 -> 0.35
- Temp
  - 0.77 -> 0.82
- Initial norm
  - 0.66 -> 0.72

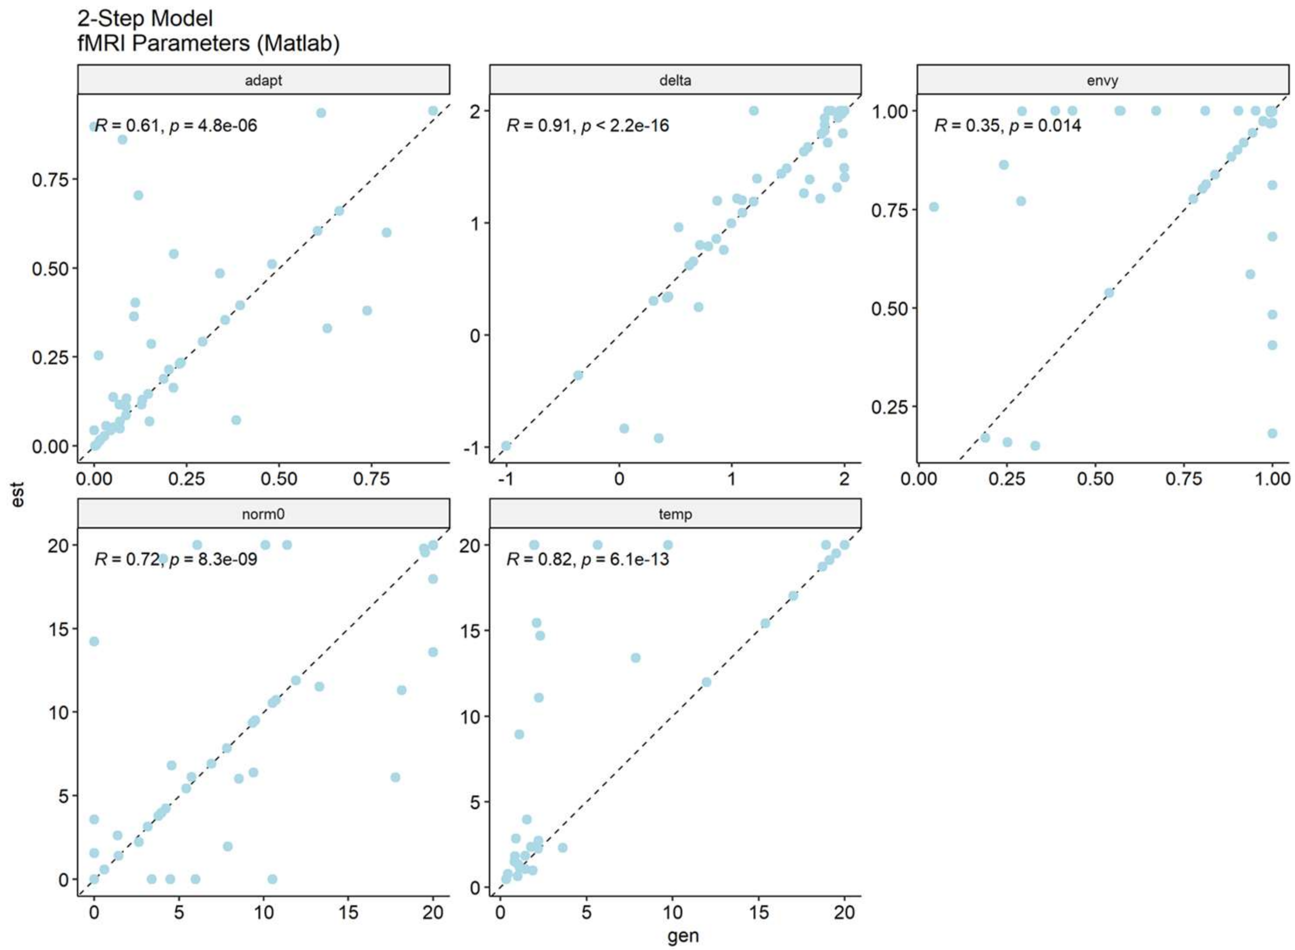

Parameter recoverability using fewer trials

Recoverability got worse!

- Adaptation rate
  - 0.39-> 0.61 -> 0.56
- Delta
  - 0.88 -> 0.91 -> 0.73
- Envy
  - 0.57 -> 0.35 -> 0.28
- Temp
  - 0.77 -> 0.82 -> 0.67
- Initial norm
  - 0.66 -> 0.72 -> 0.69

Original (fMRI) -> Replication (fMRI) -> Online Data

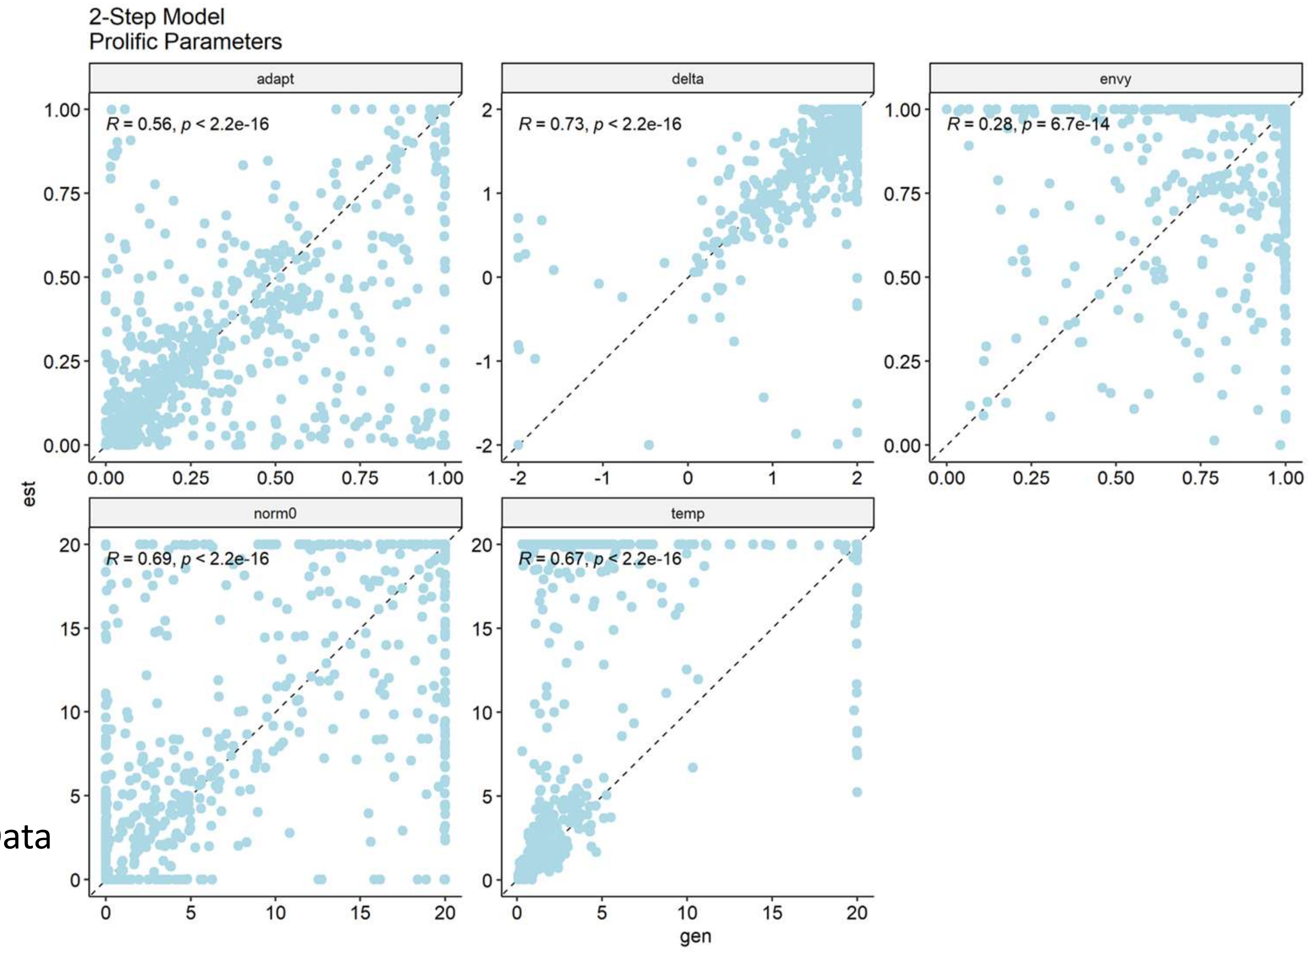

Parameter recoverability using random parameters

Recoverability got *much* worse!

- Adaptation rate
  - 0.39-> 0.61 -> 0.25
- Delta
  - 0.88 -> 0.91 -> 0.90
- Envy
  - 0.57 -> 0.35 -> 0.31
- Temp
  - 0.77 -> 0.82 -> 0.38
- Initial norm
  - 0.66 -> 0.72 -> -0.009

Original (fMRI) -> Replication (fMRI) -> Simulations

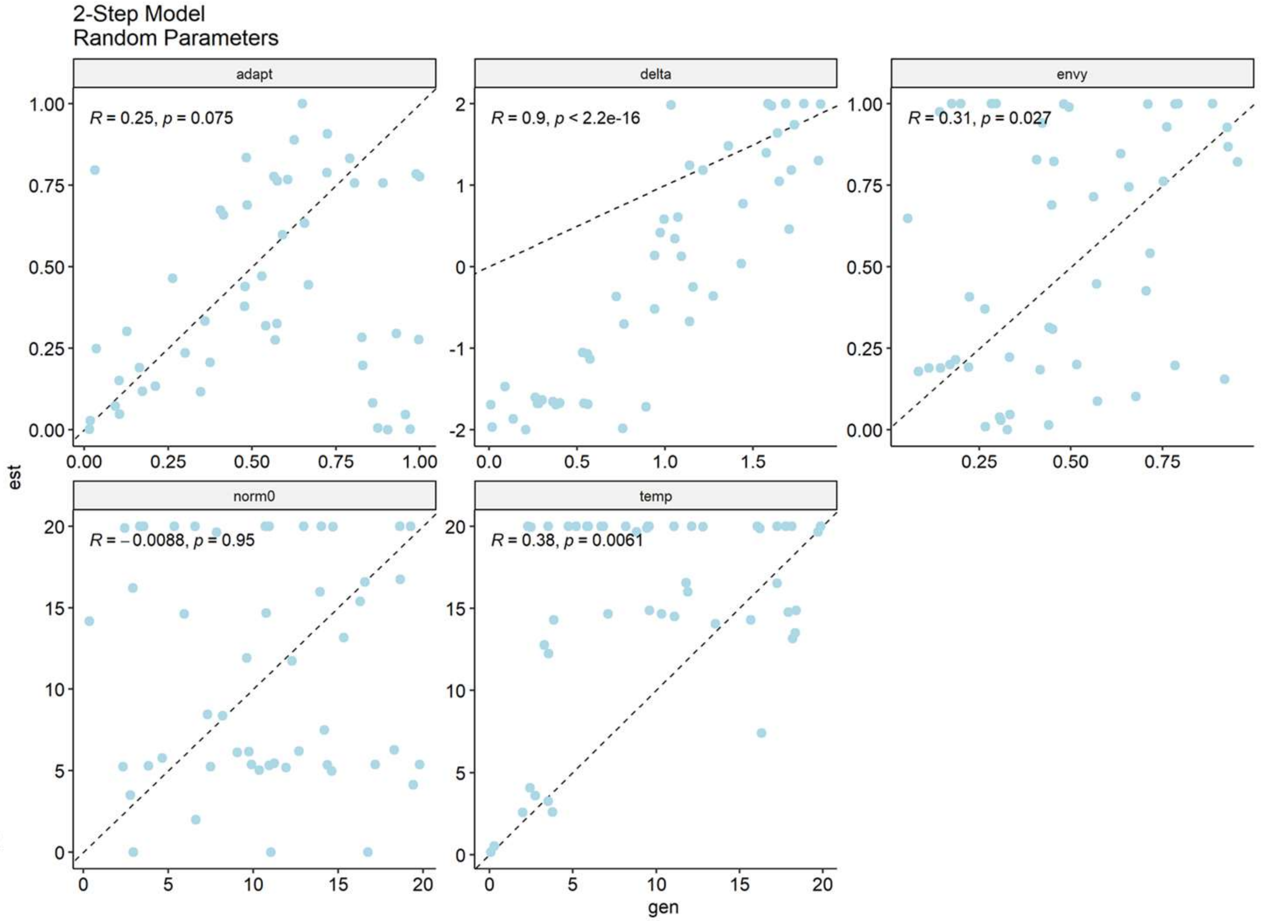

Comparing Programming Languages

- Goal: Compare parameter estimation and recoverability using Matlab and R
  - Matlab: Original code
  - R: More flexible
- Dataset
  - Round 1 of Prolific UG2 data (*N* =711)
  - Controllability condition
  - 24 trials (out of 30)

R > Matlab ????

