



# Statistics in the Service of Science: Don't Let the Tail Wag the Dog

Henrik Singmann<sup>1</sup> · David Kellen<sup>2</sup> · Gregory E. Cox<sup>3</sup> · Suyog H. Chandramouli<sup>4</sup> · Clinton P. Davis-Stober<sup>5</sup> · John C. Dunn<sup>6</sup> · Quentin F. Gronau<sup>7</sup> · Michael L. Kalish<sup>2</sup> · Sara D. McMullin<sup>5</sup> · Danielle J. Navarro<sup>8</sup> · Richard M. Shiffrin<sup>9</sup>

Accepted: 20 January 2022  
© The Author(s) 2022

## Abstract

Statistical modeling is generally meant to describe patterns in data in service of the broader scientific goal of developing theories to explain those patterns. Statistical models support meaningful inferences when models are built so as to align parameters of the model with potential causal mechanisms and how they manifest in data. When statistical models are instead based on assumptions chosen by default, attempts to draw inferences can be uninformative or even paradoxical—in essence, the tail is trying to wag the dog. These issues are illustrated by van Doorn et al. ([this issue](#)) in the context of using Bayes Factors to identify effects and interactions in linear mixed models. We show that the problems identified in their applications (along with other problems identified here) can be circumvented by using priors over inherently meaningful units instead of default priors on standardized scales. This case study illustrates how researchers must directly engage with a number of substantive issues in order to support meaningful inferences, of which we highlight two: The first is the problem of *coordination*, which requires a researcher to specify how the theoretical constructs postulated by a model are functionally related to observable variables. The second is the problem of *generalization*, which requires a researcher to consider how a model may represent theoretical constructs shared across similar but non-identical situations, along with the fact that model comparison metrics like Bayes Factors do not directly address this form of generalization. For statistical modeling to serve the goals of science, models cannot be based on default assumptions, but should instead be based on an understanding of their coordination function and on how they represent causal mechanisms that may be expected to generalize to other related scenarios.

**Keywords** Mixed models · Default Bayes Factors · Standardized effect sizes · Coordination function · Generalizability · Statistics

## Introduction

A central goal of science is to develop theories that describe causal mechanisms<sup>1</sup> in sufficient detail to derive predictions but with sufficient generality to apply across different

related scenarios. For example, we would like a theory of memory to explain why increased study time results in improved recognition performance. To advance this goal, we might build a formal model that connects independent variables like study time to dependent variables like accuracy via theoretical constructs like trace strength and processes of encoding and retrieval, describing mechanisms by which study time is causally related to accuracy.<sup>2</sup> By constructing models based on theory, we put ourselves

<sup>1</sup>Our use of “causal mechanisms” carries only a vague reading of both “causal” (to accommodate reasons as causes) and “mechanisms” (to accommodate probabilistic and computational theories) and we therefore ask the reader to avoid any kind of strong reading, keeping in mind that the issues discussed here hold regardless of one’s ontological commitments.

<sup>2</sup>At this stage, it is important to get a couple of points clear: First, when developing any kind of theoretical account, one needs to ensure that its level of specificity is commensurate with the “resolution” of the data, both in terms of degrees of freedom but also in terms of the precision of the data points being considered (for relevant discussions, see Kellen, 2019; Navarro, 2018). Second, there are a number of cases in which researchers are interested in comparisons such as “*Is treatment A more efficacious than treatment B?*”, without having any kind of deep theoretical considerations in mind. As will become clear below, it would be a mistake to think that these cases are somehow immune to the kinds of problems that will be discussed here.

Henrik Singmann, David Kellen, and Gregory E. Cox contributed equally to this work.

✉ Henrik Singmann  
singmann@gmail.com

Extended author information available on the last page of the article.

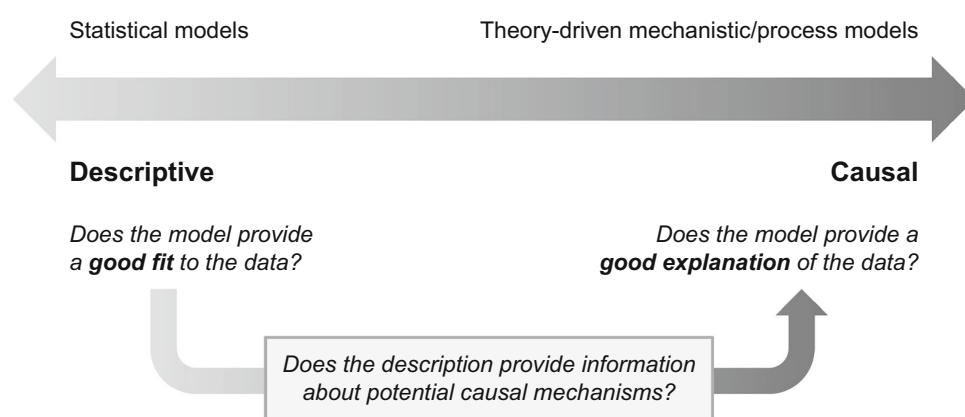
in a position to help explain effects on other dependent variables (e.g., response time), to understand the action of other related independent variables (e.g., repetition or spacing), and to suggest new avenues of research (e.g., what factors affect trace strength). In contrast, when models are constructed based on “default” assumptions that are unconnected to theory, these scientific insights are forfeit. Trying to make inferences from atheoretical modeling is, in effect, the tail (modeling) wagging the dog (theory).

Formal models are integral to drawing inferences from data in order to achieve scientific goals, contributing in a number of different ways (Cox & Shiffrin, *in press*; Luce, 1995; Navarro, 2021; Kellen, 2019). An important dimension along which modeling can vary is between *descriptive* and *causal* (Fig. 1). This dimension does not necessarily refer to any particular property of the models themselves, but rather to the purpose the model is intended to serve. As illustrated in the example above, the purpose of causal modeling is to represent causal mechanisms in a theory in a way that enables quantitative predictions to be derived and tested. While causal modeling can thus be extremely productive in terms of making scientific inferences, it is not always possible if there is limited knowledge of the phenomena to be explained in a given domain. In such a circumstance, descriptive modeling can be used to identify important quantitative features of a set of data—“signals” in the presence of “noise”. Statistical modeling typically has a descriptive purpose in this sense. But to yield useful scientific insights, the “signals” detected by descriptive modeling should be indicative of the action of potential causal mechanisms or their interactions. For this to be true, even models with a primarily descriptive purpose should be designed so as to be meaningfully linked to potential causal mechanisms. Any particular modeling endeavor is likely to have a mixture of descriptive and causal purposes. This ambiguity can lead to confusion about how to best evaluate a model: Causal modeling is successful when the model provides an insightful and productive *explanation* of patterns in data. In turn, descriptive modeling

is successful when the model “fits”, achieving a close match between predicted and observed data patterns. Descriptive modeling serves science when the description it offers provides information about potential causal factors that are relevant to explaining data patterns.

It is especially important to consider causal mechanisms—even in otherwise descriptive modeling—in the psychological sciences. Here, data exhibit a large degree of “noise” because they reflect variability from many sources, only some of which are related to the scientific questions at hand. To address this state of affairs, increasingly sophisticated statistical models are brought to bear in order to extract relevant signals from these data. One example is the general purpose statistical framework discussed by van Doorn et al. ([this issue](#)), henceforth vDAHSW, that of Linear Mixed Models (LMMs). They intend this framework to be able to describe data patterns across myriad substantive applications—to identify “signals” that represent effects and interactions regardless of the specific experimental manipulations or scales of measurement. In accord with this intent, they formulate models and compare them using a variety of “default” assumptions, including assumptions about model structure and the way prior distributions should be defined. But by adopting these defaults in order to expand the scope of their model’s descriptive abilities, they end up breaking the link between the formal representations in the model and the causal mechanisms that may be theorized to yield the observed data patterns. Admittedly, making that link is hard, but it is precisely that link that determines whether a descriptive statistical model serves its intended scientific purpose: to identify “signals” (e.g., a difference in average accuracy between conditions) that potentially reflect the action of a theorized causal mechanism. This issue is not specific to the LMMs explored by vDAHSW - it arises in any analysis entailing a complex web of assumptions. If these assumptions are chosen by “default”, then there is no guarantee that they will be at all related to the scientific theories the model is meant to embody or inform. Regardless of the goal of the model, irreflective

**Fig. 1** Modeling in science has different purposes that lie on a continuum between descriptive and causal. Considerations about model complexity (is it warranted in light of the data?) are tacitly assumed in this illustration



adoption of default assumptions renders a model scientifically meaningless.

Statistical methodology that keeps science on the path toward ‘progress’ needs to be based on an appreciation of how causal mechanisms are linked to data. As such, we believe that the call by vDAHSW for general guidelines and defaults that are divorced from this connection is, at best, problematic. In the first part of our paper, we directly address the questions posed by vDAHSW, showing among other things that the differences in Bayes Factors obtained under different aggregation regimes result from the tacit adoption of “default Bayes Factors”. These rely on the specification of priors over standardized effects, but the reported differences in conclusions dissipate when using non-standardized effects. What appeared to be a deep inferential paradox can be avoided by merely introducing subject-matter considerations—namely, the scale of measurement—into a descriptive model. In the second part of this paper, we discuss how giving careful thought to different kinds of causal considerations can help descriptive models better serve scientific goals. One consideration is the “coordination function” that specifies how the components of a model are linked to theoretical constructs and observable quantities; another is “generalization”, which focuses on how a model and/or its components are intended to be applied to other related but non-identical scenarios. Models, whether intended for primarily descriptive or causal purposes, should be tailored to the empirical substrate and theoretical statements at hand, keeping “default” assumptions to a minimum.

### Part I: The Perils of Letting the Tail Wag the Dog

LMMs provide a widely applicable descriptive modeling approach for data with complex dependency structures, such as those that often arise in the psychological sciences. LMMs distinguish between “fixed effects” due to experimental condition or group membership and “random effects” due to, for example, variability between individuals or items. Fixed effects are represented by regression coefficients that are constant (i.e., fixed) across all observations within a group or condition (just like the coefficients in ordinary regression models). Random effects are represented by regression coefficients that can differ across different clusters or groups of observations (e.g., individuals, items).<sup>3</sup> By explicitly assigning different representations to different types of variability, LMMs

<sup>3</sup>This understanding of fixed and random effects is the same as vDAHSW’s. In any case, to avoid any potential misunderstandings, we provide a complete formal specification of the linear mixed models in the [Appendix](#).

already represent some assumed aspects of a causal structure. But as illustrated in the examples by vDAHSW, many other choices must be made when specifying LMMs and these too should be guided by causal considerations.

The examples provided by vDAHSW rely on a representation of fixed and random effects that is commonly adopted when applying LMMs to typical experimental designs: Fixed effects are represented by an *intercept* representing the overall mean and a *slope* representing the average difference between conditions. Random effects are represented by allowing individuals to have different intercepts and/or slopes which are normally distributed around the group average. In a frequentist setting, there is a wealth of theoretical and practical knowledge on how to apply these kinds of LMMs, most of which are built into the `lme4` R package (Bates et al., 2015). Much of this knowledge can be boiled down to a general recommendation of using the most complete random-effects structure that can be justified by the experimental design (Baayen et al., 2008; Barr et al., 2013; Judd et al., 2012; Westfall et al., 2014).<sup>4</sup> Behind this recommendation is the idea that failing to represent a random effect that is in fact present—in other words, failing to consider as many causal factors as possible—will lead to erroneous inferences.

