Modeling individual differences in a pragmatic reference game

as a consequence of variable disengagement from unsuccessful strategies

John Duff ❖ Alexandra Mayn ❖ Vera Demberg
Saarland University, Dept. of Language Science & Technology

UNIVERSITÄT DES SAARLANDES

European Research Council
Established by the European Commission

MathPsych/ICCM Symposium on Computational Psycholinguistics

21 July 2024

jduff@lst.uni-saarland.de

Gricean pragmatics

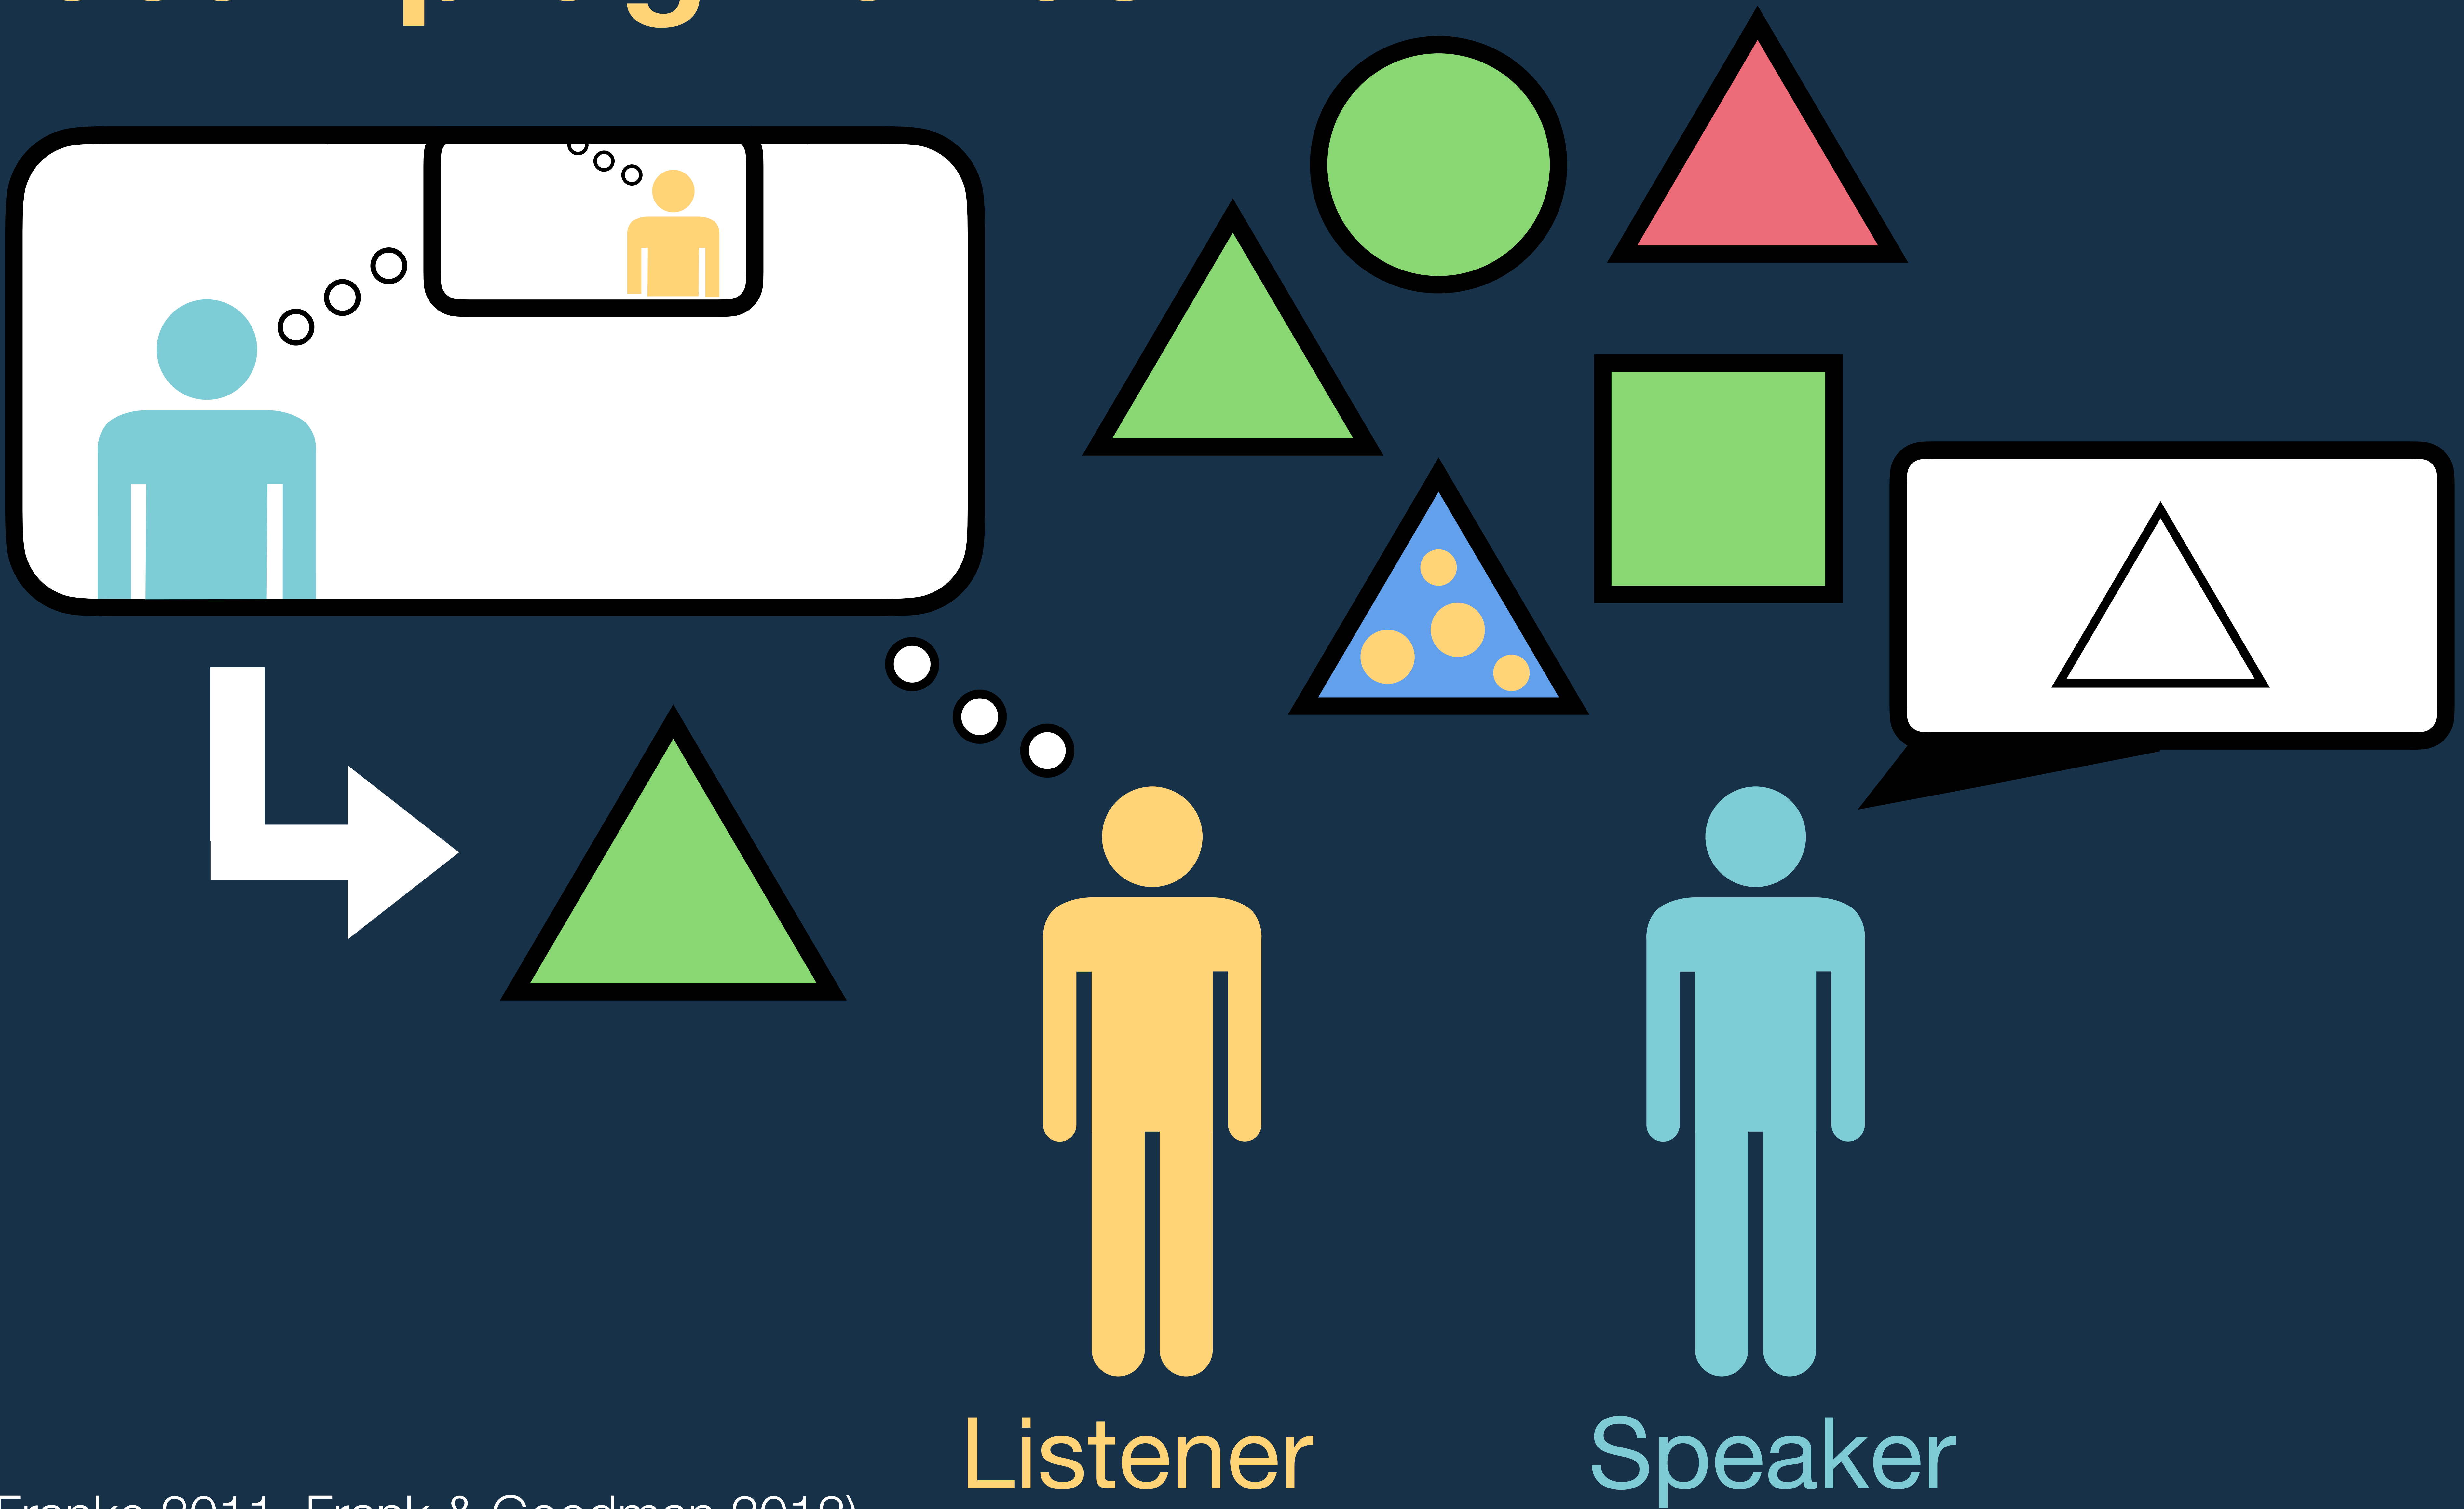

(Grice 1975, Franke 2011, Frank & Goodman 2012)

Two empirical complications

Pragmatic reasoning in games only emerges over time

(Sikos et al. 2021)

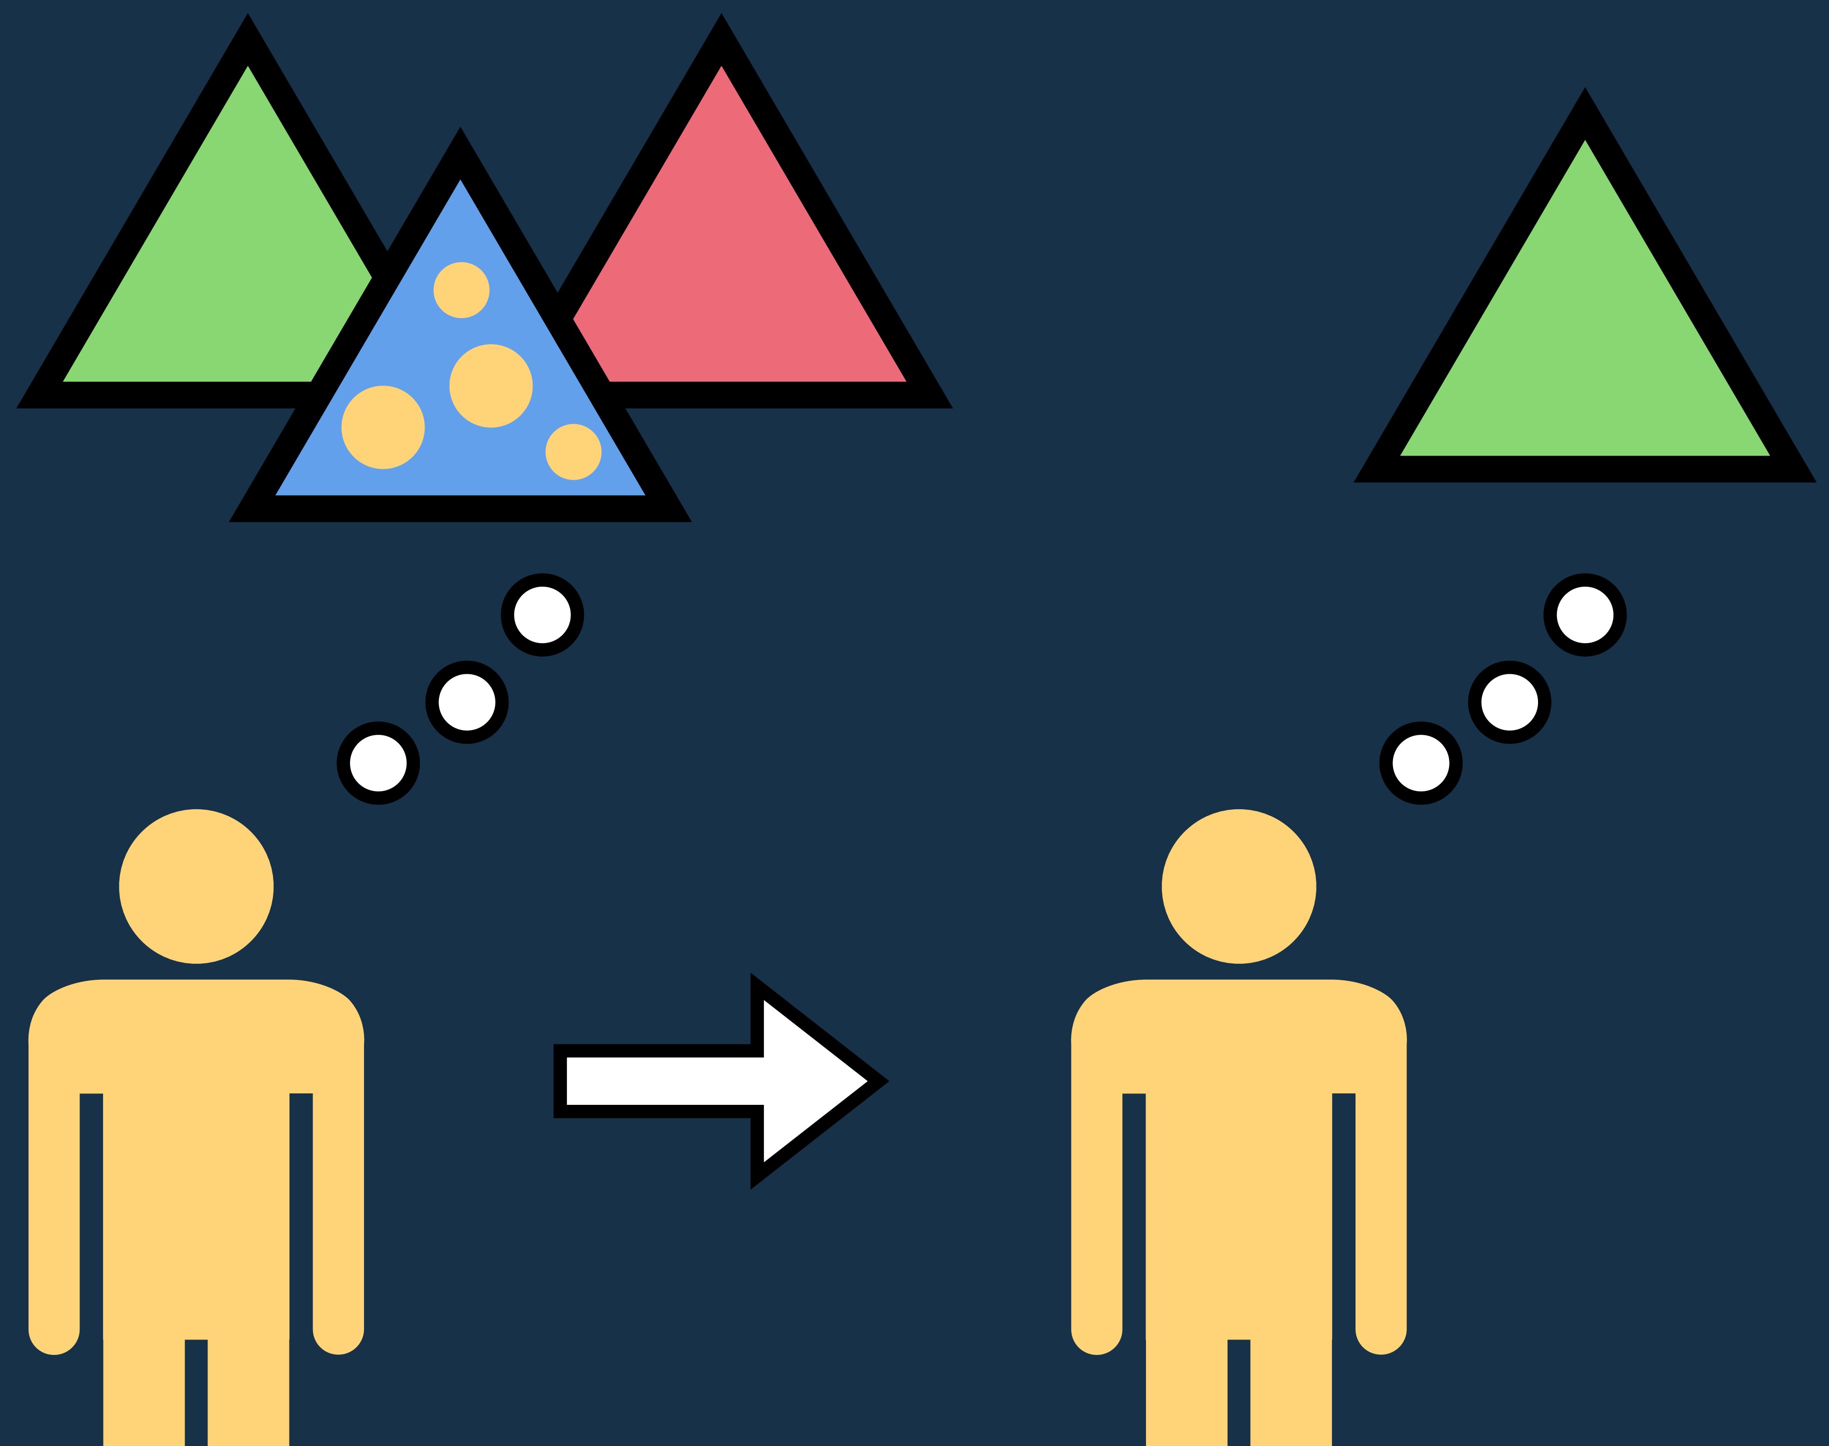

Individuals vary in their depth of pragmatic reasoning

(Franke & Degen 2016, Mayn & Demberg 2023)

Modeling performance via reinforcement learning

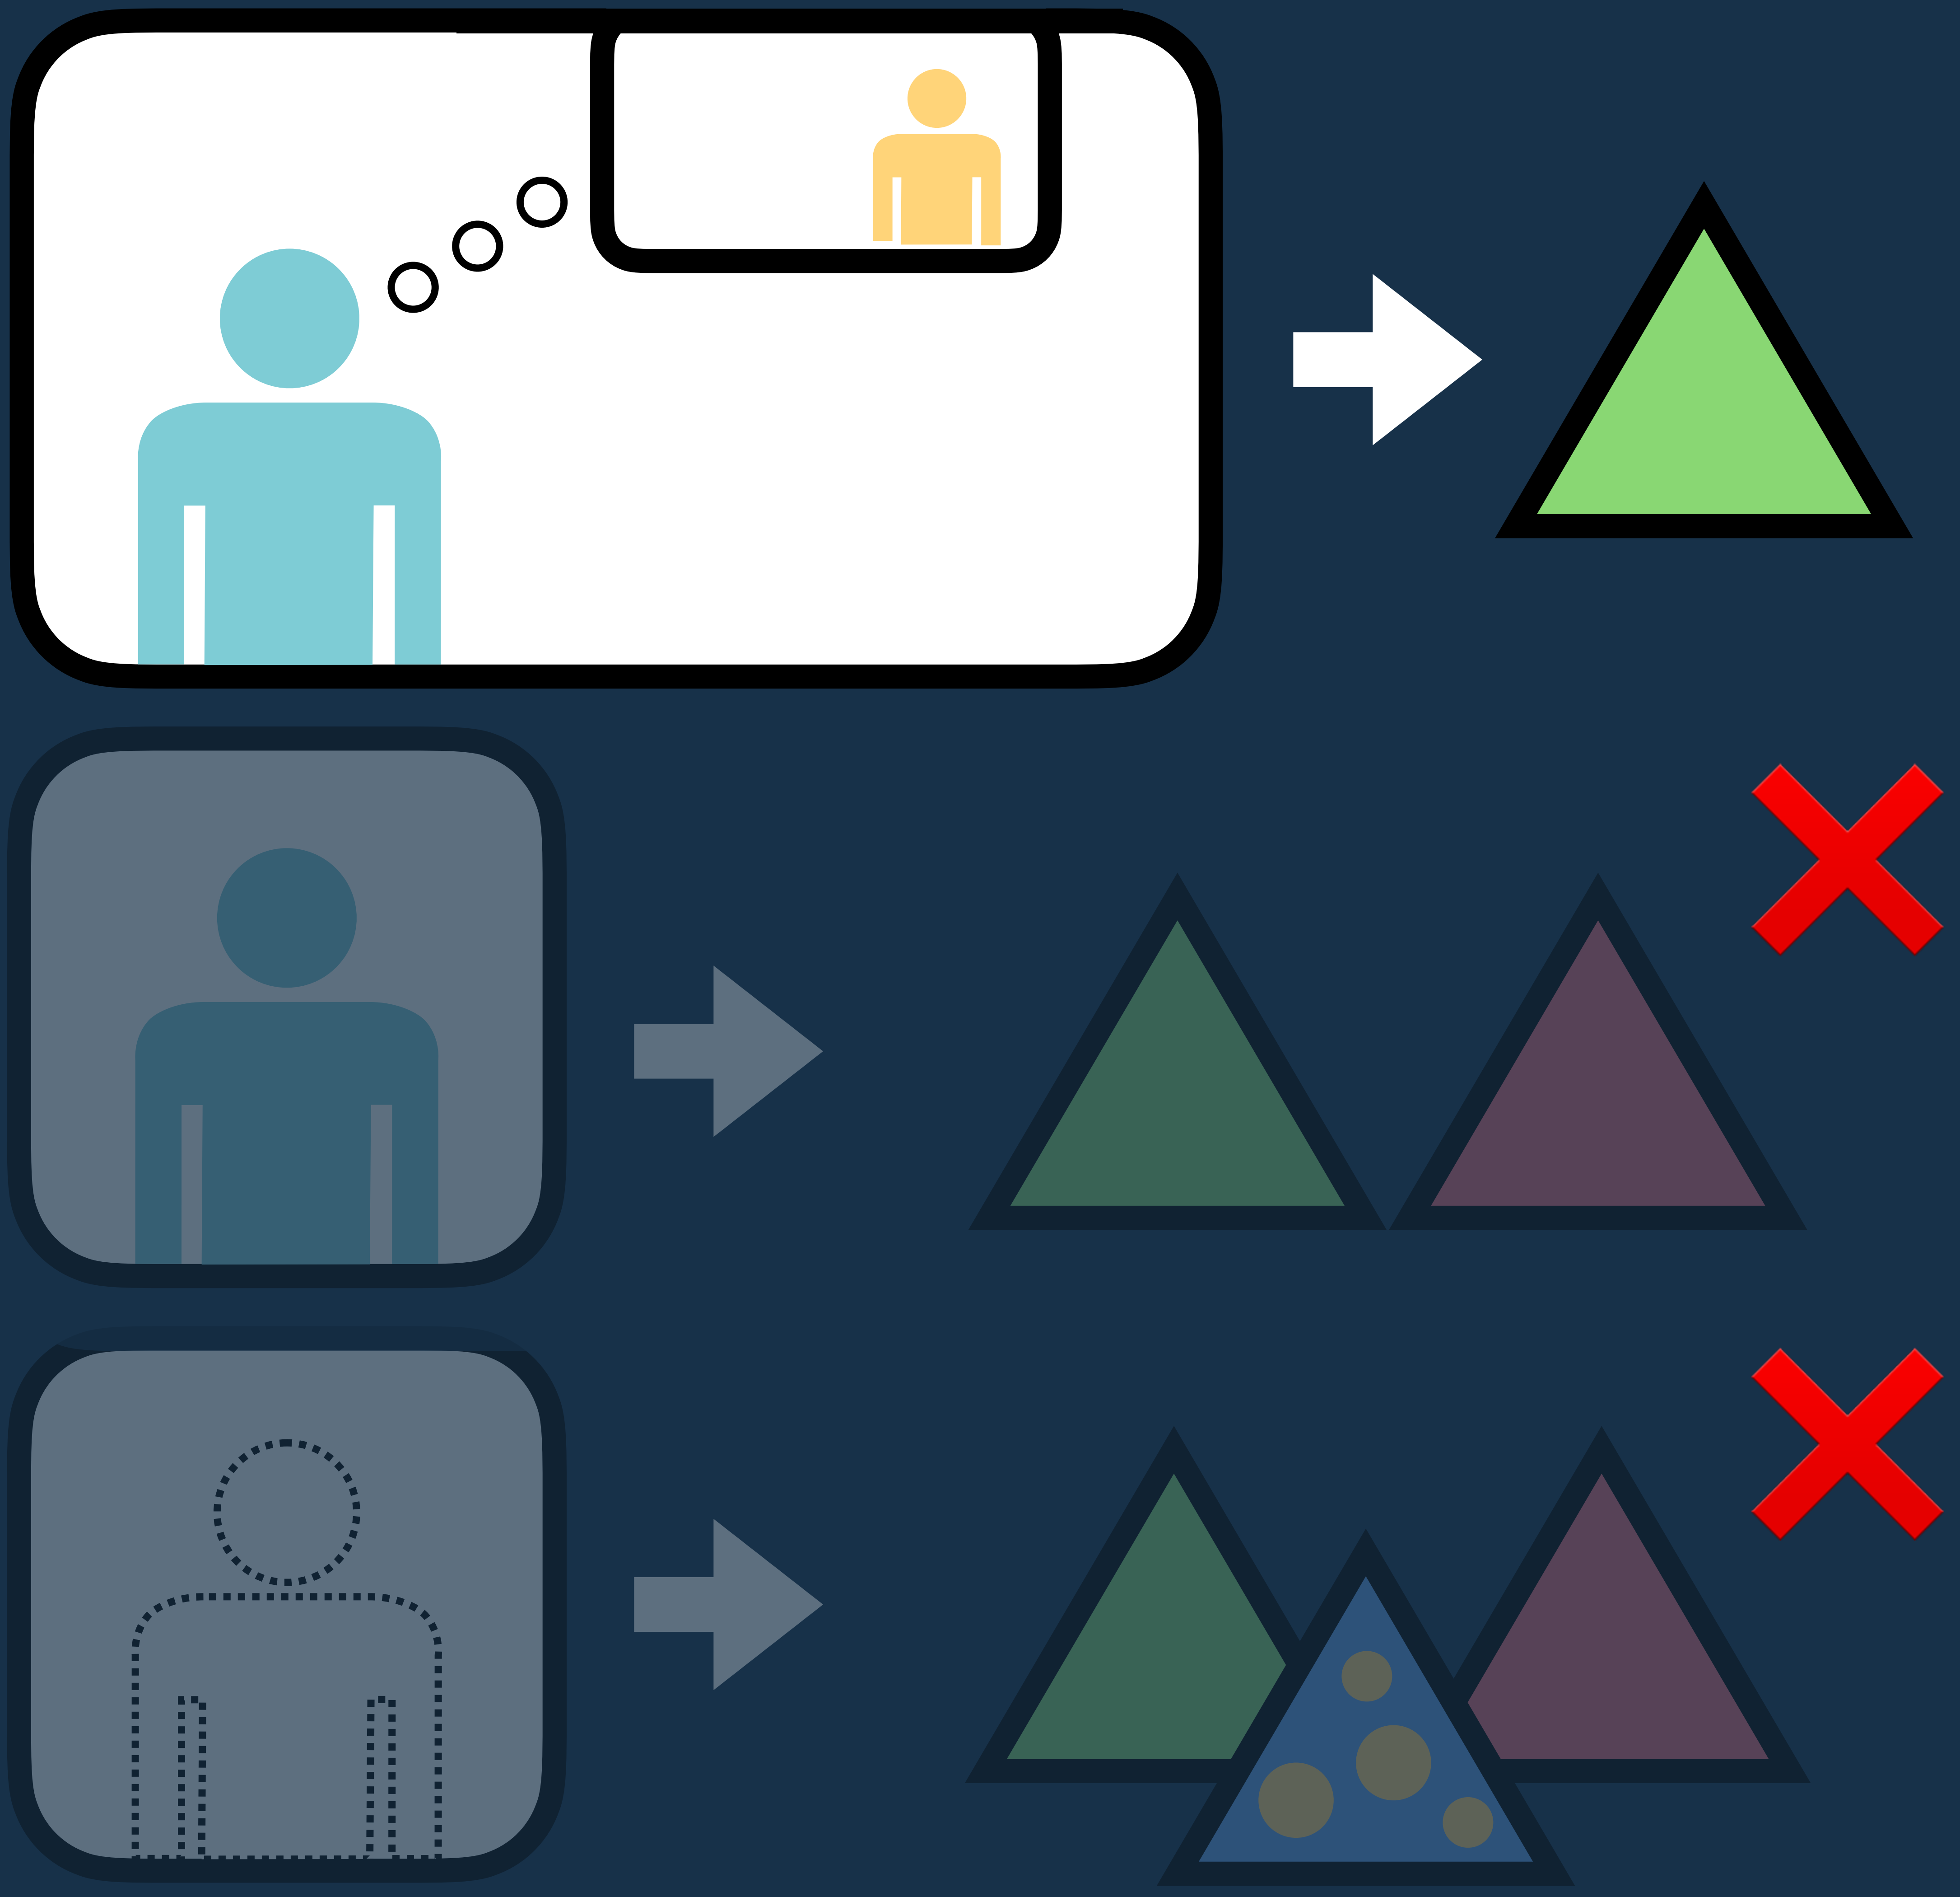

Comprehenders find an optimal strategy through exploration and failure

(cf. Stocco et al. 2021)

Roadmap

1. Background
2. Our ACT-R model
3. Modeling individual differences across tasks

Pragmatic reference game (RefGame)

(Frank & Goodman 2012 and following; cf. Wittgenstein 1953)

(“Trivial” Trial)

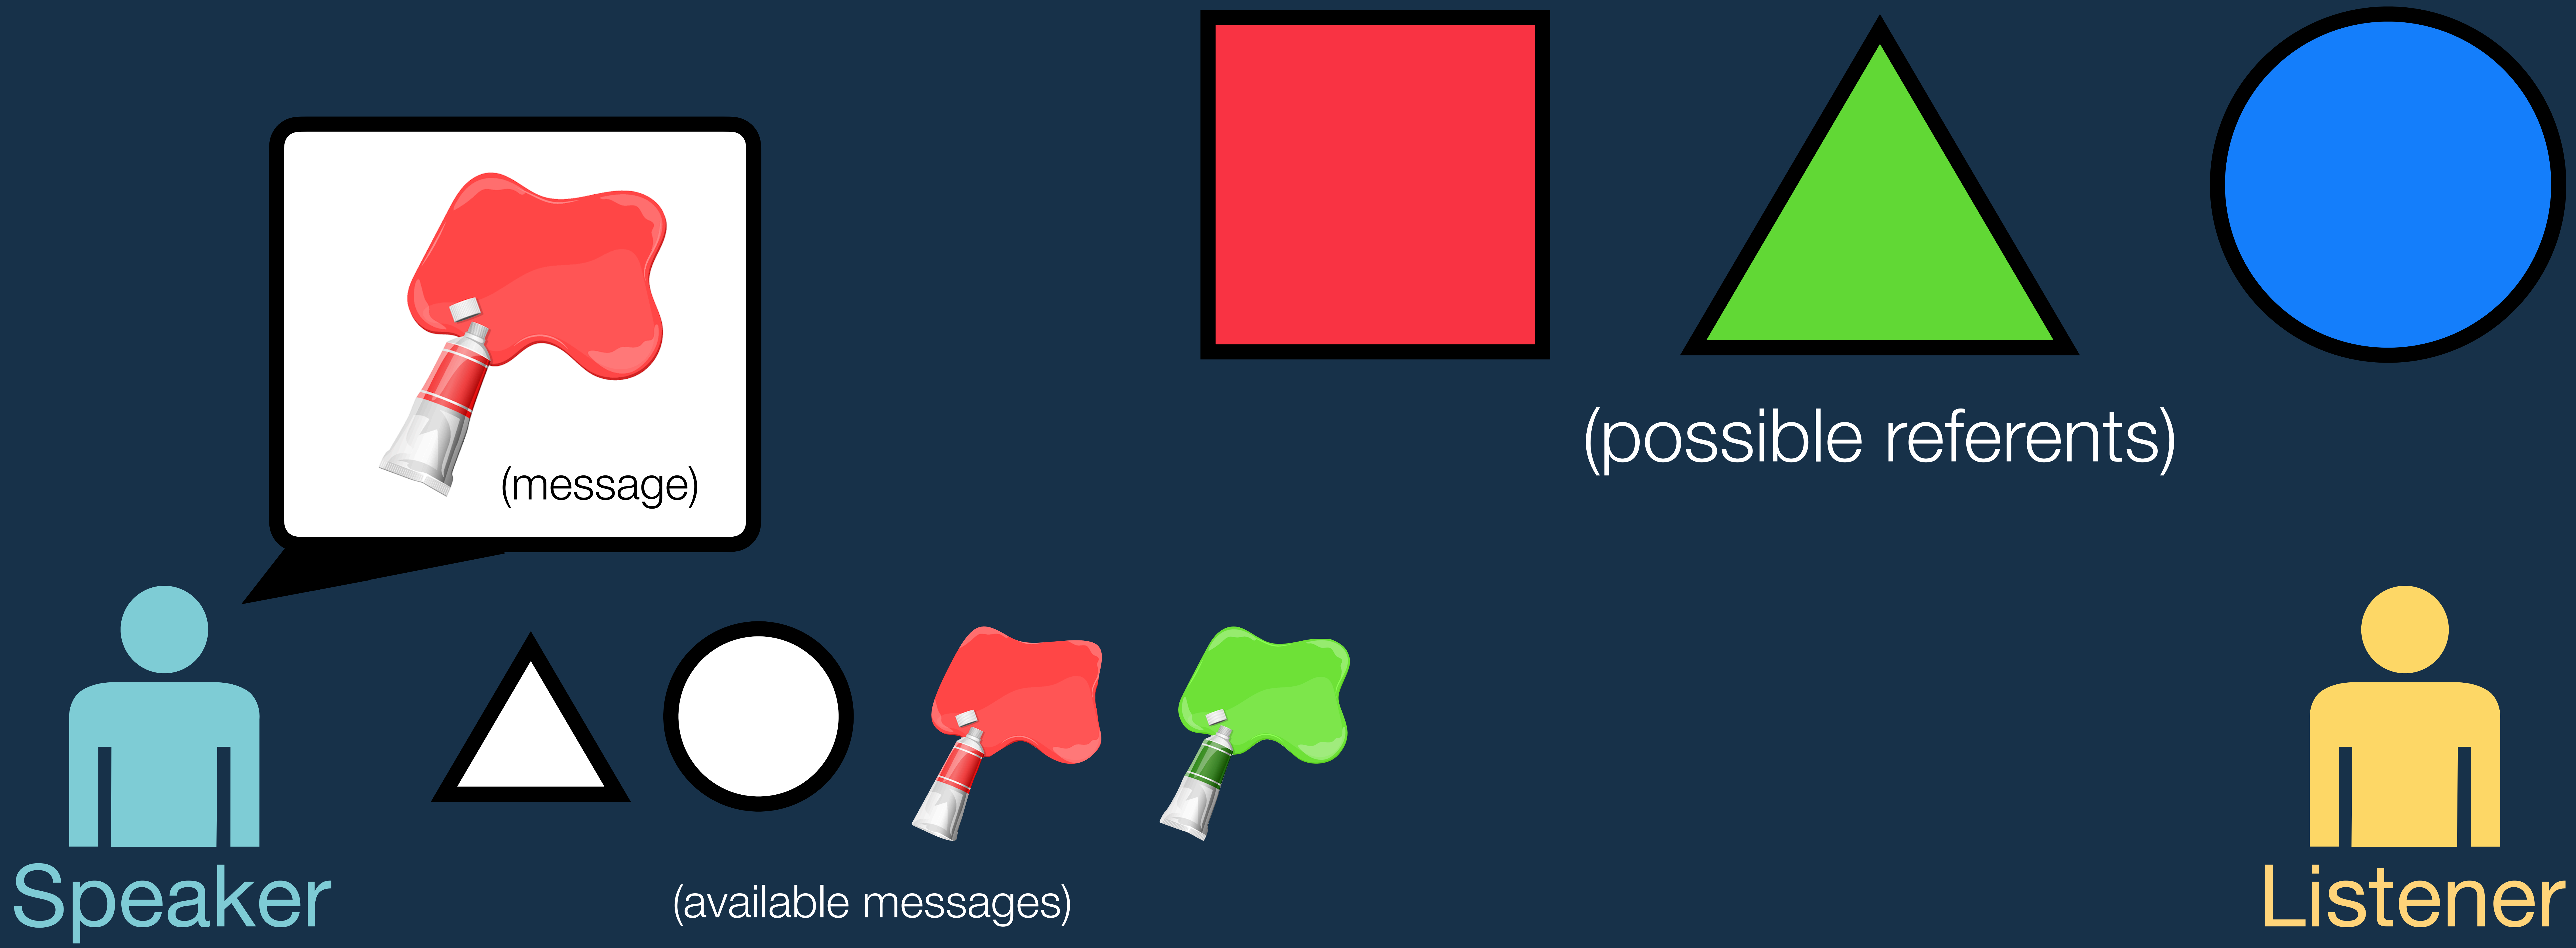

Three interpretive strategies (Franke & Degen 2016)

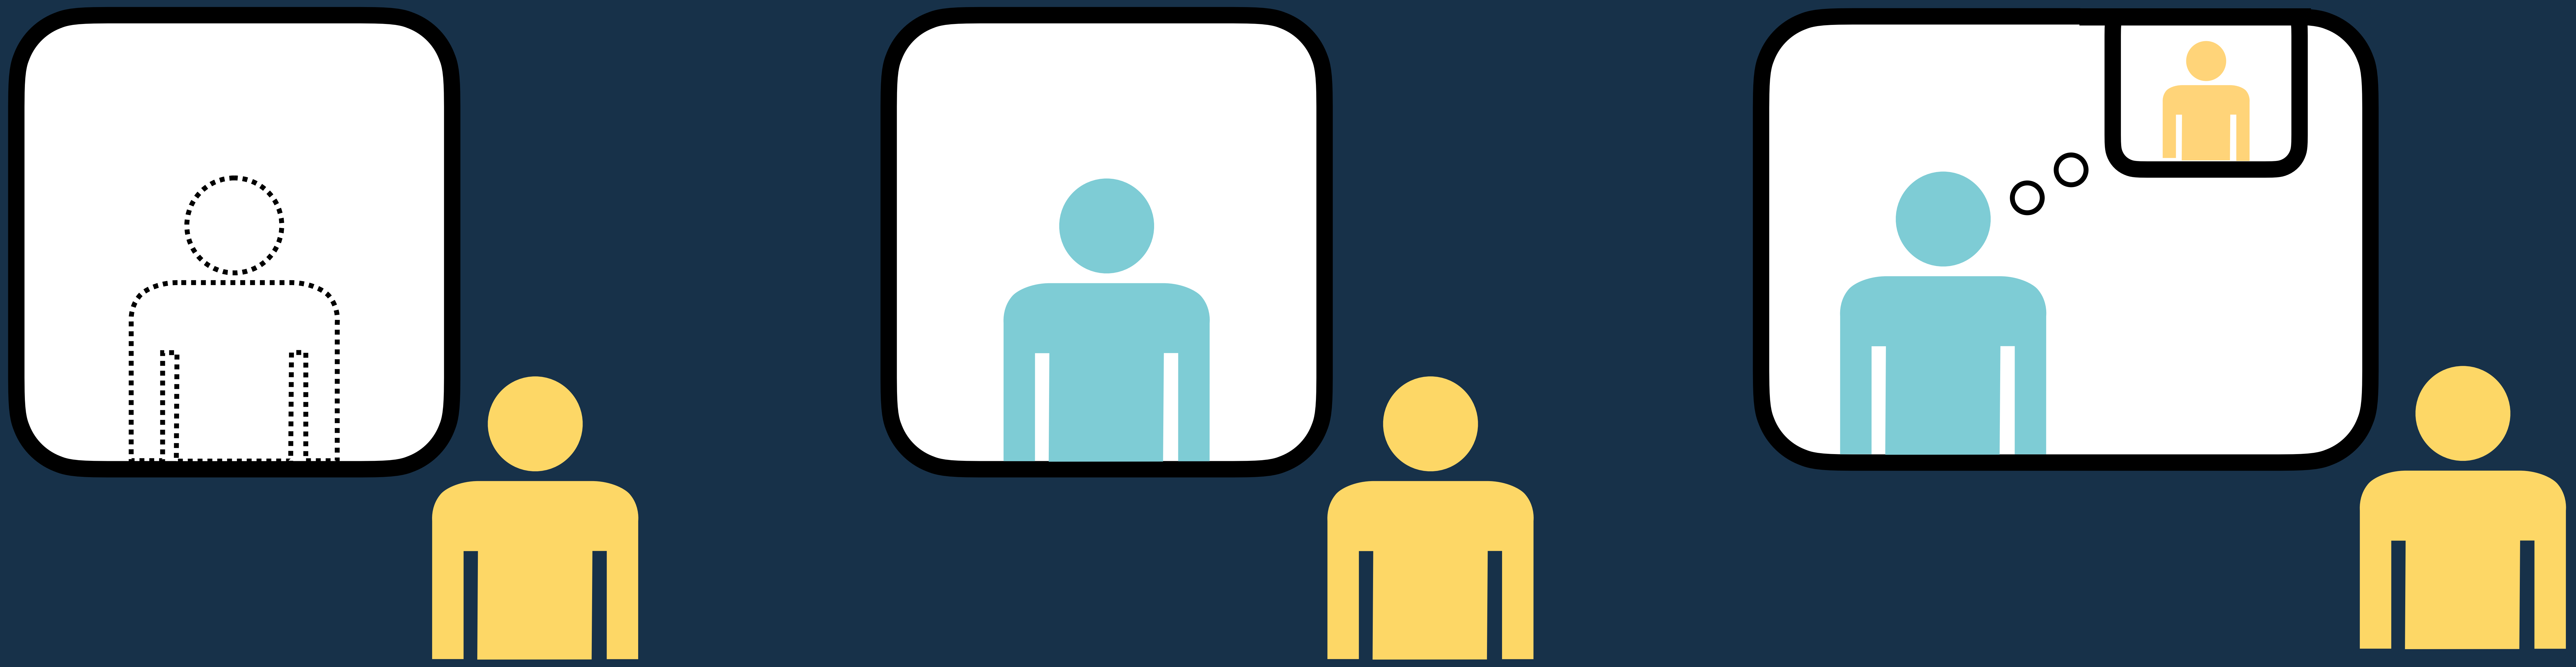

Strategy:

