Theoretically Informed Generative Models Can Advance the Psychological and Brain Sciences: Lessons from the Reliability Paradox

Nathaniel Haines*

The Ohio State University, Department of Psychology,
1835 Neil Ave., Columbus, OH., 43210
Email: haines.175@osu.edu; Website: <http://haines-lab.com/>

Peter D. Kvam

University of Florida, Department of Psychology,
945 Center Dr, Gainesville, FL, 32611
Email: pkvam@ufl.edu; Website: <https://peterkvam.com/>

Louis Irving

University of Florida, Department of Psychology,
945 Center Dr, Gainesville, FL, 32611
Email: louis.irving@ufl.edu; Website: <https://theapclab.wordpress.com/people/>

Colin Tucker Smith

University of Florida, Department of Psychology,
945 Center Dr, Gainesville, FL, 32611
Email: colinsmith@ufl.edu; Website: <https://theapclab.wordpress.com/>

Theodore P. Beauchaine

The Ohio State University, Department of Psychology,
1835 Neil Ave., Columbus, OH., 43210
Email: beauchaine.1@osu.edu; Website: <https://tpb.psy.ohio-state.edu/LAP/people.html>

Mark A. Pitt

The Ohio State University, Department of Psychology,
1835 Neil Ave., Columbus, OH., 43210
Email: pitt.2@osu.edu; Website: <https://u.osu.edu/markpitt/>

Woo-Young Ahn

Seoul National University, Department of Psychology,
1 Gwanak-ro, Gwanak-gu, Seoul, South Korea
Email: wahn55@snu.ac.kr; Website: <https://ccs-lab.github.io/>

Brandon M. Turner*

The Ohio State University, Department of Psychology,
1835 Neil Ave., Columbus, OH., 43210
Email: turner.826@gmail.com; Website: <https://turner-mbcn.com/>

*Co-corresponding authors

Abstract

Theories of individual differences are foundational to psychological and brain sciences, yet they are traditionally developed and tested using superficial summaries of data (e.g., mean response times) that are both (1) disconnected from our otherwise rich conceptual theories of behavior, and (2) contaminated with measurement error. Traditional approaches therefore lack the flexibility required to test increasingly complex theories of behavior. To resolve this theory-description gap, we present the *generative modeling approach*, which involves formally specifying how behavior is generated within people and how generative processes vary across people. Generative modeling shifts our focus away from estimating descriptive statistical “effects” toward estimating psychologically interpretable parameters, while simultaneously enhancing the reliability and validity of our measures. We demonstrate the utility of generative models in the context of the “reliability paradox”, a phenomenon wherein highly replicable group effects (e.g., Stroop effect) fail to capture individual differences (e.g., low test-retest reliability). Simulations and empirical data from the Implicit Association Test, and Stroop, Flanker, Posner, and Delay Discounting tasks show that generative models yield (1) more theoretically informative parameters, and (2) higher test-retest reliability estimates relative to traditional approaches, illustrating their potential for enhancing theory development.

Keywords: Bayesian analysis, implicit attitudes, impulsivity, individual differences, generative modeling, measurement error, reliability, self-control, theory development

Introduction

Across the social, behavioral, and brain sciences, the researcher's primary agenda is to provide an explanation for *why* observable measures (i.e., data) exhibit systematic patterns¹. Typically, this process begins with the development of a theory or hypothesis about what should happen in specific situations (e.g., experiments) if our assumptions about the phenomena are correct. To test our theory, we design an environment that places the agent under consideration (e.g., human, rat) into a specific situation, and we collect data that describe the agent's experience (e.g., decisions, neural activity). To evaluate the fidelity of our hypothesis, we use statistical models such as *t*-tests, ANOVAs, multiple (logistic) regression, and factor analytic models, to connect our explanatory theories to the data observed in our experiment²⁻⁴. We use these statistical models to make inferences, such as inferring population-level effects or estimating out-of-sample predictive accuracy using in-sample data, and relate these inferences to claims about theory. However, underlying our statistical models are causal or distributional assumptions that are too often mis-aligned with the substantive theories of interest, a situation we refer to as a *theory-description gap*. When assumptions are misaligned, theories become “divorced” from the statistical tests meant to validate or invalidate them, and the disconnect impedes progress in science⁵⁻⁷.

In this article, we demonstrate how *generative modeling*—a theory-driven data analysis approach where formal model specification is guided by hypotheses about latent processes—can alleviate the pressures associated with theory-description gaps. The goal of this demonstration is twofold. First, we show that the inferences made with the traditional statistical analyses applied in behavioral and brain sciences generate misleading conclusions about the utility of behavioral tasks. Second, we show that this problem can be addressed using modern computational

modeling and hierarchical Bayesian model fitting methods to analyze behavioral data. Put together, this paper is an appeal for wider use of these more sophisticated statistical methods, showing that the theoretical advantages of generative modeling approaches translate directly to practical benefits in terms of improved measurement properties like reliability and validity.

Specifically, we focus on how generative modeling resolves a vexing theory-description gap that is currently shaking the foundations of research on individual differences: the *reliability paradox*. The reliability paradox refers to the counterintuitive result that person-level measures of behavior across a broad range of tasks (e.g., the Implicit Association Test, Stroop, Flanker, and Posner Cueing tasks) and modalities (e.g., accuracy, response times, task-based and resting-state fMRI) show poor test-retest reliability despite task manipulations showing consistent and robust effects at the group level⁸⁻¹³. The reliability paradox is a major challenge to behavioral research because low test-retest reliability implies that the psychological constructs on which theories are based are unstable across time, making the task or modality untrustworthy for validating theories based on temporally stable individual differences^{9,14-16}. As we demonstrate, however, the reliability paradox arises from the implicit, overly restrictive assumptions that researchers make when using standardized statistical models that fail to connect their theories to data in *meaningful* ways. To illustrate how *generative models* can overcome the reliability paradox, we use response time measures as our observable data as an in-depth example. However, we additionally apply this approach to response / preference data from the Delay Discounting task, showing that the principles and benefits of generative modeling are much more general and apply to any measure extracted from data (e.g., accuracy).

As a concrete example, consider the Stroop task¹⁷ which asks participants to indicate either the word or text color of a word presented on a screen. Words and colors can be either congruent

(“red” in red text), neutral (“boat” in red text), or incongruent (“blue” in red text). In a typical experiment using the Stroop task, our observable measures are distributions of responses and response times for each person, each condition, and each perhaps session in the case of test-retest measurements. As illustrated in Figure 1, the “Stroop effect” is typically estimated for each person as the difference in mean response times (across trials) between incongruent and congruent trials, such that a positive value indicates that a person takes longer (on average) to respond to incongruent relative to congruent trials. This difference in mean response time serves as what we call the *behavioral model*, which is the mathematical model used to make inferences about individual differences in latent cognitive processes. In the case of the Stroop task, the mean difference is supposed to measure lexical interference, under the assumption that slowing in the incongruent condition is a measurement of the degree of interference induced by the word-color conflict.

In this view, the behavioral model defines the effect of interest to the extent that it can be measured with response times, and thus the reliability of the effect is determined by the reliability of the metric produced by the behavioral model (mean response time gap). To estimate test-retest reliability, the person-level Stroop effect values from one session are used as dependent variables in a secondary statistical model predicting the same participants’ Stroop effect values from a later session. If a Stroop effect (and thus the inferred degree of interference between lexical and visual processing) is stable, we should be able to use each person’s Stroop effect at one point in time to accurately predict their Stroop effect at a later point in time. Yet if such an analysis shows that our behavioral measures have low test-retest correlations—as revealed by the reliability paradox—it is difficult to imagine how we could use them to make

inference on other individual differences relevant for testing our theories (e.g., estimating correlations between behavioral and neural or self-report measures).

The method of creating a behavioral model therefore determines not only the inferences we make about the tasks and their reliability, but the inferences we make about latent constructs like lexical interference. For example, if we find low correlations between a Stroop effect on the task and a self-report measure of lexical processing (e.g., dyslexia scale), we might be inclined to conclude that they measure separate cognitive or neural processes. Such findings (and conclusions) are appearing more frequently—both the literature on impulsivity¹⁸⁻¹⁹ and risky choice²⁰⁻²² have found low reliability of behavioral measures and weak correlation to neural and self-report measures, leading them to conclude that impulsivity and risk-taking each have multidimensional factor structures. Although these conclusions may ultimately hold, they are drawn on the basis of behavioral models that define psychological constructs in terms of simple summary statistics. As we show in this paper, the degree to which these types of conclusions may be drawn rests on assumptions about these simple behavioral models and our ability to detect correlations across time and across measures. As we show below, however, this approach to analyzing behavioral data is flawed and is likely to have resulted in many incorrect, overconfident conclusions about the status of psychological constructs drawn from behavior.

Altogether, we refer to the aforementioned approach to statistical inference as the “two-stage summary approach”, given that summary statistics are estimated in one stage and then an individual difference correlation (i.e. test-retest) is estimated independently in a second stage.

Figure 1. Pathway from theory to inference with behavioral data.

Behavioral tasks are designed to test substantive theories by relating the theoretical predictions to features of observed behavior within individual participants. Figure 1 shows the “behavioral model” often assumed when analyzing Stroop data in the “Two-Stage Approach” column, which makes implicit assumptions about how response times arise that we describe in the next section. The group-level model captures variation in person-level behavioral model parameters (e.g., mean response time). Here, we illustrate a group-level model estimating test-retest reliability. The traditional two-stage summary approach treats these models in a contiguous way, estimating person-level behavioral model parameters (e.g., mean response time contrasts) and then entering the resulting point estimates into a secondary model to assess group differences, individual difference correlations (i.e. test re-test), etc. Conversely, with generative modeling we construct a single model that integrates our assumptions and hypotheses about the data generating process, spanning distributions of trial-by-trial response times within persons to the distributions of individual differences across people. This example depicts a lognormal behavioral model, which is comprised of two parameters that capture the shape of a person’s response time distribution (described in detail in the next section). Person-level parameters follow a group-level distribution that captures dependence in person-level parameters across test-retest sessions, where parameter estimates for each session and their relationship to one another are all estimated simultaneously by estimating the variance-covariance matrix among them. This approach allows information to flow between the person- and group-level models when estimating parameters, which leads to more accurate estimation of person- and group-level parameters relative to the two-stage summary approach (e.g., mean response times, individual difference correlations). Better parameter estimation allows for us to test theories with more fidelity.