As exemplified by vDAHSW, there is growing interest in applying Bayesian inferential methods (Jeffreys, 1961) such as *Bayes Factors*, in the context of LMMs. Consider data  $y$  along with two models  $\mathcal{M}_m$  and  $\mathcal{M}_n$  from which they are assumed to have arisen. The Bayes Factor captures the updating of the relative prior probability (prior odds) of the two models in light of the data

$$\frac{\overbrace{P(\mathcal{M}_m | y)}^{\text{Posterior Odds}}}{\overbrace{P(\mathcal{M}_n | y)}^{\text{Posterior Odds}}} = \frac{\overbrace{P(y | \mathcal{M}_m)}^{\text{Bayes Factor}}}{\overbrace{P(y | \mathcal{M}_n)}^{\text{Bayes Factor}}} \times \frac{\overbrace{P(\mathcal{M}_m)}^{\text{Prior Odds}}}{\overbrace{P(\mathcal{M}_n)}^{\text{Prior Odds}}} \tag{1}$$

The Bayes Factor (BF) is a ratio of *marginal likelihoods*, which together quantify the relative support for each model that is provided by the data.  $BF_{m,n}$  values larger than 1 indicate a support for model  $\mathcal{M}_m$ , whereas values between 0 and 1 indicate a support for model  $\mathcal{M}_n$  (on a logarithmic scale, positive and negative  $\log(BF_{m,n})$  indicate support for models  $\mathcal{M}_m$  and  $\mathcal{M}_n$ , respectively). For any model  $\mathcal{M}$ , the marginal likelihood corresponds to a weighted average of the likelihood of the data, with weights given by the parameter priors  $g(\theta)$ :

$$P(y | \mathcal{M}) = \int f_{\mathcal{M}}(y | \theta)g(\theta) d\theta. \tag{2}$$

<sup>4</sup>Singmann and Kellen (2019) provide an introduction to LMMs and the R package `aex`, which consolidates LMM application and frequentist statistical inference for experimental designs. Package `aex` is built on top of `lme4` but uses orthogonal contrasts per default and provides hypothesis tests and  $p$ -values for each of the model terms (i.e., main effects and interactions). For more details, see Singmann and Kellen (2019).


Spelling out the complete definition of a Bayes Factor allows us to highlight two important points: First, the Bayes Factor represents an “adjustment” to a set of prior beliefs based on data; just like a  $p$  value does not represent the probability of the null hypothesis, a Bayes Factor does not tell us the posterior odds in favor of one model over another, except in the special case that the prior odds were one (i.e., the two models were believed equally probable *a priori*). Second, the Bayes Factor marginalizes over the prior distributions on the parameters of each model  $\theta$ . As a consequence, the specification of these priors is not arbitrary—it is a core feature of the models being compared.

In the context of LMMs, Bayes Factors can be used to contrast nested models that assume different sets of fixed and/or random effects, as illustrated in vDAHSW. They computed “default Bayes Factors” using different LMMs and using data constructed under different aggregation regimes. These Bayes Factors were computed for two different kinds of model comparisons:

- **Strict Null Comparison:** The full or nesting model  $\mathcal{M}_A$  included both fixed and random slopes whereas the null or nested model  $\mathcal{M}_0$  included neither (Rouder et al., 2016). Both models included fixed and random intercepts.
- **Balanced Null Comparison:** Equivalent to the comparison above, with the exception that the null model  $\mathcal{M}_0$  included the random slopes (but not the fixed slope; see Barr et al., 2013).

Figure 2 illustrates the results obtained by vDAHSW with one of their data sets. The data in question were simulated from a model in which an effect of `condition` is present. This fact is reflected in the positive  $\log(\text{BF}_{A,0})$  values obtained across the board, all supporting the alternative model  $\mathcal{M}_A$ . However, we also see rather drastic differences between the two types of null comparisons. In the case of the strict null comparisons, in which the null model includes neither the fixed nor random slopes, the Bayes Factors are generally more extreme and become even more so as the number of observations being aggregated decreases. In contrast, aggregation has a negligible impact on the balanced null comparisons, in which the null model does include the random slopes (but not fixed slope). Based on this pattern of results (which hold for the other examples they considered), vDAHSW raised a number of questions (see their Table 3), which we distill into the two questions below:<sup>5</sup>

<sup>5</sup>vDAHSW also included an analysis of the completely aggregated (i.e., only one trials per condition remained for analysis). However, as this analysis somewhat falls outside of the LMM framework, so we omit it here.



**Fig. 2** Illustration of results from vDAHSW. Different  $\log(\text{BF}_{A,0})$  for the condition effect of the simulated data from vDAHSW (Example 2) that both shows a condition effect as well as non-zero random slopes. The two models for the “strict null comparison” differ in both fixed effects and random slopes (dashed line), the two models in the “balanced null comparison” only differ in their fixed effects (solid line). The  $x$ -axis shows the different aggregation regimes in terms of number of trials entered into the analysis (i.e., 100 trials refers to the full, disaggregated data; 50 trials refers to an analysis in which pairs of trials are aggregated, etc.). Model  $\mathcal{M}_A$  refers to the alternative (i.e., full) model, model  $\mathcal{M}_0$  refers to the null (i.e., reduced) model

- *Which models should be compared? That is, should a hypothesis test compare models which differ in both in terms of fixed effects and random effects, or only in terms of fixed effects?*
- *Should the data be aggregated if there are multiple observations per participant and cell of the design? And if so, how?*

As shown below, finding good answers to these questions is not a matter of finding a good “default”, but carefully considering the meaning of the model structure and parameters in the context of a *specific research application*. In fact, some of the issues raised below strongly suggest that a generic default approach is ultimately untenable.

## Question 1: Which Models Should Be Compared?

### Experimental Designs Involving One Factor

When using a LMM to analyze experimental data, a researcher is often hoping to answer the question, “did an experimental factor have an effect?” For example, the data in Fig. 2 shows an “effect” of `condition` on the dependent variable, and ideally a statistical model should

detect this effect. As discussed above, the full model for this design includes a fixed intercept, a fixed slope representing the effect of `condition`, as well as the corresponding by-participant random effect terms—the random intercept and the random slope. In a frequentist setting, Barr et al. (2013) showed that biased results are obtained if the individual differences in the effect of `condition` are unaccounted for by the LMM (i.e., if the model does not include the corresponding random slopes) (see also Schielzeth and Forstmeier, 2009). This general message also holds within a Bayesian setting (see Veríssimo, [this issue](#)). Therefore, it seems generally appropriate to designate the model with maximal random-effects as the alternative model  $\mathcal{M}_A$ . The remaining question then is which null model  $\mathcal{M}_0$  is most appropriate, when evaluating an effect of `condition`. vDAHSW discuss the two different variants introduced earlier, strict and balanced null comparisons, which differ in terms of whether or not they include by-participant random slopes for `condition`. Given that most experimental researchers appear to take the presence of individual differences as a given—they are just interested in the average effect (i.e., the fixed effect slope). Generally speaking, in such circumstances, the balanced null comparison strikes us as being the most reasonable one.

Contrary to this view, vDAHSW argue that the null model used in the balanced null comparison “seems implausible” (p. 12). In response, we raise two points: First, the support for  $\mathcal{M}_A$  in strict null comparisons is always *underqualified*. After all, this support might be due to (i) a non-zero fixed effect, (ii) the presence of individual differences (captured by the random slopes), or (iii) both. A “strict” null conflates these three different scenarios. Second, a null model including random slopes is often a plausible contender. For example, consider a study investigating ego depletion in a within-subjects design (e.g., Lin et al., 2020). In such a study, it seems possible to observe no overall effect of ego depletion (given the known difficulties in replicating this effect; e.g., Vohs et al., [in press](#)), but individual differences among the responses to the ego depletion manipulation. For instance, some participants might have heard of ego depletion and therefore act in line with the ego depletion hypothesis whereas other participants might believe the opposite to the ego depletion hypothesis, that initial exertion energizes (Savani & Job, 2017). If the proportion of participants for which these two opposing effects hold is *approximately* equal, then we should observe non-zero random slopes along with a mean difference of (approximately) zero. Such a scenario doesn’t strike us as being beyond the pale.<sup>6</sup> Of course, one could disagree

with our assessment by pointing out the implausibility of random effects cancelling each other out *exactly*. The problem with such an argument is that it can also be used to renders virtually any (point) null hypothesis implausible. For they state that some kind of effect attributed to a given manipulation is *exactly zero* (see Meehl, 1978). It is not clear to us how one could still defend any kind of statistical evaluation/testing of null hypotheses—the kind proposed by vDAHSW for instance—under these circumstances.

Despite our defense of balanced nulls, we can also think of a number of specific circumstances in which a strict null is perfectly sensible, as in the case of ESP (Bem, 2011). What is “plausible” depends on considering possible causal factors that may be at work in any specific application, and how best to represent these factors in a model that makes a number of abstractions (deliberate omissions) and idealizations (deliberate distortions; see Kellen, 2019). Contra vDAHSW and Rouder et al. (2016), we argue that the plausibility of such null models isn’t something that can be adjudicated in the abstract—a case-by-case assessment is necessary (for a related discussion, see Heathcote and Matzke, [this issue](#)). At the end of the day, researchers should have the ability to make a call on the null models that they consider to be plausible, as long as they are able to convincingly substantiate their decisions.

In cases where the appropriate null model is somehow a point of contention, the research question being posed might be better served by a model-selection effort involving more than two candidate models. In a Bayes Factor application, this would involve designating one of the models as the reference model and then calculating the Bayes Factor comparing it with all other candidates. For example, we might designate the model without fixed and random slopes for `condition` as the reference model and then calculate Bayes Factors with three further models: (i) a model that only includes the fixed slope of `condition` but not the corresponding random slopes, (ii) a model including the random slopes but no corresponding fixed slope, and (iii) a model that includes both fixed and random slopes. By computing the Bayes Factors associated with the three possible model pairings, one is able to get a fuller picture of how `condition` affects the dependent variable.

### Experimental Designs Involving Multiple Factors

The utility of comparing more than two models is amplified by Example 3 from vDAHSW. This example is concerned with an experimental design involving two crossed within-subjects factors. The maximal model includes (i) a fixed intercept, (ii) fixed main effects for the two factors as well as their interaction, (iii) a by-participant random intercept, and (iv) by-participant random slopes for both factors as well as their interaction.

<sup>6</sup>This is not a fringe view. As far as we know, all software solutions for implementing LMMs (e.g., R packages `lmerTest`, `car`, and `afex`) engage in balanced null comparisons.

A balanced null comparison *of the interaction* would involve a null model in which the fixed effect for the interaction is omitted. In the case *of main effects*, however, balanced null comparisons can be made in two different ways: On one hand, we can specify a null model in which everything but the fixed main effect being targeted is omitted (i.e., the fixed interaction remains). On the other, we could specify a null model in which the fixed main effect being targeted as well as the fixed interaction are omitted.<sup>7</sup> Deciding between these possible representations is made even more difficult by the fact that, in factorial designs, what constitutes a main effect and what constitutes an interaction is itself often an arbitrary choice (e.g., Brauer & Judd, 2000).

Once again, rather than commit to a single pair of models under consideration, the researcher may be better served by a model-selection exercise. In the case of two factors for which the mapping of factors is not arbitrary this could result in five different structures (all including the fixed intercept): (i) intercept only, (ii) main effect A, (iii) main effect B, (iv) both main effects, and (v) both main effects plus interaction. If the maximal model supports random slopes for all fixed effects, we could combine all possible fixed-effect structures with all possible random-effect structures, resulting in a total of  $5 \times 5 = 25$  models. For example, Singmann et al. (2014, Table 2) used this approach in a design involving crossed random effects for participants and items. In line with the notion that it is sometimes reasonable to allow for random slopes in the absence of their corresponding fixed effects, the model that was most strongly supported included random slopes for one factor but not the corresponding fixed effect.

