**Maybe maximal:
Good enough mixed models optimize power while controlling Type I error**

Running Head: Good Enough Mixed Models

Michael Seedorff
Dept. of Biostatistics, University of Iowa
Alphabet, Inc.

Jacob Oleson
Dept. of Biostatistics
University of Iowa

Bob McMurray
Dept. of Psychological and Brain Sciences
Dept. of Communication Sciences and Disorders
Dept. of Otolaryngology and
DeLTA Center
University of Iowa

Corresponding Author

Bob McMurray
278 PBSB
Dept. of Psychological and Brain Sciences
University of Iowa
Iowa City, IA 52242
319-335-2408
Bob-mcmurray@uiowa.edu

Abstract

Mixed effects models have become a critical tool in all areas of psychology and allied fields. This is due to their ability to account for multiple random factors, and their ability to handle proportional data in repeated measures designs. While substantial research has addressed how to structure fixed effects in such models, there is less understanding of appropriate random effects structures. Recent work with linear models suggests the choice of random effects structures affects Type I error in such models (Barr, Levy, Scheepers, & Tily, 2013; Matuschek, Kliegl, Vasishth, Baayen, & Bates, 2017). However, this question has not been addressed for between-subject effects, which are crucial for many areas of psychology, nor has this been examined in logistic models. Moreover, mixed models expose a number of researcher degrees of freedom: the decision to aggregate data or not, the manner in which degrees of freedom are computed, and what to do when models do not converge. However, the implications of these choices for power and Type I error are not well known. To address these issues, we conducted a series of Monte Carlo simulations which examined linear and logistic models in a mixed design with crossed random effects. These suggest that a consideration of the entire space of possible models using simple information criteria such as AIC leads to optimal power while holding Type I error constant. They also suggest data aggregation and the d.f, computation have minimal effects on Type I Error and Power, and they suggest appropriate approaches for dealing with non-convergence.

Keywords: Mixed Effects Models, Multilevel Models, Random Effects, Repeated Measures

**Maybe maximal:
Good enough mixed models optimize power while controlling Type I error**

1. Introduction

1.1 Mixed Effects Models

Mixed effects models have gained prominence in psychology and related fields (Brauer & Curtin, 2018; Meteyard & Davies, 2020, for reviews and tutorials). Psycholinguistics was an early and widespread adopter of mixed models (Baayen, Davidson, & Bates, 2008; Jaeger, 2008; Quené & van den Bergh, 2004); however, they are now important in a range of fields including clinical psychology (Field & Wright, 2011; Tasca, 2009), cognitive neuroscience (Koerner & Zhang, 2017), social psychology (Judd, Westfall, & Kenny, 2012) and in allied fields like speech pathology (Harel & McAllister, 2019; Oleson, Brown, & McCreery, 2019) and second language acquisition (Cunnings, 2012). Mixed models have also taken hold in psychological assessment, as seen in the use random effects models in item response theory (De Boeck, 2008).

Mixed models have become dominant for a variety of reasons, but a critical motivation was that they solved two important limitations of ANOVA. First, most versions of ANOVA only model a single random effect (Baayen et al., 2008; Clark, 1973; Judd et al., 2012). Second, ANOVA could not be used to analyze binary responses and traditional logistic regression is not suitable for repeated measures (Jaeger, 2008; Quené & van den Bergh, 2008). However, while mixed models help solve these problems, they raise a range of unsolved issues regarding how best to construct and use them.

1.1.1 Flexible Random Effects.

Many experiments have two (or more) sources of random sampling. For example, a psycholinguistic experiment may sample subjects from the population and words from the

language (and wish to generalize over both). Or a social psychology experiment might sample subjects from the population, and facial stimuli from a population of people. Classic approaches solved this issue of crossed random effects by combining two ANOVAS: one for subjects (collapsing over items), and a second for items (collapsing over subjects) (Clark, 1973; Raaijmakers, Schrijnemakers, & Gremmen, 1999). However, Baayen et al. (2008) proposed that mixed models could include multiple random effects in the same model. This made statistical interpretation more direct. While initially proposed for psycholinguistics where crossed random effects are common, this has proved useful in social psychology (Judd et al., 2012) and assessment (De Boeck, 2008). Moreover, crossed random effects are but one example of the flexibility of mixed models which can also handle nested random effects (e.g., capturing variance due to subject nested within classroom as in educational psychology) and other structures (Brauer & Curtin, 2018).

This flexibility in modeling random effects is seen even when there is only one source of random sampling (Snijders & Bosker, 1999). Mixed models can capture the fact that subjects' mean performance differs systematically using random intercepts, but can also capture the possibility that subjects systematically differ in their response to the experimental factors (a random slope). When combined with the ability to model crossed and nested effects, this allows one to capture the complex structure of error variance. This has the potential to improve power and reduce type I error – but only if the structure of the error variance is adequately modeled.

However, rarely in practice does the researcher know the true structure of the variance underlying the data. Instead, this must be assumed from prior work (or theoretical assumptions) or estimated from the data. Importantly, power and Type 1 error are dependent on the structure of the random effects (Barr et al., 2013; Matuschek et al., 2017). However, there is ongoing controversy about the best approach when faced with uncertainty about the true random slope

structure in the data. Until this issue is better understood, researchers may be faced with underpowered or under-conservative models, perhaps without knowing it.

1.1.2 Binary Data.

A fundamental assumption of ANOVA is that the responses are normally distributed (Jaeger, 2008). However, proportional data are not normally distributed, and thus, the interpretation of effects may not be reliable, particularly for proportions near floor or ceiling. Given the importance of measures like accuracy or two-choice decisions in psychology, this is a fundamental limitation. While transformations such as rationalized arc sine (Studebaker, 1985) or the empirical logit (Jaeger, 2008) can help, they are not ideal as they only approximate a Gaussian distribution when pooling over many trials, and behave poorly near 0 or 1.

Logistic regression correctly assumes a Bernoulli distribution, but it does not handle repeated measures. Logistic mixed models can do so, while simultaneously incorporating multiple random effects. However, there is even less guidance around fitting logistic mixed models when the researcher is uncertain as to the correct random effects. Thus, an important goal of this study is to extend the Monte Carlo evaluation of random effects selection procedures that has been conducted with linear models (Barr et al., 2013; Matuschek et al., 2017) to logistic mixed models. Logistic mixed models also raise practical questions. For example, many logistic mixed models fail to converge and there is a need to identify best practices when this occurs.

1.2 Constructing Mixed Models in Practice

Part of the flexibility of mixed models comes from the fact that they are not just modeling the design of the experiment (as in ANOVA), they are modeling the underlying factors and distributions most likely to have generated the data. With this flexibility comes the difficult task

of setting up a model to accurately reflect the generating source of the data. There has been considerable interest in optimal selection of fixed effects structure (Brauer & Curtin, 2018). However, less research has focused on selecting the random effects structure.

This manuscript conducts a more comprehensive simulation study to evaluate the consequences of several decisions for Type I error and power. We examine 1) procedures for selecting the random effects structure in the presence of uncertainty about the true generating functions; 2) the issues of aggregating data within conditions; 3) the issue of the appropriate d.f. for tests of fixed effects; 4) whether logistic mixed models perform similarly to linear models; and 5) issues of model convergence. While the first issue has engendered considerable debate (Barr et al., 2013; Matuschek et al., 2017), the others have received less attention.

1.2.1 Random Effects

In typical experiments, the fixed effects are typically well-defined (e.g., the factors that are manipulated, covariates that are measured). When the fixed effects are unknown, it is common (even recommended) practice to use the maximum reasonable fixed effects structure while selecting an optimal random effects structure (Fitzmaurice, Laird, & Ware, 2012) as this avoids pitfalls with stepwise regression (Steyerberg, Eijkemans, & Habbema, 1999; Whittingham, Stephens, Bradbury, & Freckleton, 2006). There has been less attention focused on selecting random effects. This is a primary concern as the structure of the random effects impacts Type 1 error and power for evaluating fixed effects (Barr et al., 2013; Matuschek et al., 2017). Consequently, choosing an incorrect random effects structure affects the conclusions drawn on the fixed effects.

Suppose a typical dataset in cognitive psychology consisting of response times collected in an auditory lexical decision task in varying background noise. Here, every subject has their

own true intercept capturing the fact that some respond faster or slower. Now suppose that subjects also have their own response to the background noise: with some individuals highly affected by the noise (their RTs fall off faster) while others have little or no effect.

One might attempt to fit this data with a simpler model, in which RT is a function of a fixed effect of noise (the overall trend) and a random intercept of subject (capturing subject mean differences), but with no random slopes. Because this model does not account for subject-specific response to the noise, all observations are treated as independent observations (after accounting for subject-specific intercepts) and a test of the fixed effect of noise is akin to a test performed with traditional regression. This could lead to increased Type I Error (TIE).

As this example suggests, if the random effects structure is too simple, TIE can be inflated. Due to this, Barr et al. (2013) argue that a maximal random effects structure should be used with all experimentally manipulated fixed effects included as random effects. This was supported by Monte Carlo simulations documenting increased TIE in models without random slopes. One consequence of the maximal approach is when the fitted models are more complex than the true data-generating model, power is reduced (Matuschek et al., 2017). This might be the lesser of two evils: confirmatory hypothesis testing is typically more concerned with avoiding Type I Errors (detecting a false effect) than with Type II Errors (failure to detect a true effect). Even with the cost in power, the maximal approach would seem to be a reasonable strategy. However, it is worth reconsidering for two reasons.

First, Barr's argument is in part conceptual: subjects will always vary in their response to a within-subject manipulation. Consequently, any model that does not include such a random slope is not modeling the error variance accurately. In some cases, this assumption may not be warranted. Consider cognitive control. Cognitive control has been measured for decades with tasks like the Stroop or Flanker tasks. In these tasks, the participant responds to a stimulus along

one dimension (e.g., the color of the font) while competing dimensions (the identity of the word) are congruent or incongruent. Such tasks consistently show robust effects of congruence with slower RTs for incongruent than congruent trials (Miyake & Friedman, 2012), a clear fixed effect of congruence. However, these tasks have been widely used as individual difference metrics with mixed results. In studies, often with hundreds of subjects, closely related cognitive control tasks may do not correlate with each other (Fan, McCandliss, Sommer, Raz, & Posner, 2002; Paap & Greenberg, 2013; Paap & Sawi, 2014): the congruence effect measured with Stroop does not correlate that of Flanker. And these tasks often do not always correlate *with themselves* (e.g., test/retest reliability: Enkavi et al., 2019). Here, the congruence effect for a subject is a random slope. The fact that these effects are not stable across tasks or testing sessions, suggests that the underlying shape of this data may not include random slopes (at least in many datasets). That is, subjects may vary in mean RT (the intercept); there may be a main effect of congruence; but variation in the size of this effect across subjects is not meaningful.