The two-stage summary approach makes two strong yet inappropriate assumptions about within-person behavior that invalidate the test-retest correlations that are derived from it. Although we provide more substantive mathematical justification in the Methods section, the intuition is as follows: When we compute mean response times in task conditions, take their difference, and then input these difference scores into a secondary statistical model (*t*-test, etc.), we are making the (first) assumption that the first-stage estimates (difference scores in this case) are estimated with no measurement error, or equivalently that either (1) we have an *infinite number of trials* within each condition used to estimate the mean response time difference, or (2) response times within persons show no variation from trial-to-trial²³⁻²⁶. The same critique applies any time we summarize our data before entering the resulting summaries into a secondary model, irrespective of whether these summaries are sample means or point estimates derived from a classical, machine learning, or even Bayesian models. The key point is that when we ignore the fact that *summary statistics are estimates*, and not observed measures, we ignore sources of uncertainty at one level of analysis that should be incorporated into our secondary statistical models²⁷⁻³². In the case of test-retest analyses, ignoring uncertainty dramatically attenuates test-retest correlations, pushing them toward zero. In such cases, researchers understandably (yet incorrectly) conclude that the effect of interest is inconsistent across repetitions, even when the “true” effect is stable and consistent, as we demonstrate below.

The second problematic assumption is that the sample mean of a response time distribution (or choice proportions if measuring accuracy/preference) is an accurate way of connecting the substantive theory under investigation to observed data. In fact, the use of sample means in this way assumes a specific generative model of within-person behavior—namely, that response times within each person and condition arise from a normal (i.e., Gaussian) distribution (see

Methods section). Researchers may be unaware that they assume this *implicit generative model* when using the two-stage summary approaches to assess test-retest reliability (or other individual difference correlations), and as a result they may erroneously infer that tasks or measurement modalities, rather than behavioral models, confer low test-retest reliability. To the extent that this assumption is inconsistent with empirical data, we will fail to detect individual differences that are otherwise readily apparent in empirical data. For example, Figure 2 shows that many different distributions, the shapes of which a theory must explain, can all have the same mean. Without constructing generative models of how individual people behave, we fail to capture important aspects of observed data—including variance³³, bimodality³⁴, and skew³⁵⁻³⁶—which may otherwise invalidate the conclusions we draw from our mis-specified statistical models. Indeed, a history of examples ranging from basic research in cognitive science, to more policy-level research on eyewitness detection accuracy and racial bias in police shootings shows that *atheoretical* use of summary statistics (per the implicit generative models that they assume) can even lead to conclusions that are opposite of those derived from more theoretically-motivated generative models^{3,37-41}. Notably, these shortcomings apply *irrespective to the replicability of the results* obtained using inappropriate or otherwise mis-specified statistical models⁴²⁻⁴⁴. In Supplementary Note 1, we provide a simulation that demonstrates the importance of distributional information on estimating stability in person-level behavior across time (i.e., test-retest reliability).

Figure 2. Qualitatively different distributions with the same mean.

All five distributions have the same mean (indicated by the dotted gray line) and would therefore produce the same conclusions if analyzed using only the sample mean summary statistic, regardless of how different their data-generating processes are. Statistical models that fail to fully account for the information contained in the full distribution of observed data fall victim to the theory-description gap.

In what remains of the paper, we test the hypothesis that deficiencies in behaviorally-defined constructs and their test-retest reliability can be attributed to the use of means and two-stage approach. We do so by evaluating the same data using these traditional approaches and contrasting them with generative models of each phenomenon. Generative modeling is a framework rooted in mathematical psychology and computational cognitive (neuro)science that resolves these issues by constructing statistical models that directly instantiate theoretical assumptions about how observed data are generated^{2,25,45-54}. To build a generative model, we begin with a different *behavioral model*. The behavioral model captures our assumptions about how individual people generate data and consequently how these data are distributed — in this case, a person’s distribution of response times across trials (and condition type, etc.). Our choice of behavioral model should be informed by theoretical assumptions and other background knowledge. For response time data, we know that response time distributions (1) must take on only positive real values, (2) have some central tendency with corresponding variability around this central tendency (e.g., a mean and variance), (3) are often right skewed, (4) are shifted away from 0 due to “non-decision” factors (e.g., visual encoding time, motor response time), and (5) typically show a linear relation the sample mean and standard deviation such that increases in mean response time are accompanied by increased variability of response times from trial-to-trial (the *law of response time*)⁵⁵.

There are many different distributions that can incorporate our knowledge about response times, and thus the process(es) that generated them. Here, we limit our discussion to the normal, lognormal, and shifted-lognormal distributions. We first show how the parameters of the lognormal models produce changes in predicted response time distributions (Figure 3). Figure 4 shows how the two-stage summary approach can be viewed as a highly constrained version of

the normal generative model, in addition to how the two lognormal models further improve upon shortcomings of the normal generative model. Given that these behavioral models can better characterize the full distribution of person-level response times, we can be more confident in using model parameters to infer meaningful consistencies (or inconsistencies) in behavior. In fact, it follows from the Fisher-Neyman factorization theorem that, if our assumed generative model is sufficient, variation in estimated model parameters will be *lower than or equal to* variation in summary statistics extracted from the same data across different conditions (in our case, test-retest sessions)⁵⁶⁻⁵⁸. Intuitively, because behavior is inherently probabilistic, observed response times offer only limited information regarding each person's data-generating parameters (e.g., μ and σ). So long as we use summary statistics or generative models that are not sufficient to fully capture the data-generating process (e.g., sample means or the normal generative model that misses the fact that response times are highly skewed), we ignore information that could otherwise inform or constrain parameter estimation. An implication is that, to the extent that we have more mechanistic behavioral models informed by the theory under investigation (e.g., computational models; see^{2,54,59}), they can be used to help us constrain inference in a way not afforded by the traditional two-stage summary approach.

Figure 3. Lognormal and shifted lognormal generative distributions.

(A) For the lognormal and shifted lognormal distributions, changes in the μ parameter (interpreted as “stimulus difficulty”) produce changes in both means and variances of response time distributions⁶⁰⁻⁶¹. (B) The σ parameter controls dispersion (interpreted as “decision noise”); changes in σ affect means and ranges of likely response times, but medians remain constant. (C) For the shifted lognormal distribution, the shift parameter δ translates the lognormal distribution forward in time without changing the shape of the response time distribution. The lognormal distribution is a special case of the shifted lognormal distribution wherein the shift parameter (δ) is set to 0. The shift parameter is interpreted as “non-decision time”, as it can capture individual differences in components of the data-generating process that are not relevant to decision-making but nevertheless produce variation in mean response times (e.g., visual encoding time, motor responses).

Figure 4. Building generative models consistent with theory.

The two-stage summary statistic approach is often used by default and is chosen without reference to an underlying theory. However, it is not a “safe” default—the implicit generative model it assumes is highly constrained and most likely inappropriate, assuming that either the within-person sample size (number of trials) is infinite/arbitrarily large, or that the within-person standard deviation of response times is approximately zero (the latter is depicted here; see the Methods section and Supplementary Note 2 for details). In the case of the Stroop task, these assumptions respectively correspond to us collecting an infinite number of responses in both conditions before taking the mean difference in response times, or alternatively, that every response time is exactly the same within conditions. Of course, we know *a priori* that neither of these assumptions are met, and so the generative approach begins by relaxing these overly restrictive assumptions. We begin with a simple normal model of response times, wherein we estimate both a mean and standard deviation within persons and conditions. However, the normal model predicts negative response times and does not capture the skew readily apparent in empirical response times distributions. The lognormal model improves upon these shortcomings, yet it cannot capture “non-decision” factors (e.g., stimulus encoding or motor response time) that could bias parameters of interest. The shifted lognormal model retains key properties of the lognormal model, but can also capture non-decision factors with the shift parameter (which is constrained to be equal across conditions within persons; see the Methods section for an explanation). In constructing these models, we reduce the theory-description gap by creating models that are more consistent with our theoretical and background knowledge.

Naturally, the shifted lognormal will not be the best generative model for all tasks. By its nature, a one-size-fits-all solution to behavioral modeling will not capture the minutiae of performance on all of the tasks we examine here. The goal instead is to examine the degree to which we can improve upon traditional two-stage procedures by incorporating (ostensibly) more valid assumptions about the behavior we are trying to model. The shifted log-normal provides a statistical distribution that reflects several realistic assumptions we can make: that participants will have different average processing speeds, that response times will be right-skewed and vary across trials, and that some component of response times can be attributed to processes that are not directly related to decision making. For specific tasks, we could build upon these assumptions by adding an account of accuracy (as with diffusion / accumulator models), across-trial variability in stimulus effects, collapsing response boundaries, or across trial shifts in attention, for instance. However, doing so treads out of clear-cut assumptions that should obviously be made a priori and into the weeds of model comparisons among many theories with plausible assumptions. To avoid this complication, we stop at the shifted lognormal as a simple model embodying clear-cut, desirable properties of a generative account of response times.

