Ten simple rules for the computational modeling of behavioral data

Robert C Wilson[1,2†]*, Anne GE Collins[3,4†]*

[1]Department of Psychology, University of Arizona, Tucson, United States; [2]Cognitive Science Program, University of Arizona, Tucson, United States; [3]Department of Psychology, University of California, Berkeley, Berkeley, United States; [4]Helen Wills Neuroscience Institute, University of California, Berkeley, Berkeley, United States

Abstract Computational modeling of behavior has revolutionized psychology and neuroscience. By fitting models to experimental data we can probe the algorithms underlying behavior, find neural correlates of computational variables and better understand the effects of drugs, illness and interventions. But with great power comes great responsibility. Here, we offer ten simple rules to ensure that computational modeling is used with care and yields meaningful insights. In particular, we present a beginner-friendly, pragmatic and details-oriented introduction on how to relate models to data. What, exactly, can a model tell us about the mind? To answer this, we apply our rules to the simplest modeling techniques most accessible to beginning modelers and illustrate them with examples and code available online. However, most rules apply to more advanced techniques. Our hope is that by following our guidelines, researchers will avoid many pitfalls and unleash the power of computational modeling on their own data.

*For correspondence:
bob@email.arizona.edu (RCW); annecollins@berkeley.edu (AGEC)

†These authors contributed equally to this work

What is computational modeling of behavioral data?

The goal of computational modeling in behavioral science is to use precise mathematical models to make better sense of behavioral data. The behavioral data most often come in the form of choices, but can also be reaction times, eye movements, or other easily observable behaviors, and even neural data. The models come in the form of mathematical equations that link the experimentally observable variables (e.g. stimuli, outcomes, past experiences) to behavior in the immediate future. In this sense, computational models instantiate different 'algorithmic hypotheses' about how behavior is generated.

Exactly what it means to 'make sense' of behavioral data is, to some extent, a matter of taste that will vary according to the researcher's goals (*Kording et al., 2018*). In some cases, a simple model that can explain broad qualitative features of the data is enough. In other cases, more detailed models that make quantitative predictions are required (*Breiman, 2001*). The exact form of the models, and exactly what we do with them, is limited only by our imaginations, but four uses dominate the literature: simulation, parameter estimation, model comparison, and latent variable inference.

Simulation involves running the model with particular parameter settings to generate 'fake' behavioral data. These simulated data can then be analyzed in much the same way as one would analyze real data, to make precise, falsifiable predictions about qualitative and quantitative patterns in the data. Simulation is a way to make theoretical predictions more precise and testable. (Some examples include *Cohen et al., 1990*; *Collins and Frank, 2014*; *Rescorla and Wagner, 1972*; *Farashahi et al., 2017*; *Montague et al., 1996*; *Abbott et al., 2015*; *Lee and Webb, 2005*).

Parameter estimation involves finding the set of parameter values that best account for real behavioral data for a given model. These parameters can be used as a succinct summary of a

given data set (*Ratcliff, 1978*; *Wilson et al., 2013*; *Daw et al., 2011*; *Donkin et al., 2016*), for investigating individual differences (*Frank et al., 2007*; *Starns and Ratcliff, 2010*; *Collins and Frank, 2012*; *Gillan et al., 2016*; *Somerville et al., 2017*; *Nilsson et al., 2011*) and for quantifying the effects of interventions such as drugs, lesions, illness, or experimental conditions (*Frank et al., 2004*; *Lorains et al., 2014*; *Dowd et al., 2016*; *Zajkowski et al., 2017*; *Warren et al., 2017*; *Wimmer et al., 2018*; *van Ravenzwaaij et al., 2011*).

Model comparison involves trying to compute which of a set of possible models best describes the behavioral data, as a way to understand which mechanisms are more likely to underlie behavior. This is especially useful when the different models make similar qualitative predictions but differ quantitatively (*Wilson and Niv, 2011*; *Daw et al., 2011*; *Collins and Frank, 2012*; *Collins and Frank, 2012*; *Fischer and Ullsperger, 2013*; *Steyvers et al., 2009*; *Haaf and Rouder, 2017*; *Donkin et al., 2014*).

Latent variable inference involves using the model to compute the values of hidden variables (for example values of different choices) that are not immediately observable in the behavioral data, but which the theory assumes are important for the computations occurring in the brain. Latent variable inference is especially useful in neuroimaging where it is used to help search for the neural correlates of the model (*O'Doherty et al., 2007*; *Wilson and Niv, 2015*; *Donoso et al., 2014*; *Cohen et al., 2017*), but also for electroencephalogram (EEG), electrocorticography (ECOG), electrophysiology and pupillometry among many other data sources (*O'Reilly et al., 2013*; *Collins and Frank, 2018*; *Samejima et al., 2005*; *Cavanagh et al., 2014*; *Nassar et al., 2012*).

Each of these uses has its strengths and weaknesses, and each of them can be mishandled in a number of ways, causing us to draw wrong and misleading conclusions (*Nassar and Frank, 2016*; *Palminteri et al., 2017*). Here we present a beginner-friendly, pragmatic, practical and details-oriented introduction (complete with example code available at [code]) on how to relate models to data and how to avoid many potential modeling mistakes. Our goal for this paper is to go beyond the mere mechanics of implementing models — as important as those mechanics are — and instead focus on the harder question of how to figure out what, exactly, a model is telling us about the mind. For this reason, we focus primarily on the simplest modeling techniques most accessible to beginning modelers, but almost all of our points apply more generally and readers interested in more advanced modeling techniques should consult the many excellent tutorials, didactic examples, and books on the topic (*Busemeyer and Diederich, 2010*; *Daw, 2011*; *Daw and Tobler, 2014*; *Heathcote et al., 2015*; *Huys, 2017*; *Turner et al., 2013*; *Vandekerckhove et al., 2015*; *Wagenmakers and Farrell, 2004*; *Rigoux et al., 2014*; *Nilsson et al., 2011*; *Farrell and Lewandowsky, 2018*; *Lee et al., 2019*).

For clarity of exposition, we chose to make all of the examples in this paper reflect a single narrow domain - reinforcement learning models applied to choice data (*Sutton and Barto, 2018*). We chose this domain for a few reasons. (1) Modeling is particularly popular in the field of learning. Indeed, this field benefits from modeling particularly because of the nature of the behavioral data: trials are dependent on all past history and thus unique, making classic data analysis with aggregation across conditions less successful. (2) The sequential dependency of trials in learning contexts can lead to technical challenges when fitting models that are absent in non-learning contexts. However, the same techniques are widely and successfully applied to other observable behavior, such as reaction times (*Ratcliff and Rouder, 1998*; *Viejo et al., 2015*; *Ballard and McClure, 2019*; *Wiecki et al., 2013*), and to other domains, including but not limited to perception (*Sims, 2018*), perceptual decision-making (*Ratcliff and Rouder, 1998*; *Drugowitsch et al., 2016*; *Findling et al., 2018*), economic decision-making (*van Ravenzwaaij et al., 2011*; *Nilsson et al., 2011*), visual short-term memory (*Donkin et al., 2016*; *Donkin et al., 2014*; *Nassar et al., 2018*), long-term memory (*Batchelder and Riefer, 1990*), category learning (*Lee and Webb, 2005*), executive functions (*Haaf and Rouder, 2017*; *Jahfari et al., 2019*), and so on. Thus, our hope is that, regardless of the techniques you use or the domain you model, by following these 10 simple steps (*Figure 1*), you will be able to minimize your modeling mishaps and unleash the power of computational modeling on your own behavioral data!
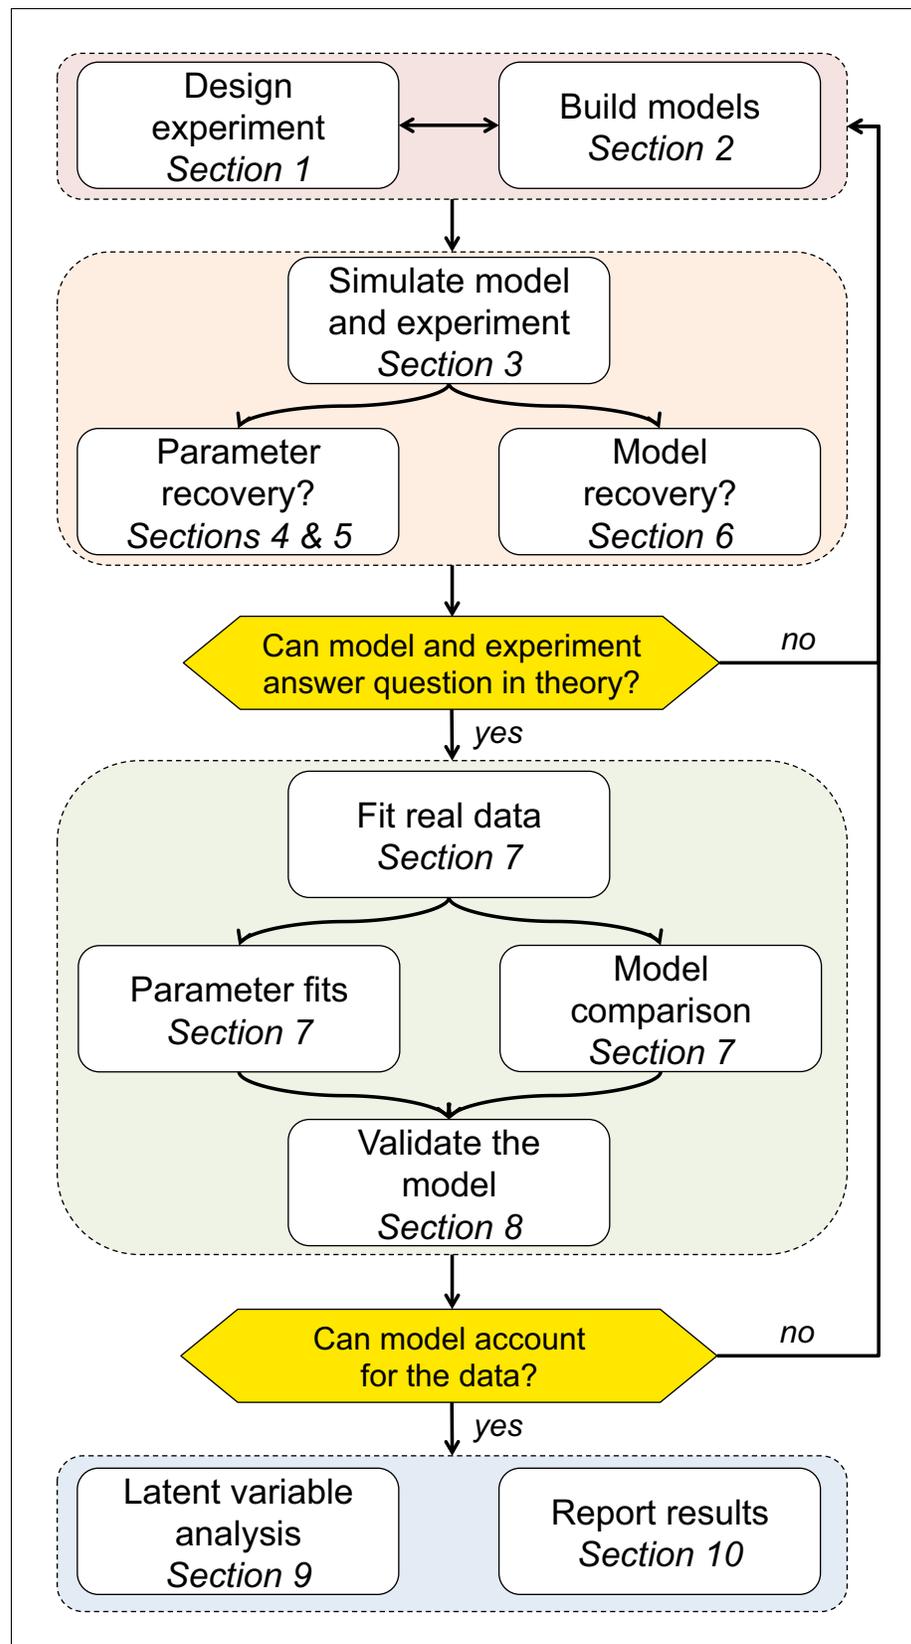
Figure 1. Schematic of the 10 rules and how they translate into a process for using computational modeling to better understand behavior.

Design a good experiment!

Computational modeling is a powerful technique, but it can never replace good experimental design. Modeling attempts to capture how information is manipulated behind the scenes to produce the behavior; thus it is fundamentally limited by the behavioral data, which is itself fundamentally limited by the experimental protocol. A researcher studying face perception would not attempt to fit Prospect Theory to a face perception task; and a researcher studying the differential effects of gain and loss would not do it in a gambling task with only gains. Although obvious in these simple cases, the question becomes more difficult as the complexity of the model increases: is a given learning protocol rich enough to allow the identification of dynamic changes in learning rate, of working memory or episodic memory contributions to learning, or of reward range adaptation? Often, the answer to these questions will be 'no' unless the protocol has been deliberately designed to provide this power.

So, how should you go about designing a good experiment with computational modeling in mind? While this process will always be something of an art form, we suggest that you ask yourself the following questions in order to optimize your experimental design:

What scientific question are you asking?

Although this sounds obvious, it is easy to get sucked into an experimental design without ever asking the most basic questions about your goals. What cognitive process are you targeting? What aspect of behavior are you trying to capture? What hypotheses are you trying to pick apart? For example, you may be trying to identify how working memory contributes to learning or how behavioral variability can be used to explore. Keeping your scientific goals in mind when you design the task can save much time later on.

Does your experiment engage the targeted processes?

This may be a difficult question to answer, and it may require expert knowledge or piloting. However, you need to know that the experimental design actually engages the processes that you are trying to model.

Will signatures of the targeted processes be evident from the simple statistics of the data?

In addition to engaging the processes of interest, the best experiments make these processes identifiable in classical analyses of the behavioral data (*Palminteri et al., 2017*). For example, if you are investigating working memory contributions to learning, you may look for a signature of load on behavior by constructing an experimental design that varies load, to increase chances of probing working memory's role in learning. Seeing signs of the computations of interest in simple analyses of behavior builds confidence that the modeling process will actually work. In our experience, computational modeling is rarely informative when there is no evidence of an effect in model-independent analyses of behavior.

