Heterogeneous Heterogeneity by Default: Testing Categorical Moderators in Mixed-effects Meta-Analysis

Josue E. Rodriguez¹, Donald R. Williams^{1, 2}, and Paul-Christian Bürkner³

¹University of California, Davis

²NWEA, Portland, Oregon

³Cluster of Excellence SimTech, University of Stuttgart, Germany

Author Note

Josue E. Rodriguez <https://orcid.org/0000-0002-9092-4869>

Donald R. Williams <https://orcid.org/0000-0001-6735-8785>

Paul-Christian Bürkner <https://orcid.org/0000-0001-5765-8995>

All authors have read and approved the final version of this article. The authors have no conflicts of interest with respect to their authorship or the publication of this article. DRW was supported by a National Science Foundation Graduate Research Fellowship under Grant No. 1650042. P.-C.B was partially funded by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2075 – 390740016. We thank Brenton Wiernik for helpful comments and discussions on earlier versions of this draft. The data that support the findings of this study are openly available on the Open Science Framework at <https://osf.io/gs6uw/>.

Abstract

Categorical moderators are often included in mixed-effects meta-analysis to explain heterogeneity in effect sizes. An assumption in tests of categorical moderator effects is that of a constant between-study variance across all levels of the moderator. Although it rarely receives serious thought, there can be statistical ramifications to upholding this assumption. We propose that researchers should instead default to assuming *unequal* between-study variances when analyzing categorical moderators. To achieve this, we suggest using a mixed-effects location-scale model (MELSM) to allow group-specific estimates for the between-study variance. In two extensive simulation studies, we show that in terms of Type I error and statistical power, little is lost by using the MELSM for moderator tests, but there can be serious costs when an equal variance mixed-effects model (MEM) is used. Most notably, in scenarios with balanced sample sizes or equal between-study variance, the Type I error and power rates are nearly identical between the MEM and the MELSM. On the other hand, with imbalanced sample sizes and unequal variances, the Type I error rate under the MEM can be grossly inflated or overly conservative, whereas the MELSM does comparatively well in controlling the Type I error across the majority of cases. A notable exception where the MELSM did not clearly outperform the MEM was in the case of few studies (e.g., 5). With respect to power, the MELSM had similar or higher power than the MEM in conditions where the latter produced non-inflated Type 1 error rates. Together, our results support that assuming unequal between-study variances is preferred as a default strategy when testing categorical moderators.

Keywords: meta-analysis, random-effects, mixed-effects, location-scale modeling, heterogeneity, moderator

Heterogeneous Heterogeneity by Default: Testing Categorical Moderators in Mixed-effects Meta-Analysis

Meta-analysis is an indispensable technique for synthesizing results from various comparable studies. Key goals include determining the average effect size (Hedges & Pigott, 2001), quantifying the extent of heterogeneity in effects (Higgins & Thompson, 2002; Ioannidis et al., 2007), and evaluating whether study characteristics moderate the effect size (Hedges & Pigott, 2005). The study of statistical methods for meta-analysis is important because systematic reviews using meta-analytic techniques tend to be influential (DeGeest & Schmidt, 2010). For example, in psychological science, meta-analysis has been used to critically evaluate prominent theories (Hagger et al., 2016), provide nuanced perspectives on phenomena of interest (van Agteren et al., 2021), and suggest potential reasons for why replication studies fail (Klein et al., 2018). It is thus paramount to use methods that provide accurate estimates and high quality inferences.

Moderator variables are often included in meta-analytic models to investigate whether particular study characteristics (e.g., measurement instrument) explain differences in effect sizes, and are typically viewed as a means to explain between-study heterogeneity (Thompson & Sharp, 1999). In a standard random-effects framework, the leftover, unexplained between-study variance is then captured by a heterogeneity parameter, τ^2 . Note that the term mixed-effects model is often used in reference to a random-effects model with moderators, and we use both terms interchangeably throughout this paper. An implicit assumption in moderator analyses that has yet to be thoroughly examined is that τ^2 is fixed across all subgroups for a categorical moderator, or across all values for a continuous moderator. That is, researchers commonly assume that, say, both subgroups of a dichotomous moderator will have the same value for τ^2 . This is analogous to the assumption of homoscedasticity in analysis of variance (ANOVA) designs. When this assumption does not hold true in the population and the subgroups of the moderator have unbalanced sample sizes, this can result in increased Type I error rates, conservative Type

I error rates, and a loss in statistical power for moderator effects (Rubio-Aparicio et al., 2017).