With the behavioral model defined, the next step is to estimate each person's behavioral model parameters. A history of work in both mathematical statistics and psychometrics shows that we can obtain more precise estimates of the “true” person-level parameters by pooling information across people using hierarchical models that assume a group-level distribution over person-level parameters⁶²⁻⁶⁴. In our case, uncertainty arises from both the inherent probabilistic nature of behavior (i.e. variation in response times within persons, or σ), and the limited number of trials that we observe within persons/conditions. Therefore, by defining an explicit group-level generative model of the person-level parameters themselves (as illustrated in the “Generative Model” row of Figure 1), the group-level model can pool information across people while accounting for uncertainty at the person level that would otherwise produce biased, attenuated individual differences correlations using the two-stage summary approach. In fact, the test-retest correlation obtained from a full generative model spanning from within- to between-person variation is akin to the “true” or “dis-attenuated” correlation that one obtains when using classical corrections for attenuation due to unreliability^{24,64}, although the generative formulation allows for much more flexibility in our choice of both group- and person-level models. Here, we assume that person-level parameters follow correlated (i.e. multivariate) normal distributions, and we use this same specification for all three behavioral models (see the Methods section for details and visual depictions).

To demonstrate the power of generative modeling in closing the theory-description gap in the context of the reliability paradox, we re-analyzed data from three prior studies. We compare results obtained from the two-stage summary approach with those using the three generative models (Figure 4) across five common behavioral tasks. We analyze data from the Stroop, Flanker, and Posner Cueing tasks¹², the Self-Concept (introversion/extraversion) and Race

(Black/White) versions of the Implicit Association Test (IAT)¹¹, and the Delay Discounting task⁶⁵. We include the Delay Discounting dataset to demonstrate that benefits of generative modeling extend beyond response time data. Individually, each of these tasks has yielded a deep body of literature—the Stroop, Flanker, and Posner Cueing tasks have been used extensively to develop theories of attention and inhibitory control, the IAT has been used to develop theories of implicit cognition and evaluations with deep societal implications, and the Delay Discounting task has been used to develop theories of impulsivity and self-control. On Google Scholar alone (as of April 2021), the collective citation count of the original papers pertaining to these tasks is over 57,000^{17,66-70}. Moreover, the five tasks cover areas of research spanning from psychology to neuroscience to behavioral economics.

Results

We discuss analysis of the Stroop data in-depth, followed by a summary of the results from the other tasks (see Supplementary Note 3 for an expanded discussion of each)

Results for the Stroop task in Study 1 of Hedge et al.¹² are shown in Figure 5. Panel A compares the estimated test-retest correlation for the two-stage approach versus each of the normal, lognormal, and shifted lognormal generative models. For the two-stage mean and standard deviation contrasts, test-retest correlations were $r = .50$ (95% CI = [.25, .69]) and $r = .07$ (95% CI = [-.22, .35]), respectively. Note that previous studies did not include standard deviation contrasts, but we include them here for comparison to the generative models. These estimates are consistent with results obtained originally by Hedge et al.¹². It is clear that the Stroop effect has less than ideal reliability when estimated using the two-stage approach: with a test-retest reliability of $r = .50$ to $r = .60$, we would need well over 200 participants to detect (with adequate power) a simple correlation between the Stroop effect and an alternative

individual difference measure with similar reliability¹². Given constraints on the number of participants available for any one design, low reliability limits the utility of the Stroop effect as a measure to advance theories that make predictions based on individual differences.

Figure 5. Test-retest correlations and model misfit for the Stroop task.

(A) Posterior distributions (red) for the test-retest correlations of each of the three generative models versus the two-stage sample mean/standard deviation approach (vertical dotted black line with corresponding horizontal 95% confidence interval) for the Stroop task in Study 1 of Hedge et al.¹². (B) Posterior predictive simulations and sample means (vertical dotted black lines) for each of the generative models for a representative participant. The lognormal models can reproduce person-level response time distributions with much more fidelity than the normal model (and subsequently, the two-stage summary approach).

Figure 5A shows *posterior probability distributions* of the model parameters obtained using Bayesian updating as opposed to the point estimates and confidence intervals obtained using the two-stage summary approach. Note that the posterior distribution can be interpreted in a variety of ways depending on our goals. For example, if one is interested in the probability that the test-retest correlation of the normal generative model is greater than the two-stage estimate of $r = .50$, this quantity can be easily computed as the proportion of the posterior distribution greater than $r = .50$. Alternatively, if we are interested in the single most likely test-retest estimate, we can simply locate the mode (or peak) of the posterior distribution. However, we are typically interested not only in a single value, such as the mode, but a range of likely values that help us convey uncertainty. Therefore, to facilitate interpretability of posterior distributions, we report the posterior mean (sometimes referred to as the posterior “expectation”) along with the 95% *highest density interval* (HDI). An HDI is a generalization of the concept of the mode, but it is an interval rather than a single value. For example, a 20% HDI would contain 20% of the area of the entire posterior distribution, where every value within the interval is more likely than every value outside of the interval. We report 95% HDIs to maintain consistency with the 95% CIs reported for the two-stage approach, although we caution readers that HDIs and CIs are different concepts that have different interpretations. As has been a focus throughout this article, a mean and CI alone may do a poor job of summarizing a skewed distribution, so we recommend that readers interpret the posterior distributions holistically to fully appreciate the generative model estimates.

For the generative models, the posterior distributions for the mean/difficulty contrast parameters ($\mu_{i,\Delta}$) across models were concentrated above the two-stage estimates (posterior mean test-retest ranging from $r = .76$ to $r = .81$). Furthermore, the 95% HDIs for the difficulty parameter in the normal (95% HDI = [.46, 1.00]), lognormal (95% HDI = [.47, 1.00]), and

shifted-lognormal (95% HDI = [.53, 1.00]) models included $r = 1.00$, indicating that we cannot rule out the possibility that there is in fact a perfect correlation in the mean/difficulty parameter contrast between retest sessions. This can be observed in the posterior distributions, which are concentrated against the upper limit of the range at $r = 1.00$. Posterior distributions for the standard deviation/dispersion parameters ($\sigma_{i,\Delta}$) were also concentrated above the two-stage estimates, although primarily for the lognormal and shifted lognormal models (posterior mean test-retest ranging from $r = .23$ to $r = .62$). In fact, the test-retest estimate for the standard deviation/dispersion parameters were much higher for the lognormal (95% HDI = [.26, .89]) and shifted-lognormal (95% HDI = [.25, .96]) models relative to the normal model (95% HDI = [-.05, .50]), which demonstrates the importance of our data-generating (distributional) assumptions when making inference on individual differences; it is clear that the assumption of within-person variation in response times is a crucial factor for obtaining more reliable estimates of the mean/central tendency of response times (refer back to Figure 3). Furthermore, this variation itself can be a stable psychological property depending on its assumed functional form (i.e. normal, lognormal, shifted lognormal). Our theories of inhibitory control should therefore account for not only the mean, but also the variability in response times across trials. More generally, our behavioral models should account for the full distribution of observed data. Generative modeling is a principled method of closing these theory-description gaps.

We can also compare the person-level parameters across models to determine if the models produce different mechanistic inferences. For example, we may be interested in the proportion of participants who show a “Stroop effect” for each model. For demonstration, here we define an effect as when the 95% HDI of the person-level posterior distribution on the contrast parameter of interest is entirely above 0. We can then identify the proportion of participants meeting this

criterion for each of the $\mu_{i,\Delta}$ and $\sigma_{i,\Delta}$ parameters. For Session 1, across all generative models, all 47 participants showed evidence for an increase in $\mu_{i,\Delta}$ in the incongruent condition. However, for $\sigma_{i,\Delta}$, 36, 27, and 20 participants showed evidence for an increase in the incongruent condition according to the normal, lognormal, and shifted-lognormal models, respectively. Results were similar for Session 2: across all models $\mu_{i,\Delta}$ showed evidence for an increase across conditions for all participants. Similarly, for $\sigma_{i,\Delta}$, 27, 24, and 18 participants showed an increase in dispersion across conditions for the normal, lognormal, and shifted lognormal models, respectively. This pattern of results suggests that changes in response times across conditions within participants may be attributable primarily to changes in $\mu_{i,\Delta}$ (difficulty) rather than $\sigma_{i,\Delta}$ (dispersion)—an inference facilitated by the lognormal models that provides clear motivation for extending the model (e.g., developing a more explicit mechanism of how $\mu_{i,\Delta}$ varies as a function of stimulus properties). Note that the addition of a shift parameter led to fewer participants showing evidence for a change in dispersion ($\sigma_{i,\Delta}$) across conditions, which emphasizes the importance of explicitly accounting for nuisance factors (i.e. non-decision time) when our goal is to interpret model parameters in psychologically meaningful ways. Without the shift parameter, we would have incorrectly concluded that the majority of participants showed an increase in dispersion within both Sessions 1 and 2, which has implications for our theories of stimulus interference and inhibitory control.

Figure 5B shows fitted model predictions compared to the observed response times for a random, representative participant. The two-stage approach is represented simply as the mean response time within each of the congruent and incongruent conditions, whereas the generative model predictions are represented by the light red curves, which are response time distributions

simulated from this participant's estimated person-level normal, lognormal, and shifted-lognormal model parameters, where variation between lines indicates uncertainty in underlying parameters. With these simulated response times, we can compare how well each model reproduces the observed response times. For this particular participant, the normal generative model reveals many shortcomings, the most obvious being failure to capture right-skew, and over-prediction of rapid response times. By contrast, the lognormal model in the middle panel provides a much better reproduction of the observed data, capturing both right-skew and the concentration of response times around the mean. Improvement offered by the shifted-lognormal model is more subtle in this example—it better captures the onset of the response time distribution (i.e. the most rapid response times) relative to the lognormal model due to the small shift, but otherwise performs similarly (Supplementary Note 6, Figures S3-S6, contains examples of where the shift makes a more noticeable difference). Note that improvement in model fit is accompanied by an increase in expected test-retest reliability for the lognormal models over the normal model, particularly for dispersion parameters. Indeed, as we described above, the normal model suggests that a large majority of participants showed increased dispersion ($\sigma_{i,\Delta}$) across conditions in Session 1, and then much fewer in Session 2. Given that the normal model poorly characterizes observed data, we may attribute this inconsistency (which was not observed to such an extent for the lognormal models) to a theory-description gap rather than to an actual change in how the experimental manipulation affected underlying cognitive processes between sessions (refer back to Figure 3). Overall, scrutinizing person-level model predictions versus observed data contextualizes group-level results and provides insight into the properties of cognitive mechanisms involved in inhibitory control.