To answer these questions, it is important to have a clear theoretical hypothesis of what phenomenon is to be modeled. In fact, although designing a good experiment is the first step, it goes hand-in-hand with designing a good model, and the two steps should ideally be done in parallel.

But what if I'm not an experimentalist?

Computational modeling is hard and many of the best modelers are specialists who never run experiments of their own. Instead these researchers test their models against published findings, publicly available datasets, or even, if they are lucky, unpublished data from their experimental colleagues. Such specialist modelers might feel that they can safely ignore this first point about experimental design and instead focus on explaining the data they can get. We strongly urge them not to. Instead we urge these specialist modelers to always be considering better ways in which their models could be tested. Such experimental thinking helps you to be more concrete in your ideas and to think about how your model might apply outside of the context for which it was designed. In addition, thinking experimentally — and even better talking with experimentalists — forces you to engage with behavior as it actually is rather than as you would like it to be, which in turn can lead to new insights. Finally, by proposing concrete experimental designs, it is easier to convince your

experimental colleagues to actually test your ideas, which is surely the goal if we are to move the field forward.

An illustrative example: the multi-armed bandit task

The ten rules in this paper are quite general, but we will illustrate many of our points using simple examples from our own field of reinforcement learning. Code for implementing all of these examples is available on GitHub (https://github.com/AnneCollins/TenSimpleRulesModeling) (*Collins and Wilson, 2019*; copy archived at https://github.com/elifesciences-publications/TenSimpleRulesModeling). The goal of these example studies is to understand how people learn to maximize their rewards in a case where the most rewarding choice is initially unknown.

More specifically, we consider the case in which a participant makes a series of T choices between K slot machines, or 'one-armed bandits', to try to maximize their earnings. If played on trial t, each slot machine, k, pays out a reward, r_t, which is one with reward probability, μ_t^k, and otherwise 0. The reward probabilities are different for each slot machine and are initially unknown to the subject. In the simplest version of the task, the reward probabilities are fixed over time.

The three **experimental parameters** of this task are: the number of trials, T, the number of slot machines, K, and the reward probabilities of the different options, μ_t^k, which may or may not change over time. The settings of these parameters will be important for determining exactly what information we can extract from the experiment. In this example, we will assume that $T = 1000$, $K = 2$, and that the reward probabilities are $\mu_t^1 = 0.2$ for slot machine 1 and $\mu_t^2 = 0.8$ for slot machine 2.

Design good models

Just as bad experiments can limit our ability to test different hypotheses, bad models – quite literally the mathematical embodiment of our hypotheses – can further limit the conclusions we can draw (*Donkin et al., 2014*). This point is especially important if we are designing new models, but even well-established computational models can be problematic in some cases (*Broomell and Bhatia, 2014*; *Nilsson et al., 2011*).

Critical to the design of the model is a clear understanding of your reason for modeling. Are you interested in a descriptive model that succinctly summarizes, but perhaps does not explain, behavioral data? A mechanistic model to tie behavior to the brain? Or an elegant mathematical model to illustrate a concept? As shown in an excellent article by Kording and colleagues (*Kording et al., 2018*), computational modelers have a wide variety of goals for their models, and understanding your own motivations is a great place to start.

More pragmatically, there are a number of different approaches for designing models that have been successfully used in the literature. Perhaps the simplest approach is to use heuristics to find a 'reasonable' way to handle information to produce the target behavior. This approach was how the delta rule (see Model three below) was first invented (*Rescorla and Wagner, 1972*). Another approach is to scour the artificial intelligence, computer science, and applied mathematics literature for algorithms that have been used to solve similar problems for artificial agents. This approach has been fruitfully applied in the field of reinforcement learning (*Sutton and Barto, 2018*), where algorithms such as Q-learning and temporal difference learning have been related to human and animal behavior and brain function (*Watkins and Dayan, 1992*; *Montague et al., 1996*). Another approach is to take a Bayes-optimal perspective, to design algorithms that perform optimally given a model of the environment and the task. Ideal observer models in vision are one example in which this approach has been applied successfully (*Geisler, 2011*). More generally, Bayes-optimal models can be further pursued by investigating simpler algorithms that approximate the ideal strategy, or by imposing bounded rationality constraints, such as limited computational resources, on ideal observer agents (*Courville and Daw, 2008*; *Nassar et al., 2010*; *Collins and Frank, 2012*; *Daw and Courville, 2007*; *Lieder et al., 2018*).

Regardless of the approach (or, better yet, approaches) that you take to design your models, it is important to keep the following points in mind:

A computational model should be as simple as possible, but no simpler

Einstein's old edict applies equally to models of the mind as it does to models of physical systems. Simpler, more parsimonious models are easier to fit and easier to interpret and should always be included in the set of models under consideration. Indeed, formal model comparison techniques (described in detail in Appendix 2) include a penalty for overly complex models, which are more likely to overfit the data and generalize poorly, and favor simpler models so long as they can account for the data.

A computational model should be interpretable (as much as possible)

In the process of developing models that can account for the behavioral data, researchers run the risk of adding components to a model that are not interpretable as a sensible manipulation of information. For example, a negative learning rate is difficult to interpret in the framework of reinforcement learning. Although such uninterpretable models may sometimes improve fits, nonsensical parameter values may indicate that something important is missing from your model, or that a different cognitive process altogether is at play.

The models should capture *all* the hypotheses that you plan to test

While it is obviously important to design models that can capture your main hypothesis, it is even more important to design models that capture competing hypotheses. Crucially, *competing models should not be strawmen* — they should have a genuine chance of relating to behavior in the task environment, and they should embody a number of reasonable, graded hypotheses. You should of course put equal effort into fitting these models as you do your favored hypothesis. Better yet, you shouldn't have a favored hypothesis at all — let the data determine which model is the best fit, not your a priori commitment to one model or another.

Box 1. Example: Modeling behavior in the multi-armed bandit task.

We consider five different models of how participants could behave in the multi-armed bandit task.

Model 1: Random responding
In the first model, we assume that participants do not engage with the task at all and simply press buttons at random, perhaps with a bias for one option over the other. Such random behavior is not uncommon in behavioral experiments, especially when participants have no external incentives for performing well. Modeling such behavior can be important if we wish to identify such 'checked out' individuals in a quantitative and reproducible manner, either for exclusion or to study the checked-out behavior itself. To model this behavior, we assume that participants choose between the two options randomly, perhaps with some overall bias for one option over the other. This bias is captured with a parameter b (which is between 0 and 1), such that the probability of choosing the two options is

$$p_t^1 = b \quad \text{and} \quad p_t^2 = 1 - b \tag{1}$$

Thus, for two bandits, the random responding model has just one free parameter, controlling the overall bias for option 1 over option 2, $\theta_1 = b$.

Model 2: Noisy win-stay-lose-shift
The win-stay-lose-shift model is one of the simplest models that adapts its behavior according to feedback. Consistent with the name, the model repeats rewarded actions and switches away from unrewarded actions. In the noisy version of the model, the win-stay-lose-shift rule is applied probabilistically, such that the model applies the win-stay-lose-shift rule with probability $1 - \epsilon$, and chooses randomly with probability ϵ. In the two-bandit case, the probability of choosing option k is

$$p_t^k = \begin{cases} 1 - \epsilon/2 & \text{if } (c_{t-1} = k \text{ and } r_{t-1} = 1) \text{ OR } (c_{t-1} \neq k \text{ and } r_{t-1} = 0) \\ \epsilon/2 & \text{if } (c_{t-1} \neq k \text{ and } r_{t-1} = 1) \text{ OR } (c_{t-1} = k \text{ and } r_{t-1} = 0) \end{cases} \tag{2}$$

where $c_t = 1, 2$ is the choice at trial t, and $r_t = 0, 1$ the reward at trial t. Although more complex to implement, this model still only has one free parameter, the overall level of randomness, $\theta_2 = \epsilon$.

Model 3: Rescorla Wagner

In this model, participants first *learn* the expected value of each slot machine based on the history of previous outcomes and then use these values to make a *decision* about what to do next. A simple model of learning is the Rescorla-Wagner learning rule (***Rescorla and Wagner, 1972***), whereby the value of option k, Q_t^k is updated in response to reward r_t according to:

$$Q_{t+1}^k = Q_t^k + \alpha(r_t - Q_t^k) \tag{3}$$

where α is the learning rate, which takes a value between 0 and 1 and captures the extent to which the prediction error, $(r_t - Q_t^k)$, updates the value. For simplicity, we assume that the initial value, Q_0^k, is zero, although it is possible to treat the Q_0^k as a free parameter of the model.

A simple model of decision making is to assume that participants use the options' values to guide their decisions, choosing the most valuable option most frequently, but occasionally making 'mistakes' (or exploring) by choosing a low-value option. One choice rule with these properties is known as the 'softmax' choice rule, which chooses option k with probability

$$p_t^k = \frac{\exp(\beta Q_t^k)}{\sum_{i=1}^{K} \exp(\beta Q_t^i)} \tag{4}$$

where β is the 'inverse temperature' parameter that controls the level of stochasticity in the choice, ranging from $\beta = 0$ for completely random responding and $\beta = \infty$ for deterministically choosing the highest value option.

Combining the learning (***Equation 3***) and decision rules (***Equation 4***) gives a simple model of decision-making in this task with two free parameters: the learning rate, α, and the inverse temperature, β. That is, in our general notation, for this model $\theta_3 = (\alpha, \beta)$.

Model 4: Choice kernel

This model tries to capture the tendency for people to repeat their previous actions. In particular, we assume that participants compute a 'choice kernel,' CK_t^k, for each action, which keeps track of how frequently they have chosen that option in the recent past. This choice kernel updates in much the same way as the values in the Rescorla-Wagner rule, i.e. according to

$$CK_{t+1}^k = CK_t^k + \alpha_c(a_t^k - CK_t^k) \tag{5}$$

where $a_t^k = 1$ if option k is played on trial t, otherwise $a_t^k = 0$, and α_c is the choice-kernel learning rate. For simplicity, we assume that the initial value of the choice kernel is always zero, although, like the initial Q-value in the Rescorla-Wagner model, this could be a parameter of the model. Note that with $\alpha_c = 1$, this model is very similar to model 2 (win-stay-lose-shift). From there, we assume that each option is chosen according to

$$p_t^k = \frac{\exp(\beta_c CK_t^k)}{\sum_{i=1}^{K} \exp(\beta_c CK_t^i)} \tag{6}$$

where β_c is the inverse temperature associated with the choice kernel.

Combining the choice kernel (***Equation 5***) with the decision rule (***Equation 6***) gives a simple model of decision-making in this task with two free parameters: the choice-kernel learning rate, α_c, and the choice-kernel inverse temperature β_c. That is, in our general notation, for this model $\theta_4 = (\alpha_c, \beta_c)$.

Model 5: Rescorla Wagner + choice kernel

Finally, our most complex model mixes the reinforcement learning model with the choice kernel model. In this model, the values update according to *Equation 3*, while the choice kernel updates according to *Equation 5*. The terms are then combined to compute the choice probabilities as

$$p_t^k = \frac{\exp(\beta Q_t^k + \beta_c CK_t^k)}{\sum_{i=1}^K \exp(\beta Q_t^i + \beta_c CK_t^i)} \tag{7}$$

This most complex model has four free parameters, i.e. $\theta_5 = (\alpha, \beta, \alpha_c, \beta_c)$.

Simulate, simulate, simulate!

Once you have an experimental design and a set of computational models, a really important step is to create *fake*, or *surrogate* data (*Palminteri et al., 2017*). That is, you should use the models to simulate the behavior of participants in the experiment, and to observe how behavior changes with different models, different model parameters, and different variants of the experiment. This step will allow you to refine the first two steps: confirming that the experimental design elicits the behaviors assumed to be captured by the computational model. To do this, here are some important steps.

Define model-independent measures that capture key aspects of the processes you are trying to model

Finding qualitative signatures (and there will often be more than one) of the model is crucial. By studying these measures with simulated data, you will have greater intuition about what is going on when you use the same model-independent measures to analyze real behavior (*Daw et al., 2011*; *Collins and Frank, 2012*; *Collins and Frank, 2013*; *Nassar et al., 2018*; *Lee and Webb, 2005*).

Simulate the model across the range of parameter values

Then, visualize behavior as a function of the parameters. Almost all models have free parameters. Understanding how changes to these parameters affect behavior will help you to better interpret your data and to understand individual differences in fit parameters. For example, in probabilistic reinforcement learning tasks modeled with a simple delta-rule model (Model 3; *Equation 3*), the learning rate parameter, α, can relate to both the speed of learning and noisiness in asymptotic behavior, as can the inverse temperature parameter, β (in *Equation 4*), as seen in *Box 2—figure 1B*.

Visualize the simulated behavior of different models

This will allow you to verify that behavior is qualitatively different for different models, making their predictions in the experimental setup different (*Box 2—figure 1A*). If the behavior of different models is *not* qualitatively different, this is a sign that you should try to design a better experiment. Although not always possible, distinguishing between models on the basis of qualitative patterns in the data is always preferable to quantitative model comparison (*Navarro, 2019*; *Palminteri et al., 2017*).

More generally, the goal of the simulation process is to clarify how the models and experimental design satisfy your goal of identifying a cognitive process in behavior. If the answer is positive — i.e. the experiment is rich enough to capture the expected behavior, the model's parameters are interpretable, and competing models make dissociable predictions — you can move on to the next step. Otherwise, you should loop back through these first three sections to make sure that your experimental design and models work well together, and that the model parameters have identifiable effects on the behavior, which is a prerequisite for the fourth step, fitting the parameters (c.f. *Figure 1*).

Box 2. Example: simulating behavior in the bandit task.

To simulate behavior, we first need to define the parameters of the task. These include the total number of trials, T (=1000 in the example), as well as the number of bandits, $K(=2)$, and the reward probability for each bandit, μ^k (0.2 and 0.8 for bandits 1 and 2, respectively). The

experiment parameters, as used in the simulation, should match the actual parameters used in the experiment.

Next we define the parameters of the model. One way to do this is to sample these parameters randomly from prior distributions over each parameter, the exact form of which will vary from model to model. These prior distributions should generally be as broad as possible, but if something is known about the distribution of possible parameter values for a particular model, this is one place to include it.