One way of dealing with heterogeneous between-study variances with categorical moderators is with subgroup analyses (Borenstein & Higgins, 2013; Schoemann, 2016), wherein a separate random-effects model is fit to each level in the moderator and separate estimates of τ^2 for each subgroup are obtained. Due to their flexibility to examine multiple predictors within a single modeling framework, a popular alternative to subgroup analyses are mixed-effects meta-regression models (Thompson & Sharp, 1999). Standard mixed-effects meta-regressions, or mixed-effects models do not permit heterogeneous variances between subgroups, but Viechtbauer and López-López (2022) recently described the mixed-effects location-scale model (MELSM) in order to overcome this limitation, while Williams et al. (2021) introduced this model under a Bayesian framework. Using a MELSM, separate estimates of τ^2 can be obtained for each level of the moderator (scale), which in turn yield modified estimates of so-called moderator effects (location). As we will show, allowing for separate heterogeneity estimates curbs the loss of power and distorted error rates when there are unequal between-study variances and unequal sample sizes in subgroups of a moderator. Unlike previous work on MELSMs, the intent of this work is not to introduce and describe the mixed-effects location-scale model and its applications, but rather to understand whether it is preferable as a default strategy over a standard mixed-effects meta-regression.

Analogy to the Two Independent Samples *t*-test

To understand why the MELSM may be suitable as a default meta-analytic model, it is informative to consider the *t*-test. There is a rich and storied literature examining the *t*-test when the population variances of the two groups are heterogeneous (Bartlett, 1936; Murphy, 1967; Welch, 1938, 1947). When the sample sizes of the groups are balanced, inferences stemming from Student's *t*-test are generally unaffected by unequal group variances, save for extreme cases (Scheffé, 1959, Ch. 10). Conversely, when the variances of

the groups and their respective sample sizes are both heterogeneous, the resulting statistics can be drastically biased and lead to invalid inferences — the well-known Behrens-Fisher problem. To overcome these issues, Welch’s t -test is commonly used, and has even been recommended to be used by default because it adjusts the degrees of freedom according to group sample sizes and variances (Delacre et al., 2017).

In the same vein, Rubio-Aparicio et al. (2020), Rubio-Aparicio et al. (2017) find that in subgroup analysis, using a pooled estimate of τ^2 (i.e., assuming equal between-study variances) yields similar inferences to using separate estimates (i.e., assuming unequal between-study variances) if the subgroup sample sizes are at least roughly equal. They conclude that pooling estimates for τ^2 is to be preferred in most scenarios — a sentiment that is echoed in the broader meta-analysis literature (e.g., Borenstein et al., 2009; Viechtbauer, 2010). We view this hesitation to assume unequal between-study variances as unwarranted because as we will show, little is lost by simply assuming an unequal variance model, but there are potentially costly consequences to assuming equal variances.

Overview

The outline of this paper is as follows. We first outline the standard mixed-effects model and the MELSM. We then describe two popular tests of moderator effects and highlight the importance of τ^2 in estimating the moderator coefficients and their respective standard errors. This section also clarifies how moderator tests are essentially weighted tests of mean differences. Next we present extensive simulation studies comparing the error rates and statistical power of moderator tests using the MELSM to those of obtained using a classical mixed-effects model. We conclude with an applied illustration of the MELSM and a brief discussion.

Mixed-effects models

In meta-analysis, researchers are routinely faced with the choice of using either a fixed-effects or random-effects model. The former provides inferences on the effect sizes of the observed studies or a set of identical studies (i.e., conditional inference). The latter

yields inferences on the parameters of a population of studies; of which the observed set is assumed to be a random sample (i.e., unconditional inference; Hedges & Vevea, 1998). In addition to sampling variability, the random-effects model considers the possibility of heterogeneous true effect sizes that arise as a result of differences in study characteristics (e.g., design, population, environmental factors, etc.; DerSimonian & Laird, 1986; Hedges, 1992). The main consideration in deciding between these two models is the inferential goal. In practice, analysts often wish to generalize beyond the observed studies to the parent population of studies, and thus random-effects models are more often employed.