Figure 6 visualizes test-retest correlations for the remaining tasks. We include detailed results and figures (akin to Figure 5) for each of these tasks in Supplementary Note 6, Figures S3-S6, along with a table with the descriptive statistics pertaining to Figure 6 (Supplementary Table 1). However, we emphasize that each of these tasks can be explored in the same or greater depth as we explored Stroop Study 1 above—in all cases, generative models readily unveil richness underlying observed data that is otherwise ignored by the coarse nature of the two-stage summary approach.

Figure 6. Test-retest correlations for all tasks and models.

Here we show means and 95% confidence intervals for the two-stage summary approach (for both sample mean and standard deviation contrasts) in black, along with the posterior means and 95% highest density intervals for the generative model parameter estimates (in various shades of red). The Implicit Association Test (IAT) datasets are from the Self-Concept (introversion/extraversion; Study 1a) and Race (Black/White; Study 2b) versions. The most striking increase in test-retest reliability is observed for μ , which is consistently higher than two-stage summary approach estimates across tasks (with an expected test-retest correlation of approximately .75 or greater for most tasks/models). In contrast, the test-retest reliabilities for σ are higher for some, but not all tasks (but note that estimation of σ is in part what is responsible for better estimation of μ ; see Methods section and Figure 4). Test-retest windows for each task were as follows: (1) 3 weeks for the Stroop, Flanker, and Posner Cueing tasks, (2) 8 weeks for both versions of the IAT, and (3) 4 weeks for the Delay Discounting task. See Supplementary Table 1 and the Supplementary Notes 5 and 6 for more detailed figures and full descriptions of each task.

There are three main take-aways from the results presented in Figure 6. First, generative models consistently infer higher test-retest correlations relative to the two-stage approach. In many cases these changes are quite substantial. For example, in Study 2 of the Flanker task, the test-retest correlation of the sample mean obtained from the two-stage summary approach was quite low, $r = -.13$, whereas the normal generative model inferred $r = .64$. For the IAT Race version, the two-stage sample mean contrast test-retest correlation was $r = .45$, whereas the normal generative model inferred $r = .83$. Such large differences have considerable implications for testing and developing theories of individual differences within each paradigm. Indeed, low test-retest correlations at the person level in the face of high group-level stability is the central paradox behind a recent influential theoretical advance within social psychology known as the “bias of crowds”⁷¹⁻⁷², which leads to the argument that IAT scores could be caused by contexts but do not exist within individuals’ minds, absent eliciting contexts⁷³. However, others have argued that measurement error in the IAT D-score is a more parsimonious solution to the apparent puzzle⁷⁴, which our generative modeling results partially corroborate.

Second, generative model estimates are highly consistent across replications of the same task, whereas the two-stage approach estimates often vary considerably (e.g., compare the two-stage and generative model estimates for Flanker Study 1 versus Study 2). For example, for the Stroop task, the two-stage standard deviation contrast is significant in Study 2 but not in Study 1. Similarly, for the Flanker task, the two-stage mean contrast is significant in Study 1 but not in Study 2. In contrast, the more theoretically informed generative model (i.e. the lognormal models) parameters replicated consistently across studies.

Third, there is variation among generative models themselves, indicating that test-retest reliability varies—sometimes substantially (e.g., compare the normal versus lognormal models

for the Stroop task and IAT Race version)—depending on the assumed behavioral model.

Variability across models suggests that we should make efforts not to overgeneralize the failings (or successes) of a single behavioral model to attributes of the behavioral task itself. In other words, we should be explicit in acknowledging that inferences are conditional on a data-generating model and not the task per se. It is important to remember that the traditional two-stage summary approach does not escape this limitation. All statistical models, from t -tests to factor analysis, carry assumptions that bias inference. With this in mind, our results run counter to mounting claims that behavioral tasks are poorly suited for developing theories of individual differences^{10-12,14,16}; the central issue is not that behavior in these tasks is unreliable, but rather that outdated and overly restrictive statistical models fail to characterize the reliable characteristics of the data-generating process underlying behavior.

Comparing Summary Statistics to Generative Model Parameters

Why are test-retest correlations often higher in generative than two-stage models? Figure 7 illustrates the main reason: *hierarchical pooling*, which refers to regression of person-level parameters toward group-level means (or the regression line in this multivariate case). We chose examples to demonstrate how hierarchical pooling within generative models affects person-level parameter estimates. In Study 1 of the Stroop task, mean contrast estimates ($\mu_{i,\Delta}$) that would ordinarily be considered outliers in the two-stage approach are pooled toward the group-mean, producing higher expected test-retest correlations. Generative model parameter estimates also reveal potential practice effects in which almost every participant's expected mean contrast (i.e., Stroop effect, $\mu_{i,\Delta}$) is lower at Session 2 relative to Session 1. In contrast, standard deviation contrast estimates ($\sigma_{i,\Delta}$) show weak pooling, as well as poor expected test-retest correlations.

The same general pattern holds in the IAT (Black/White Race version), where $\mu_{i,\Delta}$ and $\sigma_{i,\Delta}$ exhibit strong and weak pooling, respectively. However, pooled $\mu_{i,\Delta}$ estimates for the IAT show regression toward the mean but without indication of potential practice effects. For the Delay Discounting task, both discounting rate (k_i) and choice sensitivity (c_i) parameters show moderate pooling (Figure 7). Together, these examples also show that hierarchical models do not automatically confer higher test-retest correlations. Instead, pooling only occurs to the extent it is warranted by data (see also the test-retest parameter recovery simulation results in the Supplementary Note 4).

Figure 7. Relationship between two-stage estimates and generative model parameters.

For the response time models (Stroop and IAT tasks), two-stage estimates are the sample mean and standard deviation contrasts for each participant and retest session (i.e., estimates from the summary statistic approach). Generative model parameters are means of the person-level posterior distributions (i.e., posterior expectations) for each participant. Gray lines connect two-stage estimates and generative model parameters for each participant, demonstrating how hierarchical models induce regression to the group-level mean/regression line. To help visualize the low correlation for the Stroop study, the standard deviation panel is zoomed in and two participants are not shown. For the Delay Discounting task, two-stage estimates reflect maximum likelihood estimates for each participant's discounting rate (k_i) and choice sensitivity (c_i) parameters; generative model parameters are means of person-level posterior distributions for each participant given by the full generative hyperbolic model. Gray and red ellipses indicate normal distribution quantiles (10%, 20%, ..., 90%, 99%) for the two-stage and generative model estimates, respectively.

Discussion

It is clear from our results that better generative assumptions improve both the reliability of our measures and the validity of conclusions we draw from behavioral data. Comparisons across five popular tasks demonstrate how incorporating assumptions about data-generating processes allows for a more precise characterization of individual differences in behavior better isolating sources of variability across experimental conditions. By attending to data-generating processes underlying behavior, generative modeling offers a solution not only to problems of low reliability (and by extension predictive validity), but also to problems with theory-description gaps arising from use and overinterpretation of statistical models that fail to instantiate reasonable assumptions, or worse, carry misleading assumptions. In contrast, traditional two-stage summary approaches used to analyze behavioral data are largely atheoretical mechanical exercises, and researchers may be unaware of implicit data-generating assumptions they make when using such methods (refer back to Figure 4). The explicitness of generative models ensures that assumptions and consequences of model design can be easily tracked, evaluated, and iteratively expanded on in future research.

It is important to emphasize that generative modeling is an iterative process, and throughout model development, model parameters can be assessed to determine their psychometric properties. Although we focused on test-retest reliability, there are many other properties worth exploring including parameter identifiability⁷⁵, parameter recovery⁷⁶⁻⁷⁸, tests of selective influence⁷⁹ (a form of construct validity where experimental manipulations cause expected changes in parameter values), as well as parameter convergence between behavioral models and models derived at other levels of analysis^{37,80} (e.g., with trait or neural models). Bayesian analysis facilitates joint estimation of all model parameters and their hypothesized relations, thus

allowing for proper calibration of uncertainty in key parameters, such as test-retest correlations. Although we focused on what we believe are uncontroversial generative assumptions in the generative models we tested here, we also discuss extensions to these generative models, including using more sophisticated evidence accumulation models of choice and response time behavior, in the Supplementary Note 8. These models can be further improved by methods that take advantage of adaptive experimental designs to further maximize informativeness of behavioral data in Supplementary Note 9⁸¹⁻⁸³.

One potential barrier to the use of generative models is that they necessitate a single joint model across all data modalities when making inference. In the context of the current study, this means that one could not simply fit the generative models proposed in our work to new data, extract the person-level parameters, and expect that such parameters should show better correspondence with other individual difference measures. Instead, one would need to construct a joint model of the behavioral and external measures of interest, which can be challenging without extensive experience with generative models. However, advances in computational statistics have recently made generative modeling much more accessible. We anticipate that generative modeling approaches will proliferate as scientists from all backgrounds recognize their utility for rigorous theory development and testing. There are now many accessible resources and software packages available to help researchers gain a deeper understanding of generative modeling, so they can apply these techniques to their own work. Resources include introductions to the philosophy and utility of generative or computational modeling for theory development^{2,84}, tutorials on building generative models from first principles^{54,85}, practical textbooks that combine introductions to both behavioral model development and hierarchical Bayesian modeling⁸⁶⁻⁸⁷, tutorials and case examples on developing joint generative models of

behavior and brain activity⁸⁰⁻⁸⁸, and open source R and Python software packages that allow both beginners and advanced users to apply popular generative models of behavioral to their own data using hierarchical Bayesian modeling⁸⁹⁻⁹¹. These and related sources make it clear that there is much more to building sound generative models than just capturing the shape of empirical data distributions, although it is clear from our results that even minimal assumptions can drastically improve the measurement of individual differences. Computational models that incorporate domain-specific assumptions about generative processes are a good example^{2,54,59}, which facilitate a level of mechanistic inference not provided by the simple, broad-based behavioral models used in the current study. We hope our message has come through clear—that the landscape for generative model development is vast, whereas traditional two-stage summary approaches are inherently limited in scope and application.