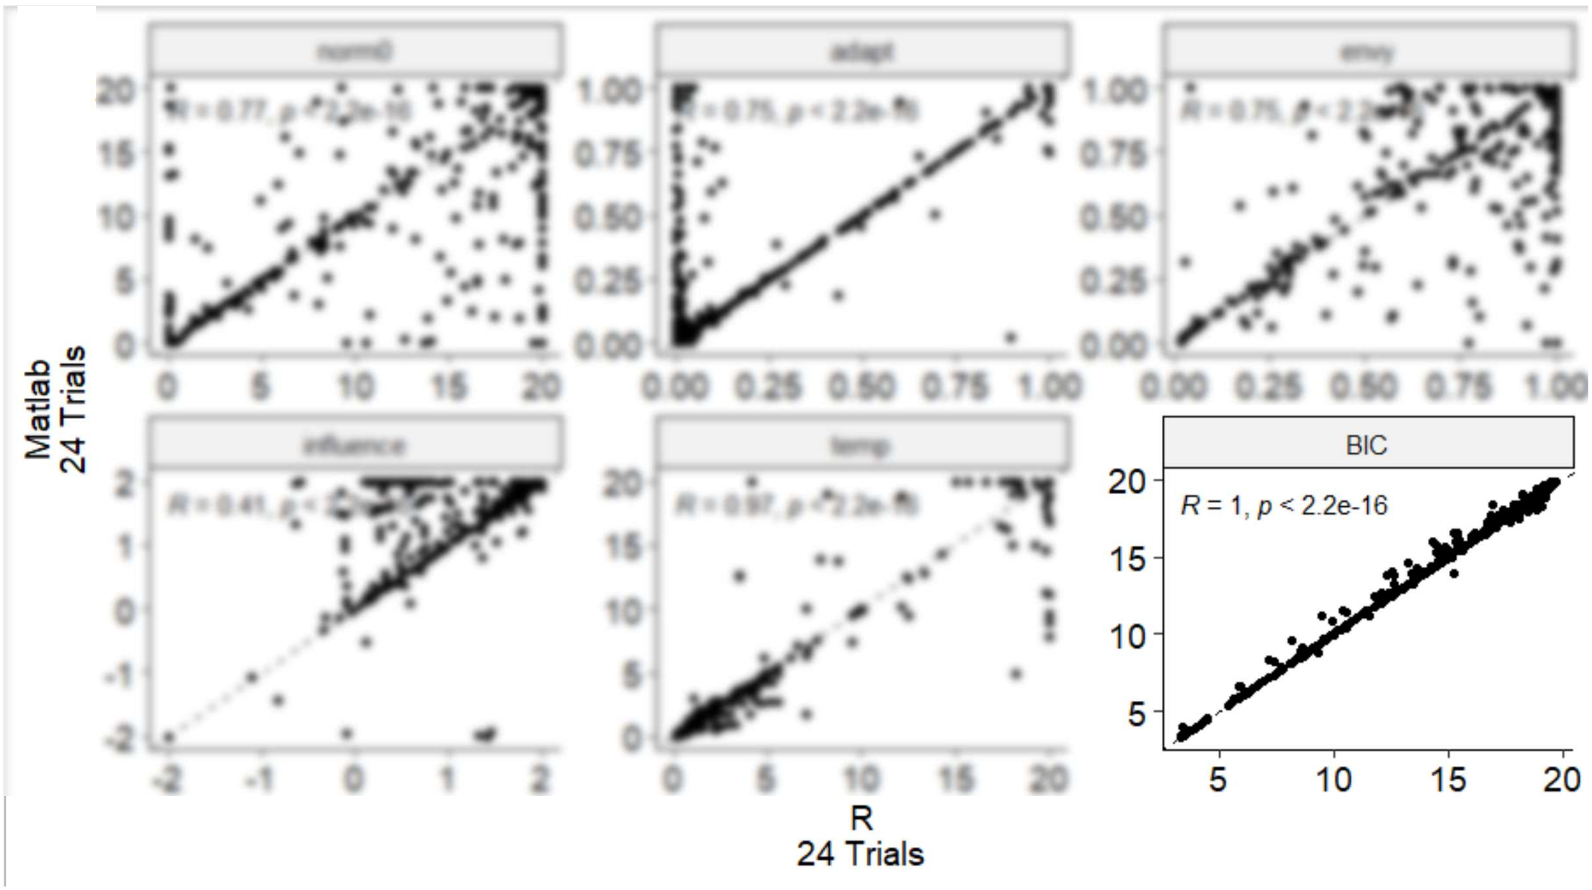

R > Matlab ????

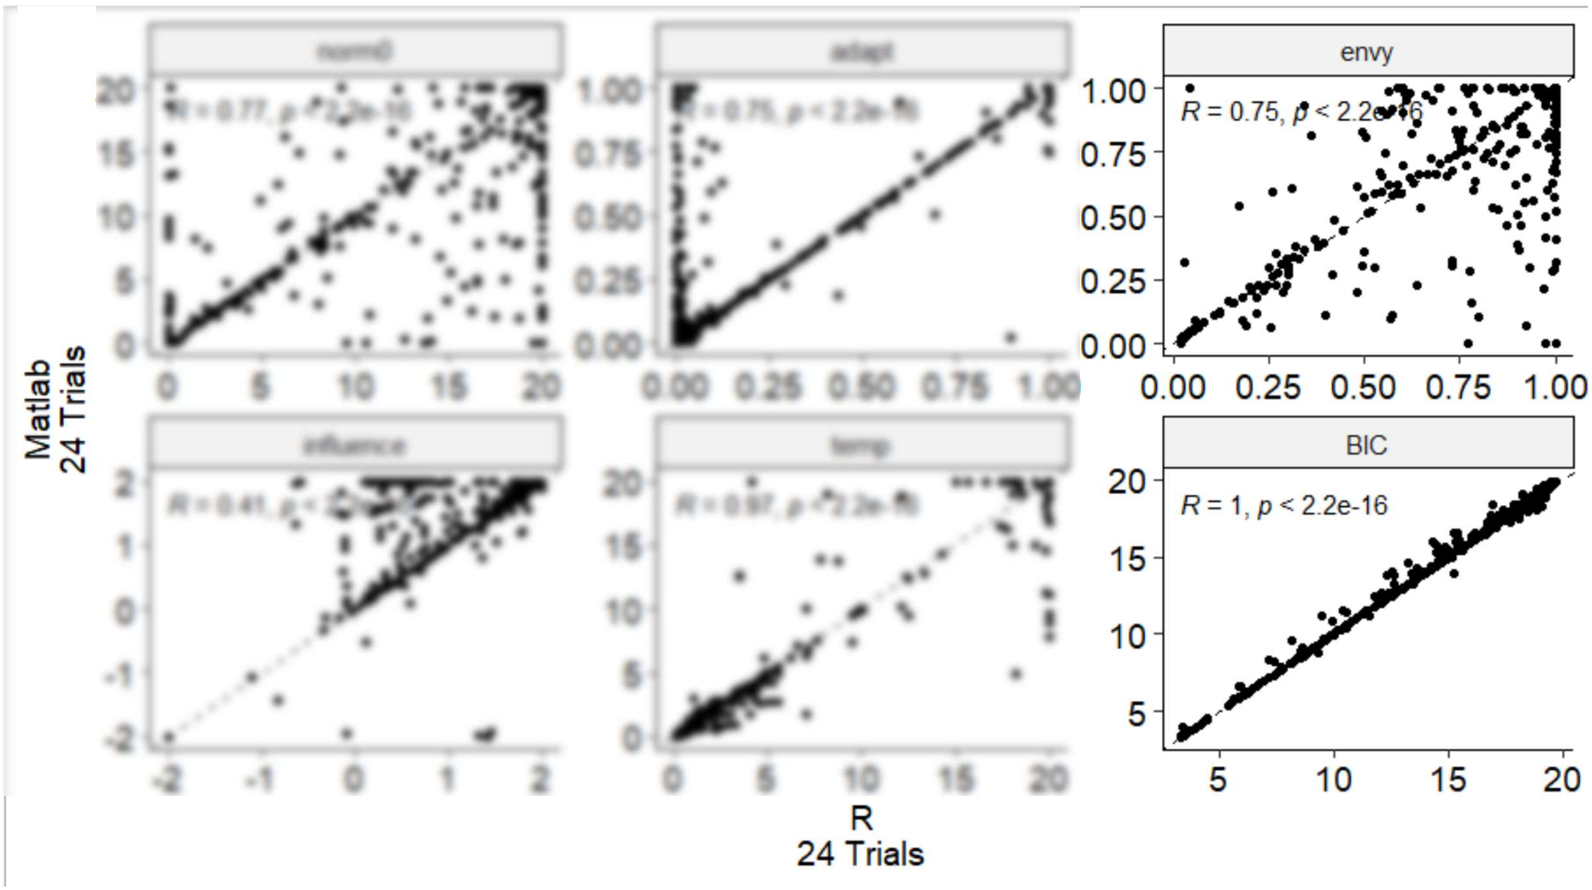

envy
R = 0.75, p < 2.2e-16
1.00
0.75
0.50
0.25
0.00
0.00 0.25 0.50 0.75 1.00
BIC
R = 1, p < 2.2e-16
20
15
10
5
5 10 15 20
Matlab
24 Trials
R
24 Trials

R > Matlab ????

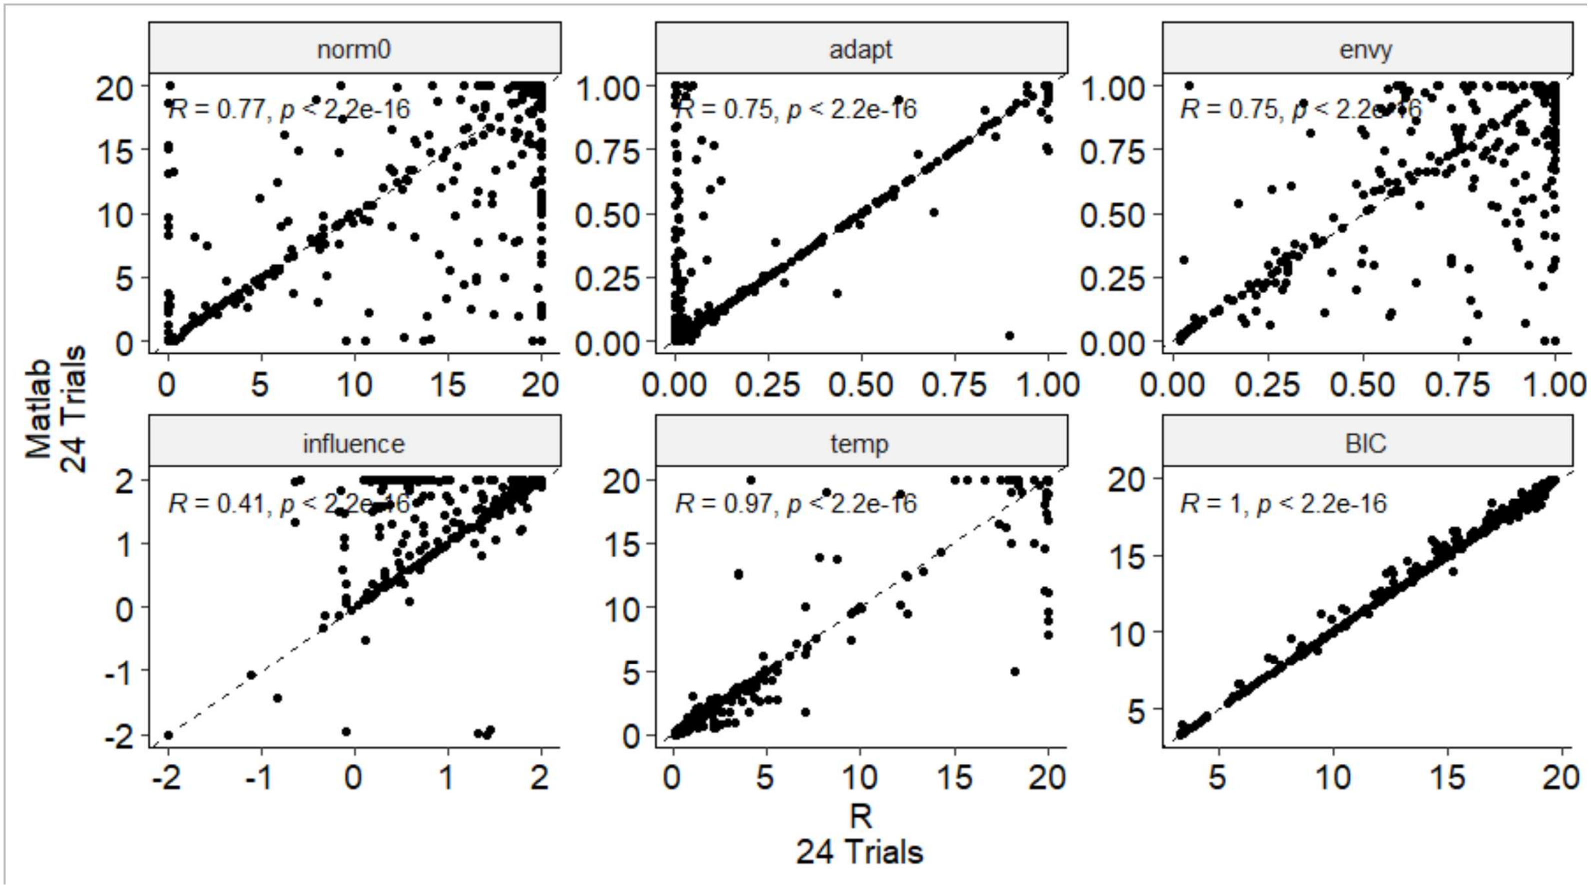

norm0
adapt
envy
influence
temp
BIC
R = 0.77, p < 2.2e-16
R = 0.75, p < 2.2e-16
R = 0.75, p < 2.2e-16
R = 0.41, p < 2.2e-16
R = 0.97, p < 2.2e-16
R = 1, p < 2.2e-16
Matlab
24 Trials
R
24 Trials