With the free parameters set, we then proceed with the simulation. First, we simulate the choice on the first trial, a_1, by assuming that the model chooses option k with probability, p_1^k. Next we simulate the outcome, r_1, of this choice. In Models 2–5, we use the action and/or outcome to update the choice probabilities for the next trial. Repeating this process for all trials up to $t = T$ completes one simulation. The simulations can then be analyzed in the same way as participants' data is, ideally with the same code taking different inputs. This process should be repeated several times, with different parameter settings, to get a handle on how the model behaves as a function of its parameters.

To illustrate how one might visualize the simulated results, we look at two model-independent measures that should capture fundamental aspects of learning: the probability of repeating an action, $p(\mathrm{stay})$ (should I change my behavior in response to feedback?), and the probability of choosing the correct option, $p(\mathrm{correct})$ (have I learned?). In **Box 2—figure 1A** below, we plot $p(\mathrm{stay})$ as a function of the reward on the last trial for each of the models with a particular set of parameters (M1: $b = 0.5$, M2: $\epsilon = 0.05$, M3: $\alpha = 0.1$, $\beta = 5$, M4: $\alpha_c = 0.1$, $\beta_c = 3$, M5: $\alpha = 0.1$, $\beta = 5$, $\alpha_c = 0.1$, $\beta_c = 1$). For some models (in particular the win-stay-lose-shift model (Model 2), we expect a strong dependence on past reward, but for others, such as the random responder (Model 1), we expect no dependence. Of course, the exact behavior of each model depends crucially on the parameters used in the simulations and care should be taken to ensure that these simulation parameters are reasonable, perhaps by matching to typical parameter values used in the literature or by constraining to human-like overall performance. Better yet is to simulate behavior across a range of parameter settings to determine how the model-independent measures change with different parameters.

A more thorough exploration of the parameter space for Model 3 is shown in **Box 2—figure 1B**, where we plot the $p(\mathrm{correct})$ in the first and last 10 trials as a function of the learning rate, α, and softmax parameter, β. Note that the 'optimal' learning rate, i.e. the value of α that maximizes $p(\mathrm{correct})$, varies between early and late trials and as a function of the softmax parameter β, where for early trials higher β implies a lower optimal α (**Daw et al., 2011**).

The question of how to choose the model-independent measures of behavior has no easy answer and calls to the domain knowledge of the specific scientific question that the modeler is attempting to answer. As a rule of thumb, the measures should capture global characteristics (e.g. overall performance) and diagnostic measures that relate to the question of interest, and may visualize different qualitative predictions of different models.
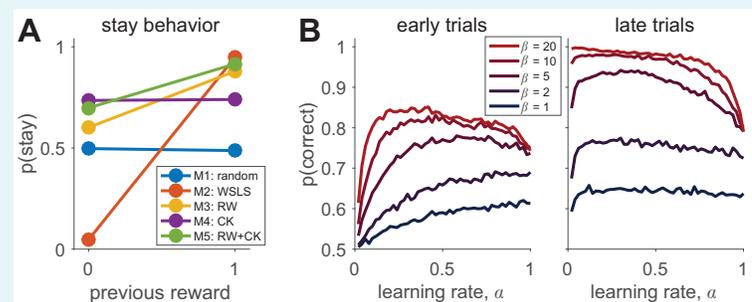
Box 2—figure 1. Simulating behavior in the two-armed bandit task. (**A**) Win-stay-lose-shift behavior varies widely between models. (**B**) Model 3 simulations (100 per parameter setting) show how the learning rate and softmax parameters influence two aspects of behavior: early performance (first 10 trials), and late performance (last 10 trials). The left graph shows that learning rate is positively correlated with early performance

improvement only for low β values or for very low α values. For high β values, there is a U-shape relationship between learning rate and early speed of learning. The right graph shows that with high β values, high learning rates negatively influence asymptotic behavior. Thus, both parameters interact to influence both the speed of learning and asymptotic performance.

Fit the parameters

A key component of computational modeling is estimating the values of the parameters that best describe your behavioral data. There are a number of different ways of estimating parameters, but here we focus on the maximum-likelihood approach, although almost all of our points apply to other methods such as Markov Chain Monte Carlo approaches (*Lee and Wagenmakers, 2014*). Mathematical details, as well as additional discussion of other approaches to model fitting can be found in Appendix 1.

In the maximum likelihood approach to model fitting, our goal is to find the parameter values of model m, $\hat{\theta}_m^{MLE}$, that maximize the likelihood of the data, $d_{1:T}$, given the parameters, $p(d_{1:T}|\theta_m, m)$. Maximizing the likelihood is equivalent to maximizing the log of the likelihood, $LL = \log p(d_{1:T}|\theta_m, m)$, which is numerically more tractable. (The likelihood is a product of many numbers smaller than 1, which can be rounded to 0 with limited precision computing. By contrast, the log-likelihood is a sum of negative numbers, which is usually tractable and will not be rounded to 0.) A simple mathematical derivation shows that this log-likelihood can be written in terms of the choice probabilities of the individual model as

$$LL = \log p(d_{1:T}|\theta_m, m) = \sum_{t=1}^{T} \log p(c_t|d_{1:t-1}, s_t, \theta_m, m) \tag{8}$$

where $p(c_t|d_{1:t-1}, s_t, \theta_m, m)$ is the probability of each individual choice given the parameters of the model and the information available up to that choice, which is at the heart of the definition of each model (for example in *Equations 1-7*).

In principle, finding the maximum likelihood parameters is as 'simple' as maximizing LL. In practice, of course, finding the maximum of a function is not a trivial process. The simplest approach, a brute force search of the entire parameter space, is occasionally useful, and may help you to understand how different parameters interact (see *Box 3—figure 1*). However, this approach is unfeasible outside of the simplest cases (e.g. one or two parameters with tight bounds) because of the high computational costs of evaluating the likelihood function at a large number of points.

Fortunately, a number of tools exist for finding local maxima (and minima) of functions quickly using variations on gradient ascent (or descent). For example, Matlab's `fmincon` function can use a variety of sophisticated optimization algorithms (e.g. *Moré and Sorensen, 1983*; *Byrd et al., 2000*) to find the minimum of a function (and other factors such as the Hessian that can be useful in some situations [*Daw, 2011*]). So long as one remembers to feed fmincon the *negative* log-likelihood (whose minimum is at the same parameter values as the maximum of the positive log-likelihood), using tools such as fmincon can greatly speed up model fitting. Even here, though, a number of problems can arise when trying to maximize LL that can be reduced by using the tips and tricks described below. Most of the tips come from understanding that optimization algorithms are not foolproof and in particular are subject to numerical constraints. They generalize to other black box optimization functions in other languages, for example the Python scipy.optimize package or the optim function in R.

Be sure that your initial conditions give finite log-likelihoods

Optimizers such as fmincon require you to specify initial parameter values from which to start the search. Perhaps the simplest way in which the search process can fail is if these initial parameters give log-likelihoods that are not finite numbers (e.g. infinities or NaNs, not a number in Matlab speak). If your fitting procedure fails, this can often be the cause.

Beware rounding errors, zeros and infinities

More generally, the fitting procedure can go wrong if it encounters infinities or NaNs during the parameter search. This can occur if a choice probability is rounded down to zero, thus making the log of the choice probability $-\infty$. Likewise, if your model involves exponentials (e.g. the softmax choice rule in *Equation 4*), this can lead to errors whereby the exponential of a very large number is 'rounded up' to infinity. One way to avoid these issues is by constraining parameter values to always give finite choice probabilities and log-likelihoods at the boundaries. One way to diagnose these issues is to include checks in the code for valid log-likelihoods.

Be careful with constraints on parameters

If the constraints are ill chosen, it is possible that the solution will be at the bounds, which is often, but not always, a red flag.

Only include parameters that have an influence on the likelihood. If only two parameters impact the likelihood, but the optimizer attempts to fit three, it will usually find the optimum for the two relevant parameters and a random value for the third; however, it will lead to slower and less efficient fitting.
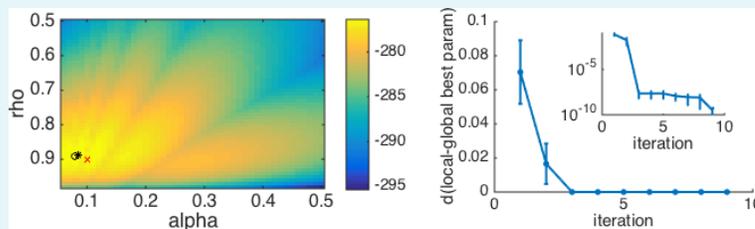
Beware local minima!

Finally, a key limitation of optimization algorithms is that they are only guaranteed to find *local* minima, which are not guaranteed to be the *global* minima corresponding to the best fitting parameters. One way to mitigate this issue is to run the fitting procedure multiple times with random initial conditions, recording the best fitting log-likelihood for each run. The best fitting parameters are then the parameters corresponding to the run with the highest log-likelihood. There is no hard-and-fast rule for knowing how many starting points to use in a given situation, besides the fact that more complex models will require more starting points. Thus, this number must be determined empirically in each case. One way to validate the number of starting points is by plotting the best likelihood score as a function of the number of starting points. As the number of initial conditions increases, the best-fitting likelihood (and corresponding the parameters) will improve up to an asymptote close to the true maximum of the function (e.g. *Box 3—figure 1*).

Box 3. Example: contending with multiple local maxima.

As a real example with local maxima, we consider the case of a simplified version of the mixed reinforcement learning and working memory model from *Collins and Frank, 2012*. For simplicity, we relegate the details of this model to Appendix 4. To appreciate the example, all one really needs to know is that in its simplest version, this model has two parameters: ρ, which captures the effect of working memory, and α, which captures the learning rate of reinforcement learning. As is seen in *Box 3—figure 1* below, this model (combined with an appropriate experiment) gives rise to a log-likelihood surface with multiple local maxima. Depending on the starting point, the optimization procedure can converge to any one of these local maxima, meaning that the 'maximum' likelihood fits may not reflect the global maximum likelihood.

To mitigate this concern, a simple and effective approach is to repeat the optimization procedure many times, keeping track of the best fitting log-likelihood and parameters in each case. An approximation to the global maximum is to take the best log-likelihood from this list of fits. The results of this multiple iteration procedure can be summarized by plotting the best log-likelihood as a function of the number of starting points, or similarly, by plotting the distance from the so-far best parameters to the final best parameters as a function of the number of starting points (*Box 3—figure 1B*). As the number of starting points increases, the best-fitting log-likelihood and parameters will converge to the global maximum. This plot also allows us to judge when we have used enough starting points. Specifically, if the best fitting parameters appear to have reached asymptote, that gives us a good indication that the fit is the best we can do.

Box 3—figure 1. An example with multiple local minima. (**Left**) Log-likelihood surface for a working memory reinforcement learning model with two parameters. In this case, there are several local minima, all of which can be found by the optimization procedure depending on the starting point. Red x, generative parameters; black circle, optimum with brute search method; black *, optimum with fmincon and multiple starting points. (**Right**) Plotting the distance from the best fitting parameters after n iterations to the best fitting parameters after all iterations as a function of the number of starting points n gives a good sense of when the procedure has found the global optimum. The inset shows the same plot on a logarithmic scale for distance, illustrating that there are still very small improvements to be made after the third iteration.

Check that you can recover your parameters

Before reading too much into the best-fitting parameter values, θ_m^{MLE}, it is important to check whether the fitting procedure gives meaningful parameter values in the best case scenario, -that is, when fitting fake data where the 'true' parameter values are known (*Nilsson et al., 2011*). Such a procedure is known as 'Parameter Recovery', and is a crucial part of any model-based analysis.

In principle, the recipe for parameter recovery is quite simple. First, simulate fake data with known parameter values. Next, fit the model to these fake data to try to 'recover' the parameters. Finally, compare the recovered parameters to their true values. In a perfect world, the simulated and recovered parameters will be tightly correlated, with no bias. If there is only a weak correlation between the simulated and recovered parameters and/or a significant bias, then this is an indication that there is either a bug in your code (which from our own experience is fairly likely) or the experiment is underpowered to assess this model.

To make the most of your parameter recovery analysis, we suggest the following tips:

Make sure your simulation parameters are in the right range

An important choice for parameter recovery is the range of simulation parameters that you wish to recover. Some models/experiments only give good parameter recovery for parameters in a particular range — if the simulation parameters are too big or too small, they can be hard to recover. An illustration of this is the softmax parameter, β, where very large β values lead to almost identical behavior in most experiments. Thus parameter recovery may fail for large β values but work well for small β values. Of course, selecting only the range of parameters that *can* be recovered by your model is not necessarily the right choice, especially if the parameter values you obtain when fitting real data are outside of this range! For this reason, we have the following recommendations for choosing simulation parameter values:

1. If you have already fit your data, we recommend matching the range of your simulation parameters to the range of values obtained by your fit.
2. If you have not fit your data but you are using a model that has already been published, match the range of parameters to the range seen in previous studies.
3. Finally, if the model is completely new and the 'true' parameter values are unknown, we recommend simulating over as wide a range as possible to get a sense of whether and where parameters can be recovered. You can rely on your exploration of how model parameters affect simulated behavior to predict a range beyond which parameters will not affect behavior much.

Note that it is not necessarily problematic if a model's parameters are not recoverable in a full parameter space, as long as they are recoverable in the range that matters for real data.

Plot the correlations between simulated and recovered parameters

While the correlation coefficient between simulated and recovered parameters is a useful number for summarizing parameter recovery, we also strongly recommend that you actually plot the simulated vs recovered parameters. This makes the correlation clear, and also reveals whether the correlation holds in some parameter regimes but not others. It also reveals any existing bias (for example, a tendency to recover higher or lower values in average).

Make sure the recovery process does not introduce correlations between parameters

In addition to looking at the correlations between simulated and recovered parameters, we also recommend looking at the correlation between the recovered parameters themselves. If the simulation parameters are uncorrelated with one another, correlation between the recovered parameters is an indication that the parameters in the model are trading off against one another (*Daw, 2011*). Such trade-offs can sometimes be avoided by reparameterizing the model (e.g. *Otto et al., 2013*) or redesigning the experiment. Sometimes, however, such trade-offs are unavoidable. In these cases, it is crucial to report the trade-off in parameters so that a 'correlation' between fit parameter values is not over-interpreted in real data.