We end with a cautionary yet hopeful note: As history has revealed, heuristic use of summary statistics absent generative models impedes scientific progress, leading us down paths we would have never explored had we been made aware of the implicit assumptions that directed us there. Although our generative models may be wrong or mis-specified, their explicitness forces us to specify theoretical assumptions regarding how behavior arises, thereby requiring us to spend time thinking about the mechanisms that underly the brain, behavior, and their inter-relations. By embracing their incompleteness, we can strive to build generative models that are precise and thus meaningfully incorrect, rather than relying on vague, heuristic theories whose verisimilitude can be deceptive because of the theory-description gap. Identifying and knowing where our assumptions are wrong provides a natural path toward deepened understanding of the mechanisms underlying behavior. In this vein, the generative models we propose here are not intended to be fit to any behavioral data wherein response time distributions are the primary

outcome measure. Instead, we encourage researchers to start from such models, and iteratively extend them to better meet the particulars of their theory, task, or application.

Method

Datasets and Behavioral Paradigms

Given that the Stroop task is the running example throughout this article, we describe the details of the Stroop task from Hedge et al.¹² below. We provide details of all other tasks and datasets in Supplementary Note 5.

For the Stroop task, two sets of participants ($n = 47$, $n = 60$ for Studies 1 and 2) performed the task in two separate sessions three weeks apart. The main effect of interest is the contrast between congruent and incongruent conditions. Specifically, participants responded to the color of a word (either red, blue, green, or yellow). The word could be the same as the font color (e.g., the word “red” in red font; congruent condition or $c = 1$), a non-color word (e.g., “ship”; neutral condition), or a color word mapping onto another response option (e.g., the word “red” colored blue, green, or yellow; incongruent condition or $c = 2$). Participants completed 240 trials in each of the three conditions.

Data Analysis

Data Preprocessing

For all tasks involving response times, we removed trials with response times recorded as < 0 , assuming that such trials could not be part of the data-generating process. For the delay discounting task, we did not remove trials. Keeping all trials (except negative response times) can help identify regions of model misfit that offer insights into cognitive mechanisms that get obscured by oversimplified preprocessing choices (e.g., removing trials with response times less than 100 milliseconds). Such heuristic preprocessing choices tend to have strong, unpredictable

effects on inference⁹². Therefore, liberal inclusion criteria not only keep our models consistent with the goals of generative modeling but also demonstrate the utility of hierarchical modeling.

Two-Stage Summary Approach

The two-stage approach proceeds by reducing behavior within each participant to a point estimate before entering the resulting point estimates into a secondary statistical model to make inference. Below, we describe its implementation for each task.

Response Time Tasks. For the IAT, Stroop, Flanker, and Posner Cueing tasks, our first analysis followed the two-stage summary approach described in the introduction. We computed mean contrasts across task conditions by taking the mean in each condition and then subtracting the congruent from incongruent condition mean response time (see Figure 1). In addition, we computed standard deviation contrasts for comparison with the generative models (i.e., standard deviations of incongruent condition response times minus standard deviations of congruent condition response times). To estimate test-retest reliabilities, we computed Pearson correlations across participants for the mean and standard deviation contrasts.

Delay Discounting Task. We used maximum likelihood estimation to estimate discounting rates (k) and choice sensitivity parameters (c) from a hyperbolic model for each participant and session, followed by Pearson correlations across participant to estimate test-retest reliabilities of model parameter point estimates (see Supplementary Notes 2 and 3 for details). We compare these estimates to a hierarchical Bayesian estimation approach described below.

Generative Modeling Approach

If the goal is to make group-level inferences, hierarchical models can appropriately account for person-level uncertainty. Further, hierarchical models can increase precision of person-level

parameter estimates. Below, we extend the concept of generative modeling from person- to group-level model parameters.

Person-level Response Time Models. To characterize response time data, a generative model must obey some of the very simple properties outlined in the introduction. The normal (Gaussian) distribution is perhaps the simplest behavioral model that can generate a full distribution of response times. Despite not capturing skew and allowing negative values, it can still be useful for shifting away from the summary statistic approach. At the minimum, the normal distribution characterizes both the central tendency and the variance or spread of the response time distribution.

For example, each person's response times in the Stroop task can be conceptualized as arising from a separate normal distribution. Parameters from each distribution (e.g., means/standard deviations) are specific to each person within each task condition. The Stroop effect can be characterized by within-participant changes in the shape of response time distributions across trials within conditions. When using a normal generative distribution, the shape of the response time distribution is characterized by changes in the mean and standard deviation parameters across congruent and incongruent condition trials for each participant. We can write the normal generative model as:

$$\mathbf{RT}_{i,c,t} \sim \mathcal{N}(\mu_{i,c,t}, \sigma_{i,c,t}) \quad (1)$$

where $\mathbf{RT}_{i,c,t}$ contains the set of response times for participant i in condition c during experimental session t . The notation $\mathbf{RT} \sim \mathcal{N}(a, b)$ signifies that the response times are drawn from a generative process of a normal distribution ($\mathcal{N}$) with mean a and standard deviation b . In Equation 1, the collection of response times in each block of our experiment are separately characterized by a specific mean ($\mu_{i,c,t}$) and standard deviation ($\sigma_{i,c,t}$). Hence, whereas the

behavioral model implied by the two-stage summary approach reduces the response time data into a single summary statistic per condition, the behavioral model in Equation 1 reduces the data into two parameters per condition—parameters which, as we discuss below, can be assessed in terms of their own mean and variance⁹³.

The two-stage summary approach implicitly assumes Equation 1 as the data-generating model, because the sample mean is the analytical maximum likelihood estimator for the μ parameter in a normal distribution. However, the two-stage approach assumes the mean is estimated without measurement error. There are only two cases in which this assumption is valid. Specifically, recall that the standard error on the sample mean ($\sigma_{\hat{\mu}}$) (and correspondingly, the μ parameter of a normal distribution) is $\sigma_{\hat{\mu}} = \frac{\sigma}{\sqrt{n}}$. In our case, σ corresponds to the standard deviation of a person's response time distribution ($\sigma_{i,c,t}$ in Equation 1), and n to the number of trials they underwent in the given condition. Assuming no measurement error is equivalent to assuming that $\sigma_{\hat{\mu}} = 0$, which only occurs as either $\sigma \rightarrow 0$ or $n \rightarrow \infty$. Of course, we know *a priori* that each person's distribution of response time will have variability from trial-to-trial (i.e. $\sigma > 0$), and that we have a limited number of trials (i.e. $n \ll +\infty$). The normal generative model captures this knowledge, allowing joint estimation of μ and σ while appropriately accounting for uncertainty given the limited number of trials collected to estimate these parameters.

To facilitate interpretation, we will introduce a relabeling of the terms in Equation 1 based on the conditions they correspond to. First, we label the congruent condition (i.e., the first condition $c = 1$) as a baseline condition, where $RT_{i,1,t} = RT_{i,baseline,t}$, characterized by a baseline mean $\mu_{i,1,t} = \mu_{i,baseline,t}$ and baseline standard deviation $\sigma_{i,1,t} = \exp(\sigma_{i,baseline,t})$. To isolate the effects of

interference, or Stroop effects, we labeled a parameter Δ to signify the change from the baseline condition to the condition of interest (e.g., incongruent condition). This means that $RT_{i,2,t}$ is characterized by a mean $\mu_{i,2,t} = \mu_{i,base,t} + \mu_{i,\Delta,t}$ and standard deviation $\sigma_{i,2,t} = \exp(\sigma_{i,base,t} + \sigma_{i,\Delta,t})$. Note that the *base* and Δ standard deviation parameters are estimated on the log scale and then exponentially transformed to ensure they are greater than 0. Therefore, the test-retest correlation for the Δ standard deviation parameters is on the log scale. See the Supplementary Note 3 for more details.

The normal generative model is better at characterizing distributional changes in response times across conditions and advancing theoretical knowledge relative to the two-stage summary approach. Nevertheless, the model is limited because it does not obey all the simple properties of response times outlined in the introduction. In particular, the normal model (1) can produce negative response times, and (2) cannot capture asymmetric variance with respect to the mean (i.e., right skew). One simple adjustment is to logarithmically transform the response time data and assume a normal model on the transformed data, a process equivalent to assuming the response time data come from a different generative model called the lognormal distribution. Given this equivalence, we can specify a more theoretically consistent generative model as:

$$RT_{i,c,t} \sim \text{Lognormal}(\mu_{i,c,t}, \sigma_{i,c,t}) \quad (2)$$

With this adjustment, parameters $\mu_{i,c,t}$ and $\sigma_{i,c,t}$ will have very different abilities when characterizing the many shapes of response time distributions. In the lognormal model, the mean and standard deviation interact according to the *law of response time*⁵⁵. That is, an increase in either parameter, holding the other constant, produces an increase in both the mean and standard deviation of the response times predicted by the model. Figures 3A and 3B show how changes in either parameter change the shape of the predicted response time data. Each possible distribution

shape represents a prediction about how a participant's response time data should look, where the possible shapes are constrained by our commitments (or hypotheses) regarding the data-generating process (i.e. the lognormal model).

The lognormal model is an improvement over the normal model but still misses one important property of response time data. It is well-established that different response modalities (e.g., key press, mouse click, or verbal response) can produce shifts in response time distributions, even when the task demands and underlying evidence accumulation dynamics are identical⁹⁴. Typically, extra time taken to interact with the stimuli and apparatus is not considered part of the decision process and is instead referred to as “non-decision” time to make the theoretical position clear. Although non-decision factors seem unimportant, they may compromise our ability to accurately characterize response time data. For example, suppose a person completes a Stroop task in two conditions where they respond either verbally or by pressing a key. Even assuming the person follows the same decision process in identifying the color of the word (i.e., they have the same $\mu_{i,c,t}$ parameter), differences in executing the response across conditions are still likely. For example, if responding verbally is quicker than via key press, we would expect the response times to be shifted higher relative to the verbal condition. In this case, fitting the lognormal distribution to the observed response times would lead to different estimates for $\mu_{i,c,t}$ across the two conditions because the simple lognormal is not specified correctly according to the demands of the experiment. Consequently, having different estimates for $\mu_{i,c,t}$ might result in different interpretations about cognitive factors across the two contexts when non-decision factors actually caused the differences.