Second, maximal models are efficient and easy to implement for a small set of factors. However, this is rarely the case. Oftentimes, researchers will include two to three experimentally derived fixed effects, along with several other nuisance or exploratory variables (e.g., trial number to capture learning). Such models often do not converge accurately. In fact, we recently reviewed every article in three volumes (113-116) of the *Journal of Memory and Language*, an early promotor of mixed models. Of the 20 articles using mixed models only three used a truly maximal random slope structure, and the most common strategy (8 articles) was one in which authors attempted a maximal slope structure and backed off due to non-convergence (and see Meteyard & Davies, 2020, for similar findings in a meta-analysis). This raises the need to better understand the consequences of random intercepts to guide these decisions. It also suggests that when slopes are not required by the data (for one factor), excluding them may

allow their inclusion other factors for which random slopes are needed.

Thus, while a fully maximal approach may be statistically conservative, it may not be necessary in all cases or for all factors. In contrast, an approach that accurately selected the random effects structure based on the observed data could offer greater flexibility and power. Such an approach would require one to accurately identify the random effects structure observed in the data. A common approach is to compare an initial model against models with increasing or decreasing complexity using likelihood ratio tests (LRTs). This proceeds through a series of nested models (e.g., each increasingly complex model includes all of the elements from the preceding simpler model) in a stepwise fashion. Note that here, the stepwise procedure is iterated over random effects not fixed effects. At each step, the new model is evaluated against the old one and more complex random effects structures are only included if they offer a significantly better fit. Model selection stops when no further gains are possible. This procedure can lead to premature stopping, for example, when the third model is not better than the second (and selection stops), but the fourth or fifth would have been (and were missed). This can also be challenging with non-nested models (e.g., separate random effects structures for subjects and items), requiring complex chains of inference.

An alternative is to use an information criterion measure for model selection (rather than a likelihood ratio test). Information criterion measures take the likelihood for a given model and apply a penalty for the number of parameters in the model. This permits an alternative strategy with greater flexibility: it can directly compare all possible models (including non-nested models). This avoids the researcher degrees of freedom involved in selecting a path for forward or backward selection procedures. Thus, a model-space procedure based on information criterion may be preferable if it performs similarly (or superior) to likelihood ratio testing.

Matuschek et al. (2017) compared several strategies for model selection. They used

Monte Carlo simulations to demonstrate that one model selection approach—backward stepwise selection using LRTs—maintains appropriate TIE for several true model conditions while consistently showing similar power to the maximal model. They also suggest that the Akaike Information Criterion (AIC) (Akaike, 1974) can inflate TIE rates above the intended level. However, this critique of AIC may be premature. This study was primarily concerned with limited sample sizes (where AIC does not perform optimally), and it didn't consider alternatives that are more suited to small sample sizes (AICc, BIC: Hurvich & Tsai, 1989; Schwarz, 1978). Importantly, information criterion measures can run into issues with mixed models. At a broader level, both this study (Matuschek et al., 2017) and the aforementioned study by Barr et al. (2013) support the inclusion of random slopes. However, both studies have a number of limitations.

First, these studies examined within-subject designs. It is unclear if these results would hold for mixed designs with between-subject effects. There has been increasing application of mixed effects models to fields in which between-subject differences are of central importance, such as clinical psychology (Field & Wright, 2011; Tasca, 2009), social psychology (Judd et al., 2012), speech pathology (Harel & McAllister, 2019) and assessment (De Boeck, 2008). It is unclear how or if the specification of random effects may impact TIE and power in such designs. Previous studies (Raaijmakers et al., 1999) using ANOVA have shown that failing to include a within-subject random effect does not alter estimates for between-subject covariates. It is unclear whether this holds for similar setups in mixed models with multiple random effects. Second, it is unclear how random effects specification extends to logistic models. There are significant differences between linear mixed models and generalized linear mixed models that make it unclear as to whether the same approaches would be useful.

1.2.2 Mixed Effects Logistic Regression

There are few techniques that are widely available for selecting random effects structures for logistic models. While a number of statistical model selection approaches are promising (Cantoni, Flemming, & Ronchetti, 2005; Hui, Muller, & Welsh, 2017; Ibrahim, Zhu, Garcia, & Guo, 2011; Jiang, Rao, Gu, & Nguyen, 2008) they cannot yet be calculated using default output in the `glmer` function in R, the standard implementation of mixed effects models in psychology.

Furthermore, logistic models are substantially more difficult to estimate computationally. It is difficult to directly compute the objective function (function that computes the likelihood of the from a given model). Thus, approximations must be made with tradeoffs between computational intensity and accuracy. In general, more computationally intensive approaches provide lower bias in the estimators but require simpler random effects structure and can yield high rates of non-convergence. While the approximation technique used should reflect the goals of the analysis (i.e., whether multiple models will be compared and the complexity of the random effects structure), the ability for one method to converge on a solution is a factor.

There are a number of reasons a mixed model may fail to converge. Sometimes non-convergence is due to an insufficient number of fitting iterations, a convergence criterion that is too stringent, poor starting parameters, or other issues unrelated to model specification. In these scenarios, solutions are typically straightforward, and one can change parameter within the fitting procedure to solve them. Sometimes, but not always, non-convergence can be solved moving to a different optimizer, or to a coarser approximation of the log-likelihood of the data given the model. In other cases, these “tricks” do not solve the problem.

It is difficult to know whether a non-converging model should be used for inference, particularly for the (non-statistician) end user. Such models can return parameter estimates and covariance structures that are not of any particular interest as they do not maximize the likelihood function but are coincidentally at the point where the algorithm failed. In other cases,

non-converging models can yield reasonable estimates and sound inference. However, most end-users lack the expertise to determine the validity of the estimates and consequently, solve this by moving to a simpler model.

Several options exist for this (see Brauer & Curtin, 2018, for an excellent discussion) (and Eager & Roy, 2017, for a Bayesian approach that reduces non-convergence). Simplifying the assumed covariance of the random-effects can often achieve convergence. This approach maintains the random slopes but drops one or more terms that capture the covariance between slopes or between slopes and intercepts. If this is insufficient, one can continue with further simplification, such as dropping a slope term entirely. Many end users are unfamiliar with the intermediate step of dropping covariance terms, and instead adopt the more problematic approach of dropping random slopes. It is unclear the degree to which these strategies affect TIE in a particular dataset that leads to non-convergence.

These approaches all deal with a scenario where the researcher knows the structure of their fixed and random effects, finds the model to be non-convergent, and wants the most similar converging model. Equally common is a situation in which the researcher is seeking an optimal structure using a model selection procedure, but some candidate models are non-convergent. Here, the straightforward approach might just be to eliminate non-converging models from consideration. This could perform fine in terms of TIE rate, if the non-convergent models tend to be those ill-fitted to the data. On the other hand, if non-convergence is primarily a factor of something other than the model's fit to the underlying structure of the data, then this approach will have increased rates of incorrect model selection and either increased TIE rates or low power. This is likely given that non-convergence is not random. More complex (conservative) random effects are the most likely to fail to converge. It is currently unclear the degree to which these choices impact TIE or power, and consequently what the "right" choice is.

1.2.3 Data Aggregation

For studies with multiple observations within groups, one can use the mean of those observations (as is in ANOVA) instead of using them as separate data points¹. This reduces the amount of data, and the computation time for should be lower on average. It's not clear how this aggregation will affect other details of mixed models. Theoretically, one would expect to get the same fixed effect estimates from both aggregated and non-aggregated models. However, degrees of freedom estimates may differ, potentially changing TIE, especially when the estimated d.f. are near the total number of observations. Reducing the number of repeated events within cells will also affect the estimate of the error variation. This also introduces a new option for dealing with non-convergence and it isn't immediately clear whether a transition from a raw to aggregated data would be preferable to reducing the complexity of the random effects. Thus, our simulations evaluate models against both aggregated and non-aggregated data.

1.2.4 Degrees of Freedom

Since the introduction of mixed models, there has been ongoing discussion around the appropriate degrees of freedom for fixed effects. The simplest approaches involve using the minimum potential degrees of freedom, using an asymptotic normal distribution to replace the t (i.e., assuming an infinite sample size), or using the total number of observations minus the number of parameters (e.g., a traditional calculation). The latter two options will be anti-conservative in nature, as they overestimate the true degrees of freedom.

Another popular option is to use a Satterthwaite approximation for the degrees of

¹ This can only be done in linear models – logistic models require individual trials to estimate the binomial distribution.

freedom (Satterthwaite, 1946), implemented in the *lmerTest* R package (Kuznetsova, Brockhoff, & Christensen, 2016). This approach attempts to calculate an effective d.f. that accounts for the correlation of the responses. At times, however, these estimates can get nonsensically large. Realistically, the d.f. itself may not be a huge problem: there is little functional difference between two t-distributions with very large degrees of freedom. Nonetheless, for smaller d.f., it may be important to get it “right”. A more recent alternative is the Kenward-Roger approximation (Kenward & Roger, 1997). Satterthwaite and Kenward-Roger both perform well in small sample sizes (Luke, 2017), although Kenward-Roger can be more computationally intensive, especially for large datasets (Halekoh & Højsgaard, 2014). Traditionally, d.f. were reported for ANOVA as a sanity check, to verify the analysis was set up correctly. No such interpretation can be had for d.f. in mixed models.

Crucially, the fact that the d.f. can become quite large raises an intriguing problem. This is particularly common in models that fail to use a random slope for a within-subject factor. Such models are known to be anti-conservative (Barr et al., 2013). It is unclear whether this is a result of an overestimation of degrees of freedom or of underestimating the standard errors. In these cases does it matter if the d.f. are large (if the model is selected appropriately)?

1.3 Overview

This study conducted a series of Monte Carlo simulations to address these fundamental issues. We investigated 1) the optimal selection of the random effects structure in both linear and logistic models for models with a between-subject factor; 2) the influence of d.f.; 3) the influence of data aggregation; and 4) the consequences of typical model simplification steps when experiencing non-convergence in logistic models.

In setting up the Monte Carlos our goal was to simulate the issue faced in actual analysis,

where the true generating model is unknown. For example, if the true generating model is the intercepts-only model, what procedures maintain TIE and do these affect power? Conversely, if the true generating model contains slopes (or slopes that do not covary), how do the different procedures perform? In these cases, the generating model was known (by us, as the builder of the simulations). However, once we have analyzed these independently, we can arrive at an accumulated decision about the performance of the procedures when the true random effects are uncertain. For instance, if a model selection criterion performs well with respect to one or two generating models but shows large inflations in TIE for another, then this approach should be ruled out as we do not know, in real analyses the true generating model.

We note that in some cases, these simulations report findings that are generally known (or are eminently reasonable) to formal statisticians (c.f., simulations on the d.f. problem). However, our experience as a writer and reviewer is that these facts are not yet widely appreciated by end users of mixed effects models in psychology. In these communities, these misunderstandings are common, leading to problems during manuscript review, and unnecessary angst amongst users. Our work on random effects selection is novel in both communities, however, our Monte Carlos are also intended to provide clear evidence and guidance on some of these less understood issues for end users in psychology.

