A Generative Model of Human Performance on an Optimal Stopping Problem

Michael Lee
Department of Psychology
University of Adelaide

with thanks to Peter Grünwald, Dan Navarro, Josh Tenenbaum, Chloë Mount, Jay Myung, Doug Vickers, Matt Welsh, Tess Gregory, Robyn Whibley, James Campbell, Chrisi Lambos, Michael Webb, Gary Ewing & Michael Paradowski

79.69

[1/5]

34.40

[2/5]

82.55

[3/5]

95.77

[4/5]

24.26

[5/5]

Why Study This Problem?

More Controlled than Knowledge-Rich

More Realistic than Knowledge-Lean

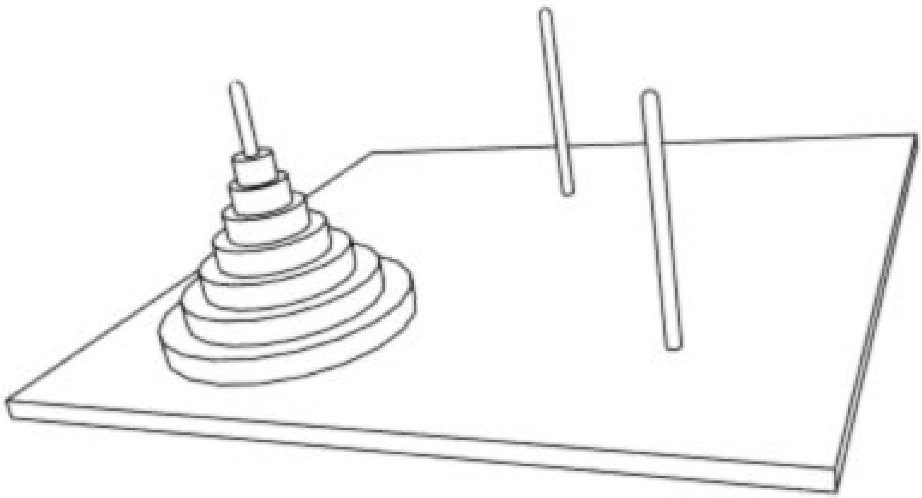

Complements Other Optimisation Problems

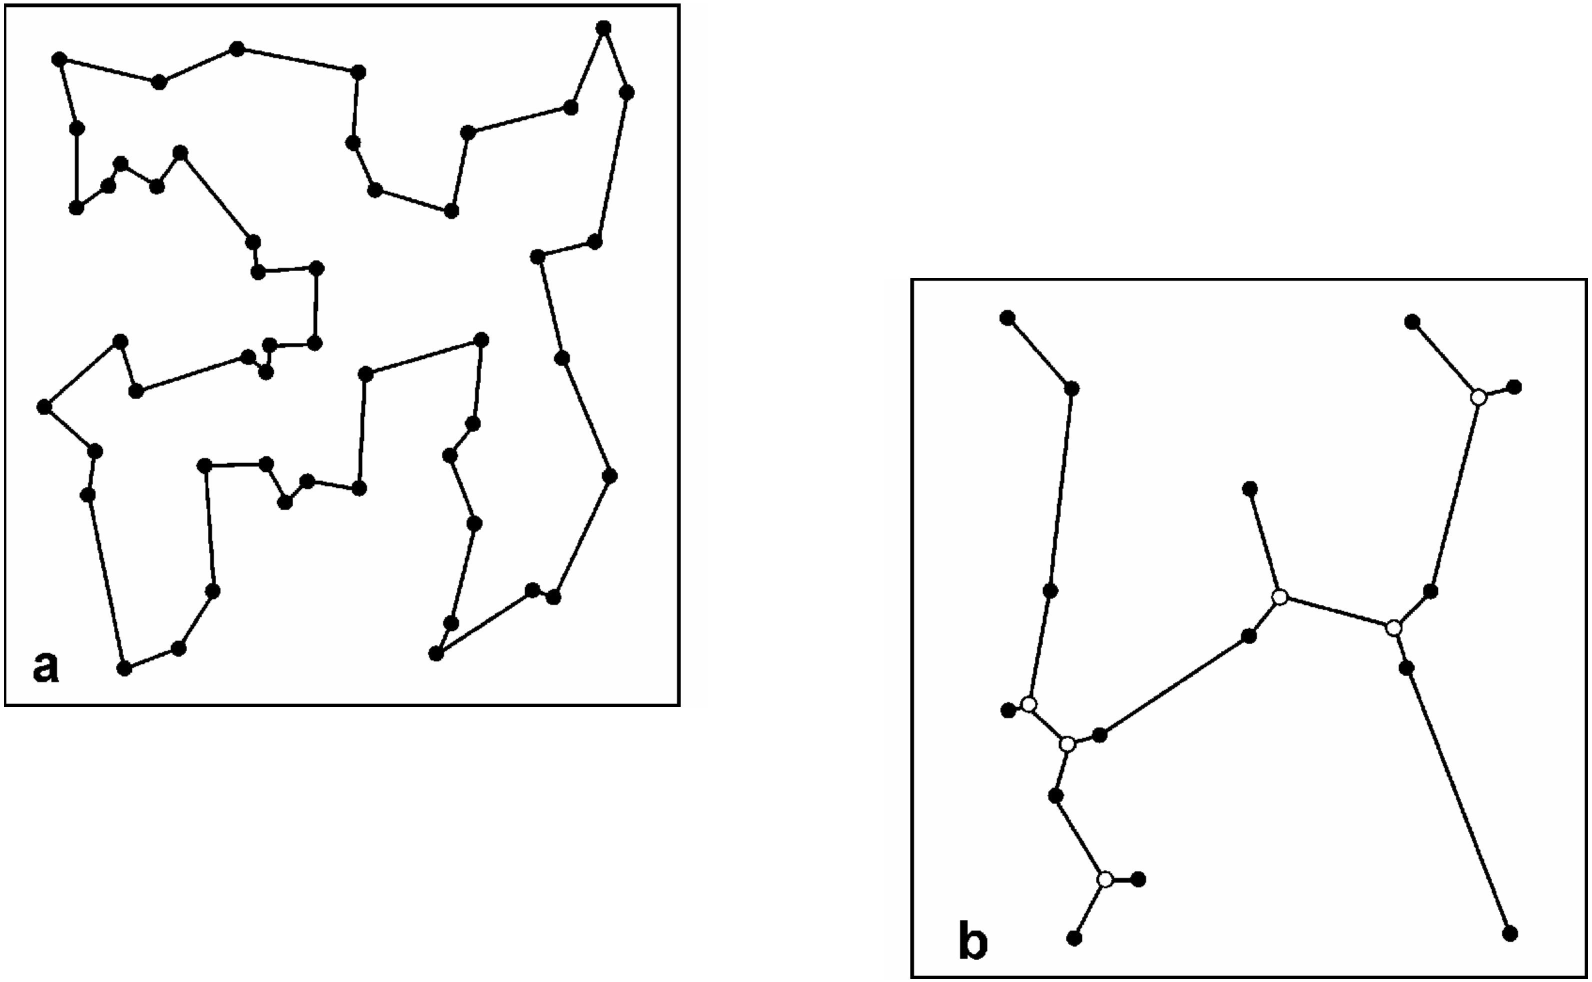

Optimal Decision Rule

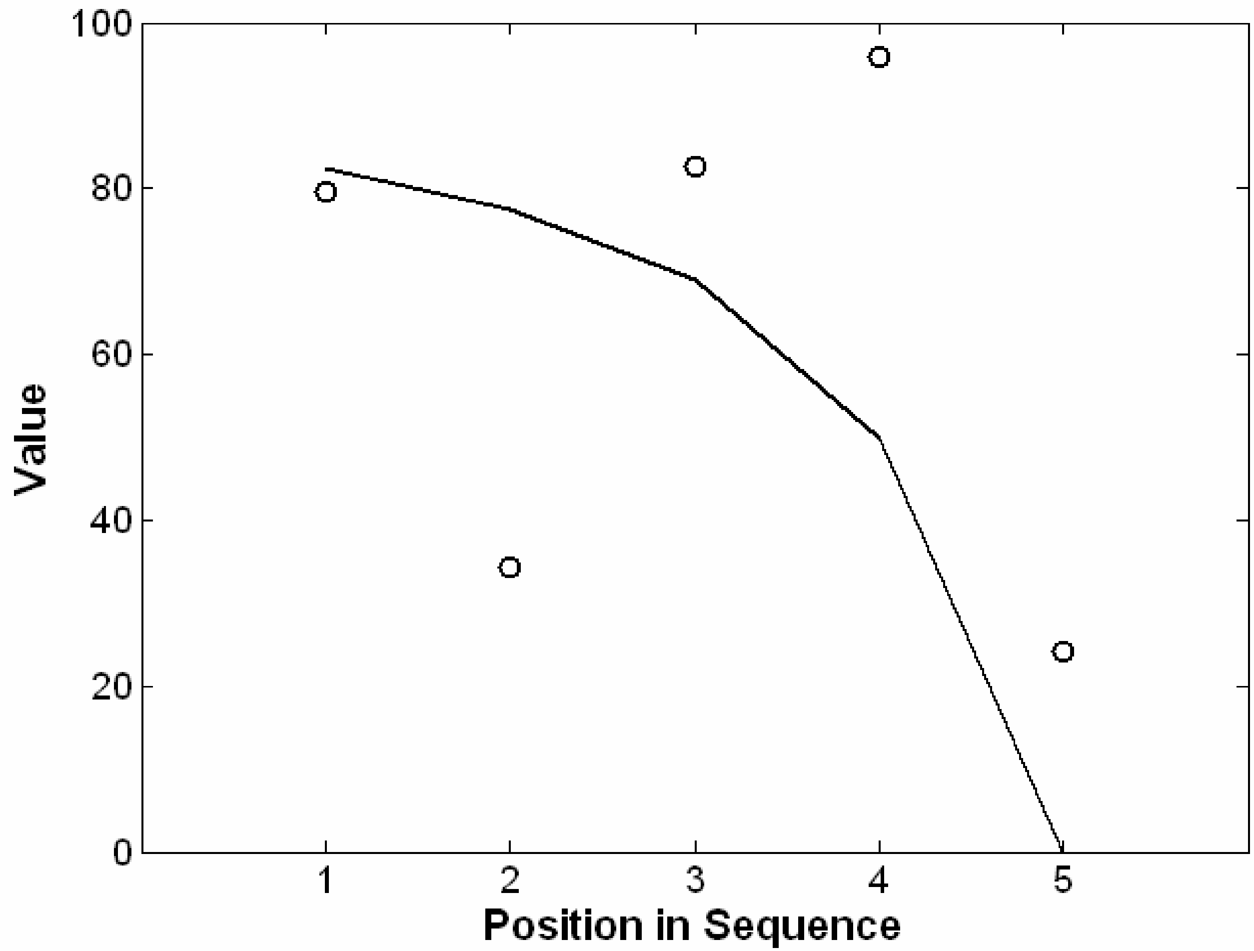

Two Optimal Decisions

follows decision rule
(good but unlucky)

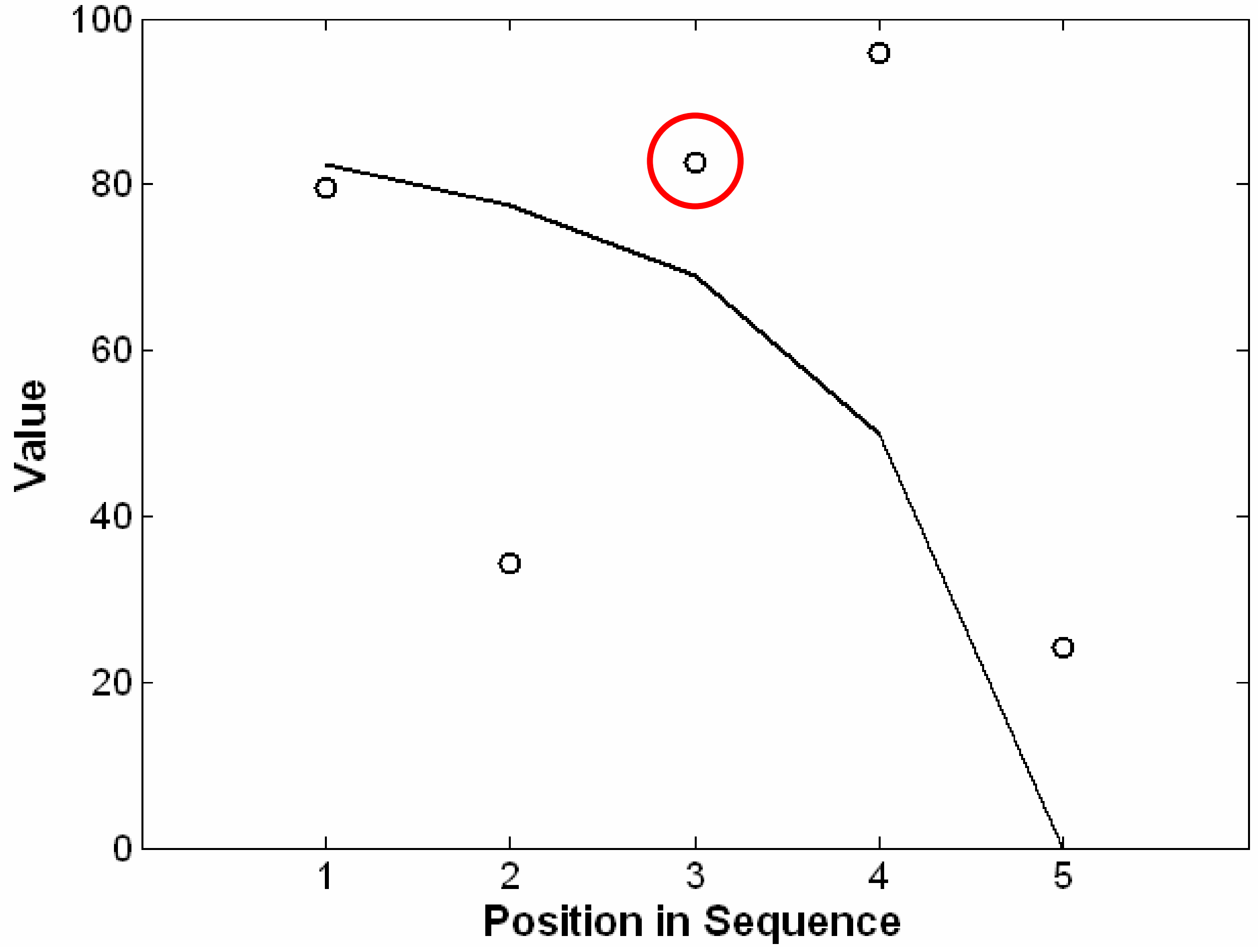

Two Optimal Decisions

violates rule, but maximum value
(bad but lucky)

