What is computational psychiatry good for?

Michael Browning[1,2,*,] Martin Paulus[3], Quentin J. M. Huys[4-6]

1 Department of Psychiatry, University of Oxford, UK.

2 Oxford Health NHS Trust, Oxford, UK.

*michael.browning@psych.ox.ac.uk

3 Laureate Institute for Brain Research, University of Tulsa, Tulsa, Oklahoma.

4 Division of Psychiatry, University College London, London, UK

5 Max Planck UCL Centre for Computational Psychiatry and Ageing Research and Wellcome Centre for Human Neuroimaging, Queen Square Institute of Neurology, University College London, London, UK.

6 Camden and Islington NHS Foundation Trust, London, UK.

It is rare in the field of biological psychiatry for hypotheses to be definitively refuted. Rather, topics of investigation drift into and out of fashion, often driven by the initial excitement of technological innovation, followed by the necessary corrective of nuanced or underwhelming clinical results. A well-known example of this is the association between depression and abnormal function of the HPA axis, as measured using the dexamethasone suppression test (DSST; 1). This observation led to a great deal of work investigating whether the association might help us identify useful subtypes of depression (2) or predict treatment response (3). As it turned out, the specificity and predictive value of the DSST was not thought to be of a level that would be useful clinically and the topic has gradually moved out of the spotlight. We are left in the familiar position of knowing that non-suppression of the DSST is associated with depression, but not being sure how we can use this knowledge to help patients.

An understandable response to this situation is to try and develop better measures—if we had a more sophisticated version of the DSST, then perhaps we would be able to realise our clinical aspirations. In this commentary we make two suggestions. First, that an alternative approach is to ask more precise questions, and second that an underappreciated application of computational techniques is that they may help us to do so.

Computational Answers: Much of the ambition of computational psychiatry to date has been on the use of formal models to find hidden answers to interesting clinical questions. For example, a recent study asked whether computational modelling might help explain previously described electroencephalography (EEG) and magnetic resonance imaging (MRI) data from patients with schizophrenia (4). In this important study, the authors replicated the group differences between patients and controls using resting state, mismatch negativity and 40-Hz auditory paradigms. They then applied a neural mass model to assess what changes at the microcircuit level might produce these effects, concluding that they could all be accounted for by reduced synaptic gain on pyramidal cells. In other words, the benefit of the modelling was that it linked EEG and MRI results to putative causal processes at the microcircuit level. In this example, the computational model is being used

analogously to the DSST task described above—as a method of measuring a hidden, but hopefully clinically important process. In the next section we consider how models may be useful in a different way; to ask, rather than answer, questions.

Computational Questions: A useful feature of computational models is that they don't just tell us what they know, they tell us what they don't know. This can be useful when we are deciding what questions we should answer. For example, turning to the example of how depressed patients learn from rewarding experiences, a common finding across the literature is that, when presented with two choices, one more rewarding than the other (Figure 1a), patients with depression will select the more rewarding choice less consistently than people who are not depressed (Figure 1c; ,5). One way of interpreting this finding is that reinforcement learning processes are disrupted in depressed individuals. Relatively simple reinforcement learning algorithms (Figures 1b) can successfully emulate the choices of both patients and controls. The first type of algorithm that was used in this situation had two free parameters (6), a learning rate controlling how quickly value learning occurs, and an inverse decision temperature, controlling how much the model uses its learnt values when selecting an action (Figure 1; b,c,d). This algorithm attributed the behaviour of depressed patients to a reduction in the inverse decision temperature. That is, patients were learning about the choices normally, but were not selecting the best option because their decisions were less influenced by the values they had learned. Later on, a second generation of algorithm was developed, in which the inverse decision temperature parameter was replaced by a reward sensitivity parameter, which controls how rewarding the reinforcer used in the task (e.g. points won, money etc) was judged to be. Using this model, the behaviour of depressed patients was explained by a reduction in reward sensitivity (7). In other words, depressed patients selected the best option less frequently due to a specific reduction in how rewarding they found the reinforcer used in the task.

These two classes of model explain the same behaviour in conceptually completely separate ways, as either an effect of how patients use information when they make decisions, or as a consequence of reduced estimates of the reward value of the reinforcer. One way to arbitrate between these competing explanations might be to use a more sophisticated model in which both the inverse decision temperature and reward sensitivity parameter are allowed to vary and ask which parameter differs in depressed patients relative to controls. The result of using this model is illustrated in Figure 1e; the model is unable to estimate reward sensitivity or inverse decision temperature as the two parameters are mathematically redundant (7) and therefore completely interchangeable. Thus, even though the parameters represent conceptually distinct hypotheses about the underlying cognitive processes, choice behaviour on these simple tasks is not able to adjudicate between them. Here the important thing the model is telling us is that it doesn't know whether the choice behaviour of patients is caused by changes in the inverse decision temperature and/or by changes in reward sensitivity.