Literal interpretation

First-order pragmatic interpretation

Second-order pragmatic interpretation

“Simple” Trials

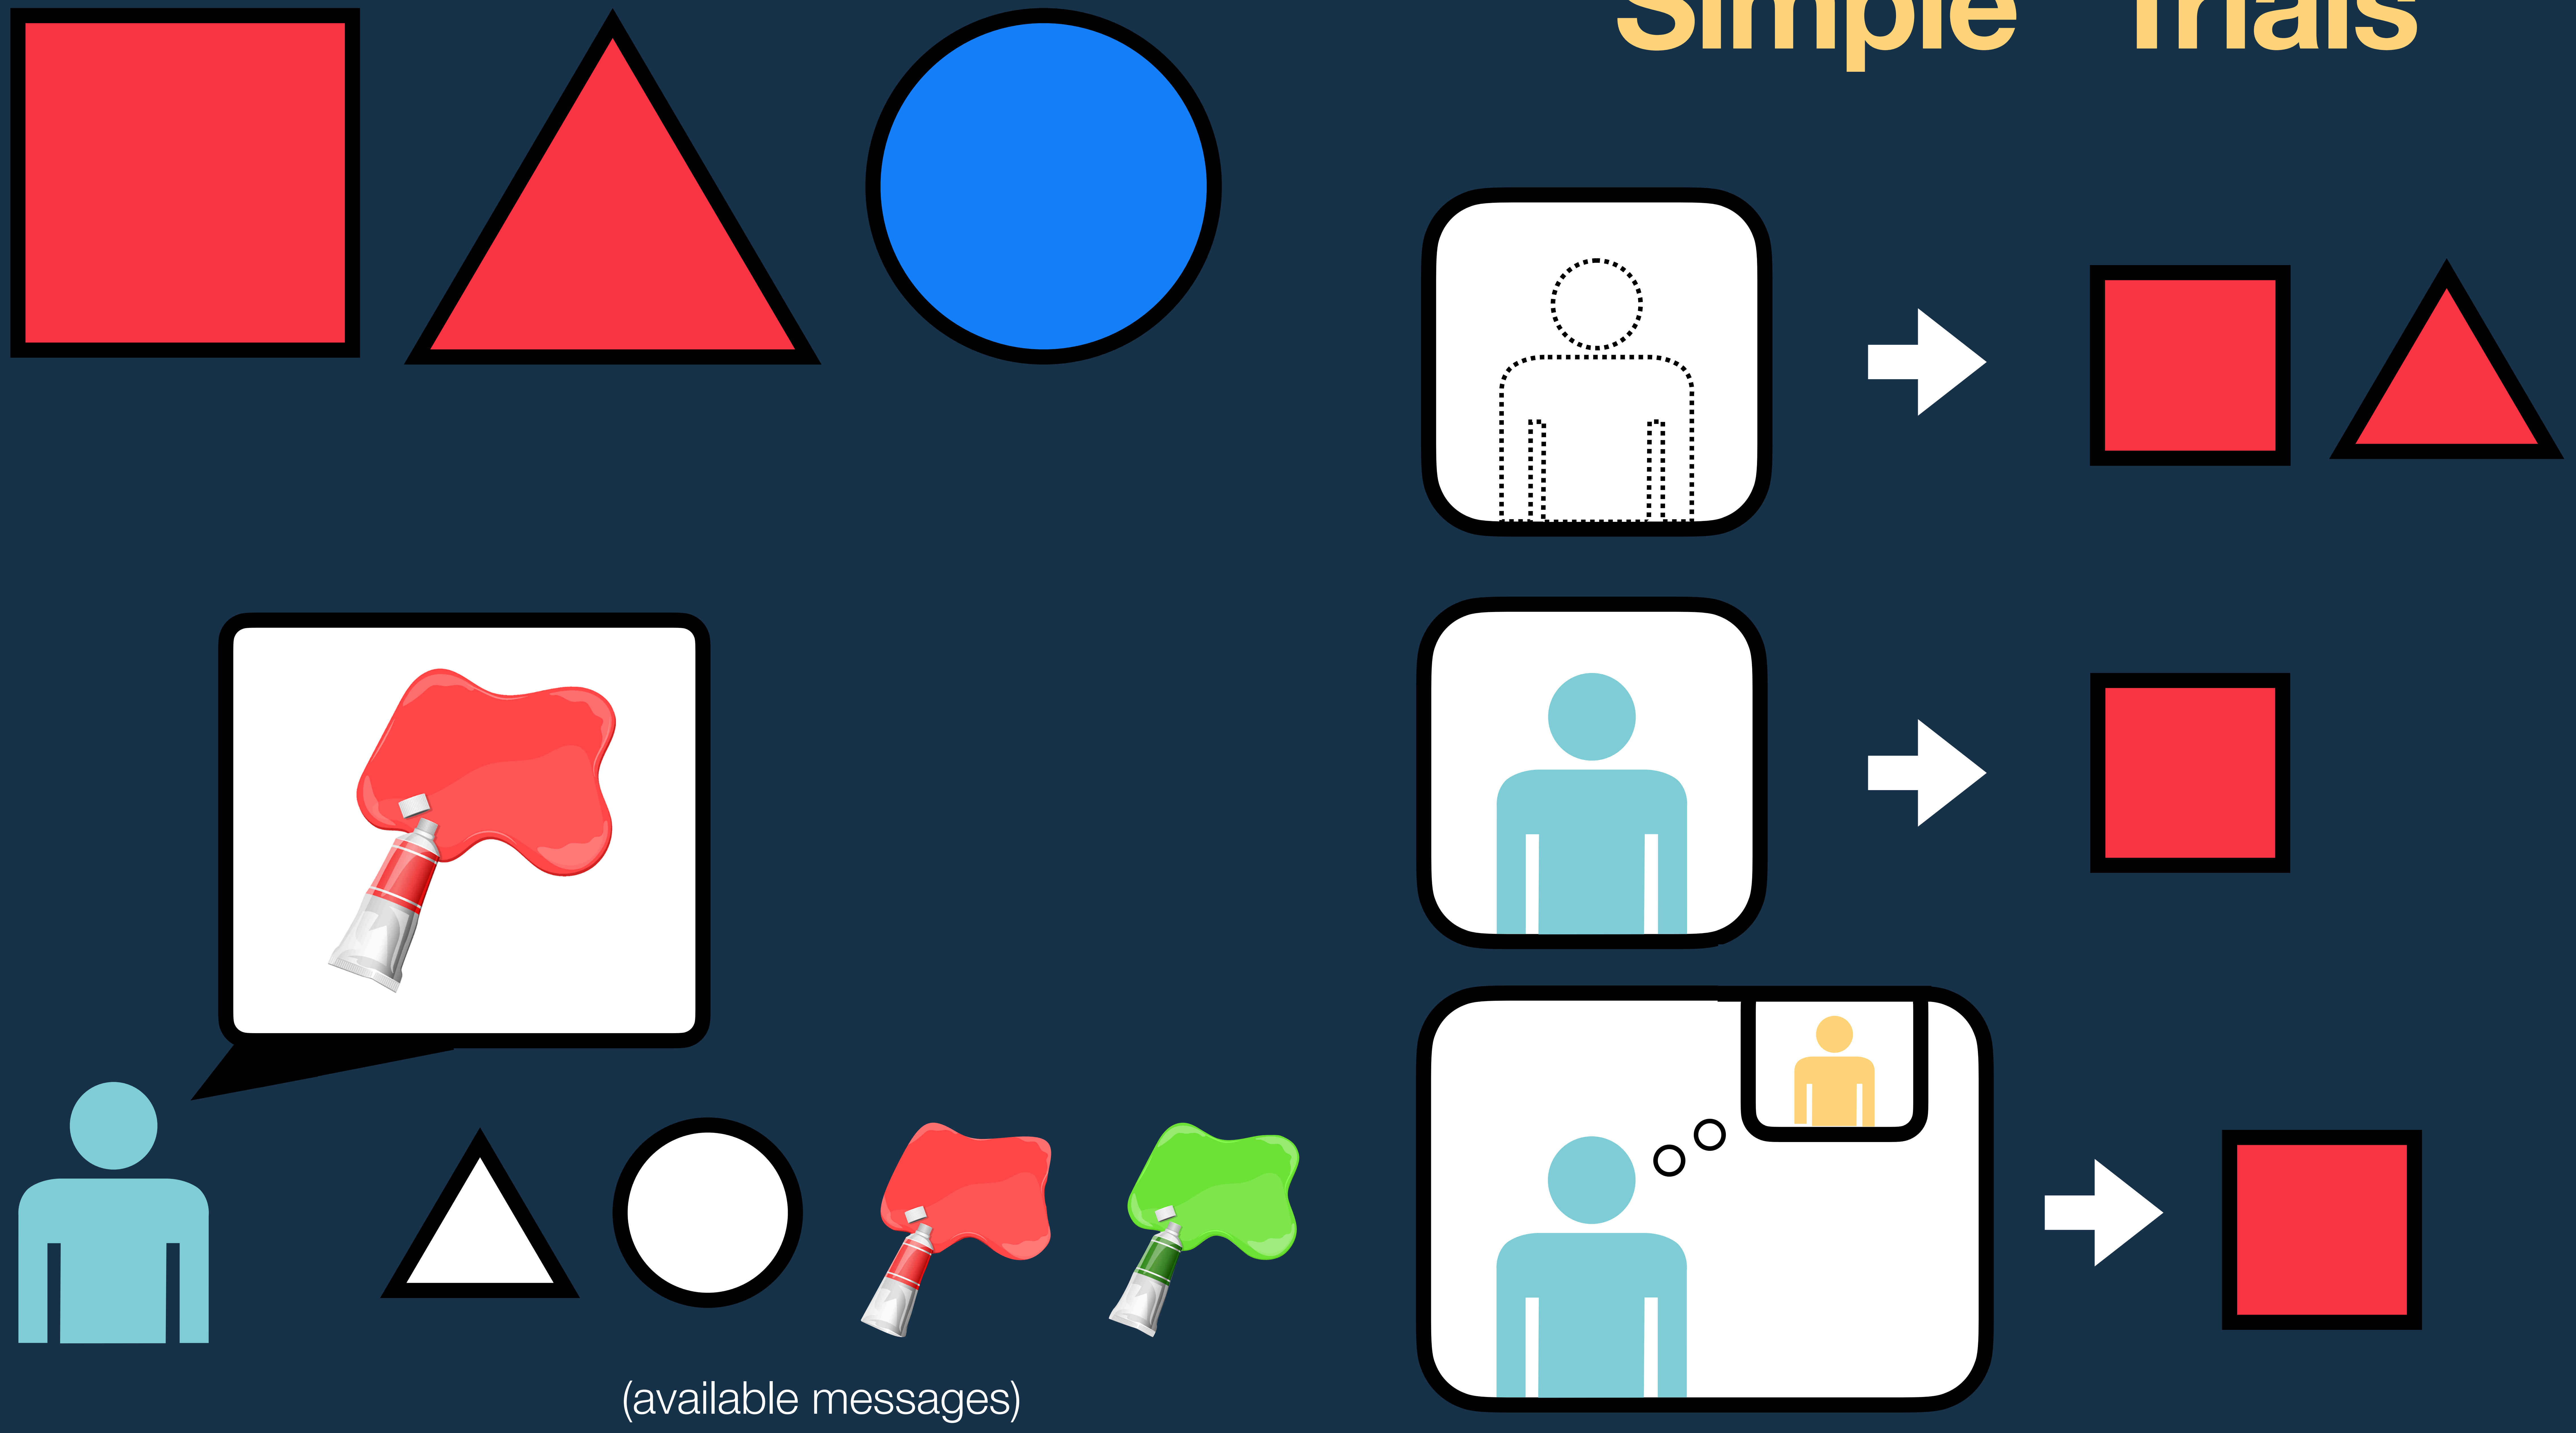

"Complex" Trials

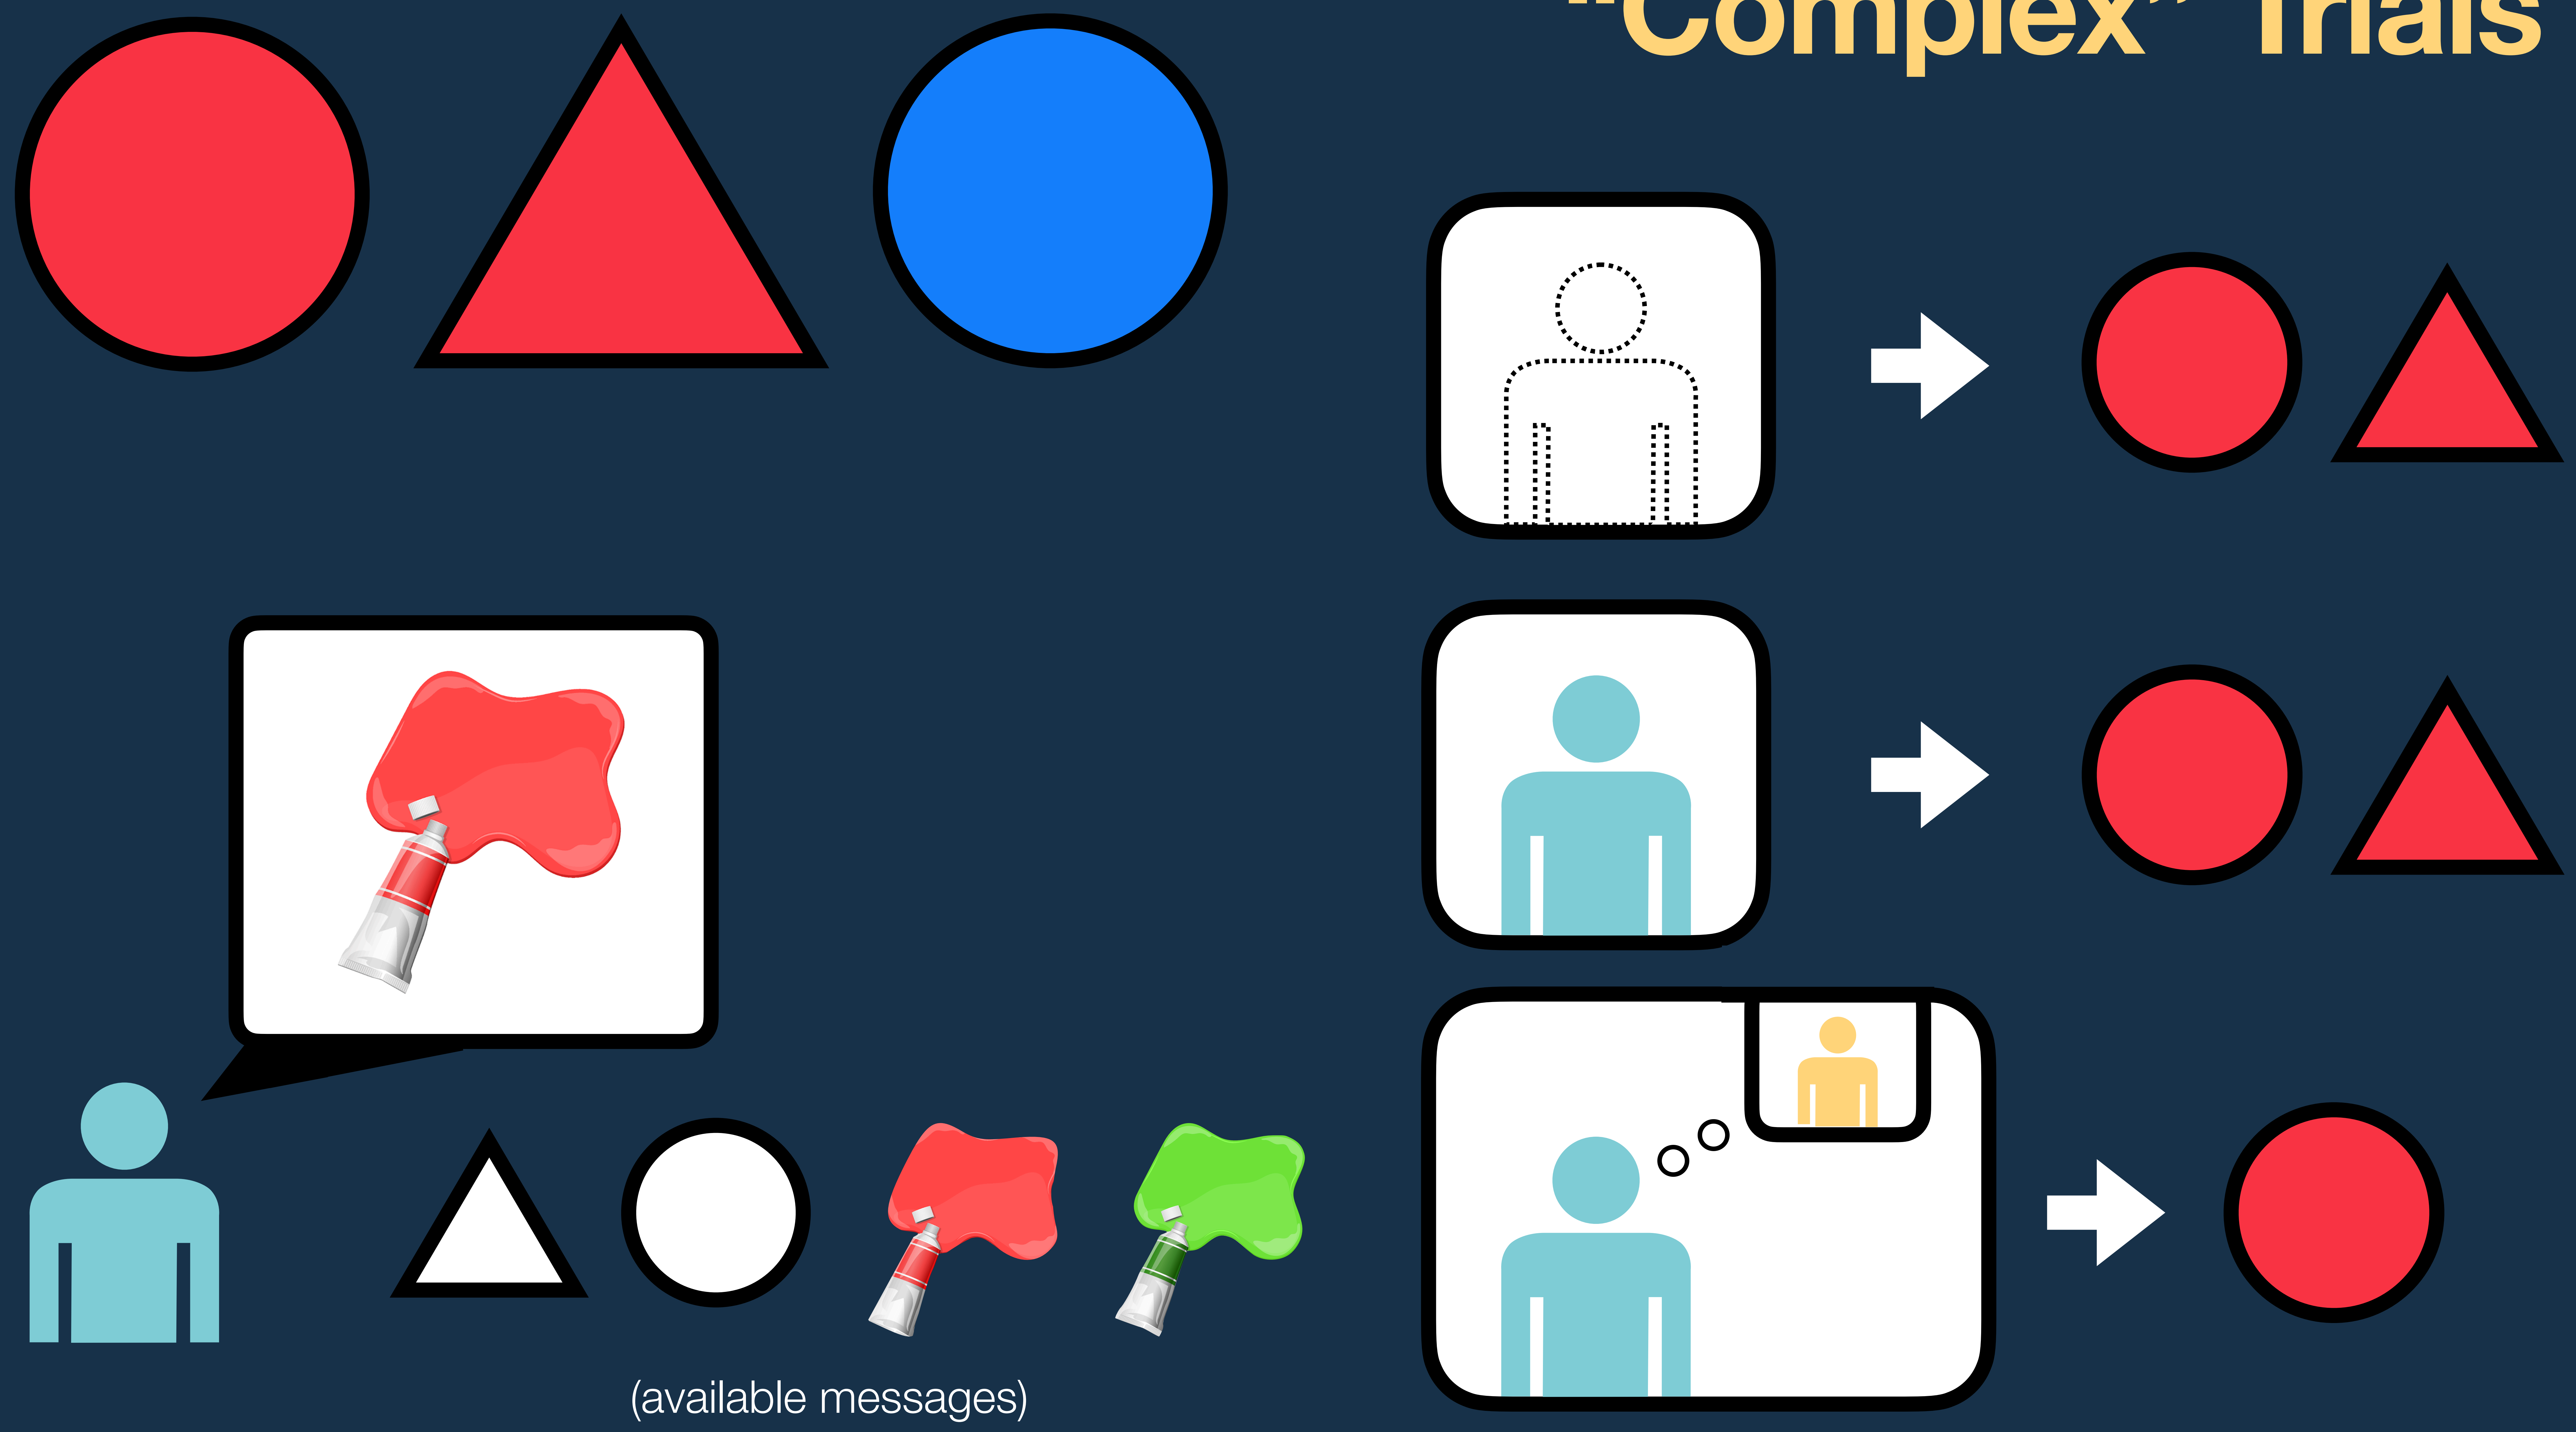

Expected success by strategy

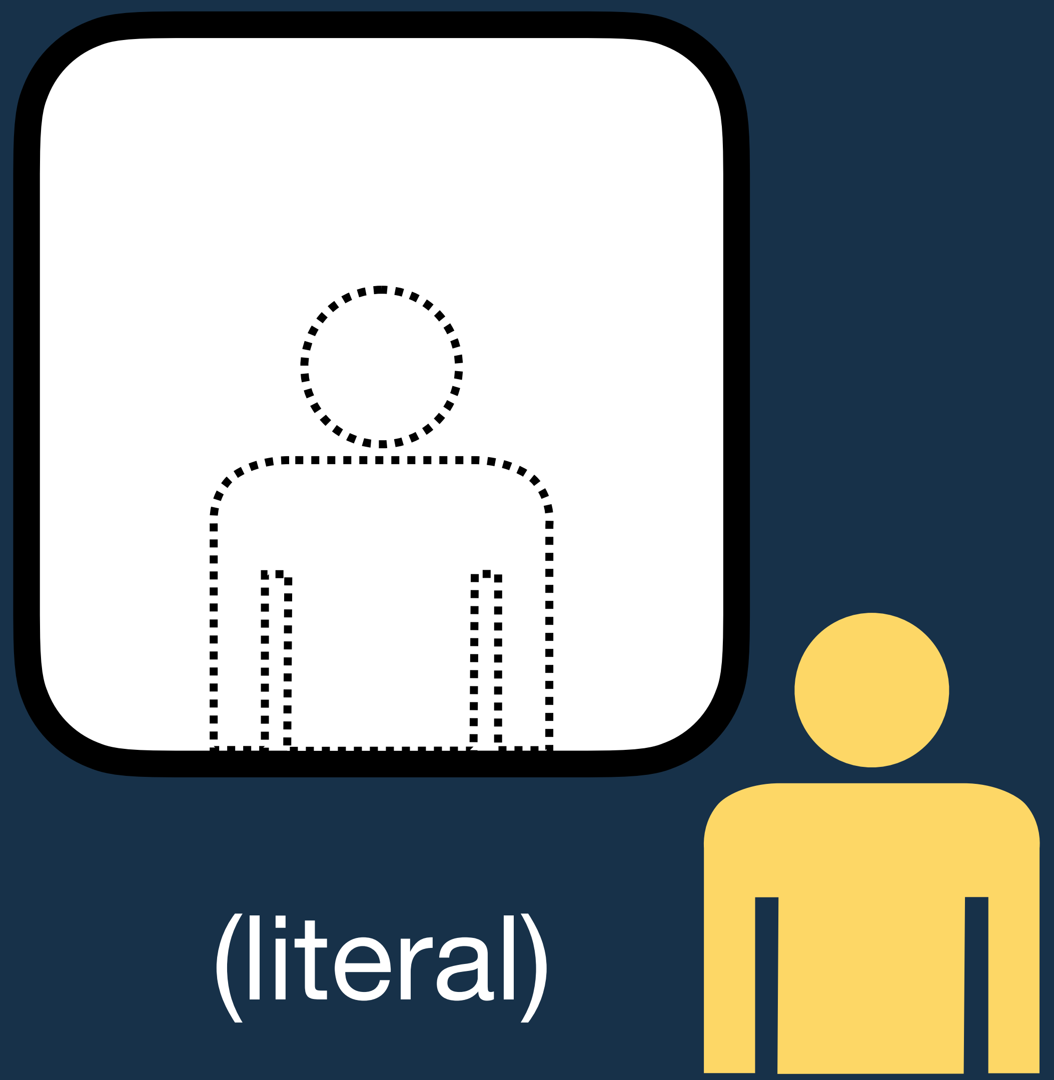

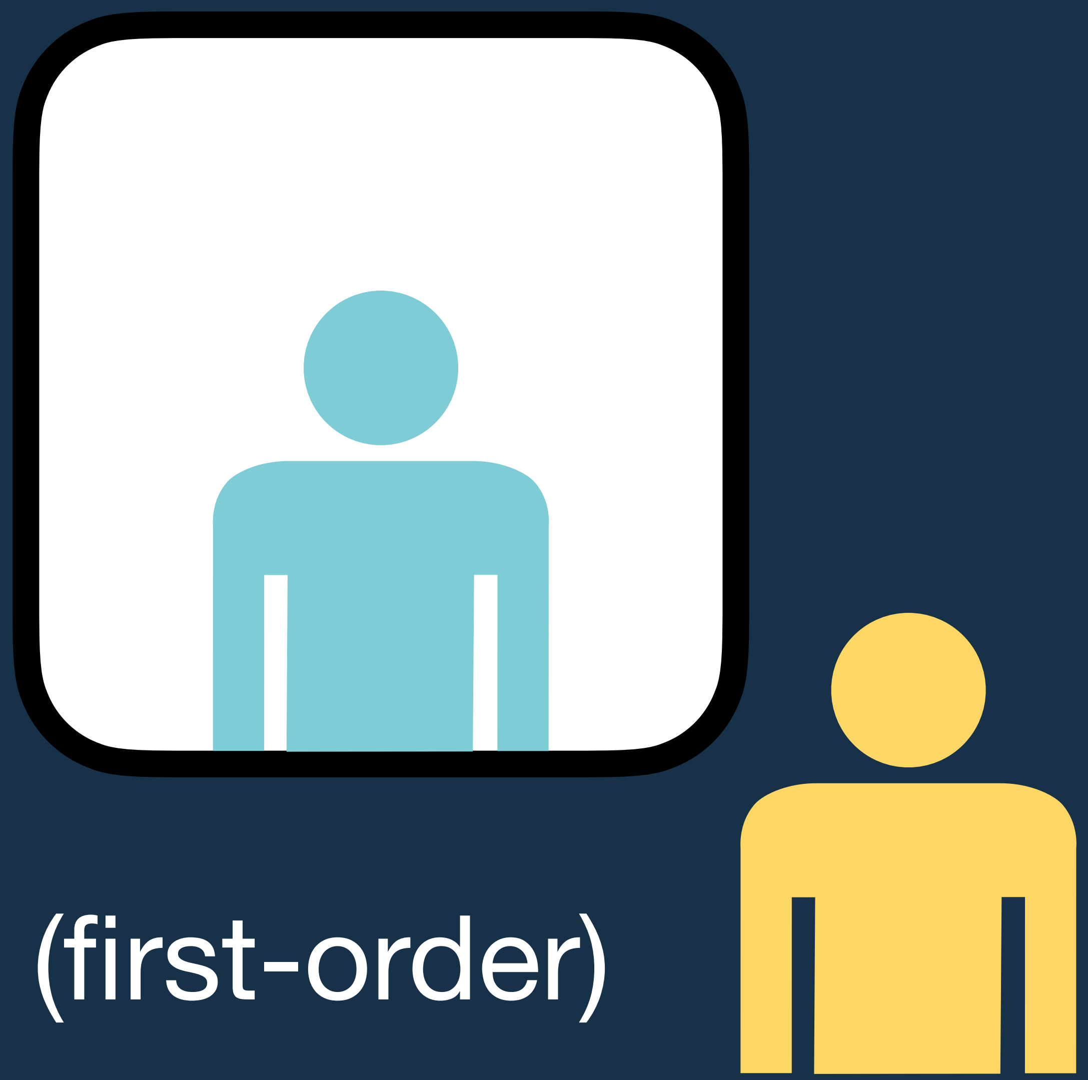

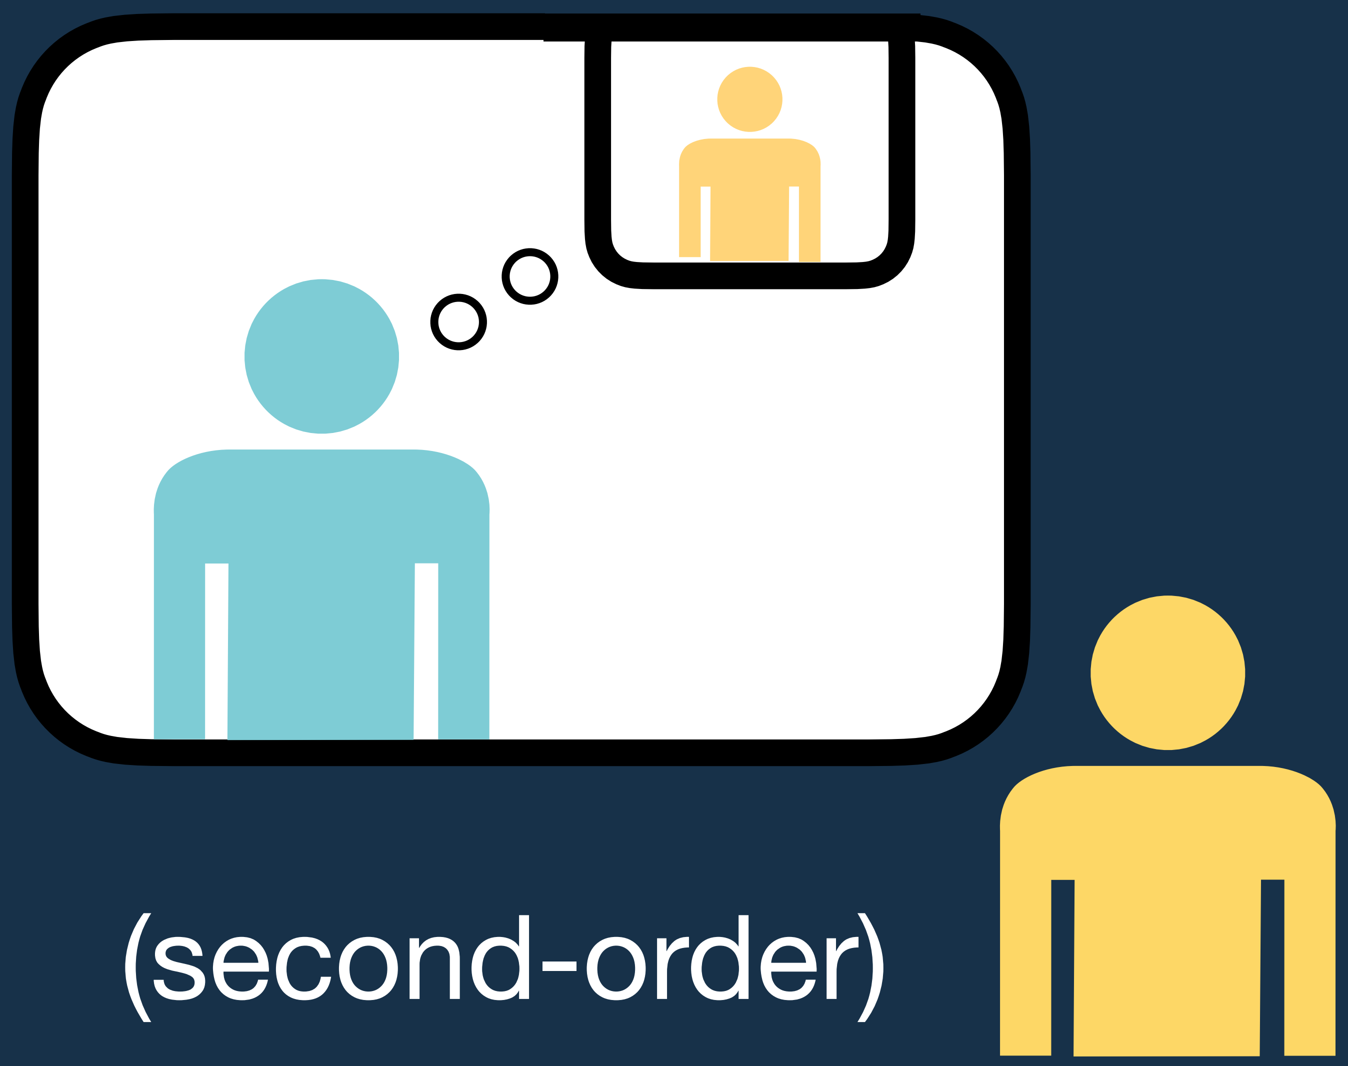

	(literal)	(first-order)	(second-order)
Picks:	Matching referents	Matching referents with fewest alternative messages	Matching referents with no more-informative messages
Trivial:	✅	✅	✅
Simple:	–	✅	✅
Complex:	–	–	✅

Observation #1: No second-order reasoning in one-shot experiments

Sikos et al. (2021)

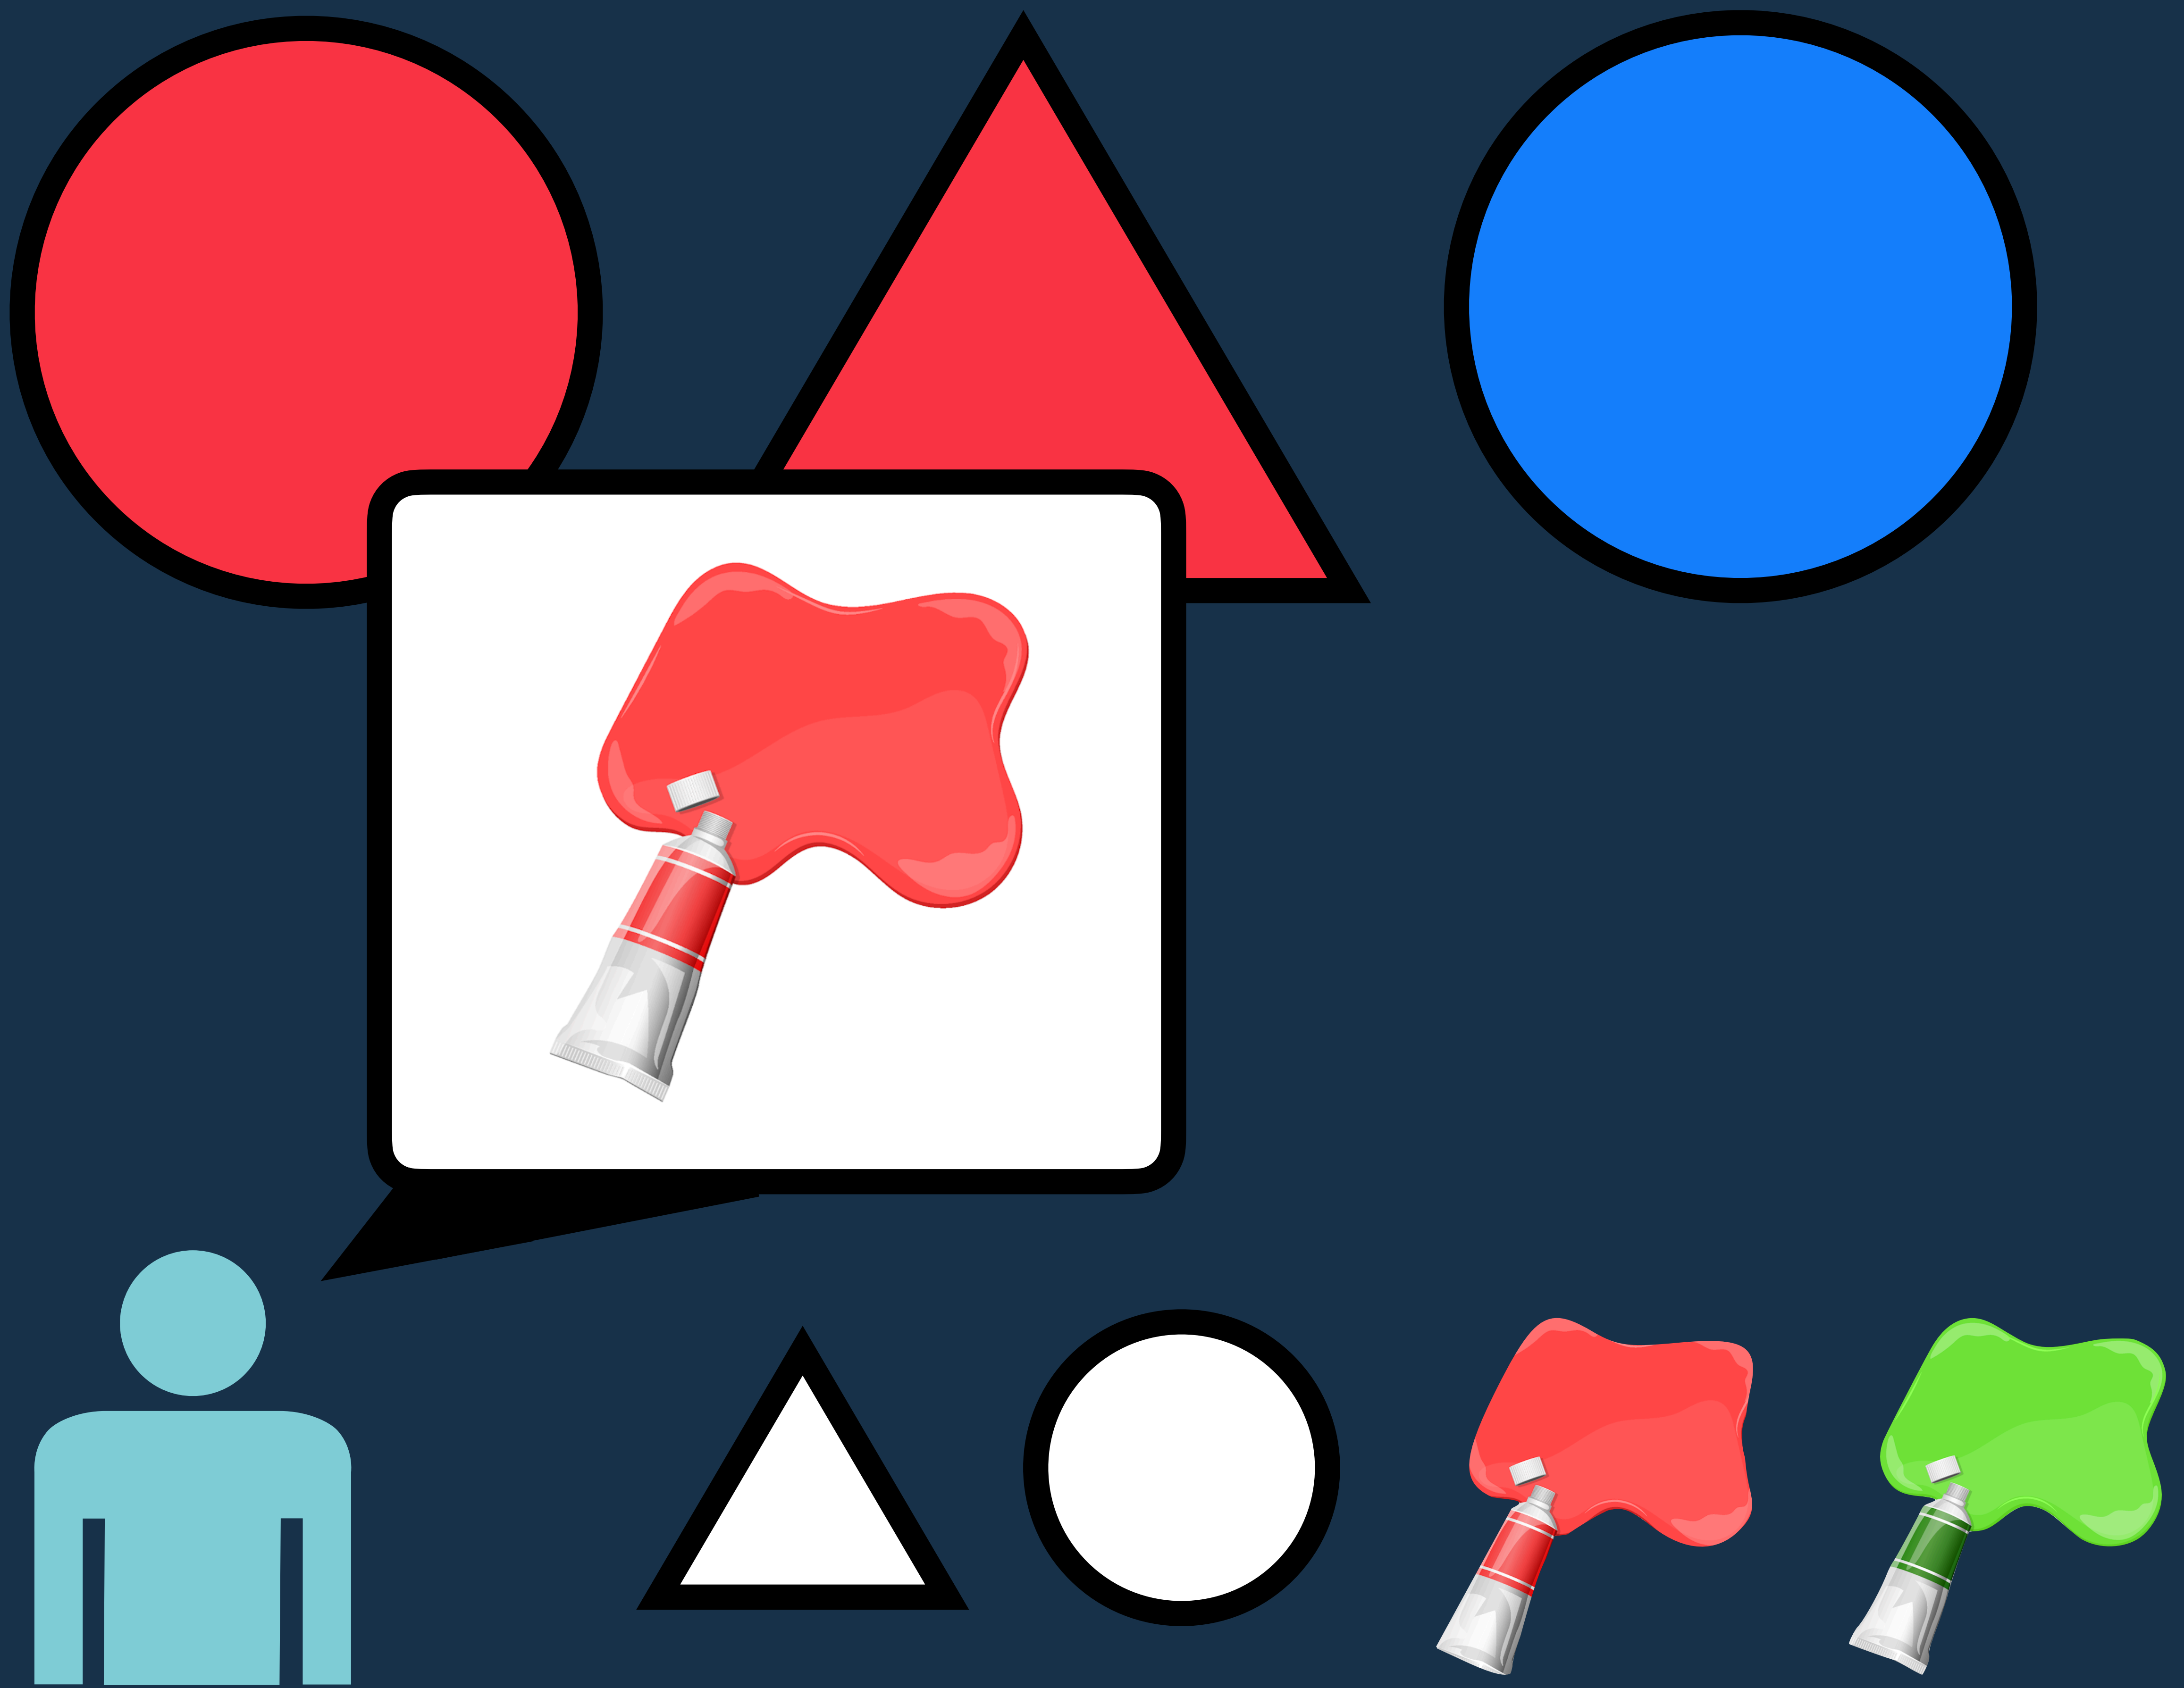

(available messages)

