Revealing Individual Differences in the Iowa Gambling Task

Lee I. Newman[1,2], Thad A. Polk[1,2] and Stephanie D. Preston[1]

Department of Psychology[1] & EECS[2], University of Michigan
Email: leenewm@umich.edu

Introduction

The Iowa Gambling Task (IGT) is a widely studied decision paradigm designed to simulate the task of deciding among competing options when the payoffs are uncertain and must be learned through experience. The relationship between brain areas involved in IGT performance[1,2,3,4] and in reward-based learning (RL)[5,6,7,8] have been established and have motivated the use of computational RL models to study the IGT. In prior work, a "standard" RL model has been fit to IGT data, and decision making in normal and patient populations has been characterized using group-averaged parameters[9,10,11,12]. In the present work we use a diverse set of RL models to investigate individual differences among healthy subjects and the efficacy of the class of RL models in capturing complex decision behavior.

Overview of the IGT

Current Profit +$260	A-	B-	C+	D+
	win +$80 lose -$250			
Consistent Gains:	$80-170	$80-170	$40-95	$40-95
Periodic Losses:	-$150-350	-$1250-2500	-$25-75	-$250-375
Loss Frequency:	50%	10%	50%	10%
Expected Net Value:	-$32	-$32	$16	$16

- Subjects are instructed to maximize net profits by choosing among four decks of cards over 100 trials.
- Payoff schedule imposes a tradeoff in decision making:
  - Decks A, B (bad decks) offer consistently larger gains, but have negative expected value.
  - Decks C, D (good decks) offer consistently smaller gains, but have positive expected value.
- Behavioral phenomena (in aggregate):
  - Initial exploratory choice to learn about payoffs.
  - Early preference for decks A and B (~ trials 10 to 30).
  - Shift to preference for decks C and D (~after trial 30).

Methods

Model Specification

Base Model

- Single-state Q-learning[13]; a standard RL model with three components:

(1) Reward Function: translates payoffs (gains, losses) experienced on trial t into an internal representation of reward r(t):

$$r(t) = W_g \cdot gain(t) + W_l \cdot loss(t) \quad (1)$$

where W_g and W_l weight the contribution of gains and losses to experienced rewards.

(2) Selection Function: computes choice probability for each deck based on stored value estimates V(t):

$$P_d(t) = \exp[\theta \cdot V_d(t)] / \Sigma_j \{\exp[\theta \cdot V_j(t)]\} \quad (2)$$

where j indexes the four decks and θ governs the sensitivity of choice to differences in the stored value estimates for each deck.

(3) Learning Function: stores and updates the behavioral value estimates V_d for each deck based on experienced payoffs:

$$V_d(t+1) = (1-\alpha) \cdot V_d(t) + \alpha \cdot r(t) \quad (3)$$

where d ∈ {A,B,C,D} and the learning rate α sets the relative influence of the current reward on the updated value estimate for the chosen deck.

Variant Models

- Theoretically motivated models chosen to explore individual differences.
- Models vary across the reward, selection, and learning components.
  - *Simple Averaging*: Greater influence of past rewards on learned value estimates than in the base model (less of a recency effect in learning).
  - *Value Decay*: Value estimates decay for decks not selected. Instantiates working memory limitations.
  - *Pursuit Selection*: Dissociates learning of values (estimation) and learning of actions (control). Instantiates the possibility that repeated actions become proceduralized with experience.
  - *Reinforcement Comparison*: Rewards represented relative to a learned reference level of average reward. Instantiates the possibility of "framing effects".
  - *Risk-Focused*: Reward is an inverse function of risk (net payoff variance, loss variance, or loss frequency) rather than net profit.
  - *Risk-Sensitive*: Reward is based on a weighted combination of net payoffs and risk. Instantiates the possibility that decision makers seek a balance between payoff maximization and risk avoidance.

Model Fitting and Comparison

Fitting

- Data from 41 healthy adults (age: M=32.4, SD=10.7); reported in [14].
- Models fit to individual trial-by-trial decision data.
- Parameters computed using numerical maximum likelihood.

Comparison

- Models compared using the Bayesian Information Criterion (BIC) which accounts for goodness of fit as well as model complexity.
- Each model compared to a nominal "dumb model" that makes choices using fixed probabilities based on each subject's selection proportions.
- Positive δBIC_m indicates that model m provides better fit to subject data than nominal model.

$$\delta BIC_m = 2\ln[L_m - L_{nom}] - (k_m - 3) \cdot \ln(100) \quad (4)$$

where L_m and L_{nom} are the likelihoods of model and the nominal model, and k_m is the number of free parameters in model m.