## Summary

In the context of LMMs, inferences are less prone to bias when models attempt to respect the potential causal factors at work by representing them as random effects, leading to so-called “balanced” null comparisons. This is the standard approach in frequentist implementations of LMMs. However, the specific model structure should be decided on a case-by-case basis. In addition, Bayes Factors can be used to evaluate more than two models, providing us with the means to conveniently investigate which candidate in a set provides a better description of the data. When establishing a set of candidates, we urge researchers to disregard blanket

statements regarding which models are plausible versus not. For example, it is possible for an experimental manipulation to produce individual-level effects that cancel each other out, resulting in a virtually zero fixed effect. Finally, in multi-factor designs, what stands for a main effect or an interaction is often arbitrary, and deciding how to represent these terms requires recourse to a (hypothesized) causal structure.


## Question 2: Should Data Be Aggregated?

vDAHSW give examples in which aggregating data prior to modeling apparently leads to different inferences than when modeling the unaggregated data. A visual inspection of Fig. 2 might suggest that only the Bayes Factors for strict null comparisons are sensitive to aggregation, but as shown in Fig. 3, even the balanced null comparison exhibits a non-monotonic pattern. This figure also shows the Bayes Factors obtained with an alternative LMM (see dashed line). Importantly, these alternative Bayes Factors are invariant to aggregation.<sup>8</sup>

When asking about the effect of aggregation, vDAHSW overlook a fundamental question: *why is aggregation impacting Bayes Factors in the first place?* This omission implies that one should somehow try to remedy or solve the issue at hand without first knowing what is behind it. Fortunately, we are not dealing with a “black box” problem but transparent formal models (Ulrich, 2009). When investigating the inner workings of the LMMs employed by vDAHSW, one can clearly identify the culprit: the specification of Bayes Factors at the level of *standardized effects*—default Bayes Factors (Ly et al., 2016). For when Bayes Factors are specified at the level of unstandardized effects, they are no longer affected by aggregation biases, as demonstrated by our reanalysis of vDAHSW’s data (see Fig. 3, dashed line). In the case of default Bayes Factors, aggregation implicitly affects the priors embedded in the models being compared, effectively changing the meaning of the model parameters. vDAHSW appear to think that these model specifications can be treated as equivalent, as

<sup>7</sup>These two possibilities correspond to what is typically referred to as Type III and Type II hypothesis tests (e.g., Singmann & Kellen, 2019). There are a number of vocal critics arguing against the use of Type III tests (most notably Venables, 1998). Interestingly, one of the main arguments being given is that models that include interactions without all of its constituting main effects are “implausible” (Rouder et al., 2016).

<sup>8</sup>All Bayesian models discussed in this section were estimated through Stan (Carpenter et al., 2017) using different add-on packages to brms (Bürkner, 2017). The standardized models were estimated using package bfrms (Singmann & Gronau, 2021), which allows us to estimate the models underlying Rouder et al.’s (2012) default Bayes Factors using Stan and produces results that are equivalent to the BayesFactor package (Morey & Rouder, 2018). The unstandardized models were estimated using package stanova (Singmann, 2021). The mathematical formulation of the standardized and unstandardized model can also be found in the Appendix. All Bayes Factors were estimated using bridgesampling (Gronau et al., 2020). The frequentist LMMs were estimated using afex (Singmann et al., 2021) using the Kenward-Roger method for obtaining *p*-values. The full R code reproducing all results and figures reported here is available on the OSF: <https://osf.io/bu9hp/>.



**Fig. 3** Two Bayes Factor variants for models differing in fixed effects. Different  $\log(BF_{A,0})$  for the condition effect of the simulated data from vDAHSW (Example 2) in which the only difference between both models is the presence or absence of the fixed effect (i.e., both models contain the corresponding random slopes). The solid line shows the default Bayes Factor (Rouder et al., 2012) used by van Doorn et al. in which the observed effect is standardized by the residual SD and which is affected by the level of aggregation in a non-monotonic way. The dashed line shows a Bayes Factor in which the effect is not standardized and which is not affected by the level of aggregation

evidenced by the fact that they *formally describe* models in terms of unstandardized effects (see their Eqs. 1–6) but end up *employing* their standardized counterparts in the actual analyses. As shown in Fig. 3, these alternative specifications are *not* equivalent.

**The Standardized Model: Default Bayes Factors**

As described in Eq. 2, the marginal likelihood—the basis of the Bayes Factor—is the average likelihood weighted by the parameter priors  $g(\theta)$ . One consequence of this weighting is that different priors can lead to radically different Bayes Factors. This feature can complicate things in practice, given that it introduces a considerable degree of “subjectivity”. vDAHSW attempt to remedy this issue by using *default Bayes Factors* (Ly et al., 2016). In order to compute default Bayes Factors, models are specified in such a way that parameters representing effects—and their associated priors—are cast on a standardized scale (Rouder et al., 2012). What this means is that the effect-parameter priors included in  $g(\theta)$  do not speak directly to the scale of the “raw” dependent variable. This detachment allows default priors to be used regardless of the original measurement scale, ostensibly removing some

of the subjectivity that is characteristic of Bayes Factor applications.

To enable this this specific model parameterization, it is necessary to define the scaling factor for the observed effects. Rouder et al. (2012) adopted the residual standard deviation  $\sigma_\epsilon$  for this role. Specifically, consider the case of a mixed model with fixed-effects model matrix  $X$ , fixed-effects parameter vector  $\theta$ , random effects model matrix  $Z$ , and random-effects parameter vector  $b$ . The linear prediction for the vector of the dependent variable  $y$  is then given by<sup>9</sup>

$$y = \mu + \sigma_\epsilon(X\theta + Zb) + \epsilon. \tag{3}$$

The main difference from the way in which LMMs are typically formulated is that  $\sigma_\epsilon$  is part of the linear prediction of the model (e.g., Singmann and Kellen, 2019).

The problem with these default Bayes Factors is that  $\sigma_\epsilon$  is only one of the many different variance components being estimated. In addition to the residual standard deviation, every random effect term in the model (i.e., each column in  $Z$ ) has its own variance component (see also Rights & Sterba, 2019). In the example data considered here, the maximal LMMs estimate three variance components, (i) the residual standard deviation, (ii) the standard deviation of the random intercept, and (iii) the standard deviation of the random slopes. The default Bayes Factor approach proposed by Rouder et al. (2012) ignores these additional variance components, standardizing the observed effect solely on the basis of the residual standard deviation  $\sigma_\epsilon$ .

**The Unstandardized Model**

As an alternative to default Bayes Factors, we calculated Bayes Factors using a pair of models in which the parameter priors are specified on an unstandardized scale matching the original scale of measurement (i.e., the linear prediction is given by Eq. 3 without  $\sigma_\epsilon$ ). In these models, we followed Rouder et al.’s (2012) contrast coding of the fixed effects—that is, we used the orthonormal contrasts with contrasts weights of  $\pm 1/\sqrt{2}$  for the `condition` effect. The prior on this effect was a zero-centered  $t$  distribution with  $df = 3$  and scale =  $1/\sqrt{2}$ . This implies a  $t$ -distribution prior for the difference between the two conditions with  $df = 3$  and scale = 1. In other words, under the alternative hypothesis we expected the difference between the two conditions to be between  $-1$  and  $1$  in units of the dependent variable with  $\approx 60\%$  probability.

<sup>9</sup>In order to facilitate reading and comprehension, we decide to rely on a simplified (but ultimately incomplete) specification of model in the main text. A detailed description of the models can be found in the [Appendix](#).

### Model Posteriors

Figure 4 shows the mean and 95% credibility intervals of the posterior distributions of different model components for the alternative model  $\mathcal{M}_A$  underlying the standardized model (i.e., the default Bayes Factor, solid line) and the unstandardized model (dashed line). These posteriors distributions explain why the Bayes Factor of the standardized model is affected by different levels of aggregation and also why this is not the case for the unstandardized model.

As shown in Fig. 4C (left most panel) aggregation has exactly the same effect on  $\sigma_\epsilon$  in both models. Under an intense data aggregation regime (i.e., 2 trials),  $\sigma_\epsilon$  takes on a lower value. This value increases when there is no aggregation and all 100 trials per condition are considered. Aggregation has no effect on any other model component of the unstandardized model. Both additional variance components—the standard deviations of the random intercepts and slopes—as well as the estimated mean difference between the two conditions (see Fig. 4A) are unaffected by aggregation. This pattern is to be expected, given that the observed mean difference, as well as individual differences in the overall mean and condition effects are the same across the different levels of aggregation. Which is exactly what the estimates of the unstandardized model show.

In contrast, all components of the standardized model are affected by aggregation. To understand why, note that in the formulation for the linear predictor of the standardized model (3), all estimated coefficients (with the exception of the fixed intercept) are scaled by  $\sigma_\epsilon$ . This means that, when  $\sigma_\epsilon$  becomes smaller, the estimated coefficients need

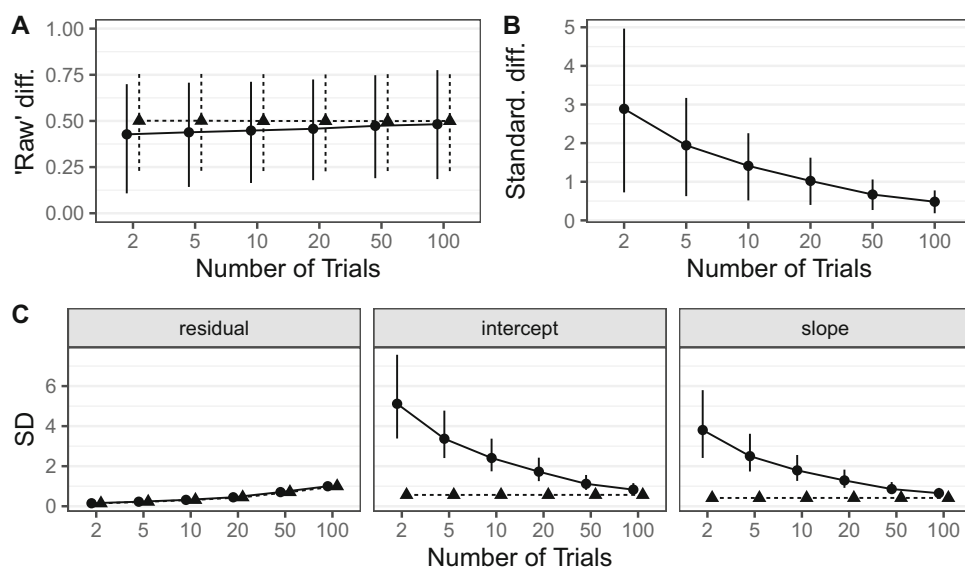
to become larger in order for the model to preserve its predictions. Given that  $\sigma_\epsilon$  scales all the random components being estimated (i.e., including the random intercept), the estimates of their respective standard deviations will be larger when  $\sigma_\epsilon$  is reduced.

Figure 4B shows the standardized difference—the standardized effect size estimate—which is only part of the standardized model and is given by the “raw” difference divided by  $\sigma_\epsilon$ . As the raw difference is only mildly affected by the number of trials compared to  $\sigma_\epsilon$ , we see the same pattern: the estimated standardized effect size is largest when  $\sigma_\epsilon$  is smallest. The estimates of the standardized difference also explain why the estimated “raw” differences in Fig. 4A are affected by the aggregation regime. The reason is that the prior is specified on the standardized-effect size scale and fixed across the different aggregation levels (it is a zero-centred Cauchy distribution with scale = 0.5). Therefore, the larger the estimated standardized difference, the larger the shrinkage due to the prior. Note that the model is closest to “true” mean difference of 0.5 only when there is no aggregation.

### Explaining the Non-monotonic Bayes Factors

The only result left unexplained is the non-monotonic pattern observed in Fig. 3. The reason behind this pattern is that the posterior of the standardized mean difference exhibits two patterns that affect the estimated Bayes Factor independently. Their interplay leads to the observed non-monotonicity. As shown in Figure 4B, the level of aggregation affects the mean of the standardized difference distribution, as well as its width.