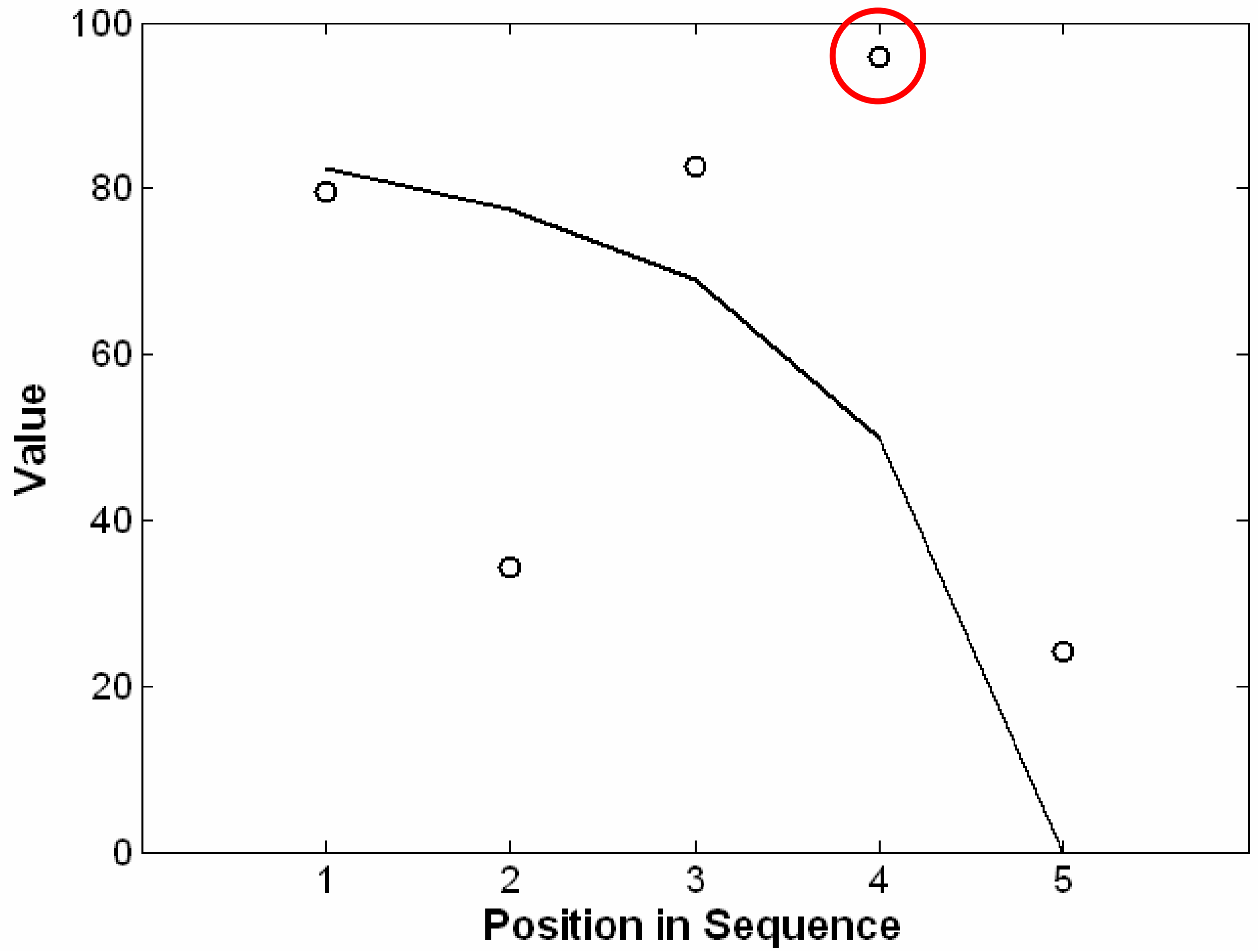

Empirical Data

Experimental Interface

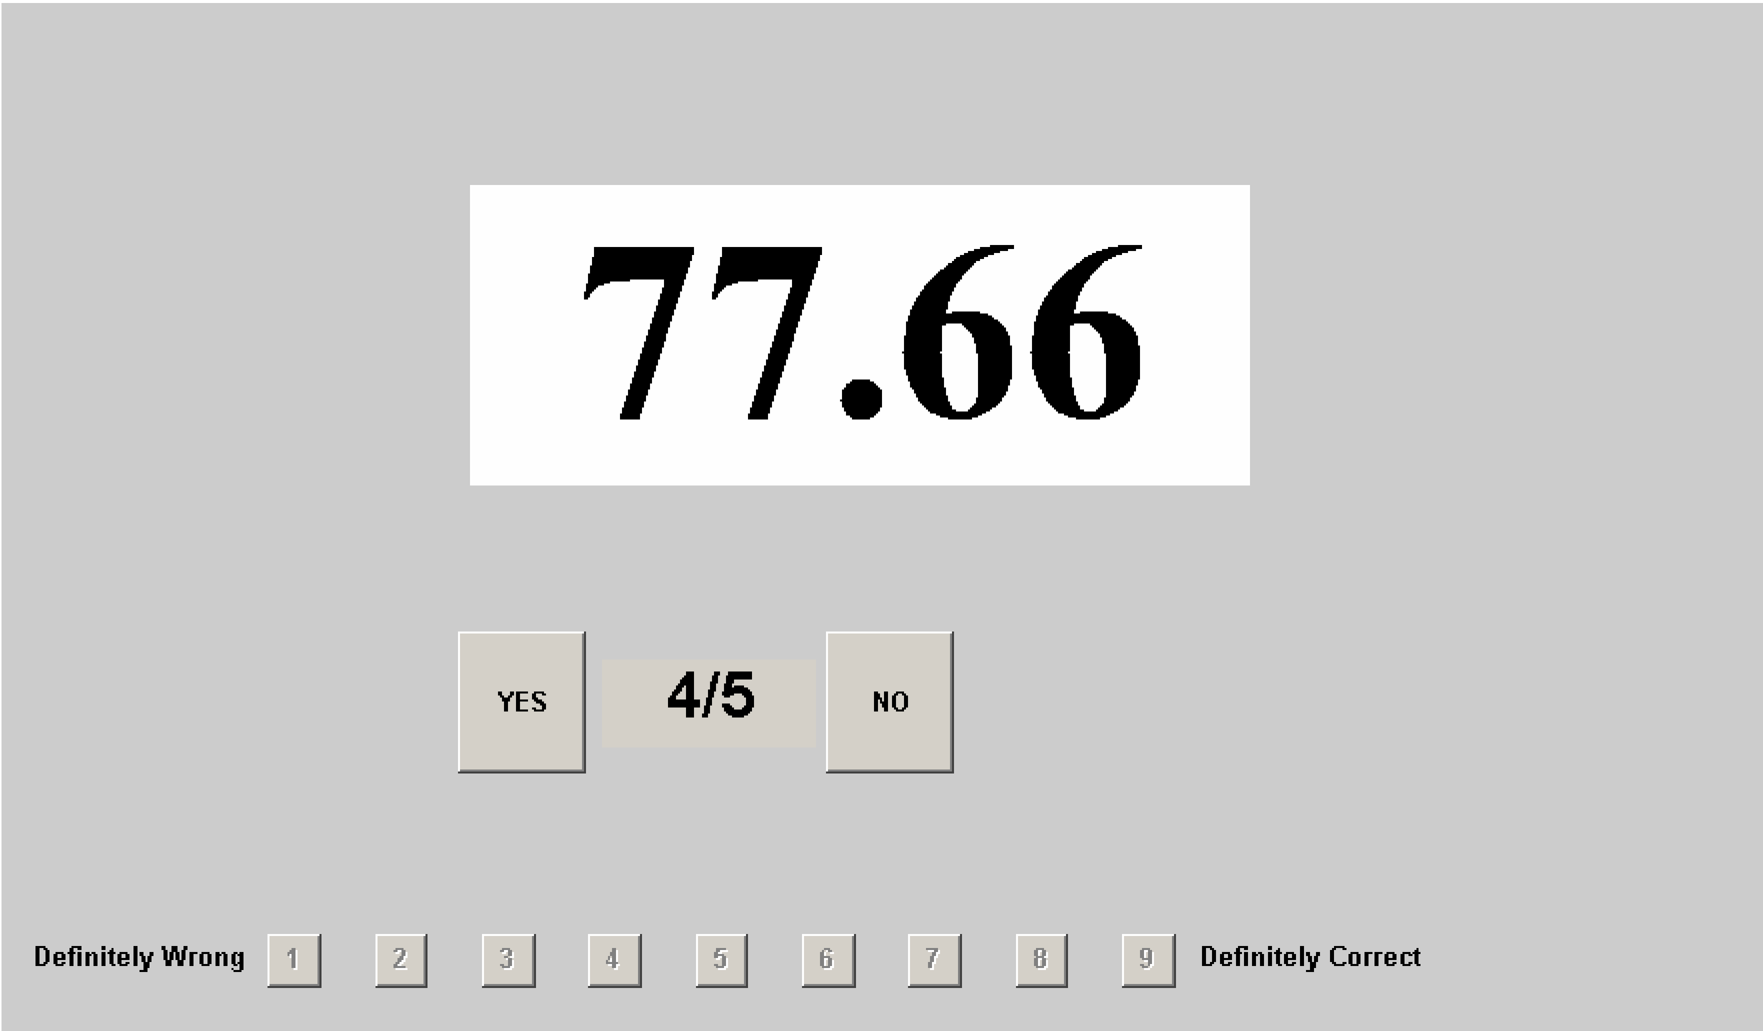

50 Subjects on 40 Five-Point Problems

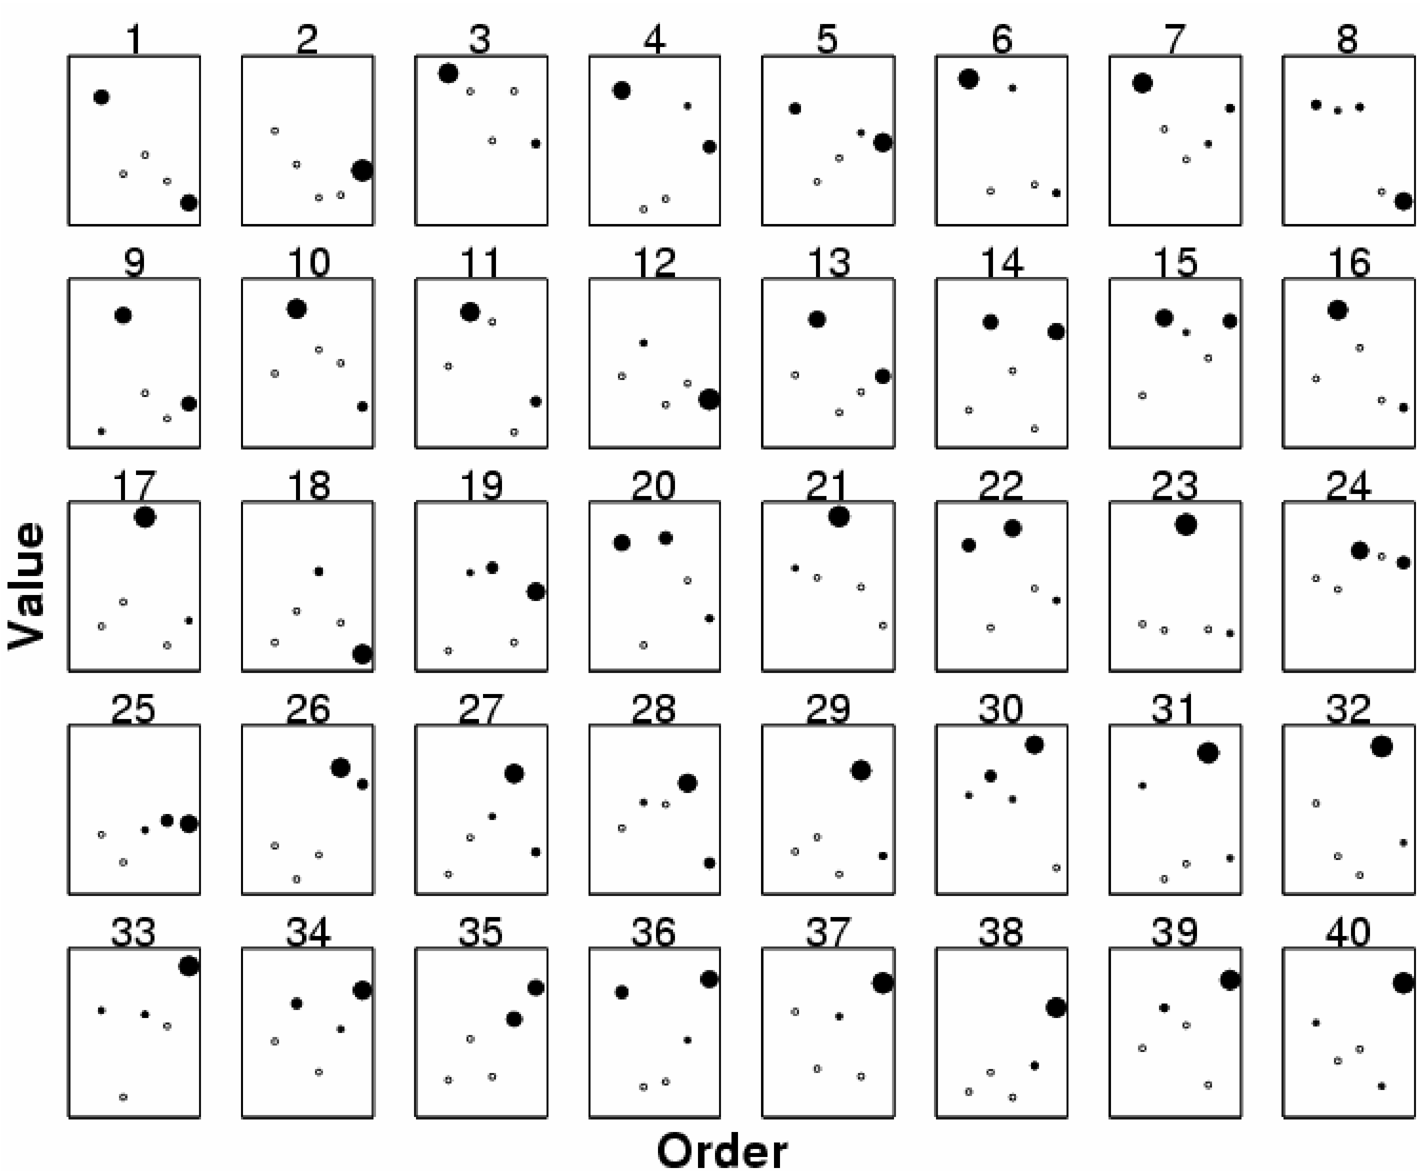

98 Subjects on 40 More Five-Point Problems

