Model parsimony and predictive power of computational models of cognition

Tilman Lesch
Department of Psychiatry
University of Cambridge
tcl36@cam.ac.uk

Mike Aitken Deakin
Department of Psychology
King's College London
michael.aitken@kcl.ac.uk

Barbara J Sahakian
Department of Psychiatry
University of Cambridge
bjs1001@medschl.cam.ac.uk

Abstract

The use of model parsimony to select appropriate models for analysing human behaviour can limit the explanatory depth and power of analysis by restricting the number of parameters included in the best-fitting model or the number of individuals included in the analysis. Here we present an extended q-learning model to investigate the effect of payoff framing on counterfactual updating. Parameter recovery is then used to determine whether preferring the simplest, plausible explanation gives the best measurement. The full and the parsimonious model recovered the original parameter values from the simulated data similarly. Although parameter values recovered by the full model were more variable it did not justify using the parsimonious model to investigate individual differences in parameter values estimated from the task behaviour. The present study provides a guideline for how parameter values based on an a priori model can be assessed to justify the use of a full model over a parsimonious.

Keywords: q-learning, model recovery, model parsimony, decision-making, uncertainty

Acknowledgements

This work was kindly funded by a Core Award from the Medical Research Council and the Wellcome Trust to the Behavioural and Clinical Neuroscience Institute (MRC Ref G1000183; WT Ref 093875/Z/10/Z). Tilman Lesch is supported by the Friedrich-Naumann-Foundation and the University of Cambridge.

1 Introduction

Mathematical models are commonly used in cognitive psychology to characterise processes in human behaviour such as behavioural strategies, susceptibility to biases or the role of environmental factors for behaviours [1], [2]. One of the challenges of such modelling is the selection of appropriate models to explain the data being analysed. Meaningful measures and criteria are needed to analyse overall suitability (absolute model fit) and to compare different models (relative model fit). Differentiating between competing models by evaluating their ability to predict behaviour is complicated by the fact that models often differ in a number of ways including number of parameters and complexity. The most commonly used methods in contemporary psychology to assess model fit are goodness of fit (GOF) and model parsimony (corrected GOF). Both GOF and model parsimony are methods of relative model comparison; they compare one model to another – relative model fit – but cannot assess absolute model fit or model validity – whether the model is appropriate to describe the observed data in the first place. Commonly used measures of goodness of fit are mean squared error or maximum likelihood: typical examples of model parsimony – corrected GOF – are the Akaike Information Criterion [3] or the Bayesian Information Criterion [4]. A measure of absolute model fit is model flexibility or generalizability; it describes the ability of a model to make valid predictions about behaviour not only for one task but also for several tasks [6]. A model that has more parameters may provide a better description of the observed data on a particular task; yet fail in terms of its ability to generalize to other tasks.

Some studies test several different models to identify the model, which best predicts the sample behaviour [5], [6]. This approach allows researchers to make inferences from the winning model's characteristics compared to the inferior models. A potential disadvantage of this approach is that the model that best fits to data might not necessarily reflect processes of psychological interest – limiting the ability to investigate group or individual differences based on model parameters. On the other hand, researchers can define an a priori model to characterise theoretical interest, and look at group or individual differences of the model parameters [7], [8]. This approach enables researchers to choose any model of interest to explain the data; however standard tools to assess model fit – GOF and model parsimony – are no longer appropriate as they assess only relative model fit. Therefore it is often unclear the extent to which the a priori model provides a useful description of the processes underlying the data.

This current study used a data set from a behavioural task typical of the application of RL model fitting in mainstream psychology. It sets out to determine whether for the case of nested models the use of the full model can be justified providing reliable parameter estimates for the analyses of group or individual differences using a novel probabilistic learning task. We specify an a-priori nested model to describe all processes of interest, and then compare parameter recovery of the full model GOF with a parsimonious model based on corrected GOF to determine which approach, if either, is preferable to recover true parameter values from simulated data.

2 A probabilistic behavioural task involving counterfactual feedback

There are considerable differences in the way in which positive and negative feedback is perceived, informs learning and leads to feelings of regret and behaviour to avoid potential negative feedback [9]. Counterfactual beliefs regarding "What could have happened" have been shown to strongly impact human behaviour besides the potential feeling of regret [10]. Behavioural data on a novel four-alternative probabilistic learning task were acquired from 83 participants. In the task individuals were instructed to maximize their reward by choosing continuously between four different stimuli. Feedback was presented not only for the selected stimulus but also for the forgone options (counterfactual feedback). The total of 220 trials was divided into three sections differing in the rate at which the probabilities changed between the four stimuli. On trials 1-40 probabilities remained constant; on trial 41 each option was assigned a different probability of reward, which remained constant between trials 41 to 120. On trials 121 through 220 the probabilities changed every 20 trials between the four stimuli. Throughout both the stable and dynamic conditions, the probabilities were programmed such that there was always one stimulus with a high (.85), one low (.15) and two with a 50/50 probability of reward. Each stimulus paid the same reward per trial if chosen and successful. There was no change in total reward for an unsuccessful choice.