Results: Group Level

(A)

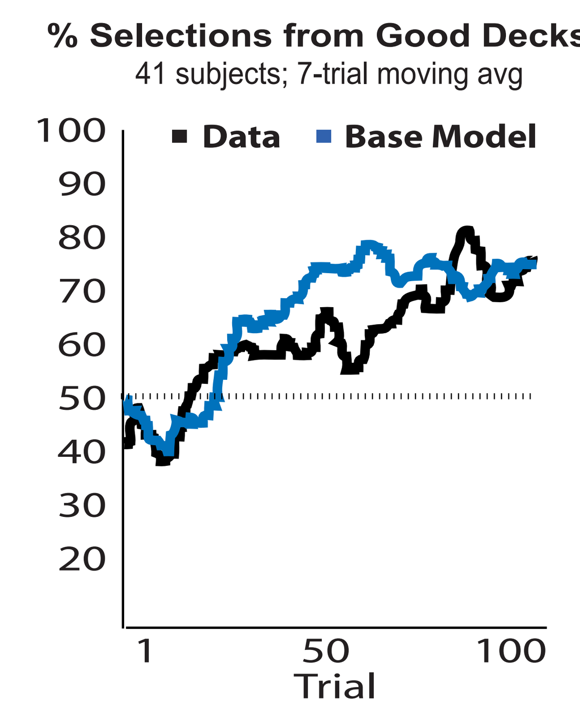

Model	Mean δBIC	%Partic. δBIC >0
Base	+13.6	61%
Risk-Sensitive	+14.2	56%
Decay	+12.8	56%
Rein. Comparison	+8.8	51%
Simple Average	+5.1	51%
Pursuit	+4.4	39%
Risk-Focused	-15.0	20%

Base Model

- Consistent with prior work, the *Base Model* reproduces core decision phenomena (Fig. A).
- Mean fit (δBIC) across subjects: +13.6,. SD=31.6 .
- Base model better than nominal model (δBIC>0) for 61% of subjects (Table, row 1).

Variant Models

- Positive fit (δBIC) for all models except *Risk-Focused* (Table, column1).
- Each model provides a positive fit (δBIC) for most/many subjects (Table, column 2).

Results: Individual Level

(B)