Standard Mixed-effects Model

When a moderator is included in a random-effects model to explain part of the variance in the observed effect sizes, it becomes a mixed-effects model. Additionally, for categorical moderators, the inferential goal is usually to contrast between two or more group average effects. In this case, the mixed-effects model can be understood as follows. Suppose we have observed $j = 1, \dots, k_i$ independent effect sizes in each of $i = 1, \dots, p$ groups, y_{ij} , and that the total number of observed effect sizes is $k = \sum_i k_i$. For simplicity, suppose that $p = 2$ (i.e., two groups). Then, the standard mixed-effects model is given by

$$y_{ij} = \theta_{ij} + \epsilon_{ij} \tag{1}$$

$$\theta_{ij} = \mu_i + u_{ij} \tag{2}$$

$$\mu_i = \beta_0 + \beta_1 x_{ij} \tag{3}$$

$$u_{ij} \sim \mathcal{N}(0, \tau^2) \tag{4}$$

$$\epsilon_{ij} \sim \mathcal{N}(0, \sigma_{ij}^2) \tag{5}$$

where θ_{ij} is the true effect size of the j th study in the i th group, β_0 is the average effect for the first group, β_1 captures the difference in average effect between the first and second group, and x_{ij} is a dummy variable indexing the group to which each observed effect size

belongs.

The terms ϵ_{ij} and u_{ij} are the within- and between-study errors, respectively. The within-study errors are assumed to be normally distributed with zero mean and variance σ_{ij}^2 . Each σ_{ij}^2 is assumed to be known in advance and is calculated according to the type of effect size index being investigated (Hedges & Olkin, 1985). The between-study errors are similarly assumed to be normally distributed with zero mean and variance τ^2 . This parameter captures the variation in the true effect sizes, θ_{ij} , that is not explained by the moderators, and can be estimated a number of ways (see Langan et al., 2017, Appendix C). In a frequentist framework, perhaps the most common estimator of τ^2 is the restricted maximum likelihood (REML; Patterson & Thompson, 1971) estimator, although τ^2 is also often estimated using Bayesian techniques (e.g., Higgins et al., 2009). Importantly, in the mixed-effects model, τ^2 does not carry a subscript and hence the assumption of equal variances amongst groups of a categorical moderator is always made.

This assumption cannot be overlooked because the role of τ^2 is crucial in calculating and testing moderator effects. The way in which the between-study variance affects the outcome of moderator tests is evident in the estimating equations for β_0 and β_1 and their standard errors. These equations are solved using weighted least-squares (Viechtbauer, 2010) where the weights are calculated based, in part, on the between-study variances. When there are two groups, the equations for testing a moderator effect (i.e., β_1) bear a striking resemblance to those of a weighted Student's t -test. For instance, the well-known Q statistic with two groups can be expressed as (Hedges & Pigott, 2005)

$$Q = \frac{(\bar{y}_1 - \bar{y}_2)^2}{w_1^{-1} + w_2^{-1}}, \quad (6)$$

where, just as for the t statistic in an independent two-sample t -test, the numerator reflects the observed mean difference in the outcome between groups and the denominator captures

the pooled variance of the numerator (see [Tests for Categorical moderators](#) for details.) In the case of more than two groups, testing a moderator effect corresponds to a weighted ANOVA (Borenstein et al., 2009). A key assumption of standard ANOVA methods is that the within-group variances are equal, and this assumption is also commonly made in meta-analysis, but with respect to the between-study variance¹. However, recall that our aim is to understand whether this assumption is warranted.

Mixed-effects Location-Scale Model

To accomplish the goal of estimating and incorporating heterogeneous variances into tests of moderator effects, we propose using a mixed-effects location-scale model (MELSM; Viechtbauer & López-López, 2022; Williams et al., 2021). This technique was first introduced outside of meta-analysis² (Hedeker et al., 2008, 2012), and has also been studied under the terms doubly hierarchical model (Lee & Nelder, 2006) and distributional regression (Bürkner, 2018). This model is formulated similarly to the standard mixed-effects model in (1), but an additional component is stipulated to permit moderators to influence the between-study variance. When a categorical moderator is included, it results in group-specific estimates for the between study-variance, τ_i^2 . If we continue with our simplifying assumption that $p = 2$, then τ_i^2 can be modeled through the log-linear equation as

$$\tau_i^2 = \exp(\gamma_0 + \gamma_1 x_{ij}) \quad (7)$$

where γ_0 captures the between-study heterogeneity for the reference group τ_1^2 , γ_1 captures the difference between τ_1^2 and τ_2^2 , and x_{ij} is the same group-indicator variable used in (3).

¹ The within-study variances in a meta-analysis are heterogeneous by default and are typically considered to be known in advance.

² Note the “location-scale” term arose because of the inclusion of a random effect for both the location and level-1 variance models in a non-meta-analytic mixed-effects model, but we use it here to refer to a mixed-effects meta-regression that allows moderators for the location and level-2 variance.