3 A q-learning model of counterfactual feedback

To model the behaviour on this task, the classic q-learning model was extended such that after a trial on which stimulus j ε {1,2,3,4} was chosen, and outcomes O_i {1 = success, 0 = failure} revealed, for *each* stimulus i ε {1,2,3,4}, the value V_i was updated in the direction of O_i using a delta rule with Learning Rate α; tendencies of participants' learning to be modulated according to whether a stimulus was selected (*Selection Bias γ)*, the outcome of a stimulus (*Outcome Bias β)*, and the outcome of the selected stimulus (*Selection-Outcome Bias δ)*, was achieved by the use of parameters which modulated the learning on each trial.

Value Function: Expected reward of stimulus i on trial $t+1$: $Q_i(t+1)=\alpha*\gamma*\beta*\delta*O_i(t)+(1-\alpha*\gamma*\beta*\delta)*Q_i(t)$

With trial t ε {1,2,…,220}, stimulus i ε {1,2,3,4}; outcome stimulus i on trial t: $O_i(t)$ ε {0,1}; expected value Q_i for option i on trial t: $Q_i(t)$; α always between 0 and 1; γ = 1 if $i=j$ (*i.e.* selected), free otherwise; β=1 if $O_i(t)$=1, free otherwise, and δ=1 if $i=j$ *or* if $i\neq j$ and $O_i(t)$=0 (i.e. selected or did not win), free otherwise. The value of α *(Learning Rate),* γ (*Selection-Bias*), β (*Outcome-Bias*), and δ (*Selection-Outcome-Bias*) was bound between zero and one.

To model participants' choices based on the fully parameterized model described above a softmax choice rule with one parameter τ (*Exploitation Rate*) for model stickiness was used – how likely individuals follow the model-based reward-maximizing prediction [11]. The probability of choosing stimulus i is increased by increased value, and decreased by increases in the value of other stimulus values.

Choice Function: probability stimulus i chosen on trial $t+1$: $P_i(t+1)=\frac{\exp(Q_i(t)*\tau)}{\sum_{k=1}^{4}\exp(Q_k(t)*\tau)}$

We followed the typical approach to fitting the model for each individual. Values of all five free parameters (*Learning Rate α, Selection Bias γ, Outcome Bias β, Selection-Outcome Bias δ* and *Exploitation Rate τ*) that maximized the likelihood of the behavioural data were determined for each participant separately [12]. Corrected goodness of model fit was assessed using the Bayesian Information Criterion [BIC, 5]. For each participant, a set of model parameters was obtained, derived on the performance on both the first and second run through the task using the full model and Maximum Likelihood Estimation [MLE, 12].

The key question is the degree to which this 'standard' approach may be considered to have correctly characterised the individual variation in the psychological processes of interest. To answer this question, we investigated performance of the modelling procedure under the assumption that it had been successful. 30 sets of pseudo-behavioural data were then generated for each individual participant's set of model parameters. Afterwards a total of eight nested models were fitted to every simulated data set and the parameter values, which maximized the likelihood of the data, were noted (for each of the 30 simulations for each of 83 participants). The eight models are as follows. Model 1 included two free parameters (Learning Rate α and Exploitation Rate τ, original q-learning with δ-rule); Models 2a, 2b and 2c included both parameters from Model 1 and one out of the three bias parameters (γ, β or δ). Models 3a, 3b, and 3c included both parameters from Model 1 and two out of three bias parameters. Model 4 included all parameters (full model); it is the same model that was used to generate the simulated data sets. All models are nested in *Model 4* by setting the appropriate bias parameters equal to one.

If the task and modelling procedure are a reasonable method for estimating an individual's parameters, then the mean for each model parameter across the 30 data sets should reflect (recover) the corresponding parameter used to generate the pseudo-behavioural data. The similarity of the 30-parameter estimates obtained in this way to the parameters *known* to have generated these data sets thus reflect a measure of the suitability of a model-fitting process (typical to contemporary psychology) as a whole to this task.

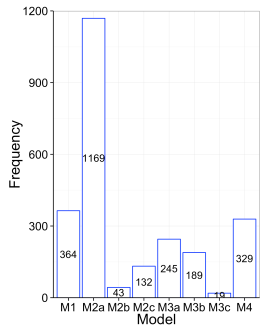

Figure 1: Frequency tables of the most parsimonious model

3 Results

In *Figure 1* the frequency of each model being the best model according to a parsimonious account was plotted across all participants and simulations. For more than half of the simulated data sets that were generated, the best model according to a parsimony criterion was not the full model that was used to generate the data, but rather Model 2a with three free parameters: Learning Rate α, Exploitation Rate τ and Selection Bias γ.