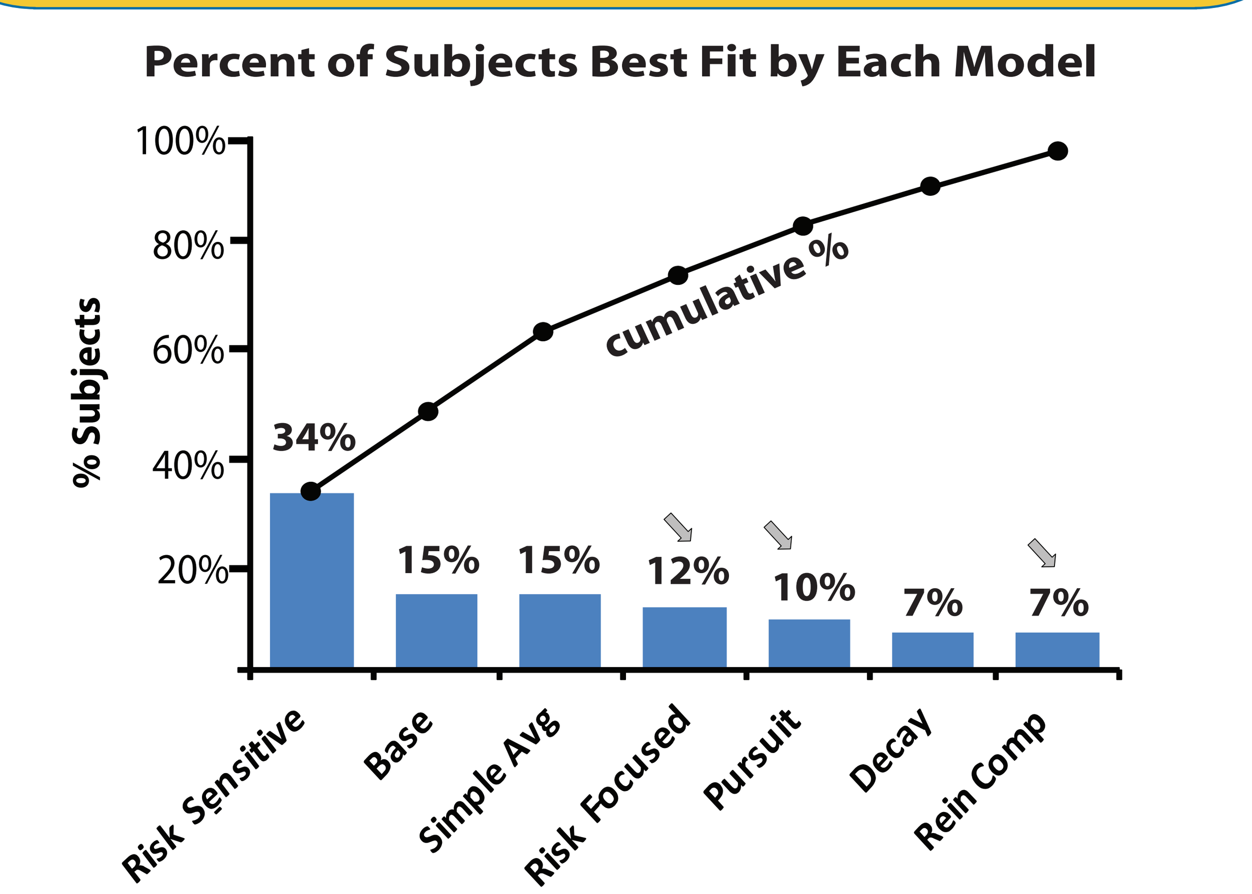

- The *Risk-Sensitive* model best fits the largest number of subjects (Fig. B), and over twice as many subjects as the base model.
- However, despite good fits at the group level (Table), no single model captures decision making for the majority of subjects; even the best *Risk-Sensitive* model provides the best fit for only 34% of the subjects.

(C)