Sensitivity to starting values

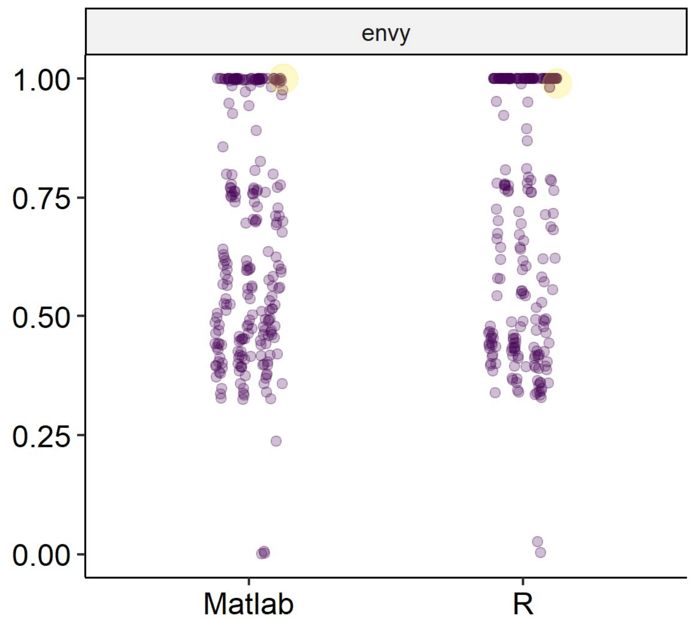

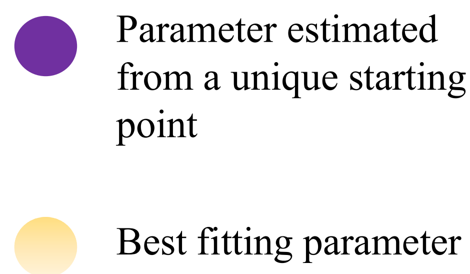

Sensitivity to starting values

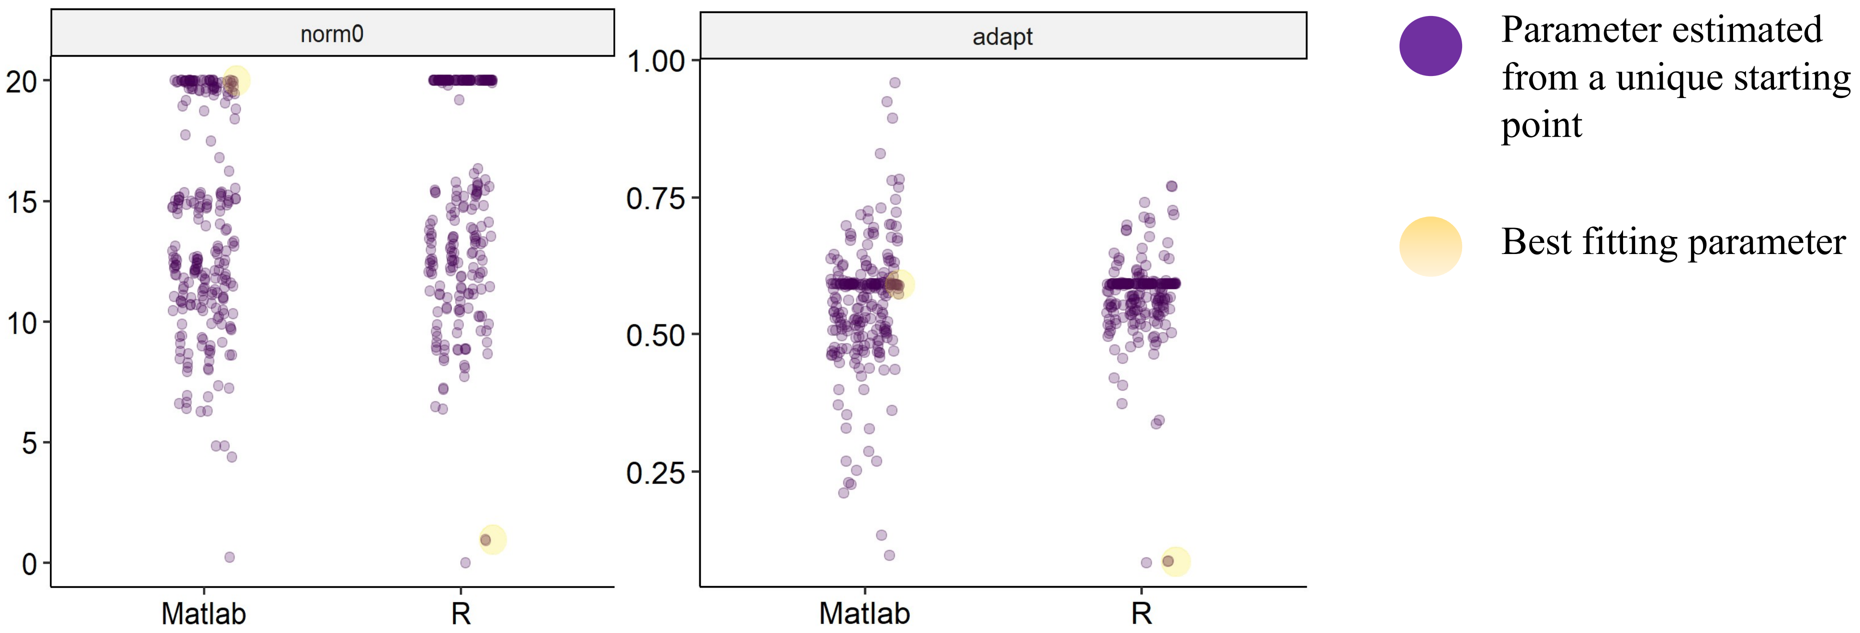

Trial Number

- In original paper, only used 20/30 possible trials
  - Exclude early, exploratory behavior (i.e., trials 1-5)
  - Avoid late, “cashing out” behavior (i.e., trials 25-30)
- In recent projects: 24 trials to increase information
  - Exclude first and last 3 trials
- What happens when we change the number of trials?