A note about parameter differences between different populations or conditions: a growing use of model fitting is to compare parameter values between populations (e.g. schizophrenia patients vs healthy controls [*Collins et al., 2014*]) or conditions (e.g., transcranial magnetic stimulation to one area or another [*Zajkowski et al., 2017*]). If your primary interest is a difference like this, then parameter recovery can be used to give an estimate of statistical power. In particular, for a proposed effect size (e.g., on the average difference in one parameter between groups or conditions) you can simulate and recover parameters for the groups or conditions and then perform statistical tests to detect group differences in this simulated data set. The power for this effect size is then the frequency with which the statistical tests detect no effect given that the effect is there.

Remember that even successful parameter recovery represents a best-case scenario!

What does successful parameter recovery tell you? That data generated by a known model with given parameters can be fit to recover those parameters. This is the best case you could possibly hope for in the model-based analysis and it is unlikely to ever occur as the 'true' generative process for behavior — that is, the inner workings of the mind and brain — is likely much more complex than any model you could conceive. There's no easy answer to this problem. We only advise that you remember to be humble when you present your results!

Box 4. Example: parameter recovery in the reinforcement learning model.

We performed parameter recovery with Model 3, the Rescorla Wagner model, on the two-armed bandit task. As before, we set the means of each bandit at $\mu_1 = 0.2$ and $\mu_2 = 0.8$ and the number of trials at $T = 1000$. We then simulated the actions of the model according to *Equations 3 and 4*, with learning rate, α, and softmax temperature, β, set according to

$$\alpha \sim U(0,1) \quad \text{and} \quad \beta \sim \text{Exp}(10) \tag{9}$$

After simulating the model, we fit the parameters using a maximum likelihood approach to get fit values of learning rate, α, and softmax parameter, β. We then repeated this process 1000 times using new values of α and β each time. The results are plotted in *Box 4—figure 1* below. As is clear from this plot, there is fairly good agreement between the simulated and fit parameter values. In addition, we can see that the fit for β is best with a range, $1 < \beta < 10$, and that outside this range, the correspondence between simulation and fit is not as good. If we further select points where parameter recovery for α is bad (i.e., when $|\alpha_{sim} - \alpha_{fit}| > 0.25$, grey dots in *Box 4—figure 1*), we find that parameter recovery for α is worse when β is outside of the range. Depending on the values of β that we obtain by fitting human behavior, this worse correspondence at small and large β values may or may not be problematic. It may be a good idea

to use the range of parameters obtained from fitting the real data to test the quality of recovery within the range that matters.
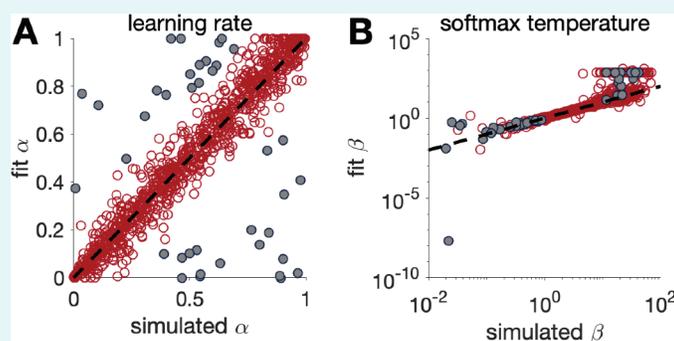
Box 4—figure 1. Parameter recovery for the Rescorla Wagner model (model 3) in the bandit task with 1000 trials. Grey dots in both panels correspond to points where parameter recovery for α is bad.

Can you arbitrate between different models?

In model comparison, our goal is to determine which model, out of a set of possible models, is most likely to have generated the data. There are a number of different ways to make this comparison (summarized in more detail in Appendix 2) that involve different approximations to the Bayesian evidence for each model (e.g., *Daw, 2011*; *Rigoux et al., 2014*). Here, we focus on the most common method which is related to the log-likelihood computed in 'Fit the parameters'.

A simplistic approach to model comparison would be to compare the log-likelihoods of each model at the best fitting parameter settings, $p(d_{1:T}|\hat{\theta}_m, m)$. However, if the data, $d_{1:T}$, used to evaluate the log-likelihood are the same as those used to fit the parameters, then this approach will lead to overfitting, as the model with the most free parameters will almost always fit this 'training' data best. As an extreme example, consider the case of a model with one 'parameter' per choice, which is the identity of the choice the person actually made. Such a 'model' would fit the data perfectly, but would of course tell us nothing about how the choices were actually determined and would make no predictions about what choices would be made in a different setting. Overfitting is a problem in that it decreases the generalizability of the model: it makes it less likely that the conclusions drawn would apply to a different sample.

One way to avoid overfitting is to perform cross-validation: by measuring fit on held-out data, we directly test generalizability. However, this is not always possible for practical reasons (number of samples) or more fundamental ones (dependence between data points). Thus, other methods mitigate the risk of overfitting by approximately accounting for the degrees of freedom in the model. There are several methods for doing this (including penalties for free parameters), which are discussed in more detail in the Appendices. There is a rich theoretical literature debating which method is best (*Wagenmakers and Farrell, 2004*; *Vandekerckhove et al., 2015*). Here, we do not position ourselves in this theoretical debate, and instead focus on one of the simplest methods, the Bayes Information Criterion, *BIC*, which has an explicit penalty for free parameters.

$$BIC = -2\log\hat{L}L + k_m\log(T) \tag{10}$$

where $\hat{L}L$ is the log-likelihood value at the best fitting parameter settings, and k_m is the number of parameters in model m. The model with the smallest *BIC* score is the model that best fits the data. Thus, the positive effect of k_m in the last term corresponds to a penalty for models with large numbers of parameters.

While *Equation 10* is simple enough to apply in order to find the model that, apparently, best fits your data, it is important to check that your model comparison process gives sensible results for

simulated data. Just as parameter fitting should be validated by parameter recovery on simulated data, so model comparison should be validated by model recovery on simulated data.

More specifically, model recovery involves simulating data from all models (with a range of parameter values carefully selected as in the case of parameter recovery) and then fitting that data with all models to determine the extent to which fake data generated from model A is best fit by model A as opposed to model B. This process can be summarized in a confusion matrix (see *Box 5—figure 1* below for an example) that quantifies the probability that each model is the best fit to data generated from the other models, that is, $p(\text{fit model} = B|\text{simulated model} = A)$. In a perfect world, the confusion matrix will be the identity matrix, but in practice, this is not always the case (e. g., *Wilson and Niv, 2011*).

When computing and interpreting a confusion matrix it is important to keep the following points in mind:

Compare different methods of model comparison

If the confusion matrix has large off-diagonal components, then you have a problem with model recovery. There are a number of factors that could cause this problem, ranging from a bug in the code to an underpowered experimental design. However, one cause that is worth investigating is whether you are using the wrong method for penalizing free parameters. In particular, different measures penalize parameters in different ways that are 'correct' under different assumptions. If your confusion matrix is not diagonal, it may be that the assumptions underlying your measures (e.g. BIC) do not hold for your models, in which case it might be worth trying another metric for model comparison (e.g., AIC [*Wagenmakers and Farrell, 2004*]; see Appendix 2).

Be careful with the choice of parameters when computing the confusion matrix

Just as parameter recovery may only be successful in certain parameter regimes, so too can model recovery depend critically on the parameters chosen to simulate the models. In some parameter regimes, two models may lead to very different behavior, but they may be indistinguishable in other parameter regimes (see *Box 5—figure 1* below). As with parameter recovery, we believe that the best approach is to match the range of the parameters to the range seen in your data, or to the range that you expect from prior work.

A note on interpreting the confusion matrix

As described above, and in keeping with standard practice from statistics, the confusion matrix is defined as the probability that data simulated by one model is best fit by another, that is, $p(\text{fit model}|\text{simulated model})$. However, when we fit a model to real data, we are usually more interested in making the reverse inference — that is, given that model B fits our data best, which model is most likely to have generated the data? This is equivalent to computing $p(\text{simulated model}|\text{fit model})$. Note that this measure, which we term the 'inversion matrix' to distinguish it from the confusion matrix, is not the same as the confusion matrix unless model recovery is perfect. Of course, the inversion matrix can be computed from the confusion matrix using Bayes rule (see Appendix 3) and it may be useful to report it in cases where the confusion matrix is not diagonal.

The elephant in the room with model comparison

As wonderful as it is to find that your model 'best' fits the behavioral data, the elephant in the room (or perhaps more correctly *not* in the room) with all model comparison is that it only tells you which of the models you considered fits the data best. In and of itself, this is rather limited information as there are infinitely many other models that you did not consider. This makes it imperative to start with a good set of models that rigorously capture the competing hypotheses (that is, think hard in Step 2). In addition, it will be essential to validate (at least) your winning model (see Step 9) to show how simulating its behavior can generate the patterns seen in the data that you did not explicitly fit, and thus obtain an *absolute* measure of how well your model relates to your data.

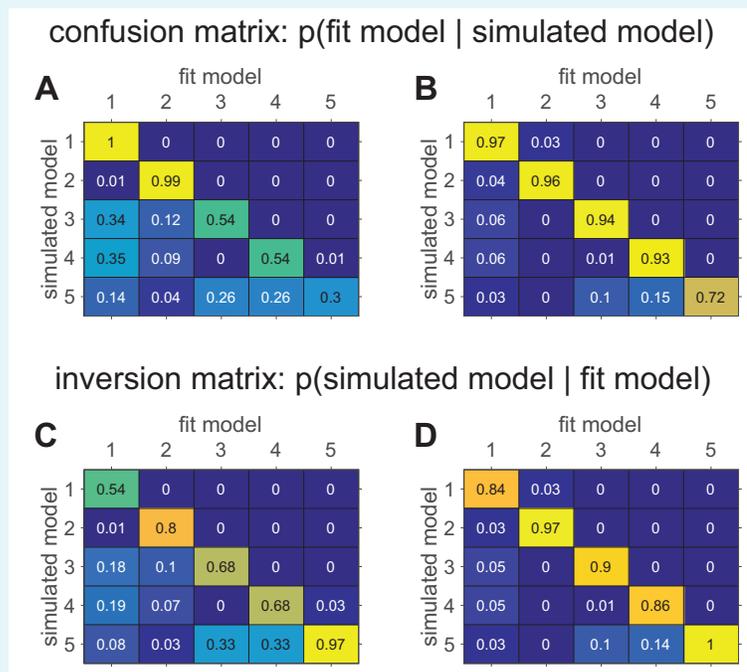## Box 5. Example: confusion matrices in the bandit task.

To illustrate model recovery, we simulated the behavior of the five models on the two-armed bandit task. As before, the means were set at $\mu_1 = 0.2$ and $\mu_2 = 0.8$, and the number of trials was set at $T = 1000$. For each simulation, model parameters were sampled randomly for each model. Each simulated data set was then fit to each of the given models to determine which model fit best (according to BIC). This process was repeated 100 times to compute the confusion matrices which are plotted below in **Box 5—figure 1A and B**.

The difference between these two confusion matrices is in the priors from which the simulation parameters were sampled. In panel A, parameters were sampled from the following priors:

Model	Priors
Model 1	$b \sim U(0,1)$
Model 2	$\epsilon \sim U(0,1)$
Model 3	$\alpha \sim U(0,1)$, $\beta \sim \mathrm{Exp}(1)$
Model 4	$\alpha_c \sim U(0,1)$, $\beta_c \sim \mathrm{Exp}(1)$
Model 5	$\alpha \sim U(0,1)$, $\beta \sim \mathrm{Exp}(1)$, $\alpha_c \sim U(0,1)$, $\beta_c \sim \mathrm{Exp}(1)$

In panel B, all of the softmax parameters β and β_c were increased by 1. This has the effect of reducing the amount of noise in the behavior, which makes the models more easily identifiable and the corresponding confusion matrix more diagonal. The fact that the confusion matrix can be so dependent on the simulating parameter values means that it is crucial to match the simulation parameters to the actual fit parameters as best as possible. Models that are identifiable in one parameter regime may be impossible to distinguish in another!

In addition to the confusion matrices, we also plot the inversion matrices in **Box 5—figure 1C and D**. These are computed from the confusion matrices using Bayes rule assuming a uniform prior on models (see Appendix 3). These matrices more directly address the question of how to interpret a model comparison result where one model fits a particular subject best.

confusion matrix: p(fit model | simulated model)

inversion matrix: p(simulated model | fit model)

Box 5—figure 1. Confusion matrices in the bandit task showing the effect of prior parameter distributions on model recovery. Numbers denote the probability that data generated with model X are best fit by model Y, thus the confusion matrix represents $p(\text{fit model}|\text{simulated model})$. (**A**) When there are relatively large amounts of noise in the models (possibility of small values for β and β_c), models 3–5 are hard to distinguish from one another. (**B**) When there is less noise in the models (i.e. minimum value of β and β_c is 1), the models are much easier to identify. (**C**) The inversion matrix provides easier interpretation of fitting results when the true model is unknown. For example, the confusion matrix indicates that M1 is always perfectly recovered, while M5 is only recovered 30% of the time. By contrast, the inversion matrix shows that if M1 is the best fitting model, our confidence that it generated the data is low (54%), but if M5 is the best fitting model, our confidence that it did generate the data is high (97%). (**D**) Similar results with less noise in simulations.

Run the experiment and analyze the actual data

Once all the previous steps have been completed, you can finally move on to modeling your empirical data. The first step to complete is of course to analyze the data *without* the model, in the same way that we recommended for model simulations in section 'Simulate, simulate, simulate!' This model-independent analysis is extremely important: you designed the experiment to test specific hypotheses, and constructed models to reflect them. Simulations showed expected patterns of behaviors given those hypotheses. If the model-independent analyses do not show evidence of the expected results, there is almost no point in fitting the model. Instead, you should go back to the beginning, either re-thinking the computational models if the analyses show interesting patterns of behavior, or re-thinking the experimental design or even the scientific question you are trying to answer. In our experience, if there is no model-independent evidence that the processes of interest are engaged, then a model-based analysis is unlikely to uncover evidence for the processes either.

If, however, the behavioral results are promising, the next step is to fit the models developed previously and to perform model comparison. After this step, you should check that the parameter range obtained with the fitting is within a range where parameter and model recovery were good. If the range is outside what you explored with simulations, you should go back over the parameter and model recovery steps to match the empirical parameter range, and thus ensure that the model fitting and model comparison procedures lead to interpretable results.