**Fig. 4** Parameter estimates of standardized and unstandardized model. Data points show the posterior means and error bars the 95% credibility intervals of the posterior distributions of different model components for the standardized (i.e., Rouder et al.,’s default Bayes Factor) and unstandardized model. A: Estimated difference between both conditions on the original scale. B: Estimated standardized difference between the conditions, which is only part of the standardized model. C: Estimates of the three variance components estimated for this data set (residual = residual SD, intercept = random intercept SD, slope = random slope SD)



Model  $\bullet$  standardized  $\blacktriangle$  unstandardized



In the present case, the computed Bayes Factor is equivalent to the Bayes Factor obtained through the Savage-Dickey method. That is, the ratio of posterior density versus prior density for an estimated standardized mean difference of 0; i.e.,  $BF_{A,0} = \frac{f_{\text{prior}}(0)}{f_{\text{posterior}}(0)}$ . This Bayes Factor will increase along with the size of the standardized difference, for a fixed posterior-distribution width. Conversely, for a fixed standardized mean difference, increasing the width of the posterior distribution will decrease the Bayes Factor. Figure 5 plots both prior and posterior distribution of the standardized mean difference, and shows how these two components can conspire to produce a non-monotonic pattern. When the aggregation is maximal (number of trials = 2) the width of the posterior is largest with the consequence that the density at 0 is comparatively large even though the mean difference is also quite extreme. Decreasing the level of aggregation leads to a tighter posterior with a lower density at 0. This lower density is counteracted by a decrease in the standardized mean difference, which ultimately results in the non-monotonic pattern shown in Fig. 3.

The Bayes Factors for the unstandardized models can also be obtained through the Savage-Dickey method. In this case, the relevant posterior refers to the raw mean difference in Fig. 4A. As this posterior is invariant to the level of aggregation, the resulting Bayes Factor is also constant. This invariance to aggregation is not something that is unique to the unstandardized Bayes Factor. When estimating a frequentist LMM and  $p$ -values across the different levels of aggregation we observe the same inferential results in each case,  $F(1, 19) = 16.27, p < .001$ . Likewise, the fixed-effect coefficient


representing the difference, its standard error, and the random-effect variance component estimates are also unaffected by aggregation—only the residual standard deviation is affected. Parameterizing the Bayesian LMMs in the same way as their frequentist counterparts leads to similar results.

### An Answer to van Doorn et al.'s Question

The analysis above explicates the relationship between the standardized model and the default Bayes Factor, and how they are affected by different levels of data aggregation. One might infer from this analysis, together with vDAHSW's work, that the use of default Bayes Factors is generally recommended when dealing with disaggregated data. We wish to push back on this interpretation: The extent to which this approach is reasonable ultimately depends on how sensible (i.e., meaningful or interpretable) the standard effect size is.

For the disaggregated data of vDAHSW, the LMM has both by-participant random intercepts and random slopes, with the residual variance representing the within-participant trial-by-trial variance (which is assumed to be the same for all participants). It follows that the standardized effect size and the corresponding prior reflect a difference between conditions that is standardized by the level of within-participant variability. These standardized differences are therefore represented in terms of within-participant variance units. It is perfectly reasonable to assume that researchers can have enough substantive knowledge to specify sensible priors in terms of such units. Aggregating trials changes the interpretation of the residual

**Fig. 5** Prior (line) and posterior (histogram) of standardized mean difference across different numbers of trials. The standardized Bayes Factor is given by the ratio of posterior density and prior density at 0. To show this more clearly, the x-axis is restricted in range



variance and thereby the interpretation of the standardized effect size. For instance, in the example case in which two trials are aggregated (number of trials = 50), the residual variance reflects the within-participant variability that is associated to the mean of *trial pairs*. We cannot see how researchers could ever be expected to specify sensible priors in terms of these scale units, and therefore cannot recommend the use of default Bayes Factors here.

However, it is important to understand that the use of default Bayes Factors on disaggregated data *is not immune to interpretability problems*. After all, there are numerous circumstances in which the residual variance terms are difficult to interpret. For instance, difficulties arise when the other random effects variances differ in terms of how well they can be estimated: In an LMM, hierarchical shrinkage is imposed over the individual-level deviations, which in turn reduces the variance estimates (e.g., Baayen, 2008). This means that if the data is relatively sparse, and therefore shrinkage is considerable, then the residual variance will not exclusively capture within-participant trial-by-trial variability. It will also capture some of the other variance components, which were not estimated with enough precision. In such situations, the residual variance estimate will be biased upwards, which in turn leads to an underestimation of the standardized effect size. In other words, whenever the random-effect terms are affected by non-negligible shrinkage, the scale of the standardized effect size will be “contaminated” by other variance components, which affects their interpretability and therefore the status of any default Bayes Factor obtained with it.

A second problematic scenario involves LMMs with more than one random-effect grouping factor. In the simulated data used by vDAHSW, all random-effects terms are by-participant terms, which makes the interpretation of the residual variance term clear. However, one of the main reasons for employing LMMs is their ability to simultaneously account for multiple sources of stochasticity using so-called *crossed random effects* (Baayen, 2008; Judd et al., 2012, 2017). In a model with crossed random effects, the residual variance term reflects all the variability that is not accounted for by the other random-effects terms. For example, in a model with by-participant and by-item random effects terms, the residual variance can be interpreted as the within-participant variability *after accounting for all item-related variability*. At first glance, this residual variance estimate can strike one as being a more principled option in the sense that the standardization obtained with it is “more pure” by virtue of excluding item-related variability. The problem is that this standardization, if used responsibly, demands an amount of background knowledge that researchers rarely have. Specifically, it requires researchers to have a strong grasp on the expected magnitude of

item-related variability, which is something that can only be achieved if one has an extensive background knowledge on some fixed experimental design. Otherwise, it becomes quite difficult, if not impossible, to assess the plausibility of any prior over standardized effects. Although it is certainly possible for researchers to obtain such background knowledge, its requirement flies in the face of the main motivation behind standardized effect sizes—specifying effects in a general way that can be easily deployed.

A similar problem arises when including additional fixed-effects covariates: Their inclusion will reduce the residual variance and therefore lead to inflated standardized effect sizes that are difficult to compare across studies that differ in terms of the covariates considered. This problem is well documented in the technical literature on effect sizes, in which we find methods that try to remedy it (Olejnik & Algina, 2003). Unfortunately, current methods for computing default Bayes Factors do not address this issue, which means that one risks obtaining inflated effects that are not easily generalizable. Again, this problem could be sidestepped if one has a strong grasp of the relative impact that different covariates have in terms of the overall variability, a requirement that flies in the face of the notion of default Bayes Factors as a widely applicable go-to procedure.

*Our answer, finally:* The only scenario in which we find the use of default Bayes Factors appropriate is when dealing with LMMs with only by-participant random effect terms, a maximal random effects structure (Barr et al., 2013) that can be estimated relatively precisely, disaggregated data, and no additional fixed-effect covariates. Only in this scenario does the residual variance term reflect the within-participant trial-by-trial variability—a sensible unit for a standardized effect size estimate. As soon as any one of the above-stated conditions is violated, the standardized effect size estimates lose their clear interpretation, leading to Bayes Factors of questionable status. Because of these issues, we generally recommend computing Bayes Factors for unstandardized LMMs. Although this is far less convenient, as it requires researchers to specify priors that speak directly to the “raw” dependent variable scales, it is the only sure way to avoid the problems we discussed above.<sup>10</sup>

<sup>10</sup>At the time of writing, there is one *practical* problem with our recommendation of only using unstandardized LMMs: To the best of our knowledge, there is no package for estimating Bayes Factors for unstandardized LMMs that is even remotely comparable to `BayesFactor` package, either in terms of efficiency or functionalities (Morey & Rouder, 2018). The results reported here were based on LMMs implemented with the general purpose sampler `Stan` (Carpenter et al., 2017), which required computation times that were at least an order of magnitude longer than when estimating the same models with the `BayesFactor` package. Our hope is that future updates will bestow the `BayesFactor` with the ability to estimate Bayes Factors for unstandardized models.

Before moving on, one aspect that we want to stress is that there is no “easy fix” for default Bayes Factors. For instance, one might be tempted to try to fix the problem by changing the standardization factor—replace the residual variance with some other variance term. In the simulated data of vDAHSW, it would be tempting to standardize the difference between conditions by the corresponding random-slope estimate. In other words, standardize the within-subject effect in terms of the across-participant variability of the effect. The problem with this approach is that all but the residual variance term are reduced by the hierarchical shrinkage, which means that the effect size estimates will be inflated. In fact, it is not uncommon for random-slope estimates to approach zero, which would result in absurdly large standardized effect size estimates. It is useful to demonstrate this problem with a concrete example: Let us consider the data reported by Freeman et al. (2010), which included a lexical-decision task with high- and low-frequency words ( $N = 25$ ). In this example, we focus on the response times for word stimuli and the effect of word frequency. As in the example of vDAHSW, the effect of word frequency is a within-subjects variable. Descriptively, the standardized effect size is 0.84 and the observed effect of word frequency is 0.074 (i.e., RTs for low-frequency words are around 75 ms slower than RTs for high-frequency words). The standard deviation for the observed difference is 0.088. If we apply an LMM to this data, the estimated effect is equal to the observed effect (i.e., no shrinkage at the level of fixed effects). However, due to the hierarchical shrinkage, the standard deviation (i.e., the random slope estimate of the frequency effect) is estimated to be only 0.027. Adopting this estimate as the standard would result in a standardized effect size of 2.69, which is absurdly large.

### Summary

Default Bayes Factors were developed with the purpose of avoiding the apparent subjectivity that is inherent to the specification of priors on the models being compared. In the case of LMMs, this approach can have unintended consequences for aggregation. By deciding whether or not to aggregate the individual data points, researchers can (strategically or unwittingly) affect the estimate of the residual standard deviation, and in turn the resulting Bayes Factor. In contrast, Bayes Factors or  $p$ -values obtained with unstandardized effects are invariant to aggregation. This illustrates how the adoption of a “default” assumption fundamentally changes the meaning of a model parameter, ultimately compromising the ability of the model to provide a clear, unambiguous characterization.

## Part II: Letting the Dog Wag Its Tail

vDAHSW’s general concern was how to best use a particular class of models, LMMs, to detect experimental effects and interactions within a mixed model design. On the one hand, the structure of LMMs provides a convenient way to characterize the data that disentangles the general consequences or “fixed effects” of varying conditions within an experimental design from the variability or “random effects” associated with different observational units such as people and items. On the other hand, the LMM characterization relies on components—linear regression coefficients—that may or may not refer to any theoretical construct of interest or how it would manifest in data. And when using Bayes Factors, assumptions must be made in order to arrive at sensible priors over these model components (i.e., parameters). The more these assumptions are adopted by “default”, the less likely they are to be meaningfully related to the constructs they are intended to represent. The adoption of “default null model comparisons” and “default priors” can lead to misleading inferences and (apparent) paradoxes.

To understand how models can move beyond defaults and contribute to broader scientific goals, we return to the continuum between *causal* and *descriptive* models (Fig. 1). The focus of vDAHSW was on descriptive modeling that typically falls under the purview of applied statistics. Statistics has been called the “science of defaults” (Gelman, 2014) and so it is not surprising that statistical models should try to find default assumptions that work well enough much of the time. But like the default settings on any piece of technology (e.g., audio balance or chair heights), the defaults are meant to be a starting point rather than an end to themselves. Even for models that are meant to be primarily descriptive, moving beyond defaults requires considering the links between theoretical constructs and observables, and between model components and the general causal mechanisms that they are meant to represent.