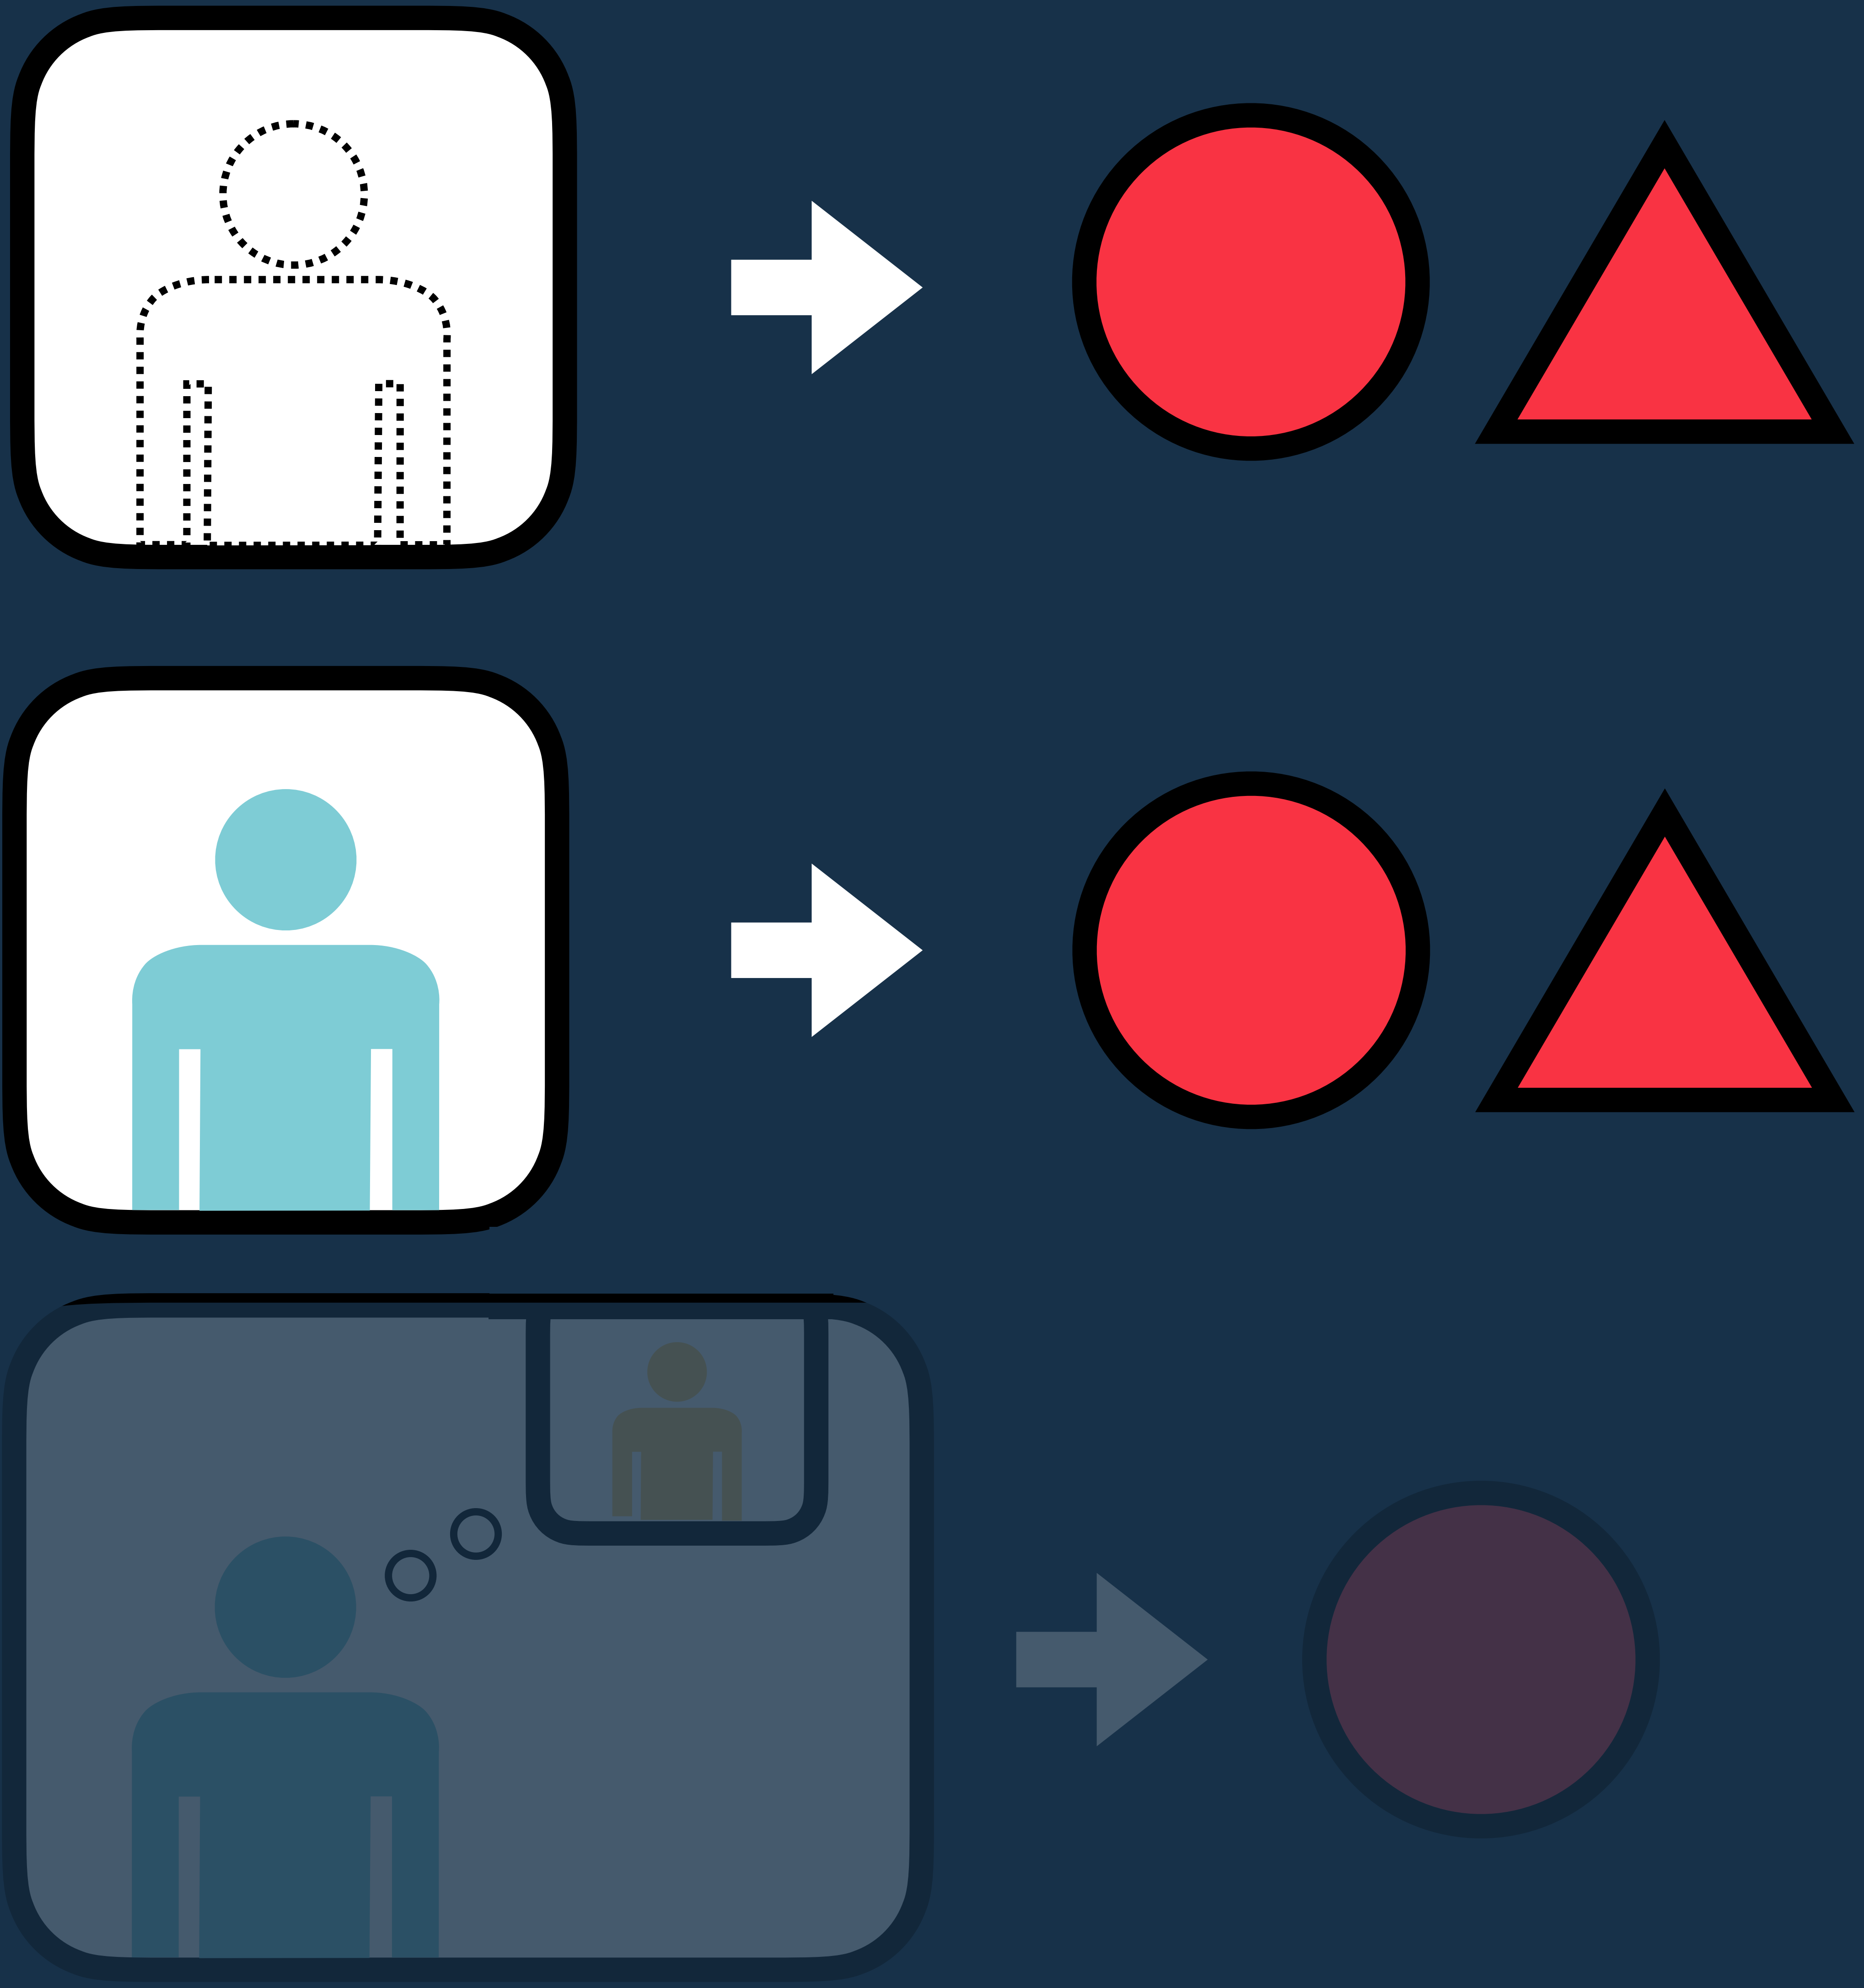

Observation #2: Individual differences in many-shot performance

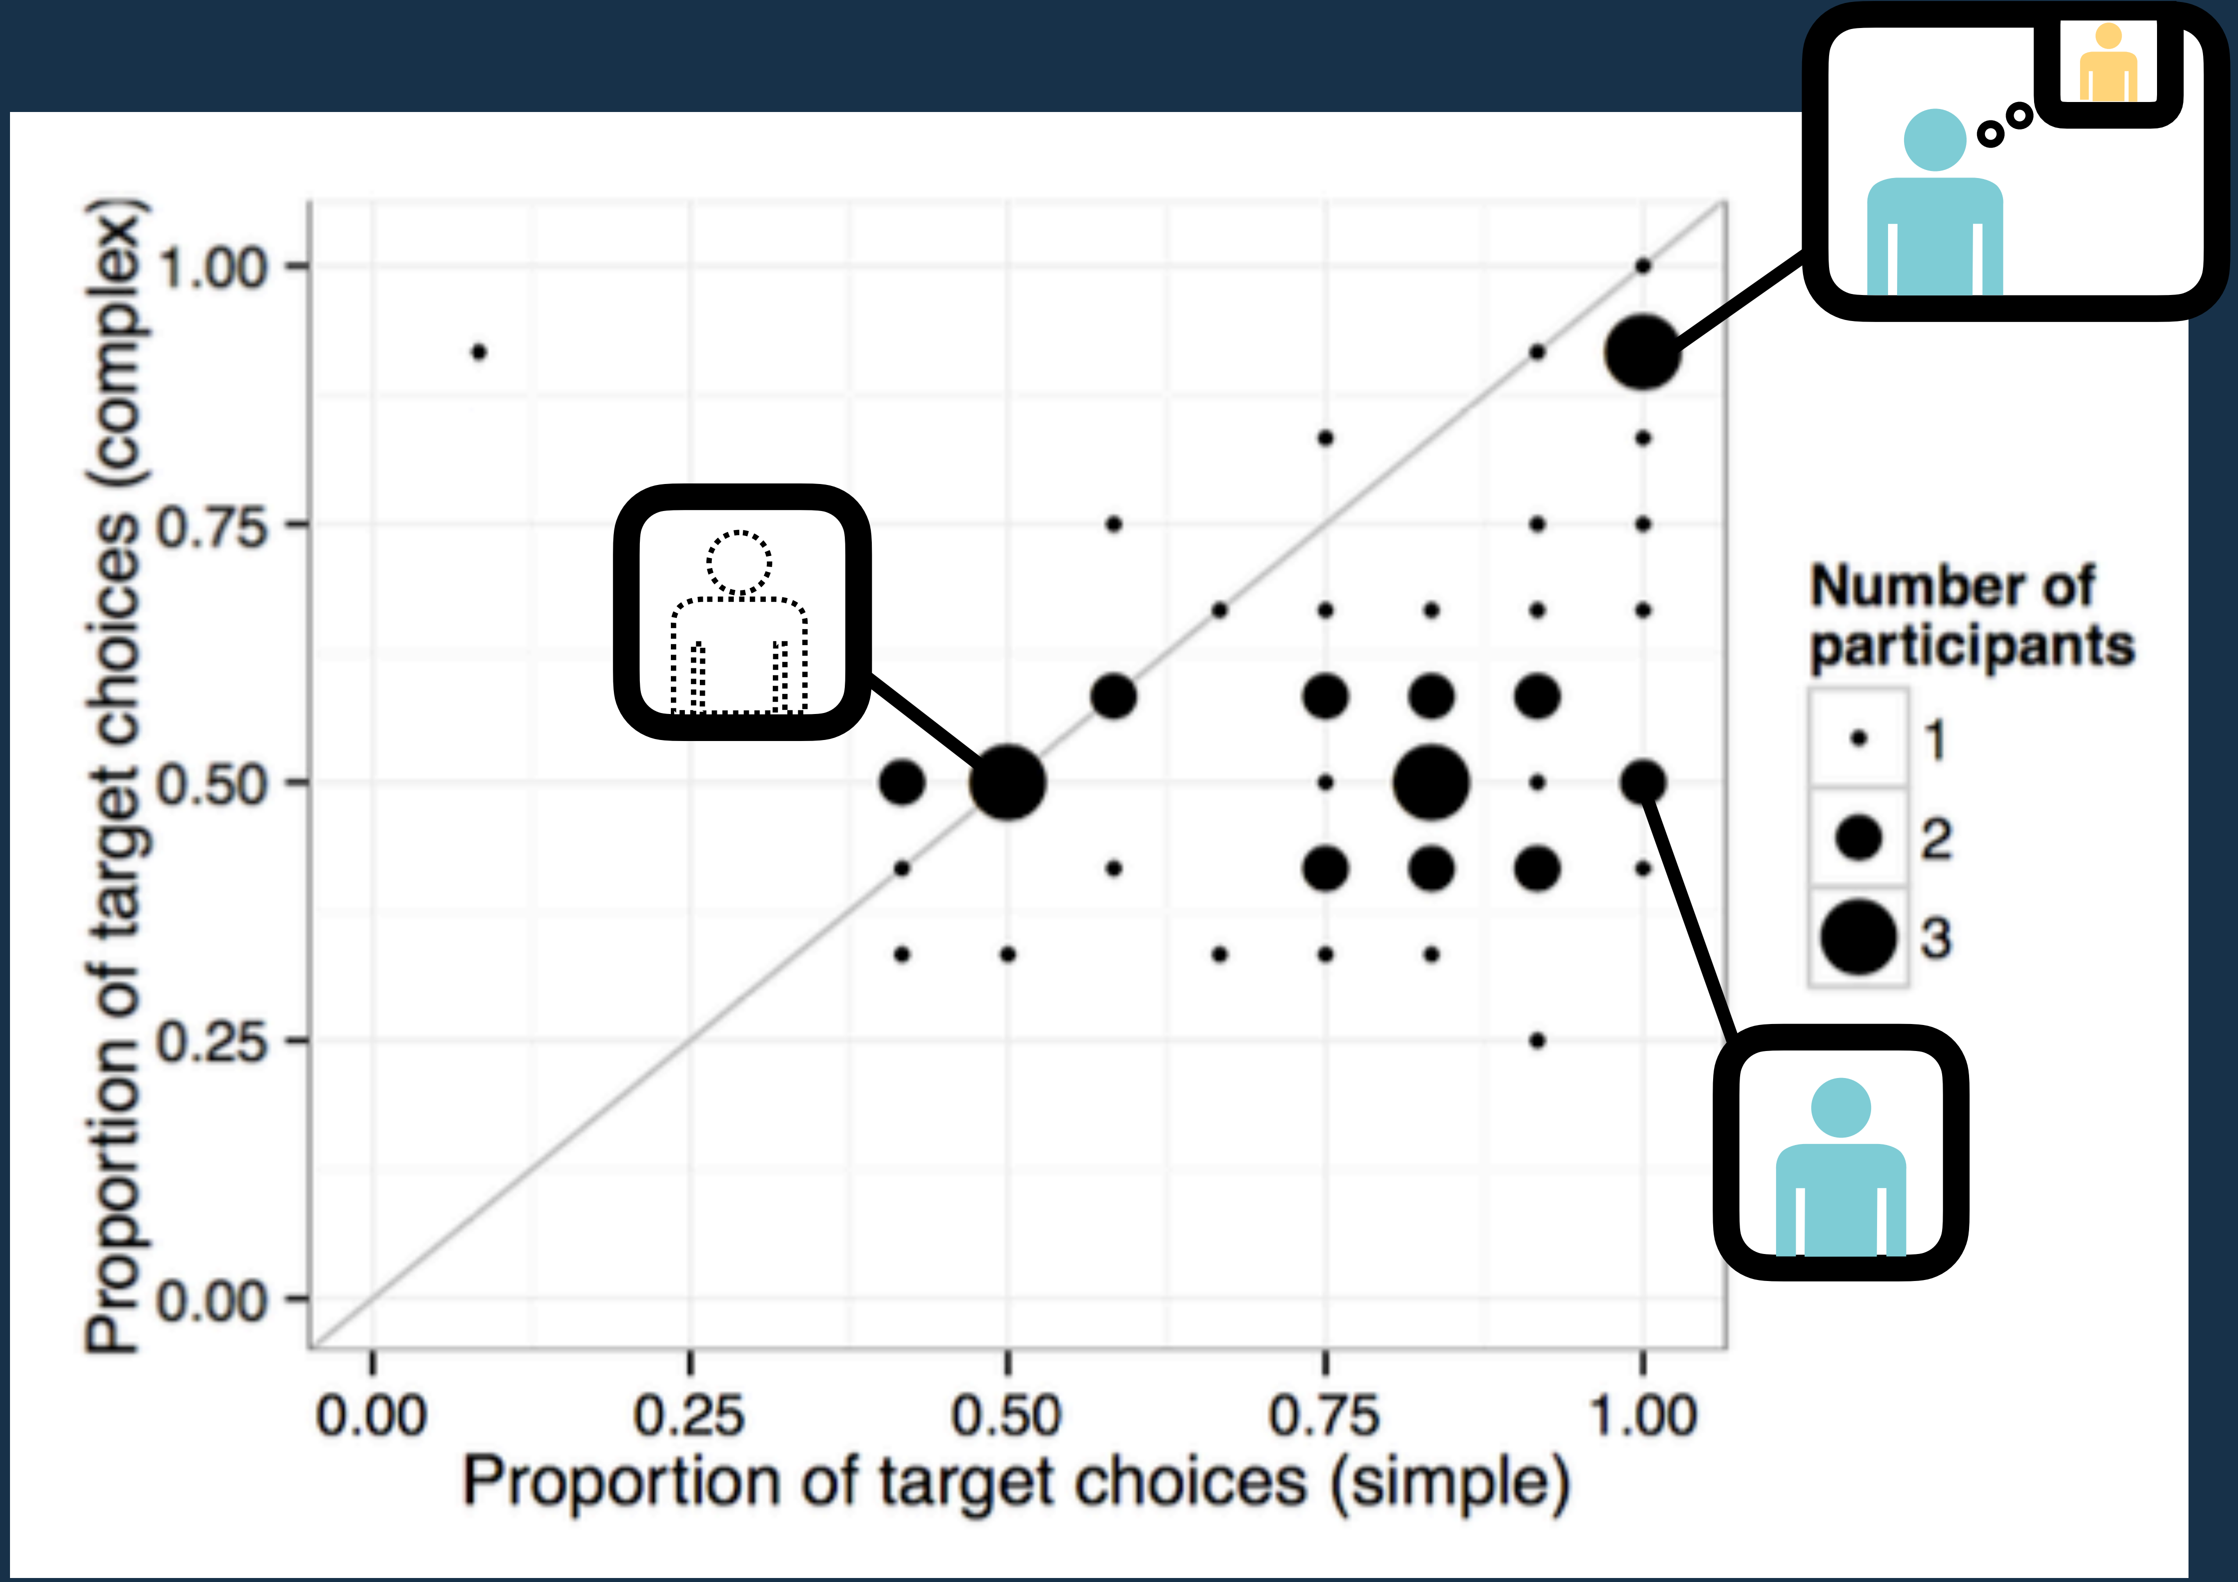

Franke & Degen (2016)

(n = 60, 12 obs/condition)

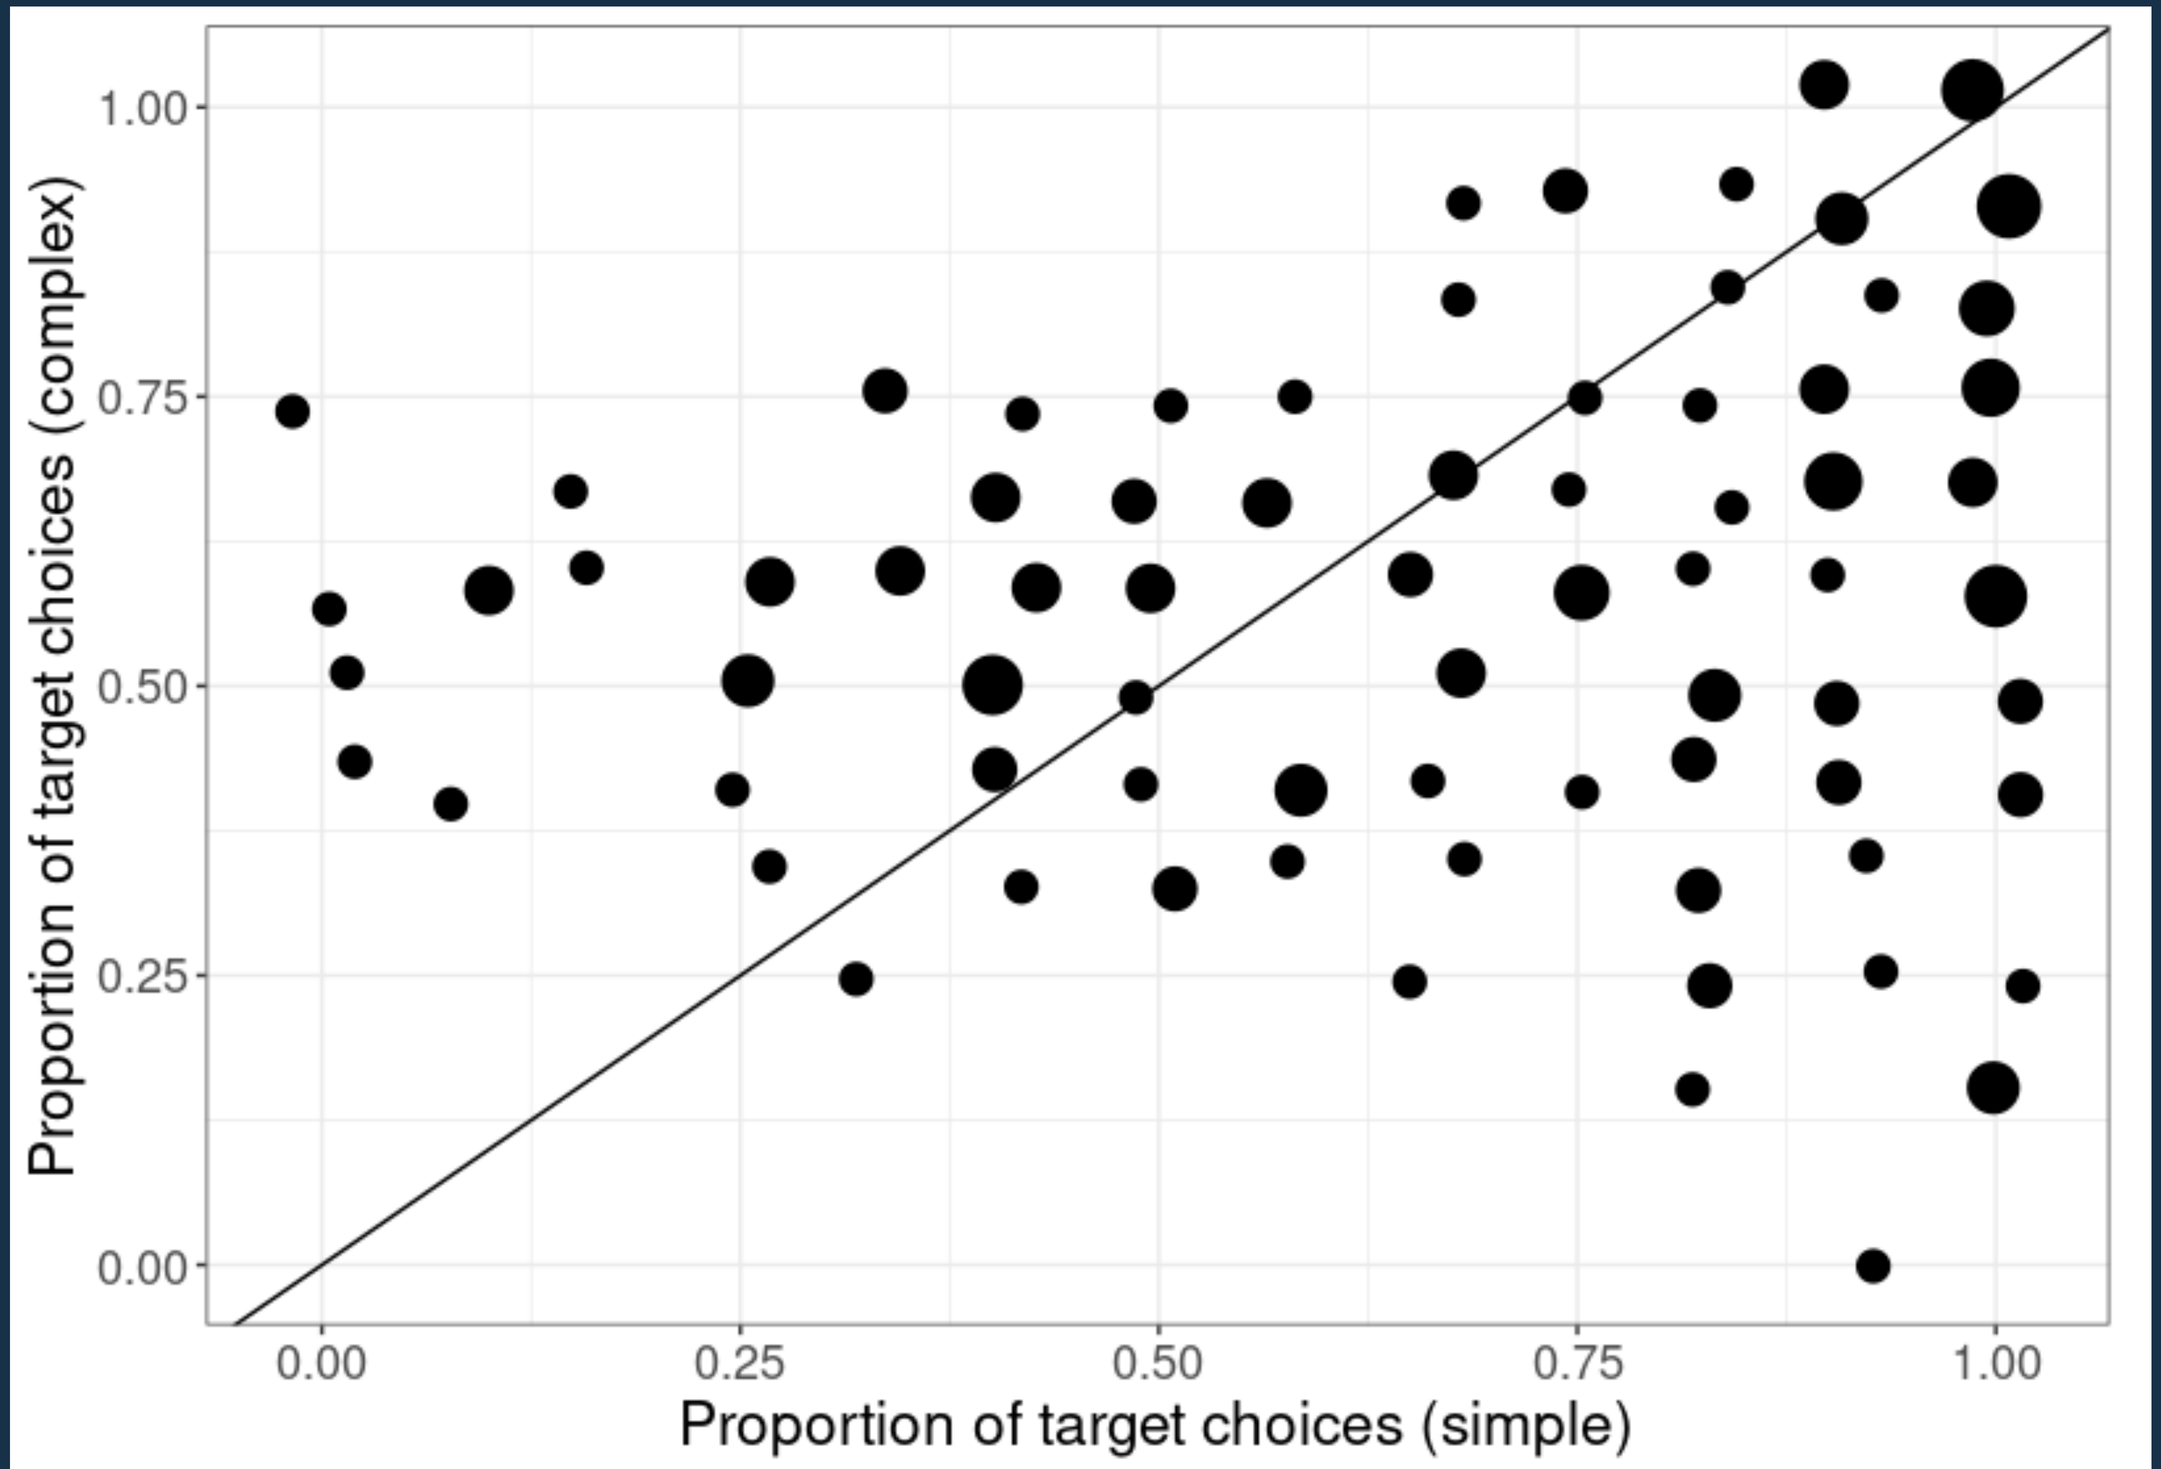

Mayn & Demberg (2023)

(n = 173, 12 obs/condition)

(debiased stimuli, cf. Mayn 2023)

Unexpected covariate: Reasoning performance

:= Raven's Matrices + Cognitive Reflection Task

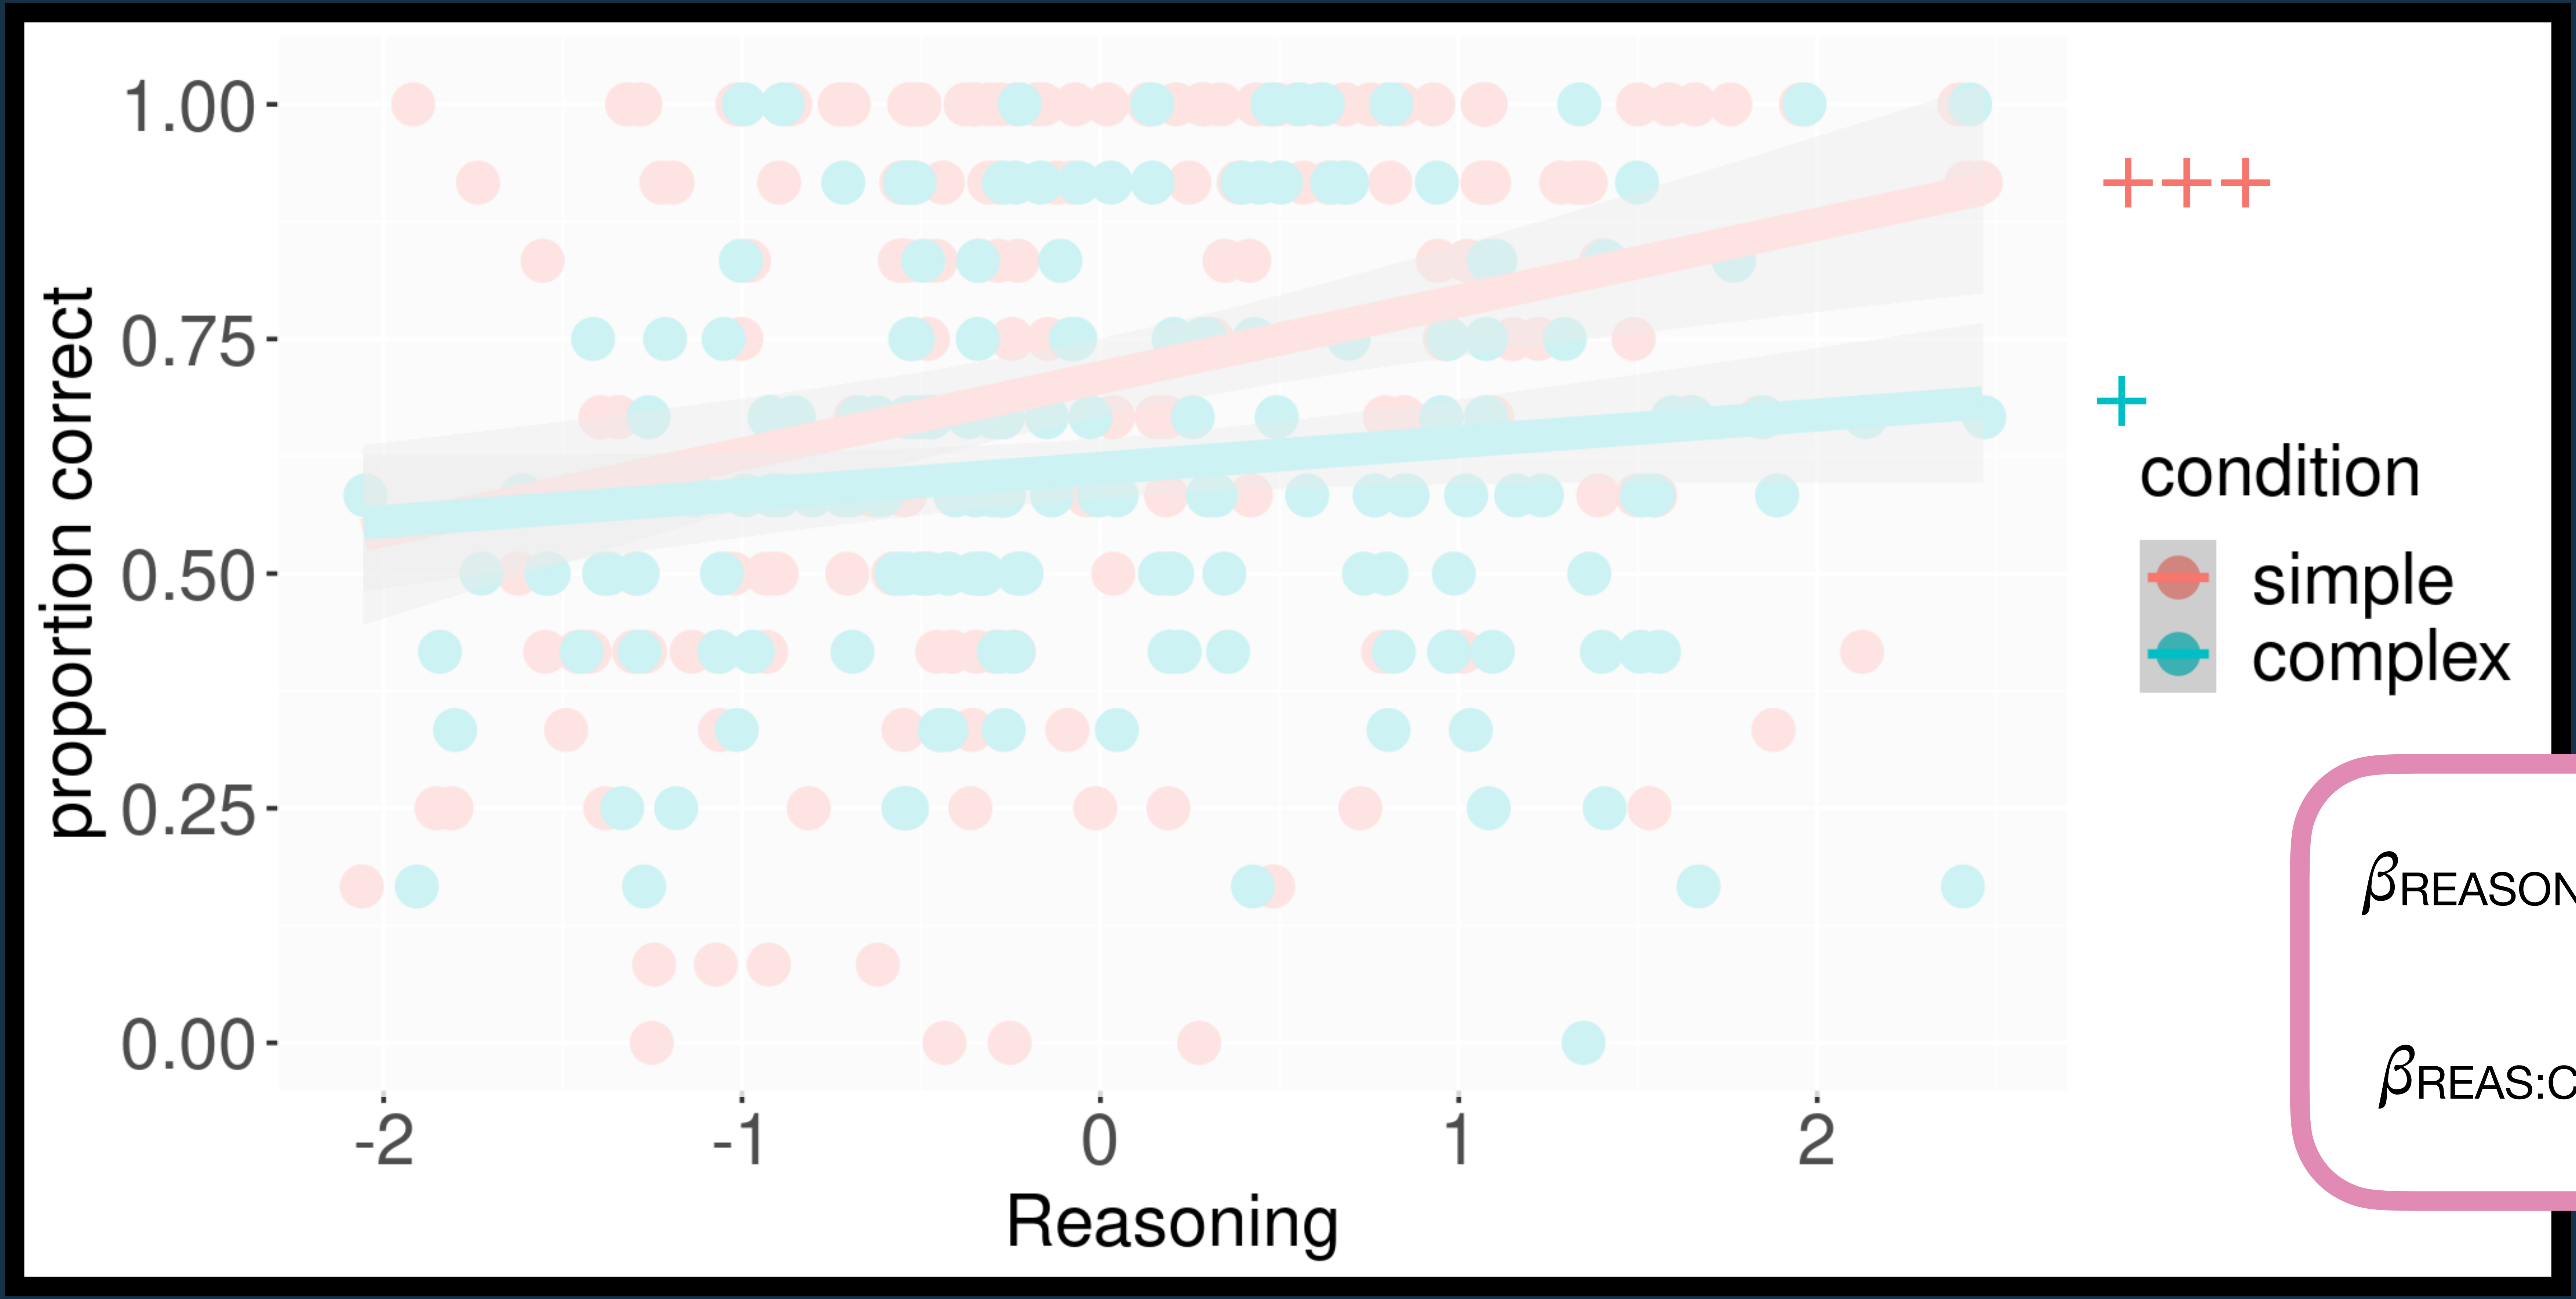

$\beta_{\text{REASONING}}$ = $(0.00, 0.31)_{95\%}$

qualified by

$\beta_{\text{REAS:COND}}$= $(0.21, 0.85)_{95\%}$

(Mayn & Demberg 2023)

(also Theory of Mind, but not Working Memory)

Raven’s Matrices

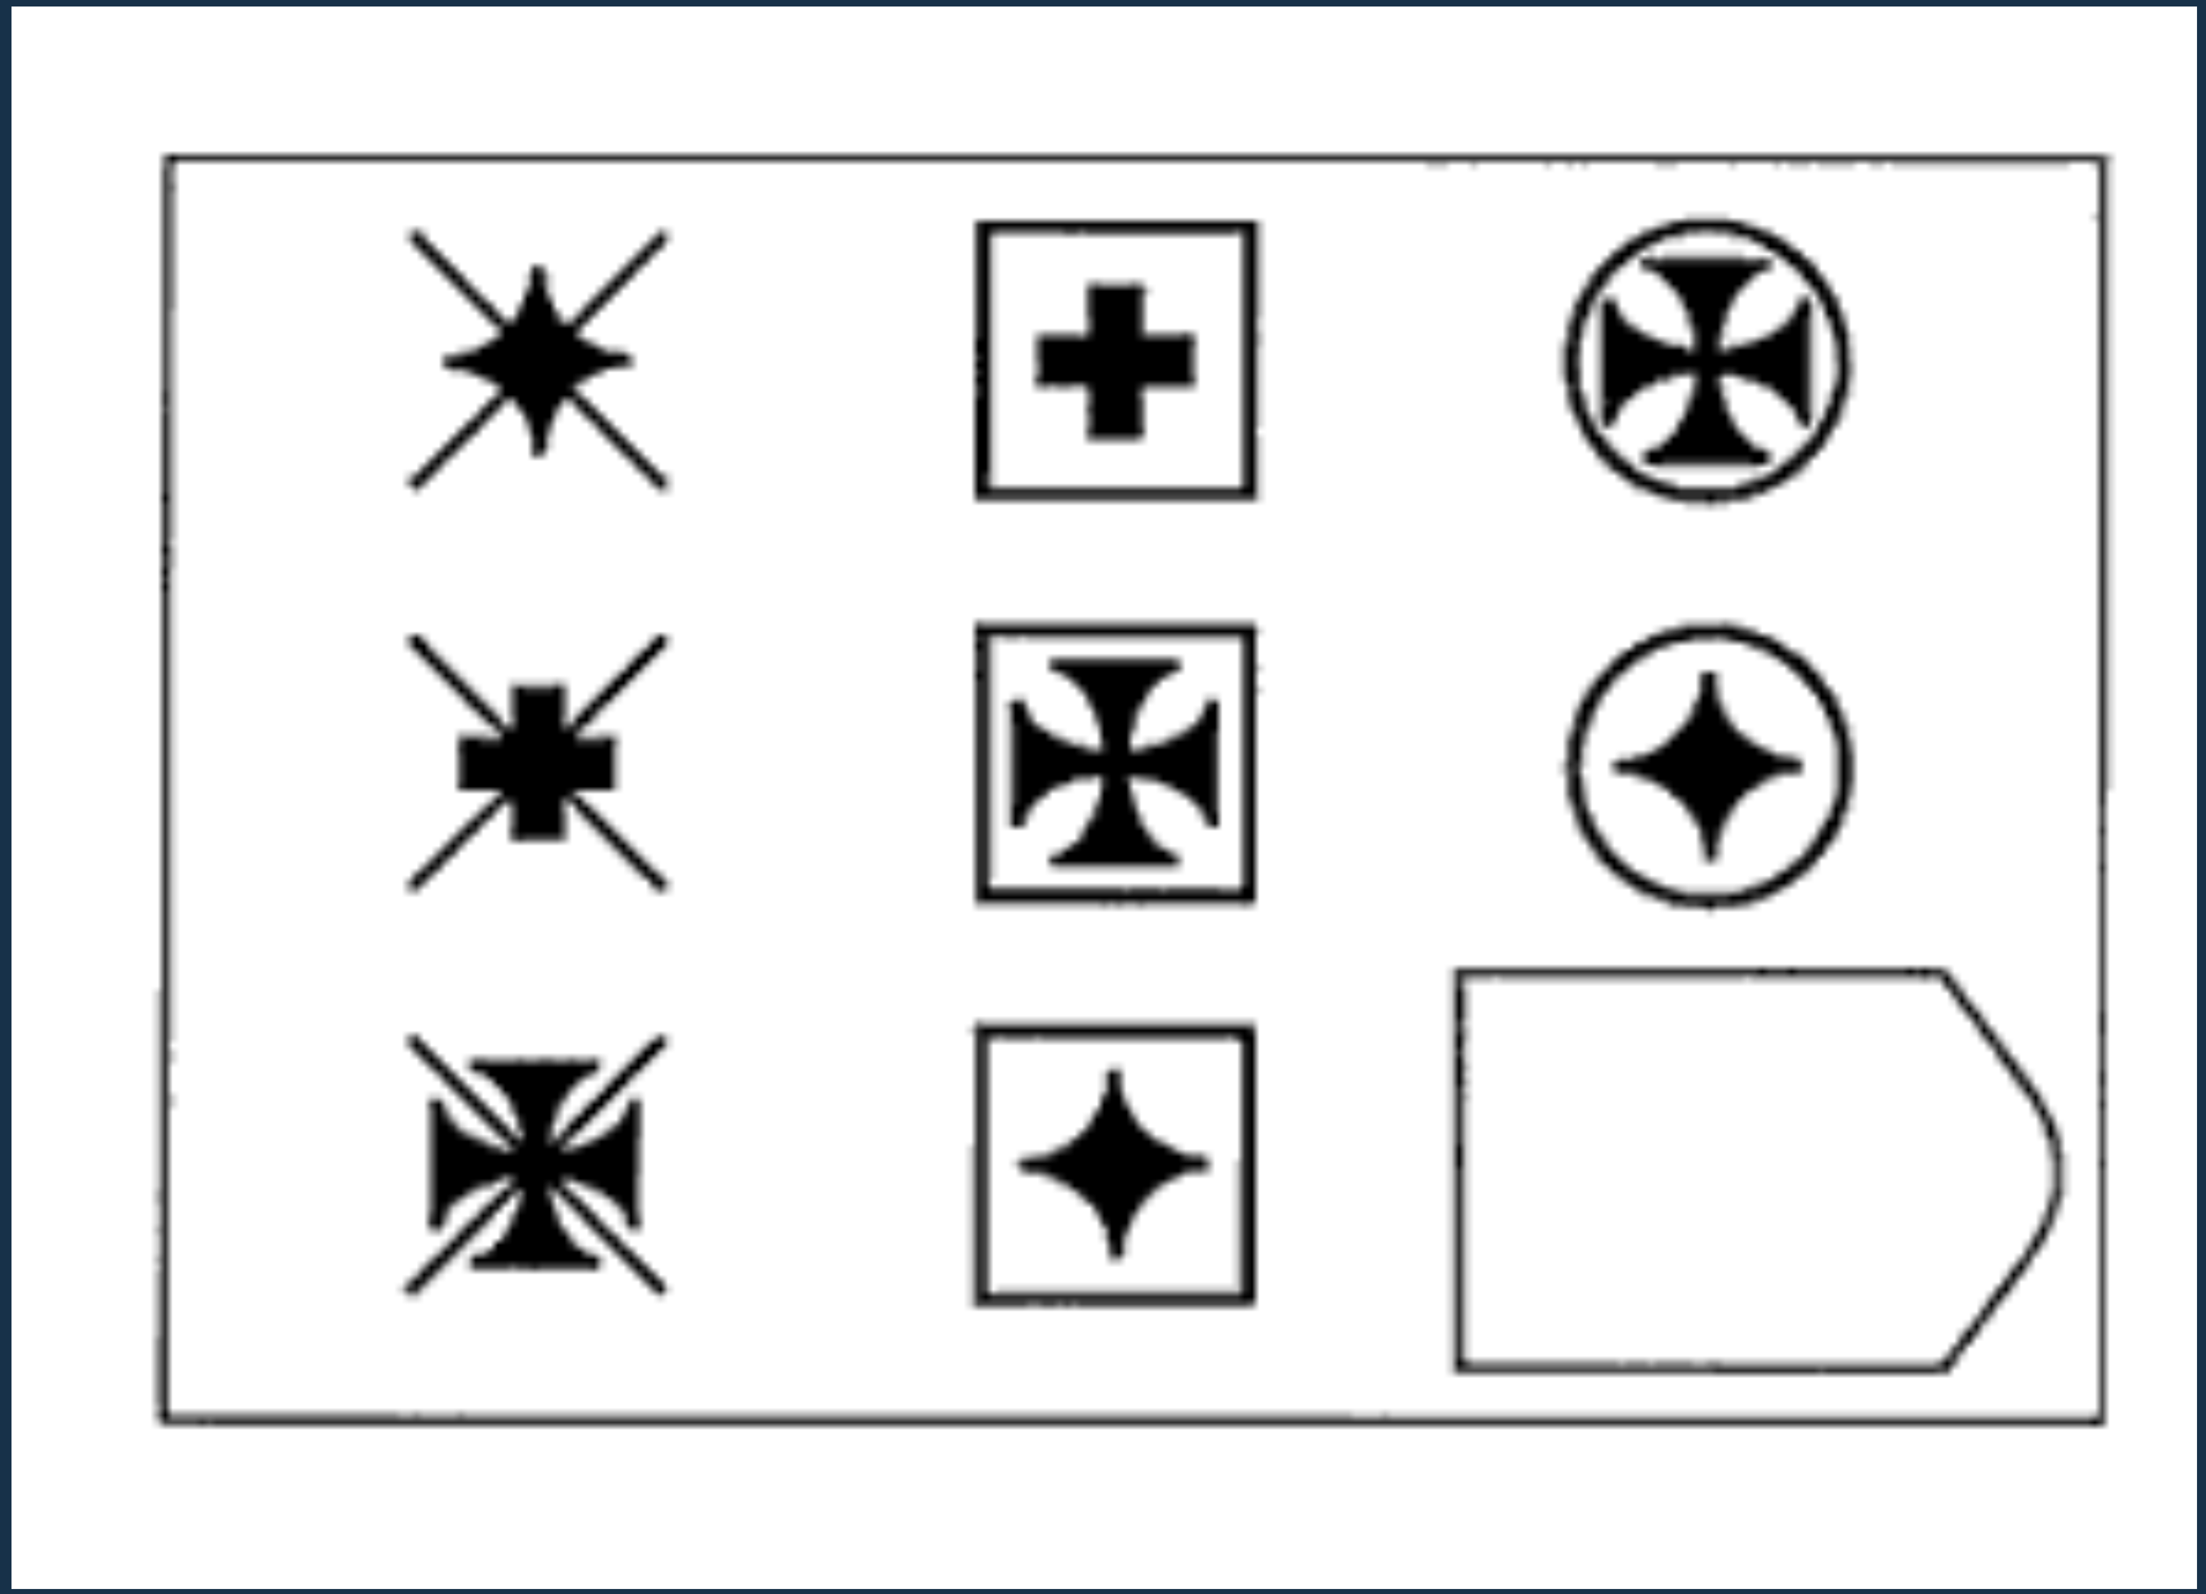

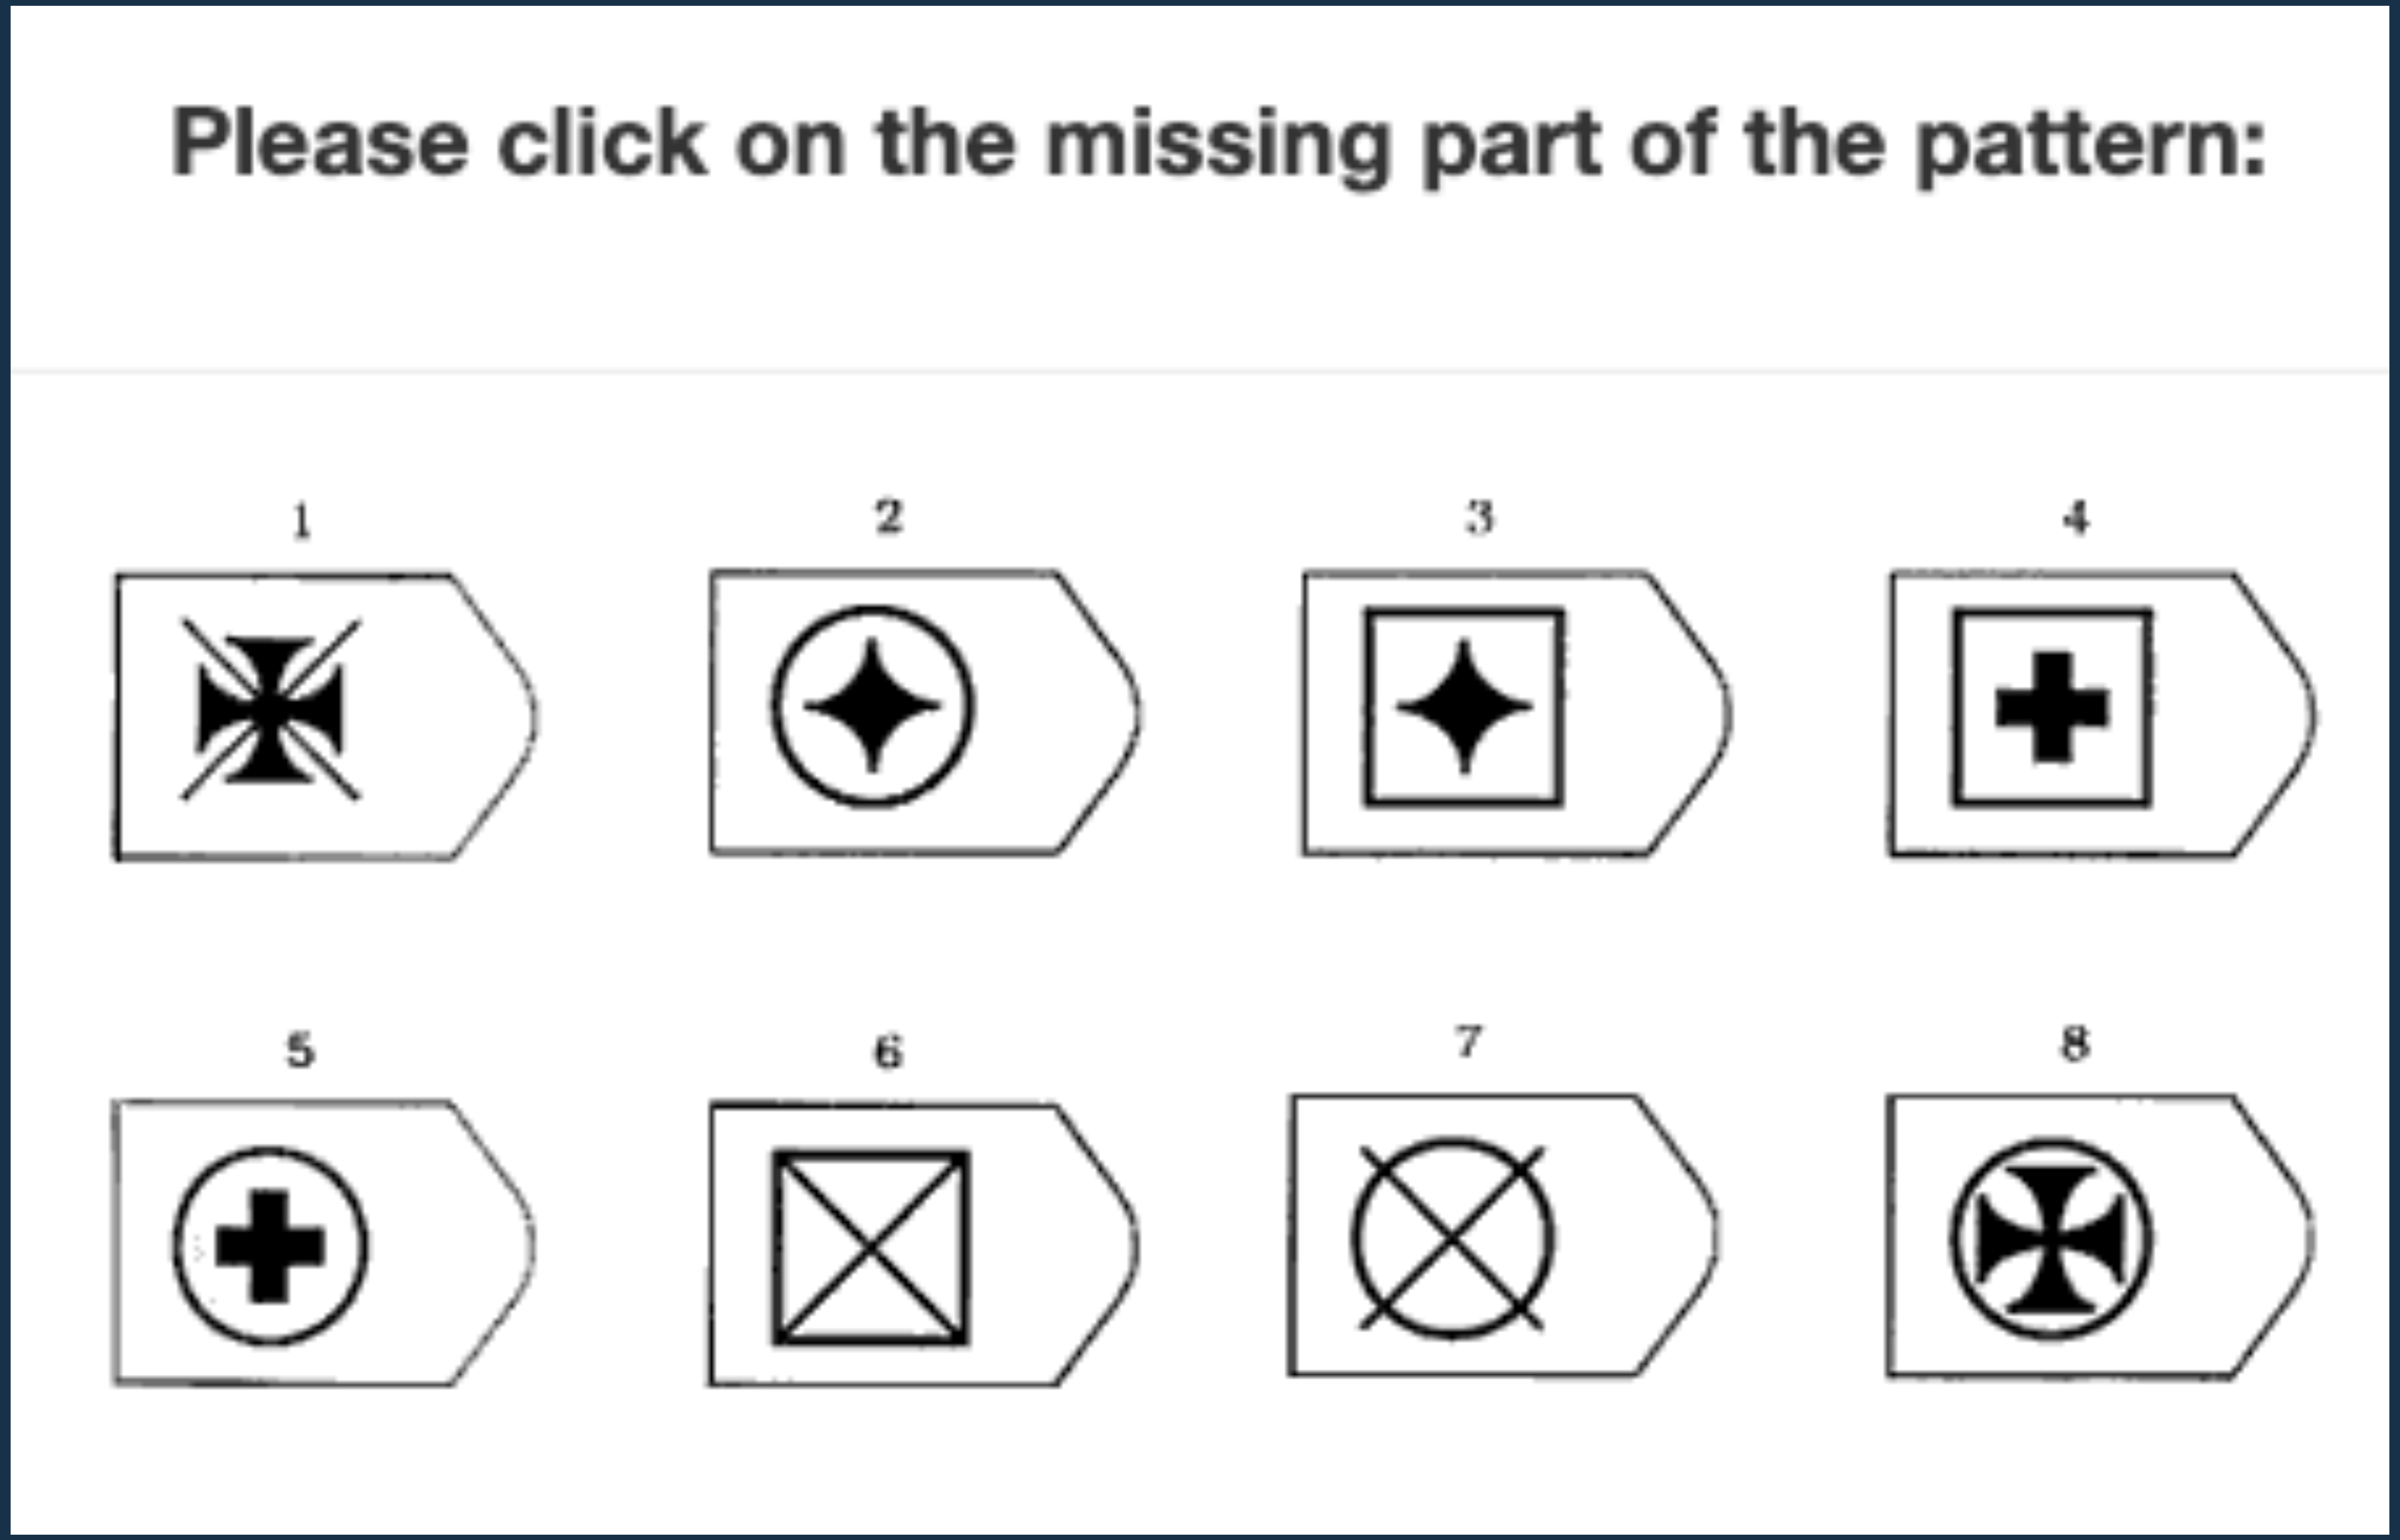

Success requires **efficient pattern induction** in a large hypothesis space.