This study is a first step. We perform extensive evaluation of these models under a small set of assumed underlying parameters (e.g., for the size of the fixed effects, variance of the random effects and so forth). With our broad goals it was infeasible to examine the full space of possible datasets. Thus, this is deep dive into a small portion of the space. We have provided all of our source code on the Open Science Framework (<https://osf.io/vj6db/>), and we hope readers will conduct their own simulations on parameters relevant to their own paradigms.

2. Methods

2.1 Linear Model Specification

Our primary models use a mixed design with one between- and one within-subjects factor. Our examination of repeated measures replicates prior work (Barr et al., 2013; Matuschek et al., 2017), and extends it to logistic regression, to issues of non-convergence and data aggregation. The between-subject factor captures issues relevant beyond cognitive psychology that have not been examined by prior studies.

We assume something resembling an auditory lexical decision experiment with one within-subject factor (e.g., background noise) and one between-subjects factor (e.g., age-group). Each has two levels (e.g., high/low noise, old/young). We assume some fixed number of items (randomly drawn from the pool of possible words). We model fully within-subject and within-item design – each item appears in both noise conditions for five repetitions (trials).

Possible random factors include participant and item (the word on that trial). All models have the same fixed effects: the between-subjects factor, the within-subjects factor, and their interaction. Several potential random effects structures are tested here. First, we consider the *intercepts-only* model (random intercepts of subjects and items). Moving to more complex structures we started by focusing on subject effects, with only item-intercepts. Here the goal is to consider the implications of model selection procedures on the between-subject effect (for linear models) and for all effects in logistic models. Thus, we next consider a *subject-slopes* model, which includes random slopes for subjects and intercepts for words (we consider both subject and item slopes in Supplementary Analysis S3). In addition, in the sloped model, the covariance terms between subjects slopes and intercepts must also be estimated. Thus, the *independent-slopes* model included random slopes and intercepts for subjects but assumed no covariance.

Model fitting was performed using the lme4 package (Bates & Sarkar, 2011) in R (R

Core Team, 2015). For the purposes of evaluating fixed effects, we used the Satterthwaite approximation to the degrees of freedom, calculated using the lmerTest package (Kuznetsova et al., 2016). The intercept-only model is described in lme4 notation in Equation (2) (see Supplement S1 for statistical equations).

$$y \sim \text{Age} * \text{Noise} + (1 \mid \text{Subject}) + (1 \mid \text{Item}) \quad (1)$$

The subject-slope model is shown in Equation (2).

$$y \sim \text{Age} * \text{Noise} + (\text{Noise} \mid \text{Subject}) + (1 \mid \text{Item}) \quad (2)$$

Finally, the independent slopes model includes random slopes on subject for noise but assumes no covariance with the random intercepts (Equation 3).

$$y \sim \text{Age} * \text{Noise} + (\text{Noise} \parallel \text{Subject}) \quad (3)$$

Our simulations assumed one of these three models as the data generating model, and then explored appropriate model selection procedures. For each data generating model, we ran two simulation studies: one assuming fixed effects are null (to assess Type I error) and one when fixed effects are non-null (to assess power). For each generating model, all three models were fit to the data. These are referred to as the *fitted models*. We then evaluated Type I error and power for several analytic choices. First, we evaluated consequences of simply assuming one of these three fitted models a priori, to understand their properties given that data. Next, we evaluated several model selection procedures for selecting the random effects structure for fitted models: These include three incremental procedures using likelihood tests: (1) forward-selection testing only the intercepts-only and the subject-slope models (LRT2, the typical practice); (2) forward-selection tests which also include the covariance terms as an option (LRT3); (3) backward-selection considering all fitted models (bLRT3). We also considered three procedures which evaluate all possible fitted models simultaneously (*model-space* procedures) using (4) AIC, (5) AICc, and (6) BIC. For AICc and BIC, the total number of observations was assumed to be the

number of subjects, N , times the number of observations per subject or n^* . For model-space procedures, the model with the lowest absolute information criteria was used.

In generating the data, we used the following values. The null effects model had a value of 1000 for β_0 (the intercept) and fixed effects of 0 for all others. The non-null effects model had values of 1000 for β_0 , 50 for β_1 (the between-subjects effect), 10 for β_2 (the within-subjects effect), and -15 for β_3 (the interaction effect). For the random effects, the variance of the random intercepts for both subjects and items was 10^4 (i.e. a standard deviation of 100 for subject-specific and item-specific intercepts), a variance of 20^2 for the random slopes (i.e. a standard deviation of 20 for the subject-specific slopes) in models where the random slopes were included, 8000 for the covariance between the subject's random intercepts and slopes (a relatively large correlation, to maximize the differences between a fully sloped model and a model that drops the covariance term) in the subject-slopes model, and 200^2 for the error variance. The simulation considered the scenario with 50 subjects (25 in each age-group), 20 items, 2 within-subject levels, and 5 repetitions per subject/item/within-subject group combination, for a total of 10,000 observations per iteration of the Monte Carlo.

This set of parameters yielded deliberately low power, with power in the range of 0.35 ~ .45 to detect a true effect. This was done to ensure sufficient room to see increases in power – if power was at the typical 0.80 there might have been ceiling effects. Sadly, too it may reflect the actual power in many experiments (Button et al., 2013; Maxwell, Lau, & Howard, 2015).

We ran 10,000 iterations for each study.

2.2 Logistic Model Specification and Simulations

Logistic models had identical fixed effects to the linear model, and we examined the same random effects structures. All three models had identical lmer notation to the linear model

with the exception that the link function is set to binomial and see Supplement S1 for formal notation. Simulations of the logistic models proceeded similarly to linear models. Both TIE and Power simulations were run for each of the three random effects structures, for a total of six total sets of simulations. Fitting was performed in R using Laplace approximation, the default choice, via the `glmer` function from the `lme4` package. The same model selection procedures as the linear model simulations were used.

The non-null effects model assumed values of 0.2 for β_1 (the between-subjects effect), 0.3 for β_2 (the within-subjects effect), and -0.3 for β_3 (the interaction effect). Values used for the random effects included 0.4^2 for the variance of the random intercepts for subjects and items, a variance of 0.2^2 for the variance of the random slope for the models that included it, and a correlation of 0.8 between the subject-specific random intercepts and slopes for the maximal model. The simulation used the same experimental design as the linear model: 50 subjects, 20 items, 5 repetitions/cell for a total of 10,000 observations.

We ran 10,000 iterations for each study. For logistic models, non-convergence was non-negligible. These models were tracked, and excluded from the primary analyses of random effects. We examine them separately to determine appropriate responses to non-convergence.

2. Estimating TIE and Power

For each simulation and model selection procedure, we calculated the overall estimated probability of TIE (for the null effects generating models) and the estimated power (for the non-null generating models). For the base model fitting with no model selection, TIE and power were calculated as the proportion of runs (of those that converged) that found a statistically significant effect for a given factor at $p < .05$ level. For model selection approaches, TIE and power were calculated as the proportion of runs where the optimal model (as chosen by the model selection

procedure) had a statistically significant effect of a given factor. For each estimated proportion (TIE and power), we computed 95% confidence intervals using the exact Clopper-Pearson interval approach. Models (or selection procedures) were considered anti-conservative if the complete coverage of the confidence interval (in the TIE models) was larger than 0.05.

There are varying degrees of harm from an incorrect model selection. Consider a case where the generating model is subject-sloped and the intercept-only fitted model had inflated TIE, while the other two held appropriate TIE with similar power. Here, an incorrect decision to select the intercept-only fitted model should be penalized more than an incorrect decision to select the independent slopes fitted model—neither is the true generating model, but the latter does not lead to inflated TIE. Thus, our goal was not to evaluate the likelihood of selecting the true model (indeed on any given run, the data may not reflect that the true model), but rather to evaluate the TIE and power of a given selection procedure. Thus, this captures to some degree not just the likelihood of the procedure selecting the true model, but also the likelihood that a given draw of the data accurately reflects the true generating function.

3 Results

We start by examining the linear model. First, we investigate random slopes on subjects, emphasizing the issue of the between-subject factors. This is followed by a discussion of the degrees of freedom and data aggregation. After that we return to random effects selection in logistic models. This is followed by an analysis of the issue of non-convergence.

3.1 Linear Model Simulation Results

3.1.1 *Random Effects Selection Procedures*

Between-Subjects Factor. The between-subjects effect does not have a corresponding

random slope, thus with regard to this effect, the random effects were correct for all fitted models. However, it is unknown whether differences in a within-subject random effects impact TIE or power for the between-subject effect.

Figure 1 shows results. Here, the left panels show TIE and the right shows power; each row represents a different generating model (intercept-only, independent slopes, or subject-slopes). Estimated TIE or power is designated with a solid black dot, confidence intervals are represented by the dashed lines, and (in the TIE panels) a light gray line is placed at 0.05 as a reference for which model selection approaches are conservative or anti-conservative.

We start by considering only the pre-specified models (the top half of each figure). Figure 1A and B shows TIE (1A) and Power (1B) for the intercepts-only generating model. All models show similar TIE and power. The independent-slopes generating model (Figures 1C, D), shows a similar pattern with no significant differences between the model selection techniques. A different story emerged for the subject-sloped generating model (Figures 1E,F). The correct fitting model (subject-slopes) has an estimated TIE very close to 0.05. However, the other two, simpler fitted models (intercepts-only and independent-slopes) are very conservative (with TIEs near 0.035, and significantly reduced power, Figure 1F). In this scenario, the simpler models do not account for within-subject variance, increasing standard errors for the between-subjects effect (average SEs for fitting with intercepts-only: 30.98, independent-slopes: 30.02, and fully-sloped: 28.72). Thus, even for between-subjects variables, accurately modeling the random effects structure is crucial.

We next turn to the model selection procedures, shown in the bottom half of each panel in Figure 1. Model-space AIC, AICc, LRT2, and bLRT3 all maintain appropriate TIE and similar power to the subject-slope fitted model. LRT3 is a bit more conservative, while model-space BIC is even more conservative, as these selection approaches are more likely to pick simpler models.

Figure 1: (A) TIE and (B) power results for the **between-subjects** effect in the random intercept only generating model. Fitting approaches are listed at the top including the intercepts-only model (Int), the independent-slopes model (Ind-slope), and the subject-sloped model (Sub-slope). We also considered the 6 model selection procedures including forward selection with LRTs (LRT2/3), backward selection with LRTs (bLRT3), and simultaneous consideration of all models using AIC, AICc, or BIC. (C) TIE and (D) power results for the between-subjects effect in the independent slopes generating model. (E) TIE and (F) power results for the between subjects-effect in the fully sloped generating model.

In sum, for a between-subject factor, TIE was generally maintained for all fitted random effects structures. However, for models generated with random slopes, power was improved when the fitted model contained slopes. Model selection procedures that were better at detecting this more complex random effect structure thus led to better power: model space AIC/AICc

procedures and some LRT approaches.