Parameter correlations

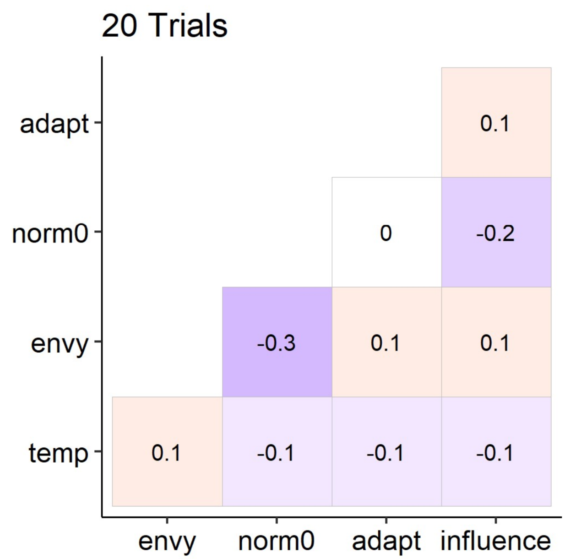

Parameter correlations

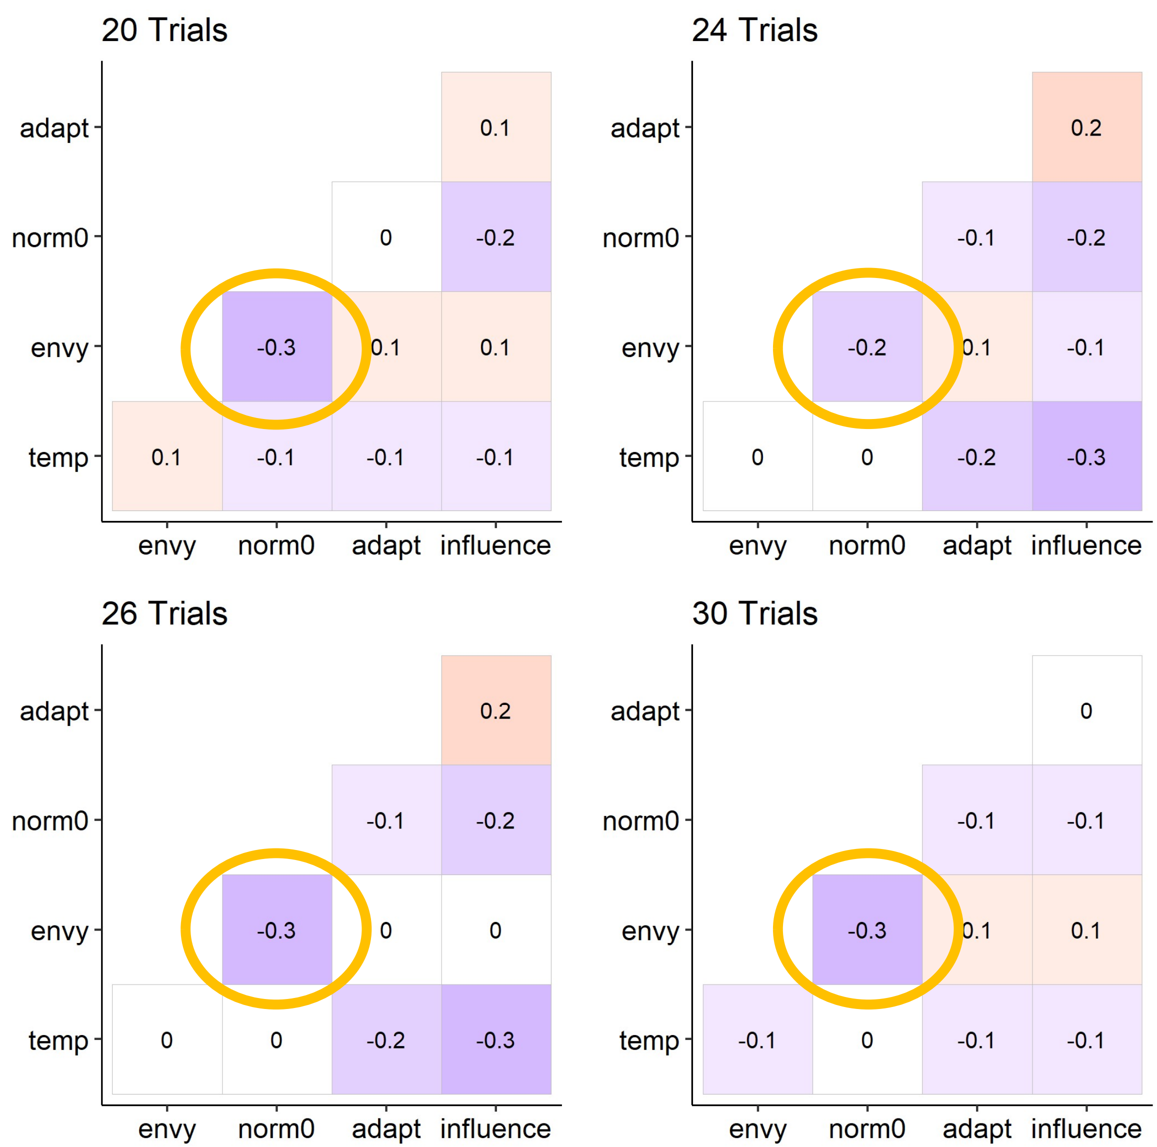

Comparing trial numbers 20 versus 24

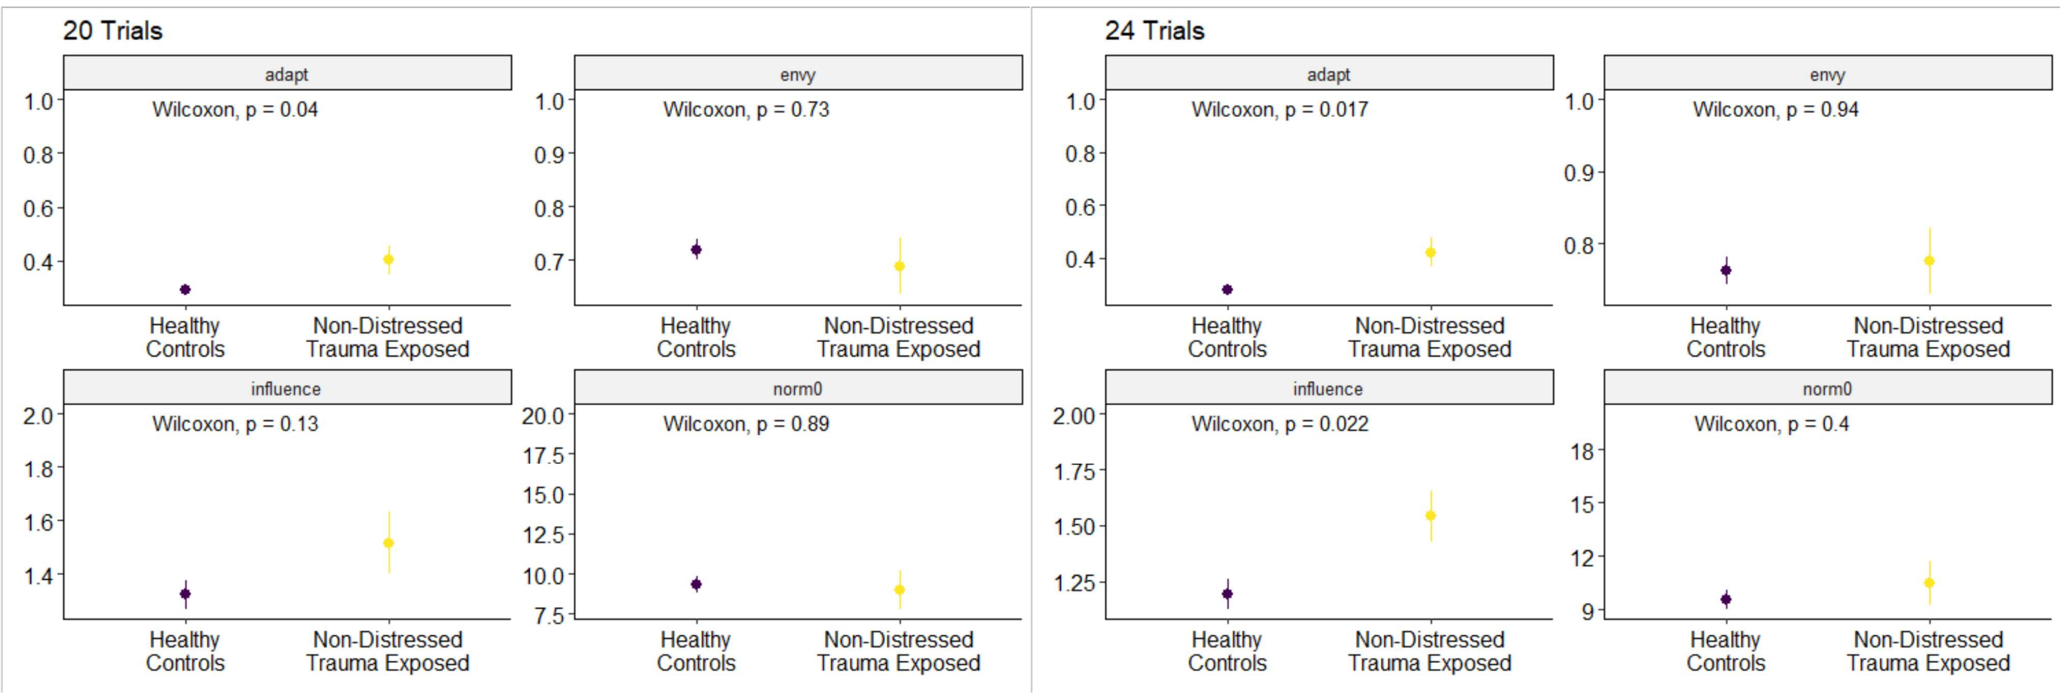

Comparing trial numbers 26 versus 30

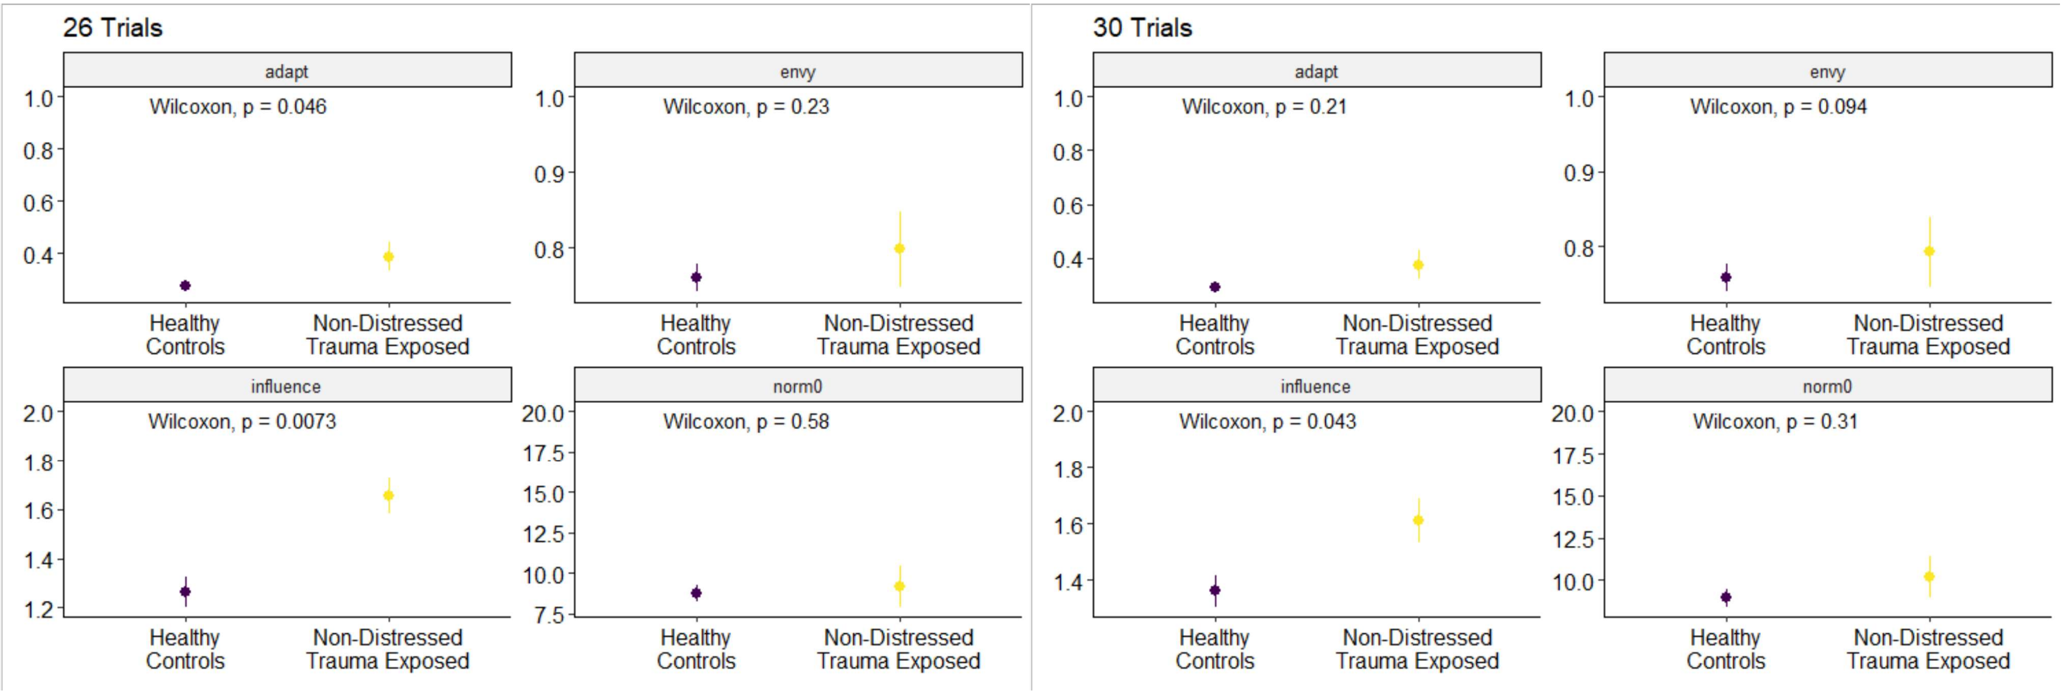

What are parameters good for?

- Parameters should reflect distinct behavioral patterns
- Simulations can illustrate parameter utility
- Exercise
  - Create baseline data
    - Simulate 1000 subjects
    - Each parameter based on empirical distributions
  - Generate manipulated data
    - Same baseline parameters
    - Iteratively change one parameter at a time
  - Plot behavioral predictions

What are parameters good for?

- Parameters should reflect distinct behavioral patterns
- Simulations can illustrate parameter utility
- Exercise
  - Create baseline data
    - Simulate 1000 subjects
    - Each parameter based on empirical distributions
  - Generate manipulated data
    - Same baseline parameters
    - Iteratively change one parameter at a time
  - Plot behavioral predictions
    - Offers over time
    - Rejections by offer amount

Baseline Simulations

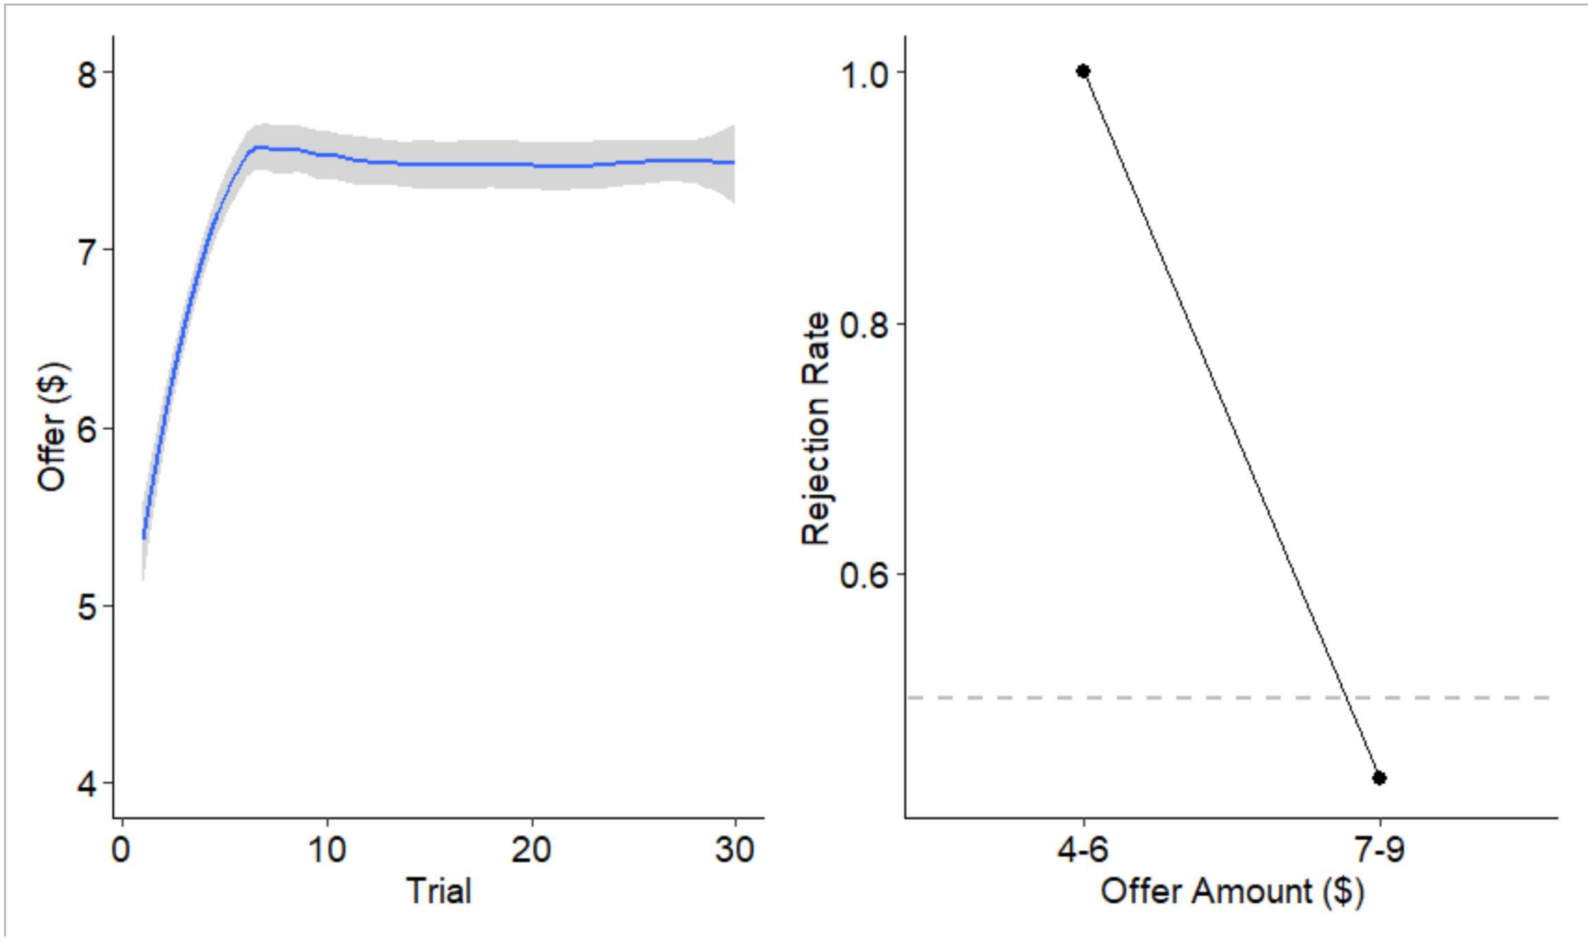

Varying Beta (Inverse Temp.)

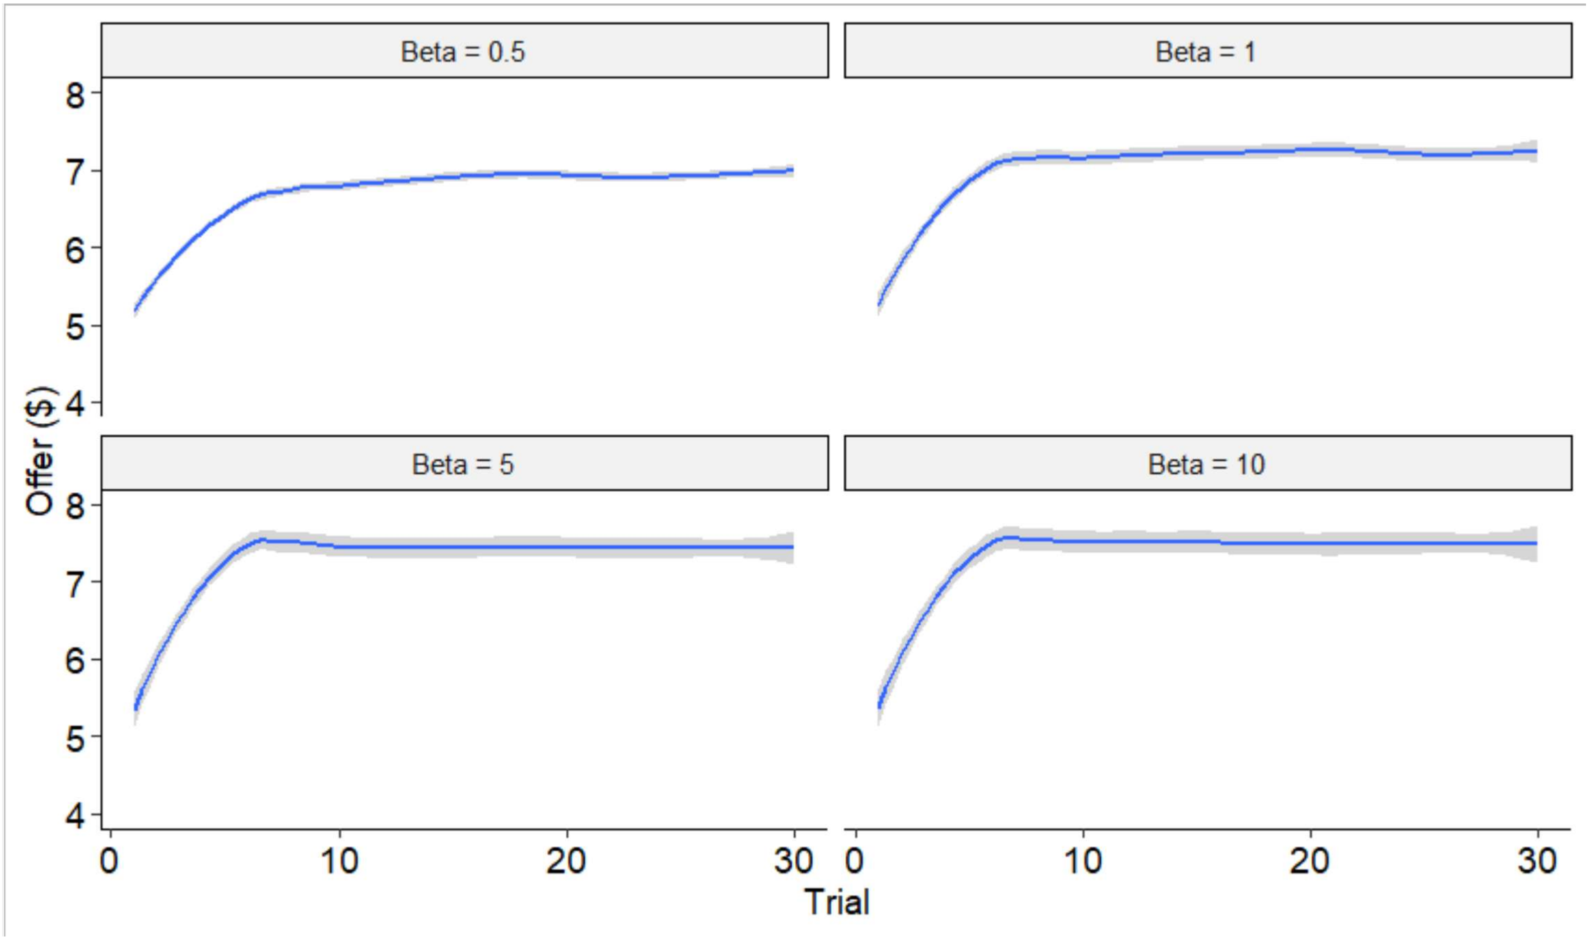

Varying Beta (Inverse Temp.)