(Carpenter et al. 1990, Gonthier & Thomassin 2015, Gonthier & Roulin 2020, Stocco et al. 2021)

Modeling individual differences in Raven's

Stocco et al. (2021):

ACT-R model for Raven's performance as rule induction via exploration and reinforcement learning

individually parameterized by:

persistence

(Eisenberger & Leonard 1980)

neg. feedback strength (F_{NEG})

(Frank et al. 2004)

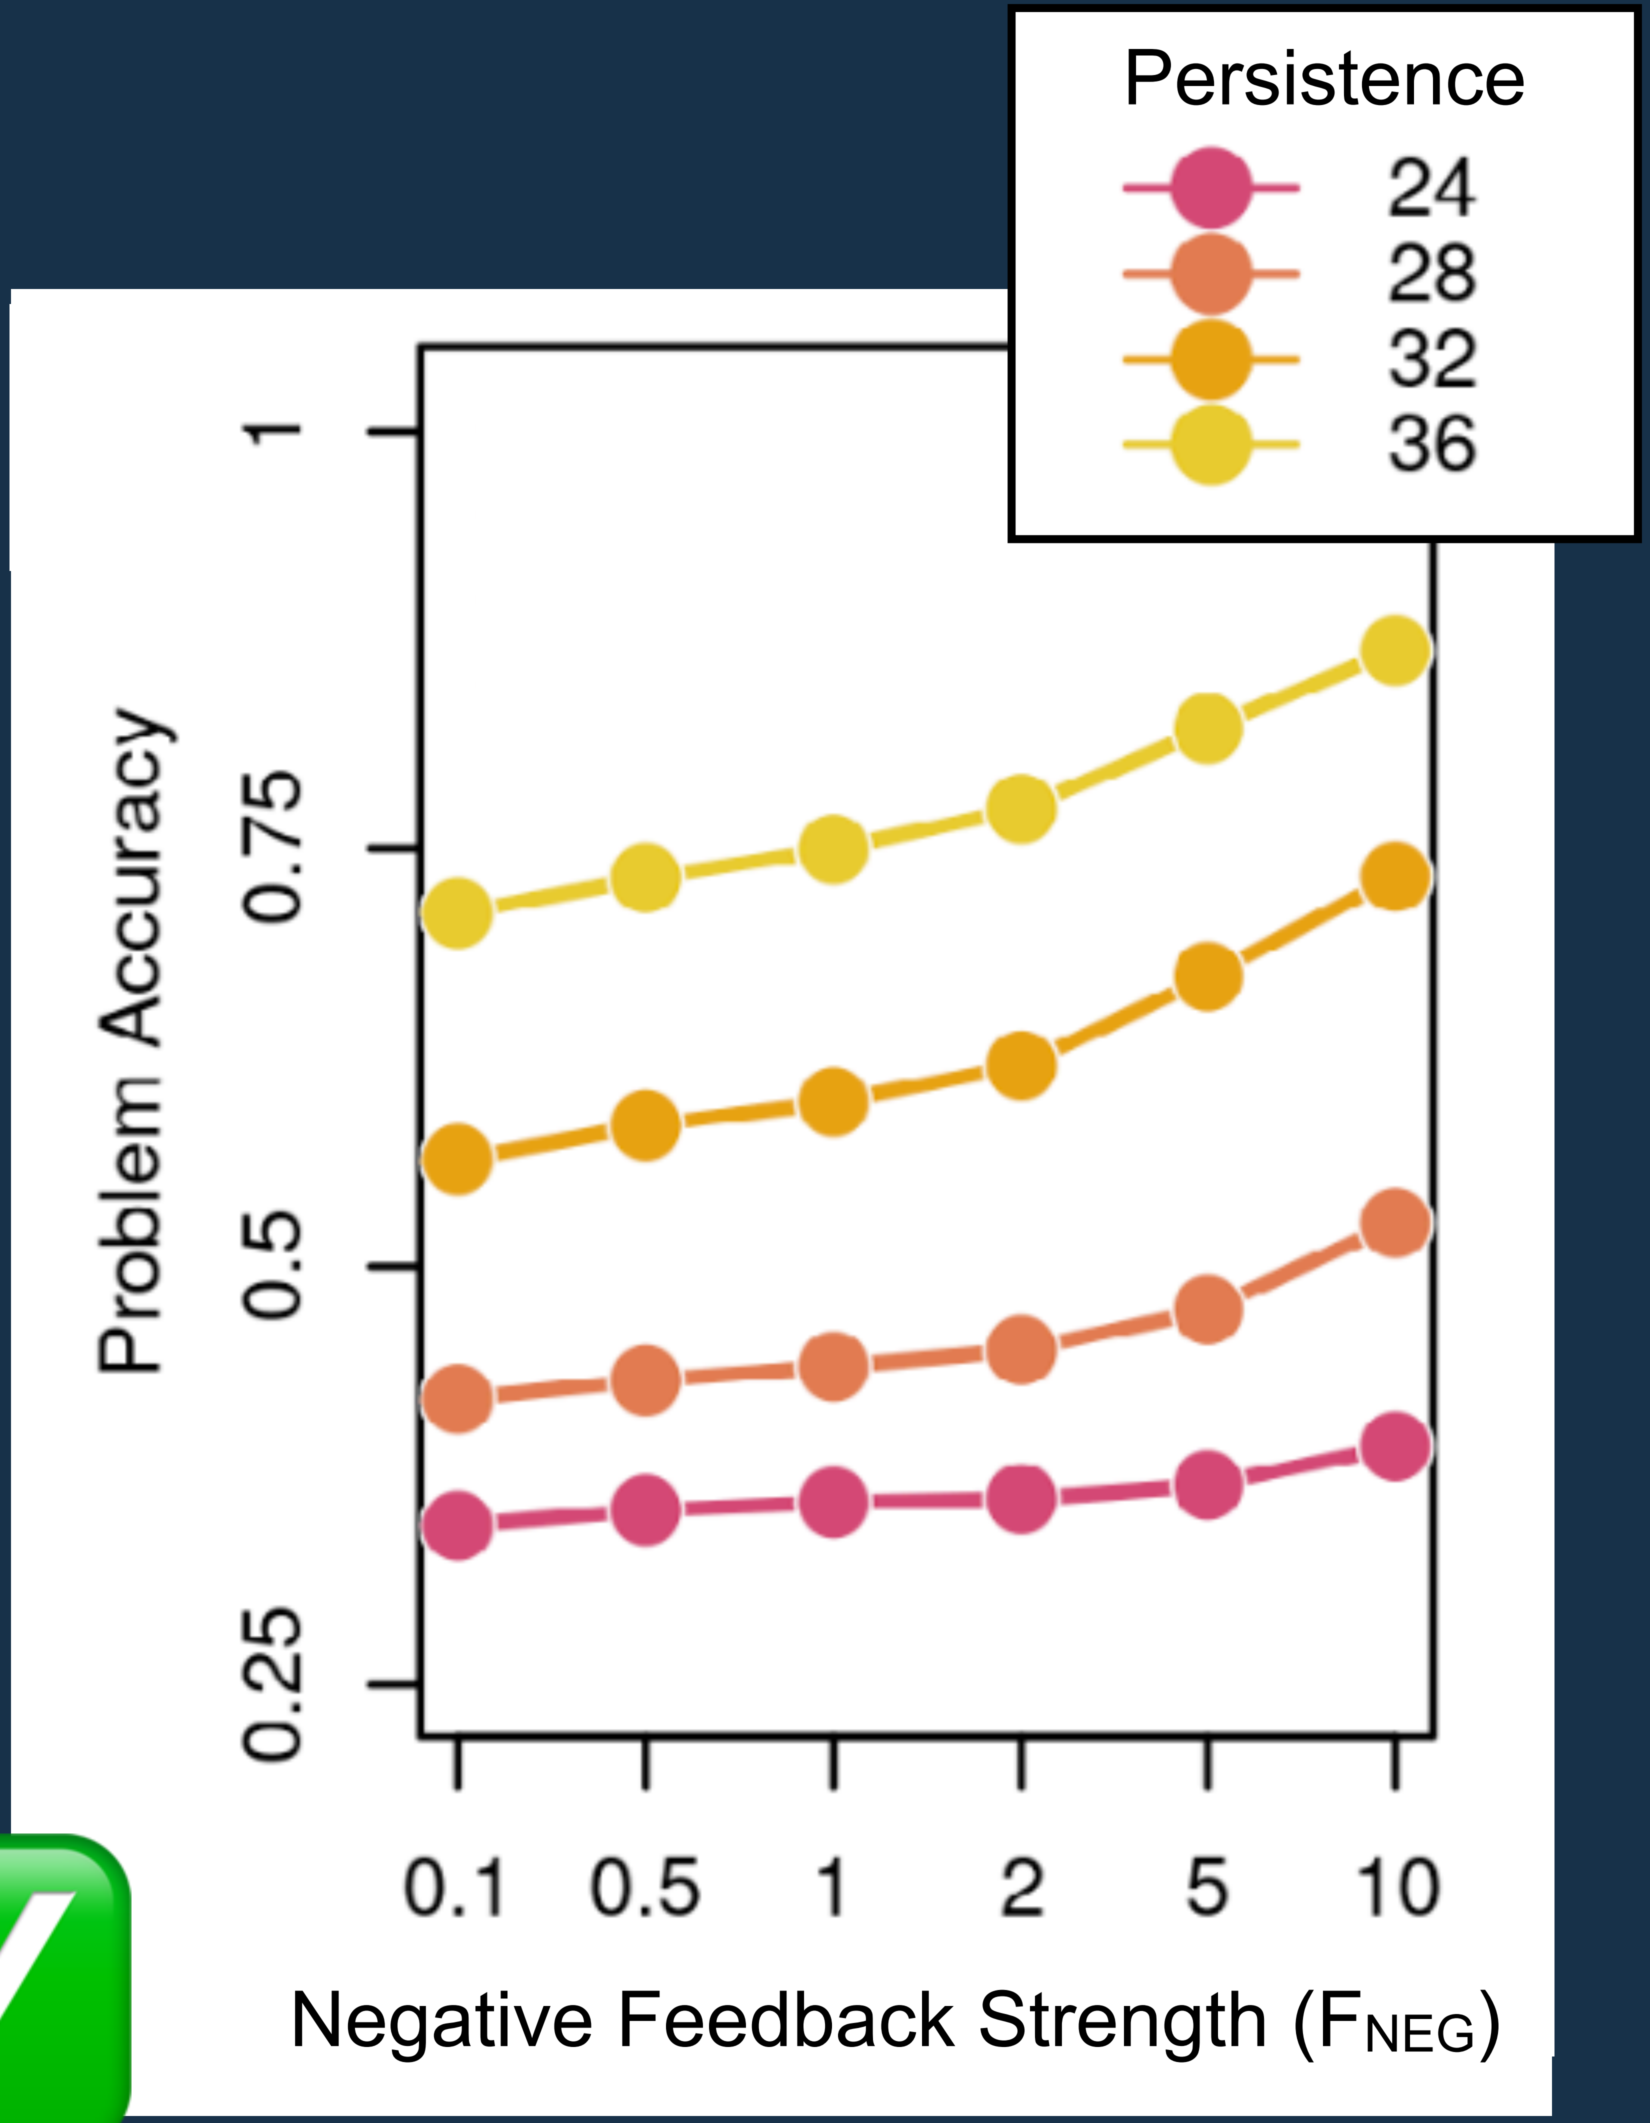

Roadmap

1. Background
2. **Our ACT-R model**
3. Modeling individual differences across tasks

RefGame as exploration

(implemented in pyactr: Brasoveanu & Dotlačil 2020)

- Attempt literal interpretation
  - Check informativity (number of matches)
  - If informative (1 match), select match
  - Else, penalize utility with F_{NEG}, return to…
- Select highest-utility strategy (with noise)
  - If already checked, penalize utility with F_{NEG}
  - Else, evaluate; select or return again
- If time ever exceeds persistence (τ), guess

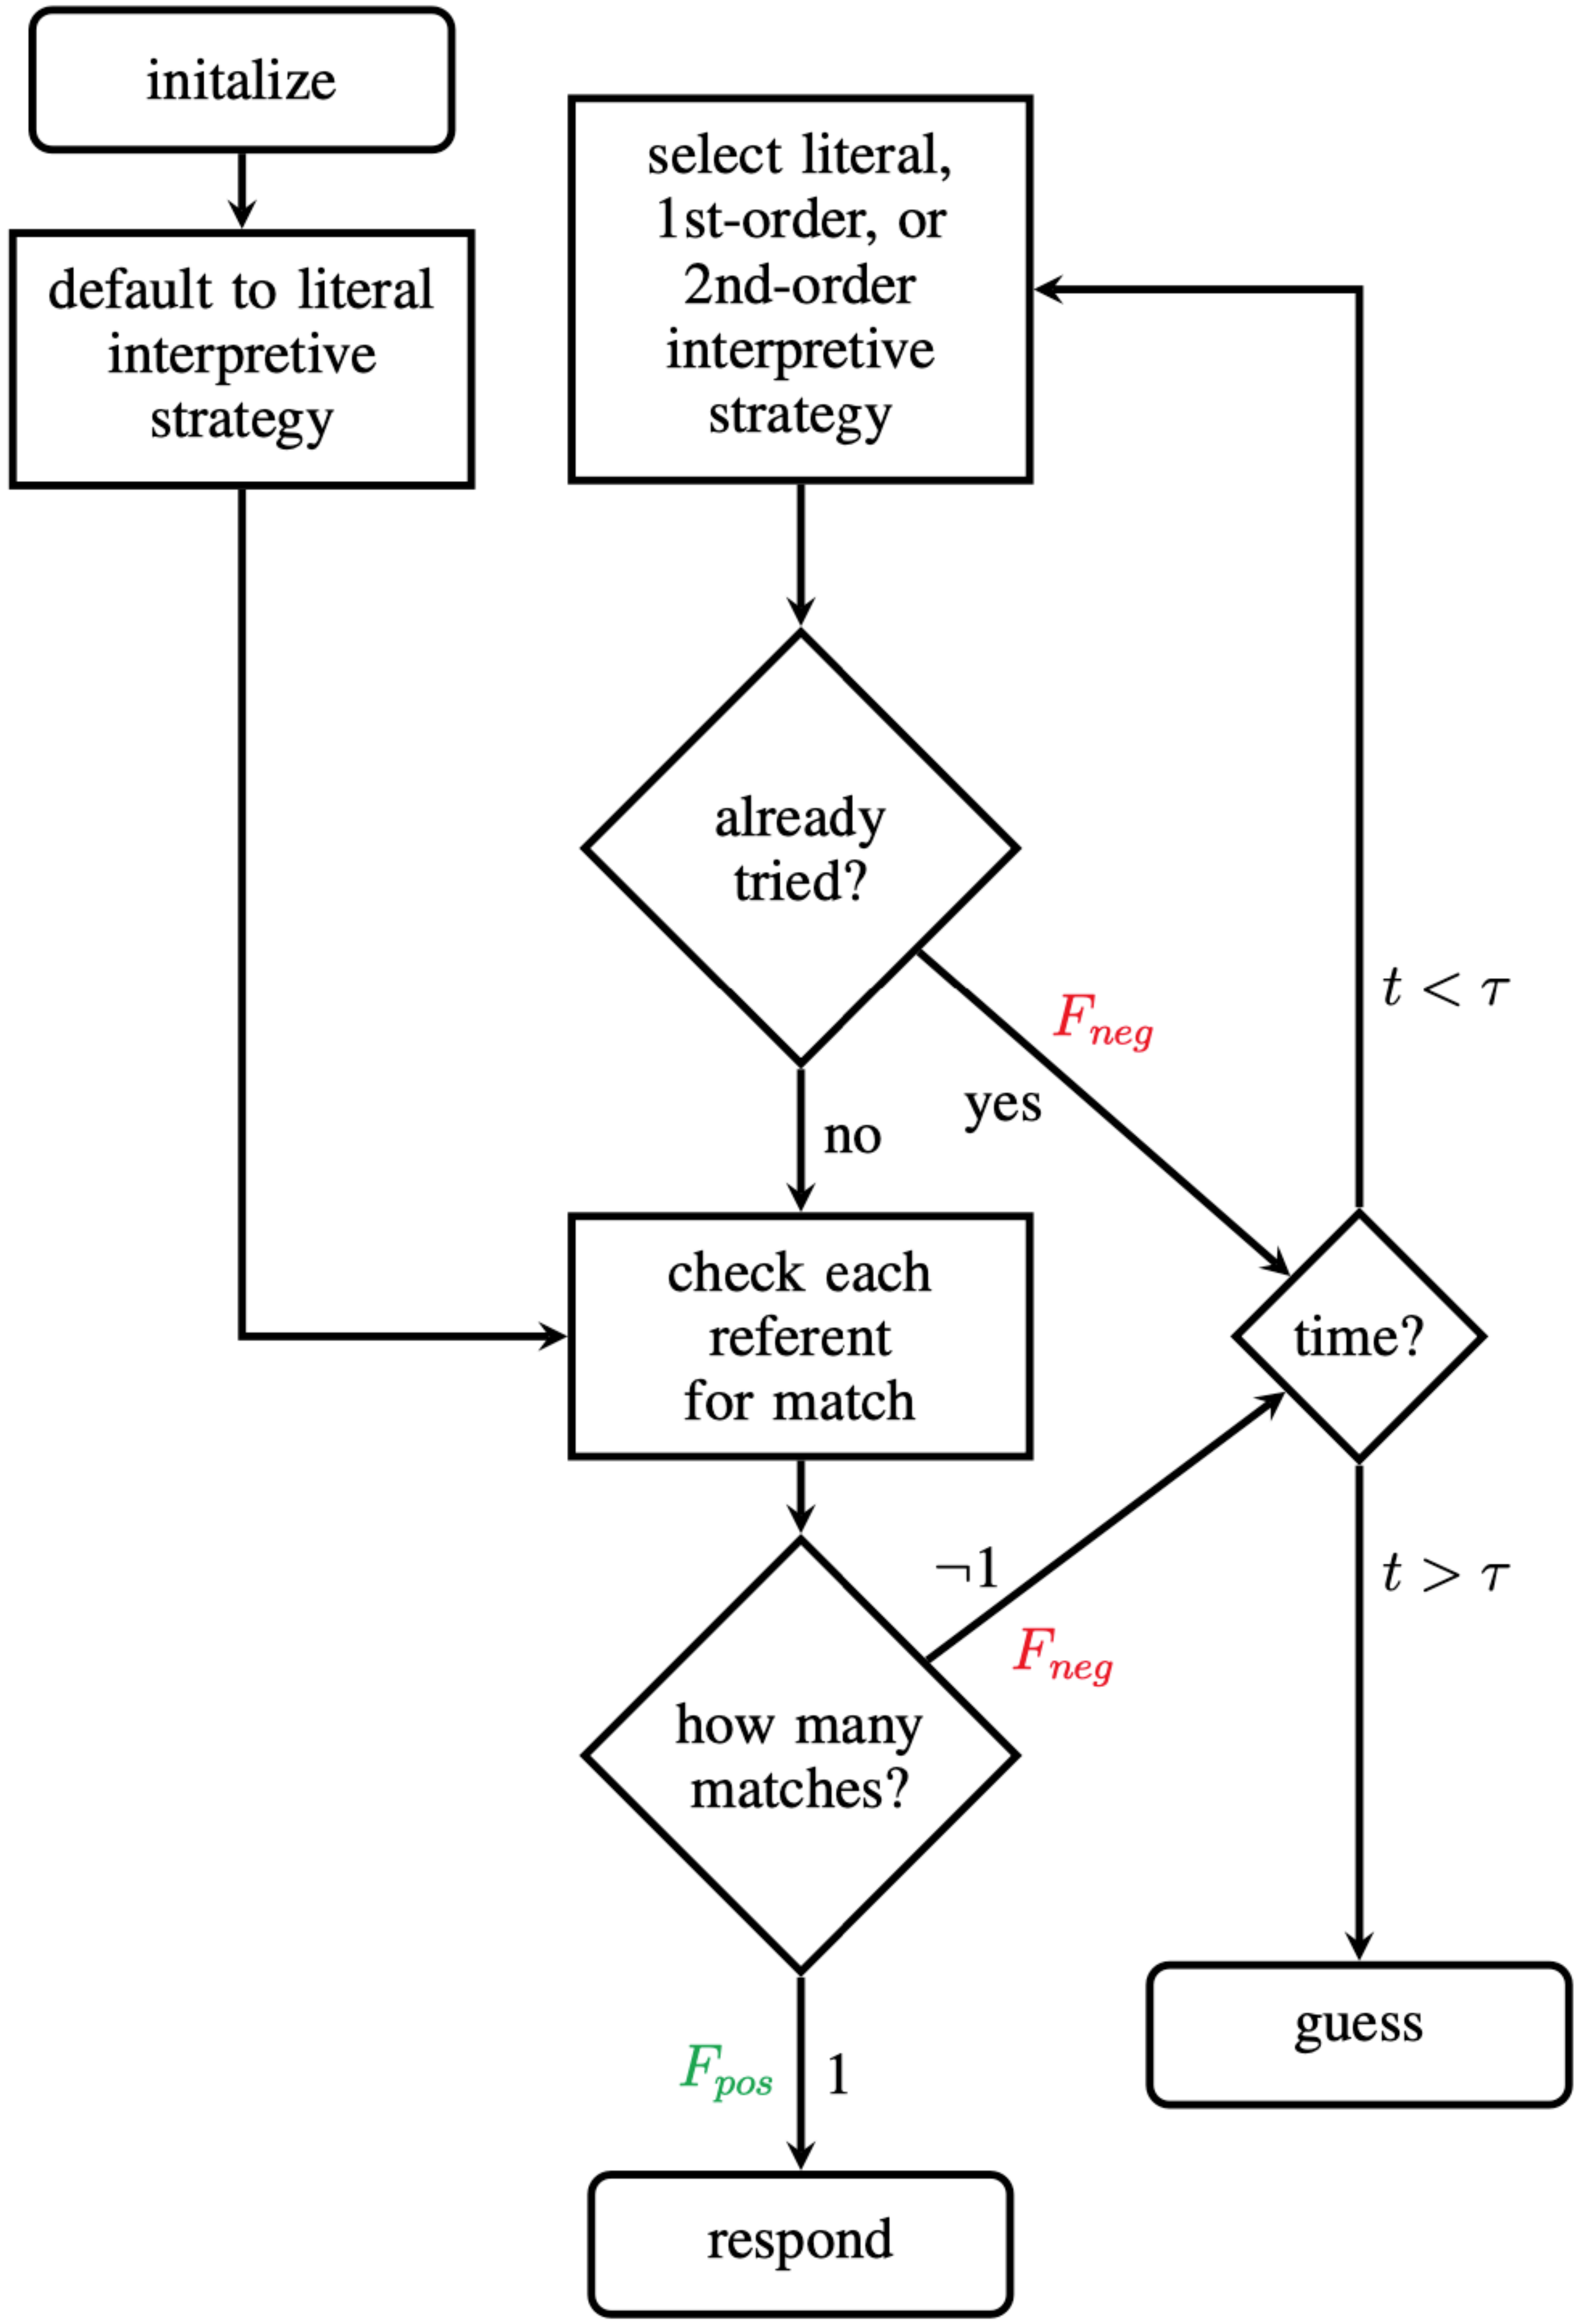

Model experiment

- Simulated task: Randomized 36-trial RefGame (16 trivial, 8 simple, 8 complex)
- Simulated participants: 10 persistence values x 20 F_{NEG} values, 25 per cell
- Critical strategy utilities begin as a fixed stair-step

Literal: 5 First-Order: -2.5 Second-Order: -5

Learning-related individual differences

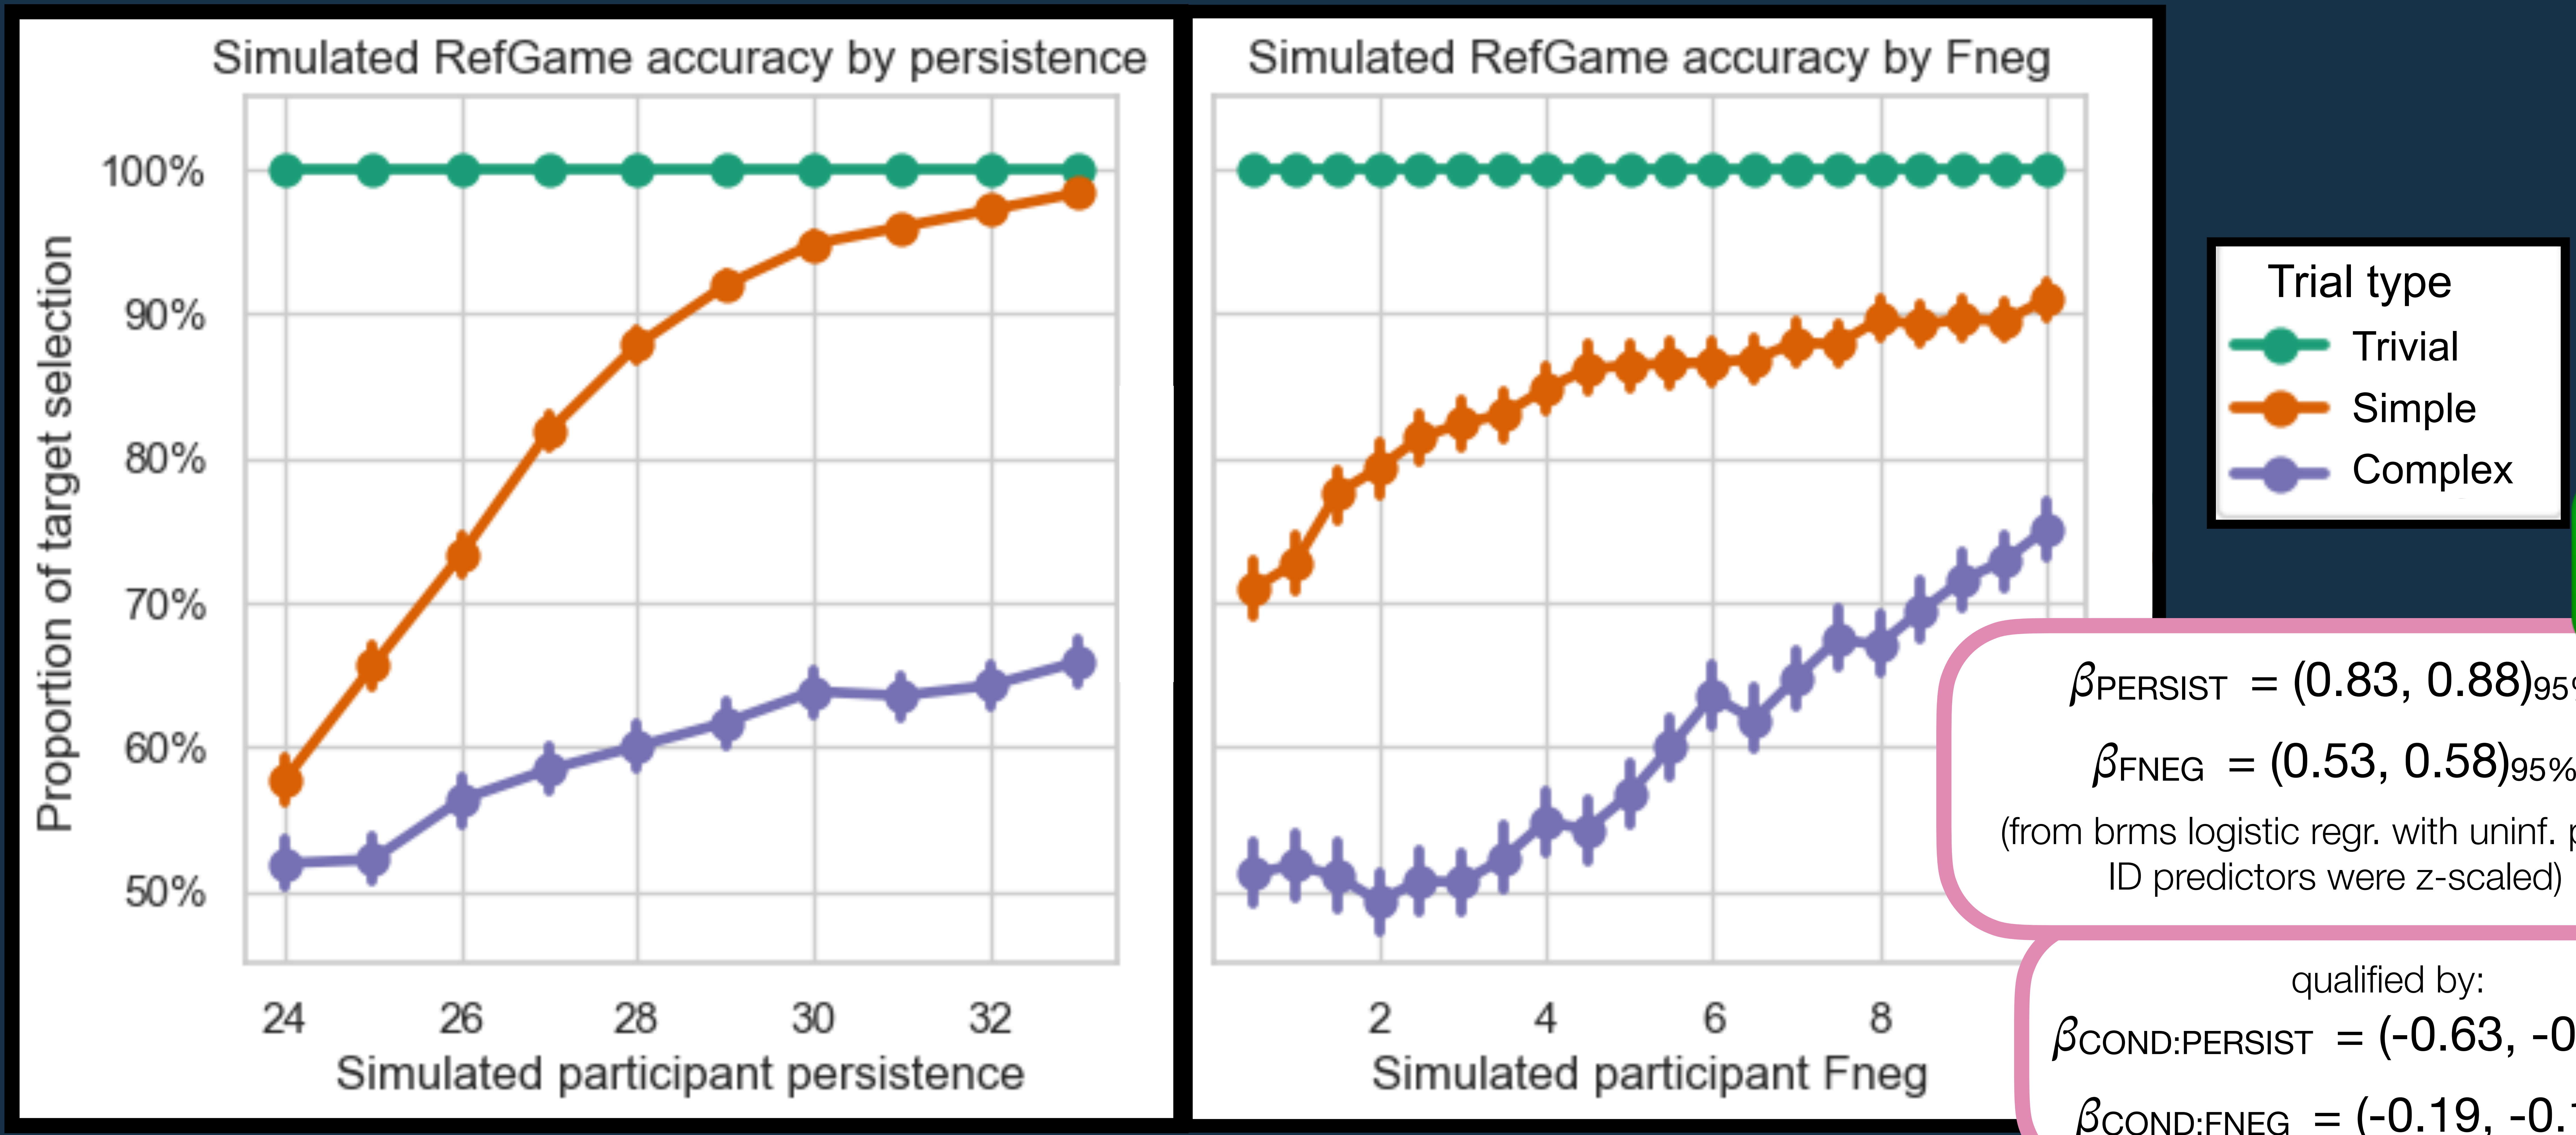

$\beta_{\text{PERSIST}} = (0.83, 0.88)_{95\%}$

$\beta_{\text{FNEG}} = (0.53, 0.58)_{95\%}$

(from brms logistic regr. with uninf. priors, ID predictors were z-scaled)

qualified by:

$\beta_{\text{COND:PERSIST}} = (-0.63, -0.59)_{95\%}$

$\beta_{\text{COND:FNEG}} = (-0.19, -0.15)_{95\%}$

Predicted learning behavior

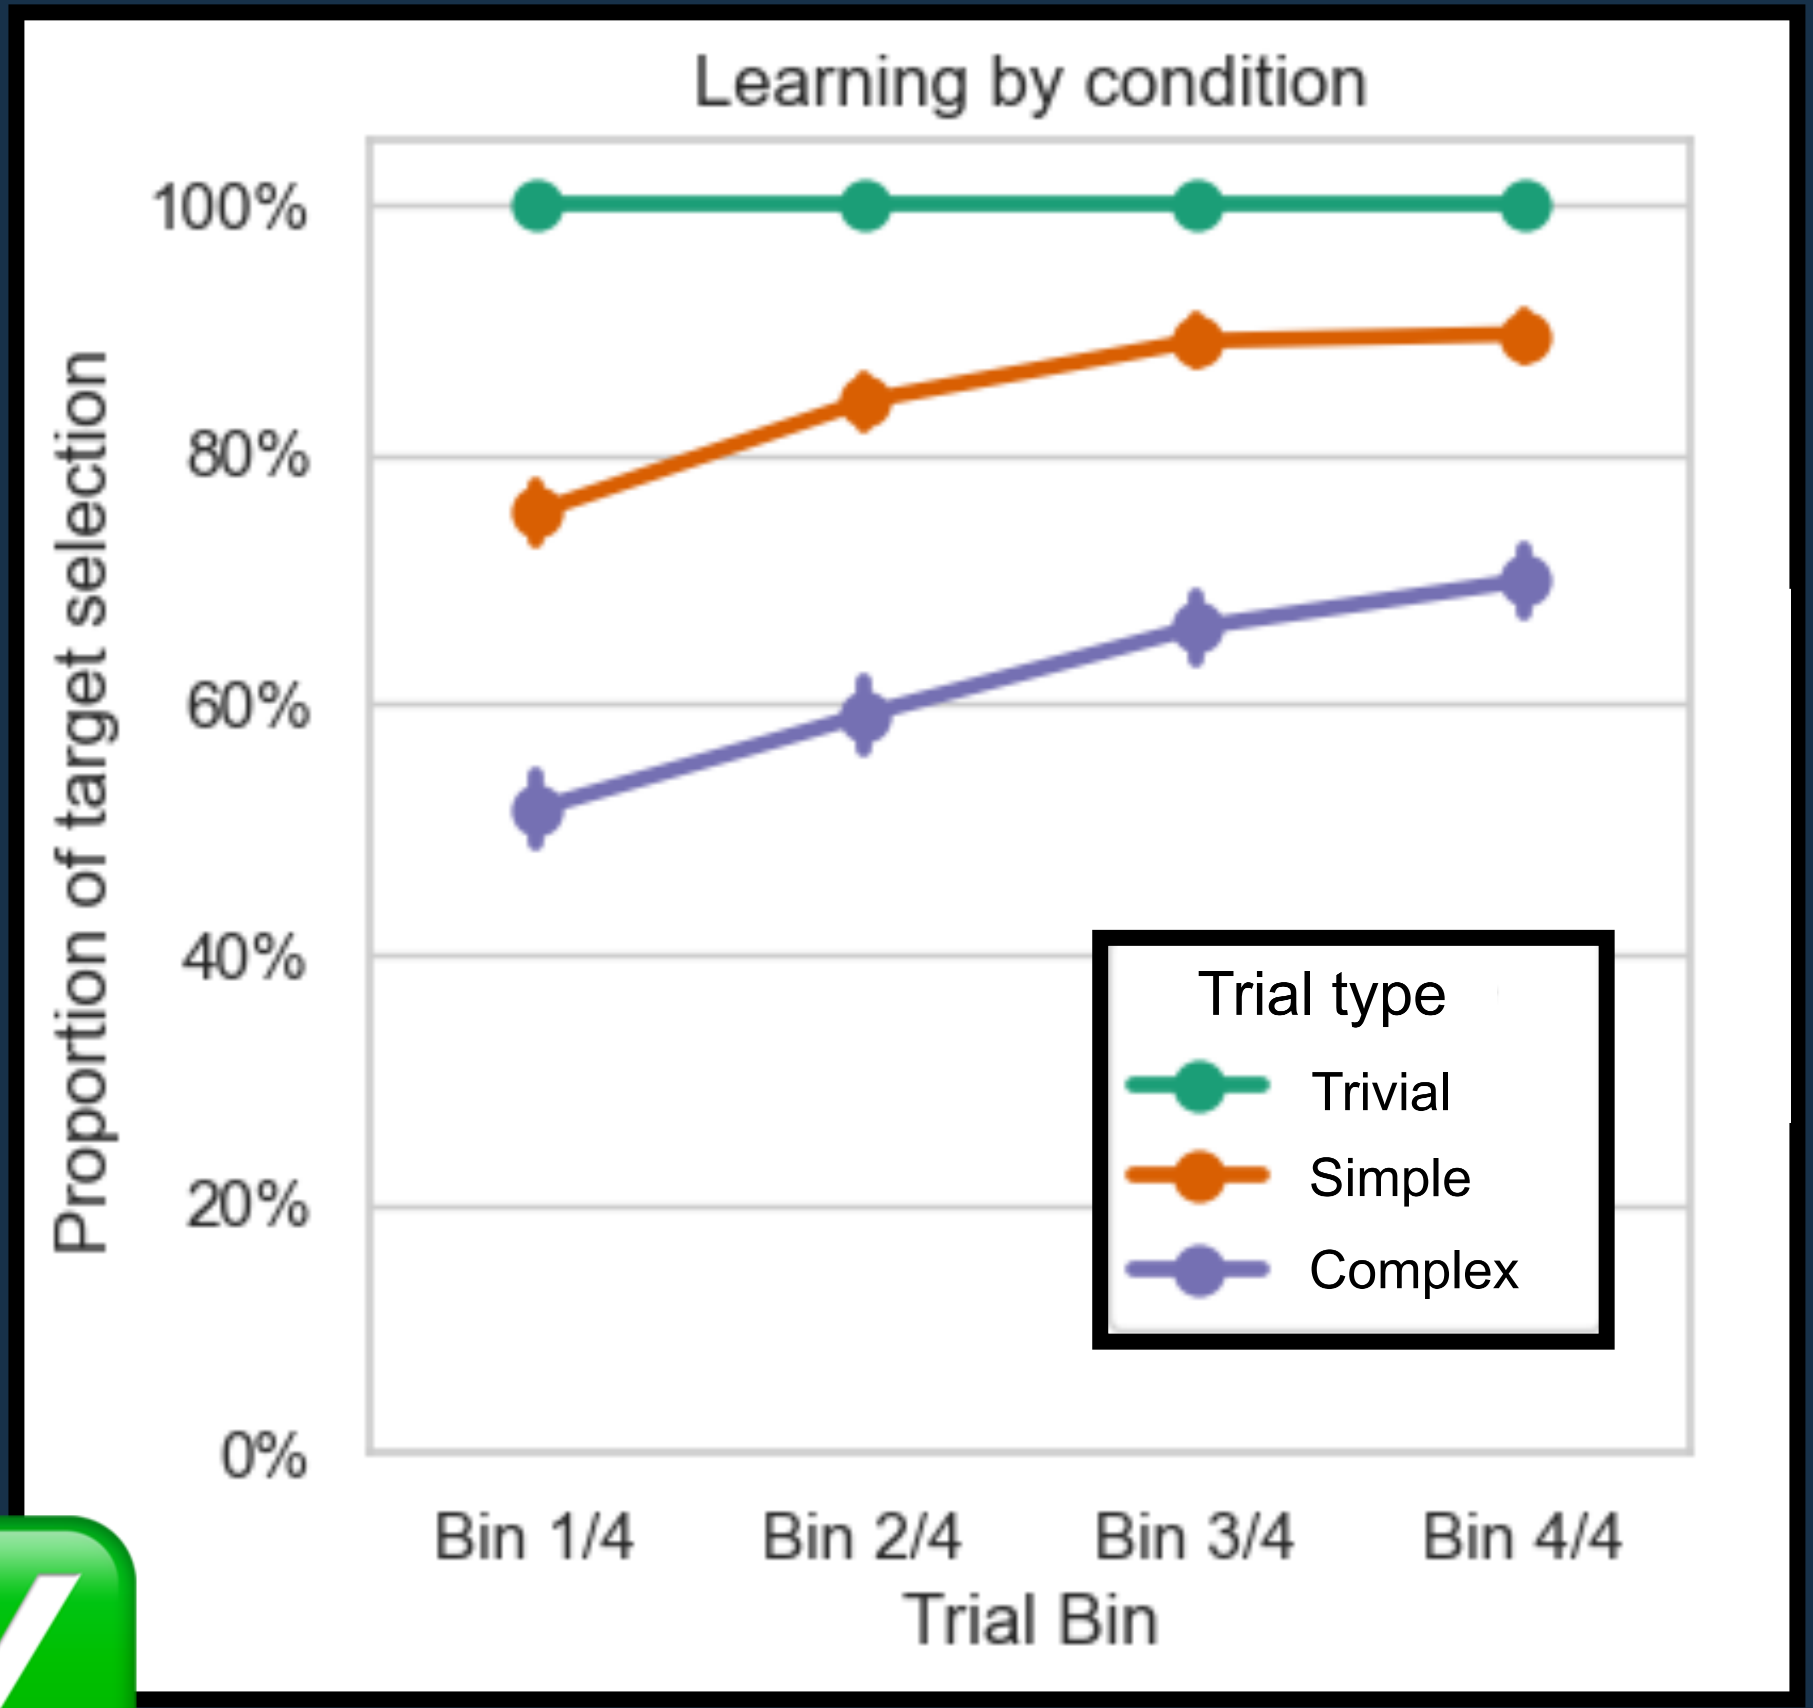

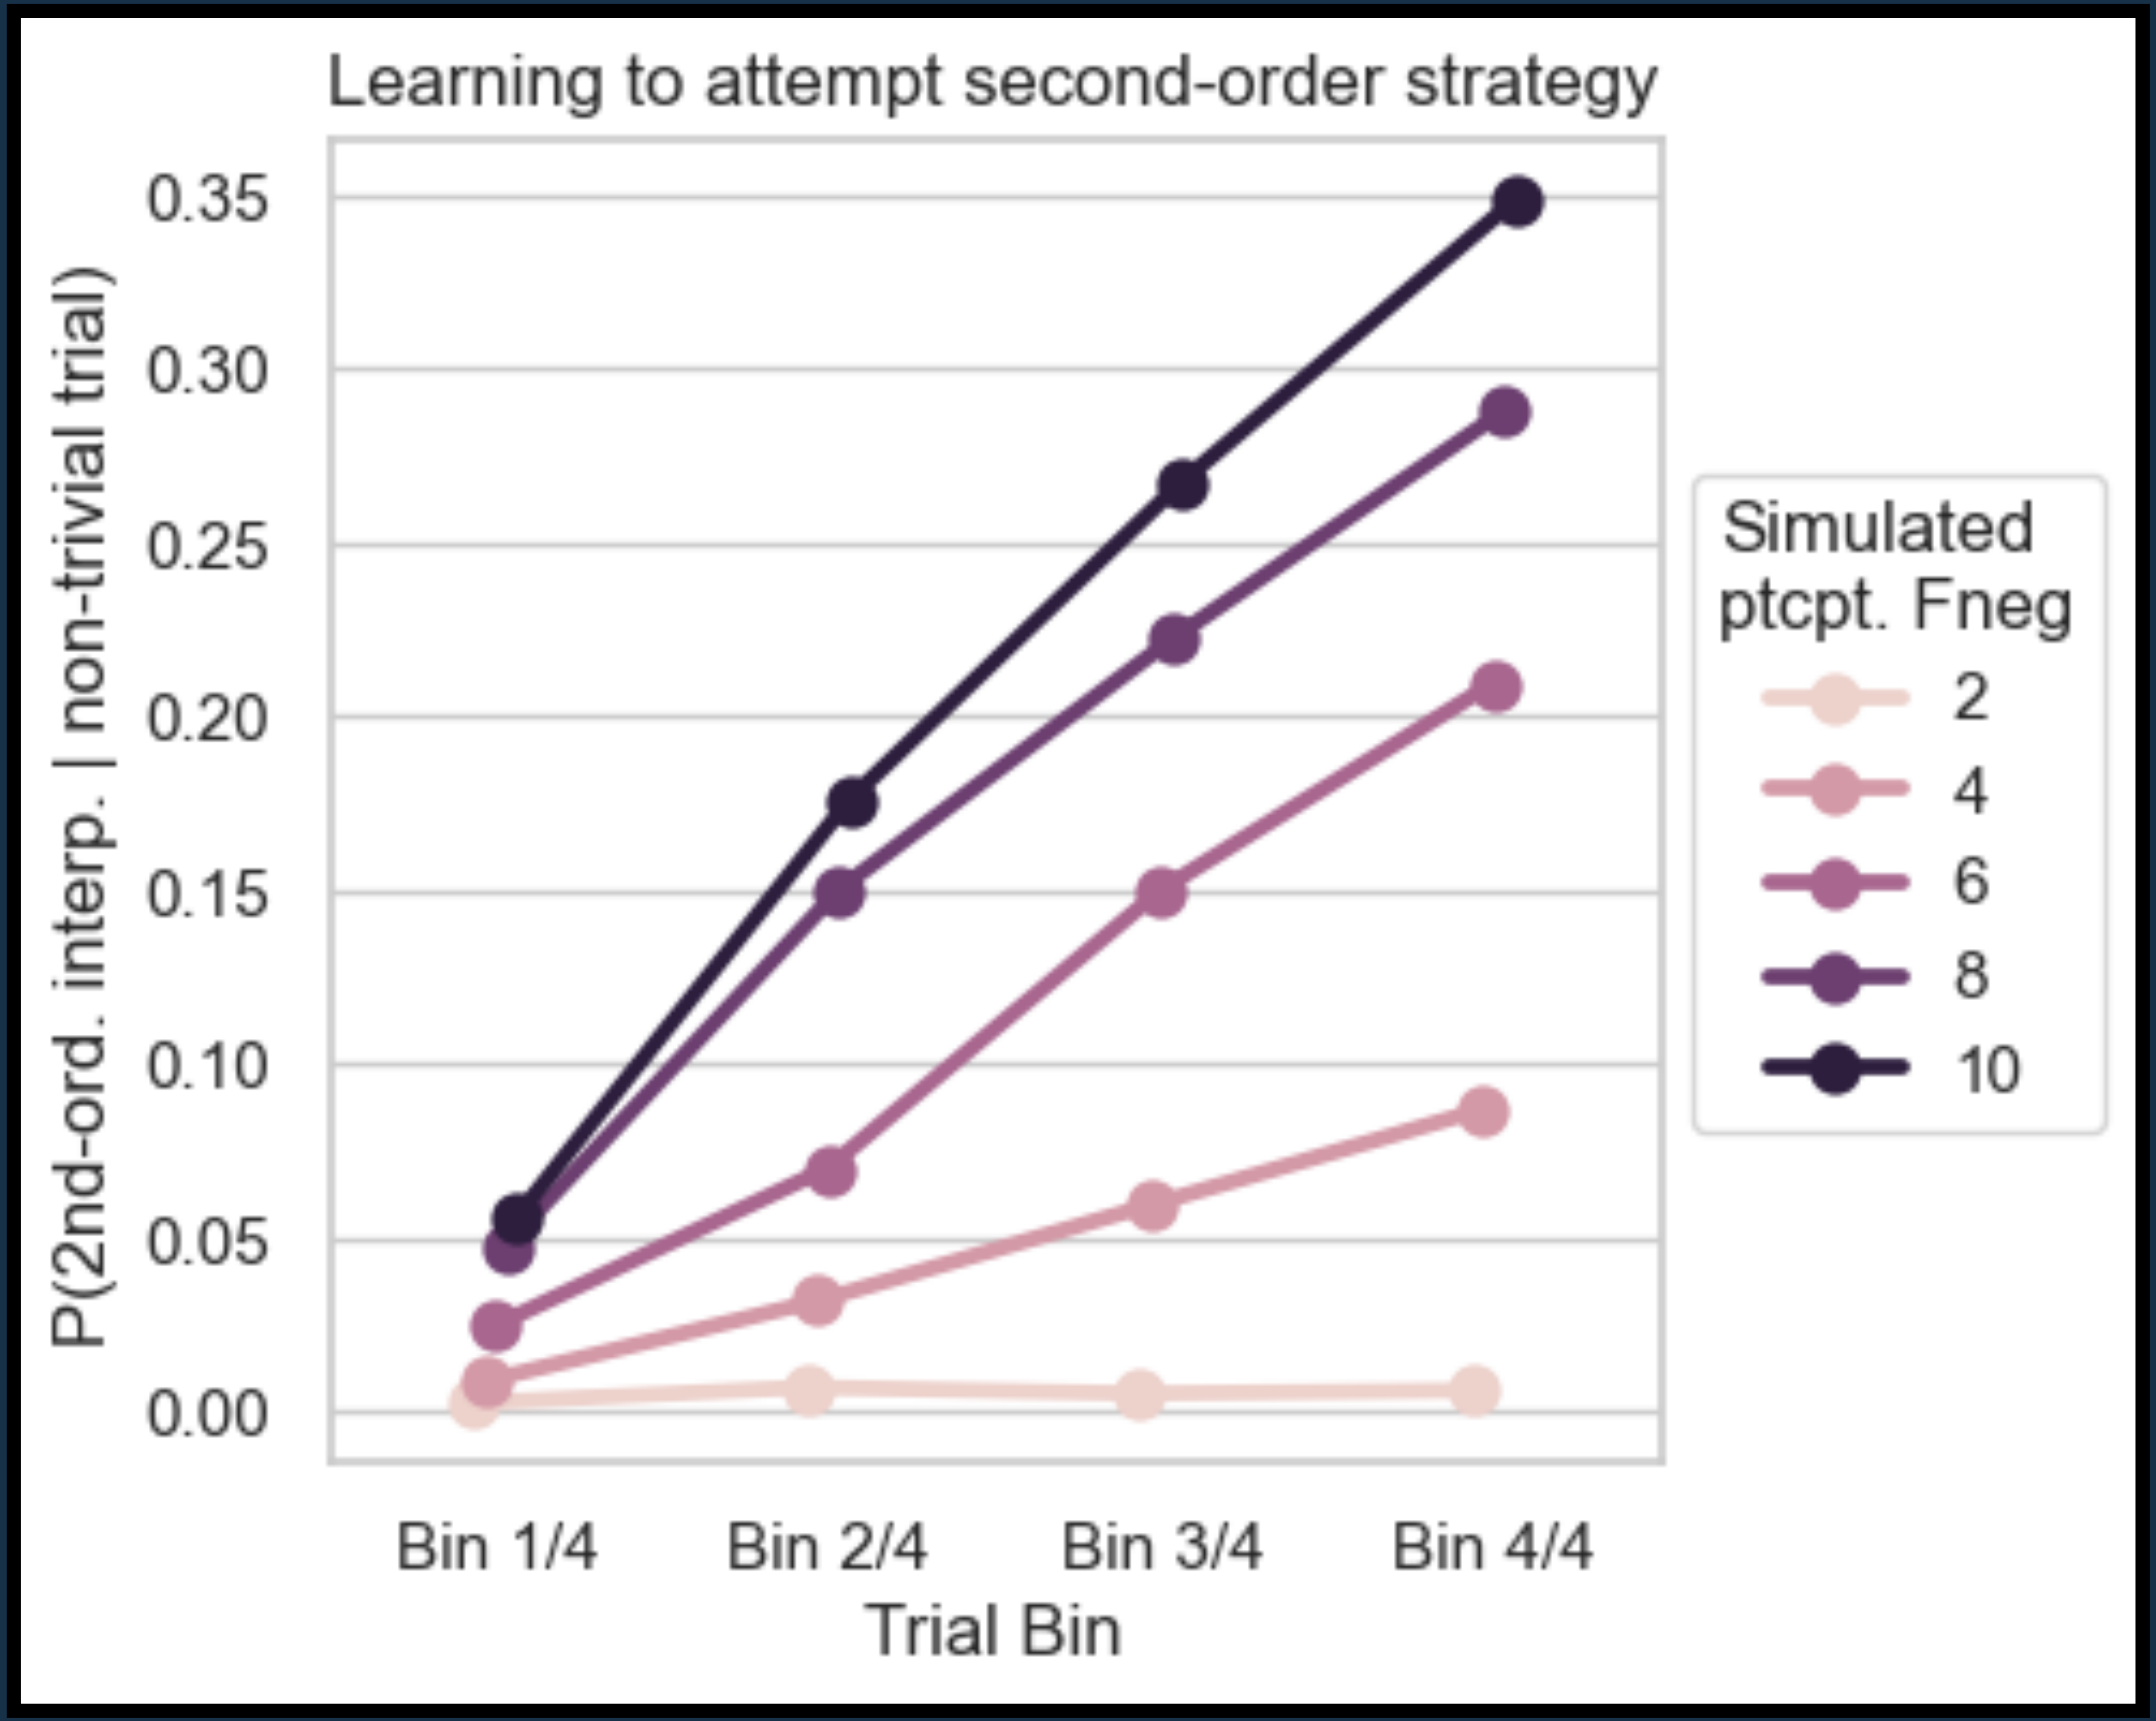

$\beta_{\text{TRIAL}} = (0.05, 0.05)_{95\%}$

(from brms logistic regr. with uninf. priors, trial was centered and not scaled)