Consider an experiment in which participants are shown a list of words of different frequencies and then asked to discriminate between members of this set and other words. The dependent variable may be the proportion of correct classifications (i.e., accuracy). The goal of applying a descriptive model like a LMM to this experiment would be to identify whether accuracy changes as a function of frequency. Any “effects” on accuracy that are identified (e.g., by a Bayes Factor supporting inclusion of a parameter representing a difference) are of interest only to the extent that they measure the actions of causal mechanisms

related to the theoretical construct of “memory”.<sup>11</sup> It is the purpose of *causal* models to represent these mechanisms and how they generate patterns in data. For example, a causal model might represent the processes by which words are encoded in memory, such that words of different frequencies are recognized at different rates (e.g., the mirror effect; see Glanzer & Adams, 1985; Shiffrin & Steyvers, 1997). The scientific value of a causal model arises from how well it explains the effects and its ability to generalize to other situations in which the same mechanisms are presumed to be operating. Meanwhile, the scientific value of a descriptive model lies in its ability to inform the development of causal models by detecting effects that are indicative of the actions of more general causal mechanisms. This is why it is crucial that descriptive models be built and applied with an appreciation of the causal models they are meant to inform.<sup>12</sup>

In the remainder of Part II, we discuss two important considerations in the development of models that better serve a scientific purpose. First, we delve more deeply into how model components are related to observable quantities in terms of the model’s *coordination functions* (Kellen et al., 2021; Chang, 2004; Van Fraassen, 2008). A coordination function is just as much a part of a scientific model as the model structure or its priors and therefore should be carefully scrutinized. Second, an important factor when evaluating causal model is its generalizability to related scenarios. Measures of relative fit, which are designed to distinguish between models based on their descriptive ability, are not often suited to assess or promote generalizability of models with a causal purpose. By appreciating the connections between models and theory and how we can assess them, we help ensure that the dog wags its tail, rather than the other way around.

## Coordination Functions

Measurement requires the (conditional) acceptance of a theory from which a *coordination function* is derived. In this way, the theory, data, and experimental design are tightly interwoven and the interpretation of effects and

interactions are dependent upon this relationship. Consider the illustration in Fig. 6, in which an observable variable  $X$  is linked to a quantitative theoretical construct  $Y$  by means of a coordination function  $f$  (and its inverse  $f^{-1}$ ) such that  $Y = f(X)$ . The coordination function, which can take on many different forms, cannot be established empirically, as, by definition, we cannot observe instances of both the observable variable and the theoretical construct. The host of challenges surrounding the establishment of a functional relationship between observables and theoretical constructs has been referred to as the “*the problem of coordination*” (Kellen et al., 2021; see also Chang, 2004; Van Fraassen, 2008).

We highlight two problems that arise when “the problem of coordination” is ignored: (1) that different dependent variables or different levels of an independent variable may have different coordination functions, and (2) that even if there is only one coordination function, it need not be linear as assumed by LMMs. We illustrate both problems using the example of a pendulum. The period of a pendulum is the time taken for a complete swing from one position and back again. The amplitude of a swing is defined as the maximum angle subtended by the arm with respect to the pivot point. Because the period is largely independent of the amplitude when the amplitude is small (i.e., it is isochronous), a pendulum can be used to measure time and the coordination function is linear. Let  $N$  be the number of swings of a pendulum in time  $t$ . Let  $T$  be the period. Then  $N = f(t) = (\frac{k}{T})t$ , where  $k$  is a constant that depends upon the length of the pendulum arm and the strength of gravity.<sup>13</sup> The function  $f$  is the coordination function linking the count of the number of pendulum swings (the dependent variable) with time (the theoretical construct).

To illustrate the first problem, suppose we conduct an experiment involving two pendula with different arm lengths, keeping in mind that the true coordination is a function of pendula arm length. On one occasion, we set them swinging for, say, 100 beats of our pulse<sup>14</sup> and on another occasion, they swing for 200 beats. We then count the number of completed swings and plot the results in Fig. 7A. This shows main effects of both independent variables (arm length and number of pulse beats) and their interaction. We can assume that these results are supported by very large Bayes Factors. If we ignore the coordination and naïvely identify the number of swings with the passage of time, we would interpret these results to mean that more time has passed over 200 pulse beats than over 100, which is a sensible conclusion. We would also

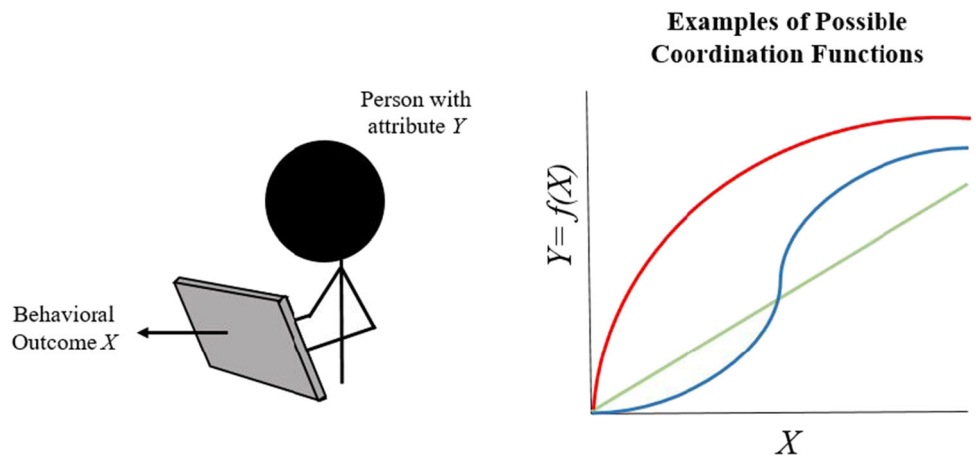
<sup>11</sup>Of course, it is possible for the dependent variable to be the focus of investigation itself and not act as a measure of any underlying construct, e.g., the number of bushels of corn produced on a specific field in a given year.

<sup>12</sup>In some cases, even the names given to different components of descriptive models like LMM can lead to confusion. For example, the role of the interaction terms in a LMM is to deal with deviations from additivity across the different factors of the experimental design. This understanding of “interaction” is not to be confused with the kinds of interactions that can be postulated by a causal model (e.g., inhibitory or facilitatory effects; see Cox & Shiffrin, 2017a, b, 2020).

<sup>13</sup>This relationship is only approximate as the effects of friction and elasticity of the pendulum need to be accounted for.

<sup>14</sup>In honor of Galileo who, it is said, established the isochrony of a pendulum in Pisa cathedral by timing it against his own pulse.

**Fig. 6** Illustration of the problem of coordination. In the left panel, we see an individual with theoretical attribute  $Y$  generating a behavioral outcome  $X$ . In the right panel, we illustrate some of the different possible shapes that the functional relationship between  $X$  and  $Y$  might take



conclude that more time has passed for the short pendulum than the long one under each pulse condition, which is absurd. Further this increase in “time” is greater for longer pulse sequences. Either we have created a clock that can manipulate time itself or we have made the mistake of ascribing the differences between the short and long pendula to the theoretical construct and not to the fact that they have different coordination functions.

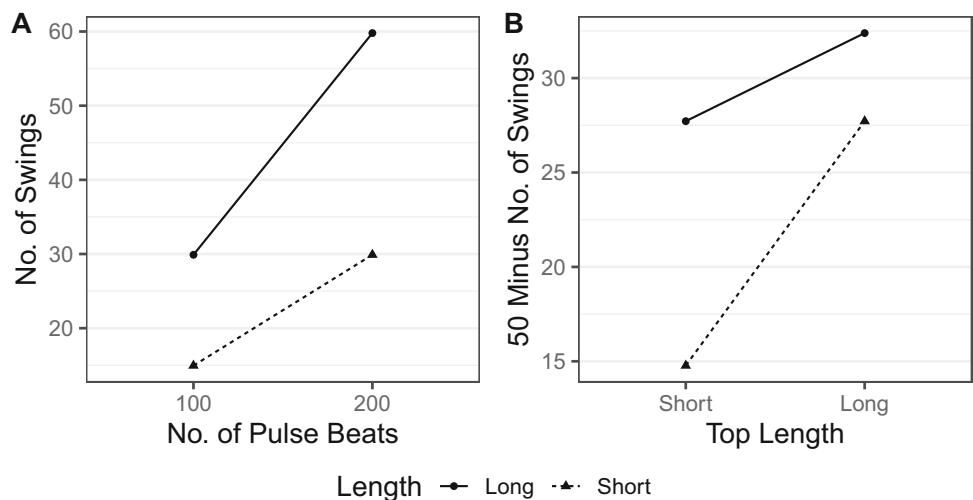
To illustrate the second problem, suppose we have learned from our earlier experiences that there is a negative relationship between the number of swings of a pendulum over a time interval and the length of its arm (see Fig. 7A) but, retaining some residual naivety, we choose as our dependent variable the difference between the number of swings and a large value (50 in this case). We then decide to investigate the effect of two independent variables on what we presume to be a measure of pendulum length. We attach either a short or long extension to the top of the pendulum arm (where it adjoins the pivot point) and either a short or long extension to the bottom of the pendulum arm (where it

swings freely). What is the effect of these variables on our measure of the total length of the arm?

Examining Fig. 7B, we see that there is an effect of adding a length to both the top and bottom ends of the arm as well as an interaction. Apparently, the effect on total length of adding an extension to the bottom of the arm is systematically reduced by the length of the extension added to the top of the arm. *Publication awaits!* The absurdity of the conclusion rests on the fact that we have ignored the possibility of a nonlinear coordination function. In fact, the number of swings ( $N$ ) is a function of the reciprocal of the square root of the arm length ( $L$ ). That is,  $N = K/\sqrt{L}$ , where  $K$  is a constant that depends on the total length of time and the strength of gravity.

As these simple examples illustrate, a purely descriptive statistical framework predicated on the identification of main effects and interactions, such as the one laid out by vDAHSW, glosses over the problem of coordination (cf. Loftus, 1978; Wagenmakers et al., 2012). This creates the risk of the components of LMMs being illegitimately

**Fig. 7** A) The number of pendulum swings as a function of pendulum length and number of pulse beats. B) A transformation of the number of pendulum swings (50 minus the number of swings) used as a measure of the total length of the pendulum arm, as an apparent function of the length of the pendulum arm and the length of an extension added to the top of the arm



reified as the “things” that researchers are ultimately targeting. The only way to avoid this risk is to place the understanding of theoretical constructs and their presumed coordinations at the forefront, a move that runs counter to the adoption of general defaults.<sup>15</sup> Doing so amounts to letting causal modeling considerations inform descriptive statistical models.<sup>16</sup>

### Assessing Model Generalizability (in Its Different Senses)

Modeling in science, whether primarily causal or descriptive, is ultimately in the service of the production of theories that can explain patterns of data using causal mechanisms that are expected to *generalize* across similar scenarios. Bayes Factors are often promoted as a method for assessing model generalizability in terms of the ability to predict new, unseen data *from the same generating process*.<sup>17</sup> This form of generalizability is, however, inextricably linked to the goals of *descriptive* rather than causal modeling. That is, the Bayes Factor assesses the relative ability of two models to predict (assign high likelihood to) data  $y$  obtained from a particular design/population, when considering the range of predictions associated with the parameter values deemed

<sup>15</sup>It should be noted that the problem of coordination can be extended to the way in which we conceptualize “noise”, “variability”, or “error” at the within and between subject levels (see Cavagnaro & Davis-Stober, 2014; Cavagnaro & Davis-Stober, 2018; Regenwetter & Davis-Stober, 2018).