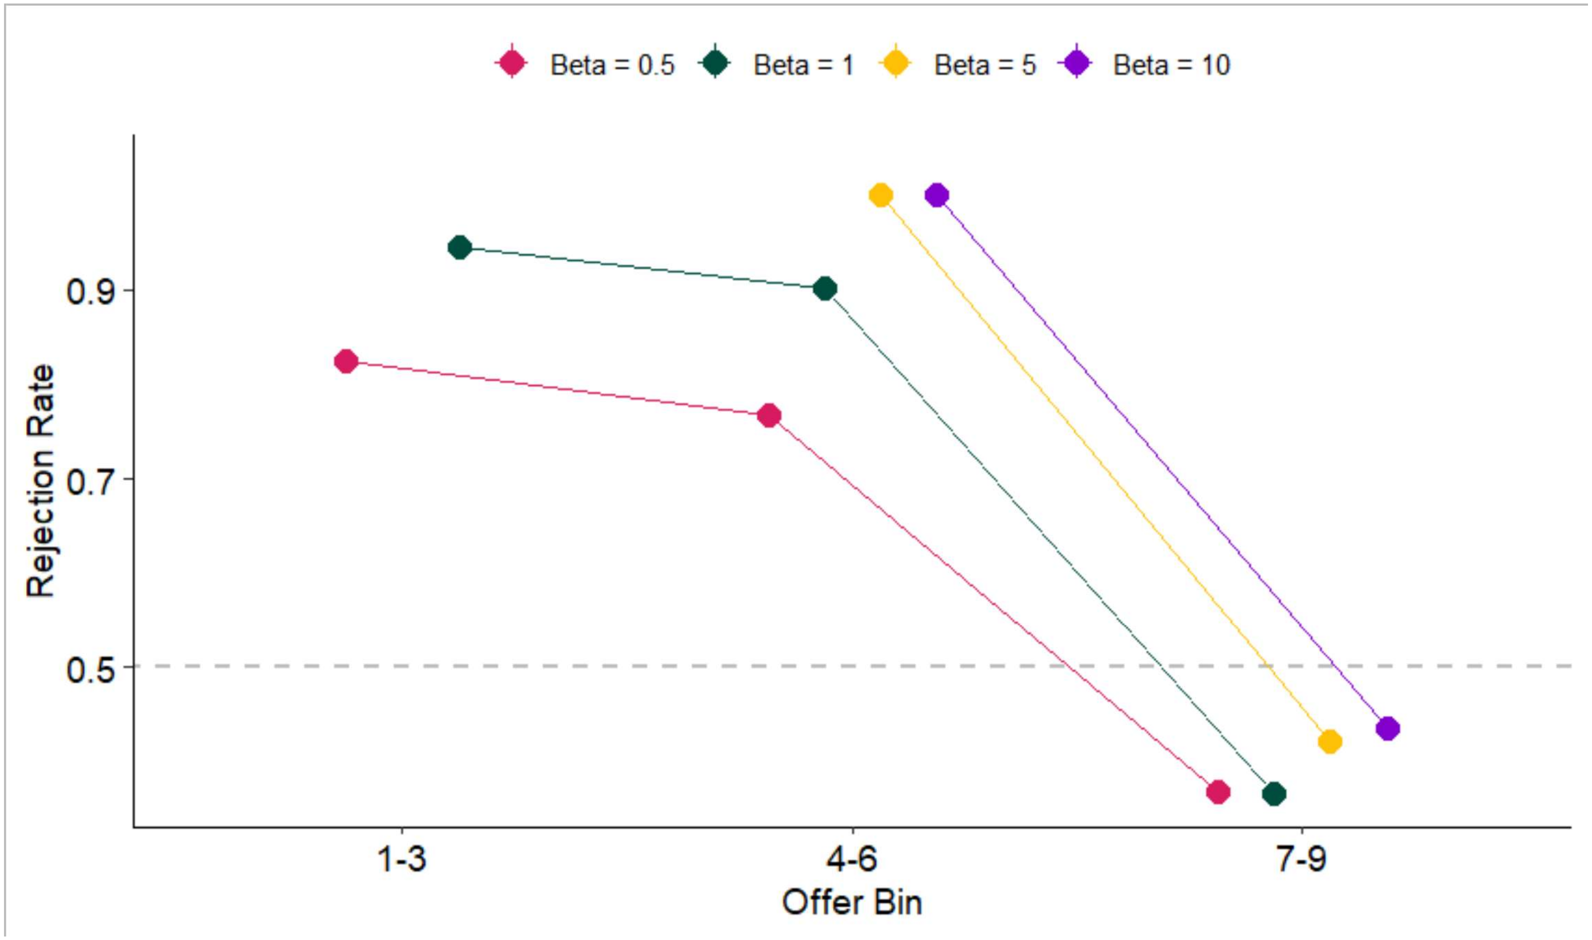

Varying Alpha (Envy)

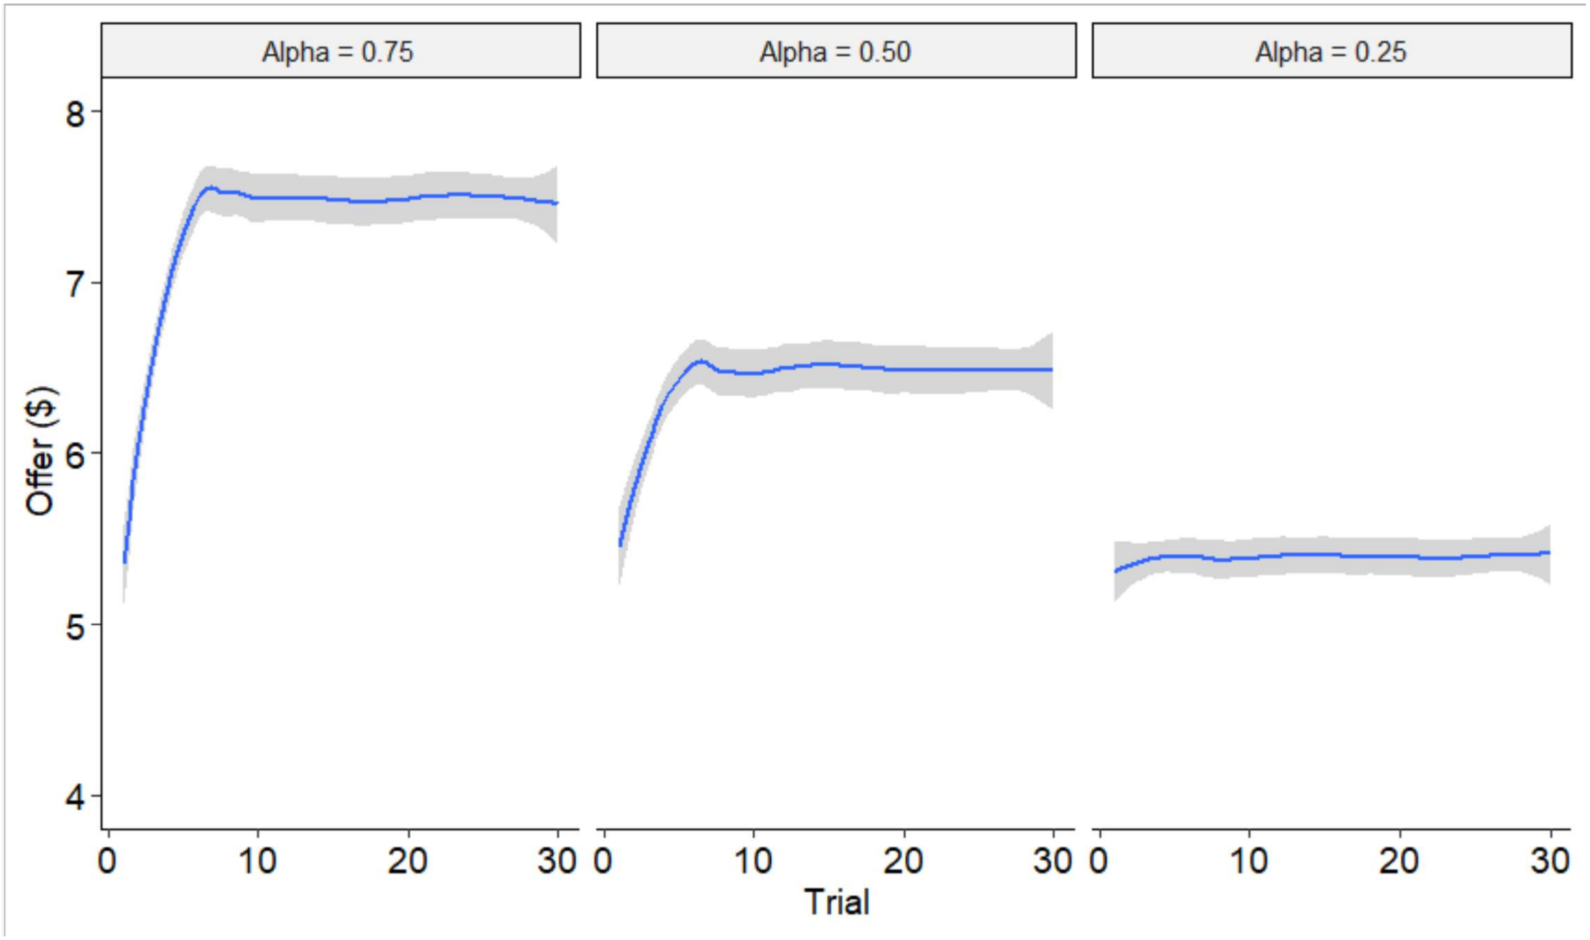

Varying Alpha (Envy)

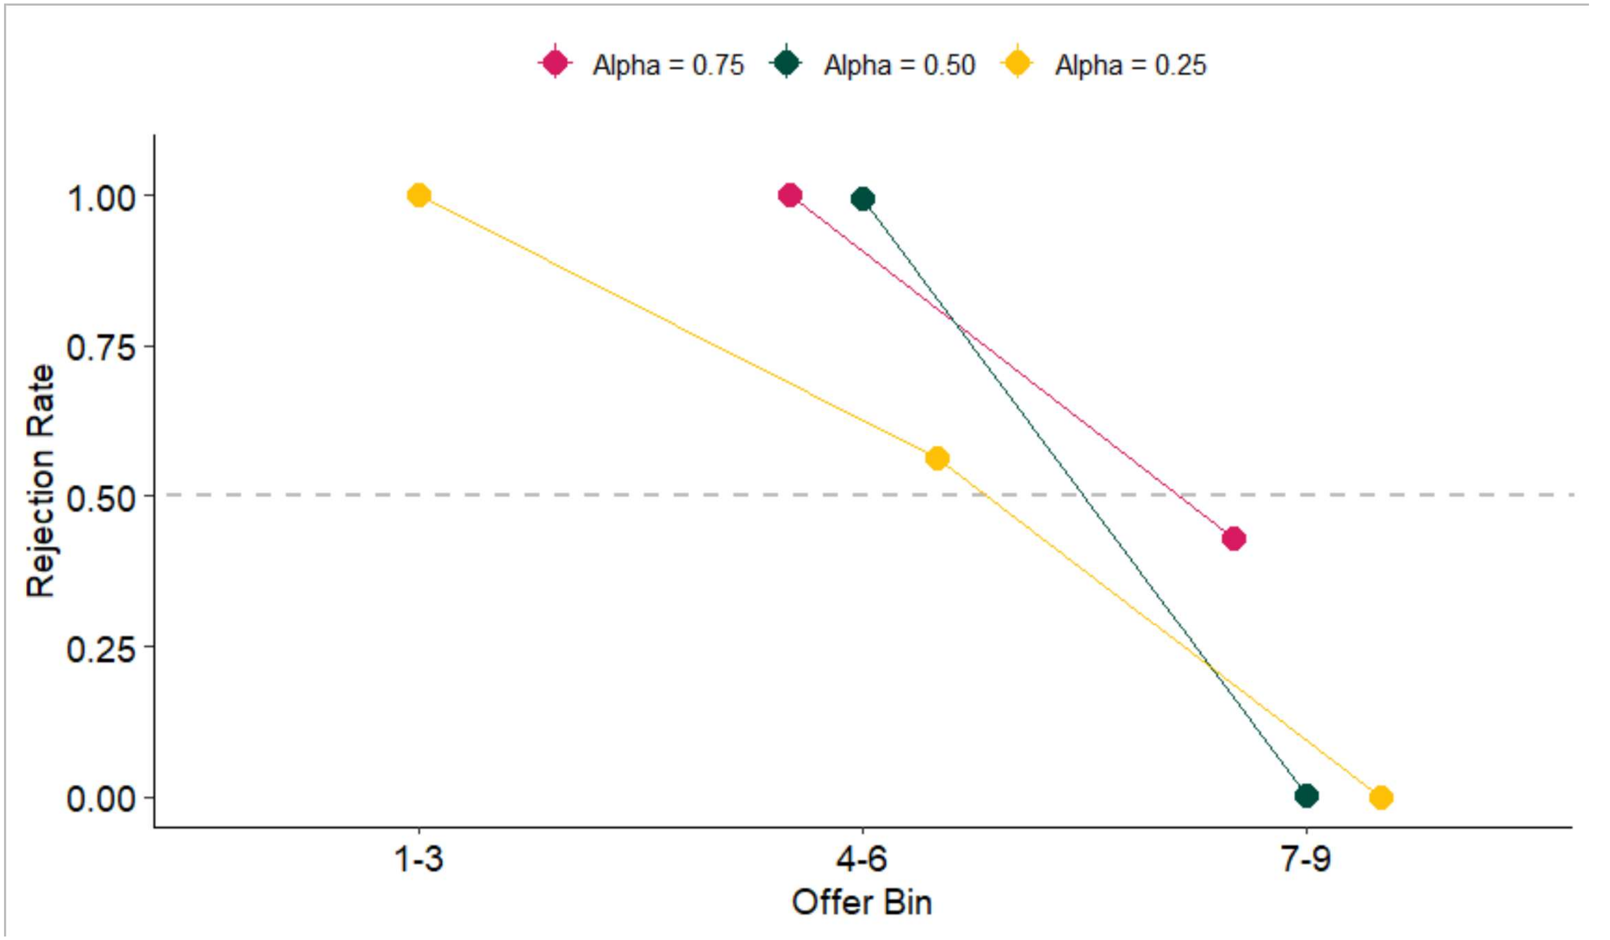

Varying f0 (Initial norm)