Roadmap

1. Background
2. Our ACT-R model
3. Modeling individual differences across tasks

Jointly modeling Raven's and RefGame

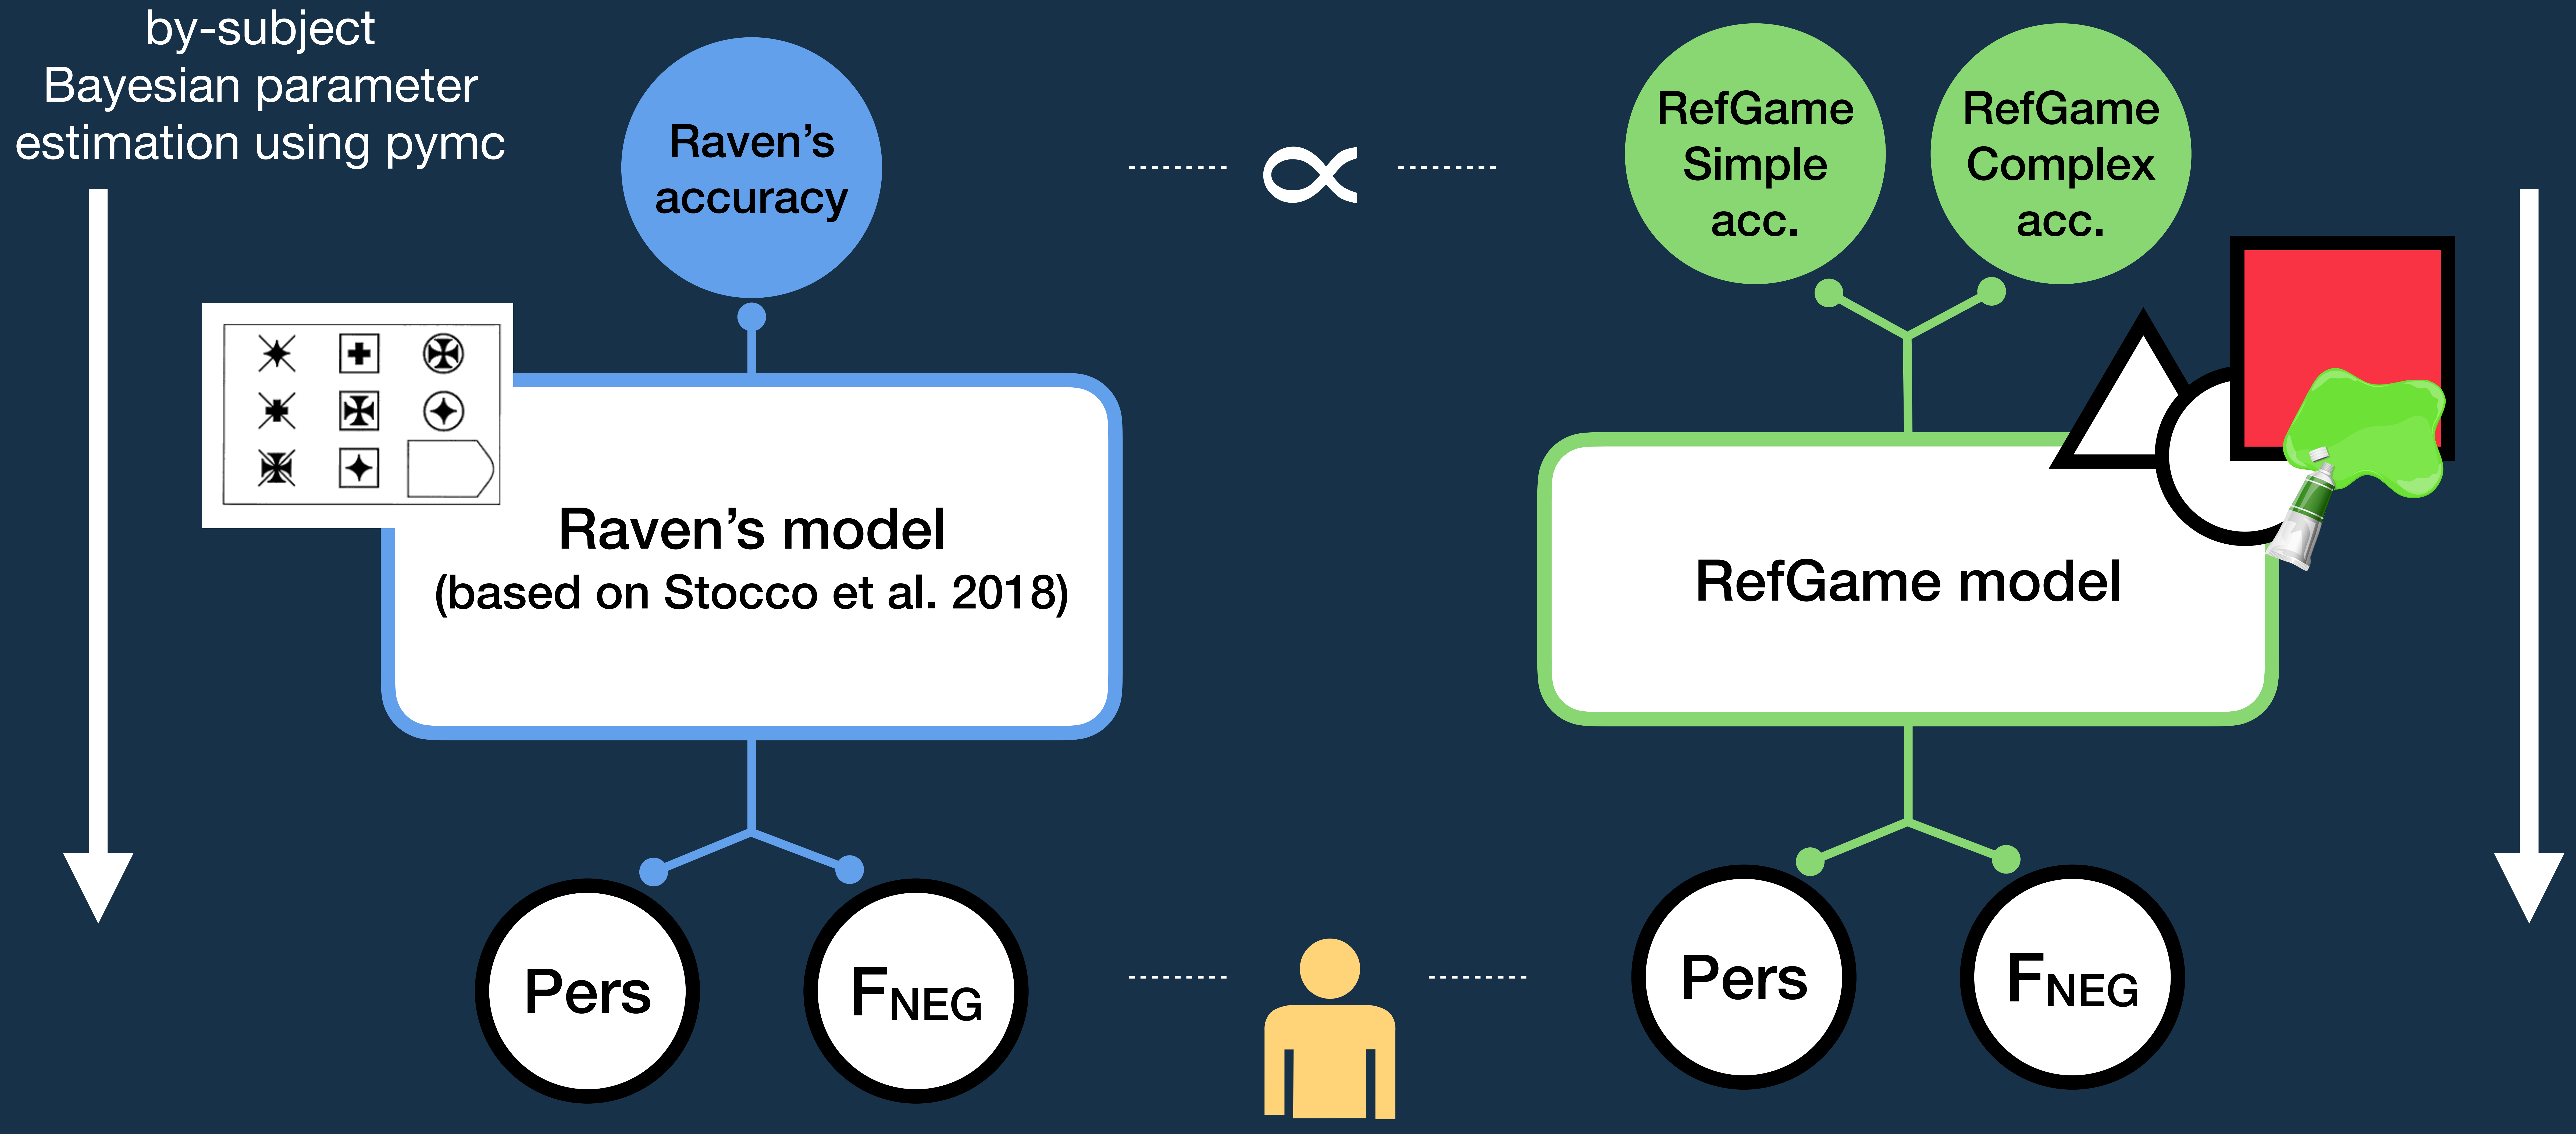

Comparing best-fit parameters across tasks

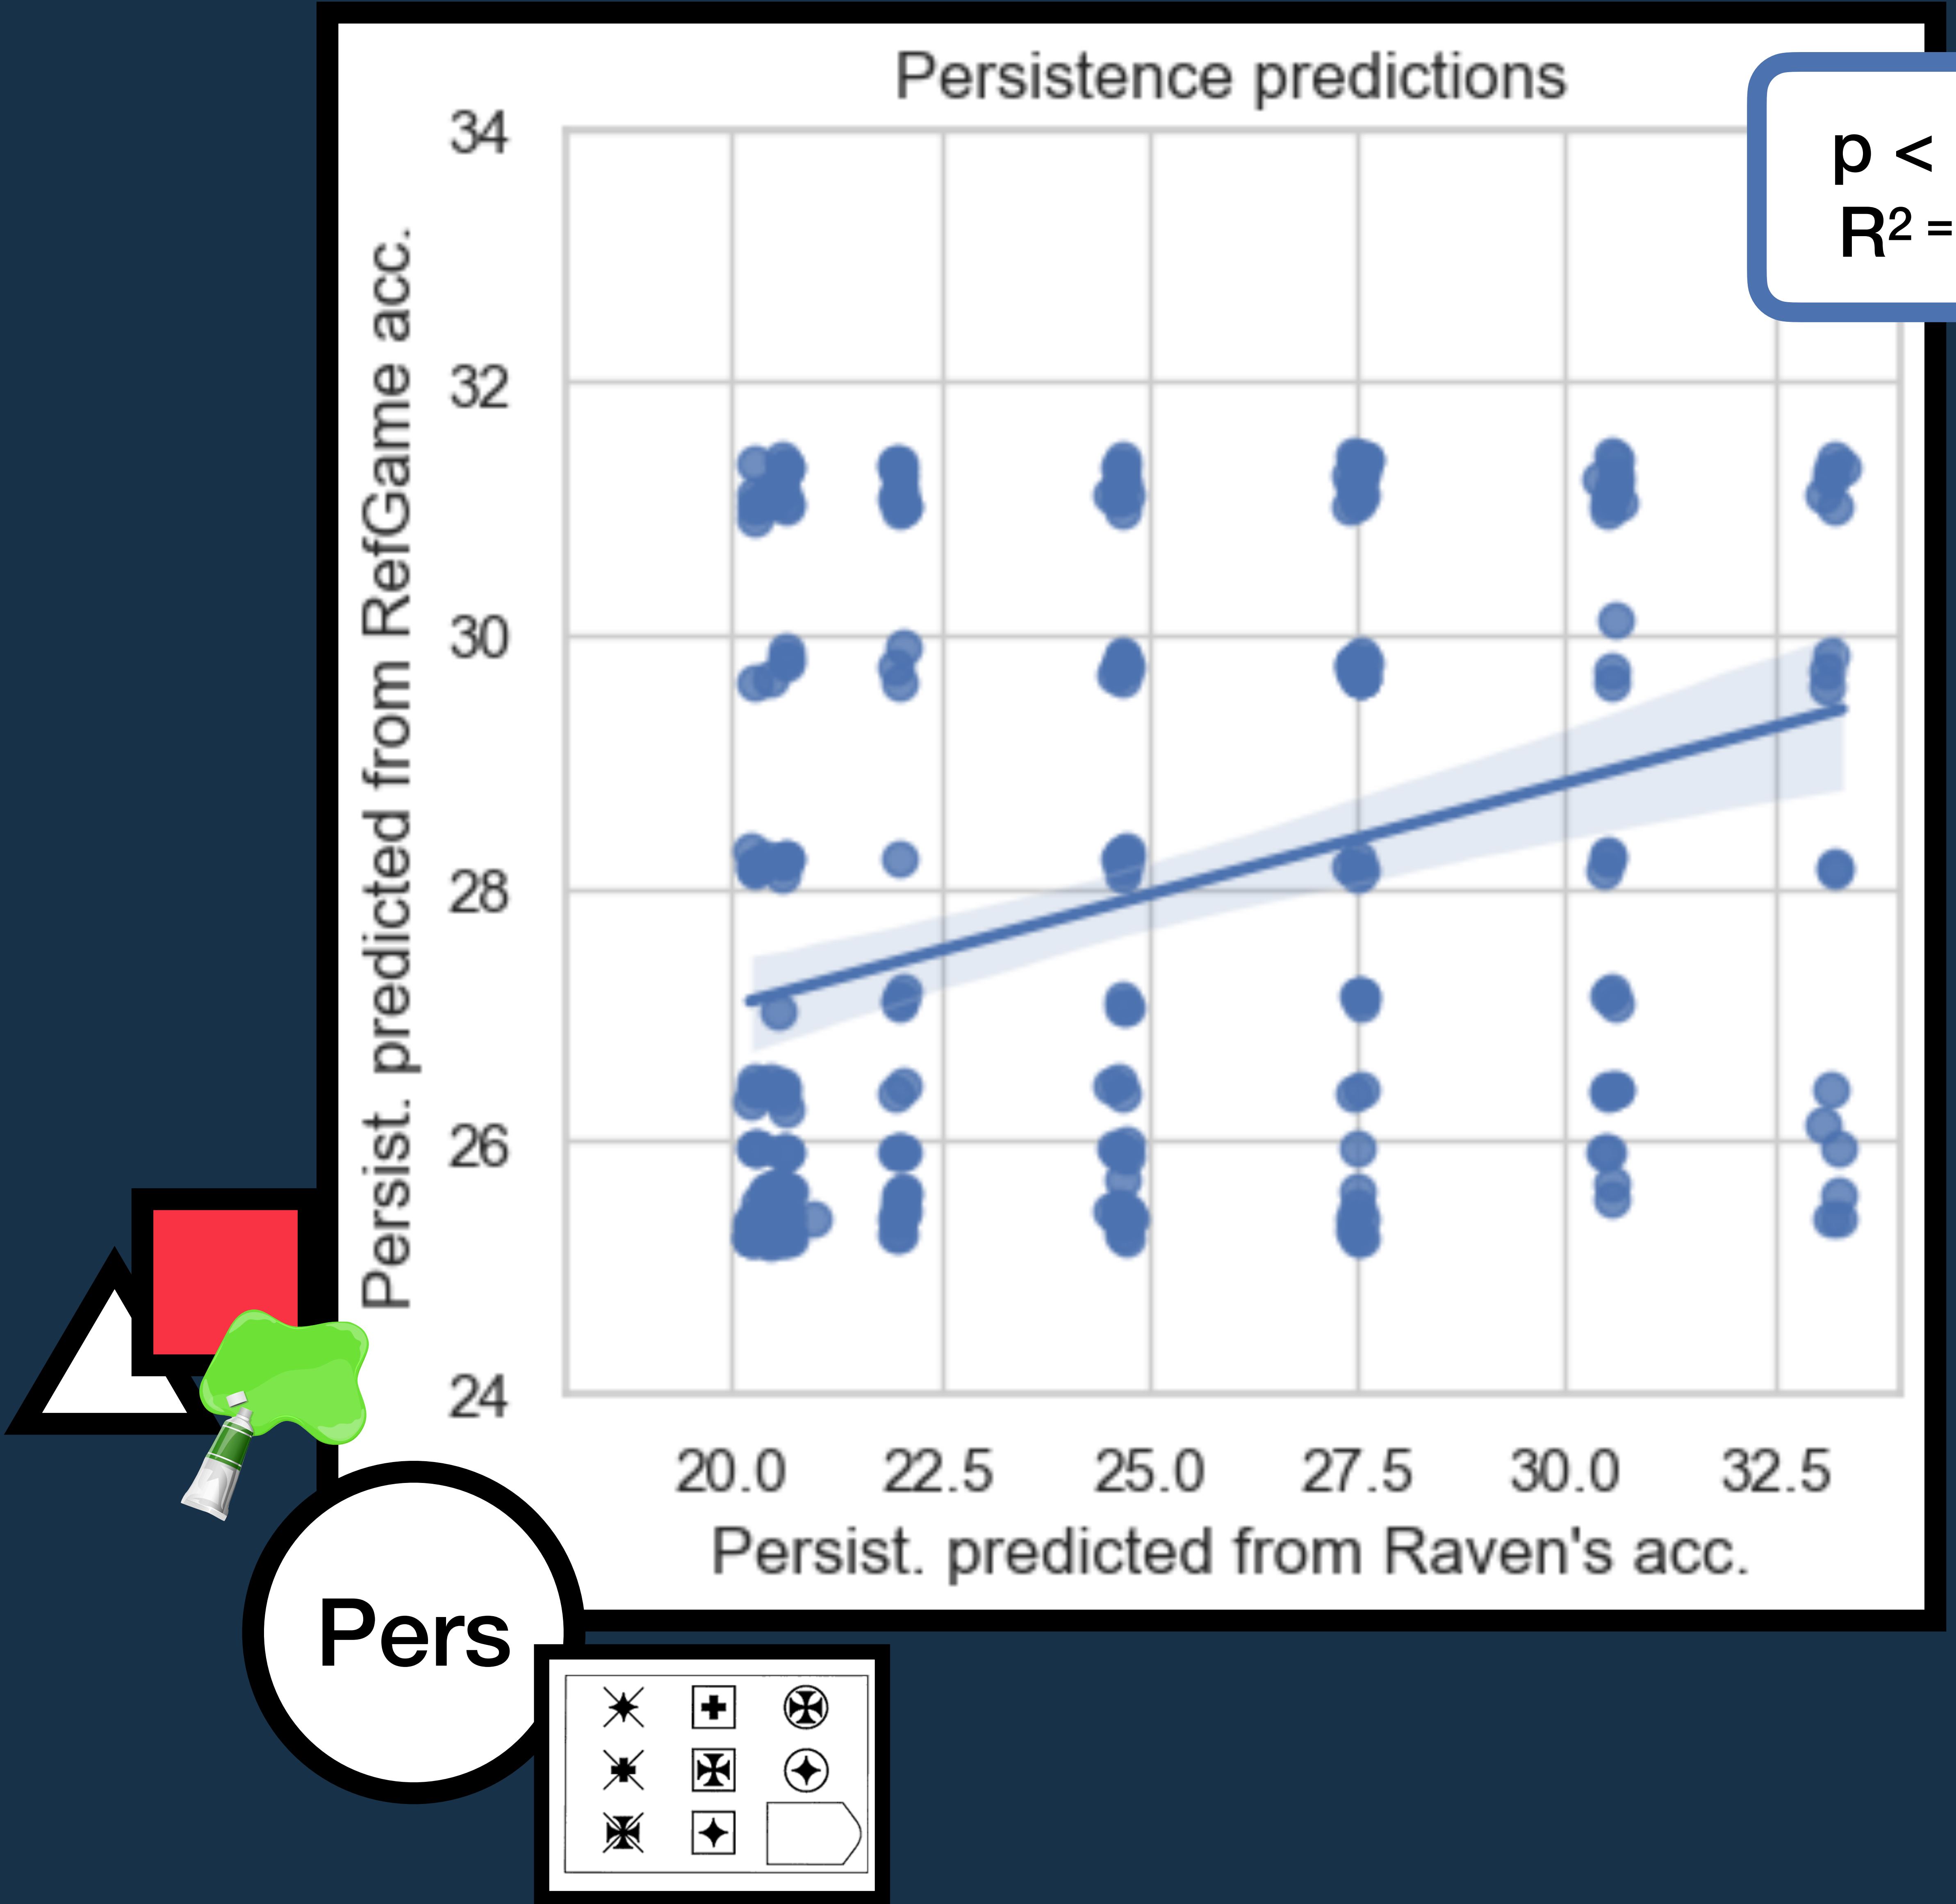

p < 0.01,
$R^2 = 0.09$

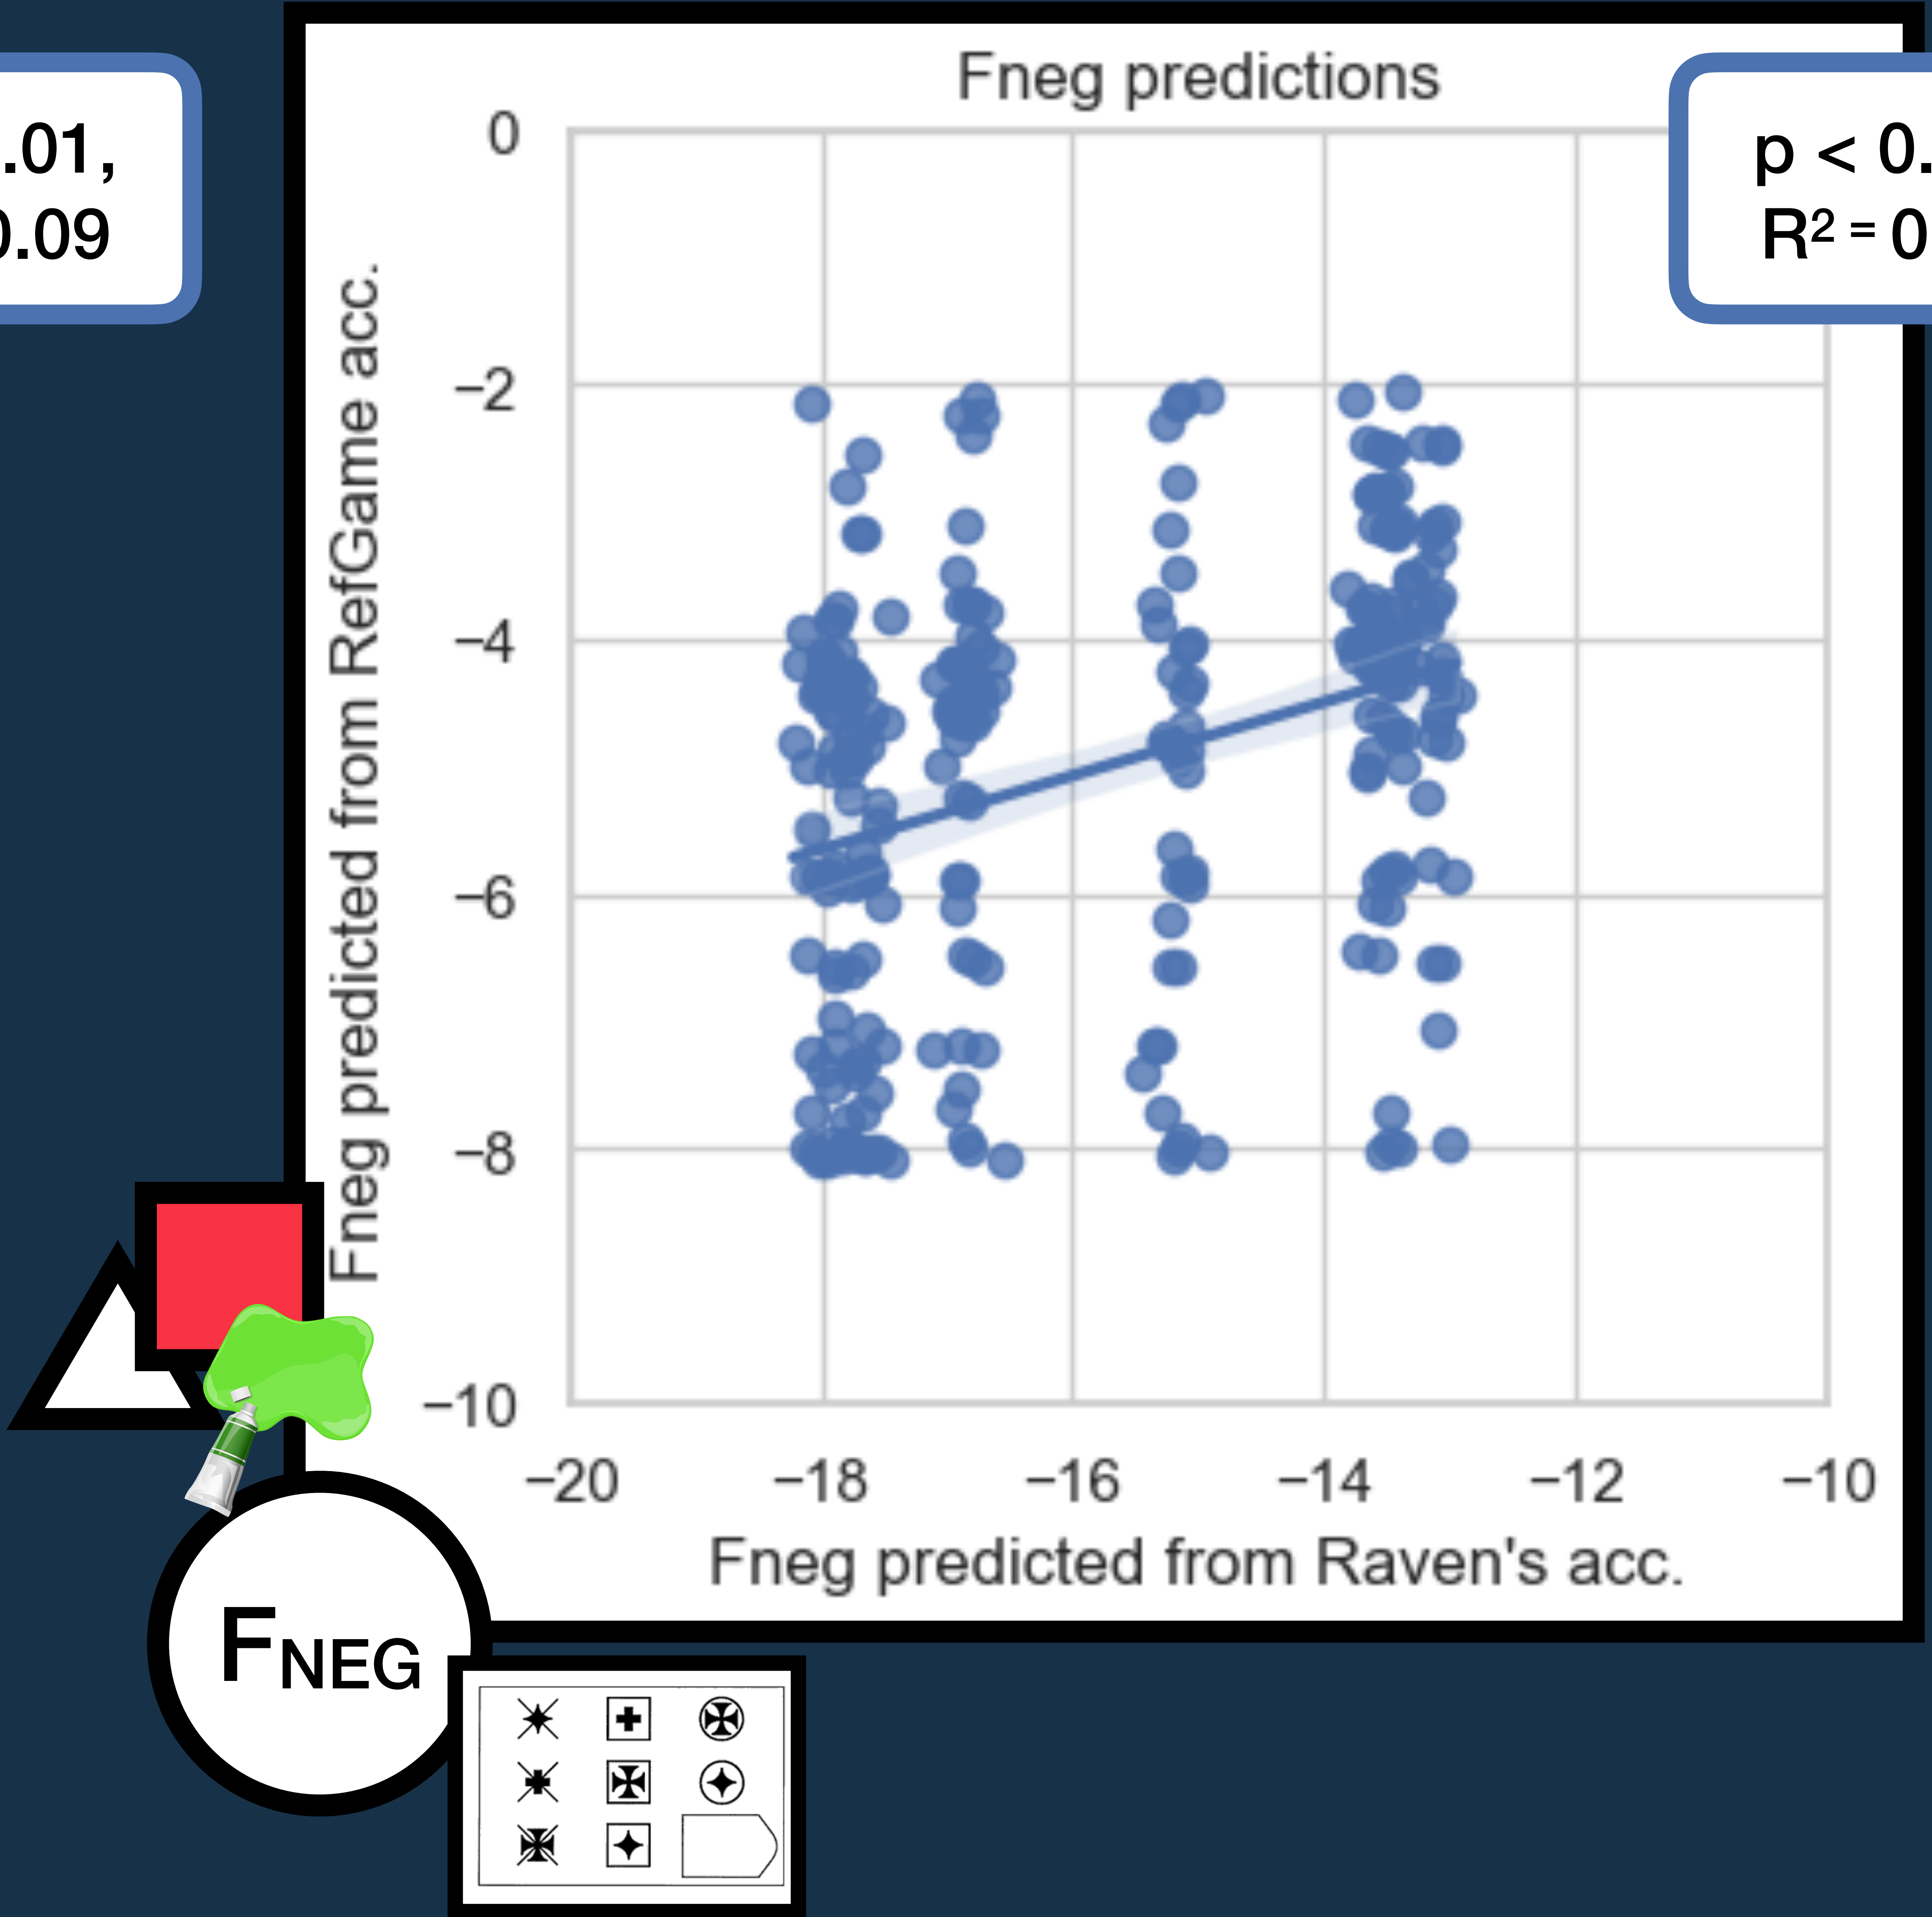

p < 0.01,
$R^2 = 0.11$

Predicting RefGame from Raven's scores

Predicting RefGame from Raven's scores

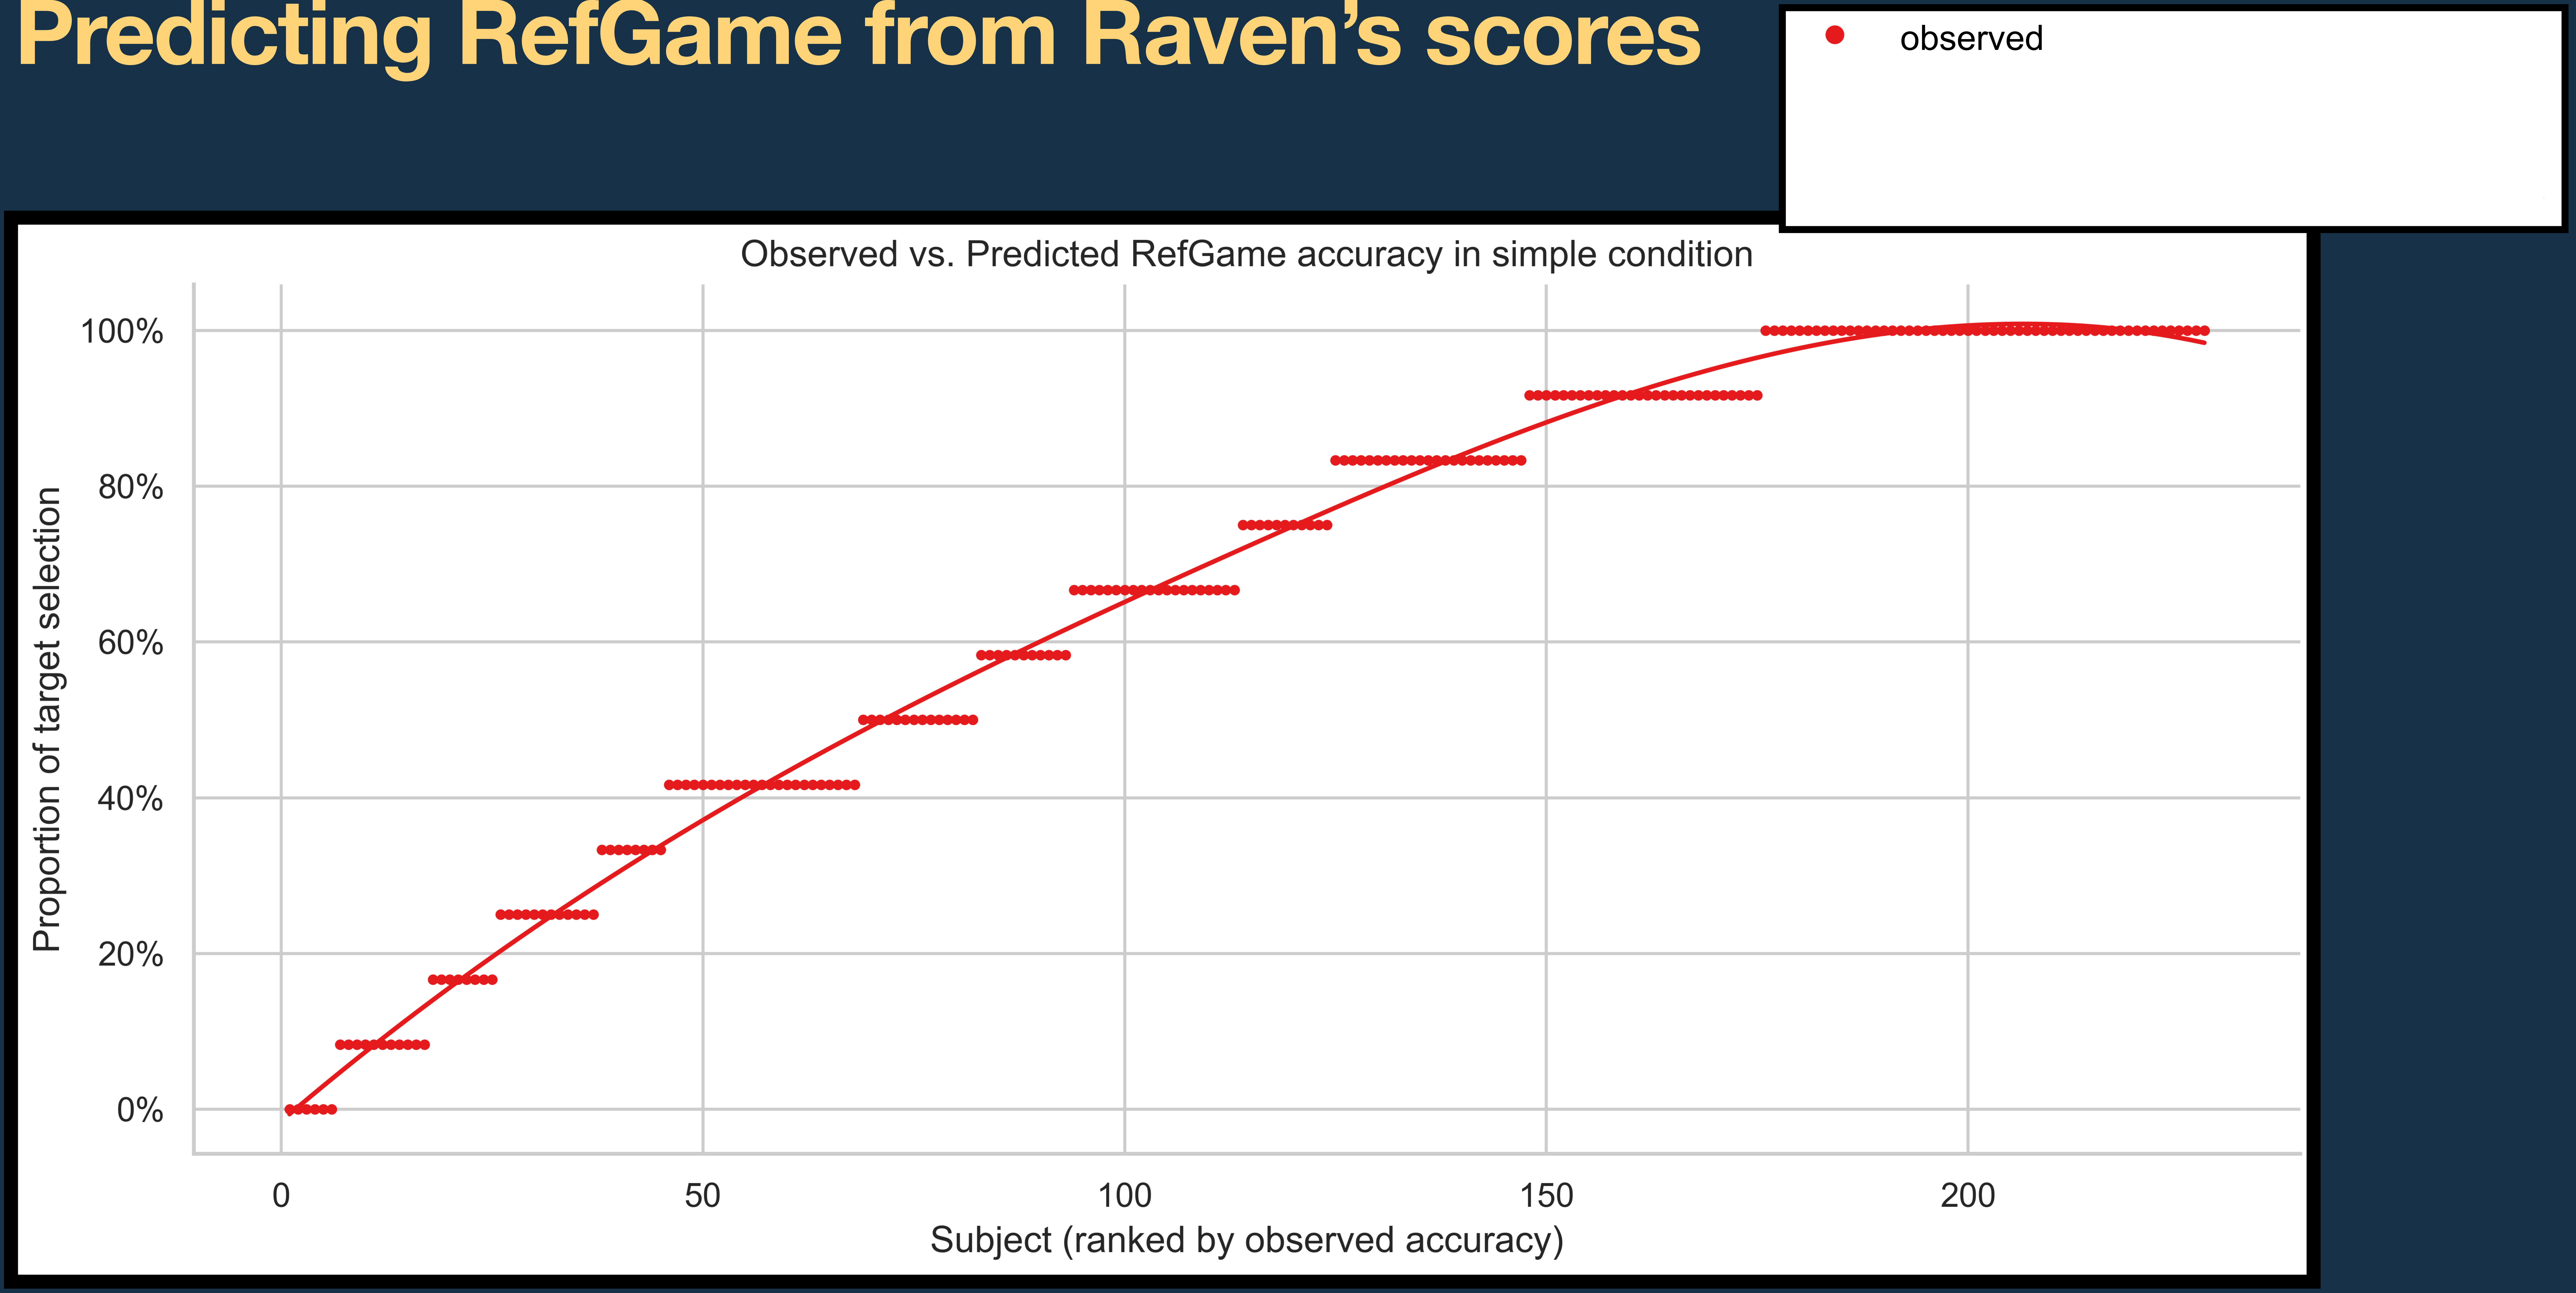

Predicting RefGame from Raven's scores

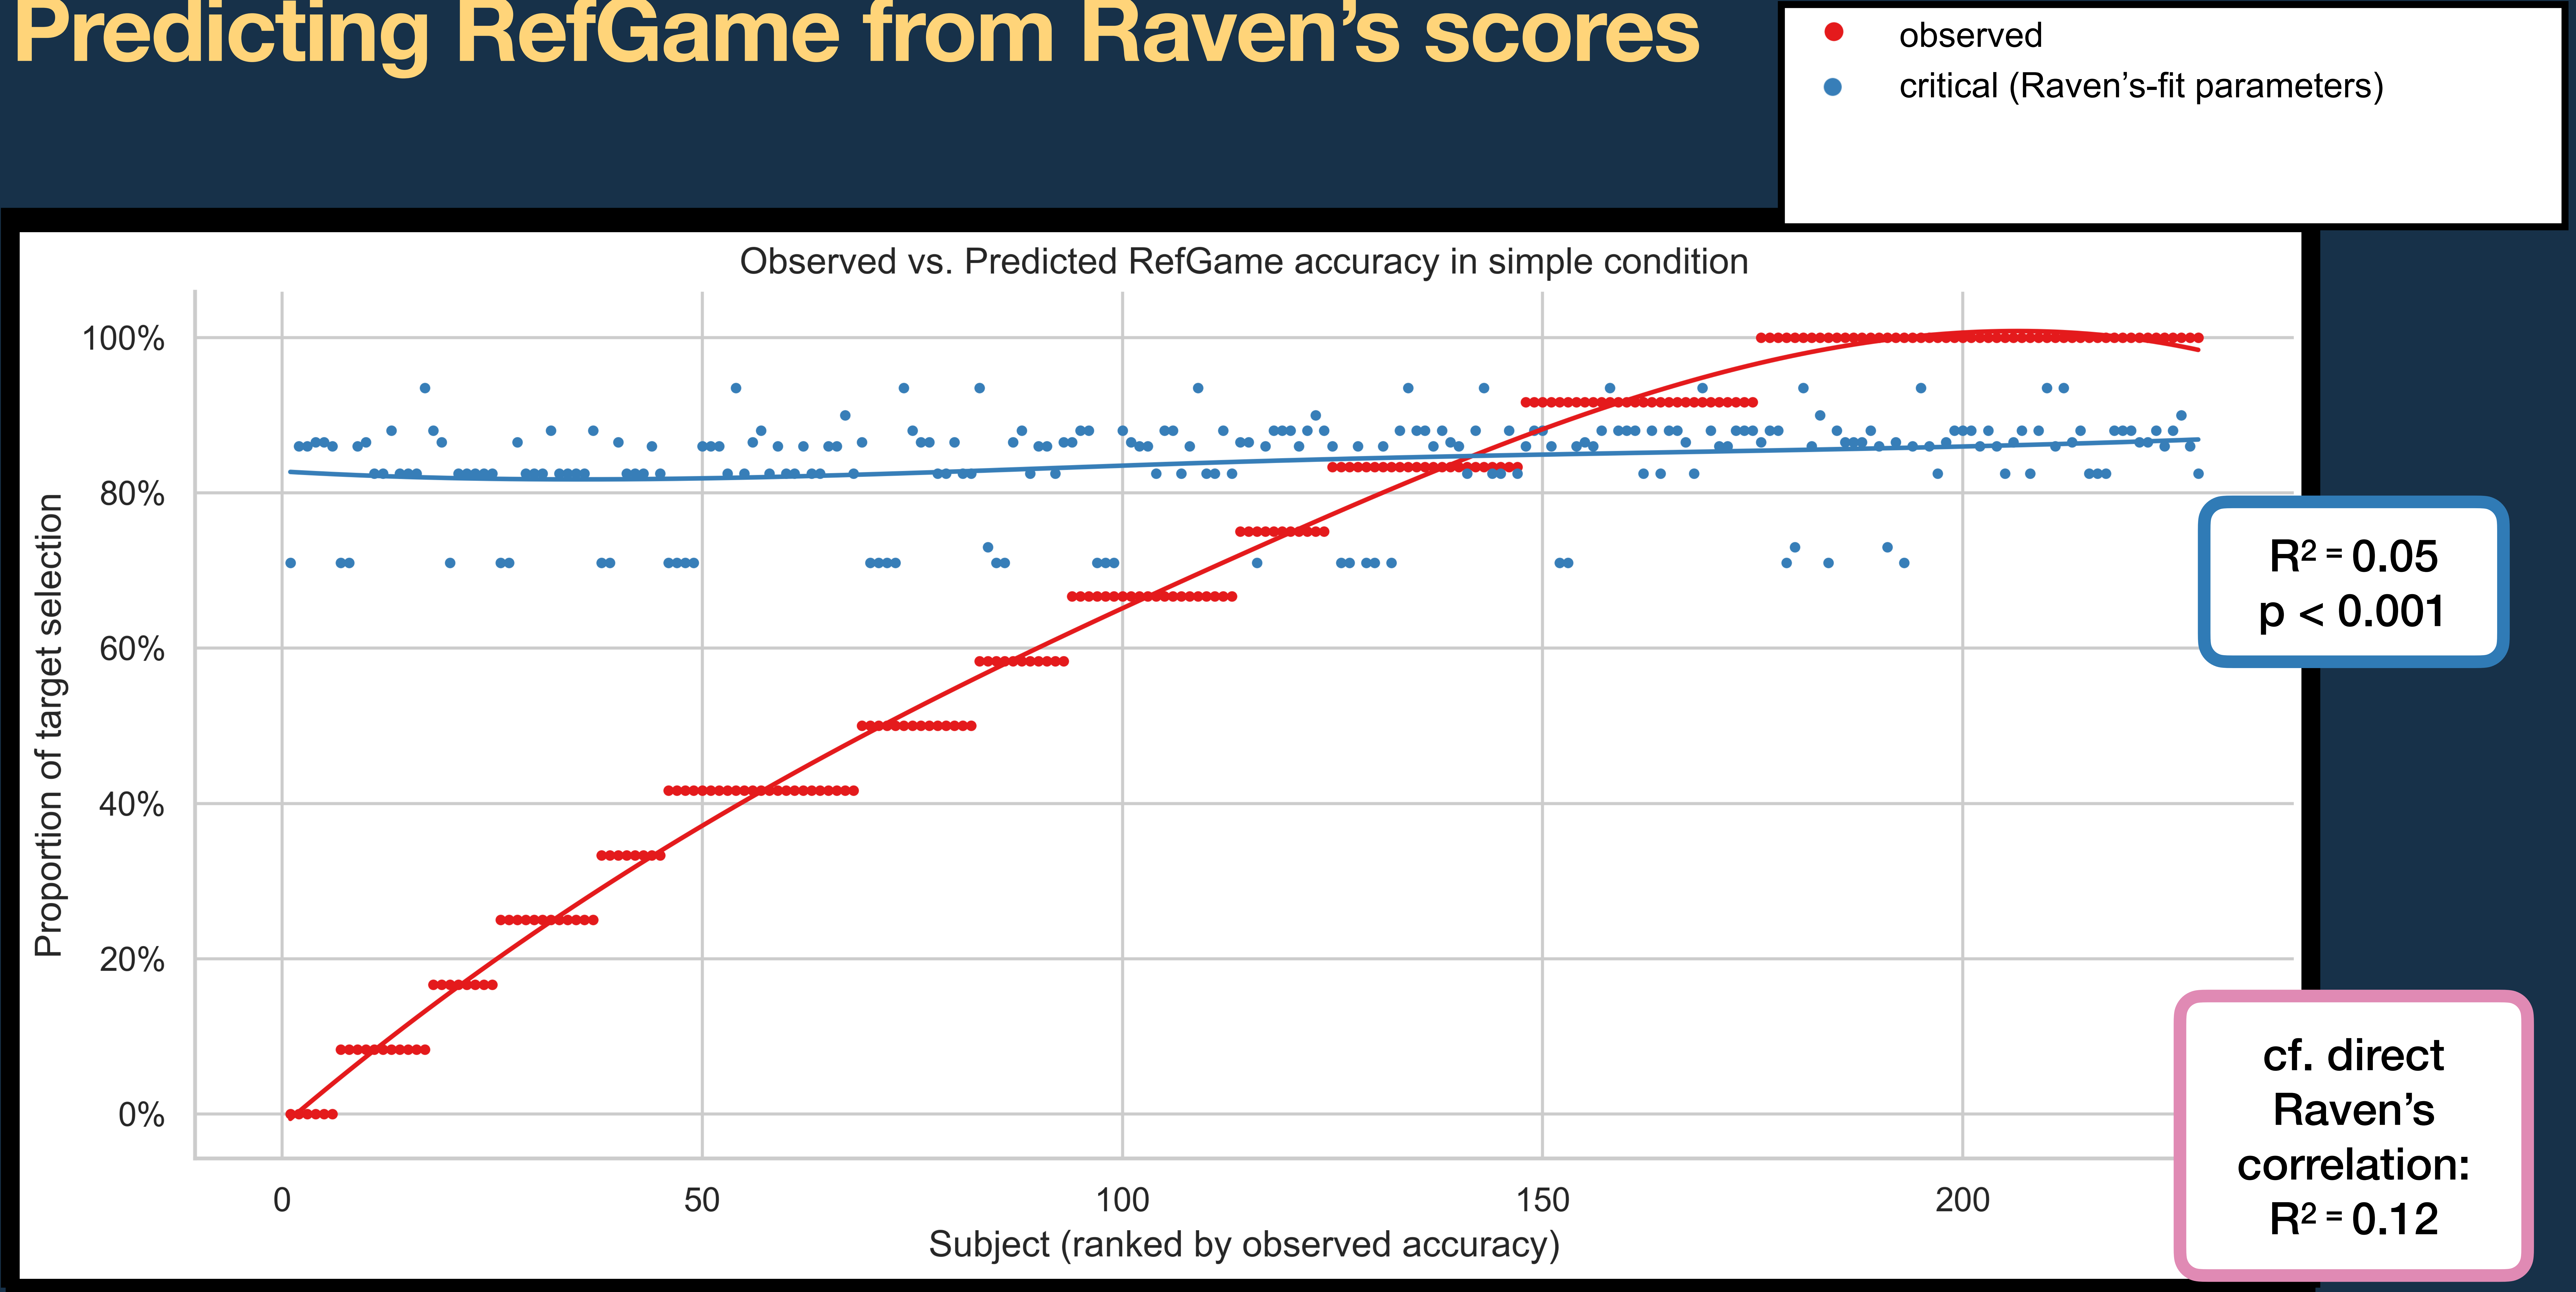

$R^2 = 0.05$
$p < 0.001$

cf. direct Raven's correlation: $R^2 = 0.12$

Deriving an upper baseline

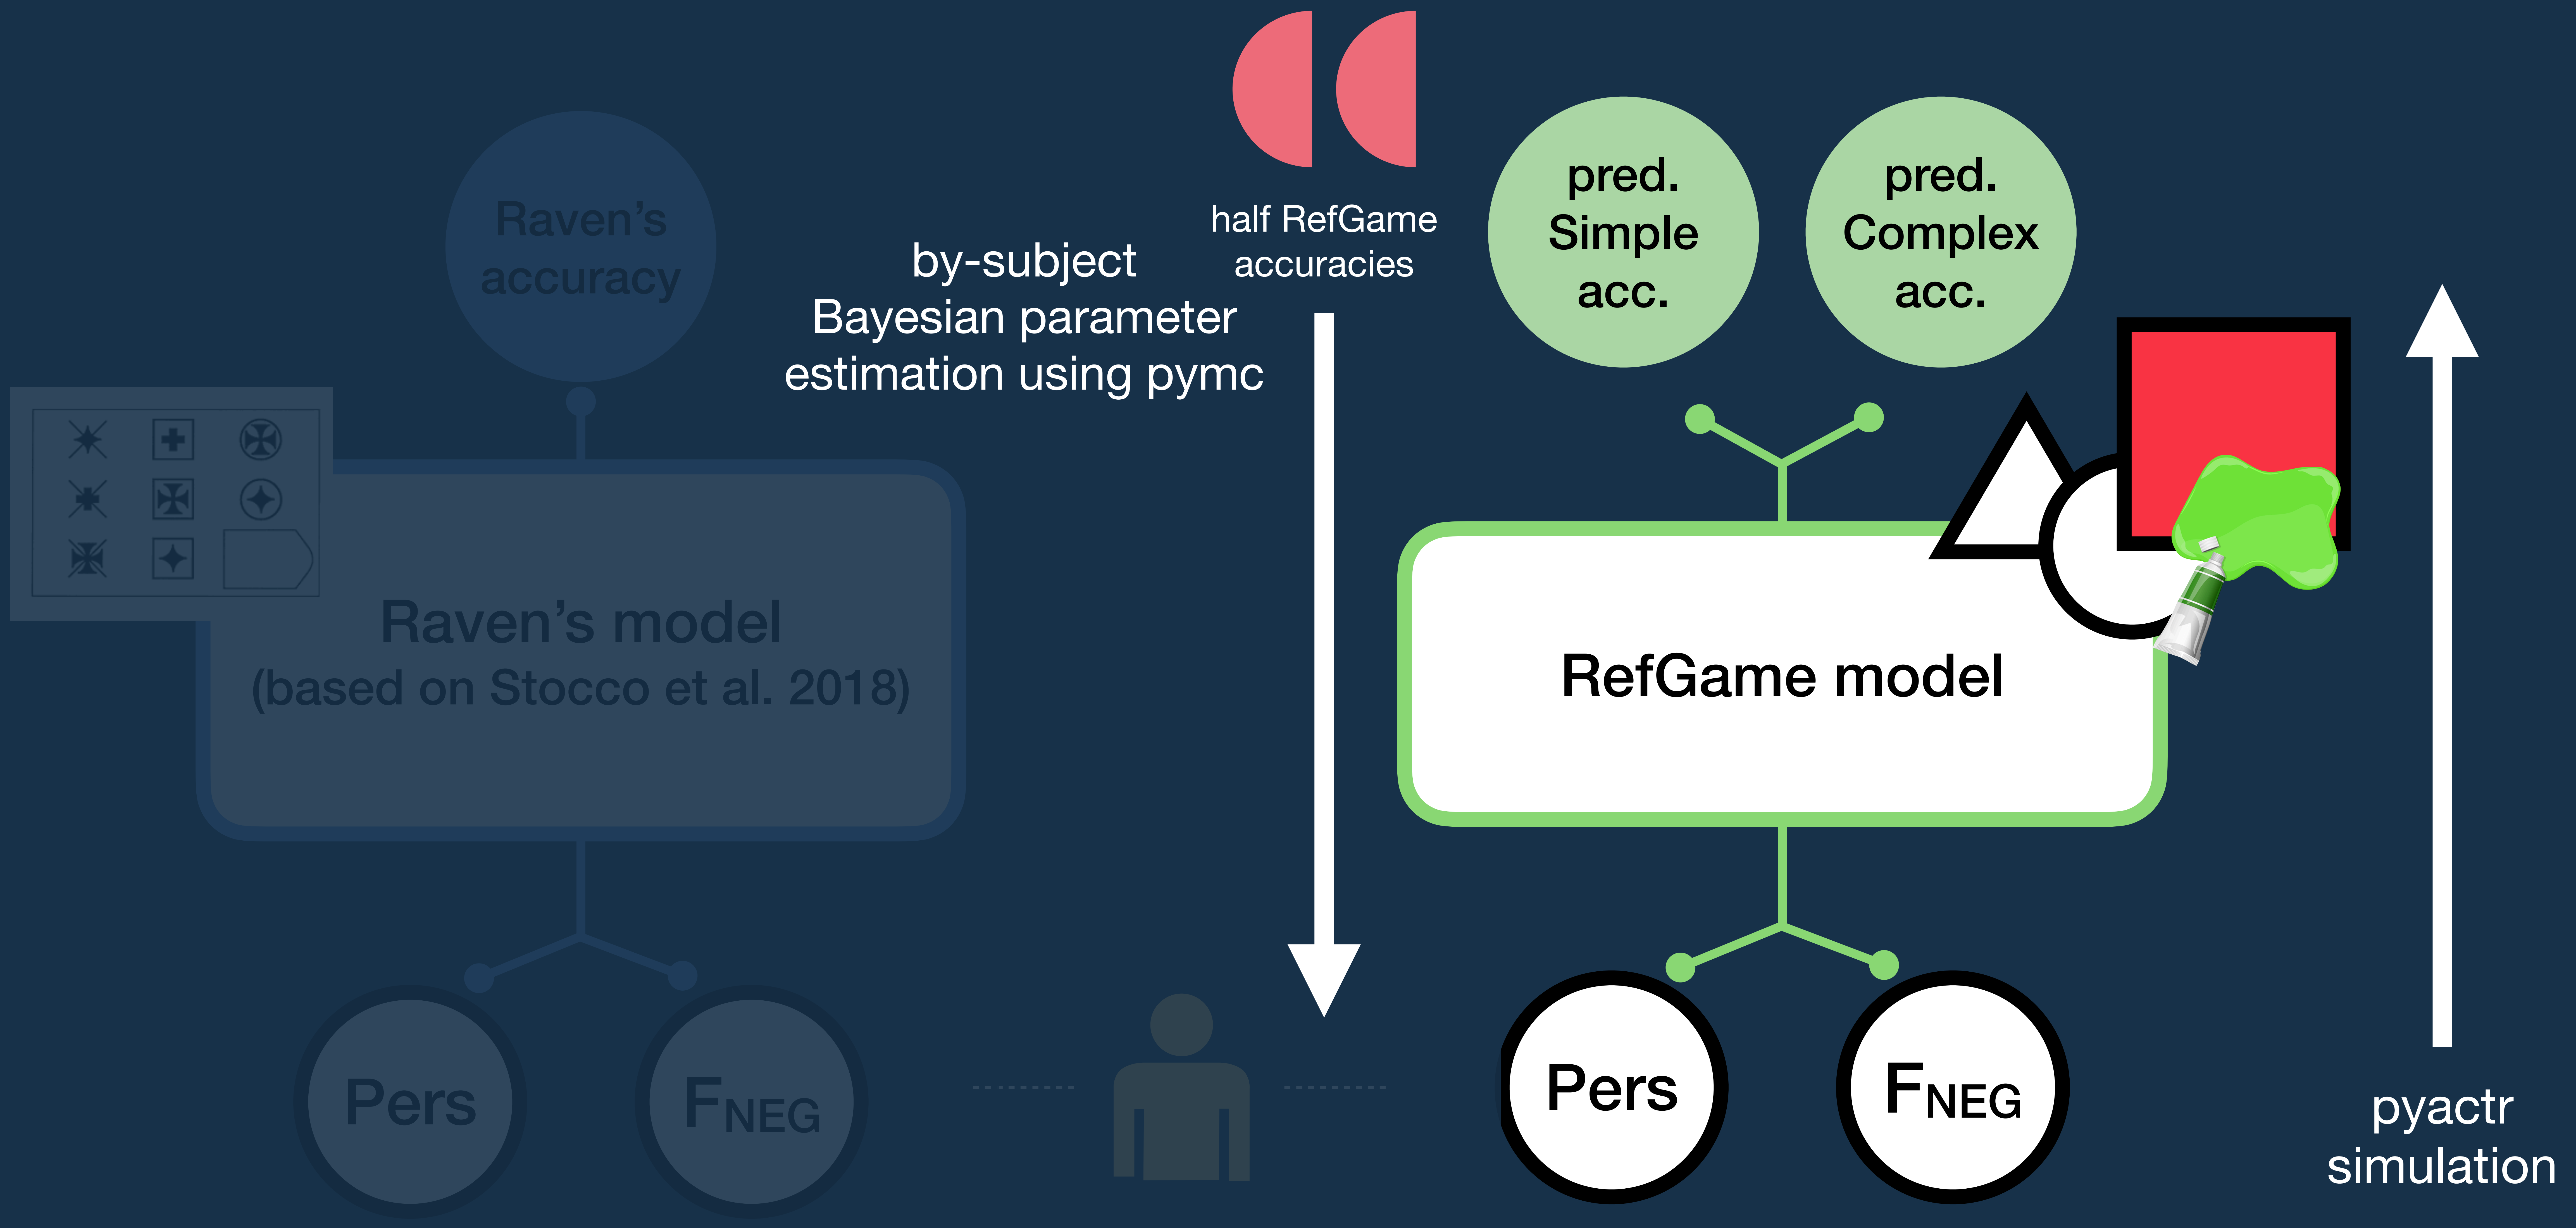

Comparing with an upper baseline

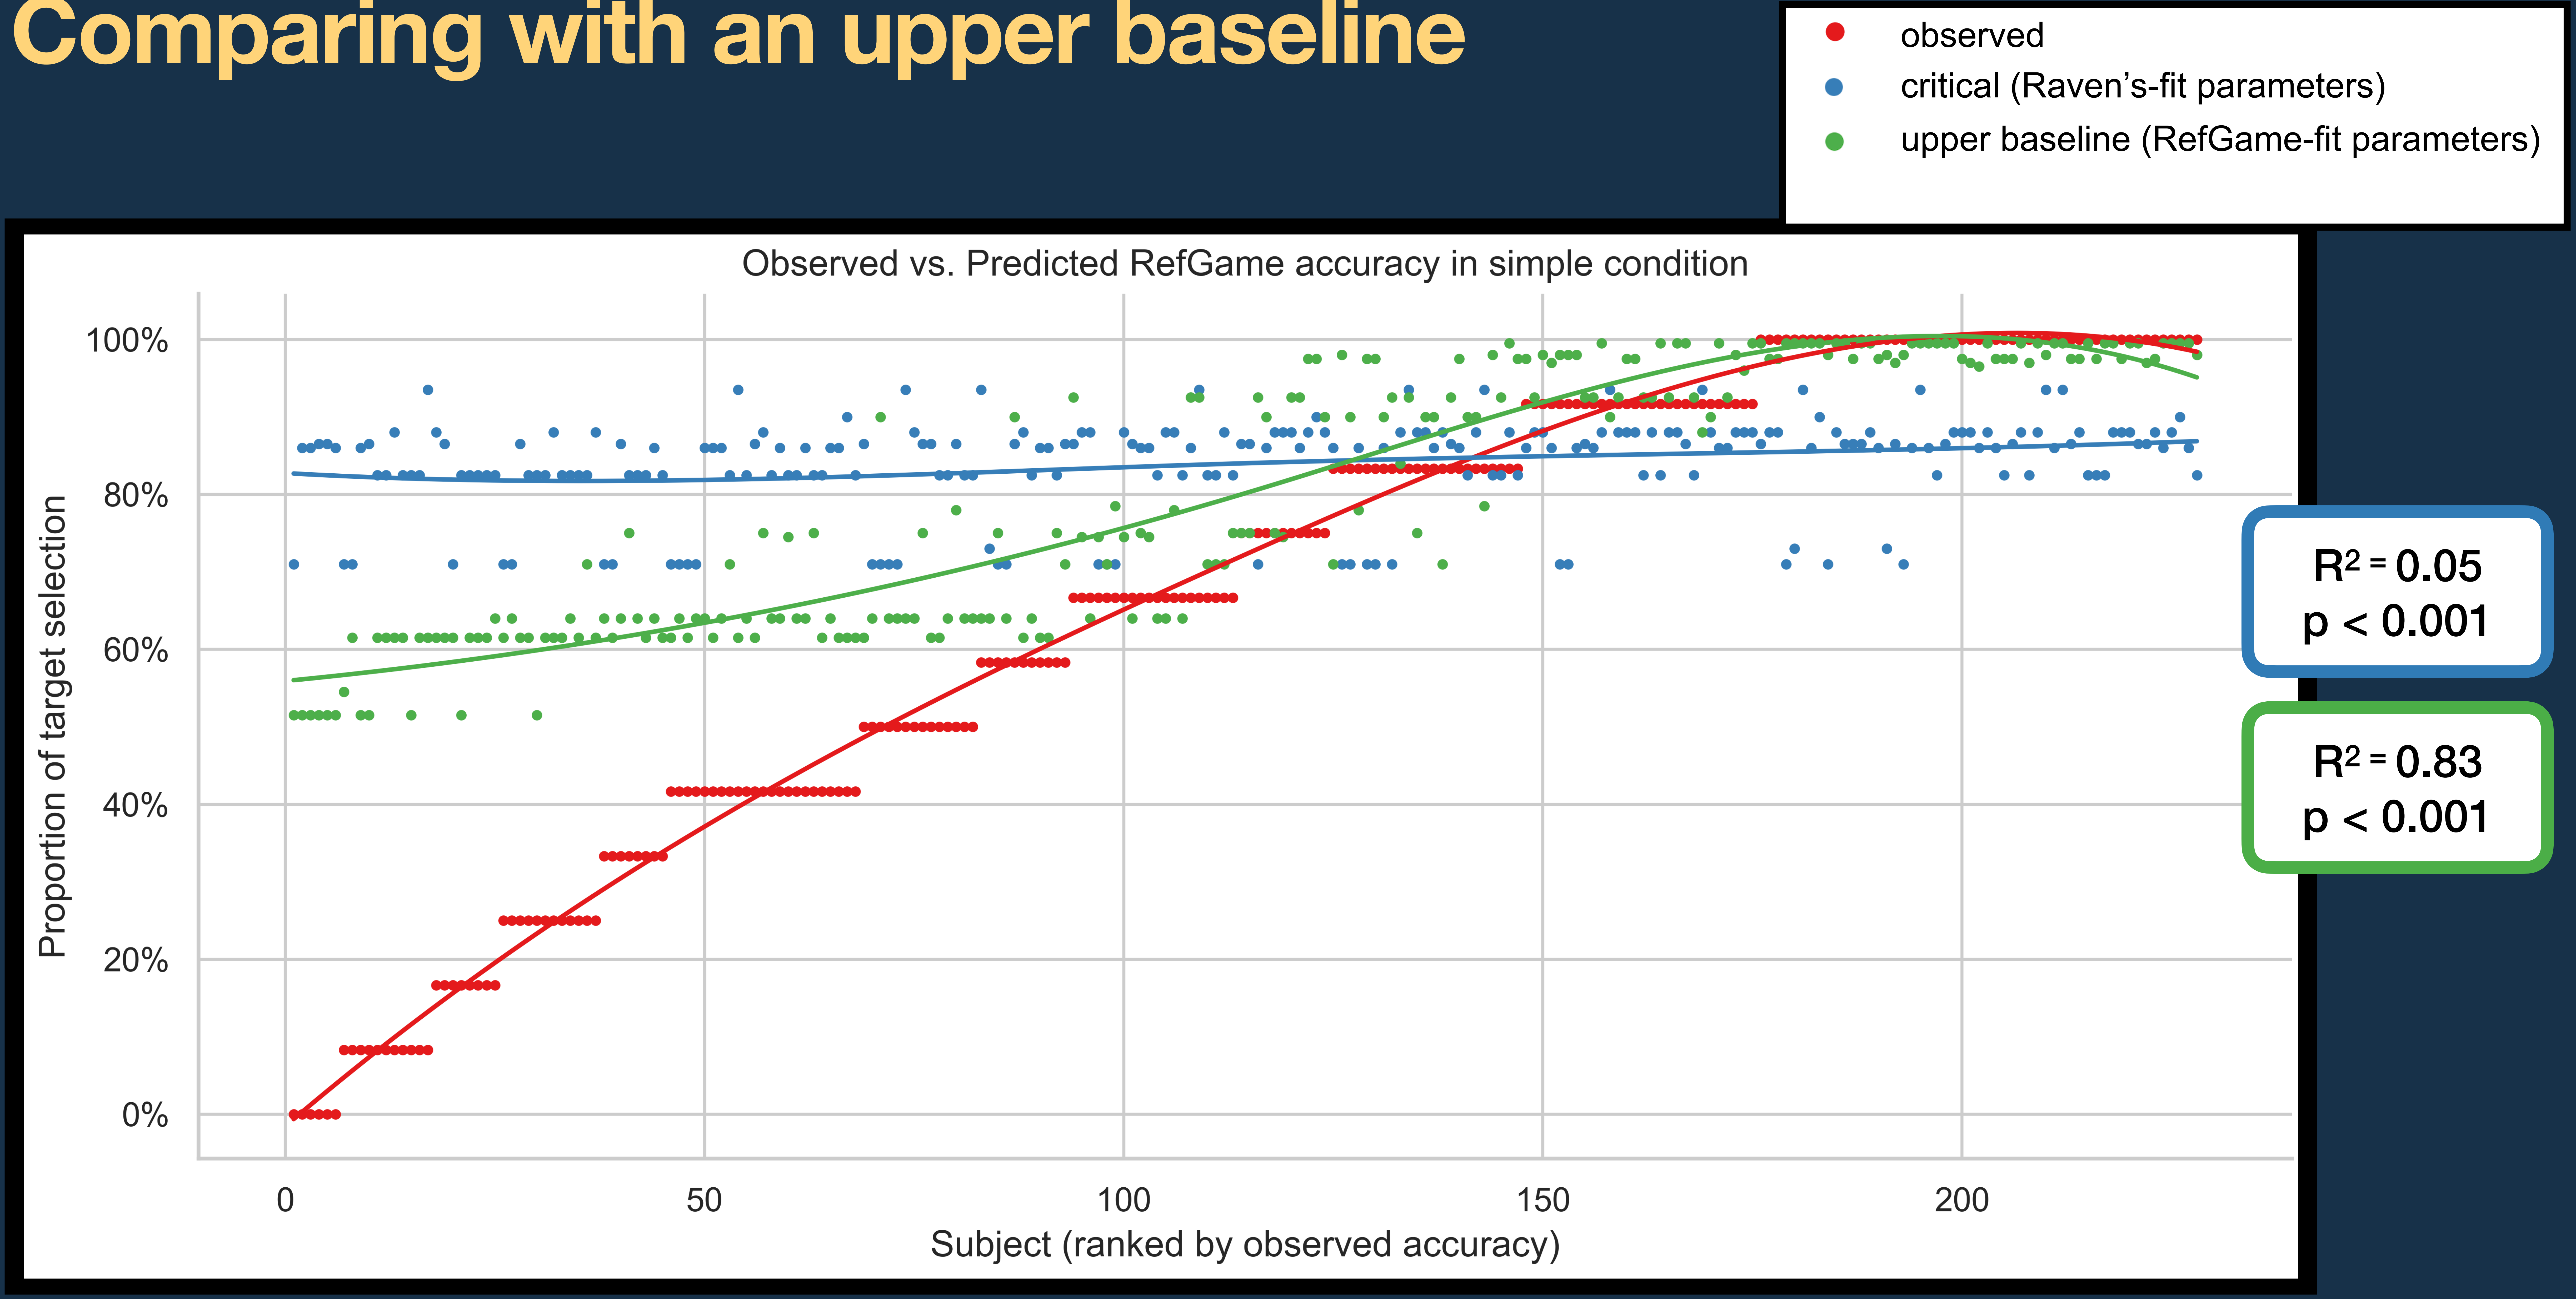

$R^2 = 0.05$
$p < 0.001$

$R^2 = 0.83$
$p < 0.001$

Deriving a lower baseline

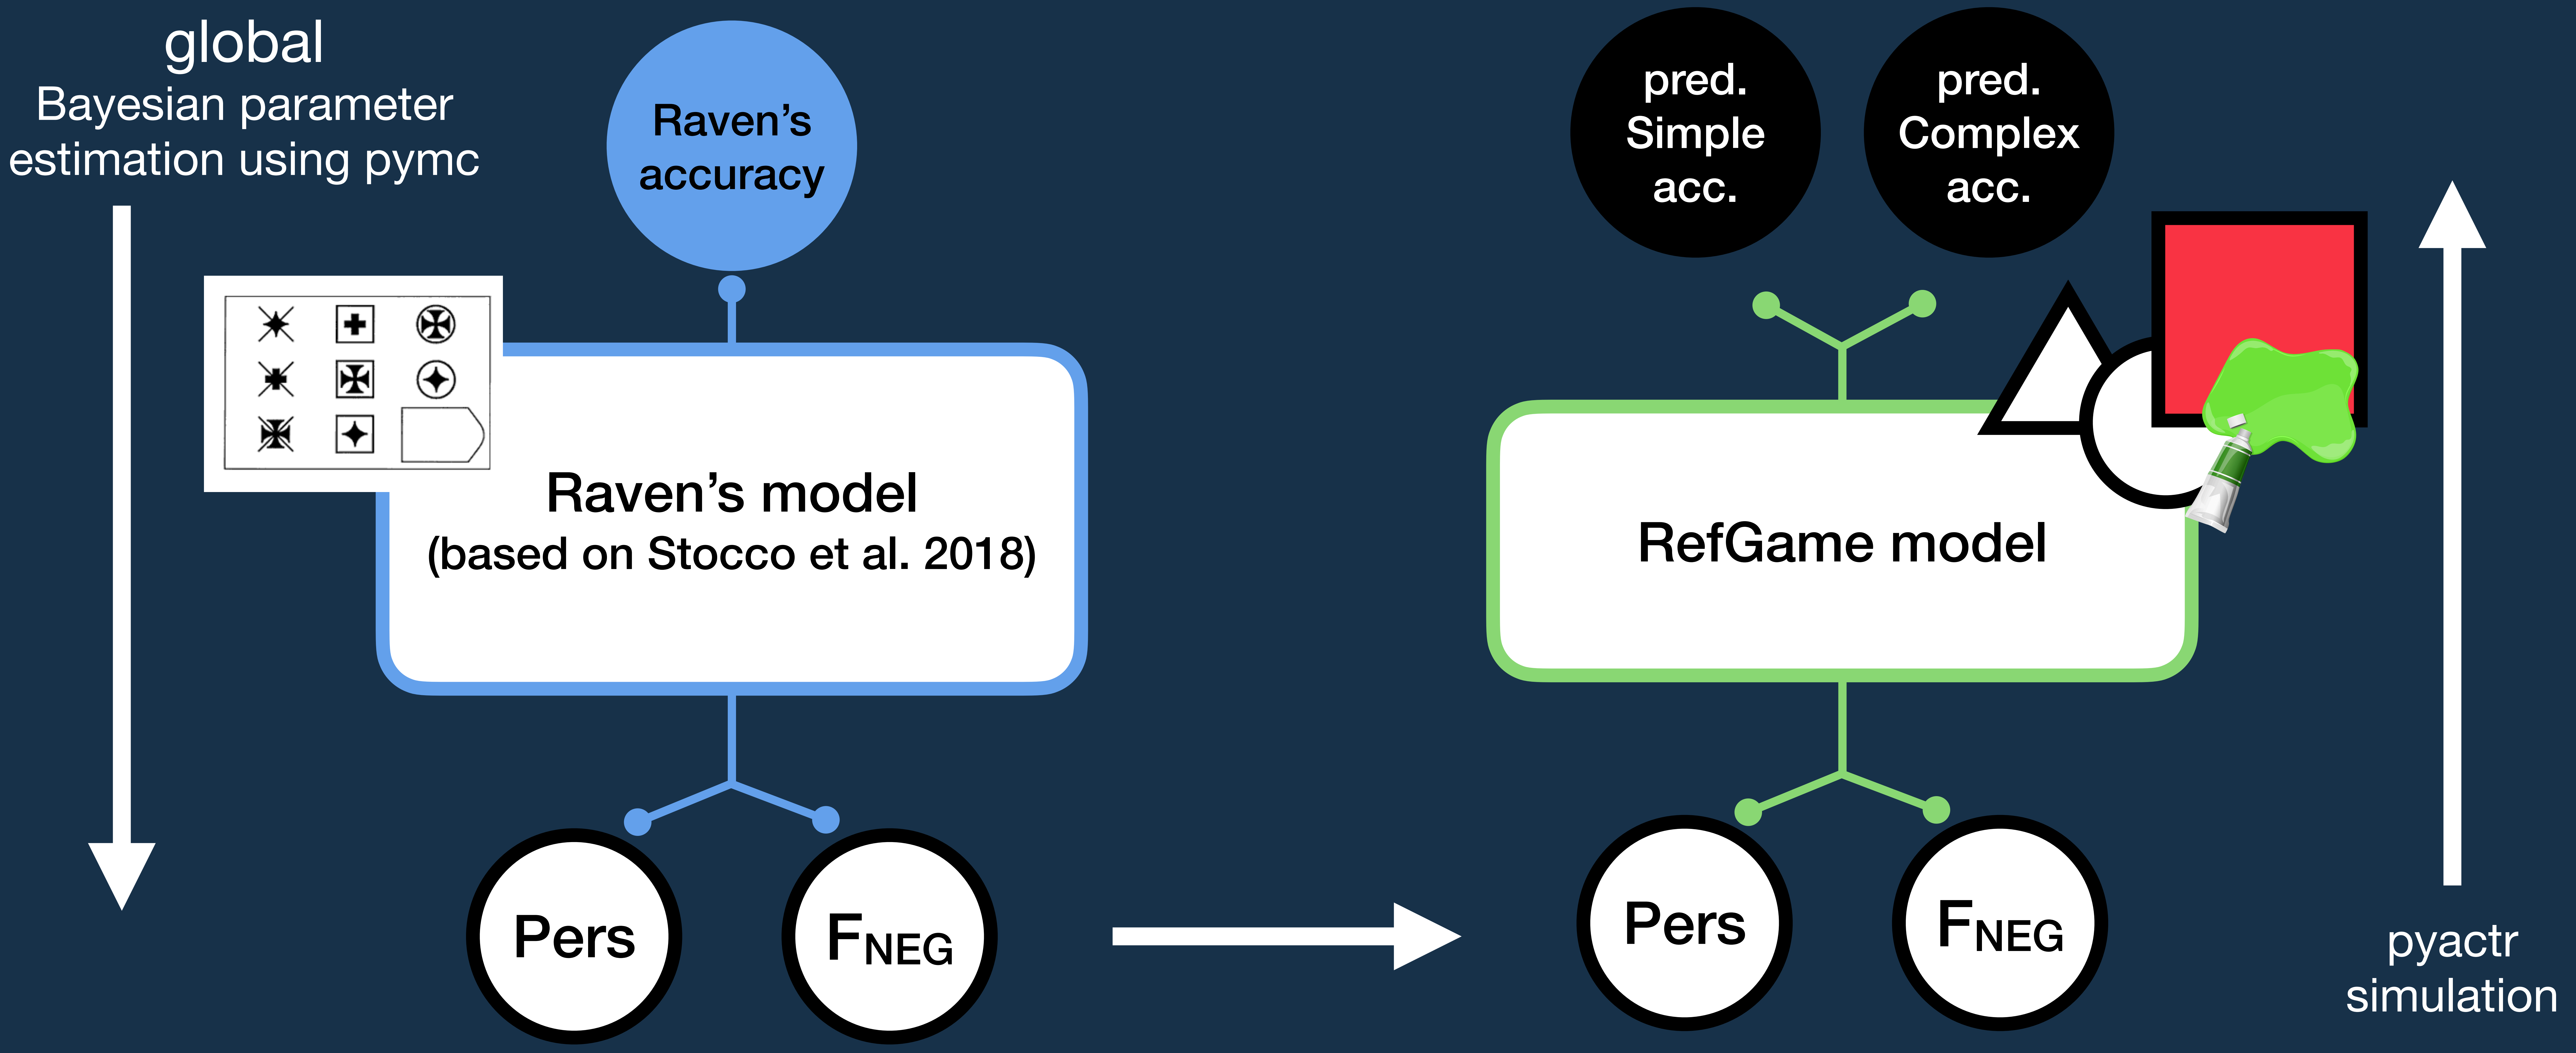

Comparing with a lower baseline

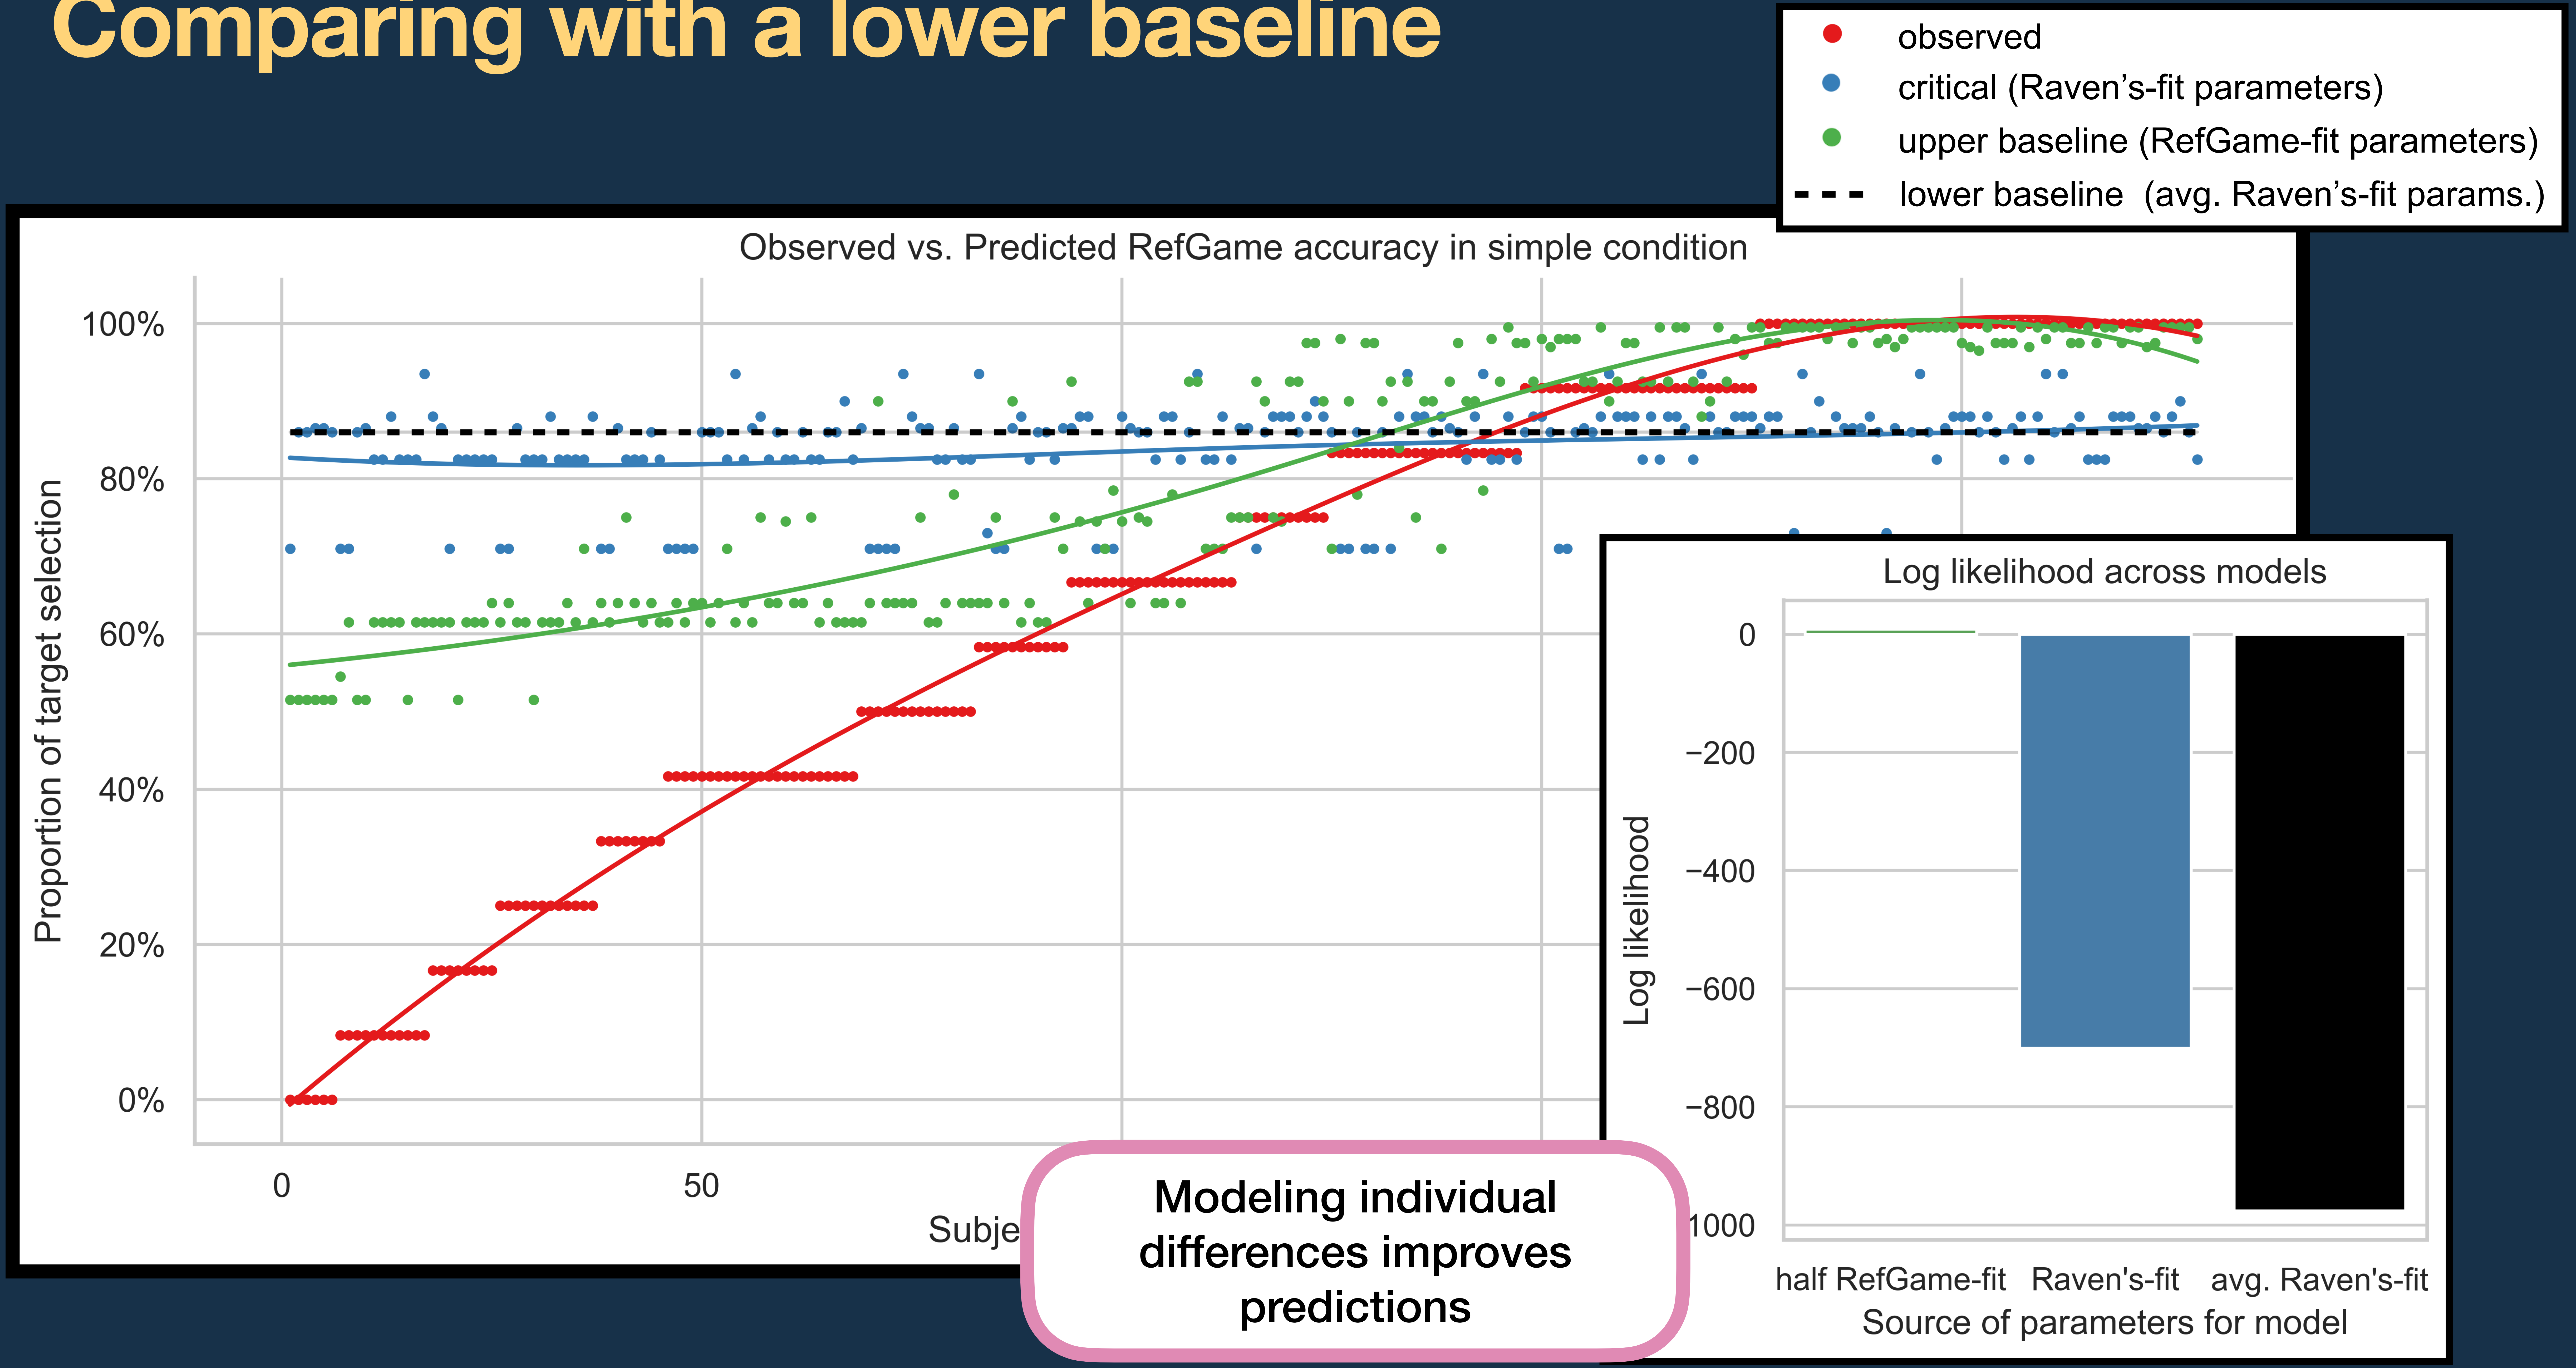

Modeling individual differences improves predictions

Introduce an ACT-R model of RefGame as a problem of strategy exploration and learning

Successfully models learning effects, individual differences, and Raven's correlation

First step towards cognitively-realistic models of pragmatic performance

Also, not shown: Experimental evidence validating the roles of persistence, F_{NEG}

In support of algorithmic-level models

- Probabilistic models of pragmatic competence (e.g. Frank & Goodman's Rational Speech Act model) have been extremely influential, but they are not models of processing

- Processing models are needed to explain a host of more complex facts:

  - On-task learning behavior

  - Evidence for inference-specific cognitive load

    (De Neys & Schaeken 2007, Marty & Chemla 2013, van Tiel et al. 2017)

  - Effects of general cognitive differences

  - Heuristics/failures of probabilistic reasoning

    (Mayn, Duff, Bila & Demberg 2024, cf. Fox et al. 2004)

Beyond the game setting

- Current model is specific to a highly controlled, novel game.
- Still, core may be plausible for ad-hoc inferences in natural comprehension:
  - Rational preference to avoid effort
  - Search for alternative meanings triggered by low informativity/relevance
  - Experience-based tuning of reasoning depth for a given interaction
- Indeed, Raven's scores also correlate with ad-hoc atypicality inferences.