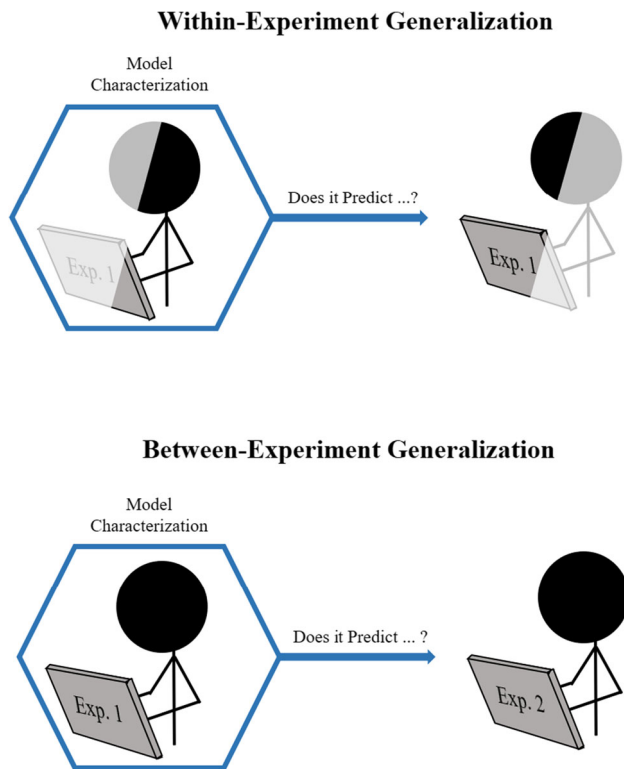
<sup>16</sup>One of the reasons researchers often give for using linear models is that they provide a convenient first approximation. Although we see this as a perfectly legitimate justification, and we do not wish to make any strong normative statements about what researchers can and cannot do (there is enough of that going around), we find this use of linear models to be somewhat “ill-fitting”. Many predictions in psychology are ordinal, in the sense that they can establish a system of inequalities. Importantly, these ordinal predictions cannot be conveniently decomposed in terms of separable main effects and interactions—both kinds of effects are part and parcel of the order constraints being imposed. It seems to us that the use of order-constrained inference methods would provide a much more convenient first approximation (see Kellen et al., 2021).

<sup>17</sup>Proponents of Bayes Factors often refer that this assessment of generalizability involves an implicit Occam’s razor that penalizes overly flexible models. To achieve this, the Bayes Factor is defined as the ratio of marginal likelihoods, which is essentially the posterior odds of the models when their prior odds are fixed at 1:1 (Shiffrin et al., 2016). But not only is this ratio arbitrary and incoherent with respect to Bayesian theory (e.g., it would set the prior odds of a very unlikely ESP hypothesis to be equal to a No ESP hypothesis), it also does not necessarily implement any kind of Occam’s razor. Due to its extreme sensitivity to priors over parameters, the Occam’s razor can be removed, and the resulting Bayes Factor over the models can be influenced by merely changing the priors over their parameters (Rasmussen & Ghahramani, 2001; Robert, 2016). A more ideal and robust scenario would be one in which an explicit notion of Occam’s razor can be defined and priors over models can be accounted for (Shiffrin et al., 2016).

probable a priori (i.e.,  $g(\theta)$ ). In other words, Bayes Factors refer to generalizability in the *specific sense* of how general the prediction of what we just observed is under each model. For any given model, the more likely  $y$  is on average—when sampling parameter values from  $g(\theta)$ —the more generalizable that model will be deemed. This is a restrictive form of generalizability that does not encompass the many other ways in which generalizations are made in science (for a related discussion, see Liu & Aitkin, 2008).

Any time we design an experiment, we engage in an act of simplification. We isolate and manipulate a small set of variables that may be relevant to a phenomenon, we control the stimuli presented to participants, and we restrict the manner in which those participants may respond. We do so to make our scientific task tractable and/or facilitate access to a phenomena of interest. Our hope is that, by designing many such experiments, each simplifying the phenomenon in different ways, we can construct causal models that simultaneously accommodate all experiments and by doing so, provide a better theory-based characterization. But for most causal models in science, we are less interested in how well it accounts for the data *that we have just observed* than we are in how well it might accommodate new data from a different experiment that relates to the same phenomenon. A different experiment that perhaps measures or simplifies it in a distinct way (Busemeyer & Wang, 2000; Schat et al., 2020) or engages with it through a different set of experimental manipulations (Baribault et al., 2018). In this sense, *every meaningful scientific inference is necessarily an act of generalization*. Even within the same experiment for a given dataset, we may be for seeking for good generalization and predictive performance at different levels. For instance, we may care about the generalization with respect to an individual for a particular session and a particular task, or we may be interested in generalization among particular groups of participants across sessions and tasks, or we may be interested in making estimates at the population level, and so on.<sup>18</sup> Figure 8 illustrates a couple of examples. Depending on the context, we may adjust the way we compare model predictions with data, which can range from coarse qualitative patterns to fine-grained quantitative predictions (Shiffrin & Steyvers, 1997). To fully engage with generalizability in all its relevant senses, researchers need a large methodological toolbox as well as the substantive background knowledge that allows for sensible decisions to be made—nothing that could ever be replaced by some default application of Bayes Factors.

<sup>18</sup>These generalizations often focus on the role of certain theoretical constructs across *different* data-generating processes. In contrast, Bayes Factors focus on generalizability with respect to the same data-generating process.



**Fig. 8** Two examples of generalization. In this example of “within-experiment generalization”, a model characterization obtained with a portion of the data observed in Experiment 1 (Exp. 1) is used to predict the remainder of the data. In the example of “between-experiment generalization”, a model characterization obtained with the (complete) data observed in Experiment 1 is used to predict the data observed in Experiment 2 (Exp. 2)

But even when restricting ourselves to generalizability as understood in the context of Bayes Factors, it is important to understand specific circumstances in which they are most informative. As is well known, Bayes Factors work best when there is at least one of the considered models that well approximates the data-generating process (often referred to as the “ $\mathcal{M}$ -closed” setting; Navarro, 2018).<sup>19</sup> In a scenario where a large amount of model misspecification is likely for both models (often referred to as the “ $\mathcal{M}$ -open” setting) it often behaves poorly. The large-sample behavior of the Bayes Factor in the latter  $\mathcal{M}$ -open setting is to select with absolute certainty that model which is closest to the truth in terms of Kullback-Leibler divergence, and as noted by Navarro (2018) this is not always well-matched to the practical goals of the scientist. In the pre-asymptotic case, recent results by Oelrich et al. (2020) highlight how this “overconfidence” can be dangerous: a Bayes Factor analysis can often appear to lend overwhelming support to a given model (when other model selection criteria do not) yet this extreme confidence

can rapidly shift with only a small amount of new data. More precisely, they derive results for the expected value and variance of the Bayes Factor in the linear regression setting (which, admittedly may differ somewhat in the LMM setting), and the conditions under which the Bayes Factor is most prone to overconfidence. Of particular note to psychologists who hope to make robust inferences from data, Theorem 1 from Oelrich et al. (2020) suggests that the Bayes Factor has highest variance (and is thus most prone to overconfidence) in situations where the models under consideration have little shared complexity, and explain the same phenomenon in vastly different ways.

From a theoretical perspective, this property of the Bayes Factor is dangerous. As scientists, we want our causal models to be highly distinguishable from one another and consequently design our experiments to *make* them distinguishable. That is, we choose experiments that render model predictions as different as possible and minimize shared complexity in the measurement context. We do so because we are hoping to end up in a situation in which one model is clearly better than the others. Sometimes we succeed, and the data are so overwhelmingly dispositive that every model-selection criterion gives similar answers (and accordingly we have no practical reason to prefer Bayes Factors over other criteria). However, when our efforts end up producing ambiguous results, the Bayes Factor can be prone to overconfidence, leading the researcher to falsely conclude that the results *are* decisive. As noted by Oelrich et al. (2020), this risk is most severe in situations when both models can only provide gross mischaracterizations, and yet there is a grain of truth in both: when two models are vastly different to each other but they’re both capturing something relevant to the phenomenon. Arguably this situation is the norm in psychological research, and yet the formal results reported by Oelrich et al. (2020) suggest that this situation is exactly when the Bayes Factors are the most overconfident, yielding extreme preferences for a specific model candidate (the one with minimal Kullback-Leibler divergence) that can swing wildly under minor changes to the experimental design or data.

Endeavors in mathematical psychology, and science more broadly, are concerned with building increasingly better models of the causal mechanisms underlying phenomena of interest. When trying to distinguish two equally well-performing models, there are two options that we can pursue: The first option is to *expand* the domain of application of the models by adding new variables to be explained (e.g., confidence as well as response times and accuracy) and/or new tasks to which they can be applied; we keep expanding the domain until the models become distinguishable once more. The second option is to *restrict* the domain of application by focusing on portions of the data for which the models do make diverging qualitative predictions

<sup>19</sup>*Minority Dissent*: Quentin Gronau would like to state that he disagrees that Bayes Factors can only be fruitfully used in the “ $\mathcal{M}$ -closed” setting (see Gronau & Wagenmakers, 2019).

that do not hinge on auxiliary assumptions (Birnbaum, 2010; Kellen et al., *in press*; Stephens et al., 2018). In either case, Bayes Factors—and quantitative model comparisons generally—are not very amenable to such a workflow. In addition to the illusion of confidence that a Bayes Factor can portray when dealing with equally good models, its handling becomes increasingly cumbersome as the models under consideration become more complex—not only does the computation of the Bayes Factor become increasingly hard, it also becomes increasingly sensitive to priors over additional parameters, as noted above. Treating the Bayes Factor as a “gold standard” for model selection can, in the worst case, discourage the use of more complex models which would in turn discourage researchers from expanding the domains of application. As a consequence, psychologists would be relegated to dealing exclusively with toy-like problems and simplistic studies rather than ecological contexts where models can be most impactful. As important as generalizability is to a successful causal model in science, many forms of generalizability cannot be assessed via quantitative model comparisons based on criteria like the Bayes Factor, which are tuned to discriminating between descriptive models.

## Final Thoughts

Ultimately, our message is simple: Don’t let the tail wag the dog. If we want descriptive modeling to serve the goals of science, critical thinking cannot be outsourced to statistical methods based on default assumptions. A main goal of science is to explain phenomena in terms of causal mechanisms that are expected to generalize across related scenarios. In service of this goal, we construct formal models which exist on a continuum from *descriptive* to *causal*. While causal models explicitly represent mechanisms and illustrate how they yield patterns in data, descriptive models are designed to identify those very same patterns. Statistical modeling, such as the LMMs illustrated by vDAHSW, serves a primarily descriptive function, but in order for inferences from descriptive models to aid scientific goals, their assumptions must be aligned with potential causal mechanisms. Adopting default assumptions impedes this alignment, leading to erroneous inferences and apparent paradoxes which are only resolved by considering how a model instantiates particular causal mechanisms. This instantiation requires one to think, among other things, about the coordination function that relates model components to observables. Finally, although statistical methods like Bayes Factors can distinguish between models in a *descriptive* sense, they are not often suited to distinguishing

between *causal* models which should be assessed based on broader criteria of generalizability.

In the last decade or so, there has been a tremendous shift toward the adoption of Bayesian methods. This shift is in part motivated by a perceived association between the use of frequentist methods and the prevalence of questionable research practices that lead to unreplicable and ungeneralizable results (e.g., Wagenmakers et al., 2011). However, there is a lingering concern that this shift will not bring an end to mindless “statistical rituals” (Gigerenzer, 2018), only that these rituals will be moved to another temple. This wouldn’t be surprising at all, when taking into consideration the numerous historical surveys arguing that psychology as a field as long been bound together by methodological standardization, and the belief that the latter is sufficient for the establishment of scientific enterprises (e.g., Danziger, 1994; Flis & van Eck, 1999; Mackenzie, 1977). One response to this concern is that Bayesian modeling requires a careful specification of one’s prior beliefs, and that the transition between prior and posterior beliefs can occur in a transparent manner. The hope is that, under the aegis of Bayes, the researcher will be strongly encouraged to think more carefully about models (e.g., their relationship with theoretical statements) and the evidence that can be derived from their contact with data. As much as we may hope that this shift in attitudes will come to pass, the reliance on default methods and standardized effects that has accompanied the recent surge in Bayesian modeling prevents us from being optimistic. The different issues raised by vDAHSW have provided us with a valuable opportunity to articulate some of our concerns in more concrete terms.