In this case, the model is useful not because it tells us what the answer is, but because it tells us, precisely, what the question should be; is reduced reward choice in depression caused by difficulty in using learnt values when making decisions or by a lower value of reinforcers? Beyond this, the model suggests how these questions might be answered; if depression is associated with difficulty in using the value of options to make decisions, then the choices of depressed patients should be less consistent even when values of options are explicitly presented, and no learning is required (8; reports evidence against this prediction). Alternatively, if patients treat reinforcers as less rewarding, then this effect should be apparent even when no decisions are required (e.g. where response to a single rewarded stimuli is measured (9)). To date, the literature is not consistent with a simple decision effect, but is generally consistent with reduction in reward sensitivity as the most likely

process associated with depression, raising interesting subsidiary questions about why this might occur (5).

In summary, computational models can be used to identify hidden processes, some of which might be useful for answering clinically interesting questions. But models also tell us when they are unable to discriminate between competing explanations, and when this occurs, are a useful way of framing the precise mechanistic questions we should be trying to answer to improve our measures.

The development of clinically useful measures from biological research requires us to ask questions that are sufficiently specific that they may be refuted. Computational models help us to ask these questions.

Acknowledgements and Disclosures

MB is supported by the Oxford Health NIHR Biomedical Research Centre. QH is supported by the UCLH Biomedical Research Centre. The views expressed are those of the authors and not necessarily those of the NHS, the NIHR or the Department of Health. MB has received travel expenses from Lundbeck for attending conferences, and has acted as a consultant for J&J, Novartis, P1vital Ltd and CHDR. He owns shares in P1vtial Products Ltd. QH has received options and consultancy fees from Aya Health and Alto Neuroscience, and a research grant from Koa Health. MPP reports the following disclosures: research/grants: the William K. Warren Foundation and NIH Grant No. R01DA016663, NIH/NIGMS Grant No. P20DA027834, and NIH Grant Nos. R01DA027797, R01DA018307, U01DA041089, and 1R01MH101453; consulting (last 3 years): has received royalties for an article about methamphetamine use disorder from UpToDate.