An important point to remember is that human behavior is always messier than the model, and it is unlikely that the class of models you explored actually contains the 'real' model that generated human behavior. At this point, you should consider looping back to Steps 2–5 to improve the models, guided by in depth model-independent analysis of the data.

For example, you may consider modeling 'unimportant parameters', representing mechanisms that are of no interest to your scientific question but that might still affect your measures. Modeling these unimportant parameters usually captures variance in the behavior that would otherwise be attributed to noise, and as such, makes for a better estimation of 'important' parameters. For example, capturing pre-existing biases (e.g. a preference for left/right choices) in a decision or learning task provides better estimation of the inverse temperature, by avoiding attributing systematic biases to noise, which then affords better estimation of other parameters like the learning rate (this is evident in *Box 6—figure 1*).

Box 6. Example: improving parameter recovery by modeling unimportant parameters.

To illustrate the effect that 'unimportant' parameters (i.e., parameters that represent mechanisms that are of no interest to your scientific question, but may still affect your measures) can have on fitting results, we model the effect of a side bias on parameter recovery in Model 3. In particular, we assume that, in addition to choosing based on learned value, the model also had a side bias, B, that effectively changes the value of the left bandit. That is, in the two-bandit case, the choice probabilities are given by

$$p_t^{left} = \frac{1}{1 + \exp\left(\beta(Q_t^{right} - Q_t^{left} - B)\right)} \qquad (11)$$

We then simulated behavior with this model for a range of parameter values and fit the model with the original version of model 3, without the bias, and the modified version of model 3, with the bias. In this simulation, agents learn for 10 independent two-armed bandits in successive 50-trial blocks, with $\mu = \{0.2, 0.8\}$ or $\mu = \{0.8, 0.2\}$ in different blocks. For simplicity, we assumed that the agent treats each block as independent, and started from the same initial values of $Q_1^{right} = Q_1^{left} = 0.5$.

As can be seen below, including the 'unimportant' bias in the fit greatly improves the extent to which both the learning rate, α, and softmax parameter, β, can be recovered.
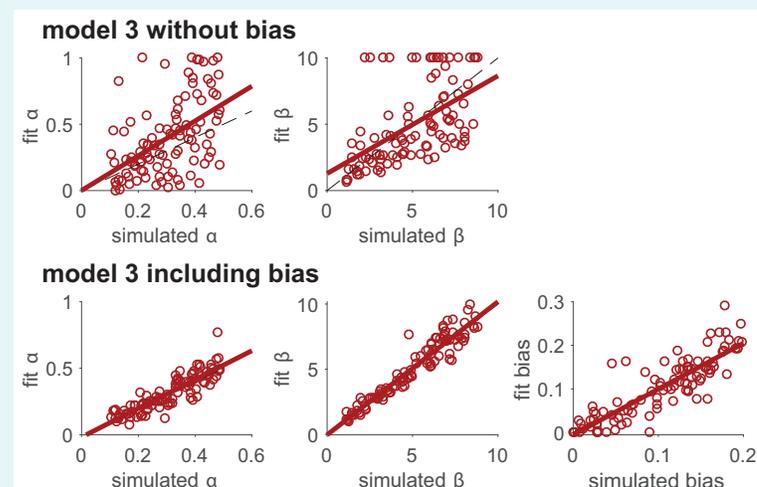
Box 6—figure 1. Modeling unimportant parameters provides better estimation of important parameters. The top row shows parameter recovery of the model without the bias term. The bottom row shows much more accurate parameter recovery, for all parameters, when the bias parameter is included in the model fits.

Validate (at least) the winning model

All the previous steps measure a *relative* goodness of fit. Does model A fit better than model B? However, before interpreting any results from a model, it is essential to ensure that the model actually usefully captures the data in an *absolute* sense. This step is called model validation, and should never be skipped: it is possible to fit a model, get high fit measures, and nevertheless completely miss the essence of the behavior.

One method for model validation is computing the average trial likelihood as an absolute measure of fit. Although this measure has some nice properties — for example, the best possible value is one when the model predicts behavior perfectly — it offers limited value when choices are actually stochastic (which may be the case in many situations; *Drugowitsch et al., 2016*) or the environment is complex. In these cases, the best possible likelihood per trial is less than 1, but it is not known what the best possible likelihood per trial could be. For this reason, although the likelihood per trial can be a useful tool for model validation (*Leong et al., 2017*), interpreting it as an absolute measure of model fit is of limited value.

A better method to validate a model is to simulate it with the fit parameter values (*Palminteri et al., 2017*; *Nassar and Frank, 2016*; *Navarro, 2019*), a procedure long performed by statisticians as part of the 'posterior predictive check' (*Roecker, 1991*; *Gelman et al., 1996*). You should then analyze the simulated data in the same way that you analyzed the empirical data, to verify that all important behavioral effects are qualitatively and quantitatively captured by the

simulations with the fit parameters. For example, if you observe a qualitative difference between two conditions empirically, the model should reproduce it. Likewise, if a learning curve reaches a quantitative asymptote of 0.7, simulations shouldn't reach a vastly different one.

Some researchers analyze the posterior prediction of the model conditioned on the past history, instead of simulated data. In our previous notation, they evaluate the likelihood of choice c_t given past data, $d_{1:t-1}$, where the past data includes choices made by the subject, *not* choices made by the model, $p(c_t|d_{1:t-1}, s_t, \theta_m, m)$. In some cases, this approach leads to very similar results to simulations, because simulations sample choices on the basis of a very similar probability, where the past data, $d_{1:t-1}$, include choices made by the *model*. However, it can also be dramatically different if the path of actions sampled by the participant is widely different from the paths likely to be selected by the model (leading to very different past histories).

Palminteri and colleagues (*Palminteri et al., 2017*) offer a striking example of this effect, where Model A fits better than Model B by any quantitative measure of model comparison, but is completely unable to capture the essence of the behavior. In their example, data are generated with a reinforcement learning agent (which takes the place of the subject) on a reversal learning task (where a choice that was previously good becomes bad, and reciprocally). These data are then fit with either a win-stay lose-shift model (model B), or a simplistic choice kernel model, which assumes that previous choices tend to be repeated (model A). Because of the autocorrelation in the choices made by the reinforcement learning agent, model A, which tends to repeat previous actions, fits better than model B, whose win-stay-lose-shift choices only depend on the action and outcome from the last trial. However, model A is completely insensitive to reward, and thus is unable to generate a reversal behavior when it is simulated with the fit model parameters. Thus, in this case, model A should be discarded, despite a greater quantitative fit. Nevertheless, the fact that the best validating model B captures less variance than model A should serve as a warning that model B is missing crucial components of the data and that a better model probably exists. This should incite the researcher to go back to the drawing board to develop a better model, for example one that combines elements of both models or a different model entirely, and perhaps a better experiment to test it.

More generally, if your validation step fails, you should go back to the drawing board! This may involve looking for a better model, as well as redesigning the task. Be careful interpreting results from a model that is not well validated! Of course, exactly what it means for a model to 'fail' the validation step is not well defined: no model is perfect, and there is no rule of thumb to tell us when a model is *good enough*. The most important aspect of validation is for you (and your readers) to be aware of its limitations, and in which ways they may influence any downstream results.

Box 7. Example: model validation where the fit model performs too well.

Most examples of model validation involve a case where a model that fits well performs poorly on the task in simulation. For example, in the *Palminteri et al. (2017)* example, the choice kernel model cannot perform the task at all because its behavior is completely independent of reward. Here, we offer a different example of failed model validation in which the model performs *better* in simulation than the predicted and observed artificial agent's behavior. Moreover, this model appears to fit data generated from a different model better than it fits data generated from itself! In this example, we imagine a deterministic stimulus-action learning task in which agents are presented with one of three stimuli (s_1, s_2, and s_3), which instruct them which of three actions (a_1, a_2, and a_3) will be rewarded when chosen. a_1 is the correct choice for both stimuli s_1 and s_2, a_3 for s_3, and a_2 is incorrect for all stimuli.

The two models that we consider are both reinforcement learning agents. The first, a 'blind' agent does not see the stimulus at all and learns only about the value of the three different actions, that is $Q(a_i)$, regardless of the stimulus. The second, a 'state-based' agent, observes the stimulus and learns a value for each action that can be different for each stimulus, that is $Q(a_i, s_i)$. Parameters in the models are set such that the learning curves for the two agents are approximately equal (*Box 7—figure 1A*). See appendices for details of the models.

We then consider how both models fit behavior simulated by either of these models. In *Box 7—figure 1B*, we plot the average likelihood with which the state-based model predicts the actual choices of the blind and state-based agents, that is the average $p(c_t|d_{1:t-1}, \theta_m, m = \text{state-based})$. As is clear from this figure, the state-based model predicts choices from the blind agent with *higher* likelihood than choices from the state-based agent! While counter intuitive, this result does not imply that the state-based model is unable to fit its own behavior. Instead, this result reflects the difference in noise (softmax parameters) between the two agents. The blind RL agent has a low noise parameter, allowing the state-based model to fit it quite well. Conversely, the state-based RL agent has a high noise parameter, meaning that the behavior is harder to predict even when it is fit with the correct model.

That the state-based model captures state-based behavior better than it fits blind behavior is illustrated in *Box 7—figure 1C*. Here, we plot the simulated learning curves of the state-based model using the parameter values that were fit to either the state-based agent or the blind agent. Although the parameters of the state-based model obtained through the fit to the state-based agent generate a learning curve that is quite similar to that of the agent (compare blue lines in *Box 7—figure 1A and C*), the state-based fit to the blind agent performs too well (compare yellow lines in *Box 7—figure 1A and C*).

Thus the model validation step provides support for the state-based model when it is the correct model of behavior, but rules out the state-based model when the generating model was different. The take-away from this example should be that measures of model-fit and model comparison cannot replace a thorough validation step, which can contradict them.
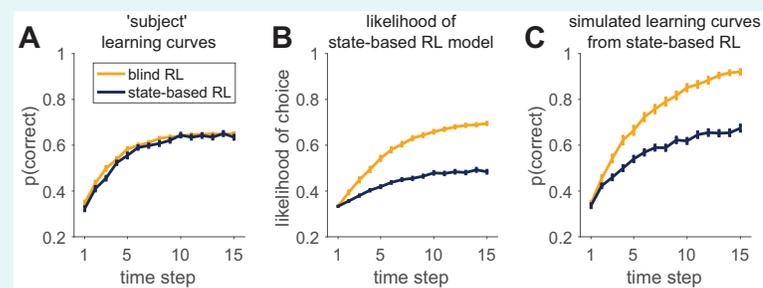
Box 7—figure 1. An example of successful and unsuccessful model validation. (**A**) Behavior is simulated by one of two reinforcement learning models (a blind agent and a state-based agent) performing the same learning task. Generative parameters of the two models were set so that the learning curves of the models were approximately equal. (**B**) Likelihood per trial seems to indicate a worse fit for the state-based-simulated data than the blind-simulated data. (**C**) However, validation by model simulations with fit parameters shows that the state-based model captures the data from the state-based agent (compare dark learning curves in panels A and C), but not from the the blind agent (yellow learning curves in panels A and C).

Analyze the winning model

To minimize risks of p-hacking, model-dependent analyses should only be performed on the winning model, after researchers are satisfied that the model captures the behavior. One particularly powerful application of model-based analysis of behavior involves estimating the *latent* variables in the model. Latent variables are the hidden components of the algorithms underlying the behavior that are not directly observable from the behavior itself. These latent variables shed light on the internal workings of the model and, if we are to take the model seriously, should have some representation in the subjects' mind and brain (*Cohen et al., 2017*; *O'Doherty et al., 2007*).

Extracting latent variables from the model is as simple as simulating the model and recording how the latent variables evolve over time. The parameters of the simulation should be the fit parameters for each subject. In most cases, it is useful to yoke the choices of the model to the choices the

participants actually made, thus the latent variables evolve according to the experience participants actually had. This is especially true if the choices can influence what participants see in the future.

Once estimated, the latent variables can be used in much the same way as any other observable variable in the analysis of data. Perhaps the most powerful application comes when combined with physiological data such as pupil dilation, EEG, and fMRI. The simplest of these approaches uses linear regression to test whether physiological variables correlate with the latent variables of interest. Such an approach has led to a number of insights into the neural mechanisms underlying behavior (*Nassar et al., 2012*; *Daw et al., 2011*; *Donoso et al., 2014*; *Collins and Frank, 2018*; *Fischer and Ullsperger, 2013*), although, as with any modeling exercise, latent variable analysis should be done with care (*Wilson and Niv, 2015*).

Other model-dependent analyzes include studying individual differences as captured by fit parameters. Fit parameters can be treated as a dependent variable in continuous analyses (e.g. correlating with age, symptom scales, and so on [*Gillan et al., 2016*]) or group comparisons (e.g. patients vs. matched controls [*Collins et al., 2014*]).

Reporting model-based analyses

Congratulations! You have developed, simulated, and fit your model (and maybe several other competing models) to your data. You have estimated parameters, computed model comparison scores, and validated whether your model can generate realistic-looking behavior. It's time to start writing! But what exactly should you report in your paper? And how should you report it?

Model selection

In many modeling papers, a key conclusion from the work is that one model fits the data better than other competing models. To make this point convincingly, we recommend including the following things in your paper, either as main results or in the supplementary material.

Model recovery analysis
Confusion matrix

Before anyone should believe your model comparison results, you need to demonstrate the ability of your analysis/experiment to distinguish between models under ideal conditions of simulated data. The best way to visualize these results is with a confusion matrix, as outlined in section 'Can you arbitrate between different models'? If the model comparison result is central to your paper, we recommend including the confusion matrix as a figure in the main text. If model comparison is less important, we recommend including it in the supplementary materials.

Number of subjects best fit by each model

The simplest way to visualize how well the winning model fits the data is with a histogram showing the number of subjects best fit by each model. Obviously if all subjects are best fit with one model, the story is simple. The more likely scenario is that some subjects will be best fit by other models. Such a result is important to acknowledge in the paper as it may reflect the use of different strategies by different people or that the 'correct' model lies somewhere in between the models you have considered.