It would be a mistake to interpret our critique as an attack on Bayesian modeling in general, the use of Bayes Factors as an inferential tool, or vDAHSW for that matter. We believe that, with informed priors, Bayes Factors can help adjudicate questions about the relative descriptive value of different models. Moreover, we believe that Bayesian methods are ideally suited to representing and propagating uncertainty in a principled way; this is especially useful in causal modeling, where joint posterior distributions over model parameters can be helpful in understanding the relationships between components of a potentially complex model and data. Rather, our target here is the oft-tacit understanding of modeling as something that can be practiced fruitfully when distanced from the data and bound by constraints that are divorced from substantive theoretical considerations. Progress cannot be achieved by global, one-size-fits-all application of methods, regardless of their particular merits: stricter  $p$ -values, standardized



effect sizes,  $p_{\text{rep}}$ , Bayes Factors, multiverse analyses, pre-registrations, etc. Regardless of the statistical framework adopted, progress can only come from a careful (and often piecemeal) scrutiny of the relationship between the different theoretical components, experiments, and data, along with clear statements of goals, values, and compromises (see Mayo, 1996; Mayo, 2018; Navarro, 2018; Kellen, 2019). W. W. Rozeboom, one of the earliest critics of  $p$ -values and proponents of Bayesian reasoning (Rozeboom, 1960) summarized it best in a related discussion:

... it's important for those who advance to graduate science to learn that holistic accept/reject statistical-test appraisals of a target theory is no more creative science than warming a frozen supermarket dinner is gourmet cooking. An alpha-grade theoretical scientist should be able to search out and devise diagnostics for basic features in a probated theory's conceptual structure that haven't yet been linked to observable consequences, and to rejoice when new procedures of data production disclose hitherto unknown data patternings that deepen our abductive access to the explanatory sources of these phenomena. [...] astute evidence appraisal focuses on select features of the hypothesis at issue with only secondary confidence adjustments, if any, in its remainder. Holistic acceptance/rejection is for amateurs. Rozeboom (2008, p. 1123).

## Appendix Linear Mixed Models: Full Specification

For all models, factors are represented through the orthonormal contrasts introduced by Rouder et al. (2012). For a factor with two levels the two contrasts codes are  $\pm 1/\sqrt{2}$ .

Furthermore, we define the following quantities that occur in both models.

- $y$ : vector of response variable (i.e., dependent variable)
- $\sigma_\varepsilon$ : residual standard deviation
- $Z$ : random-effect model matrix (picks out particular random intercepts and random slopes for each participant; e.g., Bates et al., 2015)
- $b$ : vector of random-effect parameters (i.e., zero-centered individual-level random intercepts, followed by zero-centered individual-level random slopes, ...)
- $t(d, l, s)$ : student  $t$  distribution with degrees of freedom  $d$ , location  $l$ , and scale  $s$ .
- $I$ : Identity matrix

## Standardized Model

Model equation:

$$\begin{aligned} y &= \mu \mathbf{1} + \sigma_\varepsilon (X^* \theta^* + Zb) + \varepsilon \\ \varepsilon &\sim \mathcal{N}(\mathbf{0}, \sigma_\varepsilon^2 I) \\ \theta^* &\sim \mathcal{N}(\mathbf{0}, g_{\theta^*} I) \\ g_{\theta^*} &\sim \text{Inverse-Gamma}(1/2, r_{\text{fixed}}^2/2) \\ b &\sim \mathcal{N}(\mathbf{0}, \Sigma) \\ \Sigma &= \text{diag}(g_1 \mathbf{1}, g_2 \mathbf{1}) \\ g_1 &\sim \text{Inverse-Gamma}(1/2, r_{\text{random}}^2/2) \\ g_2 &\sim \text{Inverse-Gamma}(1/2, r_{\text{random}}^2/2) \\ p(\mu, \sigma_\varepsilon^2) &\propto 1/\sigma_\varepsilon^2 \end{aligned}$$

Additional notation:

- $\mu$ : intercept parameter
- $X^*$ : fixed-effect model matrix without the intercept column (i.e., for a model with one 2-level factor consists of only one column).
- $\theta^*$ : fixed-effect parameter vector without the intercept parameter (i.e.,  $\theta = (\mu, \theta^*)$ )
- $g_{\theta^*}, g_1, g_2$ : Variances of effects, see Rouder et al. (2012)
- $r_{\text{random}}, r_{\text{fixed}}$ : Prior scales

## Unstandardized Model

$$\begin{aligned} y &= X\theta + Zb + \varepsilon \\ \varepsilon &\sim \mathcal{N}(\mathbf{0}, \sigma_\varepsilon^2 I) \\ b &\sim \mathcal{N}(\mathbf{0}, \Sigma) \\ \Sigma &= \text{diag}(\sigma_1^2 \mathbf{1}, \sigma_2^2 \mathbf{1}) \\ \theta_1 &\sim t(3, 0, 2.5) \\ \theta_2 &\sim t(4, 0, 1/\sqrt{2}) \\ \sigma_1 &\sim t^+(3, 0, 2.5) \\ \sigma_2 &\sim t^+(3, 0, 2.5) \\ \sigma_\varepsilon &\sim t^+(3, 0, 2.5) \end{aligned}$$

Additional notation:

- $X$ : full fixed-effect model matrix consisting of two columns, intercept plus condition-specific contrast codes.
- $\theta$ : fixed-effect parameter vector that includes the intercept parameter
- $\theta_1$ : fixed-effect intercept parameter
- $\theta_2$ : fixed-effect condition-specific parameter

**Author Contribution** All authors contributed to the conceptualization and writing of the manuscript.

**Availability of Data and Material** Data and scripts can be found on Open Science Framework (OSF): <https://osf.io/bu9hpf/>.

**Code Availability** See OSF repository mentioned above.

## Declarations

**Conflict of Interest** The authors declare no competing interests.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Baayen, H. (2008). *Analyzing linguistic data : a practical introduction to statistics using r*. Cambridge University Press.
- Baayen, H., Davidson, D., & Bates, D. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59(4), 390–412. [10.1016/j.jml.2007.12.005](https://doi.org/10.1016/j.jml.2007.12.005).
- Baribault, B., Donkin, C., Little, D. R., Trueblood, J. S., Oravecz, Z., Ravenzwaaij, D. v., & Vandekerckhove, J. (2018). *Metastudies for robust tests of theory*. Proceedings of the National Academy of Sciences, 201708285. Retrieved 2018-05, from <http://www.pnas.org/content/early/2018/03/08/1708285114>, <https://doi.org/10.1073/pnas.1708285114>.
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255–278. <https://doi.org/10.1016/j.jml.2012.11.001>.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1). Retrieved from <https://doi.org/10.18637/jss.v067.i01>.
- Bem, D. J. (2011). Feeling the future: Experimental evidence for anomalous retroactive influences on cognition and affect. *Journal of Personality and Social Psychology*, 100(3), 407–425. Retrieved 2013-06, from <https://doi.org/10.1037/a0021524>.
- Birnbaum, M. H. (2010). Testing lexicographic semiorders as models of decision making: Priority dominance, integration, interaction, and transitivity. *Journal of Mathematical Psychology*, 54(4), 363–386. Retrieved 2021-06, from <https://www.sciencedirect.com/science/article/pii/S0022249610000416>, <https://doi.org/10.1016/j.jmp.2010.03.002>.
- Brauer, M., & Judd, C. M. (2000). Defining variables in relationship to other variables: When interactions suddenly turn out to be main effects. *Journal of Experimental Social Psychology*, 36(4), 410–423. Retrieved 2021-06, from <https://doi.org/10.1006/jesp.2000.1420>.
- Bürkner, P.-C. (2017). brms: An r package for bayesian multilevel models using stan. *Journal of Statistical Software*, 80(1), 1–28. <https://doi.org/10.18637/jss.v080.i01>.
- Busemeyer, J. R., & Wang, Y.-M. (2000). Model comparisons and model selections based on generalization criterion methodology. *Journal of Mathematical Psychology*, 44(1), 171–189.
- Carpenter, B., Gelman, A., Hoffman, M., Lee, D., Goodrich, B., Betancourt, M., & Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, 76(1), 1–32. [10.18637/jss.v076.i01](https://doi.org/10.18637/jss.v076.i01).
- Cavagnaro, D. R., & Davis-Stober, C. P. (2014). Transitive in our preferences, but transitive in different ways: An analysis of choice variability. *Decision*, 1(2), 102–122.
- Cavagnaro, D. R., & Davis-Stober, C. P. (2018). A model-based test for treatment effects with probabilistic classifications. *Psychological Methods*, 23(4), 672–689.
- Chang, H. (2004). *Inventing temperature: Measurement and scientific progress*. Oxford: Oxford University Press.
- Cox, G. E., & Criss, A. H. (2017b). Parallel interactive retrieval of item and associative information from event memory. *Cognitive Psychology*, 97, 31–61.
- Cox, G. E., & Criss, A. H. (2020). Similarity leads to correlated processing: A dynamic model of encoding and recognition of episodic associations. *Psychological Review*, 127, 792–828.
- Cox, G. E., & Shiffrin, R. M. (2017a). A dynamic approach to recognition memory. *Psychological Review*, 124(6), 795–860. <https://doi.org/10.1037/rev0000076>.
- Cox, G. E., & Shiffrin, R. M. (in press). Computational models of event memory. In M. J. Kahana, & A. Wagner (Eds.) *Oxford handbook of human memory*. Oxford University Press.
- Danziger, K. (1994). *Constructing the subject: Historical origins of psychological research*. Cambridge University Press.
- Flis, I., & van Eck, N. J. (1999). A large-scale term co-occurrence analysis of scientific literature in psychology. *History of Psychology*, 21, 334–362.
- Freeman, E., Heathcote, A., Chalmers, K., & Hockley, W. (2010). Item effects in recognition memory for words. *Journal of Memory and Language*, 62(1), 1–18. Retrieved 2014-08, from <https://doi.org/10.1016/j.jml.2009.09.004>.
- Gelman, A. (2014). How do we choose our default methods?. In X. Lin, C. Genest, D. L. Banks, G. Molenberghs, D. W. Scott, & J. L. Wang (Eds.) *ast, present, and future of statistical science* (pp. 291–299). CRC Press. (An optional note).
- Gigerenzer, G. (2018). Statistical rituals: The replication delusion and how we got there. *Advances in Methods and Practices in Psychological Science*, 1(2), 198–218. Retrieved 2018-07, from <https://doi.org/10.1177/2515245918771329>.
- Glanzer, M., & Adams, J. K. (1985). The mirror effect in recognition memory. *Memory & Cognition*, 13(1), 8–20.
- Gronau, Q. F., Singmann, H., & Wagenmakers, E.-J. (2020). Bridgesampling: An r package for estimating normalizing constants. *Journal of Statistical Software*, 92(10). Retrieved 2020-02, from <https://doi.org/10.18637/jss.v092.i10>.
- Gronau, Q. F., & Wagenmakers, E.-J. (2019). Limitations of bayesian leaveone-out cross-validation for model selection. *Computational Brain & Behavior*, 2(1), 1–11.
- Heathcote, A., & Matzke, D. (this issue). The limits of marginality. *Computational Brain and Behavior*.
- Jeffreys, H. (1961). *Theory of probability*, 3rd edn. Clarendon Press.
- Judd, C. M., Westfall, J., & Kenny, D. A. (2012). Treating stimuli as a random factor in social psychology: A new and comprehensive solution to a pervasive but largely ignored problem. *Journal of Personality and Social Psychology*, 103(1), 54–69. [10.1037/a0028347](https://doi.org/10.1037/a0028347).
- Judd, C. M., Westfall, J., & Kenny, D. A. (2017). Experiments with more than one random factor: Designs, analytic models, and statistical power. *Annual Review of Psychology*, 68(1), 601–625. Retrieved 2017-03, from <https://doi.org/10.1146/annurev-psych-122414-033702>.
- Kellen, D. (2019). A model hierarchy for psychological science. *Computational Brain & Behavior*, 2(3), 160–165.