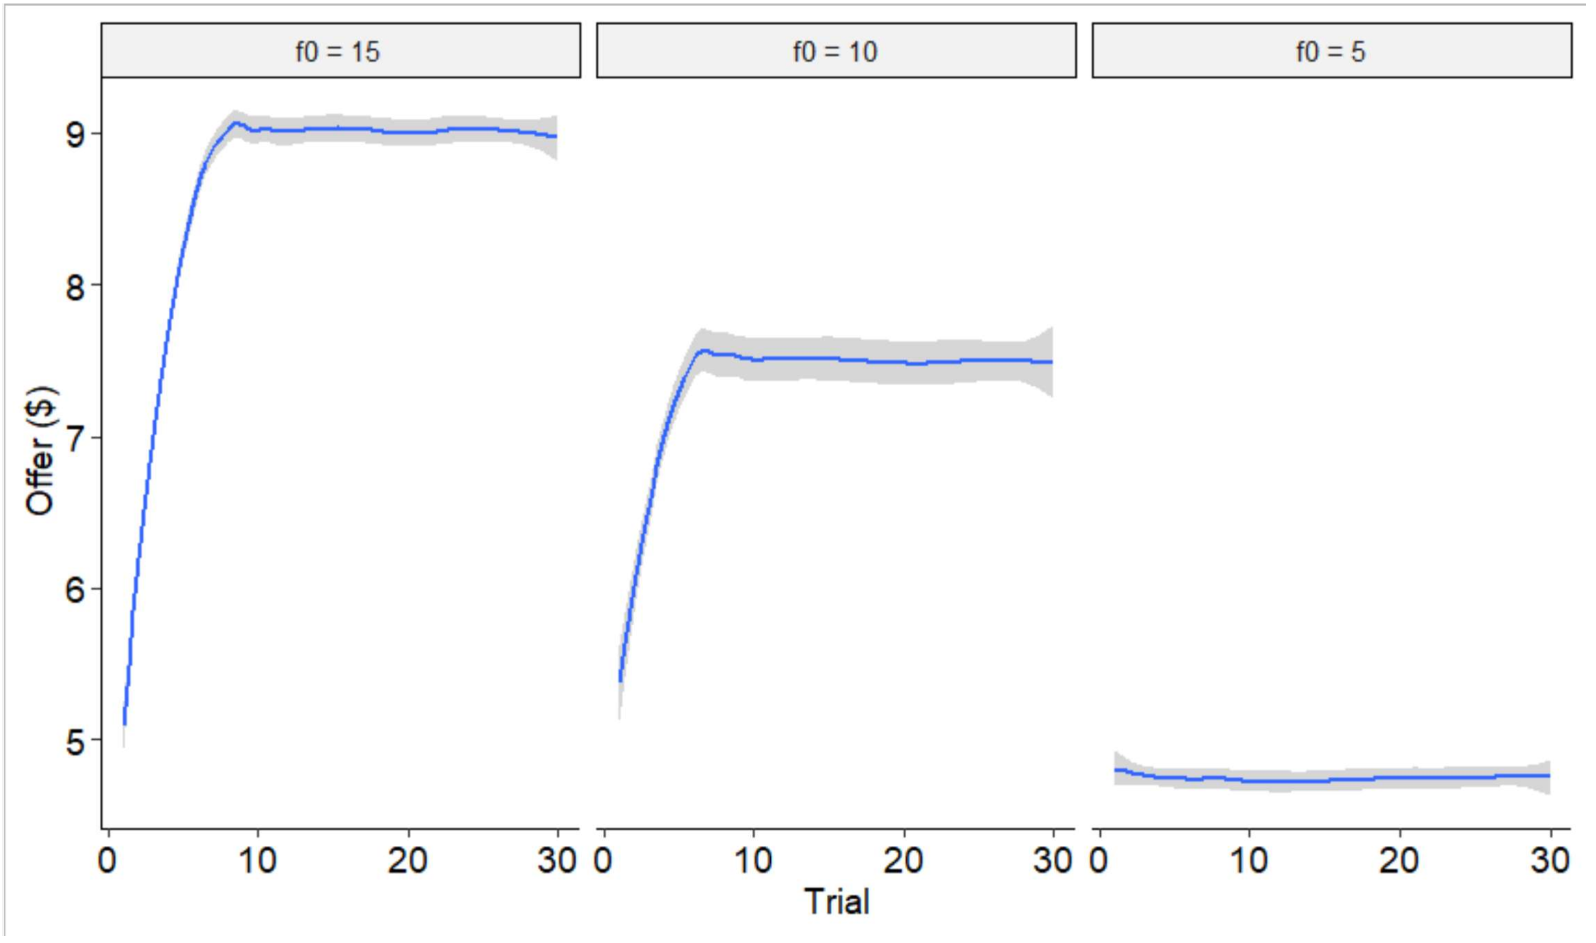

Varying f0 (Initial norm)

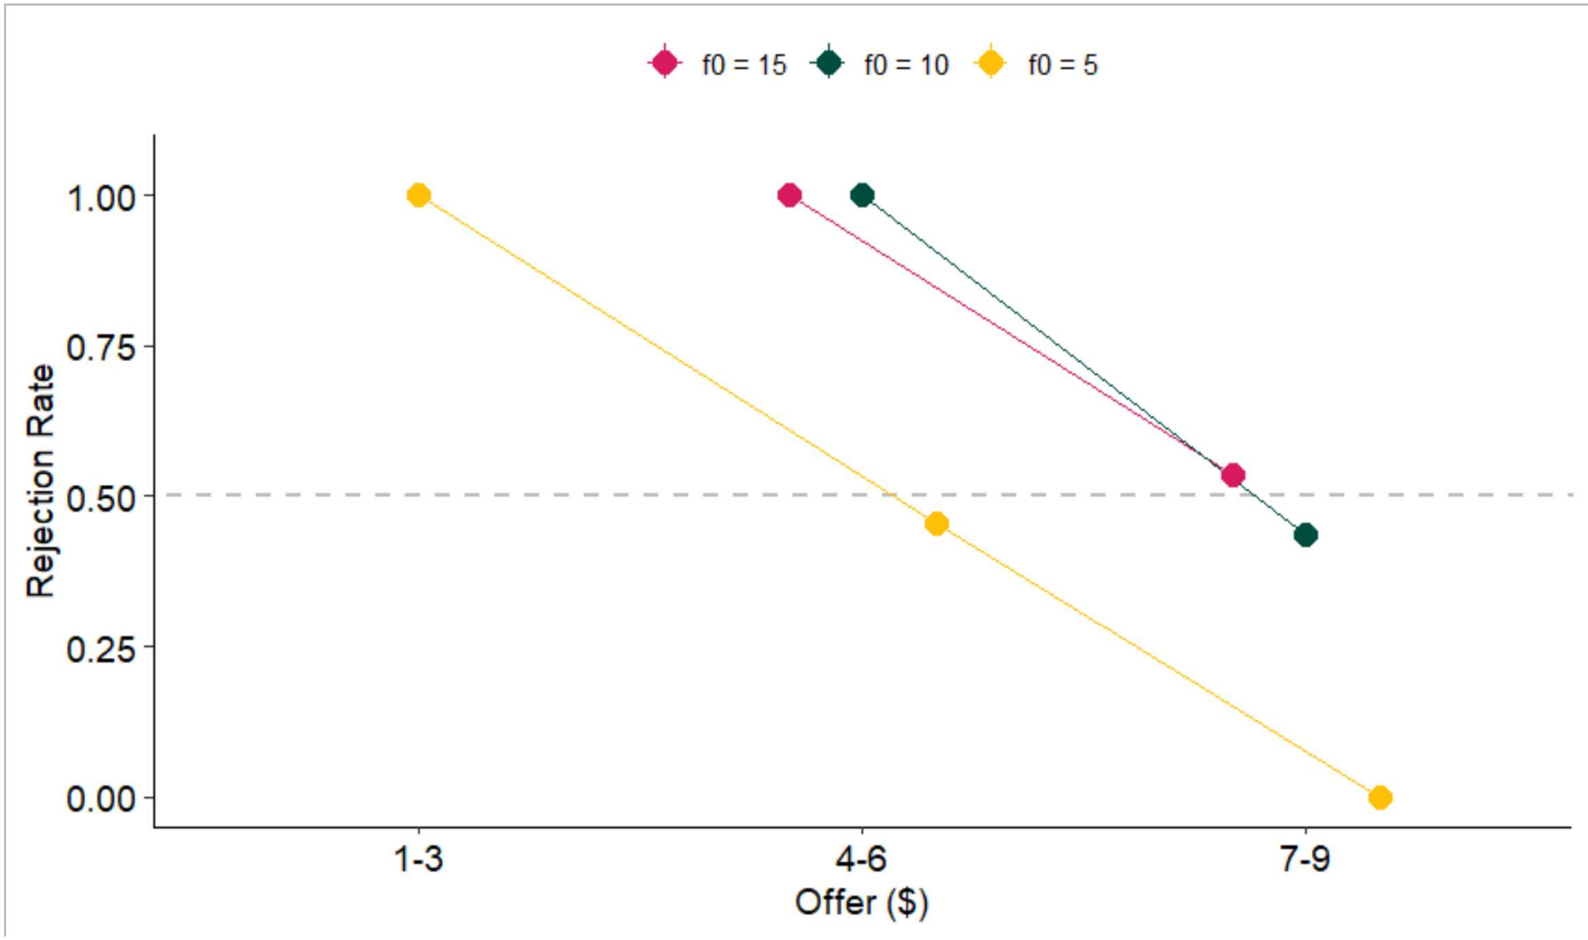

Varying epsilon (adaptation rate)

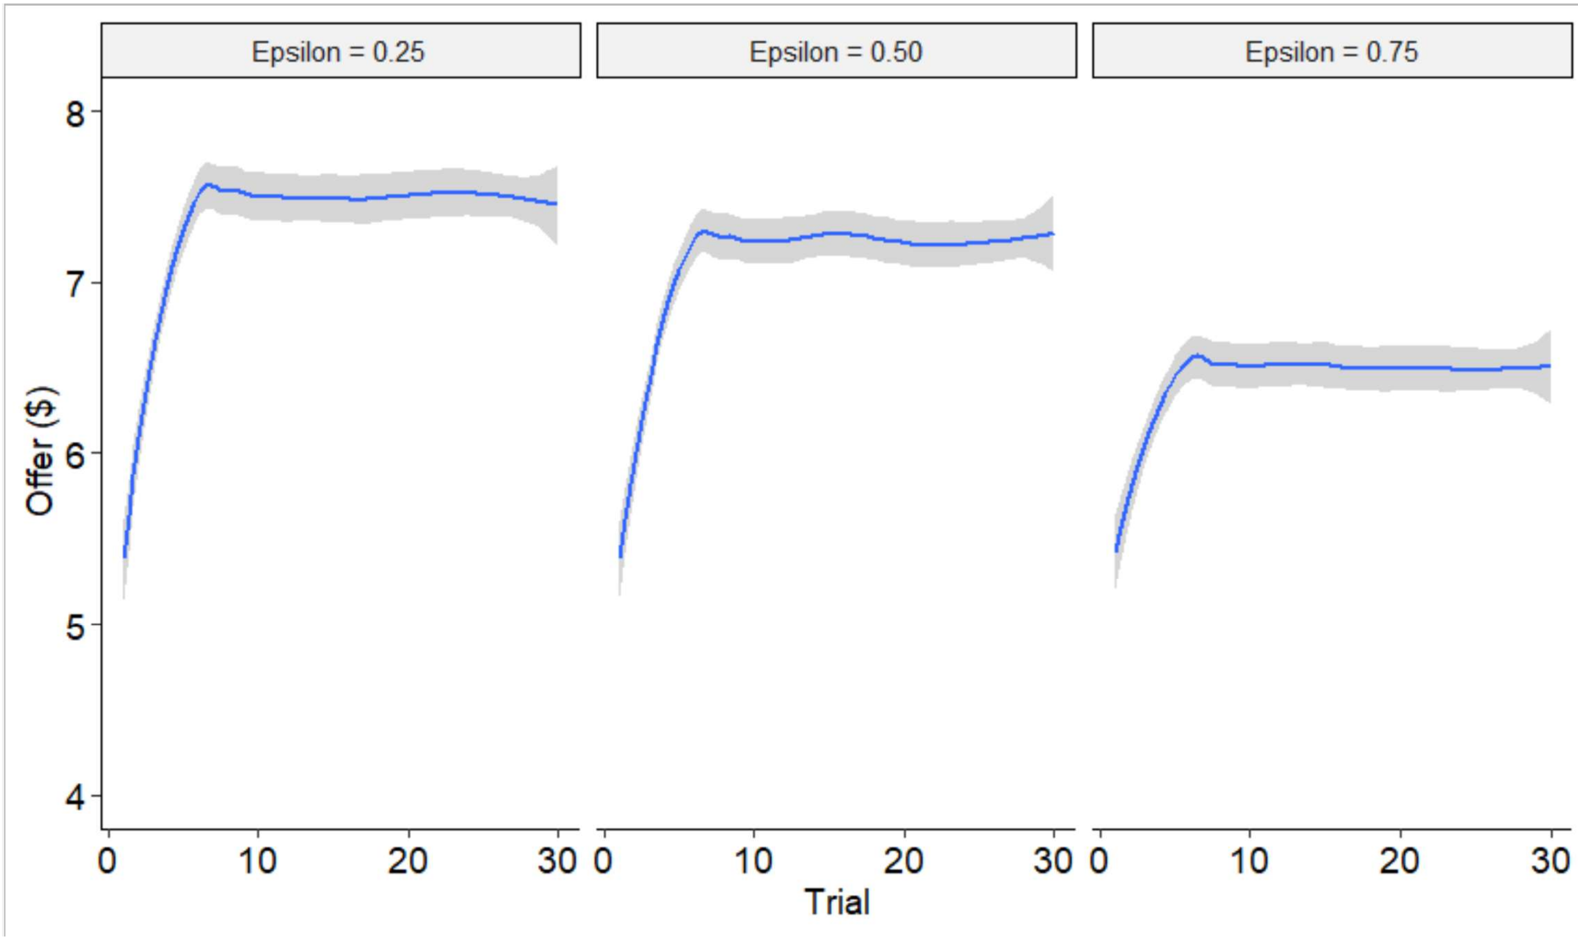

Varying epsilon (adaptation rate)