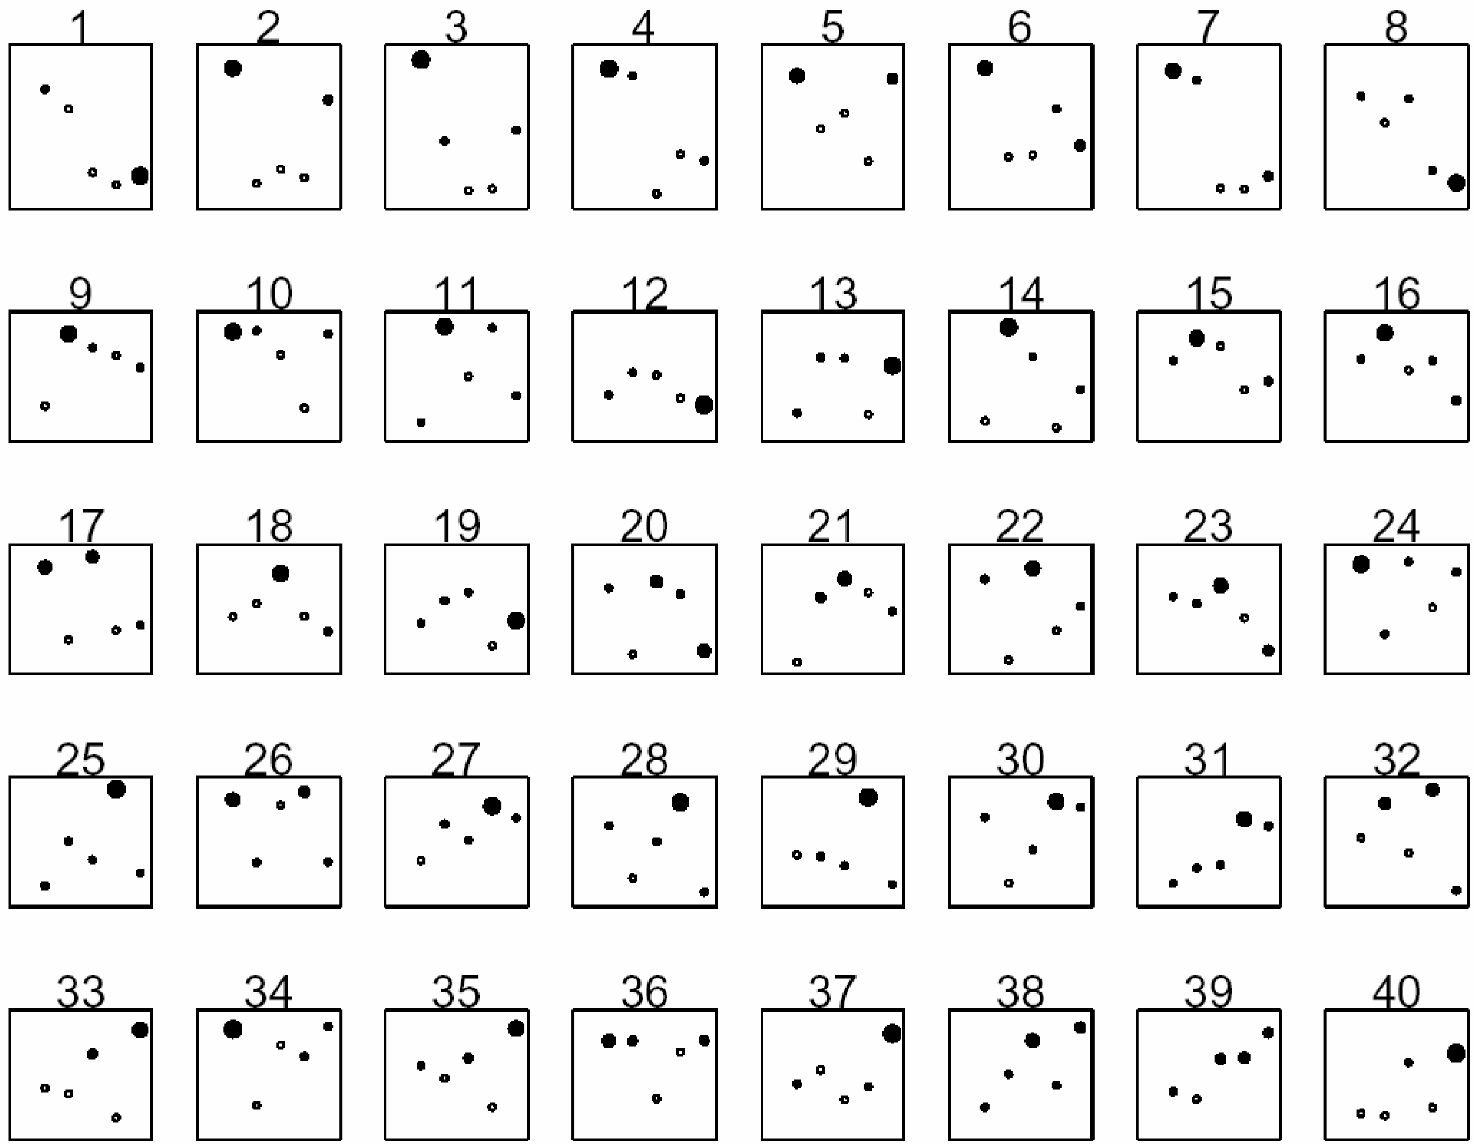

Individual Diffs in Accuracy & Confidence

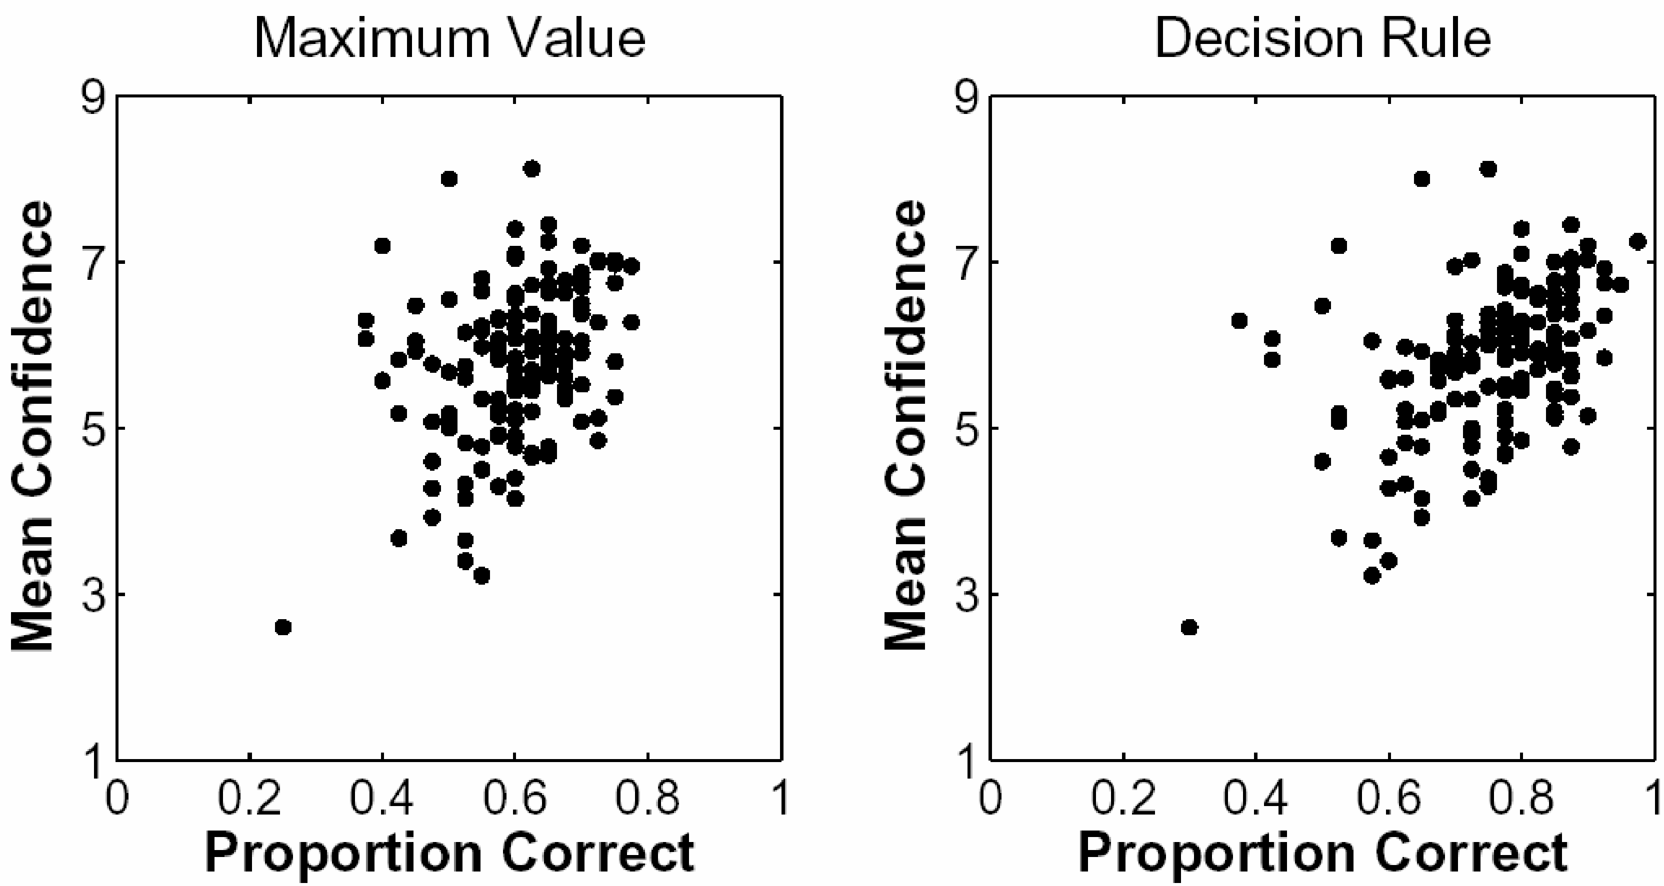

No Evidence of Learning

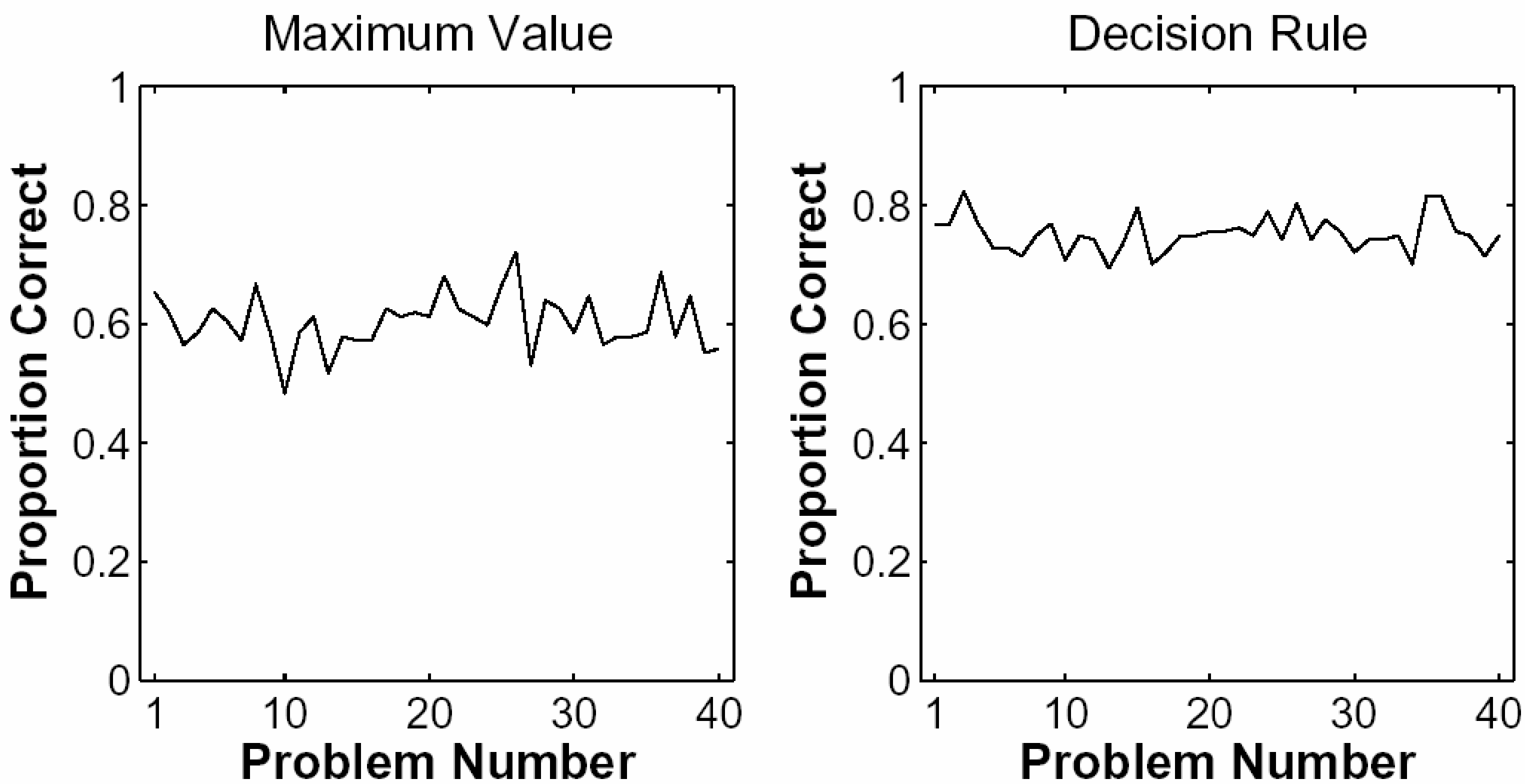

Other Values Don’t Seem to Matter