- Kellen, D., Davis-Stober, C. P., Dunn, J. C., & Kalish, M. L. (2021). The problem of coordination and the pursuit of structural constraints in psychology. *Perspectives on Psychological Science*, 1745691620974771. Retrieved 2021-03, from <https://doi.org/10.1177/1745691620974771>. (Publisher: SAGE Publications Inc).
- Kellen, D., Winiger, S., Dunn, J. C., & Singmann, H. (in press). Testing the foundations of signal detection theory in recognition memory. *Psychological Review*. Retrieved from <https://doi.org/10.31234/osf.io/p5rj9>.
- Lin, H., Saunders, B., Friese, M., Evans, N. J., & Inzlicht, M. (2020). Strong effort manipulations reduce response caution: A preregistered reinvention of the ego-depletion paradigm. *Psychological Science*, 0956797620904990. Retrieved from <https://doi.org/10.1177/0956797620904990> (Publisher: SAGE Publications Inc.)
- Liu, C. C., & Aitkin, M. (2008). Bayes factors: Prior sensitivity and model generalizability. *Journal of Mathematical Psychology*, 52, 362–375.
- Loftus, G. R. (1978). On interpretation of interactions. *Memory & Cognition*, 6(3), 312–319. Retrieved 2021-06, from <https://doi.org/10.3758/BF03197461>.
- Luce, R. D. (1995). Four tensions concerning mathematical modeling in psychology. *Annual Review of Psychology*, 46, 1.
- Ly, A., Verhagen, J., & Wagenmakers, E.-J. (2016). Harold jeffreys's default bayes factor hypothesis tests: Explanation, extension, and application in psychology. *Journal of Mathematical Psychology*, 72, 19–32. Retrieved 2017-10, from <https://doi.org/10.1016/j.jmp.2015.06.004>.
- Mackenzie, B. D. (1977). *Behaviourism and the limits of scientific method*. Humanities Press.
- Mayo, D. G. (1996). *Error and the growth of experimental knowledge*. University of Chicago Press.
- Mayo, D. G. (2018). *Statistical inference as severe testing: How to get beyond the statistics wars*. Cambridge University Press.
- Meehl, P. E. (1978). Theoretical risks and tabular asterisks: Sir karl, sir ronald, and the slow progress of soft psychology. *Journal of Consulting and Clinical Psychology*, 46, 806–834. 10.1037//0022-006X.46.4.806.
- Morey, R. D., & Rouder, J. N. (2018). BayesFactor: Computation of bayes factors for common designs. (R package version 0.9.12-4.2. <https://CRAN.R-project.org/package=BayesFactor>).
- Navarro, D. J. (2018). Between the devil and the deep blue sea: Tensions between scientific judgement and statistical model selection. *Computational Brain and Behavior*. Retrieved from <https://doi.org/10.1007/s42113-018-0019-z>.
- Navarro, D. J. (2021). If mathematical psychology did not exist we might need to invent it: A comment on theory building in psychology. *Perspectives on Psychological Science*, 1745691620974769.
- Oelrich, O., Ding, S., Magnusson, M., Vehtari, A., & Villani, M. (2020). When are Bayesian model probabilities overconfident? arXiv:2003.04026.
- Olejník, S., & Algina, J. (2003). Generalized eta and omega squared statistics: Measures of effect size for some common research designs. *Psychological Methods*, 8(4), 434–447. <https://doi.org/10.1037/1082-989X.8.4.434>.
- Rasmussen, C. E., & Ghahramani, Z. (2001). Occam's razor. *Advances in neural information processing systems* (pp. 294–300).
- Regenwetter, M., & Davis-Stober, C. P. (2018). The role of independence and stationarity in probabilistic models of binary choice. *Journal of Behavioral Decision Making*, 31, 100–114.
- Rights, J. D., & Sterba, S. K. (2019). Quantifying explained variance in multilevel models: An integrative framework for defining r-squared measures. *Psychological Methods*, 24(3), 309–338. Retrieved 2021-06, from <https://doi.org/10.1037/met0000184>.
- Robert, C. P. (2016). The expected demise of the bayes factor. *Journal of Mathematical Psychology*, 72, 33–37. Retrieved 2017-10, from <https://doi.org/10.1016/j.jmp.2015.08.002>.
- Rouder, J. N., Engelhardt, C. R., McCabe, S., & Morey, R. D. (2016). Model comparison in ANOVA. *Psychonomic Bulletin & Review*, 23(6), 1779–1786. Retrieved 2017-04, from <https://doi.org/10.3758/s13423-016-1026-5>.
- Rouder, J. N., Morey, R. D., Speckman, P. L., & Province, J. M. (2012). Default Bayes factors for ANOVA designs. *Journal of Mathematical Psychology*, 56(5), 356–374. Retrieved 2013-03-23, from <https://doi.org/10.1016/j.jmp.2012.08.001>.
- Rozeboom, W. W. (1960). The fallacy of the null-hypothesis significance test. *Psychological Bulletin*, 57(5), 416–428. Retrieved 2021-04, from <https://doi.org/10.1037/h0042040>.
- Rozeboom, W. W. (2008). The problematic importance of hypotheses. *Journal of Clinical Psychology*, 64(9), 1109–1127. Retrieved 2020-12, from <https://doi.org/10.1002/jclp.20522>.
- Savani, K., & Job, V. (2017). Reverse ego-depletion: Acts of self-control can improve subsequent performance in indian cultural contexts. *Journal of Personality and Social Psychology*, 113(4), 589–607. Retrieved 2021-06, from <https://doi.org/10.1037/pspi0000099>.
- Schat, E., van de Schoot, R., Kouw, W. M., Veen, D., & Mendrik, A. M. (2020). The data representativeness criterion: Predicting the performance of supervised classification based on data set similarity. *Plos one*, 15(8), e0237009.
- Schielzeth, H., & Forstmeier, W. (2009). Conclusions beyond support: overconfident estimates in mixed models. *Behavioral Ecology*, 20(2), 416–420. <https://doi.org/10.1093/beheco/arn145>.
- Shiffrin, R. M., Chandramouli, S. H., & Grünwald, P.D. (2016). Bayes factors, relations to minimum description length, and overlapping model classes. *Journal of Mathematical Psychology*, 72, 56–77. Retrieved 2020-10, from <https://doi.org/10.1016/j.jmp.2015.11.002>.
- Shiffrin, R. M., & Steyvers, M. (1997). A model for recognition memory: Rem—retrieving effectively from memory. *Psychonomic Bulletin & Review*, 4(2), 145–166.
- Singmann, H. (2021). stanova: Bayesian models with categorical variables. *Zenodo*. Retrieved from <https://doi.org/10.5281/ZENODO.4904823>.
- Singmann, H., Bolker, B., Westfall, J., Aust, F., & Ben-Shachar, M. S. (2021). afex: Analysis of factorial experiments. (R package. <http://cran.rproject.org/package=afex>).
- Singmann, H., & Gronau, Q. F. (2021). bfrms: Bayes factors for bayesian regression models using 'stan'. *Zenodo*. Retrieved 2021-06, from <https://doi.org/10.5281/ZENODO.4904827>.
- Singmann, H., & Kellen, D. (2019). An introduction to mixed models for experimental psychology. In D. H. Spieler, & E. Schumacher (Eds.) *New methods in cognitive psychology* (pp. 4–31). Psychology Press.
- Singmann, H., Klauer, K. C., & Kellen, D. (2014). Intuitive logic revisited: New data and a bayesian mixed model meta-analysis. *PLoS ONE*, 9(4), e94223. Retrieved 2014-05, from <https://doi.org/10.1371/journal.pone.0094223>.
- Stephens, R. G., Dunn, J. C., & Hayes, B. K. (2018). Are there two processes in reasoning? the dimensionality of inductive and deductive inferences. *Psychological Review*, 125(2), 218–244. Retrieved 2019-10, from <https://doi.org/10.1037/rev0000088>.
- Ulrich, R. (2009). Uncovering unobservable cognitive mechanisms: The contribution of mathematical models. In F. Rösler, C. Ranganath, B. Röder, & R. H. Kluwe (Eds.) *Neuroimaging of human memory: Linking cognitive processes to neural systems* (pp. 25–41). Oxford University Press.
- van Doorn, J., Aust, F., Haaf, J. M., Stefan, A., & Wagenmakers, E.-J. (this issue). Bayes factors for mixed models.

- Computational Brain and Behavior. Retrieved from <https://doi.org/10.31234/osf.io/y65h8>.
- Van Fraassen, B. C. (2008). *Scientific representation: Paradoxes of perspective*. Oxford University Press.
- Venables, W. (1998). *Exegeses on linear models*. S-plus user's conference, Washington.
- Verissimo, J. (this issue). When fixed and random effects mismatch: Another case of inflation of evidence in non-maximal models. *Computational Brain and Behavior*.
- Vohs, K., Schmeichel, B., Lohmann, S., Gronau, Q. F., Finley, A. J., Others, M., & Albarracín, D. (in press). A multi-site preregistered paradigmatic test of the ego depletion effect. *Psychological Science*. (type: article) <https://doi.org/10.31234/osf.io/e497p>.
- Wagenmakers, E.-J., Krypotos, A.-M., Criss, A. H., & Iverson, G. (2012). On the interpretation of removable interactions: A survey of the field 33 years after Loftus. *Memory & Cognition*, 40(2), 145–160. Retrieved 2016-12, from <https://doi.org/10.3758/s13421-011-0158-0>.
- Wagenmakers, E.-J., Wetzels, R., Borsboom, D., & van der Maas, H. L. J. (2011). Why psychologists must change the way they analyze their data: The case of psi: Comment on Bem (2011). *Journal of Personality and Social Psychology*, 100(3), 426–432. Retrieved 2013-06, from <https://doi.org/10.1037/a0022790>.
- Westfall, J., Kenny, D. A., & Judd, C. M. (2014). Statistical power and optimal design in experiments in which samples of participants respond to samples of stimuli. *Journal of Experimental Psychology: General*, 143(5), 2020–2045. <https://doi.org/10.1037/xge0000014>.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

## Affiliations

Henrik Singmann<sup>1</sup> · David Kellen<sup>2</sup> · Gregory E. Cox<sup>3</sup> · Suyog H. Chandramouli<sup>4</sup> · Clinton P. Davis-Stober<sup>5</sup> · John C. Dunn<sup>6</sup> · Quentin F. Gronau<sup>7</sup> · Michael L. Kalish<sup>2</sup> · Sara D. McMullin<sup>5</sup> · Danielle J. Navarro<sup>8</sup> · Richard M. Shiffrin<sup>9</sup>

David Kellen  
davekellen@gmail.com

Gregory E. Cox  
gregcox7@gmail.com

- <sup>1</sup> University College London, London, UK
- <sup>2</sup> Syracuse University, Syracuse, NY, USA
- <sup>3</sup> University at Albany, State University of New York, Albany, NY, USA
- <sup>4</sup> University of Helsinki, Helsinki, Finland
- <sup>5</sup> University of Missouri, Columbia, MO, USA
- <sup>6</sup> University of Western Australia, Perth, Australia
- <sup>7</sup> University of Amsterdam, Amsterdam, Netherlands
- <sup>8</sup> University of New South Wales, Sydney, Australia
- <sup>9</sup> Indiana University, Bloomington, IND, USA