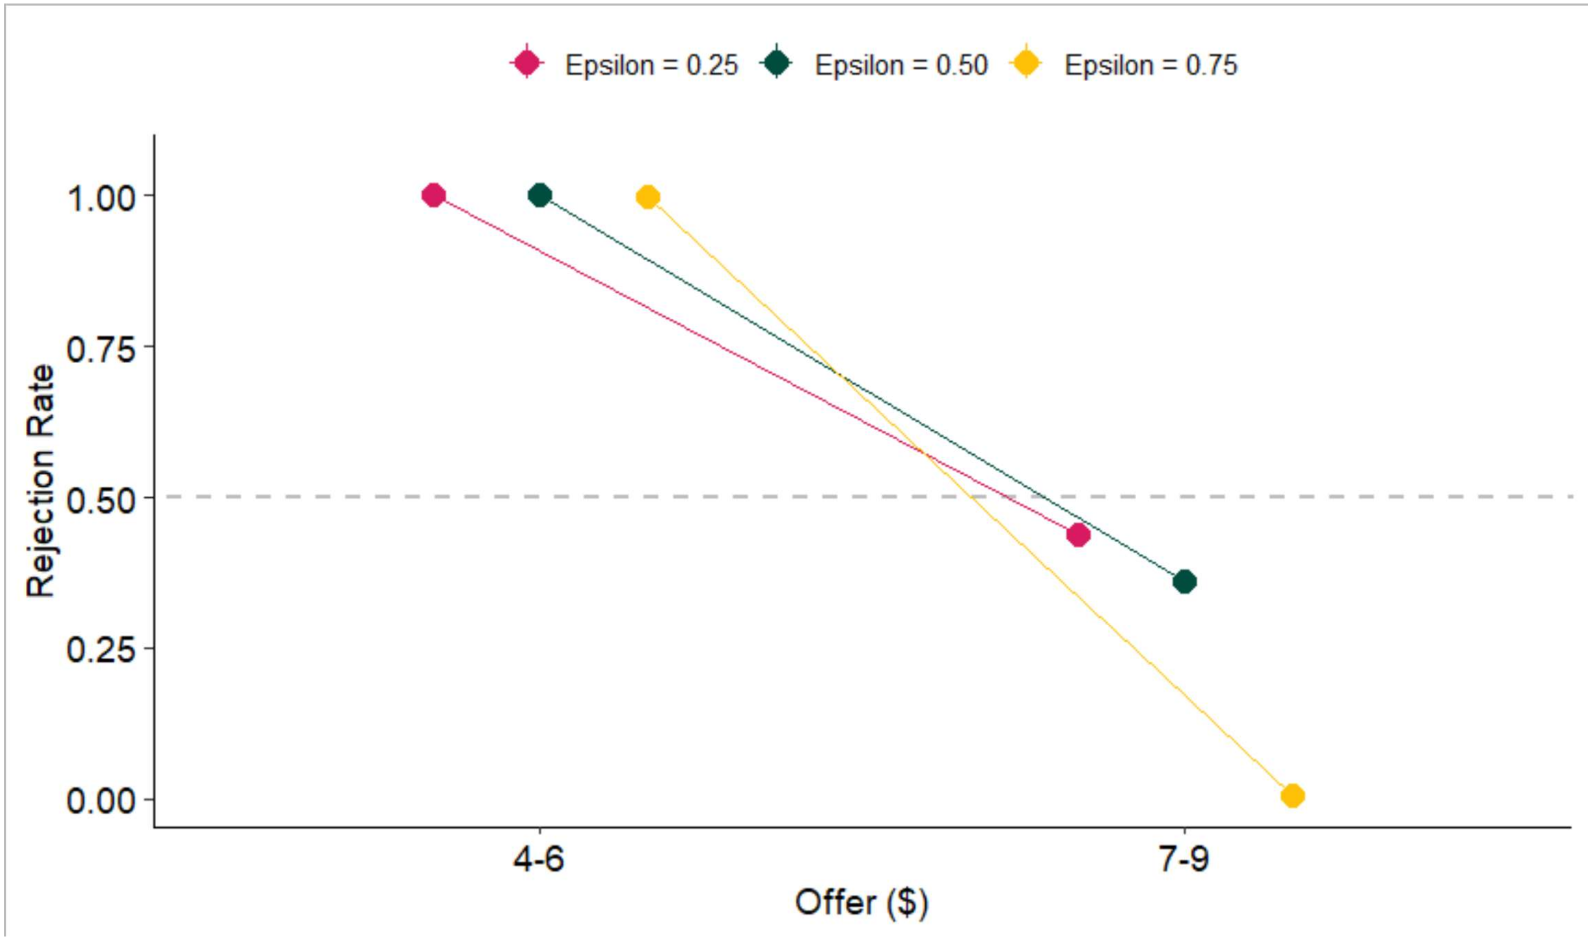

Varying delta (influence)

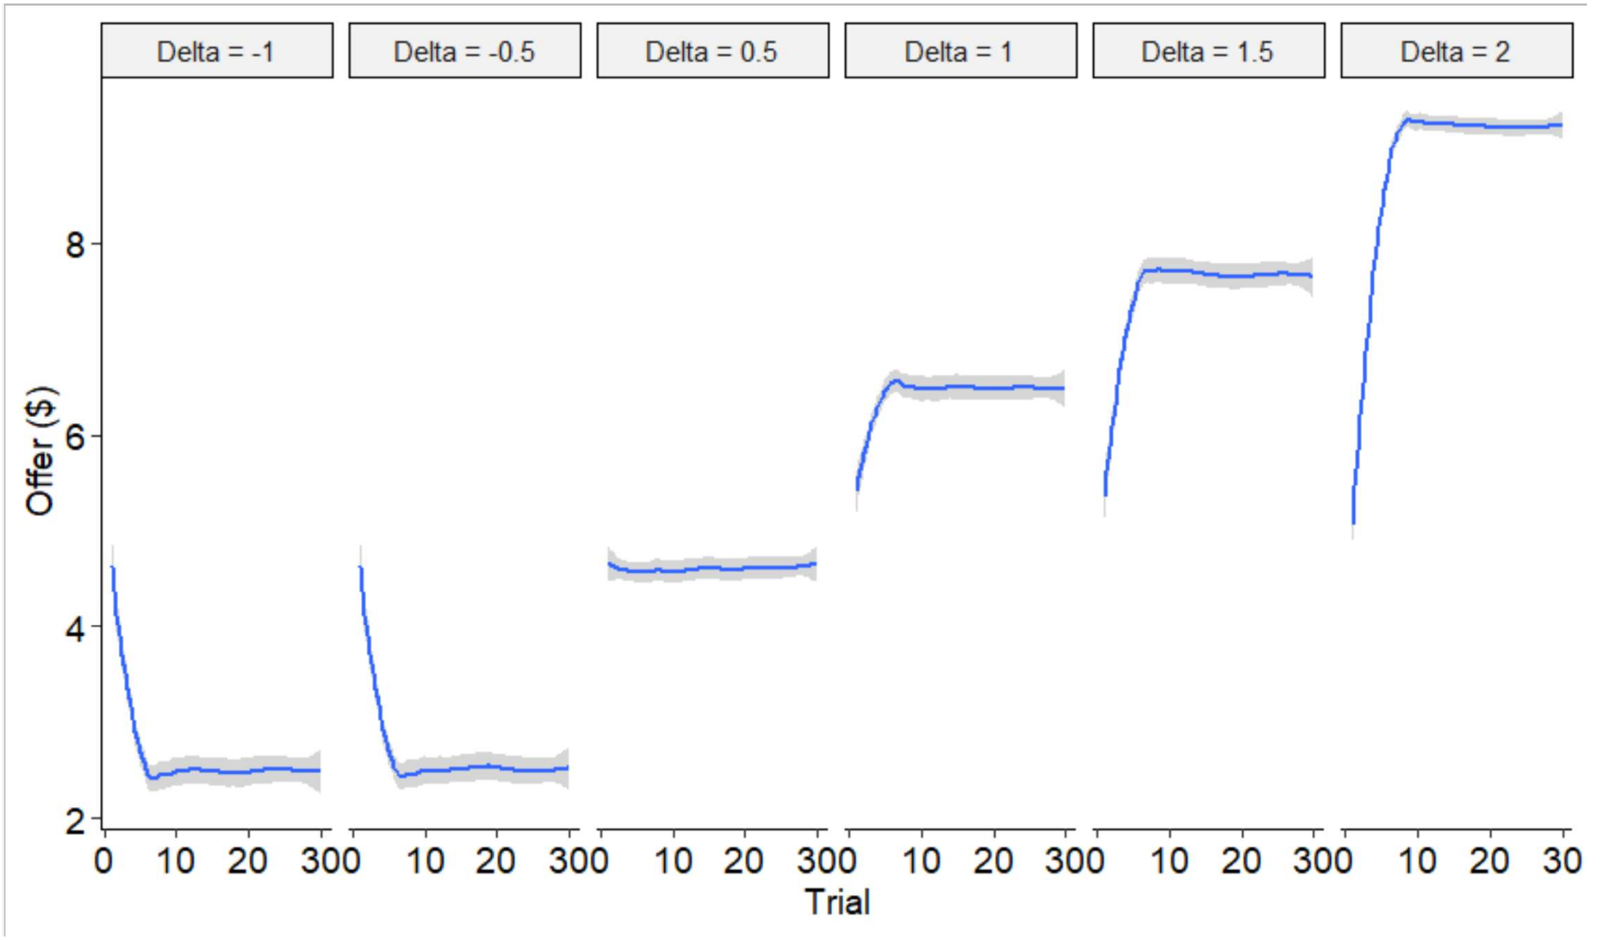

Varying delta (influence)

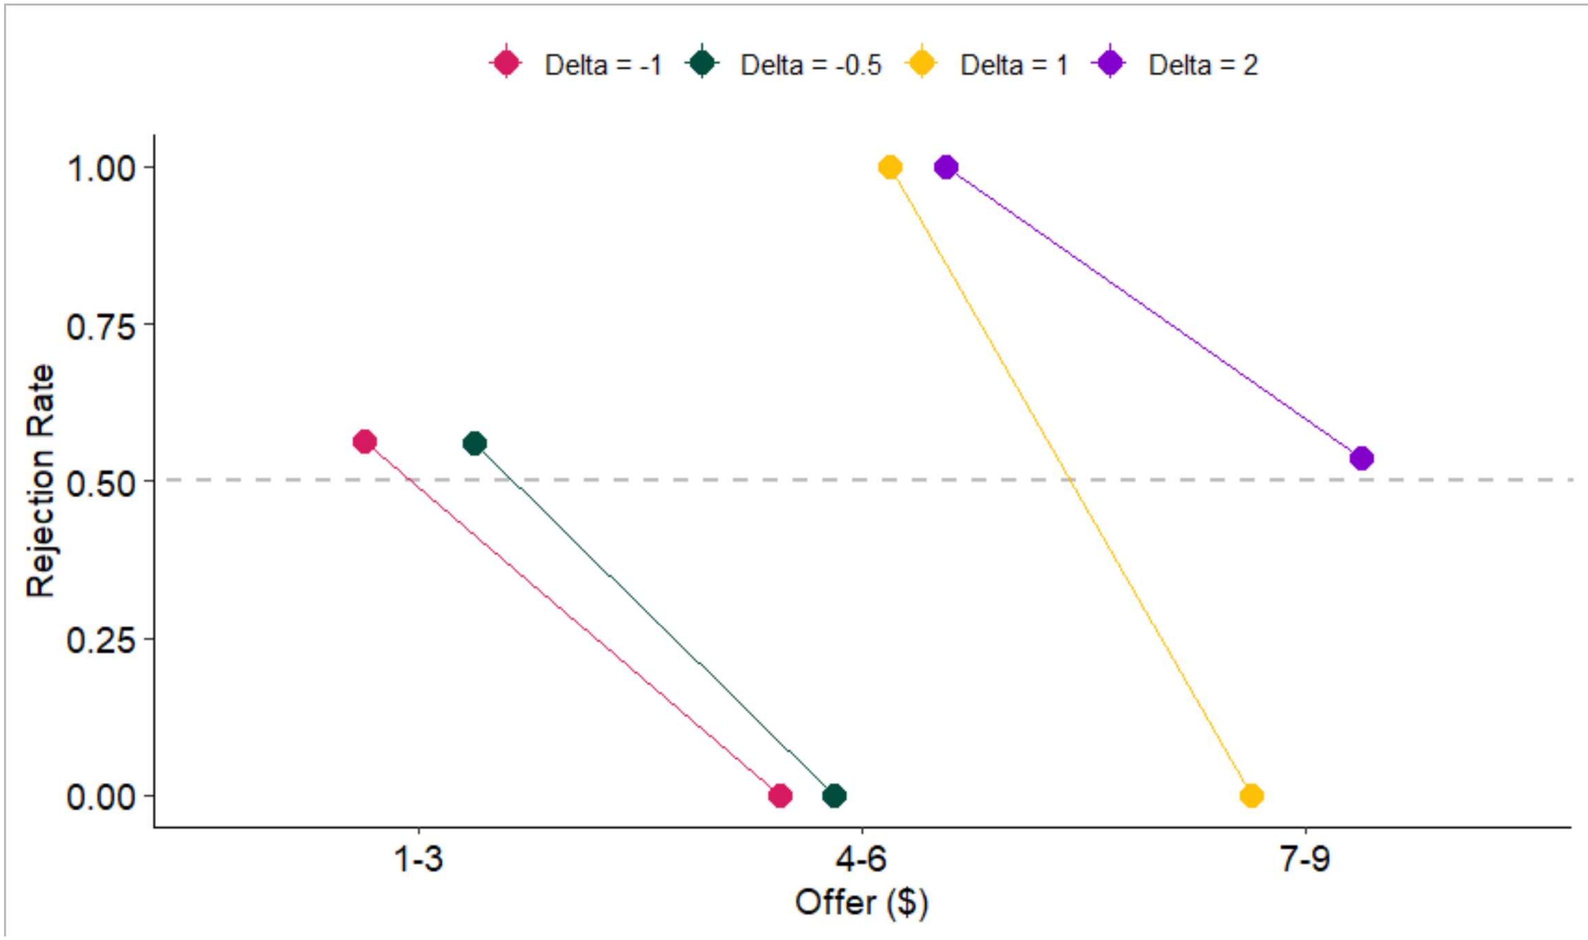

Connecting behavior to parameters

- Flat offer rate
  - Low initial norm
  - Low envy
  - Low, positive influence
- Steep offer rate
  - High envy
  - High initial norm
  - Low adaptation
  - High, positive influence
- Accept all high-offers
  - Moderate-to-low envy
  - Low initial norm
  - High learning rate
  - Moderate, positive influence

Digression:
How do model predictions stack up to human behavior?