A simple solution is to adjust the lognormal distribution by introducing an additional parameter δ to move the distribution a distance of δ away from zero. Figure 3C illustrates the

effect of δ on a specific lognormal distribution. To impose some theoretical constraints, we could assume δ is specific to each person and is unlikely to change between conditions within a behavioral task. This assumption would be inappropriate in our example above, but it would be justified for the analyses we perform in later sections due to the method of data collection. With this new shift parameter and imposed theoretical constraints, we can now specify a shifted-lognormal model as:

$$\mathbf{RT}_{i,c,t} \sim \text{Shifted-Lognormal}(\delta_{i,t}, \mu_{i,c,t}, \sigma_{i,c,t}) \quad (3)$$

where $\mu_{i,c,t}$ and $\sigma_{i,c,t}$ have the same interpretations as described for Equation 2, and $\delta_{i,t}$ indicates the amount of shift or “non-decision time” specific to each person at each of the two experimental sessions.

Group-level Response Time Models. We have now defined generative models of person-level behavior for the response time tasks (normal, lognormal, and shifted lognormal models). The next step toward building full generative models of test-retest reliability is to define group-level probability distributions for person-level parameters. Starting with the three response time models, we assume that each person-level parameter i in the congruent task condition at each of the two sessions t is drawn from a normal group-level distributions with unknown means and standard deviations:

$$\begin{aligned} \mu_{i,\text{base},t} &\sim \mathcal{N}(\mu_{\text{mean},\text{base},t}, \mu_{\text{sd},\text{base},t}) \\ \sigma_{i,\text{base},t} &\sim \mathcal{N}(\sigma_{\text{mean},\text{base},t}, \sigma_{\text{sd},\text{base},t}) \end{aligned} \quad (4)$$

The group-level normal distributions here are considered prior models (or prior distributions) on the person-level parameters. Estimating group-level parameters from prior models allows for information to be pooled across participants such that each person-level estimate influences its corresponding group-level mean and standard deviation estimates, which in turn influence all other person-level estimates. This interplay between person- and group-level parameters

produces regression of person-level estimates toward the group mean (also referred to as *hierarchical pooling*, *shrinkage*, or *regularization*), which increases precision of person-level estimates⁶³. The specific amount that person-level parameters are pooled toward the group-level mean is proportional to the uncertainty or reliability when estimating the person-level parameters. In fact, from the perspective of classical test theory, the pooling incurred through the hierarchical model is a form of “true score estimation”, wherein the hierarchical pooling factor produces person-level estimates equivalent to reliability-based corrections under certain conditions^{24,64}. Note that the two-stage summary approach can be viewed as a special case of Equation 4 in combination with the normal behavioral model, wherein we assume that $\mu_{sd,base,t} \approx +\infty$, $\sigma_{mean,base,t} \approx -\infty$, and $\sigma_{sd,base,t} \approx 0$ (keeping in mind that $\sigma_{i,base,t}$ is on the log scale). In this highly restricted case, the group-level distribution across $\mu_{i,base,t}$ is essentially uniform, and there is subsequently no pooling of $\mu_{i,base,t}$ toward $\mu_{mean,base,t}$. Further, the person-level standard deviation parameters ($\sigma_{i,base,t}$) are all assumed to be approximately 0, for reasons given in the text surrounding Equation 1 (see also Figure 4). Therefore, in the same way that the normal behavioral model can be viewed as a relaxation of the overly restrictive (and implicit) assumptions underlying the two-stage summary approach, the normal group-level model (Equation 4) relaxes the assumption that person-level parameters follow a uniform distribution.

Beyond the relation between hierarchical pooling and reliability-based corrections, the normal distributions function similarly for person-level latent parameters in Equation 4 as they do for observed response times in Equation 1. Both cases assume that a normal distribution at one level of analysis generates observed or unobserved data at another level (e.g., observed response times are generated by normal distributions within participants, with unobserved means

and standard deviations generated from normal group-level distributions). *This joint specification of relations between parameters over all levels of analysis embodies the generative perspective.*

It allows for group- and person-level parameters to be estimated simultaneously, thus avoiding the pitfalls of the two-stage approach entirely (we illustrate the effect of these generative assumptions on person-level parameters in Figures 5 and 8). Although we do not demonstrate it here, the group-level model (i.e., Equation 4) itself is of theoretical interest—it can be extended to estimate relations between personality traits and decision mechanisms³⁷, or to generalize parameter estimates beyond non-representative samples⁹⁵.

To estimate test-retest reliability, we can assume that person-level change parameters (e.g., $\mu_{i,\Delta,t}$ and $\sigma_{i,\Delta,t}$) are correlated across sessions. Staying true to the generative perspective, we can estimate this correlation by assuming scores are drawn from a multivariate normal distributions rather than independent normal distributions as in Equation 4:

$$\begin{aligned} \begin{bmatrix} \mu_{i,\Delta,1} \\ \mu_{i,\Delta,2} \end{bmatrix} &\sim \text{MVNormal} \left(\begin{bmatrix} \mu_{\text{mean},\Delta,1} \\ \mu_{\text{mean},\Delta,2} \end{bmatrix}, \mathbf{S}_\mu \right) \\ \begin{bmatrix} \sigma_{i,\Delta,1} \\ \sigma_{i,\Delta,2} \end{bmatrix} &\sim \text{MVNormal} \left(\begin{bmatrix} \sigma_{\text{mean},\Delta,1} \\ \sigma_{\text{mean},\Delta,2} \end{bmatrix}, \mathbf{S}_\sigma \right) \end{aligned} \quad (5)$$

Using a multivariate normal distribution allows us to estimate covariances ($\mathbf{S}_\mu$ and $\mathbf{S}_\sigma$ matrices) between person-level parameters across sessions that can be decomposed into group-level parameter variances and the correlation between person-level parameters across sessions—this correlation represents the test-retest reliability of the generative model parameters (see Supplementary Note 3 for mathematical details). If the correlation is zero, then Equation 5 is equivalent to Equation 4 (i.e. the normal distributions are independent). As the correlations become stronger, however, person-level parameters become increasingly pooled toward a common regression line, which in our case represents the test-retest correlation. In the same way

that Equation 4 produces person-level estimates that are “corrected” for unreliability, Equation 5 produces a test-retest correlation that is corrected for attenuation due to unreliability (or uncertainty) in person-level estimates. Therefore, correlations extracted from the full generative model indicate the “true” (where true indicates the statistical expectation) temporal stability of the person-level parameters given the assumed generative model²⁴. Similar to the discussion surrounding Equation 4, the two-stage summary approach can be viewed as a special case of Equation 5, wherein $\mu_{sd,\Delta,t} \approx +\infty$, $\sigma_{mean,\Delta,t} \approx -\infty$, and $\sigma_{sd,\Delta,t} \approx 0$. In this highly restricted case, the test-retest correlation obtained from the normal generative model is equivalent to what is obtained with the two-stage summary approach.

For the shifted lognormal model, we estimated a single shift parameter for each participant at each timepoint (assuming that shift is equivalent between task conditions). Details about the shift parameter specification and prior distributions for group-level parameters in Equations 4-5 are available in Supplementary Note 3. Figure 8 illustrates the full lognormal generative model.

Figure 8. Full lognormal generative model spanning from within- to between-person variation.

At the highest level, person-level parameters follow group-level normal distributions, which capture between-person variation. In our case, we assume that the difference in person-level ($\mu_{i,\Delta}$ and $\sigma_{i,\Delta}$) parameters across congruent and incongruent task conditions follow a multivariate normal distribution across the two test-retest sessions. Before fitting the model to data, the group-level distributions can take on many different possible shapes, and person-level difference parameters (indicated by the dark red points) therefore show no systematic correlations. As a result, the lognormal distributions implied within persons can also take on many possible shapes. These possibilities are then substantially reduced after observing empirical data. The observed response times for each person (gray histograms on the lower right panels) inform their behavioral model parameters, and the behavioral model parameters then inform the group-level model parameters (dark red points and gray histograms in the middle and top panels indicate posterior means of the person-level parameters). The test-retest correlation we obtain from the full generative model therefore represents the stability in the person-level parameters across test-retest sessions, while accounting for uncertainty within each session that arises due to a lack of information (e.g., minimal number of trials). Here, uncertainty is represented by variability in person- and group-level distributions across randomly drawn posterior samples (light red lines).

Delay Discounting Model. Delay discounting data are often modeled using a hyperbolic discounting model, which includes a discounting rate (k) that indicates how steeply a person discounts delayed rewards and a choice sensitivity parameter (c) that indicates how randomly versus deterministically a person chooses between smaller-sooner versus larger-later rewards based on their valuations for each. The mathematical details of this model are available in Supplementary Note 3. Extending the person-level hyperbolic delay discounting model to a full generative model that can estimate test-retest reliability follows the same logic as outlined for response time models. We used the same multivariate normal distribution parameterization to estimate test-retest correlations between discounting rate (k) and choices sensitivity (c) parameters.

Parameter Estimation

A benefit of Bayesian estimation is that after specifying a joint probability model (i.e. the full group- and person-level generative model), it is possible to compute conditional probabilities that determine which parameter values are most credible given the observed data. This results in *posterior distributions* over model parameters that express the probability of a specific value for the parameter given the model and data. Because computing conditional probabilities analytically requires solving complex and often intractable integrals, Bayesian model parameters are typically estimated using numerical integration methods. We estimated parameters from all models using Stan (version 2.19.2), a probabilistic programming language that uses a variant of Markov Chain Monte Carlo to estimate posterior distributions for parameters within Bayesian models⁹⁶. Details are described in Supplementary Note 3.