(Ryzhova, Mayn & Demberg 2023)

- We aim to extend our model in this direction.

Alexandra Mayn

Vera Demberg

ERC Grant #948878 to V. Demberg,
“Individualized interactions in discourse”

Thanks!

Thanks also to Sebastian Schuster, Michael Frank, and Niels Taatgen for suggestions and feedback.

Ask us about:

- New experiments validating ID effects by measuring persistence and F_{NEG} directly
- Simulated and observed response time effects
- Related work observing probability fallacies in first-order reasoning
- Details of the model

(n = 150, 8 obs./cond. + 20 trivial)

New data: Independent persistence measures

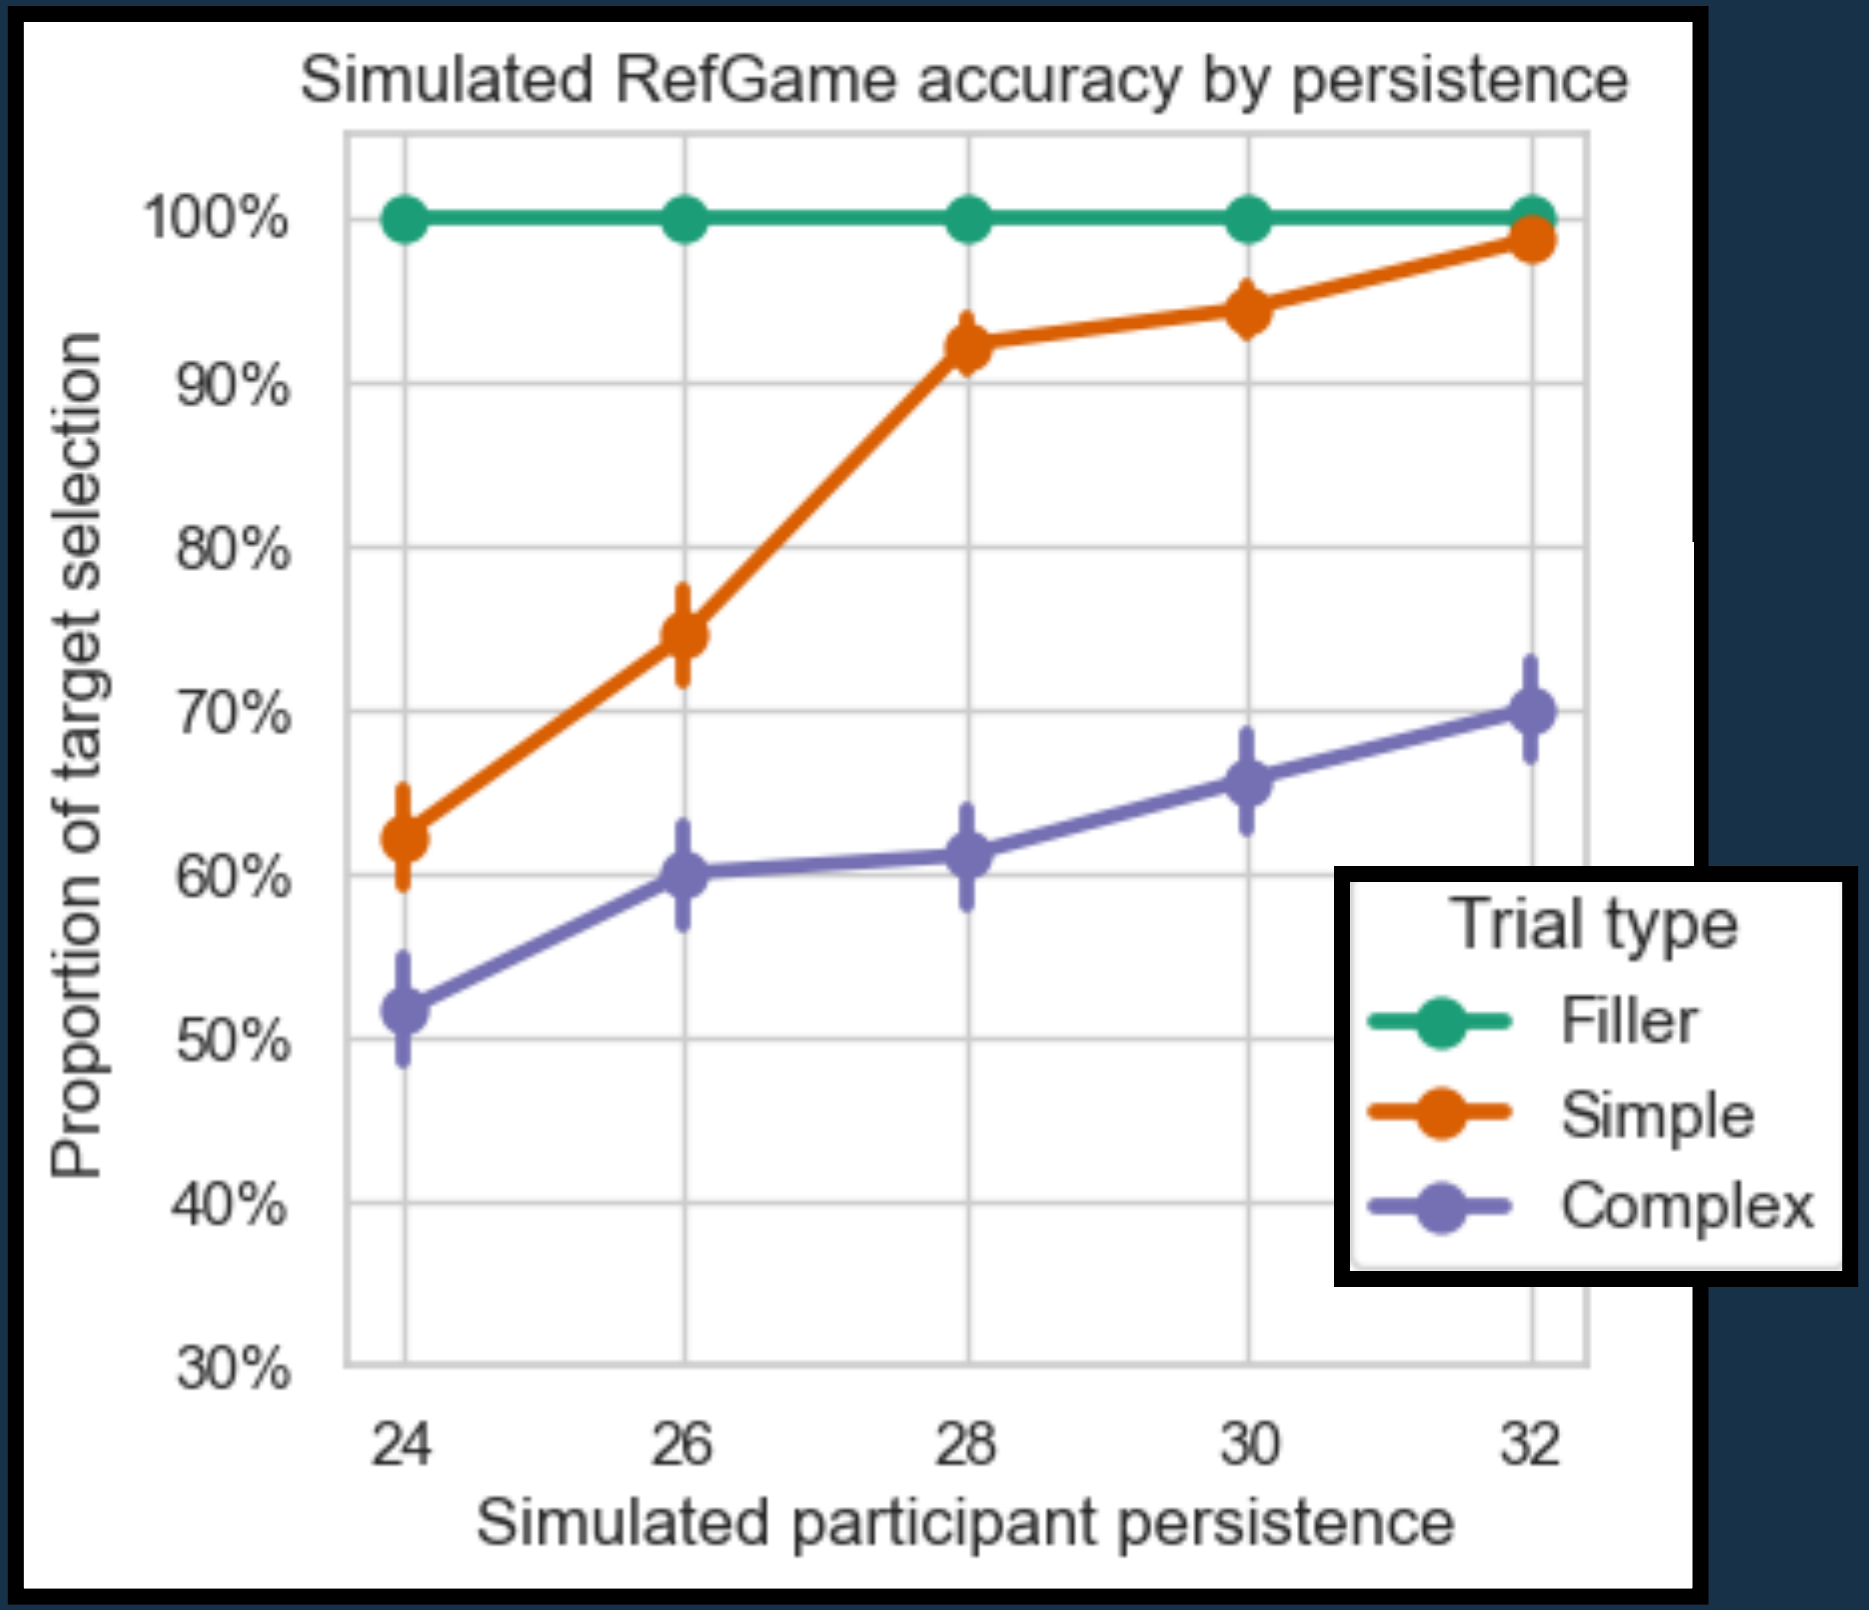

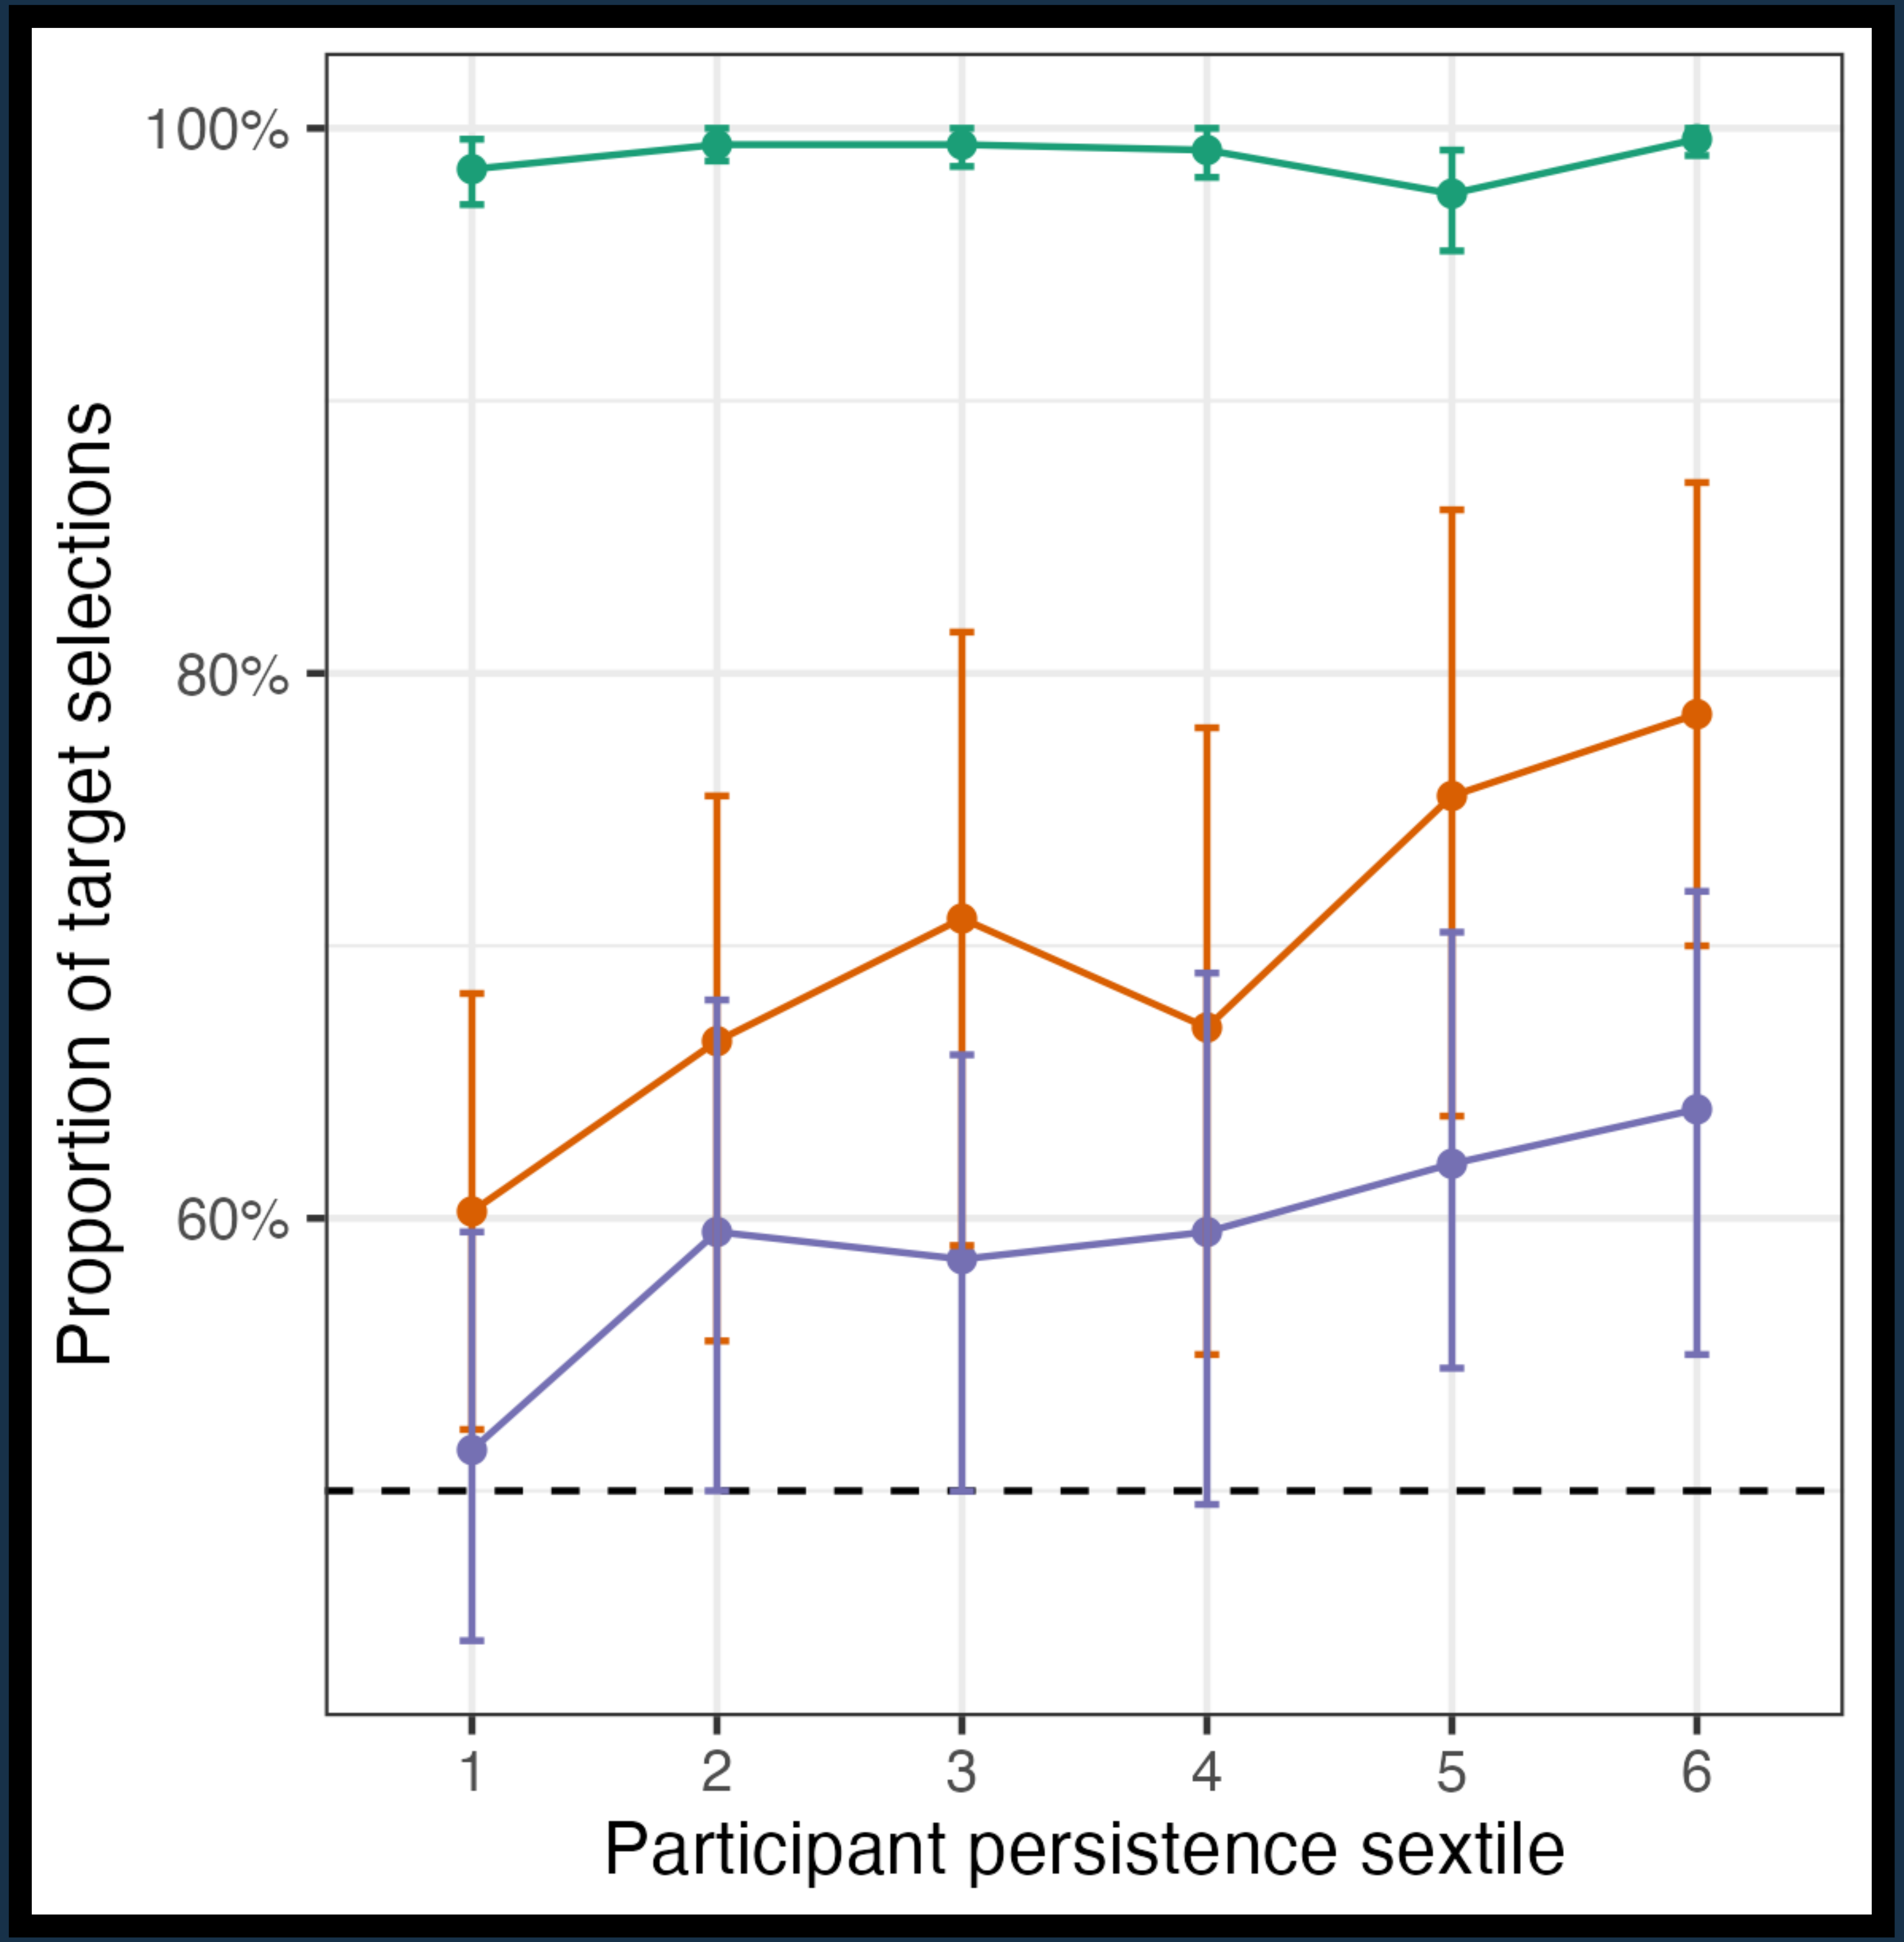

Model β_{PERSIST} = $(0.83, 0.88)_{95\%}$

Human β_{PERSIST} = $(0.08, 0.58)_{95\%}$

(from brms logistic regr. with uninf. priors, ID predictors were z-scaled)

Measuring Persistence:

Impossible Anagrams

(Ventura & Shute 2013)

(see also Eisenberg & Leonard 1980; Dale et al. 2018)

rveir	kjoer	ardot
(easy)	(hard)	(impossible)

Anagram Persistence:

$\text{SkipTime}_{\text{IMPOSS}}$ / $\text{Correct RT}_{\text{EASY}}$

- Also correlated with:
  - Time spent on (task-final) impossible Raven's problem — $R = 0.18$
    (Dale et al. 2018)
  - Grit score derived from self-assessment — $R = 0.20$
    (Duckworth & Quinn 2009)

(*n* = 150, 8 obs./cond. + 20 trivial)

New data: Independent F_{NEG} measures

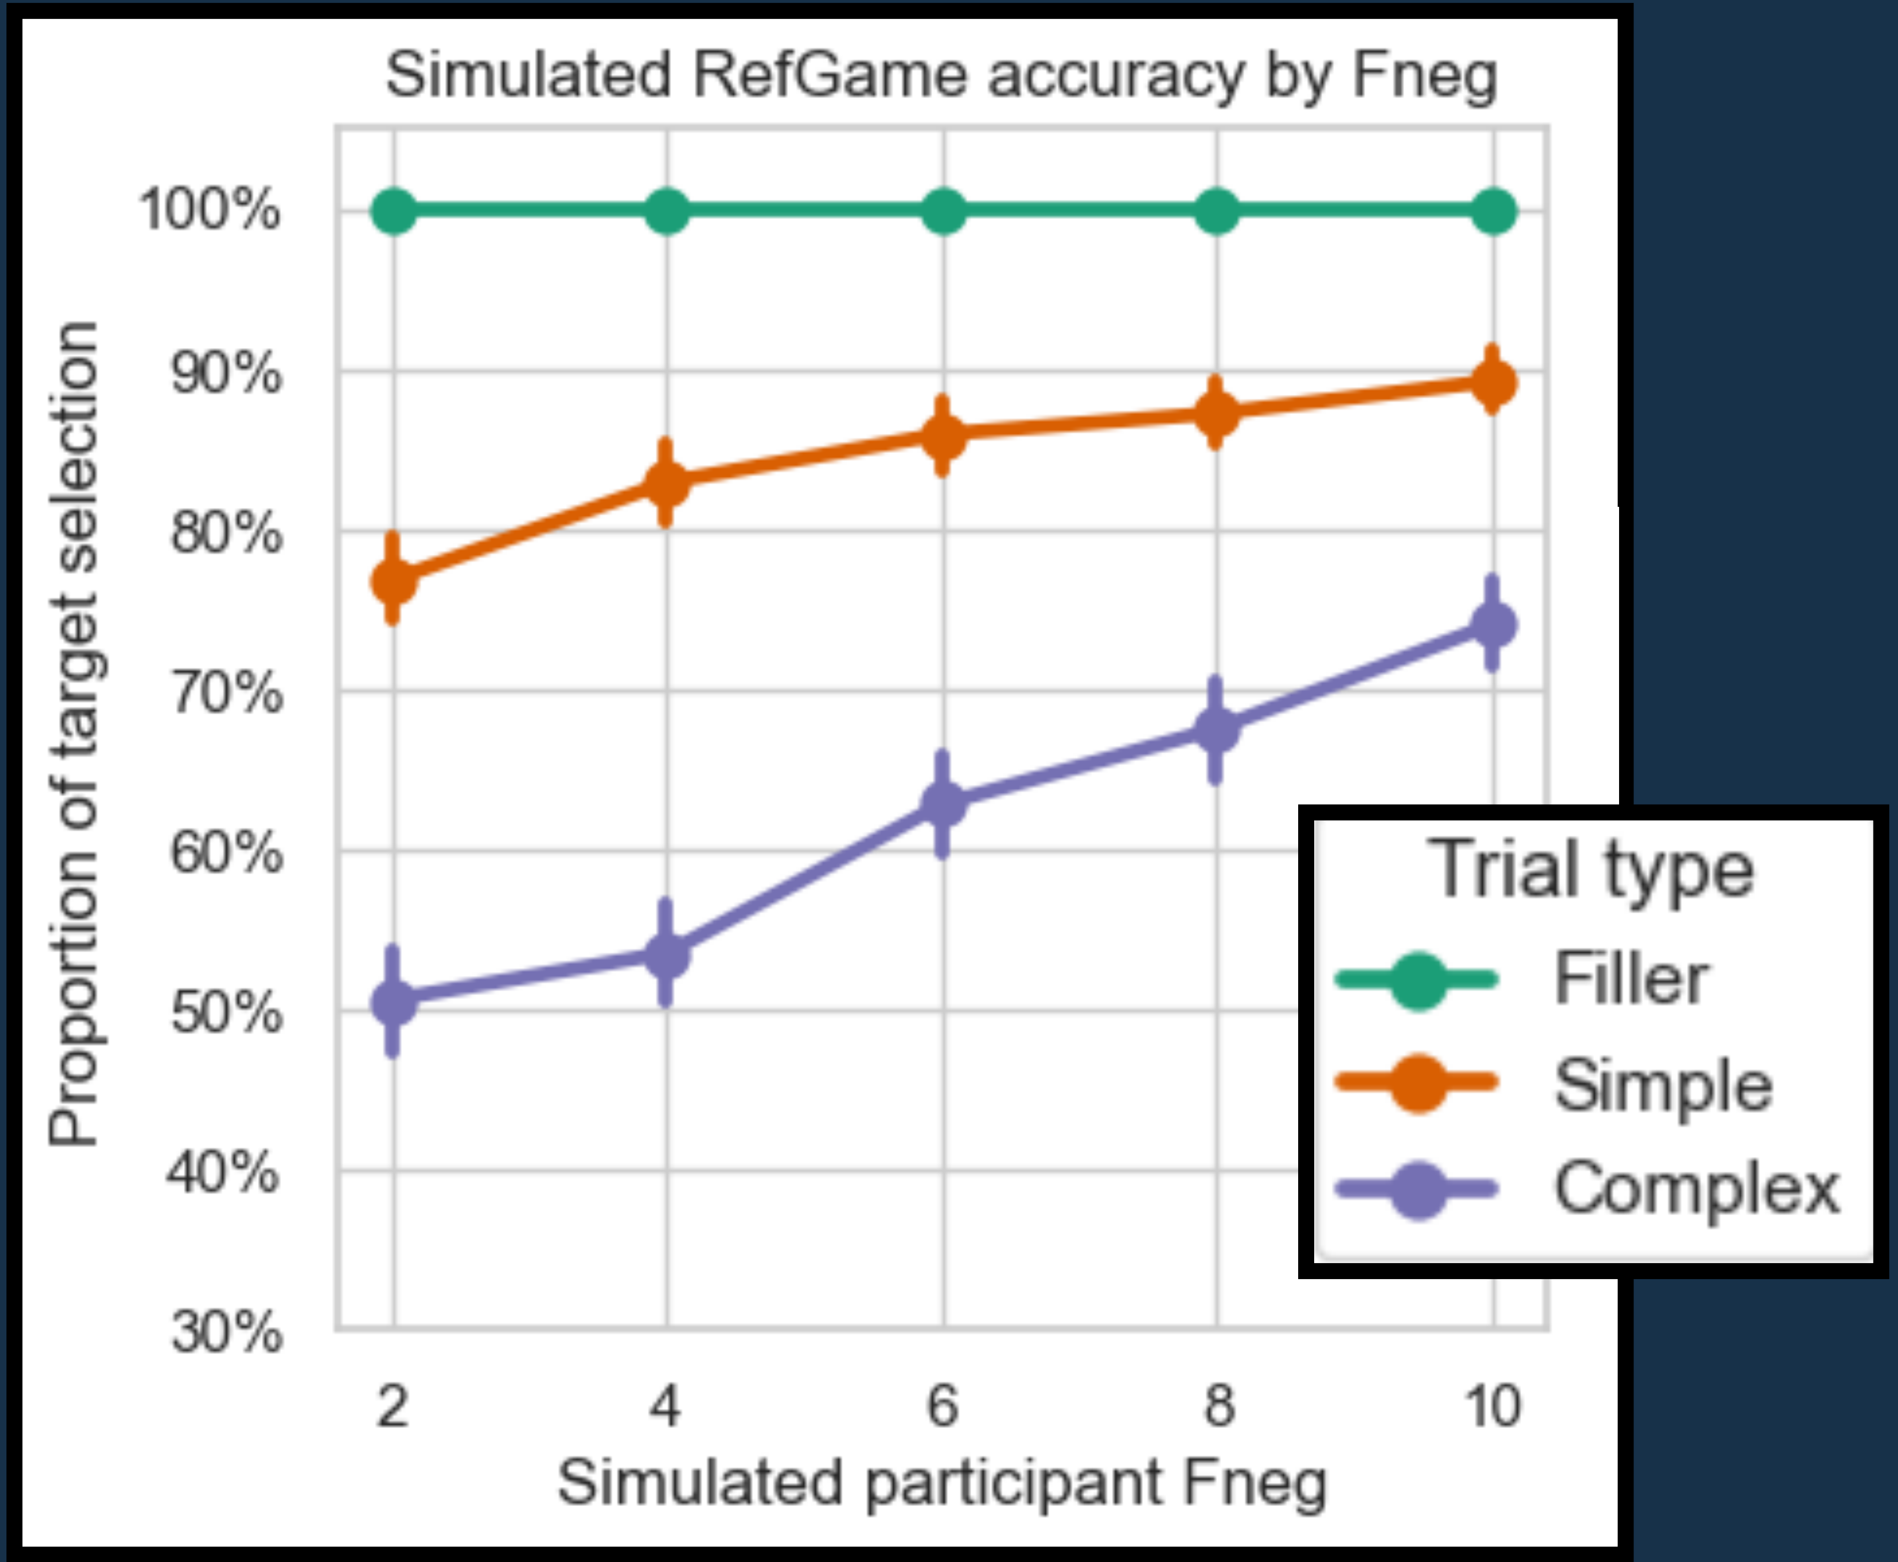

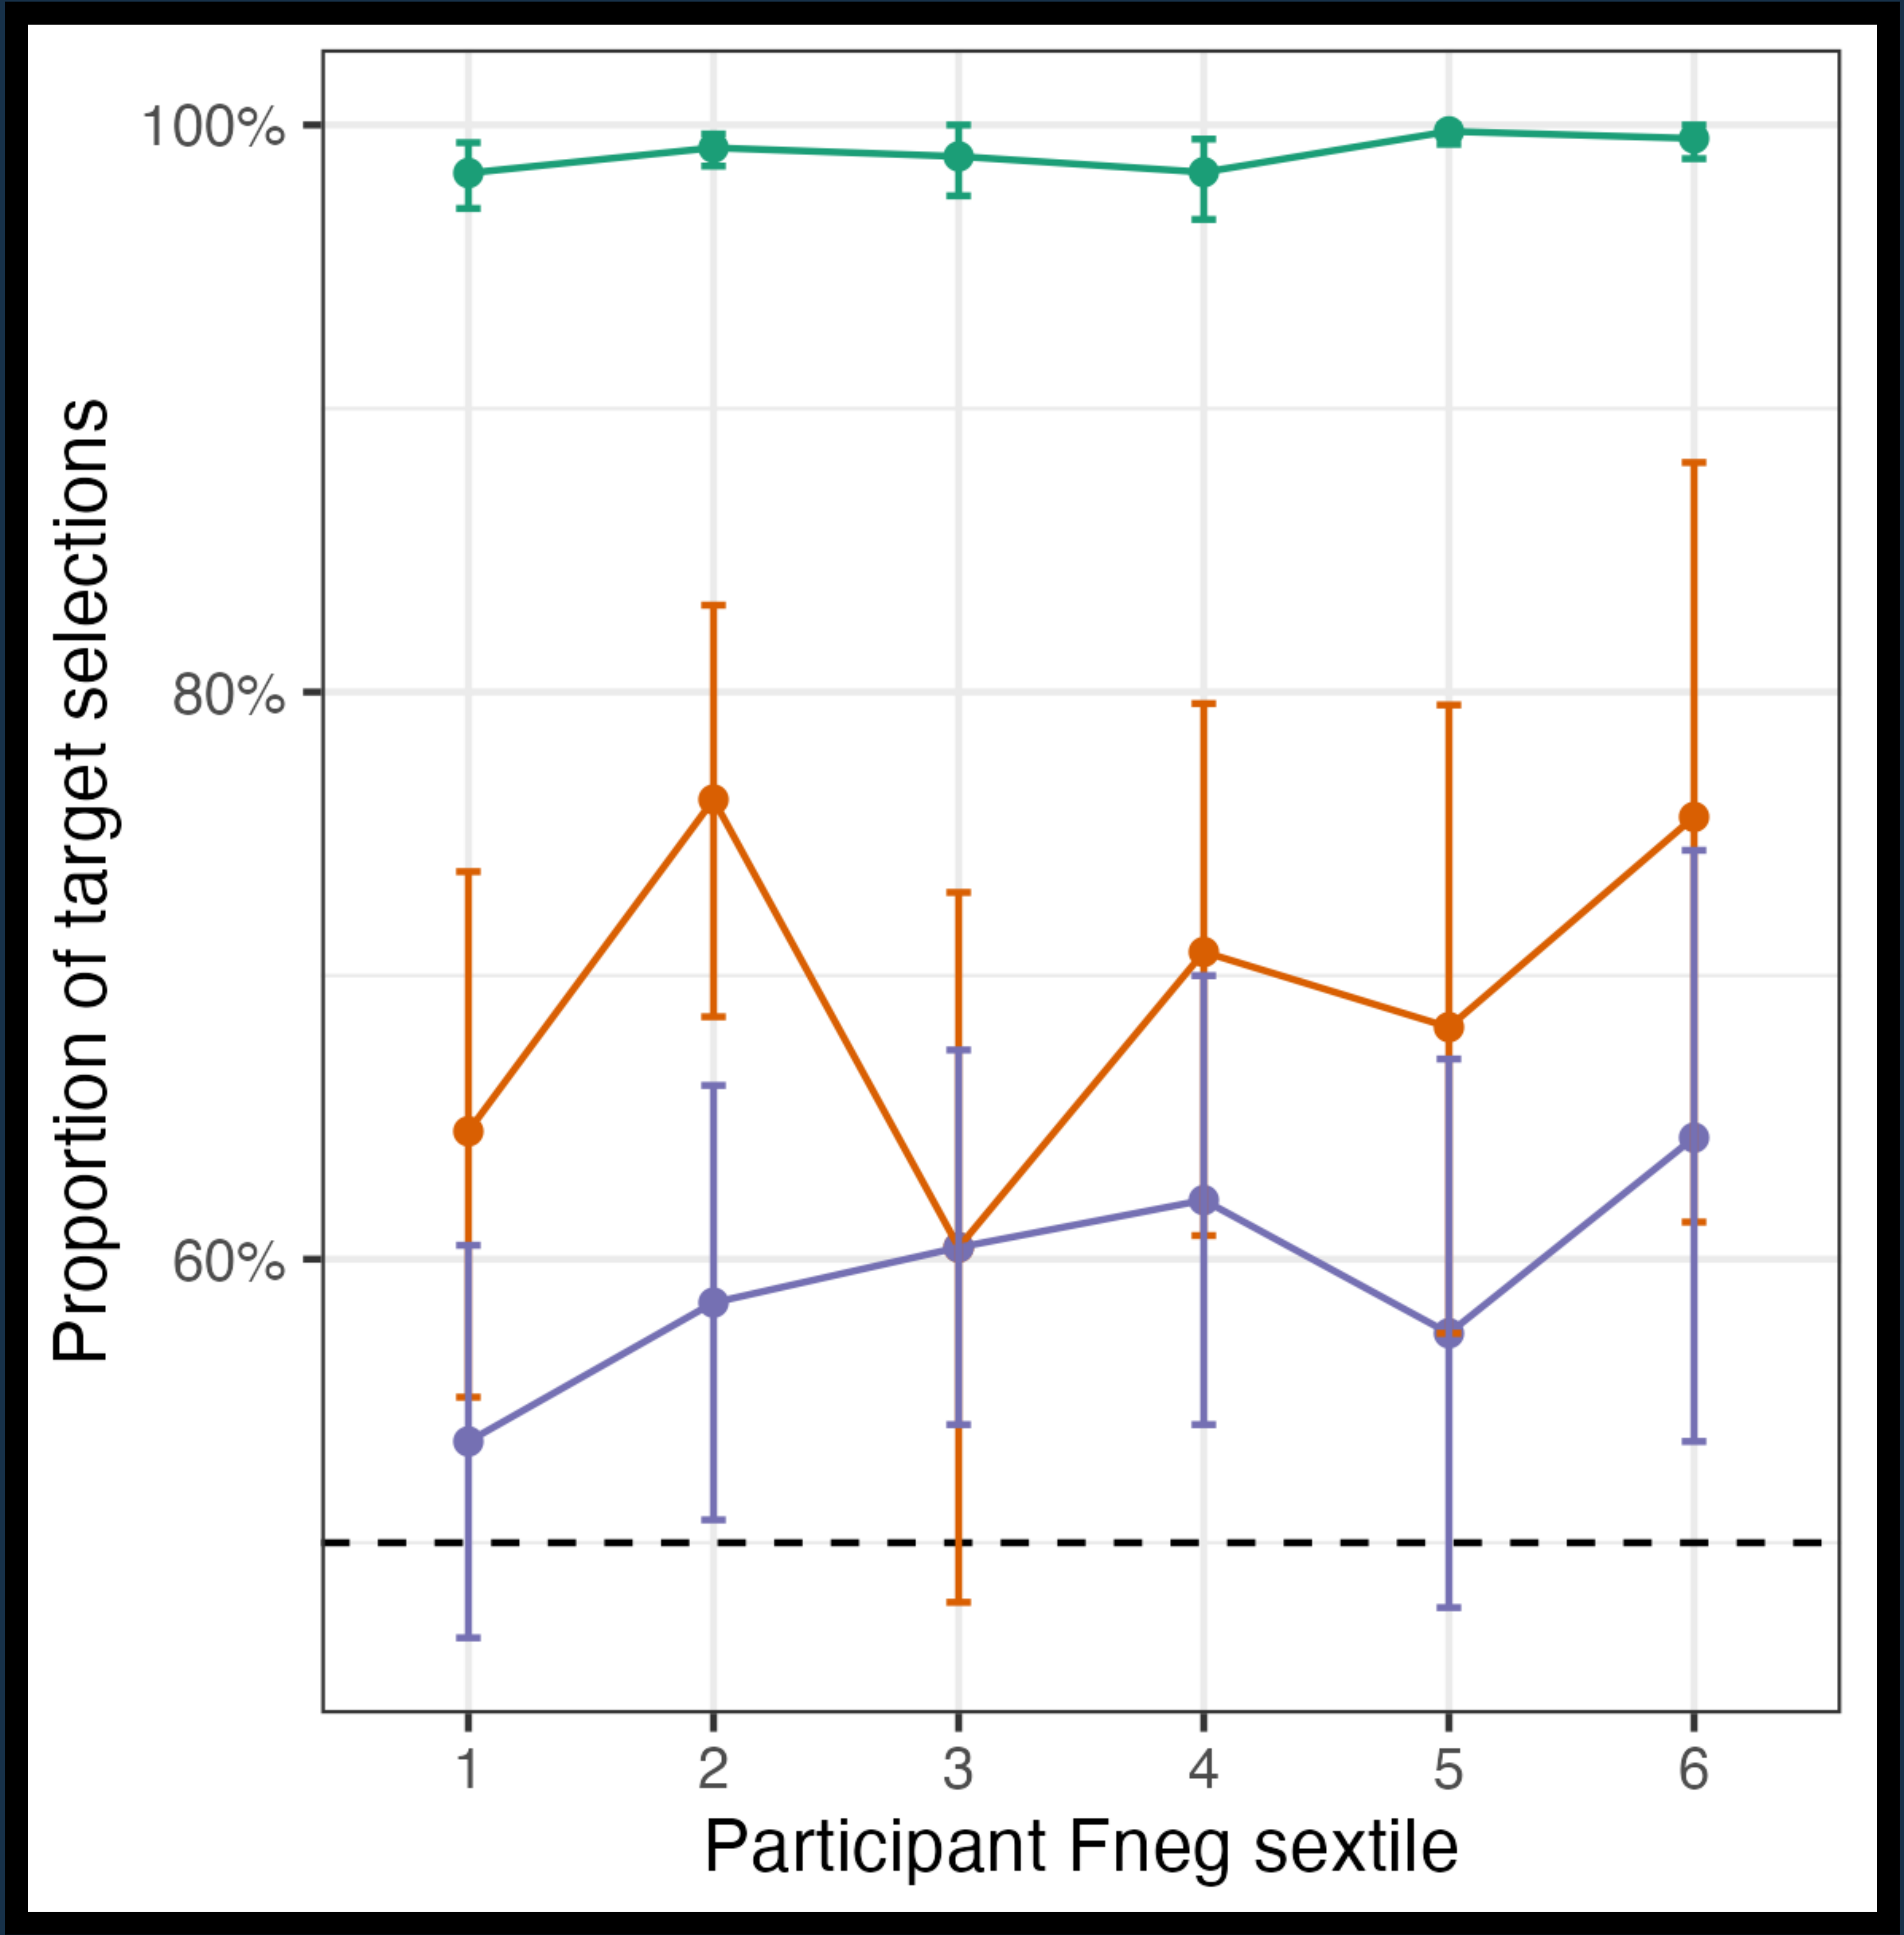

Model β_{FNEG} = $(0.53, 0.58)_{95\%}$

Human β_{FNEG} = $(-0.05, 0.40)_{95\%}$

(from brms logistic regr. with uninf. priors, ID predictors were z-scaled)

Measuring F_{NEG}:

The Probabilistic Stimulus Selection task

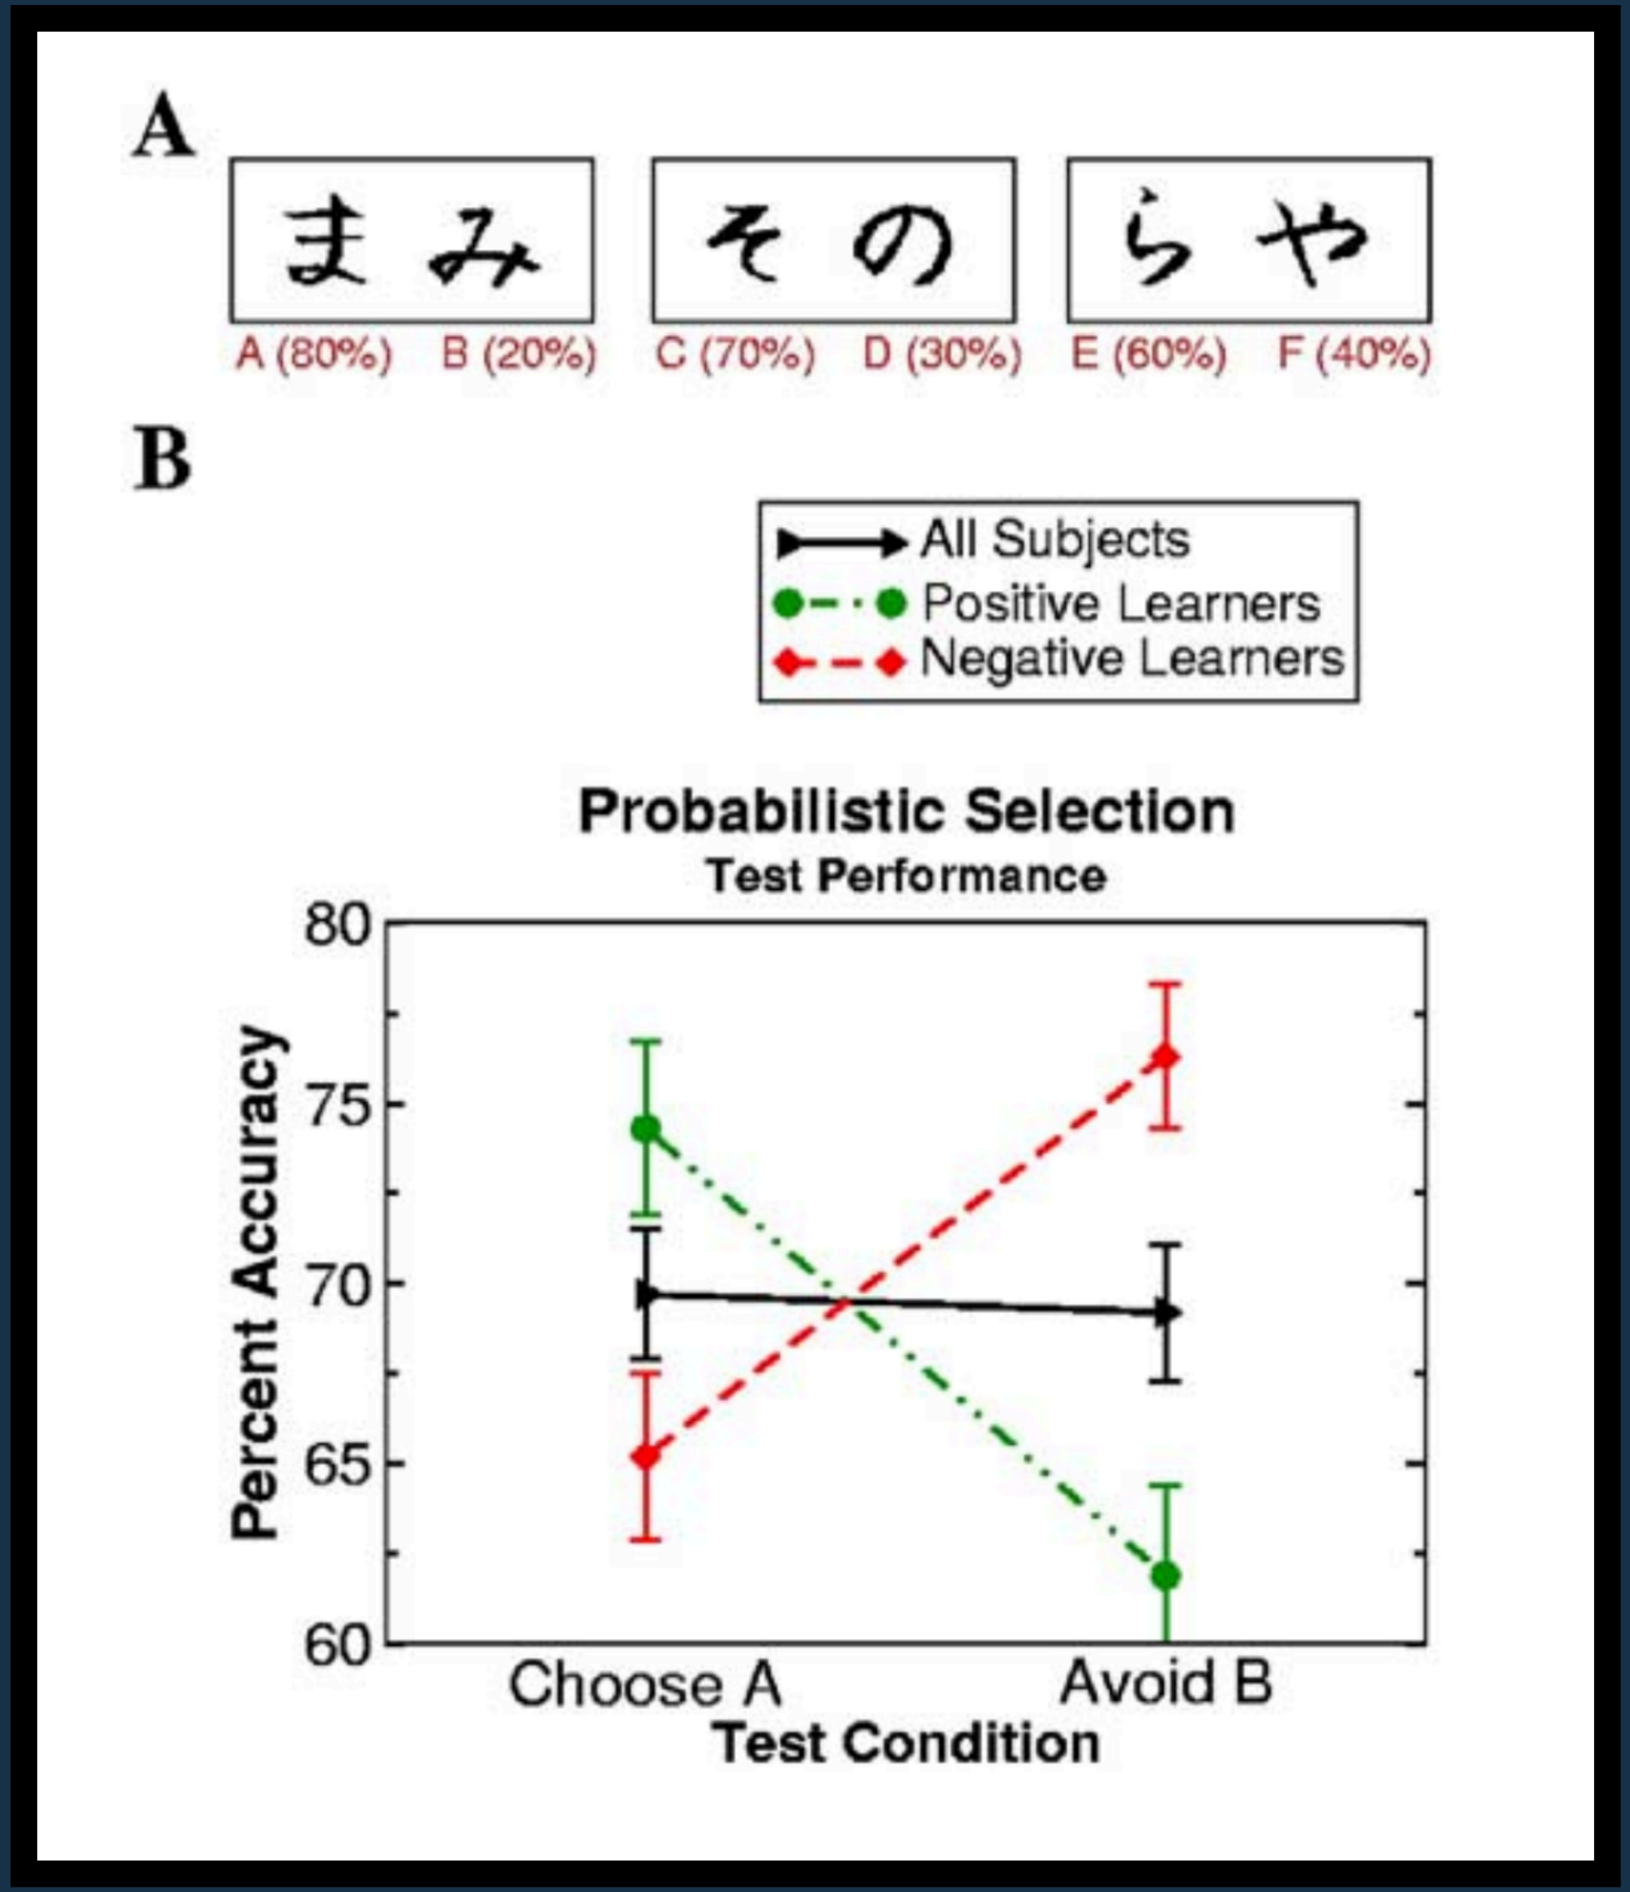

(Frank et al. 2004, 2005, 2007)

- Two pathways to learn from experiences where A is a better choice than B:
  - Learn positive value of A (via F_{POS})
  - Learn negative value of B (via F_{NEG})
  - Measure independently on test phase
- Corresponds to individual differences in dopamine levels in basal ganglia, and error-related negativity in ERPs.