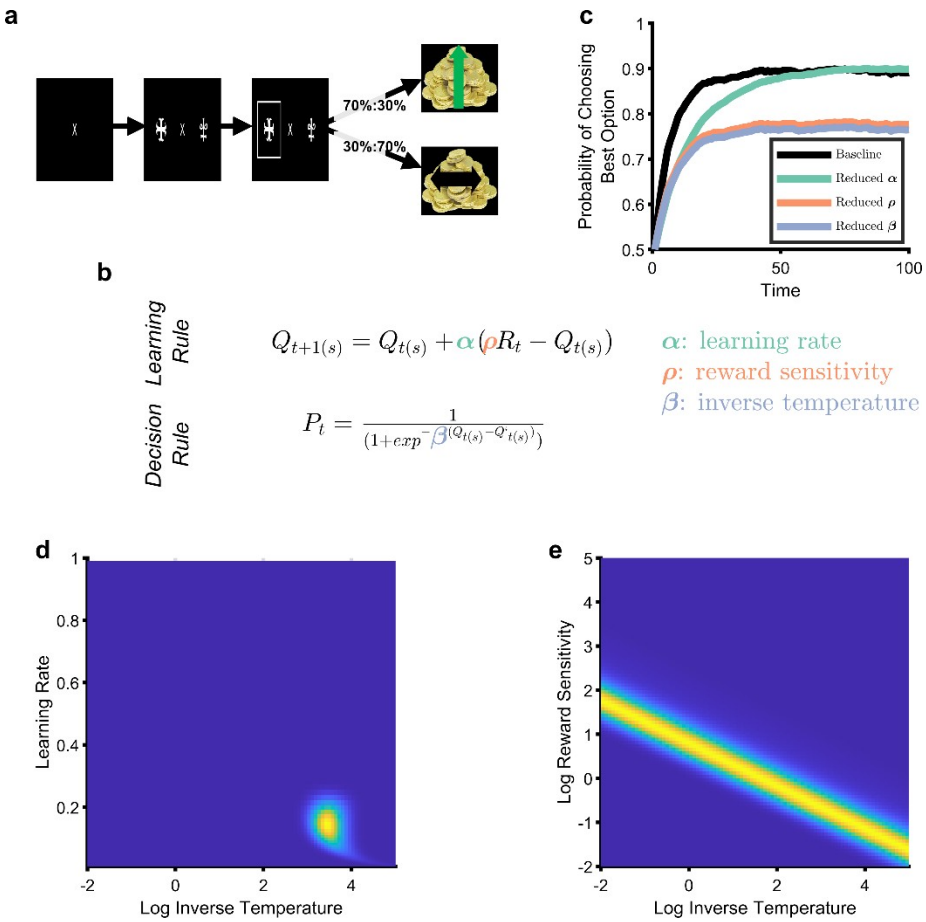

Figure 1: What computational models don't know. a) a commonly used learning task in which participants must choose between two options, one of which is associated with a higher probability of reward than the other. A commonly observed finding is that depressed patients select the most rewarded option less consistently than non-depressed participants. b) A simple computational model describes how participants may solve this task; Q_t, the value of an option, is learnt using a simple updating process and is fed into a decision rule. The behaviour of the model can be influenced by a number of different parameters including; learning rate (α), reward sensitivity (ρ) and inverse decision temperature (β). c) An illustration of the effect of changing these parameters on choice in the task; the learning rate influences the rate at which the model reaches a behavioural plateau, both the reward sensitivity and inverse decision temperature control the level of the plateau. The behaviour of depressed patients can be captured by either reducing reward sensitivity or inverse decision temperature. d) The posterior estimate of parameter values after fitting to a participant's choices, using a model in which learning rate and inverse decision temperature are allowed to vary (but reward sensitivity is fixed). Possible values of the inverse temperature parameter are represented along the x axis, possible values of learning rate are represented on the y axis. As can be seen, the parameters are precisely estimated, with learning rate lying between 0.1-0.2 and the log inverse temperature between 3.5-4. Yellow corresponds to the highest posterior probability, i.e. the most likely parameter value given the observed behaviour. e) The marginal posterior estimate of reward sensitivity and inverse temperature when all three parameters are allowed to vary and the model is fitted to the same data as d (for simplicity, the marginal probability of learning rate is not shown). Because the model is unable to discriminate between reward sensitivity and inverse decision temperature it is unable to estimate either—any value of inverse temperature or reward sensitivity is possible. In other words, the model doesn't know whether participant choices are influenced by changes in reward sensitivity or decision temperature.

References

1. Carroll BJ, Martin FI, Davies B (1968): Resistance to suppression by dexamethasone of plasma 11-O.H.C.S. levels in severe depressive illness. *Br Med J* 3: 285–287.
2. Schatzberg AF, Rothschild AJ, Stahl JB, Bond TC, Rosenbaum AH, Lofgren SB, *et al.* (1983): The dexamethasone suppression test: identification of subtypes of depression. *Am J Psychiatry* 140: 88–91.
3. Rush AJ, Weissenburger J, Vasavada N, Orsulak PJ, Fairchild CJ (1988): Dexamethasone suppression test status does not predict differential response to nortriptyline versus amitriptyline. *J Clin Psychopharmacol* 8: 421–425.
4. Adams RA, Pinotsis D, Tsirlis K, Unruh L, Mahajan A, Horas AM, *et al.* (2022): Computational Modeling of Electroencephalography and Functional Magnetic Resonance Imaging Paradigms Indicates a Consistent Loss of Pyramidal Cell Synaptic Gain in Schizophrenia. *Biological Psychiatry* 91: 202–215.
5. Huys QJM, Browning M (in press): A Computational View on the Nature of Reward and Value in Anhedonia. *Anhedonia: Preclinical, Translational, and Clinical Integration*.
6. Kunisato Y, Okamoto Y, Ueda K, Onoda K, Okada G, Yoshimura S, *et al.* (2012): Effects of depression on reward-based decision making and variability of action in probabilistic learning. *J Behav Ther Exp Psychiatry* 43: 1088–1094.
7. Huys QJM, Pizzagalli DA, Bogdan R, Dayan P (2013): Mapping anhedonia onto reinforcement learning: a behavioural meta-analysis. *Biol Mood Anxiety Disord* 3: 12.
8. Rutledge RB, Moutoussis M, Smittenaar P, Zeidman P, Taylor T, Hrynkiewicz L, *et al.* (2017): Association of Neural and Emotional Impacts of Reward Prediction Errors With Major Depression. *JAMA Psychiatry* 74: 790–797.
9. Knutson B, Bhanji JP, Cooney RE, Atlas LY, Gotlib IH (2008): Neural responses to monetary incentives in major depression. *Biological Psychiatry* 63: 686–692.