Within Subjects Factor. We next turn to the within-subject factor. Figure 2A/B display the results for the intercepts-only generating model. As expected, more complex models (independent slopes and subject-sloped) show reduced power while all model selection procedures perform comparably to the true fitted model. This is driven by correct selection of the intercepts-only fitted model with rates of 97.8% for LRT3, 94.3% for bLRT3, 97.6% for LRT2, 88.9% for AIC, 88.9% for AICc, and 99.9% for BIC.

For the independent-slopes generating model (Figures 2C/D), the intercepts-only fitted model is extremely anti-conservative with TIE over 0.1. In contrast, both fitting models with random slopes perform appropriately. This drives anti-conservative behavior in the model-space BIC, LRT2, bLRT3, and LRT3 approaches as they are more likely to choose a simpler model, while model-space AIC and AICc maintain an appropriate TIE. Rates of selection of the intercepts only fitted model were 45.3% for LRT3, 43.0% for bLRT3, 54.1% for LRT2, 25.2% for AIC, 25.2% for AICc, and 79.1 % for BIC.

Finally, the subject-sloped generating model (Figure 2E, F) shows similar results: the intercepts-only fitting model is anticonservative, and this leads to anti-conservativeness in model selection approaches that over-select it. The exception to this is LRT2 and bLRT3, which provide similar results to AIC/AICc (both perform similarly to the fully sloped fitted model).

In sum, model space AIC/AICc approaches maintain TIE for the within-subject factor, for all generating model scenarios. This is in comparison to LRT2, bLRT3, and LRT3, which are more commonly used selection approaches. These show an increase in TIE based on increased rates of choosing a model that omits random slopes entirely. Moreover, when the generating model does contain slopes, AIC/AICc shows significantly higher power in detecting an effect than solely using a maximal random effects structure for fitting.

Figure 2: (A) TIE and (B) power results for the **within-subjects effect** in the random intercept only generating model. (C) TIE and (D) power results for the within-subjects effect in the independent slopes generating model. (E) TIE and (F) power results for the within-subjects effect in the fully sloped generating model.

The Interaction. Results for the interaction mirrored those of the within-subject effect.

These are reported in full in Supplement S2).

Summary. Model space AIC/AICc approaches maintain TIE for the within subjects factor, for all generating model scenarios. This is in comparison to LRT2, bLRT3, and LRT3, which are more commonly used and show increased TIE rate, since they are more likely to

choose a model that omits random slopes entirely. Additionally, when the true generating model contains slopes, AIC/AICc shows significantly better power in detecting an effect than when solely using maximal random effects. Supplement S3 extends this to crossed random effects and shows identical results.

3.1.2 Degrees of Freedom

We next evaluated the d.f. output by the Satterthwaite procedure. D.f. was stable for the between-subject factor, with estimates from 48 to 50, roughly the number of subjects. However, for the within-subject factor, d.f. varied greatly ranging from about 50 to upwards of 9000, and it was occasionally nonsensically above 10,000 (the number of observations). Detailed analyses (Supplement S4) suggested this could appear even when the data was generated with random slopes, and fit with the appropriate model, or when data came from an intercepts-only model but was fit with random slopes. Thus, it is unwise to over rely on d.f. as a sanity check for an appropriately conservative model. Moreover, Supplement S4 shows that the increase in TIE error observed for intercept-only models does not derive from mis-specified d.f., but rather from mis-specified error terms. Thus, again, the d.f. should not be used to evaluate suitability.

3.1.3 Data Aggregation

We next asked whether the choice of averaging the data within cell-means, or using raw (e.g., trial-by-trial) data impacted TIE and power. We focus here on a situation in which each item was encountered on multiple repeated trials. Thus, we can aggregate within-item and within-subjects, preserving the random effects structure.

For every simulation iteration, we reran all analyses using aggregated data (the mean of the 5 responses within each subject $\times$ item $\times$ within-subject-factor cell). This reduced the per-

Figure 3: (A) Non-convergence rates for when fitting with the independent slopes model. (B) Non-convergence rates when fitting with the fully sloped model.

subject sample from 200 to 40 observation and the total number from 10,000 to 2,000. We repeated the same fitting approach for each potential fitted model.

Aggregation may be beneficial as an alternative approach for dealing with non-converging models. Non-convergence rates were extremely low in general (under 0.5%) for the linear model (aggregation is not possible for logistic models where this is more of an issue). When fitting using the intercepts-only model, non-convergence was never encountered in either condition. However, sloped fitted models had greater rates of non-convergence.

Figure 3 compares non-convergence rates between models fitted to raw and aggregated data. Non-convergence varied as a function of both the underlying data (x-axis) and the fitted model (Figure 3A, B). This suggests non-convergence cannot be treated as a random event. However, for both fitted models, non-convergence rates were consistently lower when fitting with aggregated data. These reductions were statistically significant (using McNemar's test) for the intercepts-only generating model/fully sloped fitted model ($p = 0.015$) and the fully sloped generating model/fully sloped fitted model ($p = 0.0009$).

We also examined the standard error of the estimates across runs of the model to

Figure 4: TIE rates for raw and aggregated data as a function of the fitted model (Y) axis, and the generating model: (A) Intercepts-Only; (B) Independent-Slopes; (C) Subject Slopes.

determine if aggregation results in artificially low SE estimates, or alters the stability of the parameter estimates across runs. This generally showed few differences between aggregated and raw data, with the exception that parameter estimates were slightly more variable for aggregated data with the intercept-only fitted model (but not more complex fitted models).

In terms of impact on the actual TIE rates, however, the decision to aggregate made no difference. Figure 4 shows TIE for each generating model (separate panels) and each fitted model. In all cases, the aggregated data (open circles) showed nearly identical TIE rates to the raw data (filled circles) and confidence intervals fully overlapped. Thus, the difference between performing model fitting using raw data or aggregated data has little effect on the conclusions drawn from the data. Because of this, it makes sense to use aggregated data to reduce the likelihood of non-convergence.

3.2 Logistic Model Simulations

We next turned to the logistic model. As with the linear model, our primary interest is the ability of the model to maintain an appropriate TIE rate while maximizing power. Analysis of the

logistic model selection results follows in a similar fashion to the linear model.

3.2.1 Model Selection

Between-Subjects Factor. We began by examining results for the between-subjects factor. For the intercepts-only generating model (Figure 5A, B), all fitting models and all model selection procedures are practically indistinguishable in terms of TIE and power. Figures 5C, D show results for the independent slopes generating model. All models maintain appropriate TIE. When it comes to power, while confidence intervals clearly overlap, there was a non-significant decrease in power when using an intercepts-only fitted model, and a similar decrease in power for model space selection using BIC. While, the estimated power is similar between subject-sloped and independent-slopes fitting models. Of the selection procedures, model-space AIC/AICc appeared to best maintain power. This is likely because AIC/AICc was more resistant to selecting the low-power intercept-only fitted model, choosing it in only 31.9% of the iterations. In comparison, it was chosen on 53.1% of runs by LRT3, 50.6% of runs by bLRT3, 61.2% of runs by using LRT2, and 83.9% of runs with BIC.

Finally, the subject-sloped generating model is shown in Figure 5E,F. Both incorrectly specified models show significantly decreased power (though they unexpectedly also show reduced TIE). This decreased for the LRT3 and model-space BIC procedures which favored these simpler models. In contrast, the full model showed the best power. This tended to be selected by model-space AIC/AICc procedures, and by the LRTs. The correct, subject-slopes fitted model was selected on 50.8% of runs by LRT3, 69.0% of runs by bLRT3, 81.6% of runs by LRT2, 80.6% of runs by AIC, 80.6% of runs by AICc, and 17.8% runs with BIC.

In sum, for the between-subject factor, logistic models performed similarly to the linear

Figure 5: (A) TIE and (B) power results of the logistic mixed model for the *between-subjects effect* in the intercept-only generating model. (C) TIE and (D) power results for the between-subjects effect in the independent slopes generating model. (E) TIE and (F) power results for the between subjects-effect corresponding to the subject-slope generating model.

models. TIE was held to constant (or conservative) levels across all models. However, when the data was generated with random slopes, power was improved when the fitted models captured this. Procedures like model space AIC/AICc and the LRTs appeared to do this appropriately.

Within-Subjects Factor. For the intercepts-only generating model (Figure 6A,B), all procedures maintained appropriate TIE. Both of the incorrectly described fitting models (fully

Figure 6: (A) TIE and (B) power results of the logistic mixed model for the *within-subjects effect* with the intercept-only generating model. (C) TIE and (D) power results for the within-subjects effect with the independent slopes generating model. (E) TIE and (F) power results for the within-subjects effect with the fully sloped generating model.

sloped and independently sloped) had decreased power, although not significantly (based on confidence intervals). Several selection procedures show power estimates close to the intercept only fitting model: LRT3, bLRT3, LRT2, and model space BIC.

Figures 6C, D show results for the generating model with independent slopes. Similar to the linear model, when the fitted model contained intercepts only, we observed large anti-

conservatism (though also high power). This increases the TIE rates of LRT3, bLRT3, LRT2, and model space BIC selection procedures which tend to favor the selection of intercepts-only models. Model space AIC/AICc had a slightly increased estimated TIE rate compared to the true generating model, although this was not significant.

Figures 6E,F show the results for the subject-sloped generating model. These follow the linear simulations closely: we find anti-conservatism in the intercepts-only fitting model, conservatism in the independent-slopes fitting model. We find TIE and power to be pretty consistent across model space AIC/AICc, LRT2, bLRT3, and the subject-sloped fitted model. Contrary to the linear model results, we do not see inflated TIE rates for LRT3 (at a statistically significant level). The model space BIC selection procedure is anti-conservative again, driven by its propensity to choose the simplest fitted model (random intercepts only).

Interaction. The interaction of the between- and within-subject factors showed a nearly identical pattern of results to the purely within-subjects factor (Supplement S2).

Summary. In sum, for within-subject effects, intercept-only models were anti-conservative when the underlying data included with random slopes (consistent with Barr et al., 2013). However, fitted models with subject-slopes and independent-slopes maintained TIE (for both generating models that included slopes). For generating models with slopes, power was roughly equal for most selection procedures (with the exception of the anti-conservative BIC).

3.2.2 Model Convergence

While non-convergence was rare in linear mixed models, non-convergence rates for logistic models were much higher (Figure 7). All intercepts-only fitted models converged, no matter the generating model. However, high rates of non-convergence were observed when the fitted model included random slopes, and the generating model did not (8.6% for the

independent-slopes fitted model; 10.1% for subject-sloped). When data came from the subject-sloped model, there was a low non-convergence rate of 0.3% when fitting with an independent slope model, but a high rate of 8.0% when fitting with the correct, subject-slopes model.