To further assess parameter recovery the original parameter values and the recovered parameter values for the full model (M4) as well as the original parameter values and the parsimonious model were plotted (Figure 2); the left plot depicts the results for the "Full Model" and right plot for the "Parsimonious Model". On the X-axis the parameter values used to generate the pseudo-behavioural data; black dots represent the recovered parameter value for each of the simulations per individual participant. The black line is the linear regression line with original parameter values as predictor variables and the simulated parameter values across all simulations as predicted values. The red line is the angle bisector in the first quadrant. The graph for the parsimonious model contains only the free parameters of the parsimonious model; parameters not in the parsimonious model or equal to one were excluded. If there were perfect recovery and minimal variance, all parameter values would haven been scattered around the angle bisector (red line, ideal) in the first quadrant.

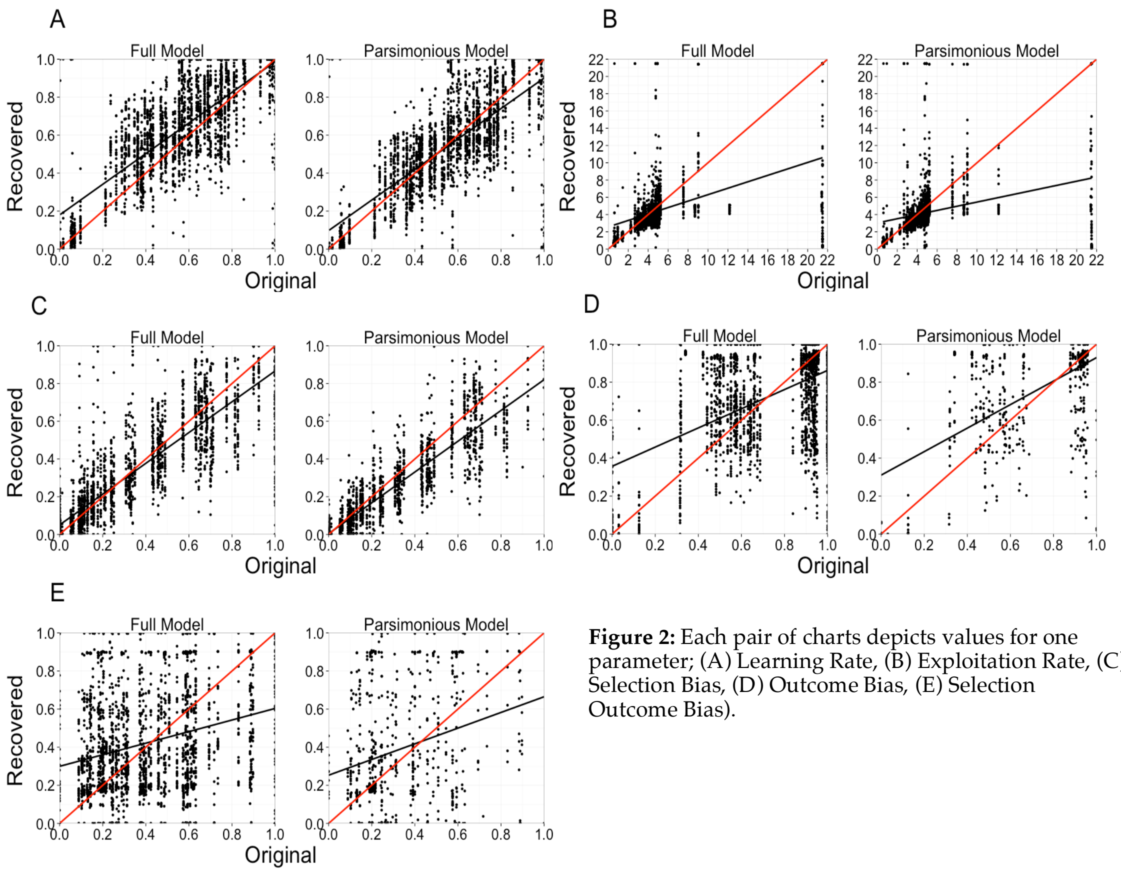

Figure 2: Each pair of charts depicts values for one parameter; (A) Learning Rate, (B) Exploitation Rate, (C) Selection Bias, (D) Outcome Bias, (E) Selection Outcome Bias).

To quantify the discrepancy between parameter values used to simulate the data and parameter estimates recovered from pseudo-behavioural data, the bias and deviation for the full model and parsimonious model were calculated (Table 1). Bias is the mean difference between the recovered parameter estimate and the original parameter value averaged across all simulations. Deviation is the square root of the mean squared difference between the recovered parameter estimates and the original parameter value used to generate the data averaged across all simulations.