New data: Learning

(n = 150, 8 obs./cond. + 20 trivial)

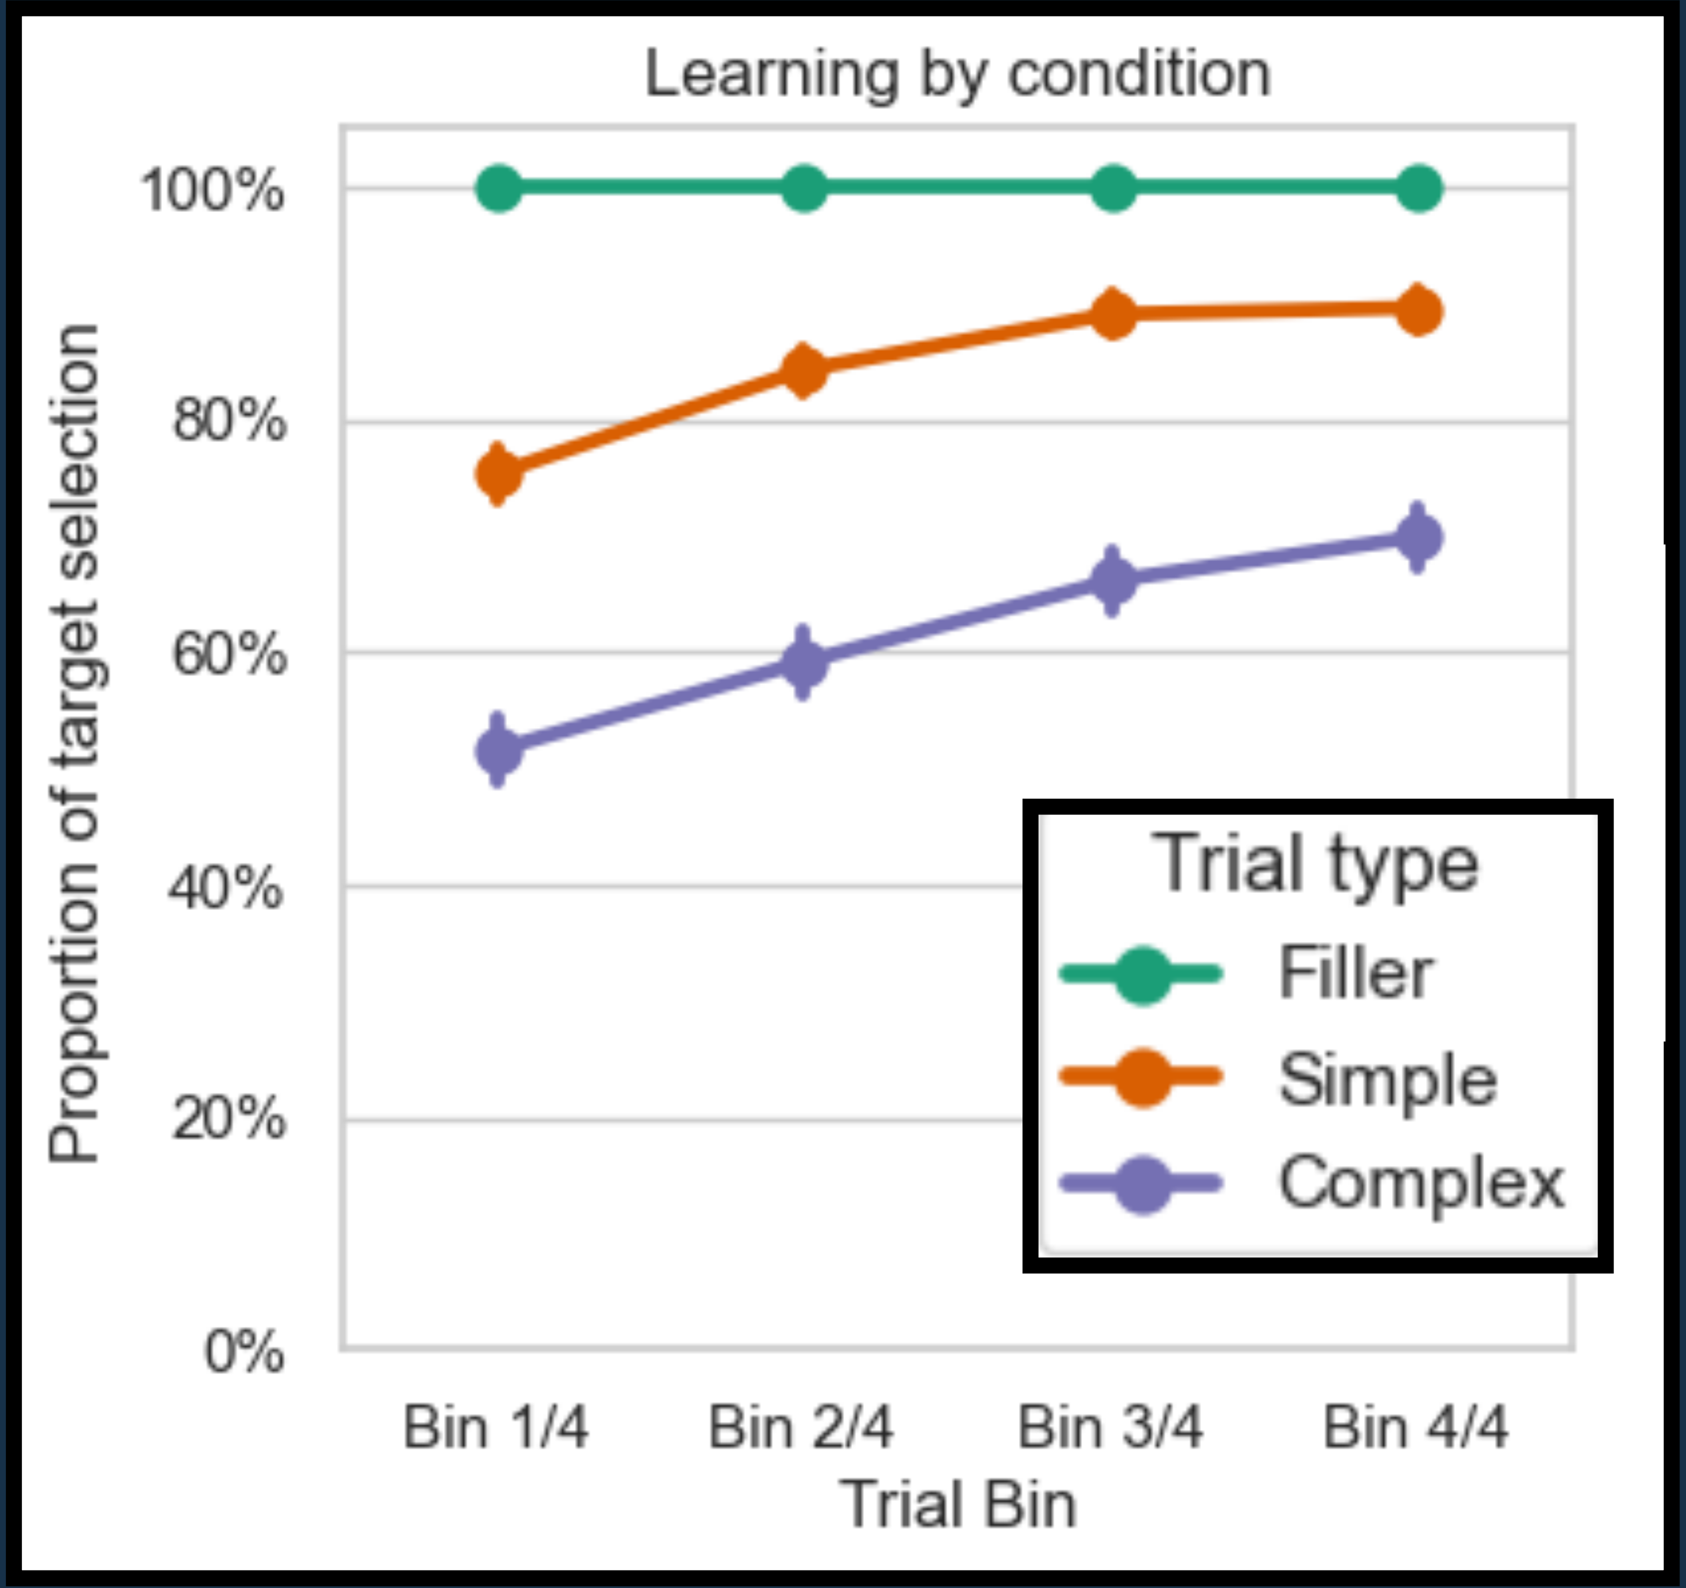

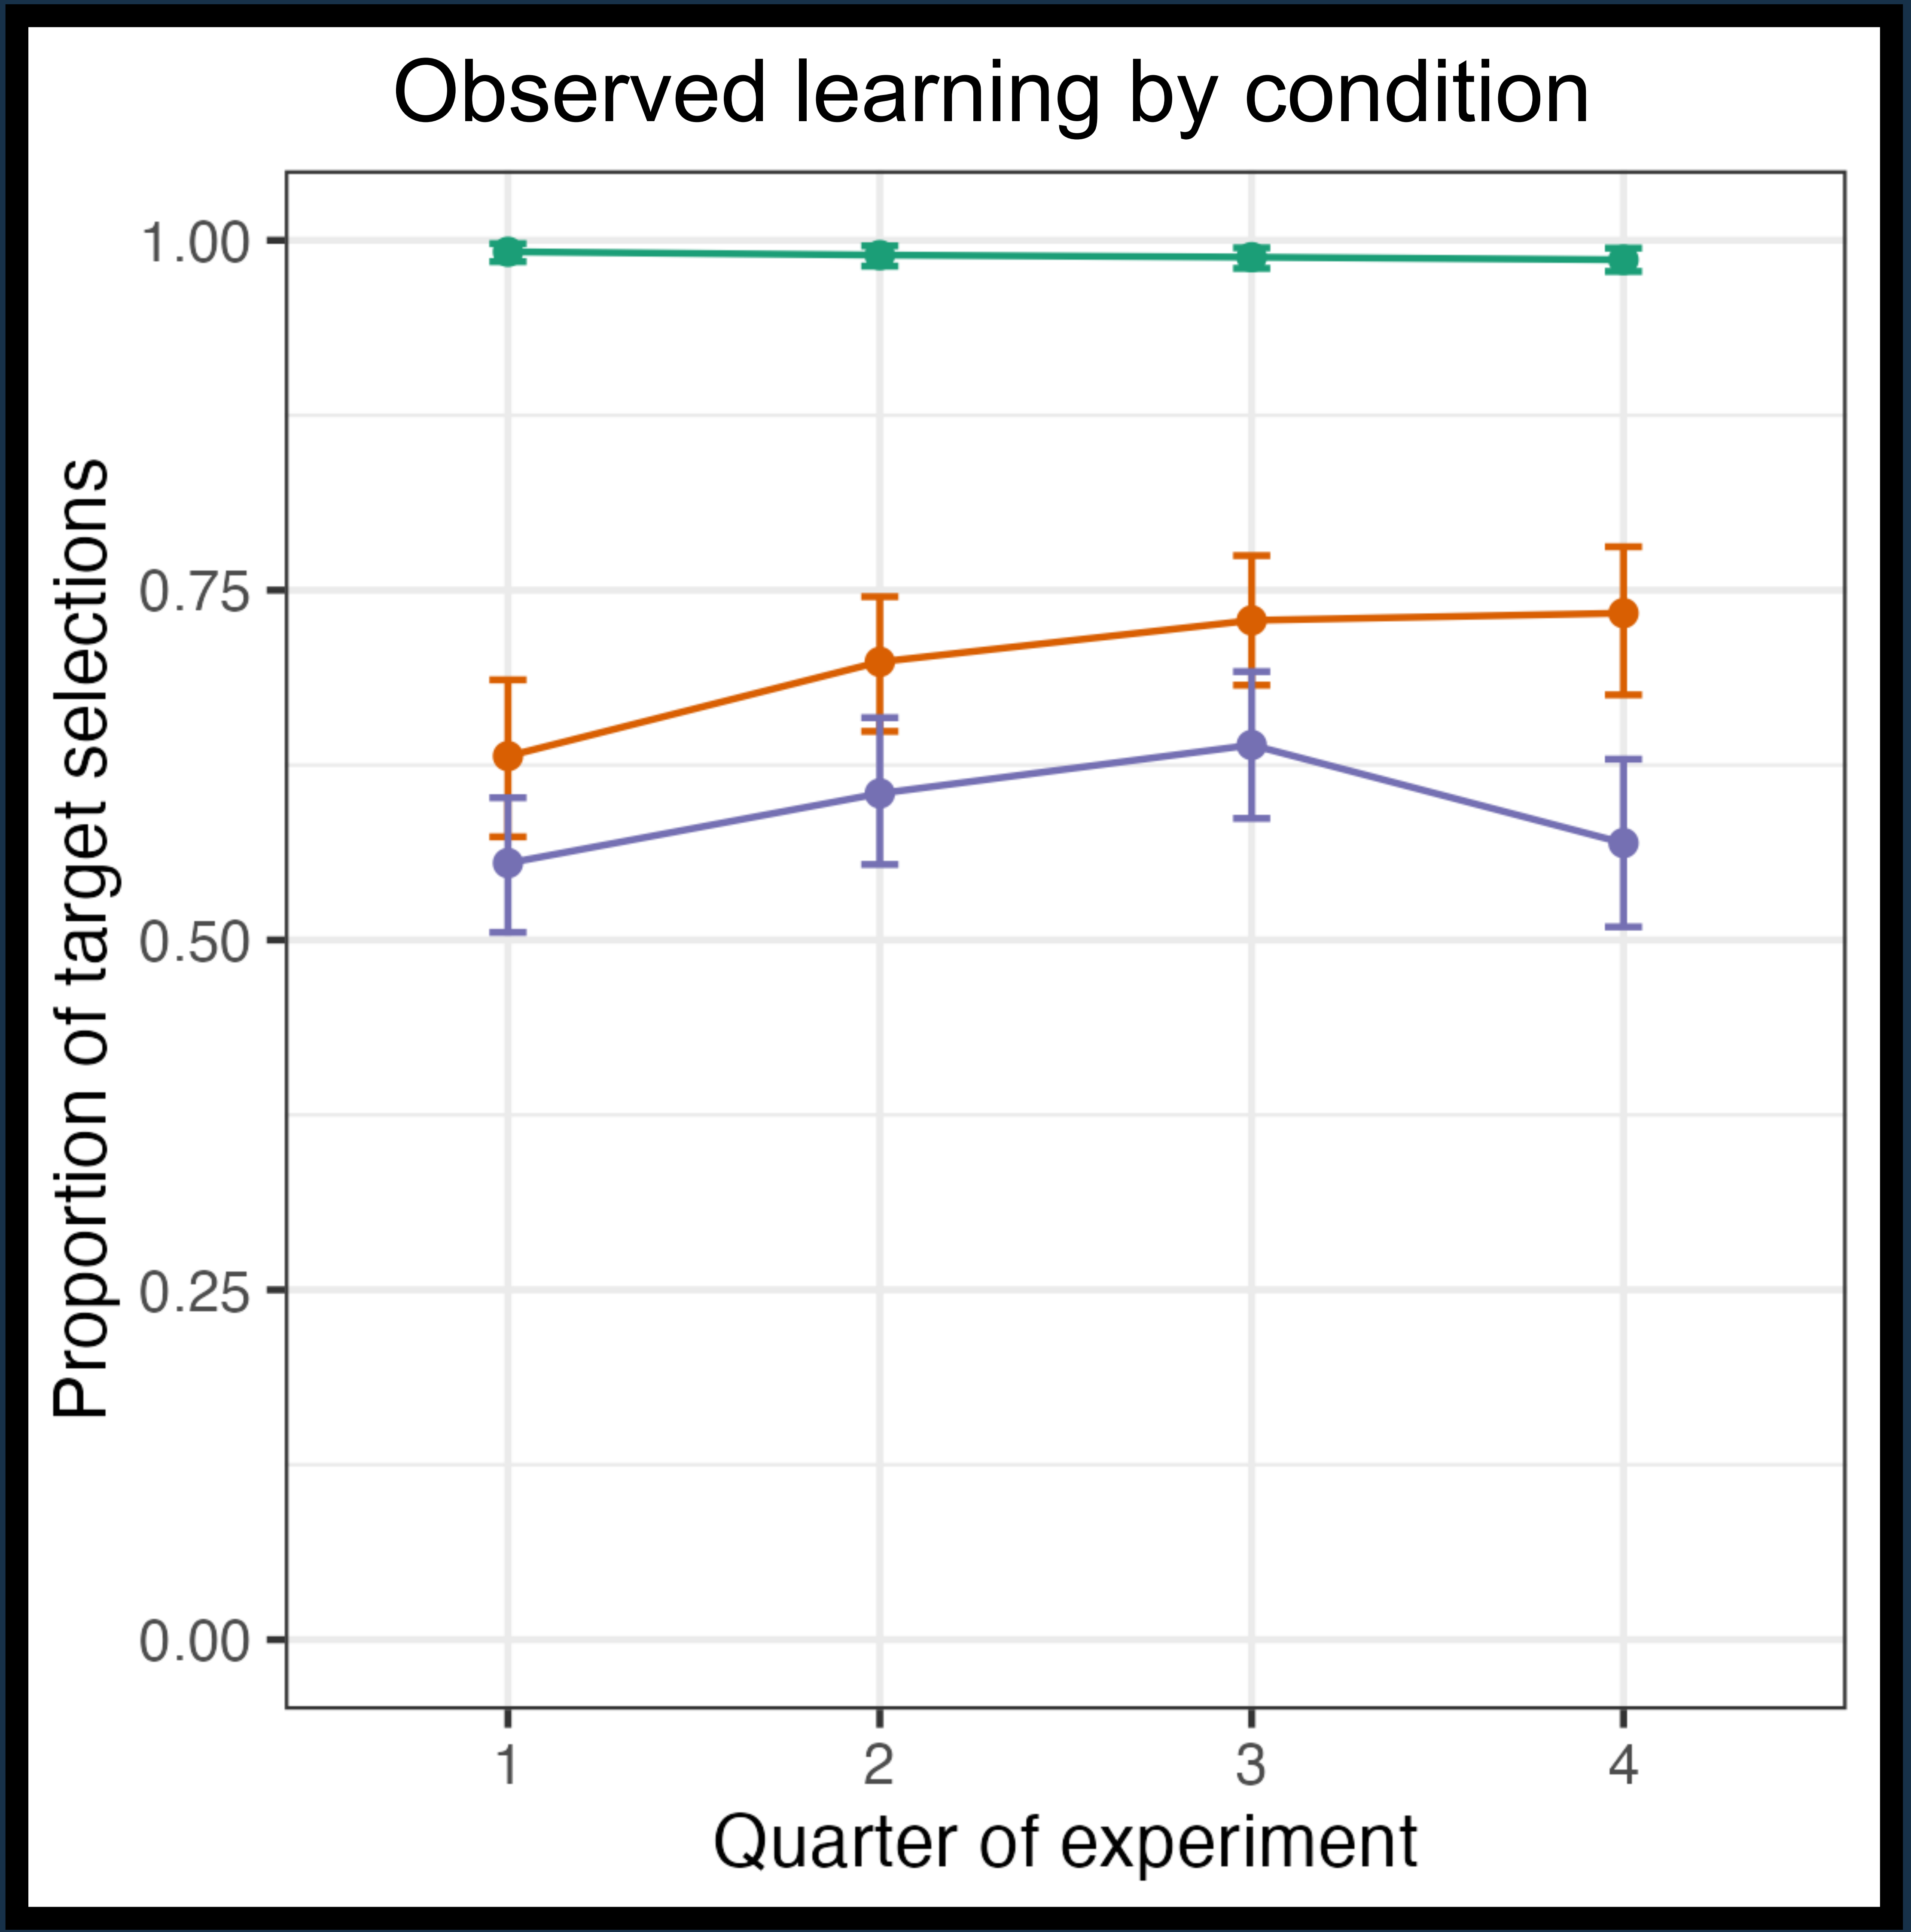

Model β_{FNEG} = $(0.05, 0.05)_{95\%}$

Human β_{FNEG} = $(0.01, 0.03)_{95\%}$

(from brms logistic regr. with uninf. priors, trial was centered and not scaled)

New data: RTs

(n = 150, 8 obs./cond. + 20 trivial)

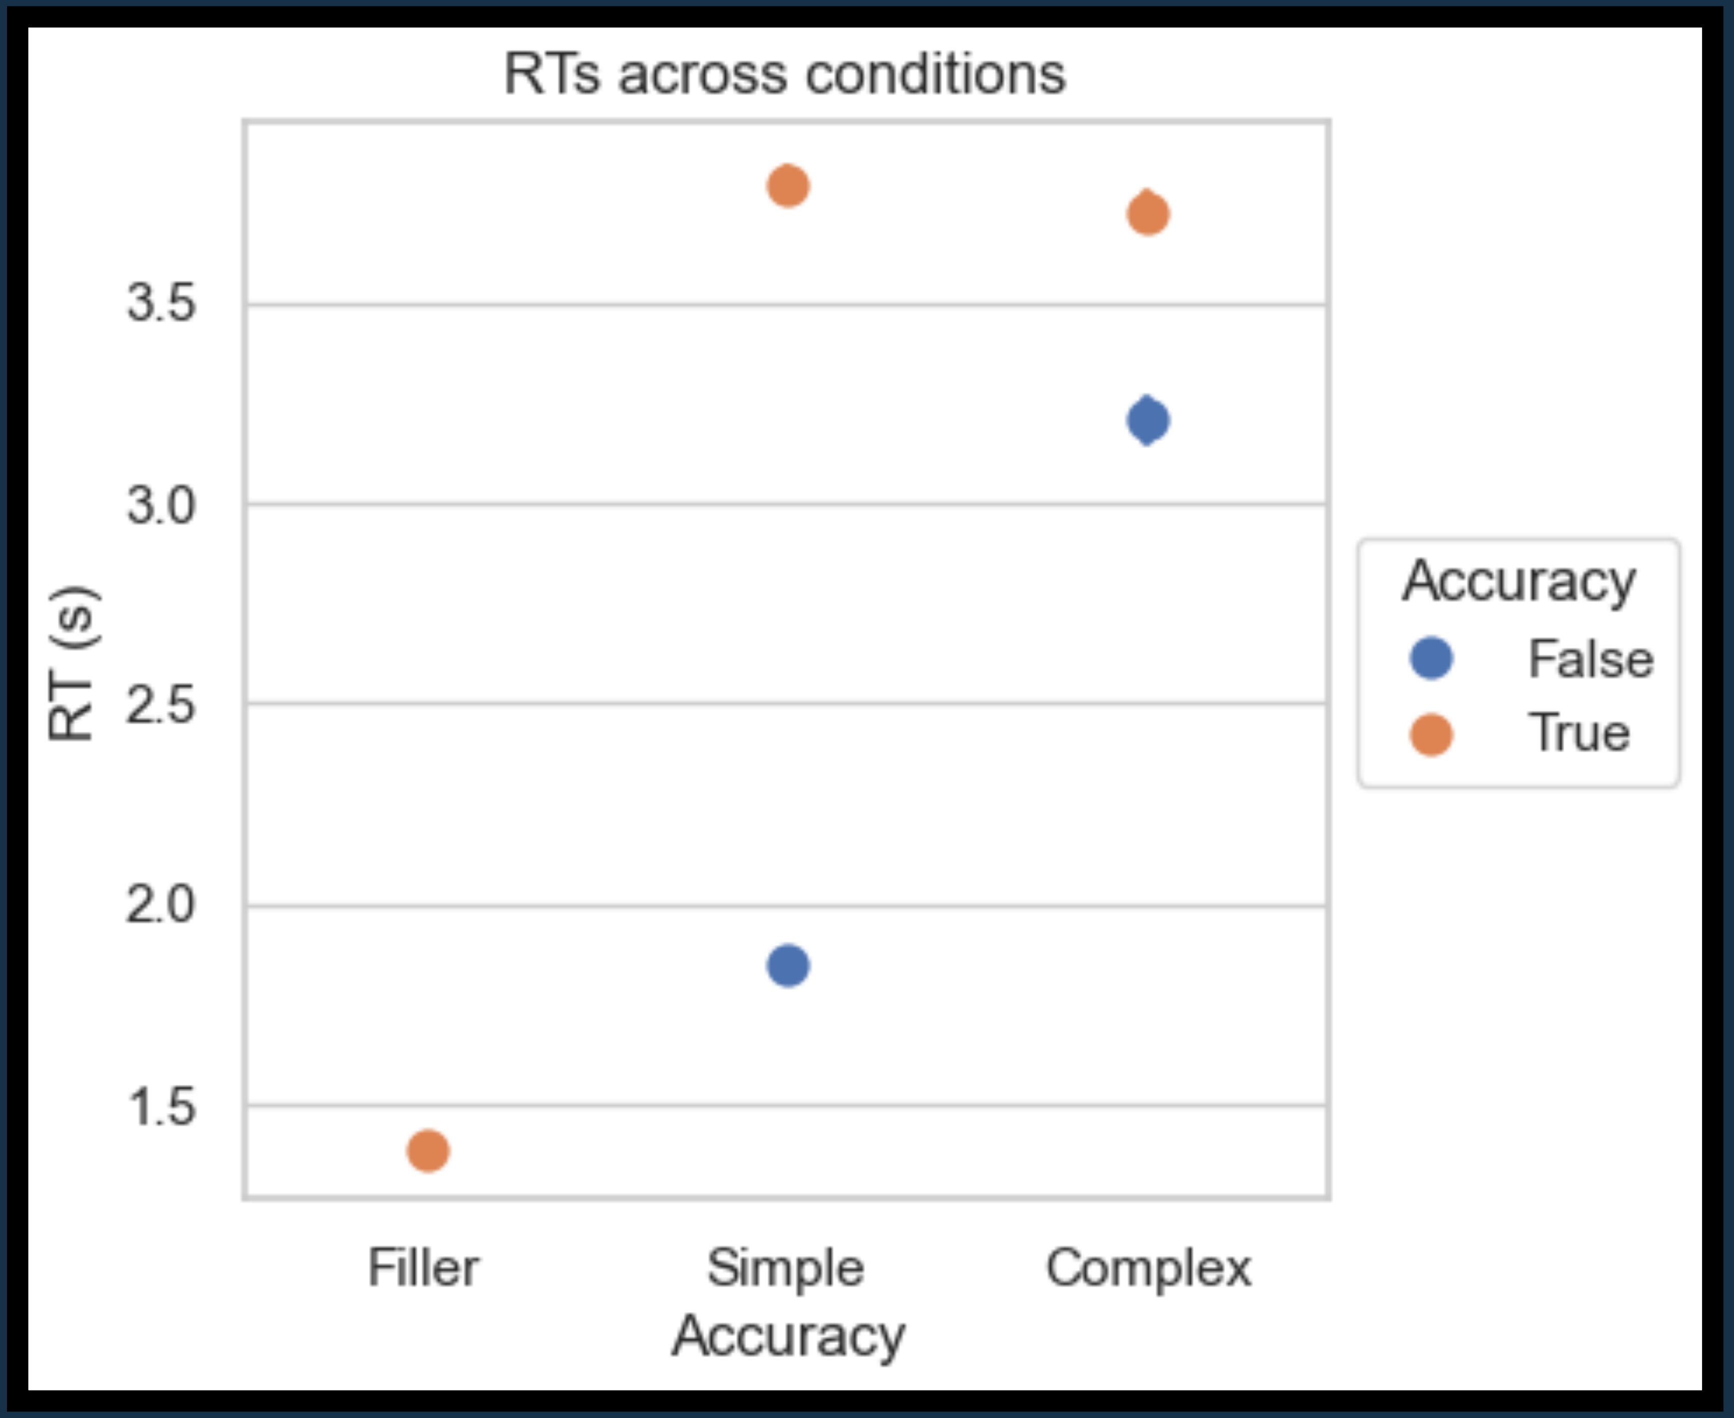

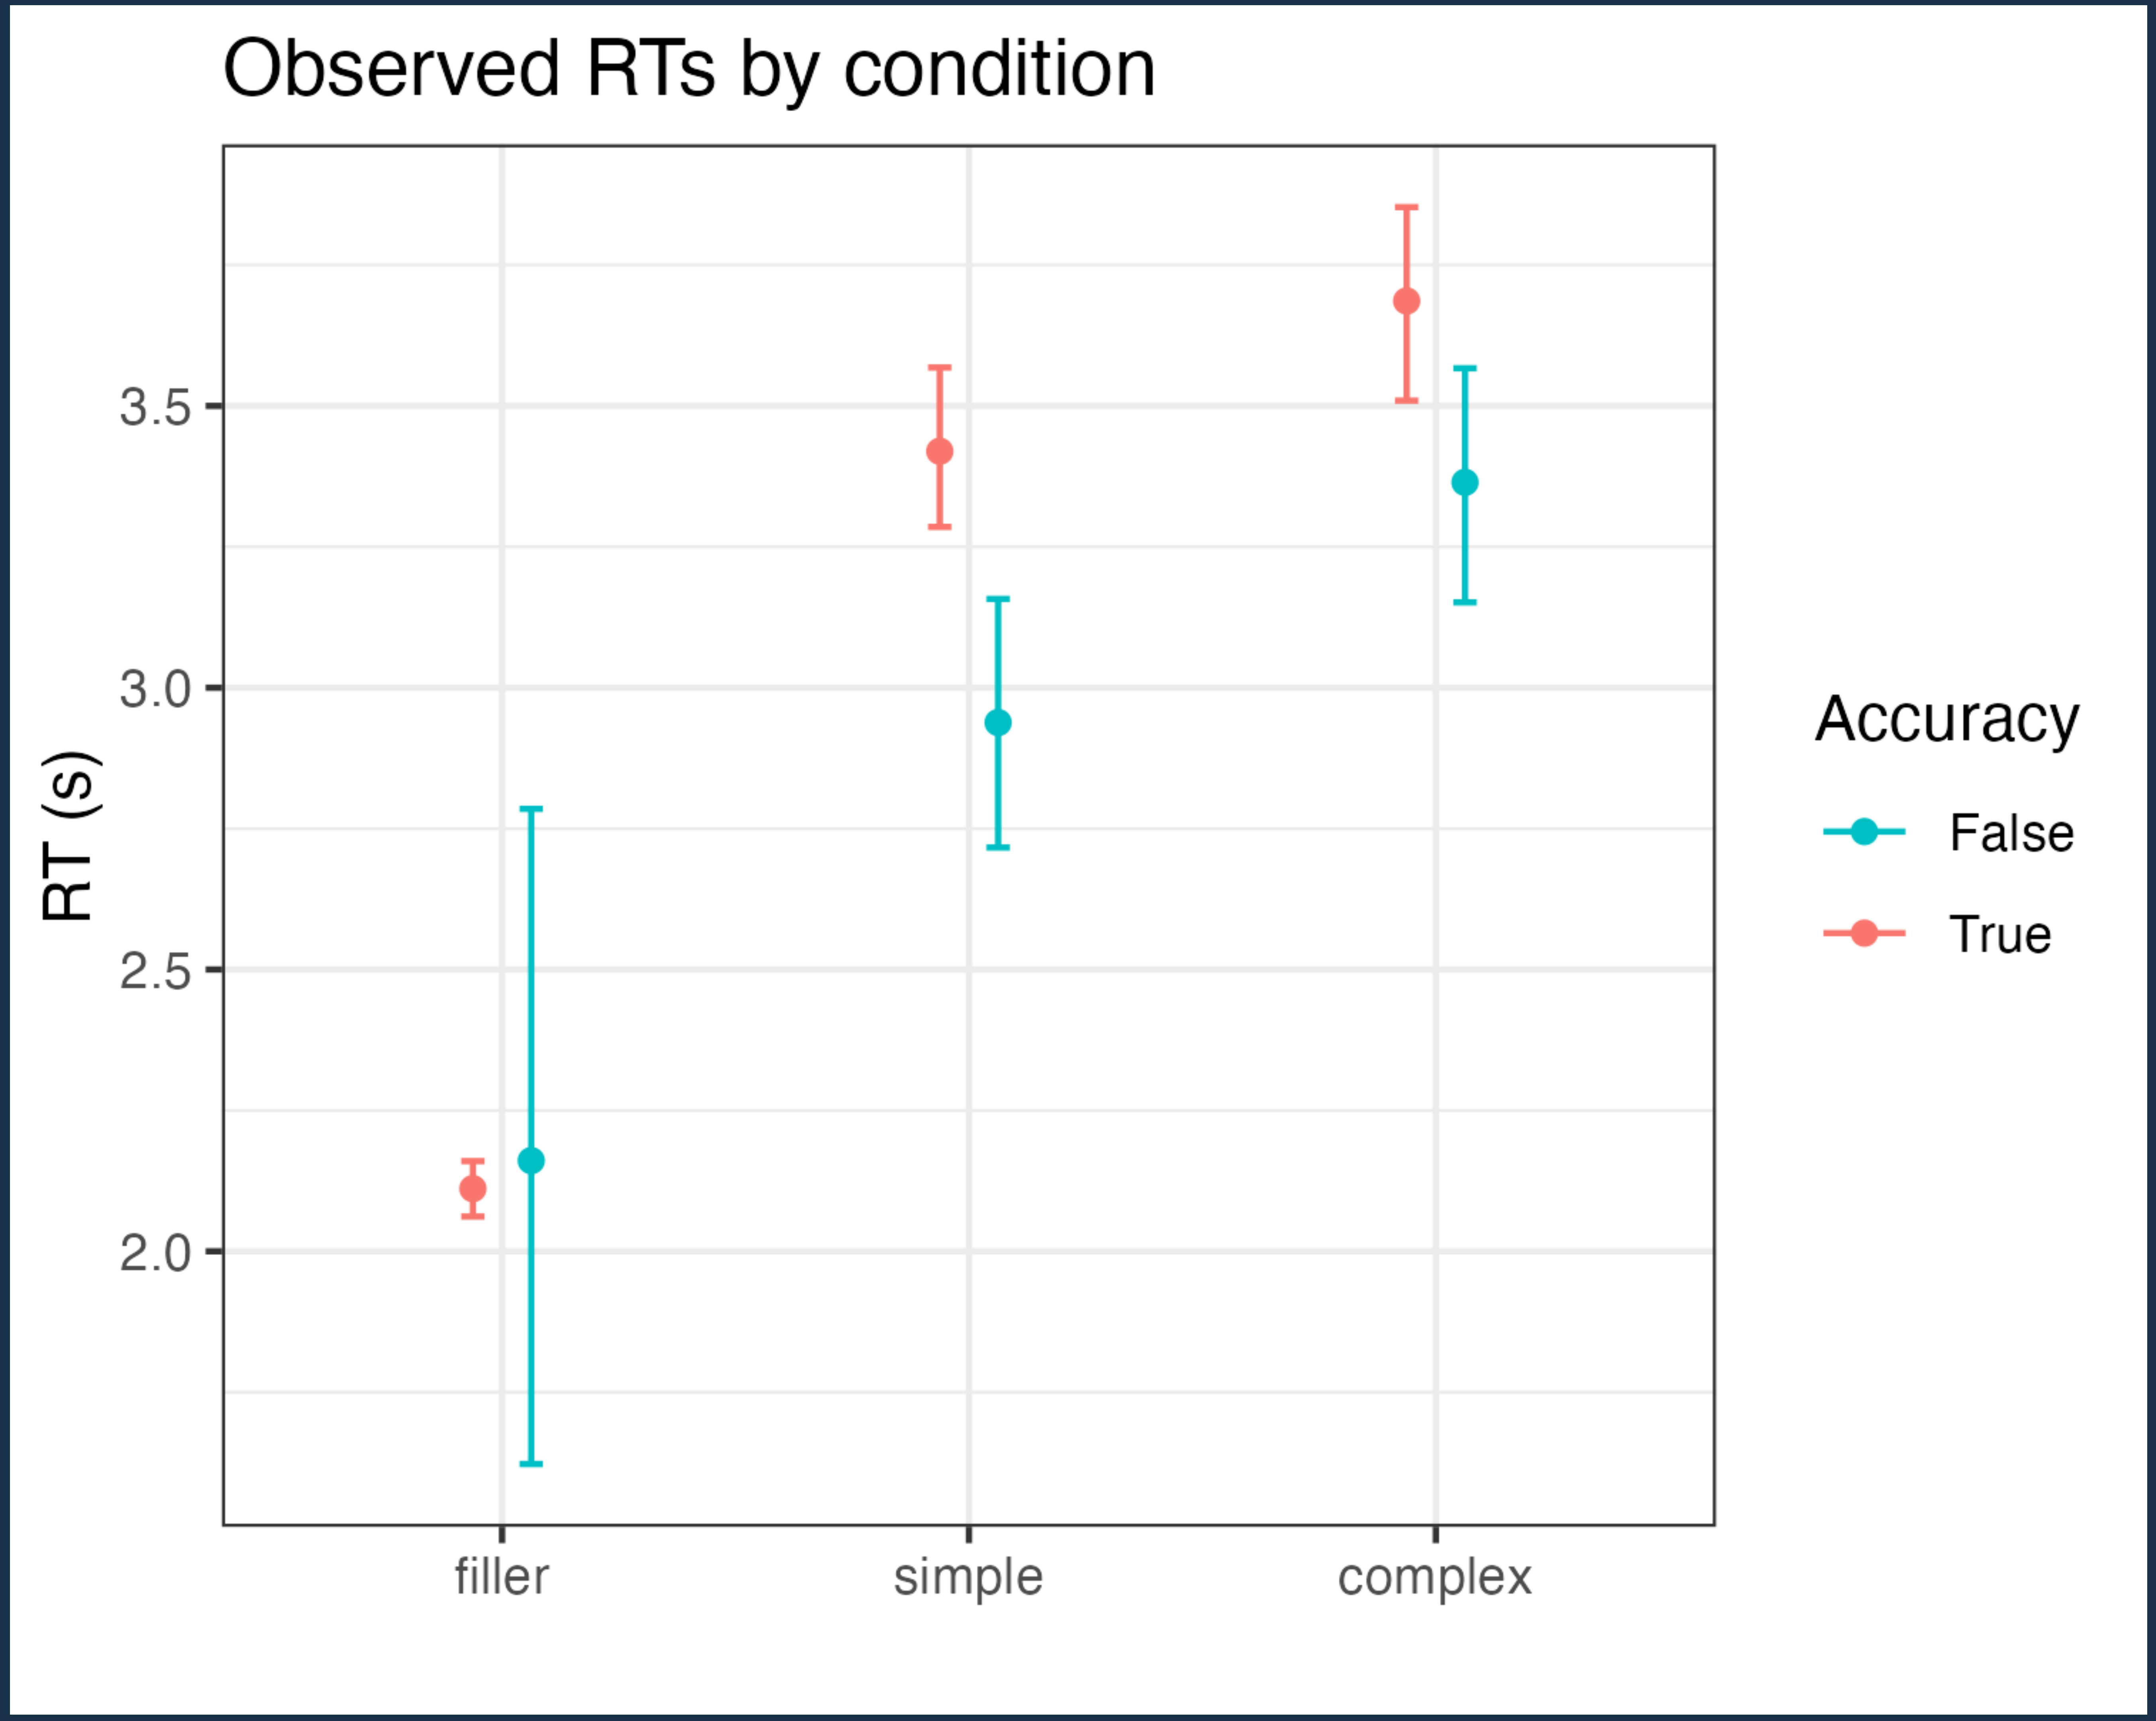

Probability fallacies in 1st-order reasoning

(Mayn, Duff, Bila & Demberg 2024)

- 1st-order pragmatic reasoning can solve "Simple" trials even with an **actual** literal (e.g. computer) speaker.

- Either 1st-order reasoning is never used, or participants apply it poorly.

(cf. Fox et al. 2004; Starns et al. 2019)

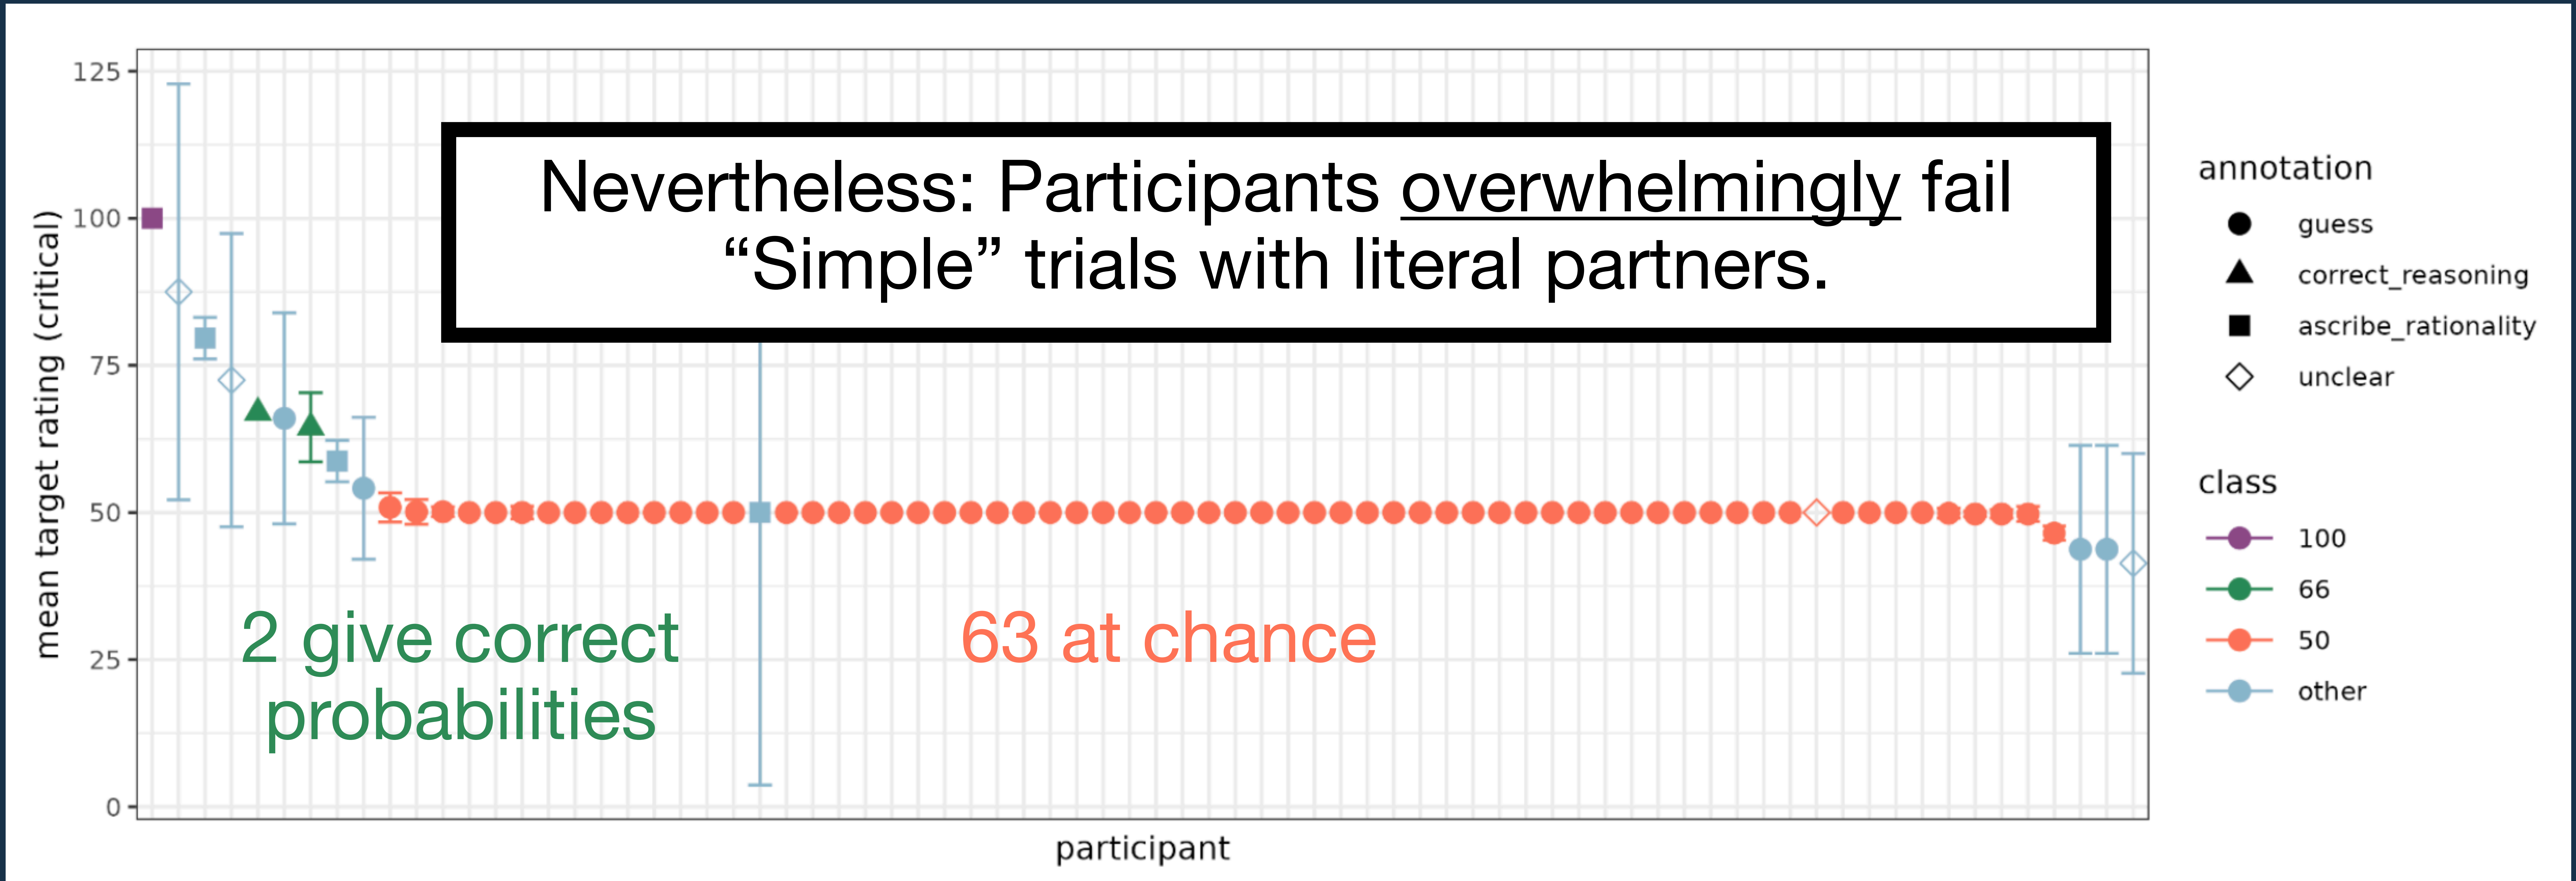

Atypicality inferences

(Ryzhova, Mayn & Demberg 2023)

`Mary went to a restaurant. She ate there!`

Mary must typically not eat when she goes to a restaurant.

- Participants with higher Raven's scores generated these inferences more often.
- Perhaps again, faster disengagement is supporting successful identification of a plausible candidate inference.