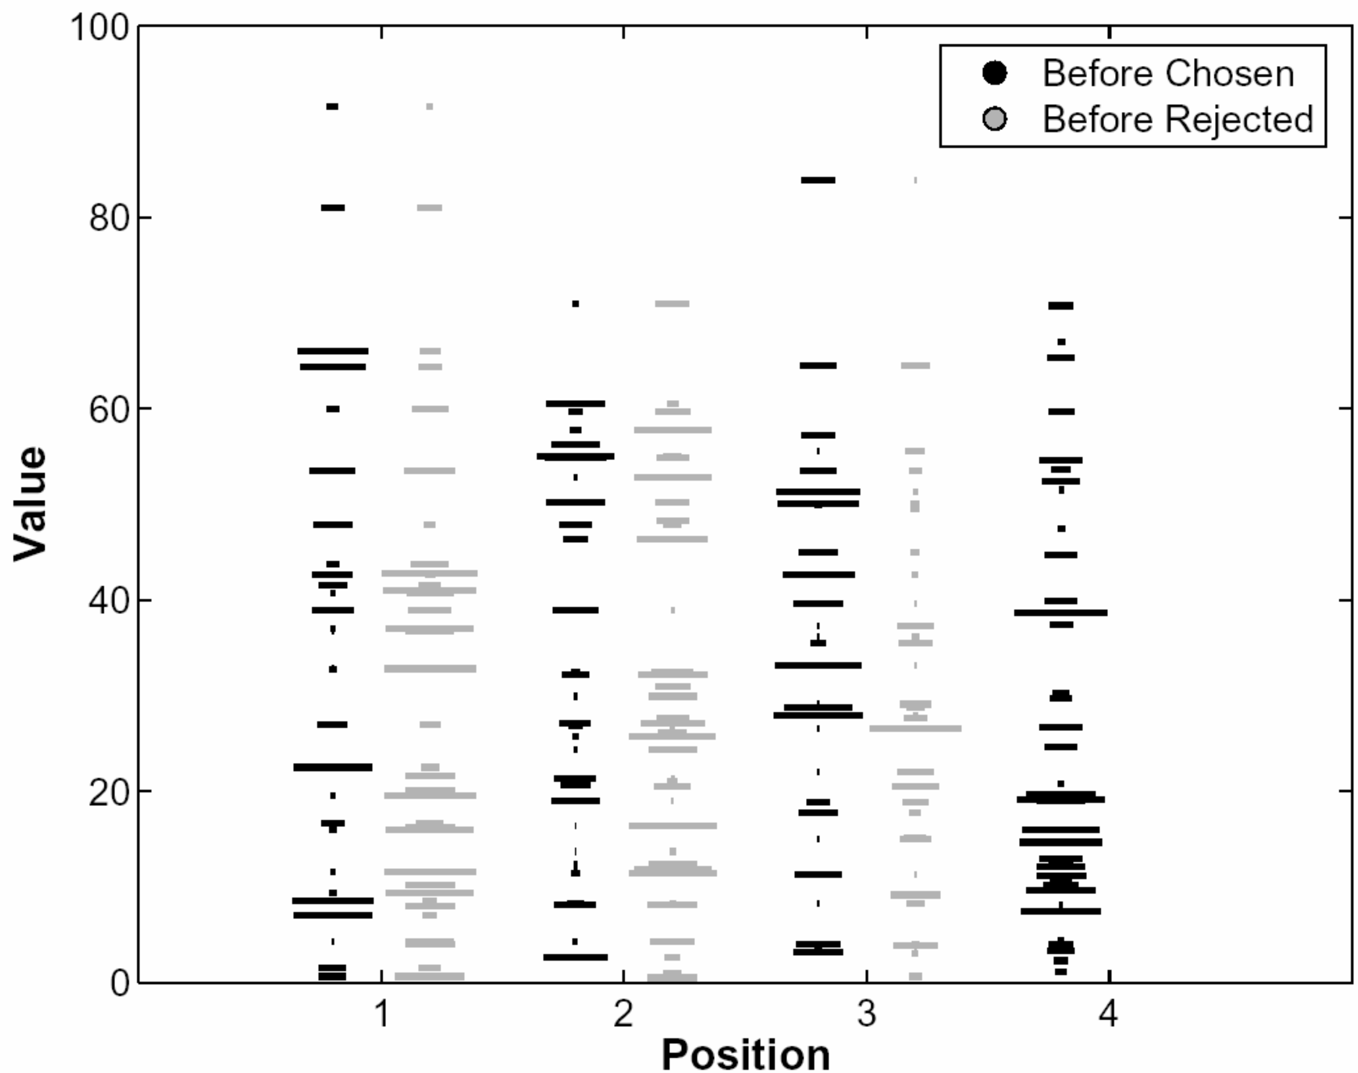

Value in Position Matters

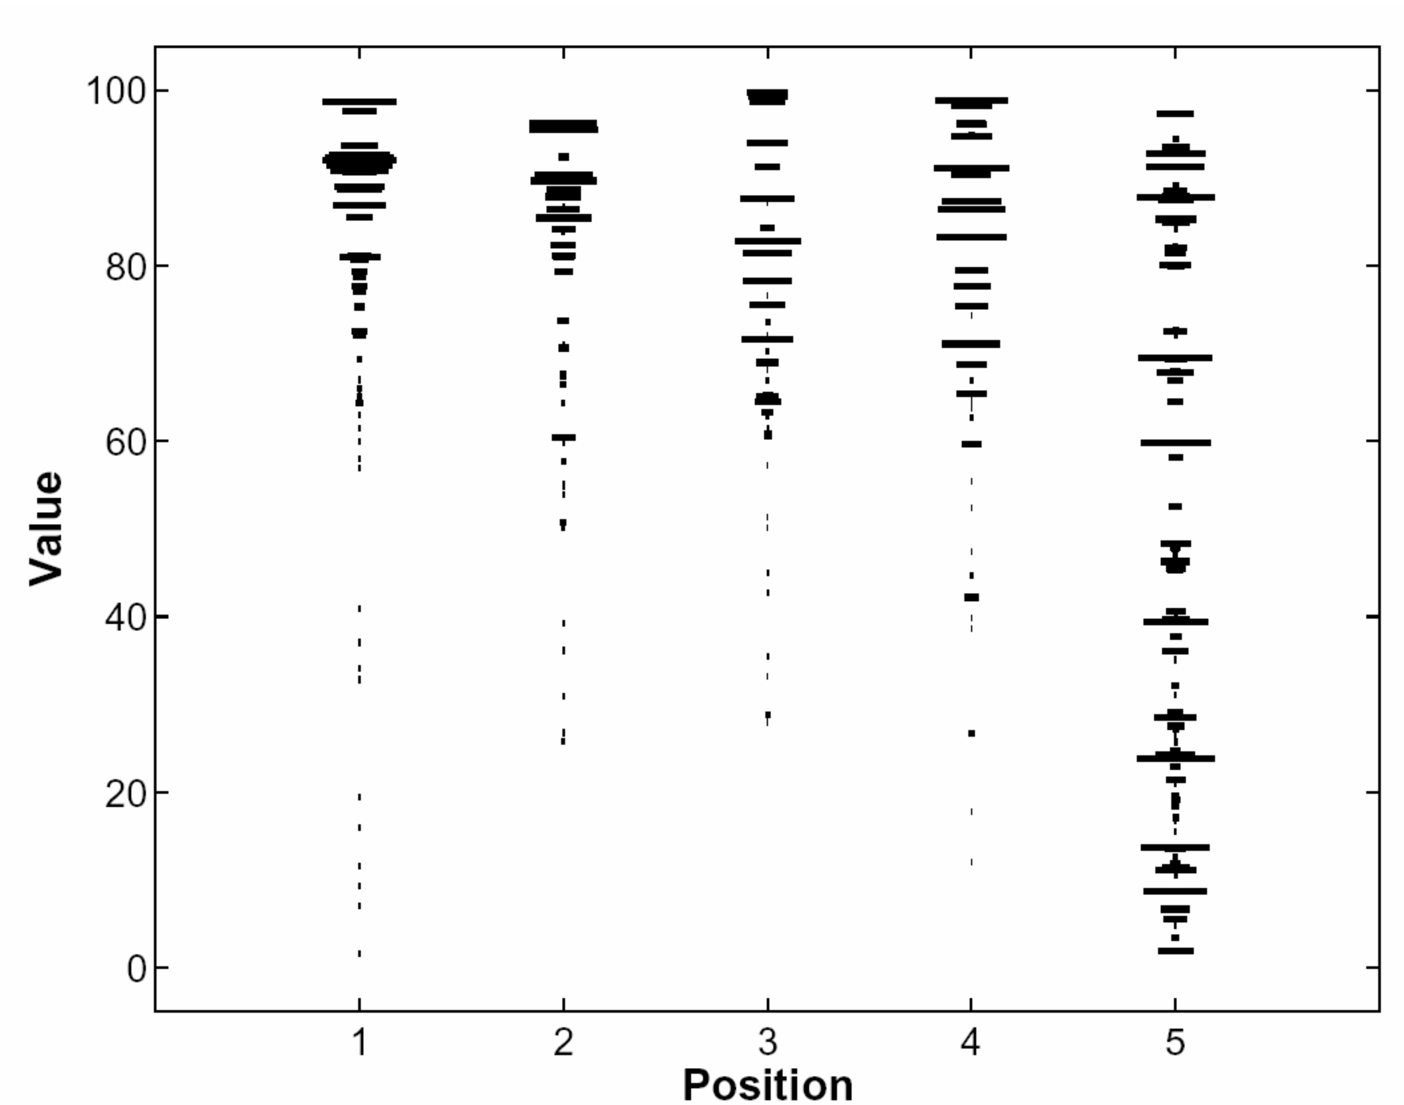

Being Currently Maximal Really Matters

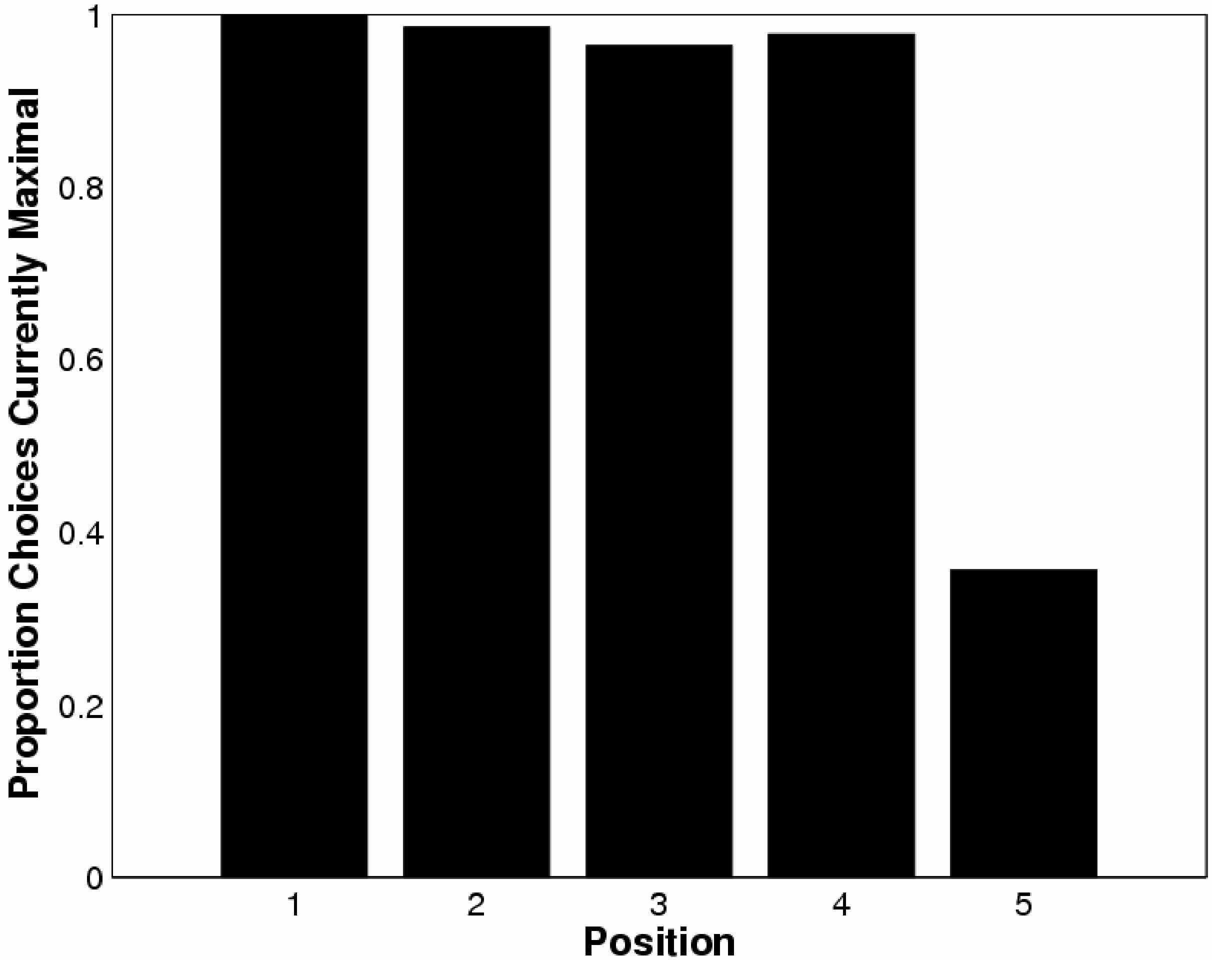

A Hierarchical Bayesian Generative Model

Why HBG Models?