A common belief among end users is that non-convergence is a sign of over specification of the random effects. However, these simulations suggest this does not always hold. Referencing the subject-sloped fitted models (the third column), there was a smaller rate of non-convergence when it was under-specified (data came from an independent-slopes

Figure 7: Observed non-convergence rates from logistic mixed model simulations.

generating model) than when it was correctly specified (fully sloped generating model; $p < 0.0001$). Thus, non-convergence is not necessarily a sign of a mis-specified model. Rather it may occur for other reasons (see Section 1.2.2).

The critical question is what to do when non-convergence occurs. While the best action is to determine the reasons for non-convergence, the kinds of analytic approaches necessary for this may not be available to typical end-users who often default to simplifying the model. Do these simulations offer any guidance?

Importantly, independent-slopes fitted model almost always converged – even for the most complex data structure (the subject-sloped data). Our previous results suggest that the independent and fully sloped models both hold TIE to appropriate levels (under either generating model), and that they have reasonable power. This suggests that the independent slopes model may be a reasonable back-off option when a fully sloped model does not converge.

To validate this, we examined the most common issue that many researchers report: when a fully sloped model might appear needed but does not converge. Here, a subject-sloped fitted model failed to converge 794 times in the TIE simulations and 347 times in the power simulations (out of 10,000). While this is a relatively small number of runs, it was sufficient to compute the TIE and power when the (correct) subject-sloped model fails to converge and alternative models were attempted on the same dataset (as they were here).

Of the models that failed to converge, the bulk could be fit with independent slopes (TIE: 793 / 794; power: 345 / 347) and all could be fit with intercepts only. This suggests that backing off to simpler models can help achieve convergence. However, it is not clear whether backing off to either model might affect TIE and power.

We start with the between-subjects factor. In the original models (Figure 5E,F), TIE rate for the correct—fully sloped—model was 0.054 ($CI_{95\%} \pm 0.0046$). When these models failed to converge, the independent-slopes and intercepts-only models were both more conservative (independent: $p(TIE)=0.032 \pm 0.012$; intercepts-only: $p(TIE)=0.028 \pm 0.011$). These compare favorably with the estimates for all datasets (independent: $p(TIE)=0.037 \pm 0.0037$; intercepts only: $p(TIE)=0.029 \pm 0.0033$). A similar story was seen with power. The subject-slopes model showed power of 0.350 ($CI_{95\%}: \pm 0.0095$). When this failed to converge, the independent-slopes and intercepts-only models showed reduced power (independent: $\beta=0.320 \pm .049$; intercepts-only: $\beta=0.282 \pm 0.047$). Thus, as in the primary analyses all models held TIE below 0.05 for between-subject factors, while less complex random effects models for the within subjects effects tended to show lower power.

We next analyzed the within-subjects factors (for results ignoring non-convergence, see Figure 6E,F). Here, the subject-slopes model held TIE near .05 at 0.0549 ($CI_{95\%}: \pm 0.0047$). For the 794 models that did not converge, the independent slopes model showed lower TIE at 0.0366

(± 0.013), which was similar to the TIE of that model applied to all of the datasets ($p(\text{TIE})=0.042 \pm 0.0039$). Backing off to intercept-only models, however, had much more deleterious consequences with a significantly higher TIE ($p(\text{TIE})=0.068 \pm 0.0175$; compared to $p(\text{TIE})=0.096 \pm 0.0058$ for all datasets). Similarly, power was reduced in the back off options. While the typical power for subject-sloped models was $0.81 (\pm 0.0079)$, when this model failed to converge, the independent-slopes model showed a reduced power of $0.69 (\pm 0.049)$; compared to power of 0.77 ± 0.0089 for all datasets), while the less conservative intercept-only model showed similar power of $0.81 (\pm 0.042)$; compared to power of 0.88 ± 0.0063).

Thus, these results suggest that for between-subject effects, one can back off to either simpler model with no loss of power or TIE. However, for within-subject effects, backing off to models with independent slopes (no covariance) maintains TIE with only a moderate reduction of power, while intercept only models may inflate TIE beyond 0.05 and should not be used.

4 General Discussion

The need to maintain TIE at $\alpha=.05$ has led to recommendations that a maximal random effects structure always be used for fitting mixed models (Barr et al., 2013). Our simulations agree that this is an optimal choice for between-subject effects, as the maximal model maintained appropriate TIE rates for all generating scenarios while having similar power to the true generating model in all instances (for both linear and logistic models). Even if the effect of primary interest is between-subject, power is maximized when other within-subject effects are properly modeled (in the random effects structure).

For within-subject effects (and interactions of between- and within- subject effects), we also agree that the maximal model is the most conservative. However, in contrast to Barr et al. (2013), our results raise the possibility that the maximal model is disadvantageous. If the true

generating structure included random slopes, a maximal model does maintain TIE at 0.05. However, it also reduces power in situations where the true generating model is simpler. This suggests that a model selection approach that selects the random effects structure most likely given the data may be more appropriate. This was seen in simulations with and without item-level slopes (Supplement S3).

Across model selection approaches, model-space techniques almost uniformly outperformed incremental ones. This benefit is certainly enhanced in more complex models than those used here. More complex models may heighten the likelihood of premature stopping, and are often non-nested, making incremental procedures difficult or impossible to design. Within these model-space techniques, both AIC and AICc proved to be the best measures of fit. Both maintain appropriate TIE for generating models with complex random effects structures, while maximizing power for generating models with intercepts only. This was true for linear and logistic models.

The choice of selection criteria should be situation dependent. AIC, for example, performs well with larger sample sizes, and is a simple calculation and computationally efficient. AICc is a better alternative when fitting with more limited sample sizes, as it provides a correction for finite sample sizes (Azari, Li, & Tsai, 2006; Sugiura, 1978). While existing (lme4) implementations output AIC for models, it is trivial to compute AICc from AIC² and this may be advantageous for small samples. Despite its generally poor performance here, BIC may also be of use in a model-space selection scheme (for datasets with other properties). However, BIC is highly dependent on sample size, and in mixed models the sample size is often unclear – is it the number of data points, or subjects (and this in turn may change if data is aggregated)?

² Simply add $(2k^2 + 2k) / (n - k + 1)$ to the original AIC. Here k is the number of free parameters and n is the sample size.

Ultimately, it remains to be seen which information criteria is most suitable for any given situation. Our results favor AIC or AICc in situations that resemble typical laboratory experiments in cognitive psychology and/or work with developmental or clinical populations. Readers should conduct their own simulations on research designs more similar to their own, and the code provided on our OSF site is intended to make such simulations relatively simple.

A historically difficult issue with regards to testing fixed effects within a linear mixed effects framework is the calculation of degrees of freedom. There were instances in our simulations when the Satterthwaite approximation yielded d.f. that were larger than the sample size—even with random slopes—and occasionally larger than the total number of observations. Often this was true even when the model was correctly specified (e.g., a subject-slopes model fitted to subject-slopes generated data). When the d.f. appear unreasonable, it may be appropriate to move to the Kenward-Roger approximation if one desires more accurate d.f.. However, in the whole, these large d.f. had little impact on TIE and power, which were driven more by the standard error (which in turn derives from the random effects specification). This makes sense, in most cases, correctly specified d.f., would be near 50, while incorrectly specified ones may be as high as 9000 (for this dataset). While this seems like a large difference, bear in mind that the critical value of t (for $\alpha=.05$) for 50 d.f. is 2.01, while for 9000 d.f., it is 1.96—the chance of obtaining a T in this range (where significance would be dependent on the d.f.) is likely to be very small. Thus, an unreasonably high d.f. may not be a sole reason for rejecting a model a priori, nor does it necessarily raise concerns that the model is anti-conservative.

We saw minimal differences between raw and aggregated data for hypothesis testing. However, non-convergence rates were systematically lower with aggregated data, for all fitting models that included random slopes. For large datasets, computation time can be significantly reduced with aggregated data. On the other hand, small datasets might see issues with growth in

d.f. approximations that lead to lower power, and with a limited number of factors, fully saturated models (fixed effects, random slopes and intercepts) may include more free parameters than the size of the aggregated data. As it may not be wise to drop random effects to deal with this, fully aggregated data may present a problem in this case. However, when there are multiple factors and multiple random effects this is not usually a problem. Note that we have modeled an experiment in which each item was encountered on multiple repeated trials. As a result, data can be aggregated within-subject and item, preserving the random effects structure. We do not recommend aggregating across items (or subjects) as this can lead to error.

Given the lack of gross effects of aggregation on TIE and power, aggregation may help deal with non-convergence. In this case, it may be appropriate to move to the alternative data option and refit before changing the model structure. However, when model non-convergence is encountered during model selection, it is important to note that all models must be refit using the same data option (aggregated or not) as models with different contributing data cannot be compared. Despite this potential benefit, we note that aggregate data tended to lead to larger SEs of the estimates for intercepts-only fitted models (Supplement 5), even as TIE and power were held constant). Thus, the bias of the researcher should be to use raw data where possible.

Model non-convergence tends to be a much larger issue with logistic mixed models than linear mixed models. As expected, we saw the highest rates of non-convergence when fitting with more complex random effects structures, while we saw no issues when fitting with a minimal model. In general, attempting to overfit the data using a model that includes random slopes when none were truly in the generating model caused higher non-convergence rates. However, non-convergence rates were entirely related to the complexity of the data or the model, and non-convergence may not be a straightforward sign that the data supports a simpler model.

If the non-converging fitted model includes a random slope and covariance terms, our

simulations suggest that backing off to a model that drops the covariance term (but maintains the random slopes) maintains $\alpha=.05$ and has only a moderate decrement in power. Across analyses, such models had similar TIE and power regardless of the underlying generating model, and for both logistic and linear models. However, dropping a random slope term entirely should be avoided as this will increase TIE. If dropping covariance terms is not an option, researchers may consider other options such as switching to an alternative optimizer (e.g., the `optimx` package). It may also be useful to consider an alternative analysis approach such as Generalized Estimating Equations (Burton, Gurrin, & Sly, 1998), which are particularly appropriate if between-subject factors are of primary importance.

We should note the most important caveat to our conclusions as a whole is that it is unclear whether these conclusions hold beyond the experimental design implemented here in the generating functions. We have attempted to create a design that is analogous to typical designs in psychology, and we have extended prior work by considering between-subjects factors (critical in social/personality psychology and clinical fields). It is important to extend these results to more design families (e.g., multiple covariates, nested random effects, autocorrelated DVs and so forth) to draw more robust conclusions. To this end, we have posted the source code for these investigations at <https://osf.io/vj6db/> and we encourage people to use it to explore new designs, and to share the results with us so that we may post them on the OSF page for this project.