Group level statistics
Exceedance probabilities

A more sophisticated and less biased (*Piray et al., 2018*) way to report model comparison results is by computing the probability that a single model best describes all the data. This is clearly an assumption whose merits should be discussed in your paper. In cases where it is valid, the method of *Rigoux et al. (2014)* computes these 'Exceedance Probabilities', the probability that each model generated all the data. These probabilities can also be reported in histogram or table form.

Model-independent measures of simulated data. The cleanest way to demonstrate the superiority of one model is if that model can account for qualitative patterns in the data that are not captured by other models (see section 'Validate (at least) the winning model').

Parameter fits

Many modeling papers involve fitting parameters to behavioral data. In some cases this is the main point of the paper, for example to show that parameter values differ between groups or treatments, in other cases parameter fitting is secondary to model comparison. In all cases, we recommend reporting the fit parameter values in as transparent a way as possible (i.e. more than just the means and standard errors).

Report distributions of parameter values

The simplest way to report parameter fits is to plot a distribution of all fit parameter values, for example in the form of a histogram (e.g. Figure S1 in *Wilson et al., 2013* and *Nassar et al., 2018*) or a cloud of points (e.g. Figure 5 in *Huys et al., 2011*). This gives a great sense of the variability in each parameter across the population and can also illustrate problems with fitting. For example, if a large number of fit parameters are clustered around the upper and lower bounds, this may indicate a problem with the model.

Plot pairwise correlations between fit parameter values

A deeper understanding of the relationships *between* fit parameters can be obtained by making scatter plots of the pairwise correlations between parameters. As with histograms of individual parameters, this approach gives a sense of the distribution of parameters, and can provide evidence of problems with the model; for example, if two parameters trade off against one another, it is a sign that these parameters may be unidentifiable in the experiment.

Report parameter recovery

Finally, all parameter fit analyses should sit on the shoulders of a comprehensive parameter recovery analysis with simulated data. If parameters cannot be recovered in the ideal case of simulated data, there is little that they can tell us about real behavior.

Share your data and code!

The most direct way to communicate your results is to share the data and code. This approach encourages transparency and ensures that others can see *exactly* what you did. Sharing data and code also allows others to extend your analyses easily, by applying it to their own data or adding new models into the mix.

Ideally the data you share should be the raw data for the experiment, with minimal or no preprocessing (apart from the removal of identifying information). The code you share should reproduce all steps in your analysis, including any preprocessing/outlier exclusion you may have performed and generating all of the main and supplementary figures in the paper. In a perfect world, both data and code would be shared publicly on sites such as GitHub, DataVerse and so on. However, this is not always possible, for example, if the data come from collaborators who do not agree to data sharing, or if further analyses are planned using the same data set. In this case, we recommend having a clean set of 'shareable' code (and hopefully data too) that can be sent via email upon request.

Should you always report all of your modeling results?

Finally, if you are using an established model, it can be tempting to skip many of the steps outlined above and report only the most exciting results. This temptation can be even greater if you are using code developed by someone else that, perhaps, you do not fully understand. In our opinion, taking shortcuts like this is dangerous. For one thing, your experiment or population may be different and the model may perform differently in this regime. For another, quite often 'established' models (in the sense that they have been published before), have not been validated in a systematic way. More generally, as with any research technique, when using computational modeling you need to demonstrate that you are applying the method correctly, and the that steps we outline here can help. In conclusion, even if developing the model is not the central point of your paper, you should report all of your modeling results.

What now?

Looping back

A modeler's work is never done. To paraphrase George Box, there are no correct models, there are only useful models (*Box, 1979*). To make your model more useful, there are a number of next steps to consider to test whether your model really does describe a process in the mind.

Improve the model to account for discrepancies with your existing data set

Model fits are never perfect and, even in the best cases, there are often small discrepancies with actual data. The simplest next step is to try to address these discrepancies by improving the model, either by including additional factors (such as side bias or lapse rates) or by devising new models entirely.

Use your model to make predictions

The best models don't just explain data in one experiment, they *predict* data in completely new situations. If your model does not easily generalize to new situations, try to understand why that is and how it could be adjusted to be more general. If your model does generalize, test its predictions against new data — either data you collect yourself from a new experiment or data from other studies that (hopefully) have been shared online.

Using advanced techniques

Another potential next step is to use more powerful modeling techniques. We focused here on the simplest techniques (maximum likelihood estimation and model comparison by BIC) because of their accessibility to beginners, and because most of the advice we give here generalizes to more advanced techniques. In particular, no matter how advanced the modeling technique used, validation is essential (*Palminteri et al., 2017*; *Nassar and Frank, 2016*; *Huys, 2017*). Nevertheless, the simple methods described here have known limitations. More advanced techniques attempt to remedy them, but come with their own pitfalls. A complete review of these advanced techniques is beyond the scope of this paper; instead we provide pointers to a few of the most interesting techniques for the ambitious reader to pursue.

Compute maximum a posteriori (MAP) parameter values

Perhaps the simplest step for improving parameter estimates is to include prior information about parameter values. When combined with the likelihood, these priors allow us to compute the posterior, which we can use to find the maximum a posteriori (MAP) parameter values. Although they are still point estimates, with good priors, MAP parameters can be more accurate than parameters estimated with maximum likelihood approaches (*Gershman, 2016*; *Daw, 2011*), although when the priors are bad, this method has problems of its own (*Katahira, 2016*).

Approximate the full posterior by sampling

Point estimates of model parameters, such as those obtained with MLE or MAP, lose interesting information about uncertainty over the parameter distribution. Sampling approaches (such as Markov Chain Monte Carlo or MCMC) provide this richer information; furthermore, they allow modelers to investigate more complex assumptions. For example, hierarchical Bayesian approaches make it possible to fit all participants simultaneously, integrating assumptions about their dependence (e.g. one single group, multiple groups, effects of covariates of interest such as age and so on; *Lee, 2011*; *Lee and Wagenmakers, 2014*; *Wiecki et al., 2013*).

Advanced optimizers and approximate likelihood

Some models have intractable likelihoods, for example if the choice state has too many dimensions, as in continuous movements, or if the model included unobservable choices. There exist methods to approximate likelihoods to relate them quantitatively to data, such as the ABC method (*Turner and Sederberg, 2012*; *Sunnåker et al., 2013*). There are also advanced methods for finding best fit parameters in a sample-efficient manner when computing the likelihood is expensive (*Acerbi and Ji, 2017*; *Acerbi, 2018*).

Model selection

Bayesian model selection provides less biased, statistically more accurate ways of identifying which model is best at the group level (*Rigoux et al., 2014*). This may be particularly important when comparing model selection between groups, for example between patients and controls (*Piray et al., 2018*).

Incorporating other types of data

We focused on modeling a single type of observable data, choices. However, there is a rich literature on fitting models to other measurements, such as reaction times (*Ratcliff, 1978*; *Ratcliff and Rouder, 1998*), but also to eye movements and neural data (*Turner et al., 2016*). Furthermore, fitting more than one measurement at a time provides additional constraints to the model, and as such may provide better fit (*Ballard and McClure, 2019*). However, fitting additional data can increase the complexity of the model-fitting process and additional care must be taken to determine exactly how different types of data should be combined (*Viejo et al., 2015*).

Epilogue

Our goal for this paper was to offer practical advice, for beginners as well as seasoned researchers, on the computational modeling of behavioral data. To this end, we offered guidance on how to generate models, simulate models, fit models, compare models, validate models, and extract latent variables from models to compare with physiological data. We have talked about how to avoid common pitfalls and misinterpretations that can arise with computational modeling, and lingered, quite deliberately, on the importance of good experimental design. Many of these lessons were lessons we learned the hard way, by actually making these mistakes for ourselves over a combined 20+ years in the field. By following these steps, we hope that you will avoid some of the errors that slowed our own research, and that the overall quality of computational modeling in behavioral science will improve.

Acknowledgements

Additional information

Funding

Author ORCIDs

Robert C Wilson https://orcid.org/0000-0002-2963-2971
Anne GE Collins https://orcid.org/0000-0003-3751-3662

References

Abbott JT, Austerweil JL, Griffiths TL. 2015. Random walks on semantic networks can resemble optimal foraging. *Psychological Review* **122**:558–569. DOI: https://doi.org/10.1037/a0038693, PMID: 25642588

Acerbi L. 2018. Variational bayesian monte carlo. Advances in Neural Information Processing Systems 8213–8223. https://papers.nips.cc/paper/8043-variational-bayesian-monte-carlo.

Acerbi L, Ji W. 2017. Practical bayesian optimization for model fitting with bayesian adaptive direct search. Advances in Neural Information Processing Systems 1836–1846. https://papers.nips.cc/paper/6780-practical-bayesian-optimization-for-model-fitting-with-bayesian-adaptive-direct-search.

Akaike H. 1974. A new look at the statistical model identification. *IEEE Transactions on Automatic Control* **19**: 716–723. DOI: https://doi.org/10.1109/TAC.1974.1100705

Ballard IC, McClure SM. 2019. Joint modeling of reaction times and choice improves parameter identifiability in reinforcement learning models. *Journal of Neuroscience Methods* **317**:37–44. DOI: https://doi.org/10.1016/j.jneumeth.2019.01.006, PMID: 30664916

Batchelder WH, Riefer DM. 1990. Multinomial processing models of source monitoring. *Psychological Review* **97**:548–564. DOI: https://doi.org/10.1037/0033-295X.97.4.548

Box GE. 1979. Robustness in the strategy of scientific model building. In: *Robustness in Statistics*. Elsevier. p. 201–236. DOI: https://doi.org/10.1016/b978-0-12-438150-6.50018-2

Breiman L. 2001. Statistical modeling: the two cultures (with comments and a rejoinder by the author). *Statistical Science* **16**:199–231. DOI: https://doi.org/10.1214/ss/1009213726

Broomell SB, Bhatia S. 2014. Parameter recovery for decision modeling using choice data. *Decision* **1**:252–274. DOI: https://doi.org/10.1037/dec0000020

Busemeyer JR, Diederich A. 2010. *Cognitive Modeling*. Sage.

Byrd RH, Gilbert JC, Nocedal J. 2000. A trust region method based on interior point techniques for nonlinear programming. *Mathematical Programming* **89**:149–185. DOI: https://doi.org/10.1007/PL00011391

Cavanagh JF, Wiecki TV, Kochar A, Frank MJ. 2014. Eye tracking and pupillometry are indicators of dissociable latent decision processes. *Journal of Experimental Psychology: General* **143**:1476–1488. DOI: https://doi.org/10.1037/a0035813

Cohen JD, Dunbar K, McClelland JL. 1990. On the control of automatic processes: a parallel distributed processing account of the stroop effect. *Psychological Review* **97**:332–361. DOI: https://doi.org/10.1037/0033-295X.97.3.332, PMID: 2200075

Cohen JD, Daw N, Engelhardt B, Hasson U, Li K, Niv Y, Norman KA, Pillow J, Ramadge PJ, Turk-Browne NB, Willke TL. 2017. Computational approaches to fMRI analysis. *Nature Neuroscience* **20**:304–313. DOI: https://doi.org/10.1038/nn.4499, PMID: 28230848

Collins AG, Brown JK, Gold JM, Waltz JA, Frank MJ. 2014. Working memory contributions to reinforcement learning impairments in schizophrenia. *The Journal of Neuroscience* **34**:13747–13756. DOI: https://doi.org/10.1523/JNEUROSCI.0989-14.2014, PMID: 25297101

Collins AG, Frank MJ. 2012. How much of reinforcement learning is working memory, not reinforcement learning? A behavioral, computational, and neurogenetic analysis. *European Journal of Neuroscience* **35**:1024–1035. DOI: https://doi.org/10.1111/j.1460-9568.2011.07980.x, PMID: 22487033

Collins AG, Frank MJ. 2013. Cognitive control over learning: creating, clustering, and generalizing task-set structure. *Psychological Review* **120**:190–229. DOI: https://doi.org/10.1037/a0030852, PMID: 23356780

Collins AG, Frank MJ. 2014. Opponent actor learning (OpAL): modeling interactive effects of striatal dopamine on reinforcement learning and choice incentive. *Psychological Review* **121**:337–366. DOI: https://doi.org/10.1037/a0037015, PMID: 25090423

Collins AG, Frank MJ. 2018. Within-and across-trial dynamics of human EEG reveal cooperative interplay between reinforcement learning and working memory. *PNAS*:201720963.

Collins AG, Wilson RC. 2019. TenSimpleRulesModeling. 3a01850. Github. https://github.com/AnneCollins/TenSimpleRulesModeling

Courville AC, Daw ND. 2008. The rat as particle filter. Advances in Neural Information Processing Systems 369–376. https://papers.nips.cc/paper/3205-the-rat-as-particle-filter.

Daw ND. 2011. Trial-by-trial data analysis using computational models. *Decision Making, Affect, and Learning: Attention and Performance XXIII* **23**:3–38. DOI: https://doi.org/10.1093/acprof:oso/9780199600434.003.0001

Daw ND, Gershman SJ, Seymour B, Dayan P, Dolan RJ. 2011. Model-based influences on humans' choices and striatal prediction errors. *Neuron* **69**:1204–1215. DOI: https://doi.org/10.1016/j.neuron.2011.02.027, PMID: 21435563

Daw ND, Courville AC. 2007. The pigeon as particle filter. Advances in Neural Information Processing Systems.