- Hierarchical
  - represent knowledge at different levels of abstraction
  - formalize the structure of relationships between levels
- Generative
  - provide an account of how cognitive models are instantiated and bounded
  - quarantine core issues of modeling from secondary questions of inference
- Bayesian
  - provides a (the?) complete and coherent approach to statistical inference with models

Generative Framework

Generative Process

Model Family

Indexed Sequences

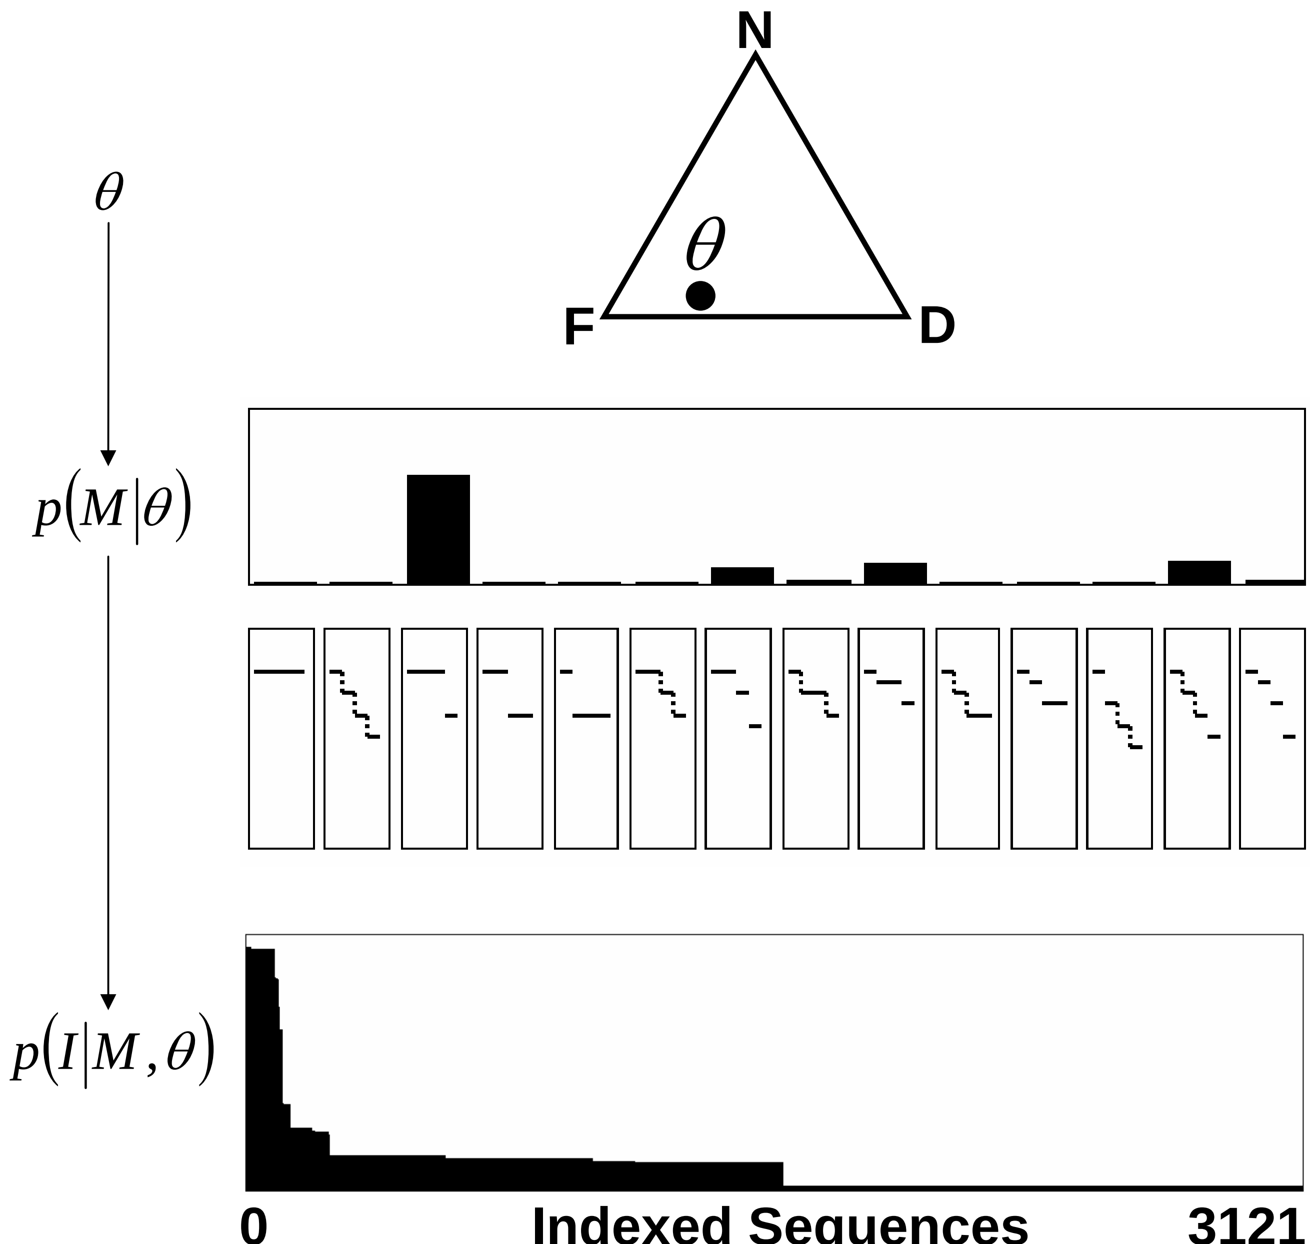

Making Inferences About Generative Process

Generative Process

Model Family

Indexed Sequences

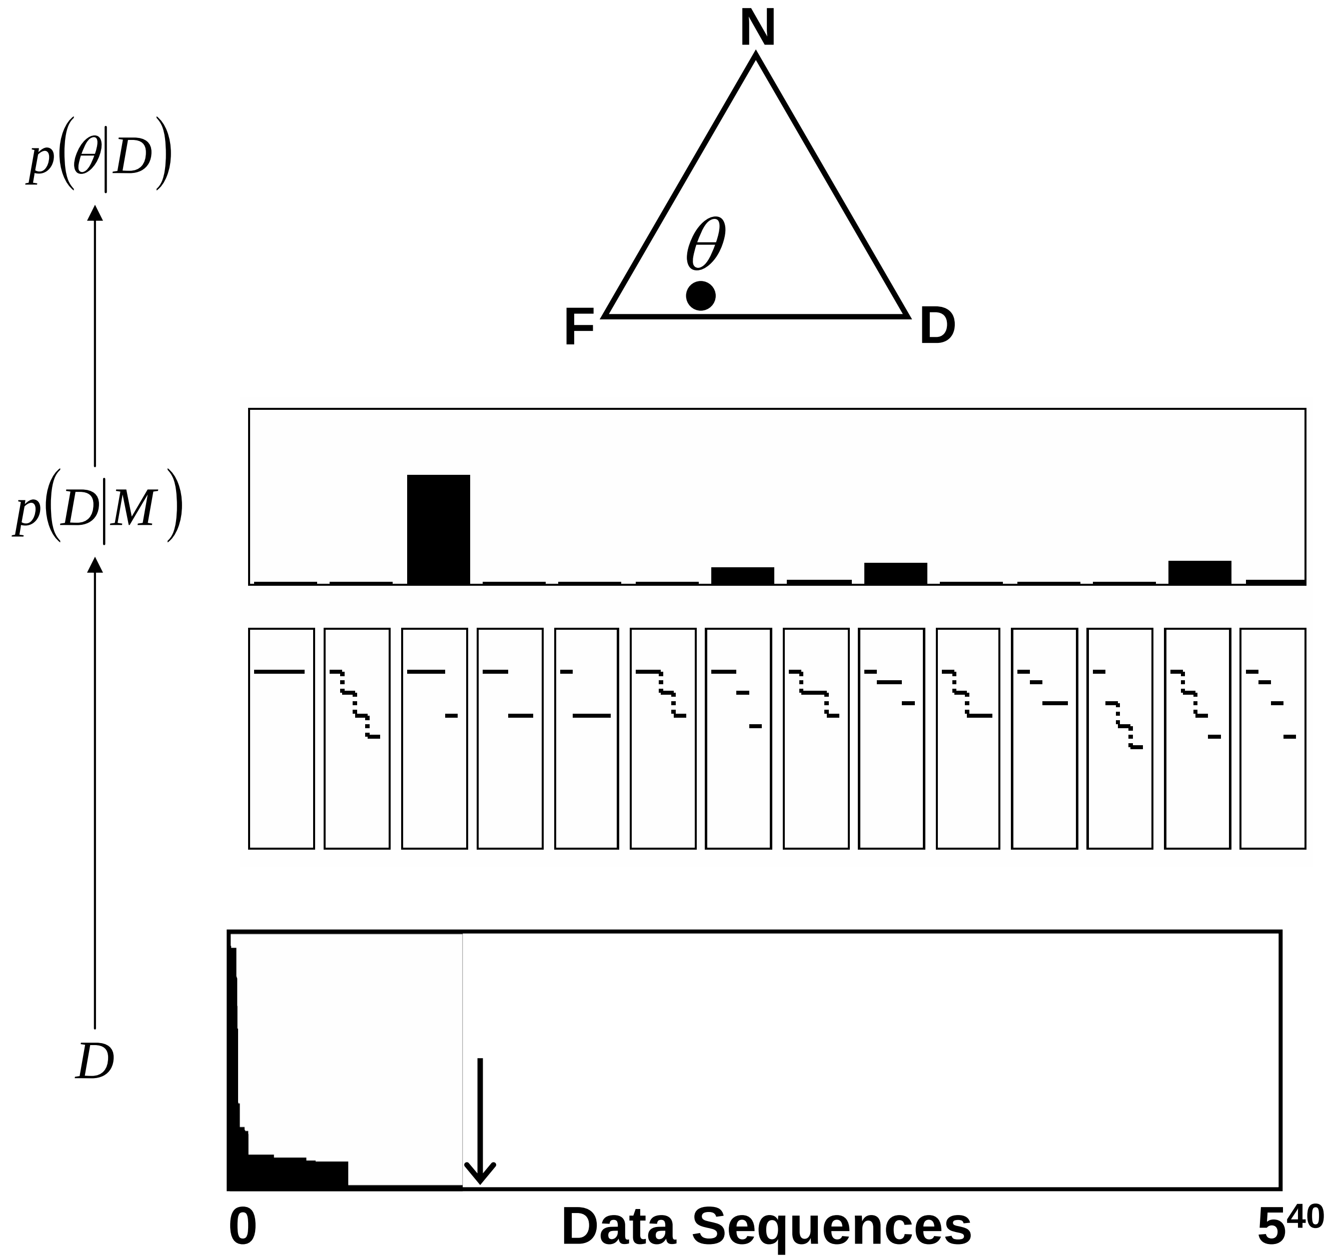

Results for One Subject

Inference from Data

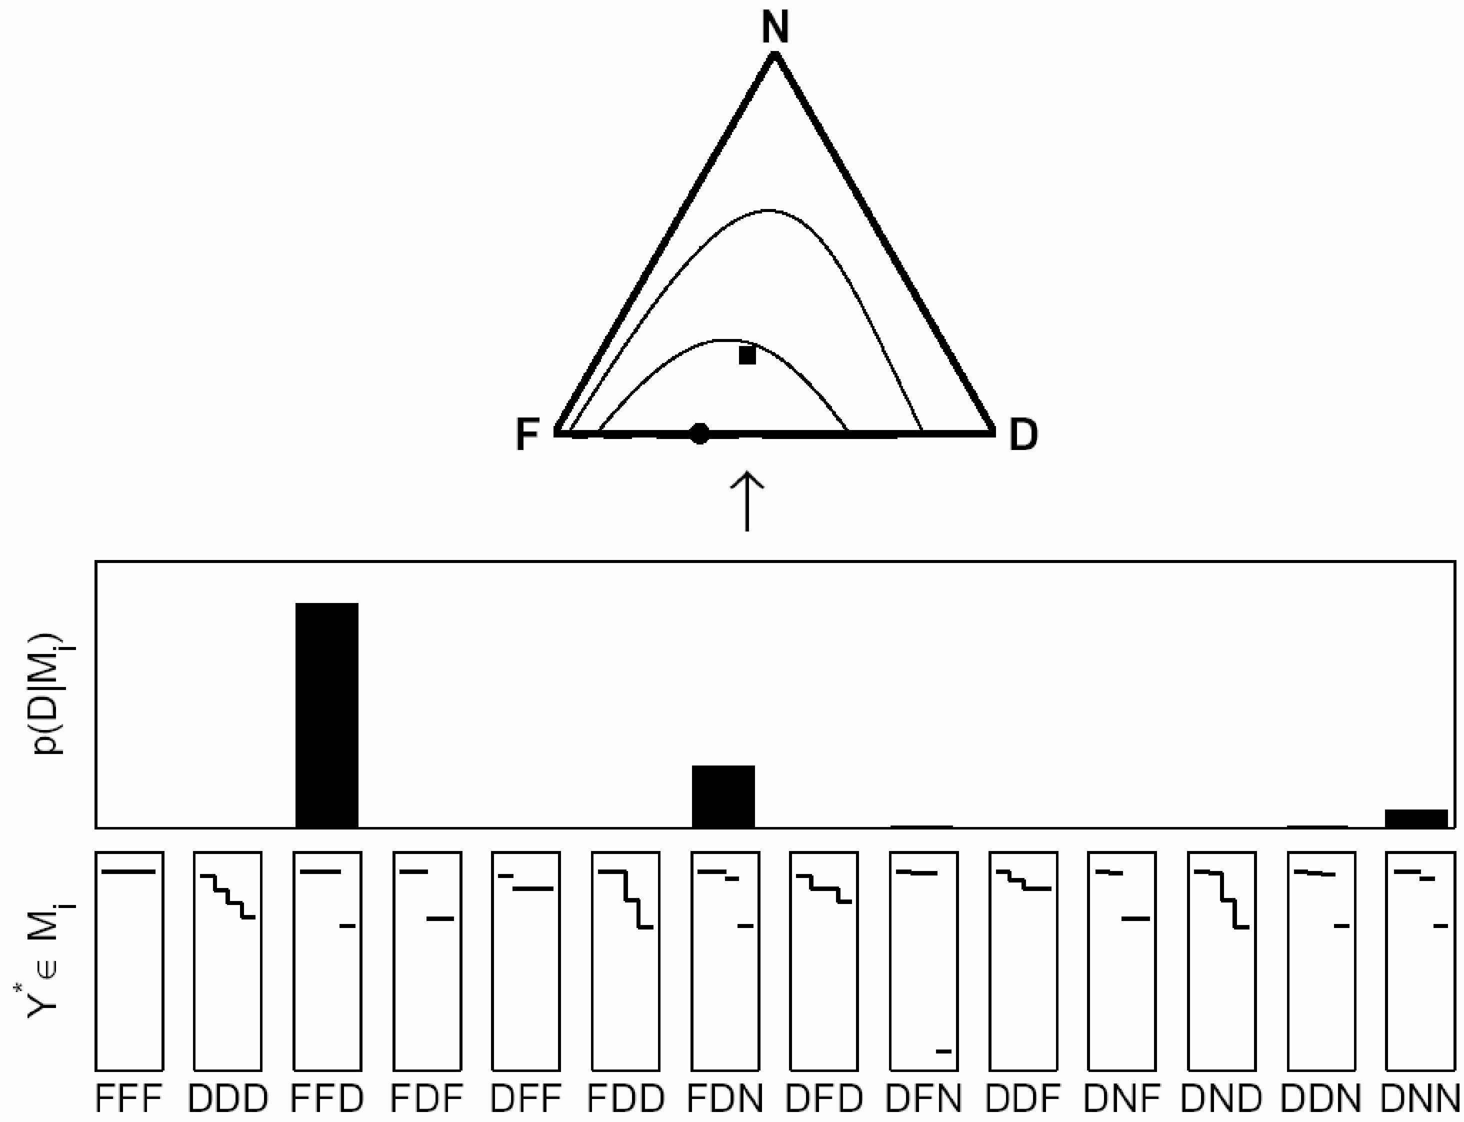

Building Predictive Models

$$p(M_i|D) = \frac{p(D|M_i)p(M_i)}{\sum_{j=1}^{14} p(D|M_j)p(M_j)}$$

Model Average

MAP **ML**

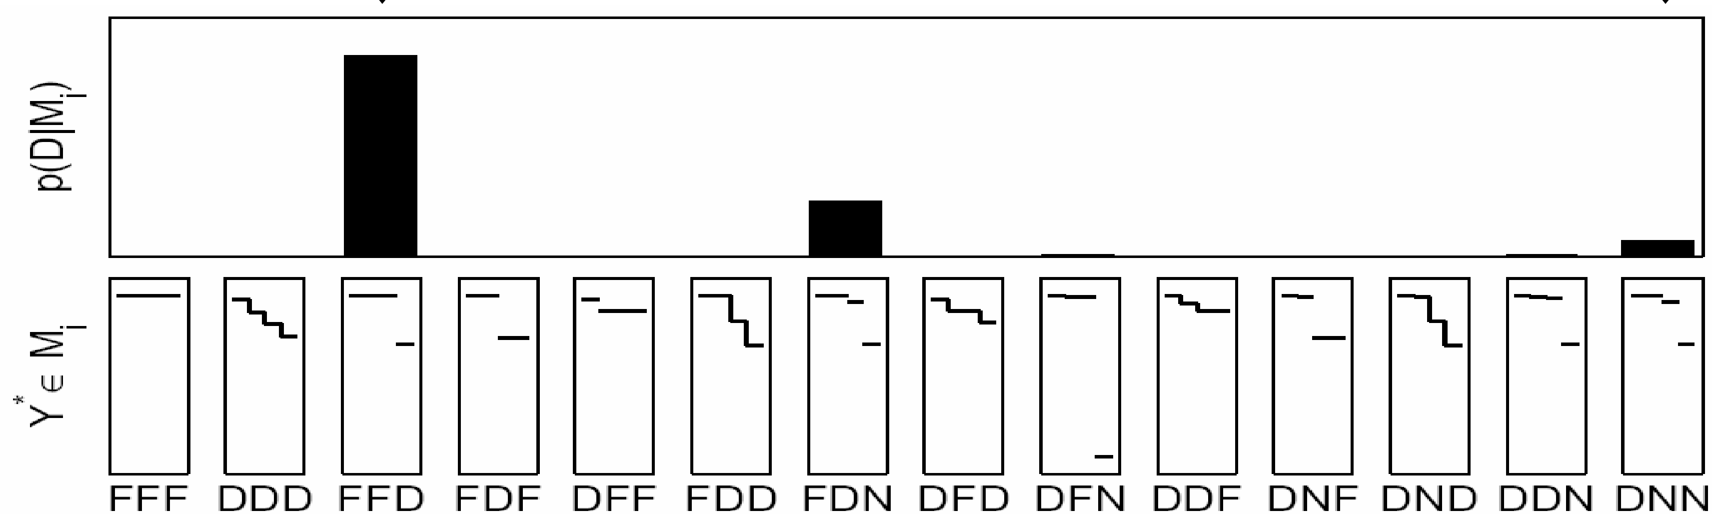

Posterior Predictive Models

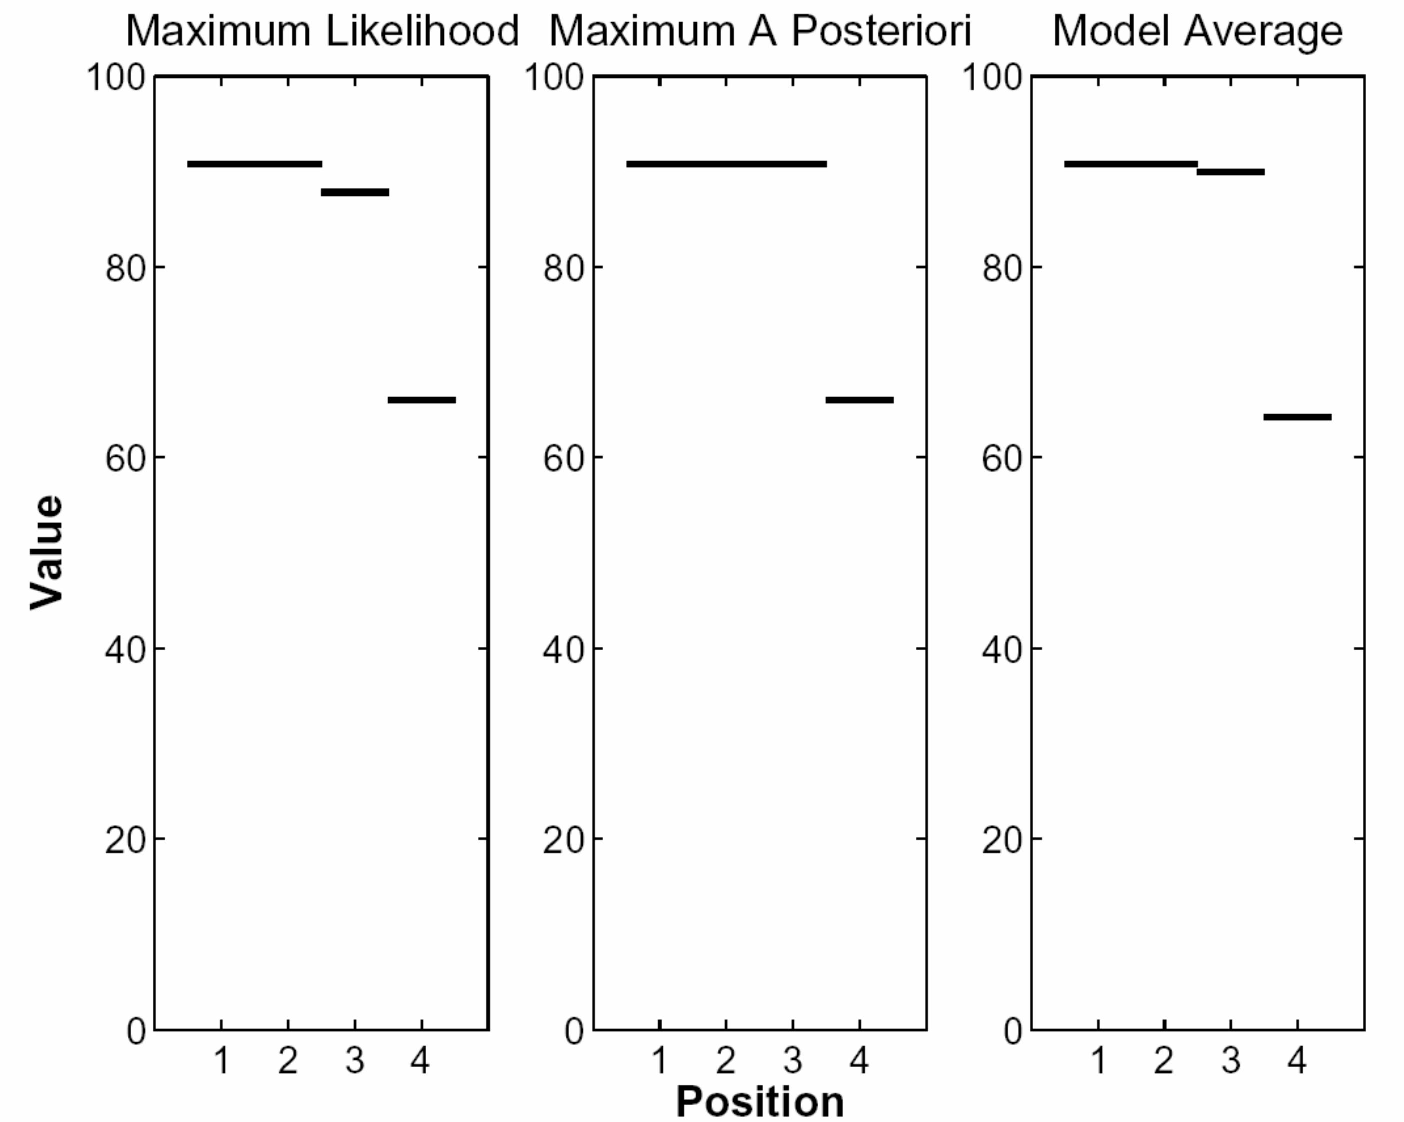