Finally, we note that these simulations also highlight the large uncertainty in the consequences of seemingly simple design choices for TIE and power in mixed models (a level of uncertainty also reflected in researcher's feelings about them: Meteyard & Davies, 2020). We have examined several common researcher decisions and our results show that some have large effects on TIE and power (while others may not). Moreover, these results show that we do not always know the consequences of even relatively basic and/or simple decisions for TIE and

power, and that methodological work is needed to pin these factors down. There are enormous researcher d.f. in mixed models, and many opportunities for mistakes (see Brauer & Curtin, 2018; Meteyard & Davies, 2020). This study as well as others (Barr et al., 2013; Matuschek et al., 2017) suggest these researcher d.f. affect the quality of inferences that are made. Thus, we caution readers to avoid jumping into these complex methods without training and guidance. More importantly, we caution Reviewer 2 (you know who you are) to resist the temptation to require authors to “just use a mixed model” as a condition of publication. While mixed models are more robust and statistically better suited to much data in psychology, it is not clear that the likelihood of inference error from a perfect—but poorly used—mixed model is greater than the likelihood of inference error from a slightly less suitable—but correctly implemented—ANOVA.

Conclusion. In general, these simulations suggest that both linear and logistic mixed models are robust across many common implementations such as the choice of aggregating the data or not, and the manner in which d.f. are computed. The random slope structure impacts both TIE and power. For within-subject contrasts (and within- by between-subject interactions), the inclusion of random slopes is necessary for maintaining TIE at .05 when the underlying data came from such a distribution. However, when the data was generated from a simpler distribution this reduces power. Moreover, for between-subject factors, proper characterization of the within-subject variance, increases of power. Both issues can be solved with an appropriate model selection scheme that examines the full space of possible models.

Acknowledgements

The authors would like to thank Teresa Treat, Grant Brown, Daniel Sewell, and Brian Smith for helpful comments on this project. We thank Marc Brysbaert and Lotte Meteyard for helpful comments on an earlier draft of this manuscript. This project was a component of Michael Seedorff’s doctoral dissertation in Biostatistics at the University of Iowa. This project was supported by P50DC000242 awarded to B. Gantz and B. McMurray and DC 008089 to B. McMurray. A prior version of this article was posted to PsyArxiv <https://psyarxiv.com/xmhfr/>. It reports the same simulations

Author Contributions and Disclosures

MS and BM conceptualized the questions and designed the simulations with consultation from JO. MS implemented the simulations and analyzed the results with consultation from BM and JO. All three authors contributed to writing the manuscripts. The authors declare no conflicts of interest with respect to authorship or the publication of this article.

Open Practices Statement

All of the relevant source code for these simulations is provided at <https://osf.io/vj6db/>. We encourage users to adapt the code to examine situations that resemble their own experimental designs and covariance structures, and to leave a squib at the site.

References

- Akaike, H. (1974). A new look at the statistical model identification. *IEEE transactions on automatic control*, 19(6), 716-723.
- Azari, R., Li, L., & Tsai, C. L. (2006). Longitudinal data model selection. *Computational statistics & data analysis*, 50(11), 3053-3066.
- Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59(4), 390-412. doi:10.1016/j.jml.2007.12.005
- 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.
- Bates, D., & Sarkar, D. (2011). lme4: Linear mixed-effects models using Eigen and Eigenpack. GNU Public License.
- Brauer, M., & Curtin, J. J. (2018). Linear mixed-effects models and the analysis of nonindependent data: A unified framework to analyze categorical and continuous independent variables that vary within-subjects and/or within-items. *Psychological Methods*, 23(3), 389-411.
- Burton, P., Gurrin, L., & Sly, P. (1998). Tutorial in biostatistics. Extending the simple linear regression model to account for correlated responses: an introduction to generalized estimating equations and multi-level mixed modeling. *Stat Med*, 17(11), 1261-1291.
- Button, K. S., Ioannidis, J. P. A., Mokrysz, C., Nosek, B. A., Flint, J., Robinson, E. S. J., & Munafò, M. R. (2013). Power failure: why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience*, 14(5), 365-376. doi:10.1038/nrn3475
- Cantoni, E., Flemming, J. M., & Ronchetti, E. (2005). Variable selection for marginal longitudinal generalized linear models. *Biometrics*, 61(2), 507-514.
- Clark, H. H. (1973). The language-as-fixed-effect fallacy: A critique of language statistics in psychological research. *Journal of Verbal Learning and Verbal Behavior*, 12(4), 335-359. doi:10.1016/s0022-5371(73)80014-3
- Cunnings, I. (2012). An overview of mixed-effects statistical models for second language researchers. *Second Language Research*, 28(3), 369-382. doi:10.1177/0267658312443651
- De Boeck, P. (2008). Random Item IRT Models. *Psychometrika*, 73(4), 533. doi:10.1007/s11336-008-9092-x

- Eager, C., & Roy, J. (2017). Mixed Effects Models are Sometimes Terrible. *PsyArxiv preprint, arXiv:1701.04858*.
- Enkavi, A. Z., Eisenberg, I. W., Bissett, P. G., Mazza, G. L., MacKinnon, D. P., Marsch, L. A., & Poldrack, R. A. (2019). Large-scale analysis of test–retest reliabilities of self-regulation measures. *Proceedings of the National Academy of Sciences*, 116(12), 5472–5477. doi:10.1073/pnas.1818430116
- Fan, J., McCandliss, B. D., Sommer, T., Raz, A., & Posner, M. I. (2002). Testing the efficiency and independence of attentional networks. *Journal of Cognitive Neuroscience*, 14, 340–347.
- Field, A. P., & Wright, D. B. (2011). A Primer on Using Multilevel Models in Clinical and Experimental Psychopathology Research. *Journal of Experimental Psychopathology*, 2, 271–293.
- Fitzmaurice, G. M., Laird, N. M., & Ware, J. H. (2012). *Applied longitudinal analysis* (Vol. 998): John Wiley & Sons.
- Halekoh, U., & Højsgaard, S. (2014). A Kenward-Roger approximation and parametric bootstrap methods for tests in linear mixed models—the R package pbkrtest. *Journal of Statistical Software*, 59(9), 1–30.
- Harel, D., & McAllister, T. (2019). Multilevel models for communication sciences and disorders. *Journal of Speech, Language and Hearing Research*, 62(4), 783–801.
- Hui, F. K. C., Muller, S., & Welsh, A. H. (2017). Hierarchical selection of fixed and random effects in generalized linear mixed models. *Statistica Sinica*, 27, 501–518.
- Hurvich, C. M., & Tsai, C. L. (1989). Regression and time series model selection in small samples. *Biometrika*, 76(2), 297–307.
- Ibrahim, J. G., Zhu, H., Garcia, R. I., & Guo, R. (2011). Fixed and random effects selection in mixed effects models. *Biometrics*, 67(2), 495–503.
- Jaeger, T. F. (2008). Categorical data analysis: Away from ANOVAs (transformation or not) and towards logit mixed models. *Journal of Memory and Language*, 59(4), 434–446.
- Jiang, J., Rao, J. S., Gu, Z., & Nguyen, T. (2008). Fence methods for mixed model selection. *The annals of statistics*, 36(4), 1669–1692.
- 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. doi:10.1037/a0028347
- Kenward, M. G., & Roger, J. H. (1997). Small sample inference for fixed effects from restricted maximum likelihood. *Biometrics*, 983–997.
- Koerner, T. K., & Zhang, Y. (2017). Application of Linear Mixed-Effects Models in Human Neuroscience Research: A Comparison with Pearson Correlation in Two Auditory Electrophysiology Studies. *Brain sciences*, 7(3), 26. doi:10.3390/brainsci7030026
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2016). Package ‘lmerTest’: GNU Public License.
- Luke, S. G. (2017). Evaluating significance in linear mixed-effects models in R. *Behavior Research Methods*, 49(4), 1494–1502.
- Matuschek, H., Kliegl, R., Vasishth, S., Baayen, H., & Bates, D. (2017). Balancing Type I error and power in linear mixed models. *Journal of Memory and Language*, 94, 305–315.
- Maxwell, S. E., Lau, M. Y., & Howard, G. S. (2015). Is psychology suffering from a replication crisis? What does “failure to replicate” really mean? *American Psychologist*, 70(6), 487–498. doi:10.1037/a0039400

- Meteyard, L., & Davies, R. A. I. (2020). Best practice guidance for linear mixed-effects models in psychological science. *Journal of Memory and Language*, 112, 104092. doi:<https://doi.org/10.1016/j.jml.2020.104092>
- Miyake, A., & Friedman, N. P. (2012). The Nature and Organization of Individual Differences in Executive Functions: Four General Conclusions. *Current Directions in Psychological Science*, 21(1), 8-14. doi:10.1177/0963721411429458
- Oleson, J., Brown, G., & McCreery, R. (2019). The Evolution of Statistical Methods in Speech, Language, and Hearing Sciences. *Journal of Speech Language and Hearing Research*, 62, 498-506.
- Paap, K. R., & Greenberg, Z. I. (2013). There is no coherent evidence for a bilingual advantage in executive processing. *Cognitive Psychology*, 66(2), 232-258. doi:<http://dx.doi.org/10.1016/j.cogpsych.2012.12.002>
- Paap, K. R., & Sawi, O. (2014). Bilingual advantages in executive functioning: problems in convergent validity, discriminant validity, and the identification of the theoretical constructs. *Frontiers in Psychology*, 5(962). doi:10.3389/fpsyg.2014.00962
- Quené, H., & van den Bergh, H. (2004). On multi-level modeling of data from repeated measures designs: a tutorial. *Speech Communication*, 43(1), 103-121. doi:<https://doi.org/10.1016/j.specom.2004.02.004>
- Quené, H., & van den Bergh, H. (2008). Examples of mixed-effects modeling with crossed random effects and with binomial data. *Journal of Memory and Language*, 59(4), 413-425. doi:<https://doi.org/10.1016/j.jml.2008.02.002>
- R Core Team. (2015). R: A language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from <https://www.R-project.org/>
- Raaijmakers, J. G. W., Schrijnemakers, J. M. C., & Gremmen, F. (1999). How to Deal with "The Language-as-Fixed-Effect Fallacy": Common Misconceptions and Alternative Solutions. *Journal of Memory and Language*, 41(3), 416-426. doi:<http://dx.doi.org/10.1006/jmla.1999.2650>
- Satterthwaite, F. E. (1946). An approximate distribution of estimates of variance components. *Biometrics bulletin*, 2(6), 110-114.
- Schwarz, G. (1978). Estimating the dimension of a model. *The annals of statistics*, 6(2), 461-464.
- Snijders, T., & Bosker, R. (1999). *Multilevel analysis: An introduction to basic and advanced multilevel modeling*. London: Sage Publications.
- Steyerberg, E. W., Eijkemans, M. J. C., & Habbema, J. D. F. (1999). Stepwise selection in small data sets: A simulation study of bias in logistic regression analysis. *Journal of Clinical Epidemiology*, 52(10), 935-942. doi:10.1016/s0895-4356(99)00103-1
- Studebaker, G. A. (1985). A "Rationalized" Arcsine Transformation. *Journal of Speech, Language and Hearing Research*, 28, 455-462.
- Sugiura, N. (1978). Further analysts of the data by akaike's information criterion and the finite corrections. *Communications in Statistics - Theory and Methods*, 7(1), 13-26.
- Tasca, G. A. (2009). Multilevel modeling of longitudinal data for psychotherapy researchers: II. The complexities AU - Gallop, Robert. *Psychotherapy Research*, 19(4-5), 438-452. doi:10.1080/10503300902849475
- Whittingham, M. J., Stephens, P. A., Bradbury, R. B., & Freckleton, R. P. (2006). Why do we still use stepwise modelling in ecology and behaviour? *Journal of Animal Ecology*, 75(5), 1182-1189. doi:10.1111/j.1365-2656.2006.01141.x

**Maybe maximal:
Good enough mixed models optimize power while controlling Type I error**

ONLINE SUPPLEMENTAL MATERIALS

Michael Seedorff
Dept. of Biostatistics, University of Iowa
Alphabet, Inc.