Daw ND, Tobler PN. 2014. Value learning through reinforcement: the basics of dopamine and reinforcement learning. In: *Neuroeconomics*. Second Edition. Elsevier. p. 283–298. DOI: https://doi.org/10.1016/b978-0-12-416008-8.00015-2

Donkin C, Tran SC, Nosofsky R. 2014. Landscaping analyses of the ROC predictions of discrete-slots and signal-detection models of visual working memory. *Attention, Perception, & Psychophysics* **76**:2103–2116. DOI: https://doi.org/10.3758/s13414-013-0561-7

Donkin C, Kary A, Tahir F, Taylor R. 2016. Resources masquerading as slots: Flexible allocation of visual working memory. *Cognitive Psychology* **85**:30–42. DOI: https://doi.org/10.1016/j.cogpsych.2016.01.002

Donoso M, Collins AGE, Koechlin E. 2014. Foundations of human reasoning in the prefrontal cortex. *Science* **344**:1481–1486. DOI: https://doi.org/10.1126/science.1252254

Dowd EC, Frank MJ, Collins A, Gold JM, Barch DM. 2016. Probabilistic reinforcement learning in patients with schizophrenia: Relationships to anhedonia and avolition. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging* **1**:460–473. DOI: https://doi.org/10.1016/j.bpsc.2016.05.005

Drugowitsch J, Wyart V, Devauchelle AD, Koechlin E. 2016. Computational precision of mental inference as critical source of human choice suboptimality. *Neuron* **92**:1398–1411. DOI: https://doi.org/10.1016/j.neuron.2016.11.005, PMID: 27916454

Farashahi S, Rowe K, Aslami Z, Lee D, Soltani A. 2017. Feature-based learning improves adaptability without compromising precision. *Nature Communications* **8**:1768. DOI: https://doi.org/10.1038/s41467-017-01874-w, PMID: 29170381

Farrell S, Lewandowsky S. 2018. *Computational Modeling of Cognition and Behavior*. Cambridge University Press. DOI: https://doi.org/10.1017/CBO9781316272503

Findling C, Skvortsova V, Dromnelle R, Palminteri S, Wyart V. 2018. Computational noise in reward-guided learning drives behavioral variability in volatile environments. *bioRxiv*. DOI: https://doi.org/10.1101/439885

Fischer AG, Ullsperger M. 2013. Real and fictive outcomes are processed differently but converge on a common adaptive mechanism. *Neuron* **79**:1243–1255. DOI: https://doi.org/10.1016/j.neuron.2013.07.006, PMID: 24050408

Frank MJ, Seeberger LC, O'reilly RC. 2004. By carrot or by stick: cognitive reinforcement learning in parkinsonism. *Science* **306**:1940–1943. DOI: https://doi.org/10.1126/science.1102941, PMID: 15528409

Frank MJ, Moustafa AA, Haughey HM, Curran T, Hutchison KE. 2007. Genetic triple dissociation reveals multiple roles for dopamine in reinforcement learning. *PNAS* **104**:16311–16316. DOI: https://doi.org/10.1073/pnas.0706111104, PMID: 17913879

Geisler WS. 2011. Contributions of ideal observer theory to vision research. *Vision Research* **51**:771–781. DOI: https://doi.org/10.1016/j.visres.2010.09.027, PMID: 20920517

Gelman A, Meng XL, Stern H. 1996. Posterior predictive assessment of model fitness via realized discrepancies. *Statistica Sinica* **6**:733–760.

Gershman SJ. 2016. Empirical priors for reinforcement learning models. *Journal of Mathematical Psychology* **71**:1–6. DOI: https://doi.org/10.1016/j.jmp.2016.01.006

Gillan CM, Kosinski M, Whelan R, Phelps EA, Daw ND. 2016. Characterizing a psychiatric symptom dimension related to deficits in goal-directed control. *eLife* **5**:e11305. DOI: https://doi.org/10.7554/eLife.11305, PMID: 26928075

Haaf JM, Rouder JN. 2017. Developing constraint in bayesian mixed models. *Psychological Methods* **22**:779–798. DOI: https://doi.org/10.1037/met0000156, PMID: 29265850

Heathcote A, Brown SD, Wagenmakers EJ. 2015. An introduction to good practices in cognitive modeling. In: *An Introduction to Model-Based Cognitive Neuroscience*. Springer. p. 25–48. DOI: https://doi.org/10.1007/978-1-4939-2236-9_2

Huys QJM, Cools R, Gölzer M, Friedel E, Heinz A, Dolan RJ, Dayan P. 2011. Disentangling the roles of approach, activation and Valence in instrumental and pavlovian responding. *PLOS Computational Biology* **7**:e1002028. DOI: https://doi.org/10.1371/journal.pcbi.1002028

Huys QJM. 2017. Bayesian Approaches to Learning and Decision-Making. In: *Computational Psychiatry: Mathematical Modeling of Mental Illness*. Academic Press. p. 247–271. DOI: https://doi.org/10.1016/b978-0-12-809825-7.00010-9

Jahfari S, Ridderinkhof KR, Collins AGE, Knapen T, Waldorp LJ, Frank MJ. 2019. Cross-Task contributions of frontobasal ganglia circuitry in response inhibition and Conflict-Induced slowing. *Cerebral Cortex* **29**:1969–1983. DOI: https://doi.org/10.1093/cercor/bhy076, PMID: 29912363

Kass RE, Raftery AE. 1995. Bayes factors. *Journal of the American Statistical Association* **90**:773–795. DOI: https://doi.org/10.1080/01621459.1995.10476572

Katahira K. 2016. How hierarchical models improve point estimates of model parameters at the individual level. *Journal of Mathematical Psychology* **73**:37–58. DOI: https://doi.org/10.1016/j.jmp.2016.03.007

Kording K, Blohm G, Schrater P, Kay K. 2018. Appreciating diversity of goals in computational neuroscience. *OSF Preprints*. https://osf.io/3vy69/.

Lee MD. 2011. How cognitive modeling can benefit from hierarchical bayesian models. *Journal of Mathematical Psychology* **55**:1–7. DOI: https://doi.org/10.1016/j.jmp.2010.08.013

Lee MD, Criss AH, Devezer B, Donkin C, Etz A, Leite FP, Matzke D, Rouder JN, Trueblood J, White C. 2019. Robust modeling in cognitive science. *PsyArXiv*. https://psyarxiv.com/dmfhk/.

Lee MD, Wagenmakers EJ. 2014. *Bayesian Cognitive Modeling: A Practical Course*. Cambridge university press. DOI: https://doi.org/10.1017/CBO9781139087759

Lee MD, Webb MR. 2005. Modeling individual differences in cognition. *Psychonomic Bulletin & Review* **12**:605–621. DOI: https://doi.org/10.3758/BF03196751, PMID: 16447375

Leong YC, Radulescu A, Daniel R, DeWoskin V, Niv Y. 2017. Dynamic interaction between reinforcement learning and attention in multidimensional environments. *Neuron* **93**:451–463. DOI: https://doi.org/10.1016/j.neuron.2016.12.040

Lieder F, Griffiths TL, M Huys QJ, Goodman ND. 2018. Empirical evidence for resource-rational anchoring and adjustment. *Psychonomic Bulletin & Review* **25**:775–784. DOI: https://doi.org/10.3758/s13423-017-1288-6, PMID: 28484951

Lorains FK, Dowling NA, Enticott PG, Bradshaw JL, Trueblood JS, Stout JC. 2014. Strategic and non-strategic problem gamblers differ on decision-making under risk and ambiguity. *Addiction* **109**:1128–1137. DOI: https://doi.org/10.1111/add.12494, PMID: 24450756

MacKay DJ. 2003. *Information Theory, Inference and Learning Algorithms.* Cambridge University Press.

Montague PR, Dayan P, Sejnowski TJ. 1996. A framework for mesencephalic dopamine systems based on predictive hebbian learning. *The Journal of Neuroscience* **16**:1936–1947. DOI: https://doi.org/10.1523/JNEUROSCI.16-05-01936.1996, PMID: 8774460

Moré JJ, Sorensen DC. 1983. Computing a trust region step. *SIAM Journal on Scientific and Statistical Computing* **4**:553–572. DOI: https://doi.org/10.1137/0904038

Nassar MR, Wilson RC, Heasly B, Gold JI. 2010. An approximately bayesian delta-rule model explains the dynamics of belief updating in a changing environment. *Journal of Neuroscience* **30**:12366–12378. DOI: https://doi.org/10.1523/JNEUROSCI.0822-10.2010, PMID: 20844132

Nassar MR, Rumsey KM, Wilson RC, Parikh K, Heasly B, Gold JI. 2012. Rational regulation of learning dynamics by pupil-linked arousal systems. *Nature Neuroscience* **15**:1040–1046. DOI: https://doi.org/10.1038/nn.3130, PMID: 22660479

Nassar MR, Helmers JC, Frank MJ. 2018. Chunking as a rational strategy for lossy data compression in visual working memory. *Psychological Review* **125**:486–511. DOI: https://doi.org/10.1037/rev0000101, PMID: 29952621

Nassar MR, Frank MJ. 2016. Taming the beast: extracting generalizable knowledge from computational models of cognition. *Current Opinion in Behavioral Sciences* **11**:49–54. DOI: https://doi.org/10.1016/j.cobeha.2016.04.003, PMID: 27574699

Navarro DJ. 2019. Between the Devil and the deep blue sea: tensions between scientific judgement and statistical model selection. *Computational Brain & Behavior* **2**:28–34.

Nilsson H, Rieskamp J, Wagenmakers E-J. 2011. Hierarchical Bayesian parameter estimation for cumulative prospect theory. *Journal of Mathematical Psychology* **55**:84–93. DOI: https://doi.org/10.1016/j.jmp.2010.08.006

O'Doherty JP, Hampton A, Kim H. 2007. Model-based fMRI and its application to reward learning and decision making. *Annals of the New York Academy of Sciences* **1104**:35–53. DOI: https://doi.org/10.1196/annals.1390.022, PMID: 17416921

O'Reilly JX, Schuffelgen U, Cuell SF, Behrens TEJ, Mars RB, Rushworth MFS. 2013. Dissociable effects of surprise and model update in parietal and anterior cingulate cortex. *PNAS* **110**:E3660–E3669. DOI: https://doi.org/10.1073/pnas.1305373110

Otto AR, Raio CM, Chiang A, Phelps EA, Daw ND. 2013. Working-memory capacity protects model-based learning from stress. *PNAS* **110**:20941–20946. DOI: https://doi.org/10.1073/pnas.1312011110, PMID: 24324166

Palminteri S, Wyart V, Koechlin E. 2017. The importance of falsification in computational cognitive modeling. *Trends in Cognitive Sciences* **21**:425–433. DOI: https://doi.org/10.1016/j.tics.2017.03.011, PMID: 28476348

Piray P, Dezfouli A, Heskes T, Frank MJ, Daw ND. 2018. Hierarchical bayesian inference for concurrent model fitting and comparison for group studies. *bioRxiv.* DOI: https://doi.org/10.1101/393561

Ratcliff R. 1978. A theory of memory retrieval. *Psychological Review* **85**:59–108. DOI: https://doi.org/10.1037/0033-295X.85.2.59

Ratcliff R, Rouder JN. 1998. Modeling response times for Two-Choice decisions. *Psychological Science* **9**:347–356. DOI: https://doi.org/10.1111/1467-9280.00067

Rescorla RA, Wagner AR. 1972. A theory of pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. *Classical Conditioning II: Current Research and Theory* **2**:64–99.

Rigoux L, Stephan KE, Friston KJ, Daunizeau J. 2014. Bayesian model selection for group studies - revisited. *NeuroImage* **84**:971–985. DOI: https://doi.org/10.1016/j.neuroimage.2013.08.065, PMID: 24018303

Roecker EB. 1991. Prediction error and its estimation for Subset-Selected models. *Technometrics* **33**:459–468. DOI: https://doi.org/10.1080/00401706.1991.10484873

Samejima K, Ueda Y, Doya K, Kimura M. 2005. Representation of action-specific reward values in the striatum. *Science* **310**:1337–1340. DOI: https://doi.org/10.1126/science.1115270, PMID: 16311337

Schwarz G. 1978. Estimating the dimension of a model. *The Annals of Statistics* **6**:461–464. DOI: https://doi.org/10.1214/aos/1176344136

Sims CR. 2018. Efficient coding explains the universal law of generalization in human perception. *Science* **360**:652–656. DOI: https://doi.org/10.1126/science.aaq1118, PMID: 29748284

Somerville LH, Sasse SF, Garrad MC, Drysdale AT, Abi Akar N, Insel C, Wilson RC. 2017. Charting the expansion of strategic exploratory behavior during adolescence. *Journal of Experimental Psychology: General* **146**:155–164. DOI: https://doi.org/10.1037/xge0000250

Starns JJ, Ratcliff R. 2010. The effects of aging on the speed-accuracy compromise: boundary optimality in the diffusion model. *Psychology and Aging* **25**:377–390. DOI: https://doi.org/10.1037/a0018022, PMID: 20545422

Steyvers M, Lee MD, Wagenmakers E-J. 2009. A bayesian analysis of human decision-making on bandit problems. *Journal of Mathematical Psychology* **53**:168–179. DOI: https://doi.org/10.1016/j.jmp.2008.11.002

Sunnåker M, Busetto AG, Numminen E, Corander J, Foll M, Dessimoz C. 2013. Approximate bayesian computation. *PLOS Computational Biology* **9**:e1002803. DOI: https://doi.org/10.1371/journal.pcbi.1002803, PMID: 23341757

Sutton RS, Barto AG. 2018. *Reinforcement Learning: An Introduction.* MIT press.