Model:	Full Model		Parsimonious	
	Bias	Deviation	Bias	Deviation
Learning Rate α	-.072	5.014	**.013**	**4.485**
Exploitation Rate τ	**.668**	**59.984**	.888	68.513
Selection Bias γ	**.026**	3.618	.045	**2.574**
Outcome Bias β	**.012**	6.274	-.0445	**2.893**
Selection-Outcome Bias δ	**.021**	8.513	-.054	**4.309**

Table 1: Bias and Deviation between original parameter values and recovered parameter values using the best fitting model and the parsimonious model; **BOLD** = minimum value.

For all parameters there was a positive linear trend between the generating parameter and the recovered parameter, confirming that the parameters of the generating model were, to an extent, recoverable by the model fitting used. However, there was considerable variability in the parameter estimates recovered from pseudo-behavioural data for each set of parameters. The positive linear relationship was clearest for Learning Rate α and Selection Bias γ suggesting the most reliable recovery of those parameter values.

The two approaches (full versus parsimonious model estimation) performed in a similar fashion, with bias being smaller for the true model relative to the parsimonious model, and the deviation was smaller for the parsimonious model compared to the full models. Larger variability implies that parameter estimates based on a single behavioural test sessions would be less well able to characterise the processes underlying behaviour of a specific individual. However, while the parsimony-based estimates gave less discrepant recovery of the true parameters, it must be noted that this approach produced no estimates of the final two parameters (β and δ) for over 50% of simulated data sets.

4 Concluding Remarks

The current study used parameter recovery from simulated data to investigate the reliability of simple RL modelling fitting as frequently performed within the mainstream psychology literature. The results showed that using an *a priori* model with high explanatory power – and a larger number of parameters than a more parsimonious model - may be justified where estimates for all parameters are required. Selecting estimates only from the most parsimonious models tended, unsurprisingly, to generate estimates closer to the generating parameter: however, this approach has the disadvantage that many individuals in the sample could not be used for analysis of group or individual differences in model-based parameter values, and there is a suggestion the parsimony restriction could generate a bias in the estimated parameters.

References

[1] J. R. Busemeyer, J. C. Stout, and P. R. Finn, "Using Computational Models to Help Explain Decision Making Processes of Substance Abusers," in *Cognitive and Affective Neuroscience of Psychopathology*, D. Barch, Ed. Oxford University Press, 2003, pp. 1–41.

[2] A. D. Redish, "Addiction as a computational process gone awry," *Science (80-.).*, vol. 306, no. 5703, pp. 1944–7, Dec. 2004.

[3] H. Akaike, "A new look at the statistical model identification," *IEEE Trans. Automat. Contr.*, vol. 19, no. 6, pp. 716–723, 1974.

[4] G. Schwarz, "Estimating the dimension of a model," *Ann. Stat.*, vol. 6, no. 2, pp. 461–464, 1978.

[5] M. a Pitt, I. J. Myung, and S. Zhang, "Toward a method of selecting among computational models of cognition.," *Psychol. Rev.*, vol. 109, no. 3, pp. 472–491, 2002.

[6] J. R. Busemeyer and J. C. Stout, "A contribution of cognitive decision models to clinical assessment: Decomposing performance on the Bechara gambling task.," *Psychol. Assess.*, vol. 14, no. 3, pp. 253–262, 2002.

[7] E. Yechiam and G. Hochman, "Loss-aversion or loss-attention: the impact of losses on cognitive performance.," *Cogn. Psychol.*, vol. 66, no. 2, pp. 212–31, Mar. 2013.

[8] J. C. Stout, S. L. Rock, M. C. Campbell, J. R. Busemeyer, and P. R. Finn, "Psychological processes underlying risky decisions in drug abusers.," *Psychol. Addict. Behav.*, vol. 19, no. 2, pp. 148–57, Jun. 2005.

[9] E. Yechiam and G. Hochman, "Losses as modulators of attention: review and analysis of the unique effects of losses over gains.," *Psychol. Bull.*, vol. 139, no. 2, pp. 497–518, Mar. 2013.

[10] G. Coricelli, H. D. Critchley, M. Joffily, J. P. O'Doherty, A. Sirigu, and R. J. Dolan, "Regret and its avoidance: a neuroimaging study of choice behavior.," *Nat. Neurosci.*, vol. 8, no. 9, pp. 1255–62, Sep. 2005.

[11] R. D. Luce, *Individual choice behavior*. New York, NY: John Wiley, 1959.

[12] I. J. Myung, "Tutorial on maximum likelihood estimation," *J. Math. Psychol.*, vol. 47, pp. 90–100, 2003.