Jacob Oleson
Dept. of Biostatistics
University of Iowa

Bob McMurray
Dept. of Psychological and Brain Sciences
Dept. of Communication Sciences and Disorders
Dept. of Otolaryngology and
DeLTA Center
University of Iowa

Supplement S1: Formal Specifications of the models.

This supplementary note details the formal specifications of each the generating and fitted models used in the main text.

S1.1 Linear Models.

The Intercepts only model is reprinted in lmer (R) notation in Equation 1, and in formal notation in Equation 2. Definition of these terms are given in Table S1.

$$y \sim \text{Age} * \text{Noise} + (1 | \text{Subject}) + (1 | \text{Item}) \quad (1)$$

$$Y_{c,s,i} = \beta_0 + S_{0,s} + I_{0,i} + \beta_1 \text{Age}_s + \beta_2 \text{Noise}_{c,s,i} + \beta_3 \text{Age}_s \text{Noise}_{c,s,i} + \epsilon_{c,s,i} \quad (2)$$

$$S_{0,s} \sim N(0, \sigma_{0,s}^2)$$

$$I_{0,i} \sim N(0, \sigma_{0,i}^2)$$

$$\epsilon_{c,s,i} | S_{0,s}, I_{0,i} \sim N(0, \sigma_{\epsilon}^2)$$

The subject-slope model is reprinted in lmer notation in shown in Equation (3), and in formal terms in Equation (4). It allows non-zero covariance between the random intercept and random slope.

$$y \sim \text{Age} * \text{Noise} + (\text{Noise} | \text{Subject}) + (1 | \text{Item}) \quad (3)$$

$$Y_{c,s,i} = \beta_0 + S_{0,s} + I_{0,i} + \beta_1 \text{Age}_s + (\beta_2 + S_{1,s}) \text{Noise}_{c,s,i} + \beta_3 \text{Age}_s \text{Noise}_{c,s,i} + \epsilon_{c,s,i} \quad (4)$$

Table S1. Definition of terms in the Equations

Term	Definition
$Y_{c,s,i}$	Dependent variable Y in condition (noise level) c , for subject s , word (item) i .
$P_{c,s,i}$	Dependent variable p (probability of response) for logistic models.
β_0, β_i (etc)	Coefficients of fixed effects estimated by the model. β_0 indicates the (fixed) intercept of the model.
Age_s	Age group of subject s (specified in the data)
$\text{Noise}_{c,s,i}$	Noise level on the trial specified by c, s, i (specified in the data).
$S_{0,s}$	Random effect of subject (intercept indicated by the $_0$). Assumed to derive from a Normal distribution with a mean of 0 and a variance of $\sigma_{0,s}^2$.
$I_{0,i}$	Random effect of Item (word) (intercept indicated by the $_0$). Assumed to derive from a Normal distribution as with $S_{0,s}$.
$\epsilon_{c,s,i}$	Noise on that trial.

For models with random slopes (e.g., equation 3), a subject or item specific slope is added to the fixed effect (e.g., β_1) before multiplying it by *Noise* or *Age*

$S_{1,s}$	Random slope of noise on subject (slope indicated by the $_1$)
$I_{1,i}$	Random slope of noise on item (word) (slope indicated by the $_1$)

$$\begin{aligned} \begin{bmatrix} S_{0,s} \\ S_{1,s} \end{bmatrix} &\sim N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_{0,s}^2 & \rho\sigma_{0,s}\sigma_{1,s} \\ \rho\sigma_{0,s}\sigma_{1,s} & \sigma_{1,s}^2 \end{bmatrix}\right) \\ I_{0,i} &\sim N(0, \sigma_{0,I}^2) \\ \epsilon_{c,s,i} | S_{0,s}, S_{1,s}, I_{0,i} &\sim N(0, \sigma_\epsilon^2) \end{aligned}$$

Here, the effective effect of Noise is a product of both the fixed effect (β_1) and a subject-specific slope ($S_{1,s}$), which is drawn from a random normal distribution with a mean of 0 and a specified variance (the third line of Equation 4). The second line of Equation 4 shows the covariance among the slopes and intercepts.

Finally, the independent slopes model is reprinted in lmer notation in Equation 5, and in formal notation in Equation 6.

$$y \sim \text{Age} * \text{Noise} + (\text{Noise} \parallel \text{Subject}) \quad (5)$$

$$Y_{c,s,i} = \beta_0 + S_{0,s} + I_{0,i} + \beta_1 \text{Age}_s + (\beta_2 + S_{1,s}) \text{Noise}_{c,s,i} + \beta_3 \text{Age}_s \text{Noise}_{c,s,i} + \epsilon_{c,s,i} \quad (6)$$

$$\begin{aligned} \begin{bmatrix} S_{0,s} \\ S_{1,s} \end{bmatrix} &\sim N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_{0,s}^2 & 0 \\ 0 & \sigma_{1,s}^2 \end{bmatrix}\right) \\ I_{0,i} &\sim N(0, \sigma_{0,I}^2) \\ \epsilon_{c,s,i} | S_{0,s}, S_{1,s}, I_{0,i} &\sim N(0, \sigma_\epsilon^2) \end{aligned}$$

S1.2 Logistic Models

The intercept-only logistic model is the same as Equation 1 in lmer notation with the exception that link function is binomial. It is specified formally in Equation 7.

$$p_{c,s,i} = \text{logistic}(\beta_0 + S_{0,s} + I_{0,i} + \beta_1 \text{Age}_s + \beta_2 \text{Noise}_{c,s,i} + \beta_3 \text{Age}_s \text{Noise}_{c,s,i}) \quad (7)$$

$$\begin{aligned} y_{c,s,i} &= \text{Bernoulli}(\pi = p_{c,s,i}) \\ S_{0,s} &\sim N(0, \sigma_{0,s}^2) \\ I_{0,i} &\sim N(0, \sigma_{0,I}^2) \end{aligned}$$

Terms are specified in the Table S1, and the logistic formula is given in 8.

$$\text{logistic}(x) = \frac{1}{1+e^{-x}} \quad (8)$$

The subject slopes model is given in Equation (9)

$$p_{c,s,i} = \text{logistic}(\beta_0 + S_{0,s} + I_{0,i} + \beta_1 \text{Age}_s + (\beta_2 + S_{1,s}) \text{Noise}_{c,s,i} + \beta_3 \text{Age}_s \text{Noise}_{c,s,i}) \quad (9)$$

$$\begin{aligned} y_{c,s,i} &= \text{Bernoulli}(\pi = p_{c,s,i}) \\ \begin{bmatrix} S_{0,s} \\ S_{1,s} \end{bmatrix} &\sim N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_{0,s}^2 & \rho\sigma_{0,s}\sigma_{1,s} \\ \rho\sigma_{0,s}\sigma_{1,s} & \sigma_{1,s}^2 \end{bmatrix}\right) \end{aligned}$$

$$I_{0,i} \sim N(0, \sigma_{0,I}^2)$$

Finally, the independent slopes model is given in (10).

$$\begin{aligned}
p_{c,s,i} &= \text{logistic}(\beta_0 + S_{0,s} + I_{0,i} + \beta_1 \text{Age}_s + (\beta_2 + S_{1,s}) \text{Noise}_{c,s,i} + \beta_3 \text{Age}_s \text{Noise}_{c,s,i}) \\
y_{c,s,i} &= \text{Bernoulli}(\pi = p_{c,s,i}) \\
\begin{bmatrix} S_{0,s} \\ S_{1,s} \end{bmatrix} &\sim N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_{0,s}^2 & 0 \\ 0 & \sigma_{1,s}^2 \end{bmatrix}\right) \\
I_{0,i} &\sim N(0, \sigma_{0,I}^2)
\end{aligned} \tag{10}$$

Supplement S2. Power and TIE for interactions.

In both the linear and logistic models, the comparison of different random effects procedures had a highly similar (if not identical) pattern for the within/between interaction as for the within-subject effect. Full results are reported here.

Linear Models.

Figure S2.1 reports results for the within/between interaction in the linear model. Results largely

Figure S2.1: (A) TIE and (B) power results for the *interaction effect* corresponding to the random intercept only generating model. (C) TIE and (D) power results for the interaction effect corresponding to the independent slopes generating model. (E) TIE and (F) power results for the interaction effect corresponding to the fully sloped generating

Fig. S2.2: (A) TIE and (B) power results of the logistic mixed model for the *interaction of the within and between-subjects effects* for the data generating model with intercepts only. (C) TIE and (D) power results for the within-subjects effect corresponding to the independent slopes generating model. (E) TIE and (F) power results for the within-subjects effect corresponding to the fully sloped generating model.

replicate the within-subjects effect. Figures S2.1A/B display results for the intercepts-only generating model. We see reduced power in the more complex models (independent slopes and fully sloped) while all model selection procedures perform comparably to the true intercept only. For the independent slopes generating model (Figures S2.1C/D), the intercepts only fitted model is extremely anti-conservative with TIE over 0.1. In contrast, both fitting models that include random slopes do appropriately well. This drives anti-conservative behavior in the model-space BIC selection approach, LRT2, bLRT3, and LRT3 as they are more likely to choose a simpler

model, while model-space AIC/AICc maintains an appropriate TIE.

Finally, for the fully sloped generating model in Figures S1E/F, we see similar results. The anti-conservative nature of the intercepts-only fitting model leads to anti-conservativeness in some model selection approaches except for LRT2 and bLRT3, and AIC/AICc, which perform similarly to the correct fully sloped fitted model, holding TIE below .05.

Logistic Models

The analysis of the interaction in the logistic models is shown in Figure S2.2. Like the within-subject effect reported in the main text (Figure 6), it shows reduced power and TIE when more complex models are fit to intercept-only data (Figure S2.2A,B). When the underlying data includes random slopes (Figure S2.2C-F), the application of an intercept-only fitting model dramatically increased TIE, whereas either independent-slopes or subject-slopes were sufficient to eliminate this anti-conservatism (regardless of whether the true model included covariance terms). AIC, AICc and most of the LRTs appeared appropriate, whereas BIC tended to be anti-conservative, as it was more likely to inappropriately choose the simpler model.