Overall Evaluation

Individual Differences at Generative Level

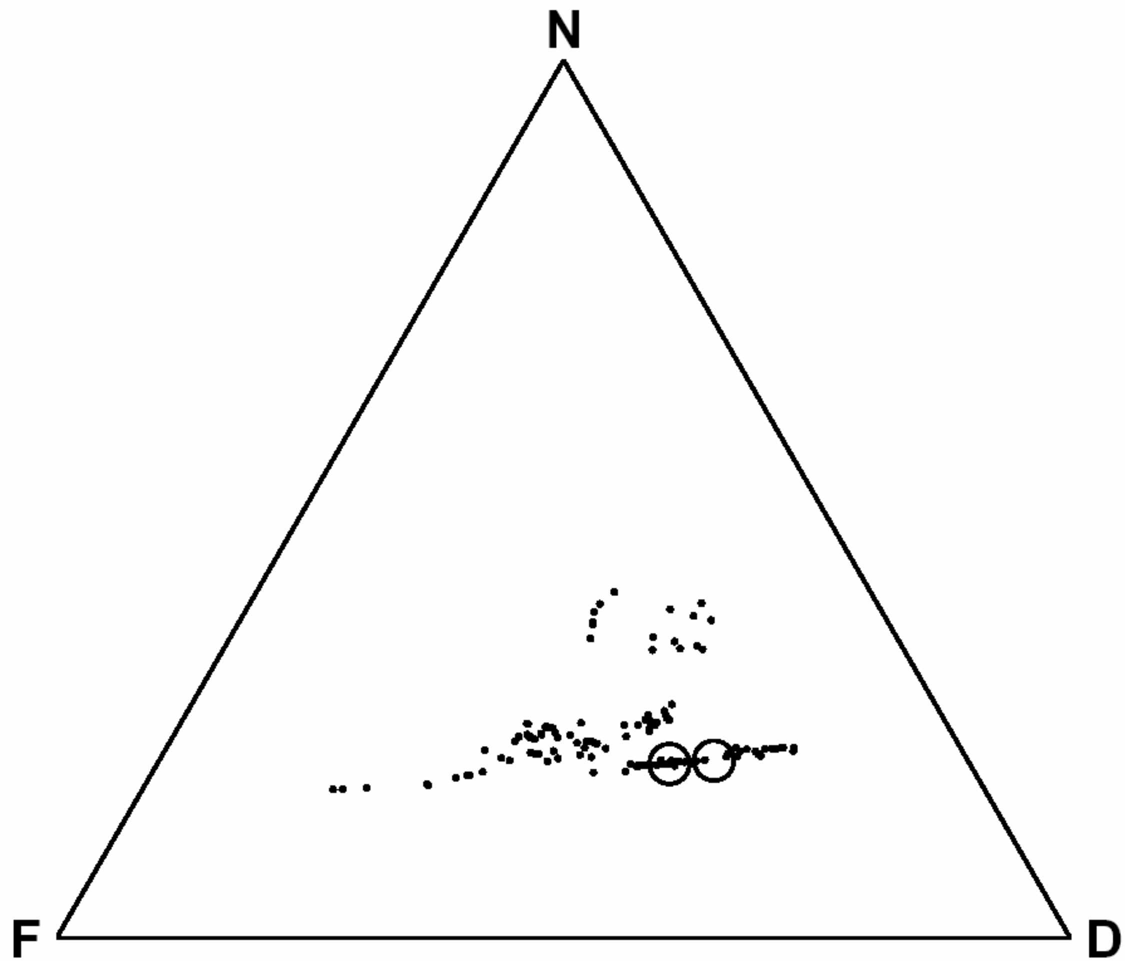

Individual Differences in Predictive Models

Cross Validation

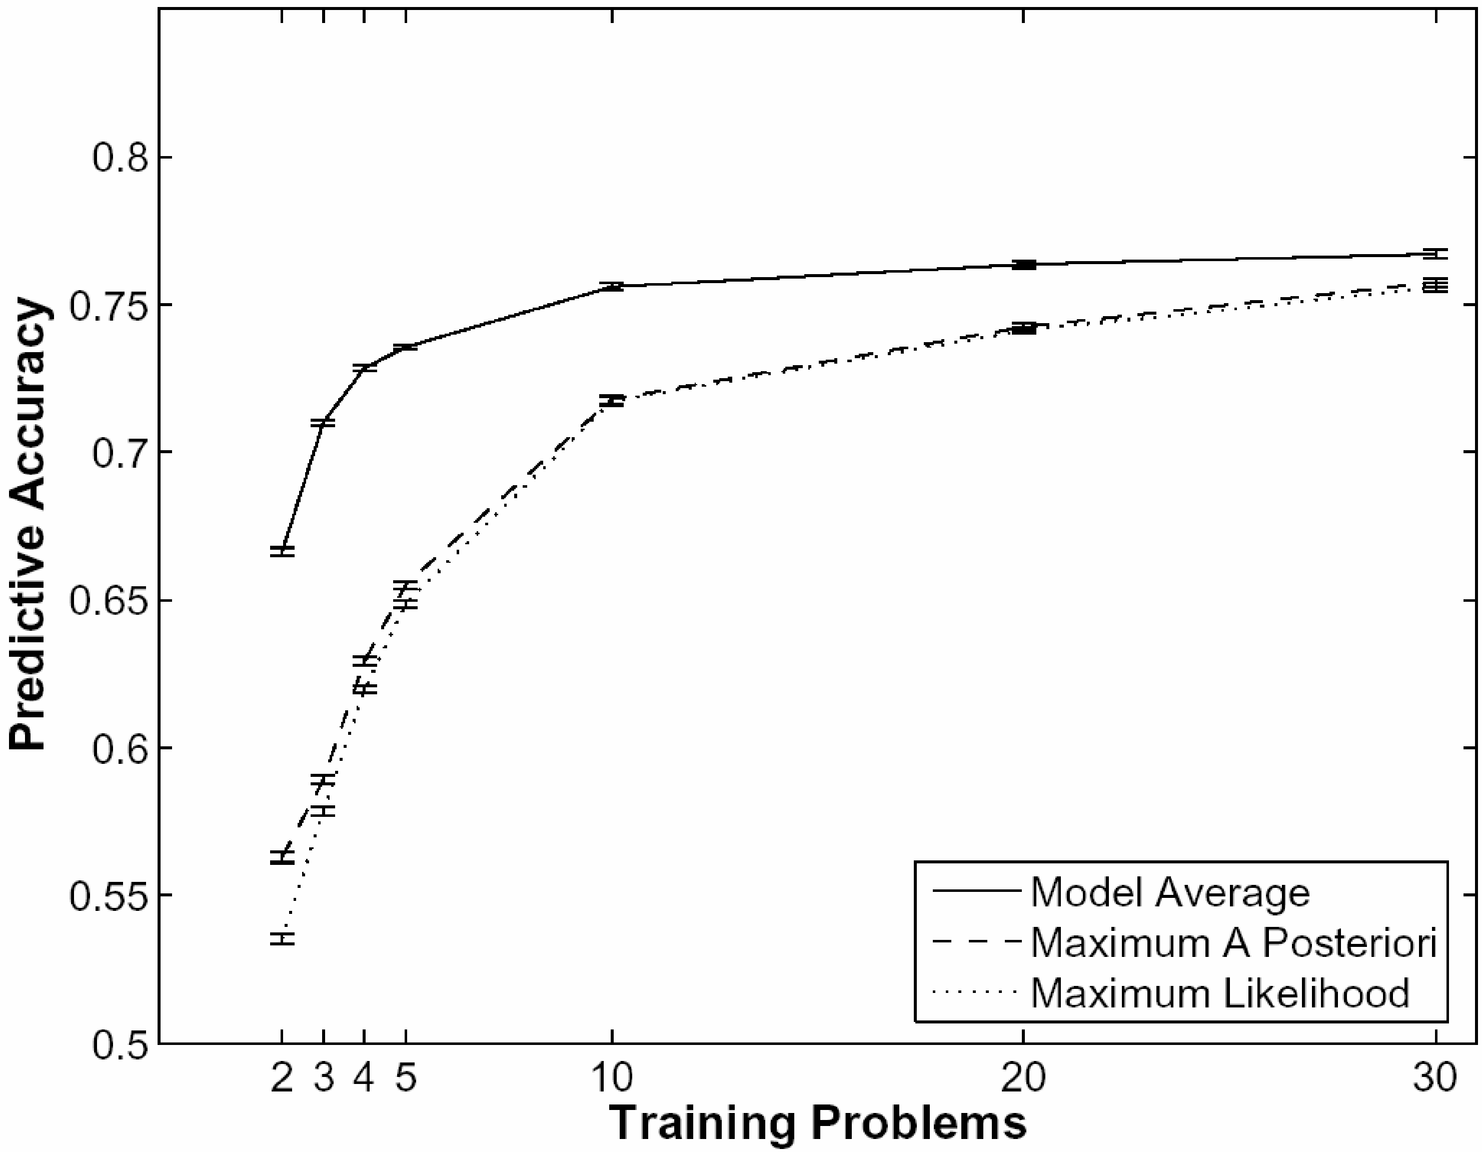

Preliminary Results of Current Work

Feedback and Reward

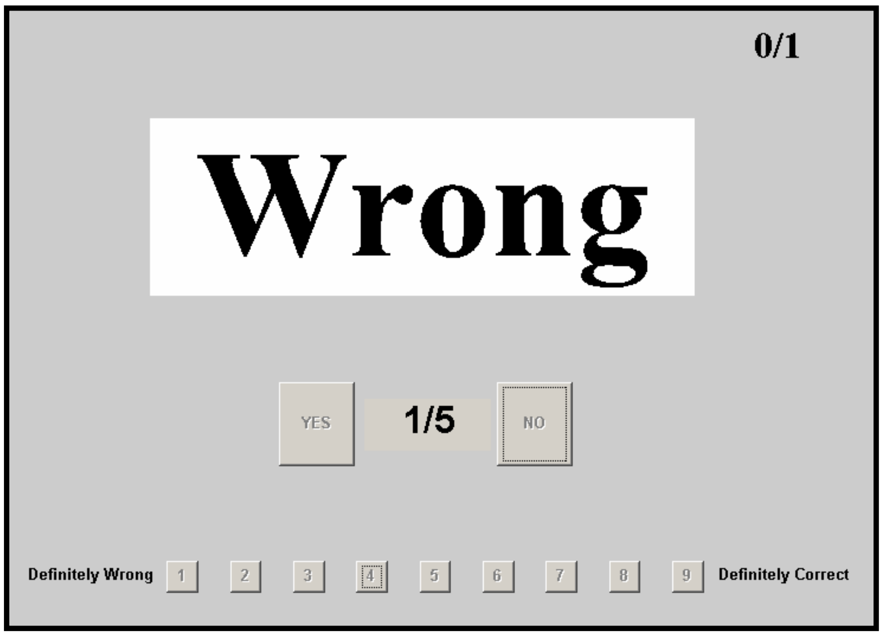

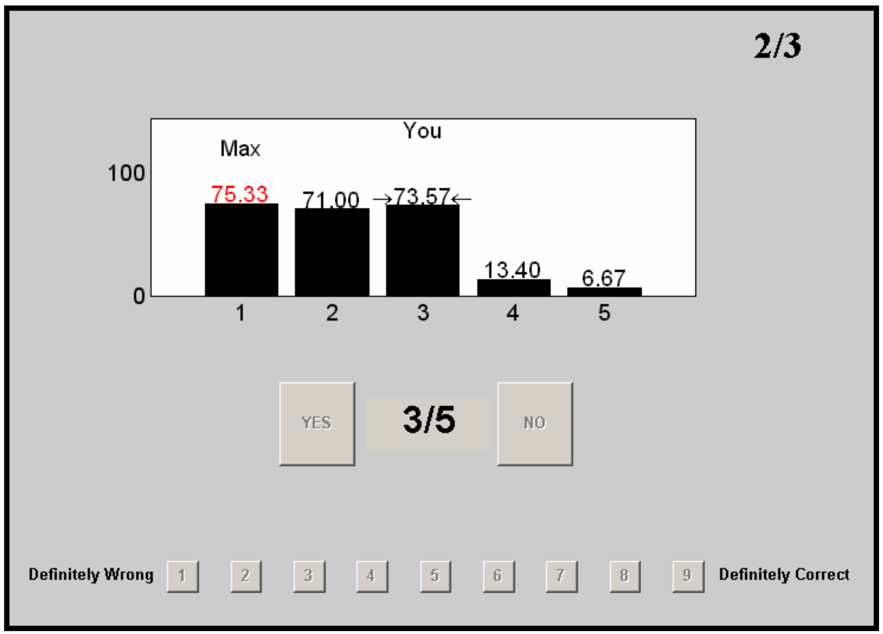

Evidence of Motivation, but not of Learning?

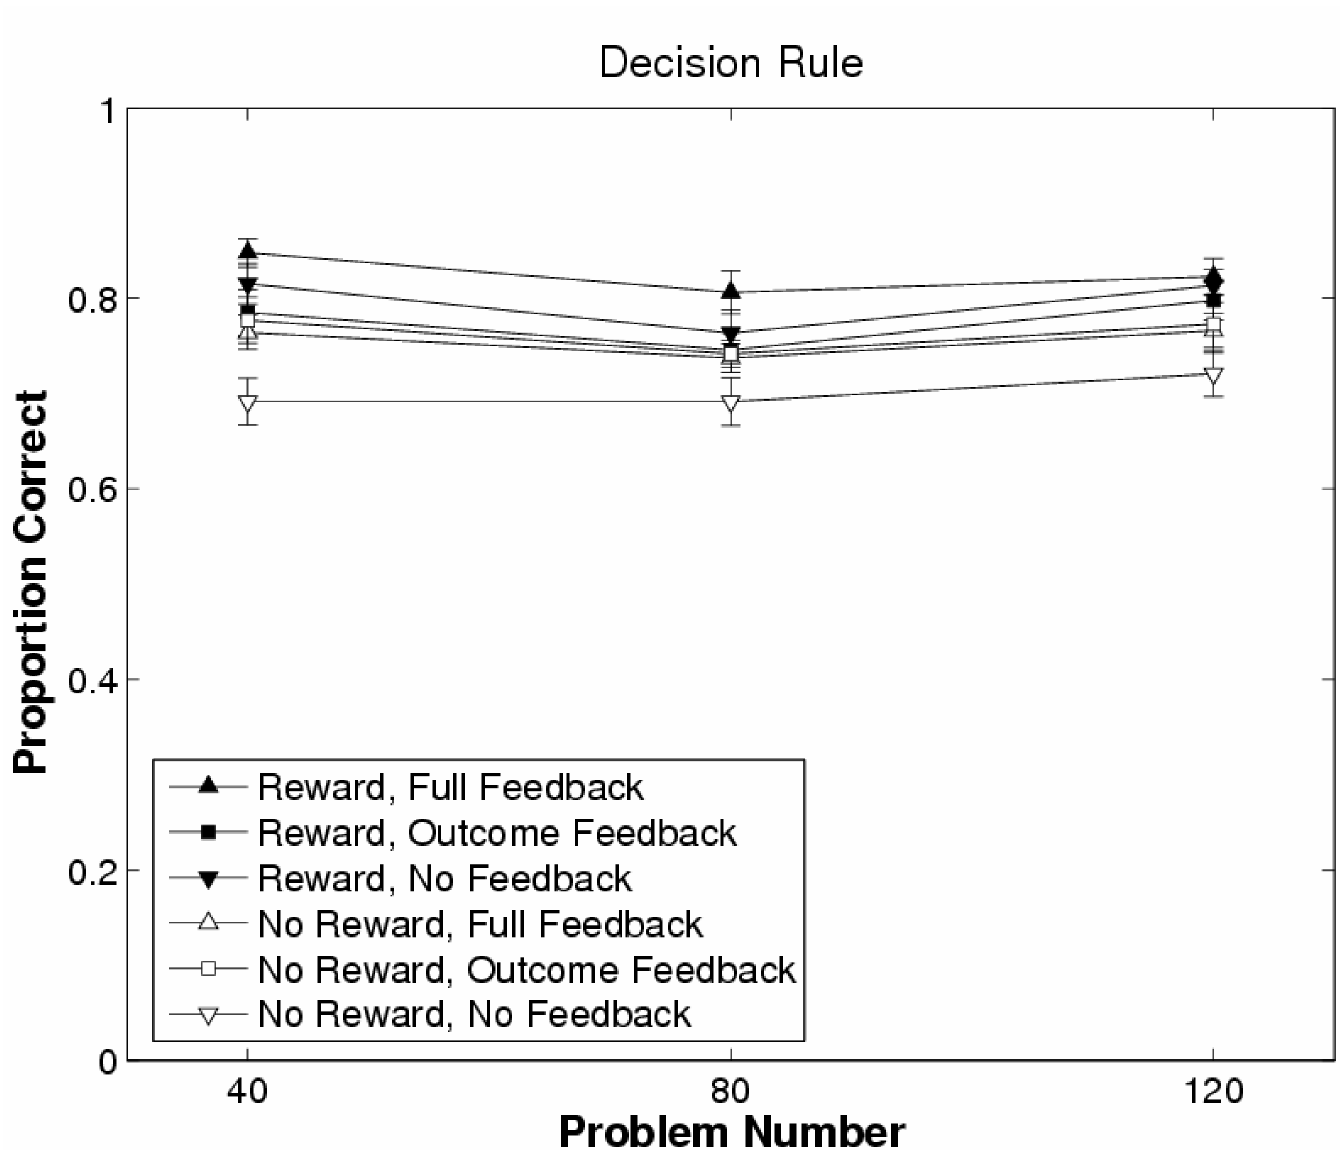

Interesting Feedback by Reward Interaction?