Sensitivity Analyses of Group-level Model Specification

We tested multiple different group-level models to determine how sensitive our results were to changes in generative assumptions. We report details in Supplementary Information, but offer a brief overview of results here. First, we used parameter recovery simulations to show that we can accurately recover the “true” underlying test-retest correlation using the full generative models as described in the main text (see Supplementary Note 4 and Figure S2). Second, we tested an alternative group-level model wherein both the person-level base and change parameters were drawn from separate multivariate normal distributions, as opposed to only the change parameters (the “Joint Separate” model in Supplementary Note 7 and Figure S8). Finally, we tested another group-level model wherein we directly estimated the $\mu_{i,c,t}$ and $\sigma_{i,c,t}$ parameters as opposed to estimating baseline and change parameters. For this model, we assumed that person-level parameters were drawn from a single multivariate normal distribution (but separate for μ versus σ) across conditions and sessions (the “Joint Single” model in Supplementary Note 7 and Figure S9). Both models produced very similar results to the model specification presented throughout the main text.

Funding Details

T.P.B. was supported in part by Grants UH2DE025980 and UL1TR002733 from the National Institutes of Health, B.M.T. was supported by a CAREER award from the National Science Foundation, and W.-Y.A. was supported by the Creative-Pioneering Researchers Program through Seoul National University and the National Research Foundation (NRF) of Korea (NRF-2018R1C1B3007313).

Disclosure Statement

The authors declare no conflicts of interest.

Data Availability Statement

All de-identified data along with the R and Stan codes used to reproduce our results and figures are available on our GitHub repository (https://github.com/Nathaniel-Haines/Reliability_2020).

References

1. Hempel, C. G. & Oppenheim, P. Studies in the Logic of Explanation. *Philosophy of Science* **15**, 135–175 (1948).
2. Guest, O. & Martin, A. E. How computational modeling can force theory building in psychological science. *PsyArXiv Preprint* 1–13 (2020). doi:10.31234/osf.io/rybh9
3. Kellen, D. A Model Hierarchy for Psychological Science. *Computational Brain & Behavior* **2**, 160–165 (2019).
4. Suppes, P. Models of Data. *Studies in Logic and the Foundations of Mathematics* **44**, 252–261 (1966).
5. Michell, J. Is Psychometrics Pathological Science? *Measurement: Interdisciplinary Research & Perspective* **6**, 7–24 (2008).
6. Szollosi, A. & Donkin, C. Neglected Sources of Flexibility in Psychological Theories: from Replicability to Good Explanations. *Computational Brain & Behavior* **2**, 190–192 (2019).
7. Yarkoni, T. The generalizability crisis. (2019).
8. Chen, B. *et al.* Individual Variability and Test-Retest Reliability Revealed by Ten Repeated Resting-State Brain Scans over One Month. *PLOS ONE* **10**, e0144963 (2015).
9. Elliott, M. L. *et al.* What Is the Test-Retest Reliability of Common Task-Functional MRI Measures? New Empirical Evidence and a Meta-Analysis. *Psychological Science* **31**, 792–806 (2020).
10. Enkavi, A. Z. *et al.* Large-scale analysis of test–retest reliabilities of self-regulation measures. *PNAS* **116**, 5472–5477 (2019).
11. Gawronski, B., Morrison, M., Phillips, C. E. & Galdi, S. Temporal Stability of Implicit and Explicit Measures. *Pers Soc Psychol Bull* **43**, 300–312 (2017).
12. Hedge, C., Powell, G. & Sumner, P. The reliability paradox: Why robust cognitive tasks do not produce reliable individual differences. *Behav Res* **103**, 1–21 (2017).
13. Noble, S., Scheinost, D. & Constable, R. T. A decade of test-retest reliability of functional connectivity: A systematic review and meta-analysis. *NeuroImage* **203**, 116157 (2019).
14. Dang, J., King, K. M. & Inzlicht, M. Why Are Self-Report and Behavioral Measures Weakly Correlated? *Trends in Cognitive Sciences* 1–3 (2020). doi:10.1016/j.tics.2020.01.007
15. Schimmack, U. The Implicit Association Test: A Method in Search of a Construct. *Perspectives on Psychological Science* **117**, 174569161986379 (2019).
16. Wennerhold, L. & Friese, M. Why Self-Report Measures of Self-Control and Inhibition Tasks Do Not Substantially Correlate. *Collabra: Psychology* **6**, 9 (2020).
17. Stroop, R. Studies of interference in serial verbal reactions. *Journal of Experimental Psychology: General* **18**, 643–662 (1935).
18. Hedge, C., Powell, G., Bompas, A., & Sumner, P. Self-reported impulsivity does not predict response caution. *Personality and Individual Differences* **167**, 110257 (2020).
19. Strickland, J. C., & Johnson, M. W. Rejecting impulsivity as a psychological construct: A theoretical, empirical, and sociocultural argument. *Psychological Review*, **128**, 336–361 (2021).

20. Frey, R., Pedroni, A., Mata, R., Rieskamp, J., & Hertwig, R. Risk preference shares the psychometric structure of major psychological traits. *Science Advances*, **3**, e1701381 (2017).
21. Arslan, R. C., Brümmer, M., Dohmen, T., Drewelies, J., Hertwig, R., & Wagner, G. G. How people know their risk preference. *Scientific Reports*, **10**, 1-14 (2020).
22. Tisdall, L., Frey, R., Horn, A., Ostwald, D., Horvath, L., Pedroni, A., ... & Mata, R. Brain-behavior associations for risk taking depend on the measures used to capture individual differences. *Frontiers in Behavioral Neuroscience*, **14**, 194 (2020).
23. Ly, A. *et al.* in *Computational Models of Brain and Behavior* (ed. Moustafa, A. A.) 467–480 (John Wiley & Sons, Ltd, 2017). doi:10.1002/9781119159193
24. Rouder, J. N. & Haaf, J. M. A psychometrics of individual differences in experimental tasks. *Psychonomic Bulletin & Review* **26**, 452–467 (2019).
25. Turner, B. M., Forstmann, B. U., Love, B. C., Palmeri, T. J. & Van Maanen, L. Approaches to analysis in model-based cognitive neuroscience. *Journal of Mathematical Psychology* **76**, 65–79 (2017).
26. Vandekerckhove, J. A cognitive latent variable model for the simultaneous analysis of behavioral and personality data. *Journal of Mathematical Psychology* **60**, 58–71 (2014).
27. Davis-Stober, C. P., Park, S., Brown, N. & Regenwetter, M. Reported violations of rationality may be aggregation artifacts. *PNAS* **113**, E4761–E4763 (2016).
28. Estes, W. K. The problem of inference from curves based on group data. *Psychological Assessment* **53**, 134–140 (1956).
29. Heathcote, A., Brown, S. & Mewhort, D. J. K. The power law repealed: The case for an exponential law of practice. *Psychonomic Bulletin & Review* **7**, 185–207 (2000).
30. Liew, S. X., Howe, P. D. L. & Little, D. R. The appropriacy of averaging in the study of context effects. *Psychonomic Bulletin & Review* **23**, 1639–1646 (2016).
31. Pagan, A. Econometric issues in the analysis of regressions with generated regressors. *Scientific American* **25**, 221–247 (1984).
32. Turner, B. M., Schley, D. R., Muller, C. & Tsetsos, K. Competing theories of multialternative, multiattribute preferential choice. *Psychological Review* **125**, 329–362 (2018).
33. Johnson, J. G. & Busemeyer, J. R. A dynamic, stochastic, computational model of preference reversal phenomena. *Psychological Assessment* **112**, 841–861 (2005).
34. Kvam, P. D. Modeling accuracy, response time, and bias in continuous orientation judgments. *Journal of Experimental Psychology: Human Perception and Performance* **45**, 301–318 (2019).
35. Kvam, P. D. & Busemeyer, J. R. A distributional and dynamic theory of pricing and preference. *Psychological Assessment* **127**, 1053–1078 (2020).
36. Leth-Steensen, C., Elbaz, Z. K. & Douglas, V. I. Mean response times, variability, and skew in the responding of ADHD children: a response time distributional approach. *Acta Psychologica* **104**, 167–190 (2000).
37. Haines, N. *et al.* Anxiety Modulates Preference for Immediate Rewards Among Trait-Impulsive Individuals: A Hierarchical Bayesian Analysis. *Clinical Psychological Science* (2020). doi:10.1177/2167702620929636