Supplement S3: Crossed Random Effects

The simulations presented in the main text simplified the problem of random effects selection to examine the consequences of various random slope selection schemes in mixed designs (as between-subject effects have not been studied). As a result, the data fitting models that were considered focused on random slopes of subject (holding the random effects of item to intercept-only). This showed that for both within- and between-subject effects, a selection scheme in which the entire space of possible models is considered (and compared using AIC or AICc) offers the best way to maximize power while controlling TIE. This was not an approach considered by either of the prior studies on this issue (Barr, Levy, Scheepers, & Tily, 2013; Matuschek, Kliegl, Vasishth, Baayen, & Bates, 2017).

However, both of these prior studies considered more complex crossed-random effects structures than what we have examined thus far. Thus, we replicated the above simulations to consider differences in both subject and item effects. Here all simulations parameters are the same except for the fact that *fully sloped* fitting and generating models include slopes of the within-subject factor on subject and item. Here, the variance of the subject slopes (across subjects) was the same as in the first set simulations, and the variance in the item slopes (across items) was equal to that of the subjects (Equation 11, in lmer format).

$$y \sim \text{Age} * \text{Noise} + (\text{Noise} \mid \text{Subject}) + (\text{Noise} \mid \text{Item}) \quad (11)$$

In order to keep the simulations tractable, we only compare three models: intercept-only, fully-sloped, and the subject-slopes model from the main text. This latter is commonly used in psycholinguistics if a fully sloped model does not converge or is otherwise not possible. We omit permutations like including slopes only on items but not subjects (which is almost never done), and did not examine dropping covariance terms as there was no reason to believe results would differ in this analysis (the reader is invited to use our publicly available source code to investigate these, if they are of interest).

Between-Subjects Factor. Figure S3.1A and B show results for the between-subject effect for the random intercepts-only generating model. In this case, all models provide similar results for TIE and power. Figs. S3.1C and D detail results for the between-subjects factor effect corresponding to the fully sloped item and subject generating model. It is clear that a random slope needs to be included in the model, but we found no differences in TIE or power whether the item random slope is included in addition to the subject random slope. Figures S3.1E and F show results from the subject-slope model. Again, we see no loss of TIE or power if including the random slope for item in the analysis – even though there are no item slopes in the underlying data. Results from the model selection procedures are shown in the bottom of each panel of Figure S3.1. All model selection procedures but BIC maintain appropriate TIE and similar power. BIC is more conservative when random slopes are included.

For models generated with random slopes, power was improved by including random slopes, but the random intercept for item did not appreciably adjust TIE or power. Model selection procedures were very similar with BIC being the most conservative.

Within Subject Factor. We next look at the results for the within-subjects factor. Figures S3.2A and B display the results for the intercept-only generating model. Applying complex models that include random slopes show reduced TIE but also reduced power. For the generating model with random slopes for item and subject (Figure S3.2C, D), the intercepts only model is anticonservative. However, both fitting models that include at least one slope do well. The best

Figure S3.1: (A) TIE and (B) power results for the *between-subjects effect* for the intercept-only generating model. (C) TIE and (D) power results in the fully sloped generating model (subject + item slopes). (E) TIE and (F) power results for the between-subjects effect in the subject-slopes generating model (subject slopes + item intercepts).

model selection procedures are AIC and AICc. Similar results are seen in Figures S3.2E and F in terms of TIE and power, again including at least one random slope appears to be important. As seen in other analyses BIC performs poorly, as it tends to favor the simpler models. However, the other model selection methods perform reasonably well.

The Interaction. The results for the interaction (Figure S3.3) are nearly identical to previous results and largely support the conclusions of the within-subject effects.

Summary. In summary, when the data generating model includes random slopes, the analysis should include a random slope. In this particular case, there appears to be no effective

Figure S3.2: (A) TIE and (B) power results for the *within-subjects effect* corresponding to the random intercept only generating model. (C) TIE and (D) power results in the fully sloped generating model (subject + item slopes). (E) TIE and (F) power results for the between-subjects effect in the subject-slopes generating model (subject slopes + item intercepts).

difference between including both random slopes (on subject or item) or just one (subject). In terms of model selection, AIC and AICc do well at maintaining TIE rates, as they do in the primary simulations. We note that this is a situation in which both primary independent variables (age and noise) were both within-item and within-subject. This is a situation in which item analyses have been argued to be less important (Raaijmakers, Schrijnemakers, & Gremmen, 1999), than in common Latin-squares or between-item designs. Thus, we would not advocate on the basis of these simulations that item slopes should be dropped in all circumstances. However,

Figure S3.3: (A) TIE and (B) power results for the **interaction effect** corresponding to the random intercept only generating model. ((C) TIE and (D) power results in the fully sloped generating model (subject + item slopes). (E) TIE and (F) power results for the between-subjects effect in the subject-slopes generating model (subject slopes + item intercepts).

broadly speaking, these simulations support the primary conclusions that a) random slopes are necessary (when the generating function includes them) to preserve TIE, and b) that model-space AIC and AICc procedures appear to appropriately control TIE while maximizing power.

S.4. Degrees of Freedom

As part of this project, we also evaluated the d.f. output by the Satterthwaite procedure. This was saved for each fixed effect for each run of the model.

D.f. was stable for the between-subject factor, with estimates from 48 to 50, roughly in line with the number of subjects.

For the within-subject factor, estimated d.f. varied greatly, depending on the generating and fitting models. With the intercepts-only fitted model, the median d.f. was 9,930 for all generating models (close to the 10,000 observations), and for 1.68% of simulations d.f. was nonsensically larger than 10,000! Alternatively, consider fitted models with random slopes. If the data suggests subject-specific slopes exist (as in the subject-sloped generating model, and occasionally with the intercepts only generating model) then the d.f., should be approximately the same as for the between-subjects factor. Indeed, when the generating- and fitted-model both included subject slopes: median d.f. ranged from 48 to 49. However, when the data was generated from the intercepts-only model but fit with a random slope model, the median increased to 9,875, closer to the number of overall observations. In this case too, we sporadically observed d.f. above 10,000: when intercepts-only data was fit with the subject-sloped model (1 instance), independent-slopes (70 instances), or even intercepts-only (162 instances). Thus, even a maximal fitted model can yield unreasonably large d.f. estimates when applied to data that does not require random slopes. Thus, even under conservative approaches (like “keep it maximal”), the d.f. may reflect the size of the dataset, not the design. Consequently, it is unwise to over rely on d.f. as a sanity check for an appropriately conservative model.

The prior section showed that when the fitting model does not include random slopes, but the generating model does, then detection of fixed effects is anti-conservative. This could derive from one of two sources. First, it could be partly due the fact that d.f. will be larger than appropriate. Second, the d.f. could be a minor factor, and the increase in TIE derives from a misspecification of the variance structure. To test this, we revised simulations in which the intercepts-only fitted model was applied to the two random slopes generating data. For these models, we recalculated the p-values using 48 d.f. (matching the design of the experiment). This does reduce observed TIE rate slightly (independent-slopes generating model: from 0.1054 to 0.0974; subject slopes: from 0.1027 to 0.0944). However, these reductions do not come close to achieving a TIE of .05. This indicates underestimation of the standard error is the primary problem, not the estimates of d.f.. More broadly, if the model specification is accurate and justified by the data, the d.f. should not be used to evaluate suitability.

Supplement S5: Effect of Data Aggregation on SE of Fixed Effects

As part of our investigation of aggregation, we asked how the choice to use raw or aggregated data affected the parameter estimates. We focused on the SE of the estimates, which are essential for determining power and TIE, and we focused on the within-subjects factor where we have seen the largest impact of various model choices.

True SE values were estimated by computing the SDs of the beta estimates across the 10,000 fitted models. We then asked if the observed SE in any given model was higher or lower than the true value on average (SE bias) in order to determine if data aggregation resulted in over-confident SEs. Next, we investigated the variation in the SE from run to run to see if it improved the stability of the estimates.

SE Bias estimates (Figure S5.1) were consistent with observed TIE rates in Section 3.1. In situations with anti-conservative TIE rates (the generating model included random slopes and the fitted model did not; the right most points of Figure S5.1A), SEs were negatively biased for both aggregated and raw data (and similarly so). That is, the SEs were lower than they should be – the source of the increased TIE. However, when the both the generating and fitted model were intercept-only, SEs were near 0. Crucially, this was not affected by the choice to aggregate the data or not. When we examined the more complex models (Figure S1.5B, C), the story also confirmed what was seen with TIE and power. In both of these models, observed SE was much closer to 0 and was larger (not smaller), suggesting a more conservative approach. However in both of these models, data aggregation made little difference. The fact that data aggregation did not affect bias was perhaps unsurprising, as these results are averaged over a large number of simulations (and this measure emphasizes the mean across runs).

We next computed the standard deviation of the SEs from run to run as a measure of how likely any given model is to estimate the true SEs. Figure S5.2 suggests that the SD of the SEs was generally lower when fitting with the intercept-only model (Panel A) than the more complex models (Panels B,C). In light of the decreased non-convergence rates for these models, this likely reflects the greater stability of these simpler models from run to run. However, within these models, the SD of the SE estimates was larger for models using aggregated data than for models using the full dataset ($p < .0001$ for all generating models). This suggests that with the removal of trial-by-trial level variance, the SEs become more variable (in the intercept-only fitted model). However, when we considered the two more complex fitting models (Figure

Figure S5.1. Bias of the standard error estimates for the within-subjects factor as a function of the true generating model. (A) For the intercept-only model; (B) For the independent slopes fitted model; and (C) for the Subject-Slopes model. Bias was defined as the difference between the average SE across runs and the SD of the estimates.

Figure S5.2. SD of the standard error estimates across runs for the within-subjects factor as a function of the true generating model. (A) For the intercept-only model; (B) For the independent slopes fitted model; and (C) for the Subject-Slopes model.

S5.2B,C), a different pattern emerged. Here, ignoring the intercept only model, SEs were generally less variable when the covariance structure of the matched that assumed by the generating model (e.g., in Panel B, the independent slopes generating model had the lowest SD, but in Panel C, the subject-slopes model showed a lower SD). However, in both of these cases, there was no observable difference between raw and aggregated data.

Thus, with respect to the SEs of the parameter estimates, the decision to aggregate or not appears to only make an impact when using the intercept-only fitted model.

References

- 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.
- Matuschek, H., Kliegl, R., Vasishth, S., Baayen, H., & Bates, D. (2017). Balancing Type I error and power in linear mixed models. *Journal of Memory and Language*, 94, 305–315.
- Raaijmakers, J. G. W., Schrijnemakers, J. M. C., & Gremmen, F. (1999). How to Deal with “The Language-as-Fixed-Effect Fallacy”: Common Misconceptions and Alternative Solutions. *Journal of Memory and Language*, 41(3), 416-426.
doi:<http://dx.doi.org/10.1006/jmla.1999.2650>