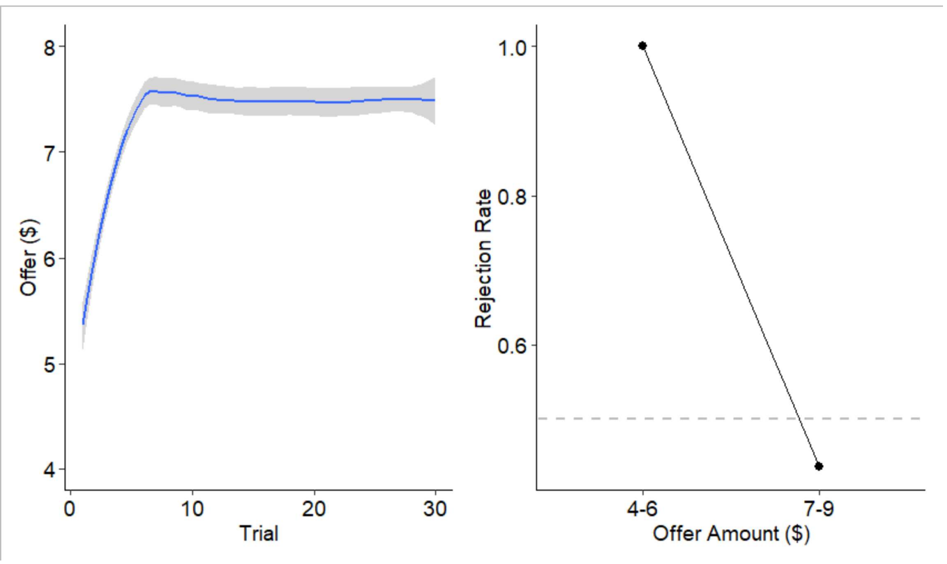

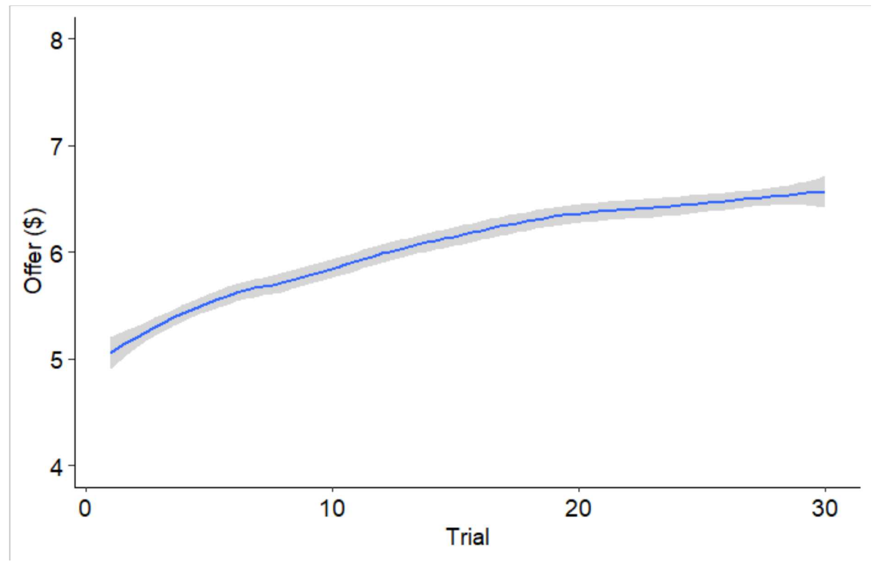

Digression:
How do model predictions stack up to human behavior?

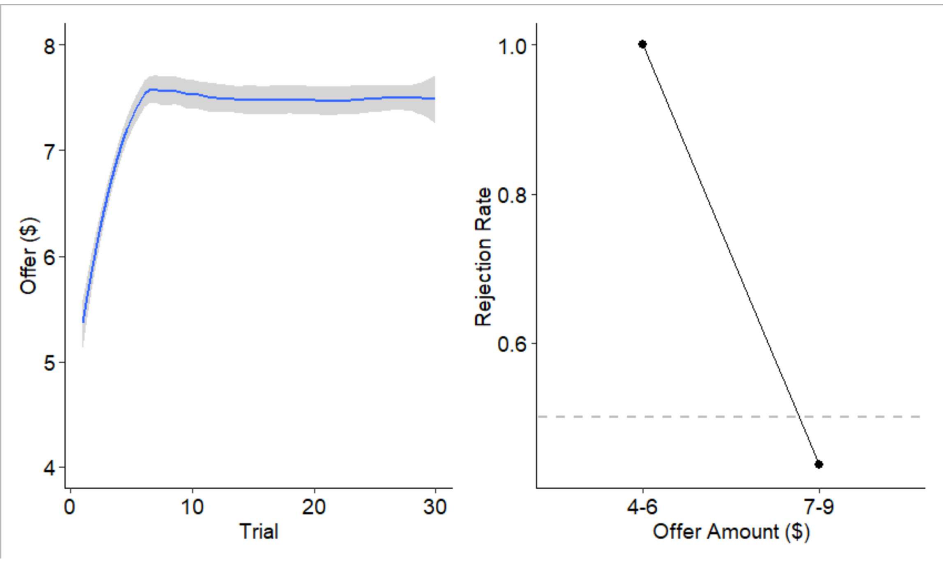

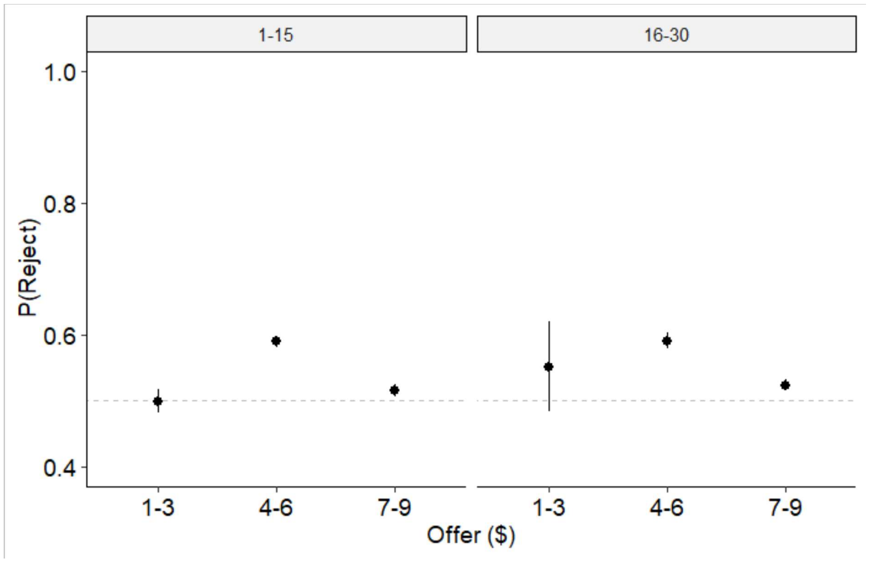

application to value-based decisions

- Value-based decisions: decisions where **each option** have **different values**
- Two-alternative case:
  - Choose between option 1 with value $R_1$and option 2 with value R_2
  - Connect to DDM by setting the drift rate proportional to difference in value

$$\boldsymbol{v} = \boldsymbol{d}(R_1 - R_2) + \varepsilon = d\Delta R + \varepsilon$$

application to value-based decisions

- Choice probabilities

$$p(\text{left}) = \frac{1}{1+\exp(2ad\Delta R)} - \frac{1-\exp(-2x_0 d\Delta R)}{\exp(2\text{a}d\Delta R)-\exp(-2\text{a}d\Delta R)}$$

- Special case with unbiased starting point ($x_0 = 0$)

$$p(\text{left}) = \frac{1}{1+\exp(2ad\Delta R)}$$

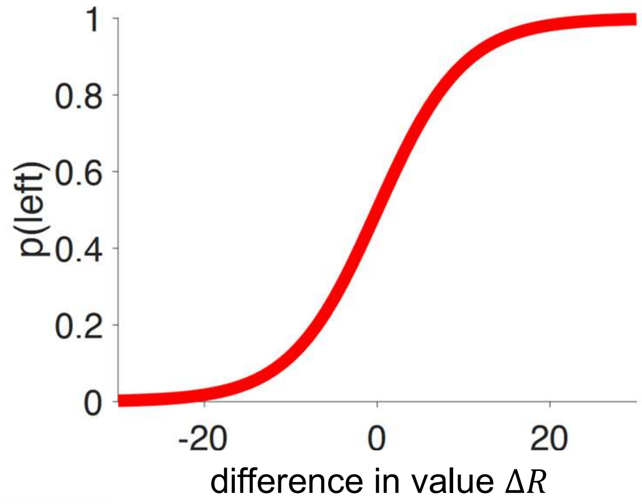

- This is the **softmax probability function!**

softmax :: DDM connection

- Compare the two:

DDM: $p(\text{left}) = \frac{1}{1+\exp(2ad\Delta R)}$ softmax: $p(\text{left}) = \frac{1}{1+\exp(\beta\Delta R)}$

- Softmax's inverse temperature parameter (β) is controlled by two DDM parameters: threshold (a) and signal-to-noise ratio (d)

$$\beta = 2ad$$

softmax :: DDM connection

- In the DDM, stochasticity in choice can be generated by:
  - Reduced signal-to-noise ratio (d)
  - Lower threshold (a)
- Different mechanisms *<u>cannot be distinguished by choices alone</u>*

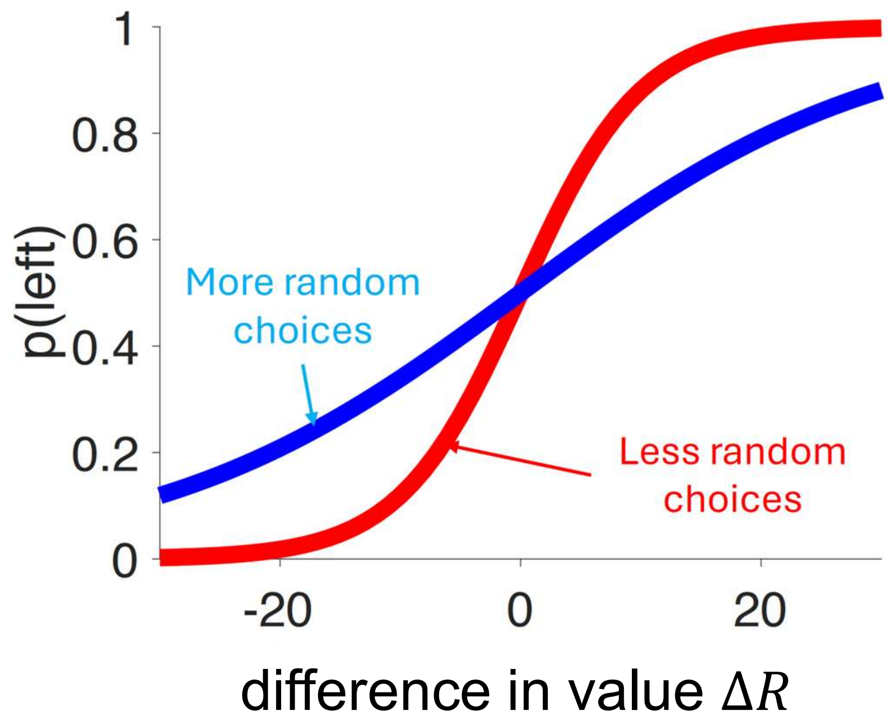

Response times in the value-based DDM

- Response time formula

$$\mathrm{RT} = t_0 + \frac{a}{\Delta R}\tanh(a\Delta R) + \frac{a}{\Delta R} \times \frac{2(1 - \exp(-2x_0\Delta R))}{\exp(2a\Delta R) - \exp(-2a\Delta R)} - x_0\Delta R$$

- Special case with unbiased starting bias $x_0 = 0$

$$\mathrm{RT} = t_0 + \frac{a}{\Delta R}\tanh(a\Delta R)$$

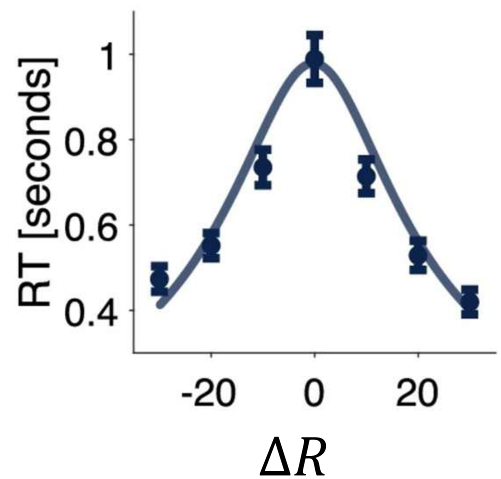

Response times in the value-based DDM

- Changes to drift rate and threshold have **opposite** effects on response times