Turner BM, Forstmann BU, Wagenmakers EJ, Brown SD, Sederberg PB, Steyvers M. 2013. A bayesian framework for simultaneously modeling neural and behavioral data. *NeuroImage* **72**:193–206. DOI: https://doi.org/10.1016/j.neuroimage.2013.01.048, PMID: 23370060

Turner BM, Rodriguez CA, Norcia TM, McClure SM, Steyvers M. 2016. Why more is better: simultaneous modeling of EEG, fMRI, and behavioral data. *NeuroImage* **128**:96–115. DOI: https://doi.org/10.1016/j.neuroimage.2015.12.030, PMID: 26723544

Turner BM, Sederberg PB. 2012. Approximate Bayesian computation with differential evolution. *Journal of Mathematical Psychology* **56**:375–385. DOI: https://doi.org/10.1016/j.jmp.2012.06.004

van Ravenzwaaij D, Dutilh G, Wagenmakers E-J. 2011. Cognitive model decomposition of the BART: assessment and application. *Journal of Mathematical Psychology* **55**:94–105. DOI: https://doi.org/10.1016/j.jmp.2010.08.010

Vandekerckhove J, Matzke D, Wagenmakers EJ. 2015. Model Comparison and the Principle of Parsimony. In: Busemeyer J. R, Wang Z, Townsend J. T, Eidels A (Eds). *The Oxford Handbook of Computational and Mathematical Psychology*. **300** Oxford University Press. p. 300319 DOI: https://doi.org/10.1093/oxfordhb/9780199957996.013.14

Viejo G, Khamassi M, Brovelli A, Girard B. 2015. Modeling choice and reaction time during arbitrary visuomotor learning through the coordination of adaptive working memory and reinforcement learning. *Frontiers in Behavioral Neuroscience* **9**:225. DOI: https://doi.org/10.3389/fnbeh.2015.00225, PMID: 26379518

Wagenmakers E-J, Lodewyckx T, Kuriyal H, Grasman R. 2010. Bayesian hypothesis testing for psychologists: a tutorial on the Savage–Dickey method. *Cognitive Psychology* **60**:158–189. DOI: https://doi.org/10.1016/j.cogpsych.2009.12.001

Wagenmakers EJ, Farrell S. 2004. AIC model selection using akaike weights. *Psychonomic Bulletin & Review* **11**:192–196. DOI: https://doi.org/10.3758/BF03206482, PMID: 15117008

Warren CM, Wilson RC, van der Wee NJ, Giltay EJ, van Noorden MS, Cohen JD, Nieuwenhuis S. 2017. The effect of atomoxetine on random and directed exploration in humans. *PLOS ONE* **12**:e0176034. DOI: https://doi.org/10.1371/journal.pone.0176034

Watkins CJCH, Dayan P. 1992. Q-learning. *Machine Learning* **8**:279–292. DOI: https://doi.org/10.1007/BF00992698

Wiecki TV, Sofer I, Frank MJ. 2013. HDDM: hierarchical bayesian estimation of the Drift-Diffusion model in Python. *Frontiers in Neuroinformatics* **7**:14. DOI: https://doi.org/10.3389/fninf.2013.00014, PMID: 23935581

Wilson RC, Nassar MR, Gold JI. 2013. A mixture of delta-rules approximation to bayesian inference in change-point problems. *PLOS Computational Biology* **9**:e1003150. DOI: https://doi.org/10.1371/journal.pcbi.1003150, PMID: 23935472

Wilson RC, Niv Y. 2011. Inferring relevance in a changing world. *Frontiers in Human Neuroscience* **5**:189. DOI: https://doi.org/10.3389/fnhum.2011.00189, PMID: 22291631

Wilson RC, Niv Y. 2015. Is model fitting necessary for Model-Based fMRI? *PLOS Computational Biology* **11**:e1004237. DOI: https://doi.org/10.1371/journal.pcbi.1004237, PMID: 26086934

Wimmer GE, Li JK, Gorgolewski KJ, Poldrack RA. 2018. Reward learning over weeks versus minutes increases the neural representation of value in the human brain. *The Journal of Neuroscience* **38**:7649–7666. DOI: https://doi.org/10.1523/JNEUROSCI.0075-18.2018, PMID: 30061189

Zajkowski WK, Kossut M, Wilson RC. 2017. A causal role for right frontopolar cortex in directed, but not random, exploration. *eLife* **6**:e27430. DOI: https://doi.org/10.7554/eLife.27430, PMID: 28914605

Appendix 1

The theory of model fitting

Formally, the goal of model fitting is to estimate the parameters, θ_m, for each model, m, that best fit the behavioral data. To do this, we take a Bayesian approach and aim to compute (or at least approximate) the posterior distribution over the parameters given the data, $p(\theta_m|d_{1:T},m)$. By Bayes' rule we can write this as

$$p(\theta_m|d_{1:T},m) = \frac{p(d_{1:T}|\theta_m,m)p(\theta_m|m)}{p(d_{1:T}|m)} \tag{12}$$

where $p(\theta_m|m)$ is the prior on the parameters, θ_m; $p(d_{1:T}|\theta_m,m)$ is the likelihood of the data given the parameters; and the normalization constant, $p(d_{1:T}|m)$, is the probability of the data given the model (which is also known as the marginal likelihood [*Lee and Wagenmakers, 2014*], more on this below). Because the probabilities tend to be small, it is often easier to work with the log of these quantities

$$\log p(\theta_m|d_{1:T},m) = \underbrace{\log p(d_{1:T}|\theta_m,m)}_{\text{log likelihood}} + \log p(\theta_m|m) - \log p(d_{1:T}|m) \tag{13}$$

The log-likelihood often gets its own symbol, $LL = \log p(d_{1:T}|\theta_m,m)$, and can be written

$$\log p(d_{1:T}|\theta_m,m) = \log\left(\prod_{t=1}^{T} p(c_t|d_{1:t-1},s_t,\theta_m,m)\right) = \sum_{t=1}^{T} \log p(c_t|d_{1:t-1},s_t,\theta_m,m) \tag{14}$$

where $p(c_t|d_{1:t-1},s_t,\theta_m,m)$ is the probability of each individual choice given the parameters of the model, which is at the heart of the definition of each model (for example in *Equations 1– 7*).

In a perfect world, we would evaluate the log-posterior, $\log p(\theta_m|d_{1:T},m)$, exactly, but this can be difficult to compute and unwieldy to report. Instead, we must approximate it. This can be done using sampling approaches such as Markov Chain Monte Carlo approaches (*Lee and Wagenmakers, 2014*), which approximate the full posterior with a set of samples. Another approach is to report a point estimate for the parameters such as the maximum of the log-posterior (the maximum a posteriori [MAP] estimate), or the maximum of the log-likelihood (the maximum likelihood estimate [MLE]). (Note that the log transformation does not change the location of the maximum, so the maximum of the log-likelihood occurs at the same value of θ_m as the maximum of the likelihood.)

$$\hat{\theta}_m^{MAP} = *arg\,max_{\theta_m} \log p(\theta_m|d_{1:T},m) \hat{\theta}_m^{MLE} = *arg\,max_{\theta_m} \log p(d_{1:T}|\theta_m,m)$$

Note that with a uniform prior on θ_m, these two estimates coincide.

These approaches for estimating parameter values each have different strengths and weaknesses. The MCMC approach is the most principled as, with enough samples, it gives a good approximation of the posterior distribution over each parameter value. This approach also gracefully handles small data sets and allows us to combine data from different subjects in a rigorous manner. Despite these advantages, the MCMC approach is more complex (especially for beginners) and can be slow to implement. On the other hand, point estimates such as the MAP and MLE parameter values are much quicker to compute and often give similar answers to the MCMC approach when the amount of data is large (which is often the case when dealing with young and healthy populations). For this reason, we focus our discussion on the point estimate approaches, focusing in particular on maximum likelihood estimation.

Appendix 2

The theory of model comparison

In model comparison, our goal is to figure out which model of a set of possible models is most likely to have generated the data. To do this, we compute (or at least try to estimate) the probability that model m generated the data, $p(m|d_{1:T})$. Note that this is the normalization constant from *Equation 12*. As with parameter recovery, this probability is difficult to compute directly and so we turn to Bayes' rule and write

$$p(m|d_{1:T}) \propto p(d_{1:T}|m)p(m) = \int d\theta_m p(d_{1:T}|\theta_m, m)p(\theta_m|m)p(m)$$

where $p(m)$ is the prior probability that model m is the correct model and $p(d_{1:T}|m)$ is the likelihood of the data given the model. In most cases, $p(m)$ is assumed to be constant and so we can focus entirely on the likelihood, $p(d_{1:T}|m)$. As before, it is easier to handle the log of this quantity which is known as the marginal likelihood (*Lee and Wagenmakers, 2014*) or Bayesian evidence, E_m (*Kass and Raftery, 1995*), for model m. More explicitly

$$E_m = \log p(d_{1:T}|m) = \log \int d\theta_m p(d_{1:T}|\theta_m, m)p(\theta_m|m) \tag{17}$$

If we can compute E_m for each model, then the model with the largest evidence is most likely to have generated the data.

Note that by integrating over the parameter space, the Bayesian evidence implicitly penalizes free parameters. This is because, the more free parameters, the larger the size of the space over which we integrate and, consequently, the smaller $p(\theta_m|m)$ is for any given parameter setting. Thus, unless the model predicts the data well for all parameter settings, it pays a price for each additional free parameter. This idea, that simpler models should be favored over more complex models if they both explain the data equally well, is known as Occam's razor (see Chapter 28 in *MacKay, 2003*).

Unfortunately, because it involves computing an integral over all possible parameter settings, computing the marginal likelihood exactly is usually impossible. There are several methods for approximating the integral based on either replacing it with a sum over a subset of points (*Wagenmakers et al., 2010*; *Lee and Wagenmakers, 2014*) or replacing it with an approximation around either the MAP or MLE estimates of the parameters. The latter approach is the most common and three particular forms are used: the Bayes Information Criterion (BIC) (*Schwarz, 1978*), Akaike information criterion (AIC) (*Akaike, 1974*) and the Laplace approximation (*Kass and Raftery, 1995*). Here, we will focus on BIC which is an estimate based around the maximum likelihood estimate of the parameters, $\hat{\theta}_m^{MLE}$,

$$BIC = -2\log\hat{\mathcal{L}} + k_m \log(T) \approx -2\log E_m \tag{18}$$

where k_m is the number of parameters in model m and $\hat{\mathcal{L}}$ is the value of the log-likelihood at $\hat{\theta}_m^{MLE}$.

Finally, we have found it useful to report the results of model comparison in terms of the likelihood-per-trial LPT, which can be thought of as the 'average' probability with which the model predicts each choice,

$$LPT = \exp\left(\frac{E_m}{T}\right) \tag{19}$$

Appendix 3

Computing the inversion matrix from the confusion matrix

In section 'Can you arbitrate between different models?' we introduced the inversion matrix, $p(\text{simulated model}|\text{fit model})$, as the probability that data best fit by one model were actually generated from another model. As shown below, this can be readily computed from the confusion matrix, $p(\text{fit model}|\text{simulated model})$, by Bayes rule. Abbreviating 'simulated model' with 'sim' and 'fit model' with 'fit' we have

$$p(\text{sim}|\text{fit}) = \frac{p(\text{fit}|\text{sim})p(\text{sim})}{\sum_{\text{sim}} p(\text{fit}|\text{sim})p(\text{sim})} \qquad (20)$$

For a uniform prior on models, computing the inversion matrix amounts to renormalizing the confusion matrix over the simulated models.

Working memory model used for local minima example

The model and experimental designs used in *Box 3—figure 1* are a simplified version of those in *Collins and Frank (2012)*. In short, the experiment attempts to parse out working memory contributions to reinforcement learning by having participants and agents learn stimulus-action contingencies from deterministic feedback, with a different number of stimuli ns being learned in parallel in different blocks. This manipulation targets WM load and isolates WM contributions; see *Collins and Frank (2012)* for details.

The simplified model assumes a mixture of a classic RL component (with parameters α and β) and a working memory component with perfect one-shot learning. The mixture is controlled by parameter ρ, capturing the prior willingness to use working memory vs. RL, and capacity parameter K, which scales the mixture weight in proportion to the proportion of stimuli that may be held in working memory: $\min(1, \frac{K}{ns})$. The original model assumes additional dynamics for the working memory policy and working memory vs. RL weights that render the model more identifiable (*Collins and Frank, 2012*).

Appendix 5

Model validation example

In this example, we imagine a deterministic stimulus-action learning task in which subjects are presented with one of three stimuli (s_1, s_2, and s_3), which instruct the subject which of three actions (a_1, a_2, and a_3) will be rewarded when chosen.

The two models that we consider are both reinforcement learning agents. The first, a 'blind' agent, does not see the stimulus at all and learns only about the value of the three different actions, i.e. $Q(a_i)$, regardless of the stimulus. The second, a 'state-based' agent, observes the stimulus and learns a value for each action that can be different for each stimulus, i.e. $Q(a_i, s_i)$.

Learning in both models occurs via a Rescorla-Wagner rule with different learning rates for positive and negative prediction errors. Thus for the blind agent, the values update according to

$$Q(a_t) \leftarrow \begin{cases} Q(a_t) + \alpha_P(r_t - Q(a_t)) & \text{if } (r_t - Q(a_t)) > 0 \\ Q(a_t) + \alpha_N(r_t - Q(a_t)) & \text{if } (r_t - Q(a_t)) < 0 \end{cases} \tag{21}$$

while for the state-based agent, values update according to

$$Q(a_t, s_t) \leftarrow \begin{cases} Q(a_t, s_t) + \alpha_P(r_t - Q(a_t, s_t)) & \text{if } (r_t - Q(a_t, s_t)) > 0 \\ Q(a_t, s_t) + \alpha_N(r_t - Q(a_t, s_t)) & \text{if } (r_t - Q(a_t, s_t)) < 0 \end{cases} \tag{22}$$

In both models, these values guide decisions via softmax decision rule with inverse temperature parameter, β.

We begin by simulating two different agents: one using the blind algorithm and the other using the state-based approach. Parameters in the models are set such that the learning curves for the two agents are approximately equal (***Box 7—figure 1A***, blind model: $\alpha_P = 0.5$, $\alpha_N = 0$, $\beta = 6.5$ state-based model: $\alpha_P = 0.65$, $\alpha_N = 0$, $\beta = 2$). In both cases, the agents start from an accuracy of 1/3 and an asymptote at an accuracy of around 2/3 — the blind agent because this is the best it can do, the state-based agent because the softmax parameter is relatively small and hence performance is limited by noise.

Next we consider how the state-based model fits behavior from these two different agents. In ***Box 7—figure 1B***, we plot the average likelihood with which the state-based model predicts the actual choices of the blind and state-based agents, that is the average $p(c_t|d_{1:t-1}, \theta_m, m = \text{state-based})$. As is clear from this figure, the state-based model predicts choices from the blind agent with *higher* likelihood than choices from the state-based agent! Although counter intuitive, this result does not imply that the state-based model is unable to fit its own behavior. Instead, this result reflects the difference in noise (softmax parameters) between the two subjects. The blind RL subject has a high β, implying less noise, allowing the state-based model to fit it quite well. Conversely, the state-based RL subject has a low β, implying more noise, meaning that the behavior is harder to predict even when it is fit with the correct model.

That the state-based model fits state-based behavior better than it fits blind behavior is illustrated in ***Box 7—figure 1C***. Here we plot the simulated learning curves of the state-based model using the parameter values that were fit to either the state-based agent or the blind agent. While the fit to the state-based agent generates a learning curve quite similar to that of the subject (compare blue lines in ***Box 7—figure 1A and C***), the state-based fit to the blind agent performs too well (compare yellow lines in panels ***Box 7—figure 1A and C***). Thus the model validation step provides support for the state-based model when it is the correct model of behavior, but rules out the state-based model when the generating model was different.