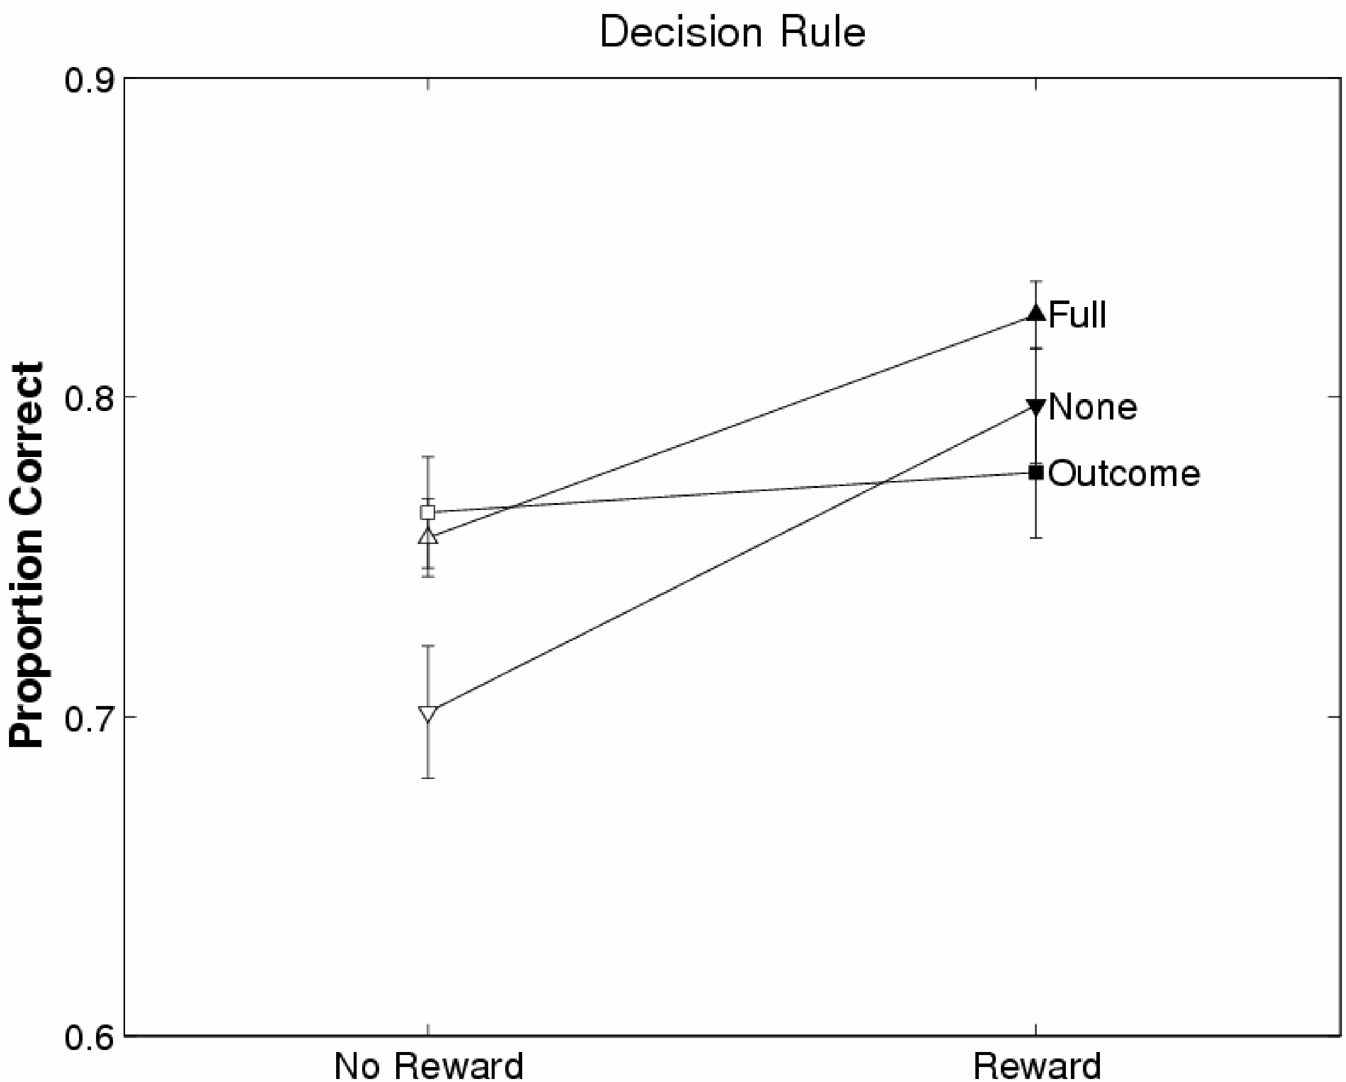

Group Decision-Making

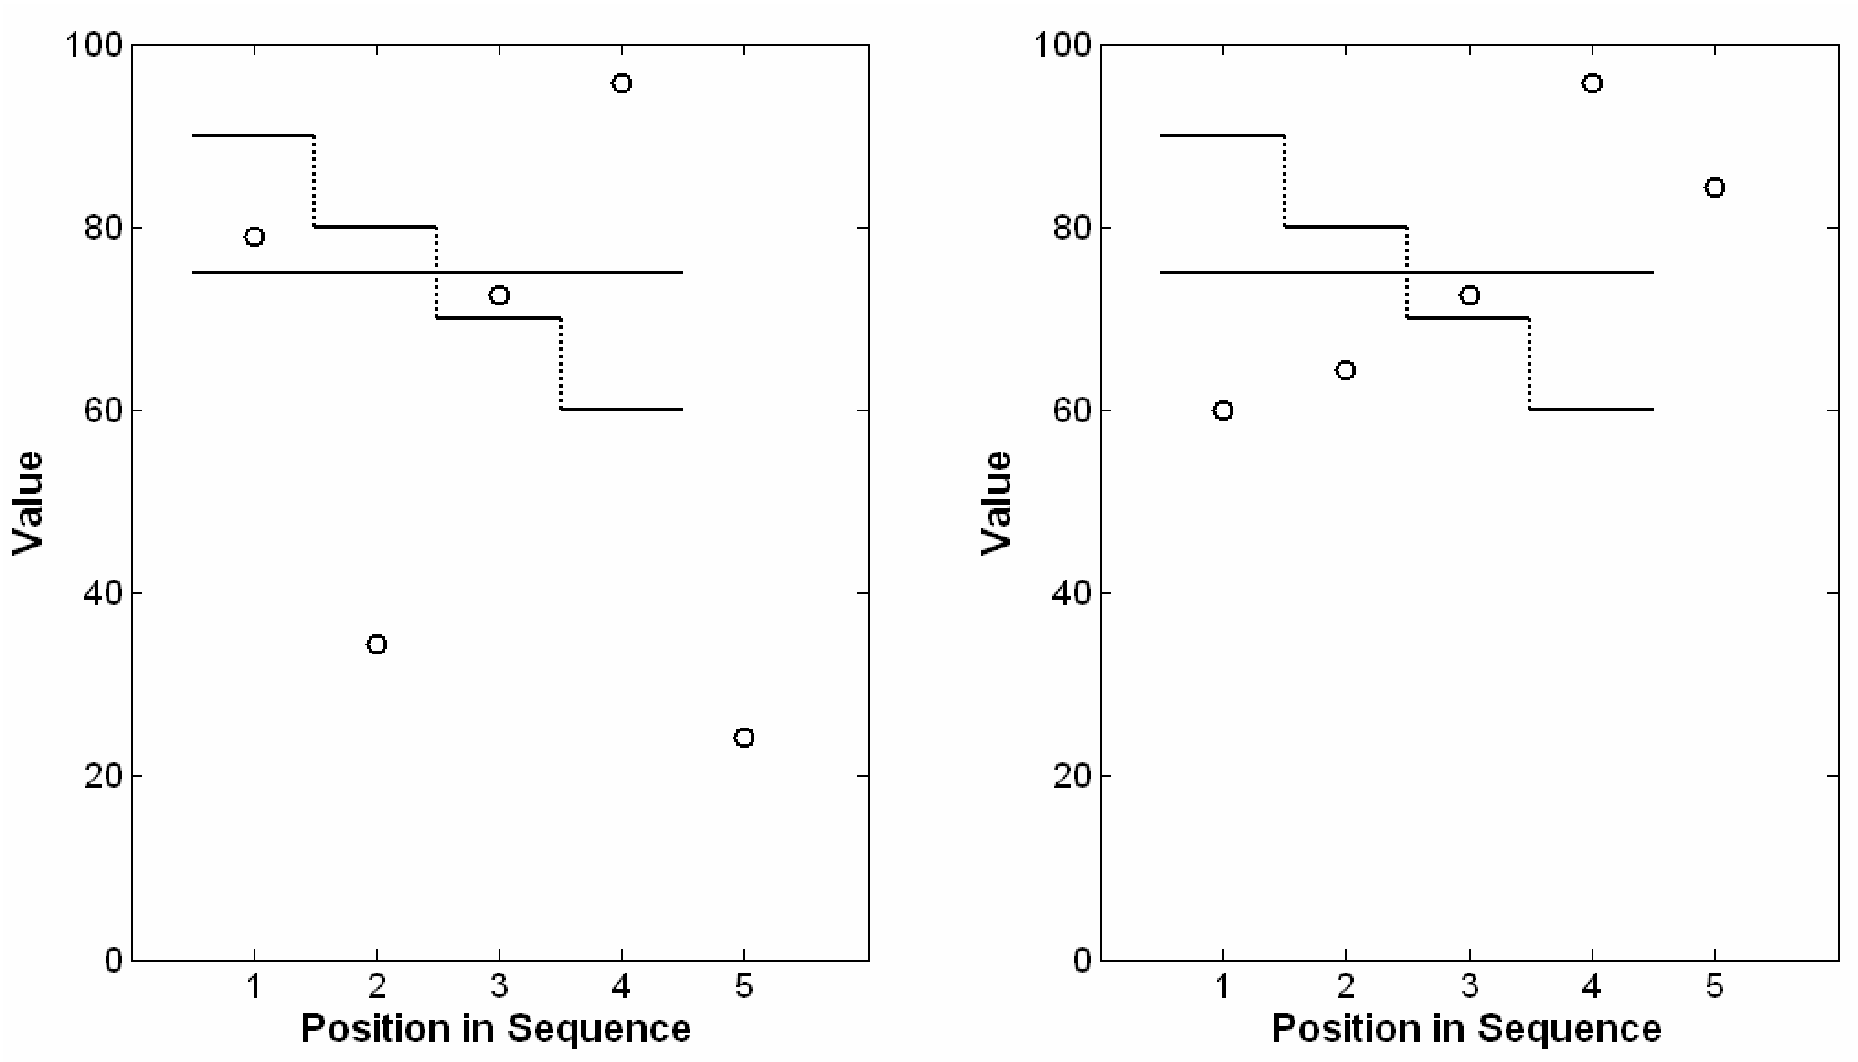

Leadership to Extremes?

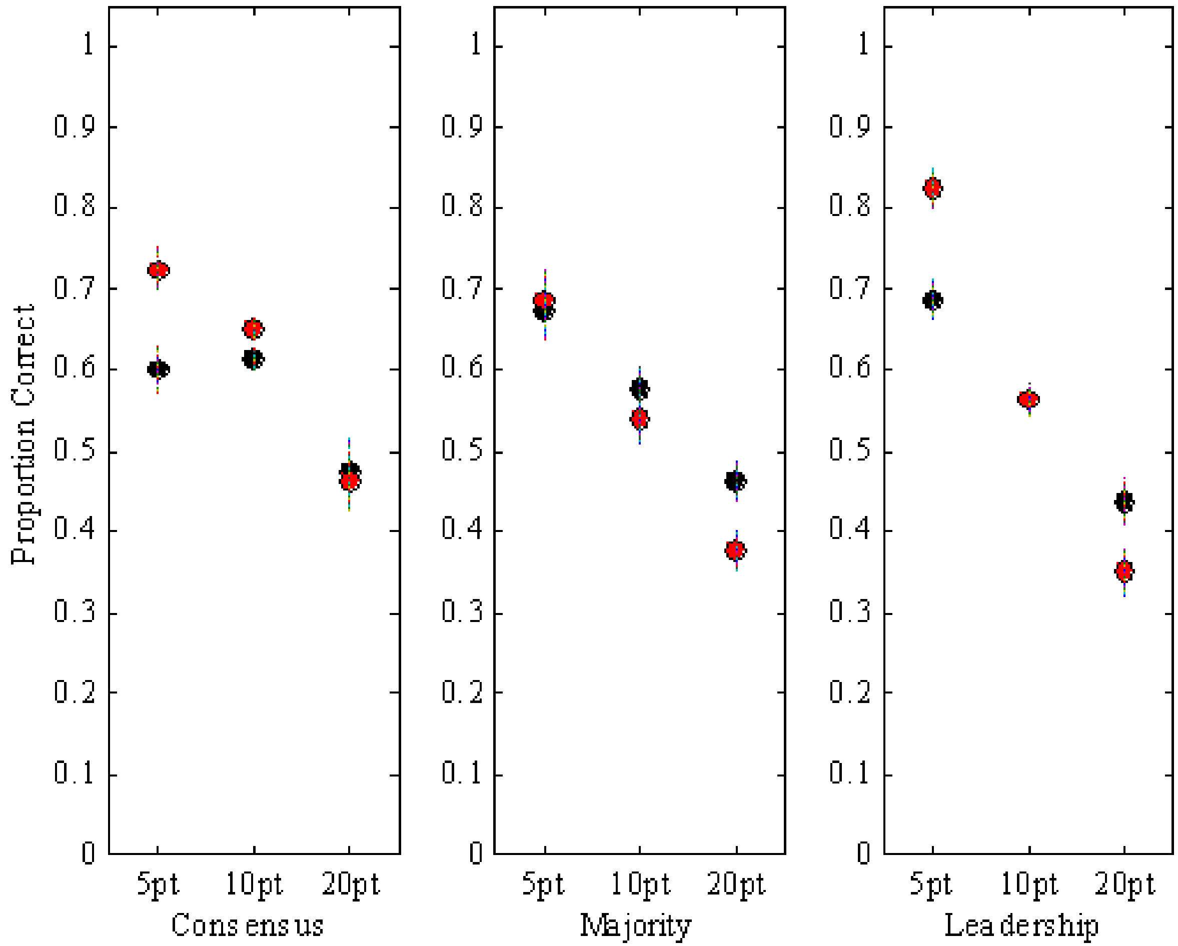

Thanks!
Questions?