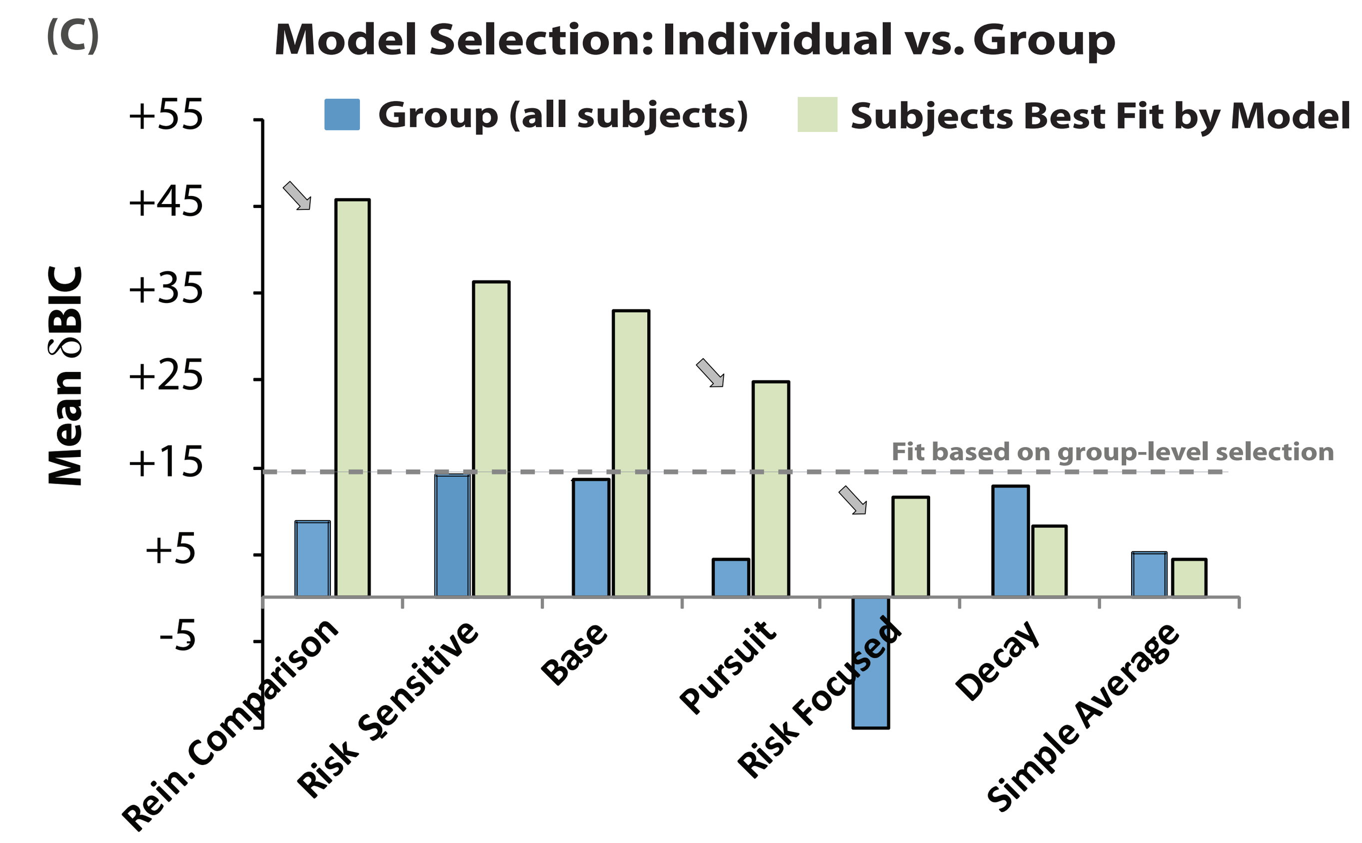

- Analysis at the invididual level reveals a set of models (Fig. C, gray arrows) that provide a good fit for subsets of subjects, despite being inferior at the group level:
  - *Pursuit*, *Reinforcement Comparison*, and the *Risk-Focused* models have inferior fits (mean δBIC) at the group level (blue bars) but provide good fits at the individual level (green bars).
  - These three models together are the best fitting models for 29% of the subjects (Fig. B, gray arrows).

Discussion

- Analyzing models at the individual level reveals subsets of decision makers not identified in prior work relying on group-level selection.
- One subset was best fit by the standard model, suggesting that choice behavior was driven by learning the expected value of the decks.
- The largest subset was best fit by a *Risk-Sensitive* model suggesting that choice behavior of these subjects was also guided by loss avoidance.
- Three models inferior at the group level revealed subsets of subjects that exhibited sensitivity to (i) the average level of experienced reward (*Reinforcement Comparison*), (ii) reward variance rather than magnitude (*Risk Focused*) and (ii) proceduralization of choice independent of estimated values of the decks (*Pursuit*).
- The *Decay* model provided a good fit for a small subset of subjects suggesting limitations in their ability to retain value estimates.
- Although there is strong evidence for the involvement of RL processes in the type of decision making engaged by the IGT, individual differences are not captured by differences in parameter values of any one RL model that has yet been applied to the task.
- We caution against characterizing IGT performance in clinical populations using the parameters of a single "best" model until more robust models are identified.

References

1. Bechara A, Tranel D & Damasio A (2000), Brain, 123, 2189-2202.
2. Bechara A, Damasio H & Damasio A (2003), ANYAS, 985, 1049-1056.
3. Clark L, Bechara A, Damasio H, et. al. (2008), Brain, 131, 1311-1322.
4. Stout J, Rodawalt W & Siemers E (2001), JINS, 7, 92-101.
5. Wallis J (2007), AnnRevNeuro, 30, 31-56.
6. Pagnoni G, Zink C, Montague R, & Berns G (2002), NatNeuro, 5, 97-98.
7. Seymour B & Dolan R (2008), Neuron, 58, 662-671.
8. Preuschoff K, Quartz S & Bossaerts P (2008), JNeuro, 28, 2745-2752.
9. Busemeyer J & Stout J (2002), PsychAssess, 14, 253-262.
10. Stout J, Busemeyer J, Lin A, et. al. (2004), PsychBullRev, 11, 742-747.
11. Yechiam E, Busemeyer J, Stout J, et. al. (2005), PsychSci, 16, 973-978.
12. Kalidindi K & Bowman H (2007), NeurNet, 20, 676-689.
13. Sutton R & Barto A (1998), Cambridge, MA, MIT Press.
14. Preston S, Buchanan T, et. al. (2007), BehavNeuro,121, 257-263.