38. Heathcote, A., Popiel, S. J. & Mewhort, D. J. K. Analysis of response time distributions: An example using the Stroop task. *Psychological Assessment* **109**, 340–347 (1991).
39. Romeu, R. J., Haines, N., Ahn, W.-Y., Busemeyer, J. R. & Vassileva, J. A computational model of the Cambridge gambling task with applications to substance use disorders. *Drug and Alcohol Dependence* **206**, 107711 (2019).
40. Ross, C. T., Winterhalder, B. & McElreath, R. Racial Disparities in Police Use of Deadly Force Against Unarmed Individuals Persist After Appropriately Benchmarking Shooting Data on Violent Crime Rates. *Social Psychological and Personality Science* **16**, 194855062091607 (2020).
41. Rotello, C. M., Heit, E. & Dubé, C. When more data steer us wrong: replications with the wrong dependent measure perpetuate erroneous conclusions. *Psychonomic Bulletin & Review* **22**, 944–954 (2014).
42. Devezzer, B., Nardin, L. G., Baumgaertner, B. & Buzbas, E. O. Scientific discovery in a model-centric framework: Reproducibility, innovation, and epistemic diversity. *PLOS ONE* **14**, e0216125 (2019).
43. Devezzer, B., Navarro, D. J., Vandekerckhove, J. & Buzbas, E. O. The case for formal methodology in scientific reform. *bioRxiv* 1–28 (2020). doi:10.1101/2020.04.26.048306
44. Regenwetter, M. & Robinson, M. M. The construct–behavior gap in behavioral decision research: A challenge beyond replicability. *Psychological Review* **124**, 533–550 (2017).
45. Ahn, W.-Y. & Busemeyer, J. R. Challenges and promises for translating computational tools into clinical practice. *Current Opinion in Behavioral Sciences* **11**, 1–7 (2016).
46. Busemeyer, J. R., Gluth, S., Rieskamp, J. & Turner, B. M. Cognitive and Neural Bases of Multi-Attribute, Multi-Alternative, Value-based Decisions. *Trends in Cognitive Sciences* **23**, 251–263 (2019).
47. Huys, Q. J. M., Maia, T. V. & Frank, M. J. Computational psychiatry as a bridge from neuroscience to clinical applications. *Nature Neuroscience* **19**, 404–413 (2016).
48. Montague, R., Dolan, R. J., Friston, K. J. & Dayan, P. Computational psychiatry. *Trends in Cognitive Sciences* **16**, 72–80 (2012).
49. Navarro, D. J. If Mathematical Psychology Did Not Exist We Might Need to Invent It: A Comment on Theory Building in Psychology:. *Perspectives on Psychological Science* (2021). doi:10.1177/1745691620974769
50. Rangel, A., Camerer, C. F., Camerer, C. & Montague, R. A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience* **9**, 545–556 (2008).
51. Townsend, J. T. Mathematical psychology: Prospects for the 21st century: A guest editorial. *Journal of Mathematical Psychology* **52**, 269–280 (2008).
52. Turner, B. M. *et al.* A Bayesian framework for simultaneously modeling neural and behavioral data. *NeuroImage* **72**, 193–206 (2013).
53. Wiecki, T. V., Poland, J. & Frank, M. J. Model-Based Cognitive Neuroscience Approaches to Computational Psychiatry: Clustering and Classification. *Clinical Psychological Science* **3**, 378–399 (2015).

54. Wilson, R. C. & Collins, A. G. Ten simple rules for the computational modeling of behavioral data. *eLife* **8**, 558 (2019).
55. Wagenmakers, E.-J. & Brown, S. On the linear relation between the mean and the standard deviation of a response time distribution. *Psychological Review* **114**, 830–841 (2007).
56. Casella, G. & Berger, R. L. *Statistical Inference*. (Duxbury Press, 2002).
57. Fisher, R. A. On the mathematical foundations of theoretical statistics. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character* **222**, 309–368 (1922).
58. Turner, B. M. & Van Zandt, T. A tutorial on approximate Bayesian computation. *Journal of Mathematical Psychology* **56**, 69–85 (2012).
59. Jarecki, J. B., Tan, J. H. & Jenny, M. A. A framework for building cognitive process models. *Psychonomic Bulletin & Review* **27**, 1218–1229 (2020).
60. Heathcote, A. & Love, J. Linear Deterministic Accumulator Models of Simple Choice. *Front. Psychol.* **3**, (2012).
61. Rouder, J. N., Province, J. M., Morey, R. D., Gomez, P. & Heathcote, A. The Lognormal Race: A Cognitive-Process Model of Choice and Latency with Desirable Psychometric Properties. *Psychometrika* **80**, 491–513 (2014).
62. Efron, B. & Morris, C. Stein's paradox in statistics. *Scientific American* **236**, 119–127 (1977).
63. Gelman, A. Multilevel (Hierarchical) Modeling: What It Can and Cannot Do. *Technometrics* **48**, 432–435 (2006).
64. Williams, D. R., Martin, S. R., DeBolt, M., Oakes, L. & Rast, P. A fine-tooth comb for measurement reliability: Predicting true score and error variance in hierarchical models. *PsyArXiv Preprint* (2020). doi:10.31234/osf.io/2ux7t
65. Ahn, W.-Y. *et al.* Rapid, precise, and reliable measurement of delay discounting using a Bayesian learning algorithm. *Scientific Reports* **10**, 1–10 (2020).
66. Eriksen, B. A. & Eriksen, C. W. Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception & Psychophysics* **16**, 143–149 (1974).
67. Green, L. & Myerson, J. A Discounting Framework for Choice With Delayed and Probabilistic Rewards. *Psychological Bulletin* **130**, 769–792 (2004).
68. Greenwald, A. G., McGhee, D. E. & Schwartz, J. L. K. Measuring individual differences in implicit cognition: the implicit association test. *Psychological Assessment* **74**, 1464–1480 (1998).
69. Mazur, J. E. in *Quantitative analyses of behavior* (eds. Commons, M. L., Mazur, J. E., Nevin, J. A. & Rachlin, H.) **5**, 55–73 (Lawrence Erlbaum Associates, Inc., 1987).
70. Posner, M. I. Orienting of Attention*:. *Quarterly Journal of Experimental Psychology* **32**, 3–25 (1980).
71. Payne, B. K., Vuletich, H. A. & Lundberg, K. B. The Bias of Crowds: How Implicit Bias Bridges Personal and Systemic Prejudice. *Psychological Inquiry* **28**, 233–248 (2017).
72. Rivers, A. M., Rees, H. R., Calanchini, J. & Sherman, J. W. Implicit Bias Reflects the Personal and the Social. *Psychological Inquiry* **28**, 301–305 (2017).
73. Jost, J. T. The IAT Is Dead, Long Live the IAT: Context-Sensitive Measures of Implicit Attitudes Are Indispensable to Social and Political Psychology:. *Current Directions in Psychological Science* **28**, 10–19 (2019).

74. Connor, P. & Evers, E. R. K. The Bias of Individuals (in Crowds): Why Implicit Bias Is Probably a Noisily Measured Individual-Level Construct. *Perspectives on Psychological Science* **15**, 1329–1345 (2020).
75. Spektor, M. S. & Kellen, D. The relative merit of empirical priors in non-identifiable and sloppy models: Applications to models of learning and decision-making. *Psychonomic Bulletin & Review* **4**, 1–22 (2018).
76. Ahn, W.-Y., Krawitz, A., Kim, W., Busemeyer, J. R. & Brown, J. W. A model-based fMRI analysis with hierarchical Bayesian parameter estimation. *Journal of Neuroscience, Psychology, and Economics* **4**, 95–110 (2011).
77. Haines, N., Vassileva, J. & Ahn, W.-Y. The Outcome-Representation Learning Model: A Novel Reinforcement Learning Model of the Iowa Gambling Task. *Cognitive Science: A Multidisciplinary Journal* **47**, 1–28 (2018).
78. Miletić, S., Turner, B. M., Forstmann, B. U. & Van Maanen, L. Parameter recovery for the Leaky Competing Accumulator model. *Journal of Mathematical Psychology* **76**, 25–50 (2017).
79. Criss, A. H. Differentiation and response bias in episodic memory: evidence from reaction time distributions. *Psychological Assessment* **36**, 484–499 (2010).
80. Turner, B. M., Forstmann, B. U., Love, B. C., Palmeri, T. J. & Van Maanen, L. Approaches to analysis in model-based cognitive neuroscience. *Journal of Mathematical Psychology* **76**, 65–79 (2017).
81. Cavagnaro, D. R., Pitt, M. A. & Myung, J. I. Model discrimination through adaptive experimentation. *Psychonomic Bulletin & Review* **18**, 204–210 (2011).
82. Myung, J. I., Cavagnaro, D. R. & Pitt, M. A. A tutorial on adaptive design optimization. *Journal of Mathematical Psychology* **57**, 53–67 (2013).
83. Yang, J., Pitt, M. A., Ahn, W.-Y. & Myung, J. I. ADOPy: a python package for adaptive design optimization. *Behav Res* 1–24 (2020). doi:10.3758/s13428-020-01386-4
84. van Rooij, I. & Baggio, G. Theory before the test: How to build high-verisimilitude explanatory theories in psychological science. *PsyArXiv* 1–24 (2020). doi:10.31234/osf.io/7qbpr
85. van Rooij, I. & Blokpoel, M. Formalizing Verbal Theories. *Social Psychology* **51**, 285–298 (2020).
86. Farrell, S. & Lewandowsky, S. *Computational Modeling of Cognition and Behavior*. (Cambridge University Press, 2018).
87. Lee, M. D. & Wagenmakers, E.-J. *Bayesian Cognitive Modeling*. (Cambridge University Press, 2014).
88. Palestro, J. J. *et al.* A tutorial on joint models of neural and behavioral measures of cognition. *Journal of Mathematical Psychology* **84**, 20–48 (2018).
89. Ahn, W.-Y., Haines, N. & Zhang, L. Revealing Neurocomputational Mechanisms of Reinforcement Learning and Decision-Making With the hBayesDM Package. *bioRxiv* **1**, 24–57 (2017).
90. Mathys, C. D. *et al.* Uncertainty in perception and the Hierarchical Gaussian Filter. *Front. Hum. Neurosci.* **8**, (2014).
91. Wiecki, T. V., Sofer, I. & Frank, M. J. HDDM: Hierarchical Bayesian estimation of the Drift-Diffusion Model in Python. *Front. Neuroinform.* **7**, 1–10 (2013).

92. Parsons, S. Exploring reliability heterogeneity with multiverse analyses: Data processing decisions unpredictably influence measurement reliability. *PsyArXiv Preprint* 1–35 (2020). doi:10.31234/osf.io/y6tcz
93. Williams, D. R., Mulder, J., Rouder, J. N. & Rast, P. Beneath the surface: Unearthing within-person variability and mean relations with Bayesian mixed models. *Psychological Methods* **26**, 74–89 (2021).
94. Gomez, P., Ratcliff, R. & Childers, R. Pointing, looking at, and pressing keys: A diffusion model account of response modality. *Psychological Assessment* **41**, 1515–1523 (2015).
95. Kennedy, L. & Gelman, A. Know your population and know your model: Using model-based regression and poststratification to generalize findings beyond the observed sample. *arXiv* (2019).
96. Carpenter, B. *et al.* Stan: A Probabilistic Programming Language. *Journal of statistical software* **76**, 1–32 (2017).