The right hand side of (7) is exponentiated to ensure only positive values are estimated for τ_i^2 . If only an intercept is included in predicting between-study variance, then this results in an estimate of the average heterogeneity across all k studies, or the usual estimate of τ^2 . In this sense, the mixed-effects model can be considered a special case of the MELSM. Additionally, the coefficients γ_0 and γ_1 are typically estimated using either REML (Viechtbauer, 2021) or Bayesian methods (Williams et al., 2021).

Critically, because the MELSM allows each group to have its own value for τ_i^2 , it follows that a different set of weights are used in estimating the coefficients β_0 and β_1 (see Hedges & Pigott, 2005, p. 442), as well as their respective standard errors. So just as the moderator tests using the mixed-effects models are akin to weighted tests of group means with equal variances, moderator tests using the MELSM can be likened to weighted tests of group means with unequal variances.

Tests for Categorical moderators

There are several popular methods in meta-regression for testing moderator effects. Of these, perhaps the most widely used is the Wald-type z -test (Borenstein et al., 2009). This test, however, is known to not adequately control the Type I error rate because it ignores the uncertainty involved in estimating the between-study variance (Knapp & Hartung, 2003; Viechtbauer et al., 2015). As we will see, this in turn affects the estimates and standard errors of the moderator effect. An alternative method originally proposed by Knapp and Hartung (2003) amends this deficiency by applying a correction factor to the variance of the moderator effect (e.g., Rubio-Aparicio et al., 2020). For these reasons, we examine how the weights resulting from the mixed-effects model and the MELSM affect both the z -test and Knapp and Hartung (KH) method.

To crystallize our ongoing analogy between classical tests for mean differences and tests of moderator effects, the expressions we use are for the t -test and ANOVA-like counterparts of the z -test and KH method, but they are equivalent to the meta-regression equations when there are two groups under consideration. That is, the omnibus Q -test is

equivalent to the z -test for meta-regression coefficients when there are two levels in a categorical moderator (cf. Hedges & Pigott, 2005, pp. 434–435 and pp. 442–443).

Similarly, the KH method for regression coefficients (Knapp & Hartung, 2003) is equivalent to the F -test described in Hartung et al. (2001) in the case of two groups. Describing the tests in these forms will help us clearly show that, in fact, assuming unequal between-study variances in meta-analyses is no different than assuming unequal variances in classical tests of mean differences. Importantly, the latter is already routine practice in psychological science.

Q statistic

In this section and the one below, we focus on frequentist estimation of moderator effects, but it should be mentioned that the same logic we use would similarly hold for Bayesian estimation (Williams et al., 2021). Testing a moderator effect entails computing a series of *weighted* means. The first step is to calculate the weighted mean effect size for each group, where the weights are given by the inverse of the total variance in y_{ij} . Under the standard mixed-effects model (MEM), we have the weights

$$w_{ij}^* = \frac{1}{\sigma_{ij}^2 + \tau^2}. \quad (8)$$

Under the MELSM, the weights are given by (Williams et al., 2021)

$$w_{ij}^{**} = \frac{1}{\sigma_{ij}^2 + \tau_i^2}. \quad (9)$$

As aforementioned, the key difference in the weights used by the MEM and the MELSM is that the former assumes a common τ^2 across all i groups whereas the MELSM can incorporate group-specific heterogeneity parameters, τ_i^2 . Using these weights, the mean for

each group can be calculated as

$$\bar{y}_i = \frac{\sum_{j=1}^{k_i} w_{ij} y_{ij}}{\sum_{j=1}^{k_i} w_{ij}}. \quad (10)$$

Note that here and below, we use the term w_{ij} in a generic sense to indicate either w_{ij}^* or w_{ij}^{**} , depending on whether equal or unequal variances are assumed. If there are no differences in between-study variances ($\tau_1^2 = \dots = \tau_p^2 = \tau^2$), then the weights in (9) are equal to those in (8). Additionally, the τ_i^2 are not usually known in practice, and so are usually replaced with their estimates, $\hat{\tau}_i^2$.

The i th group mean is assumed normally distributed with mean μ_i and variance $\sigma_i^2 = \left(\sum_{j=1}^{k_i} w_{ij}\right)^{-1}$. These variances can in turn be used to obtain the weighted grand mean of effect sizes across groups

$$\bar{y} = \frac{\sum_{i=1}^p w_i \bar{y}_i}{\sum_{i=1}^p w_i} \quad (11)$$

$$w_i = 1/\sigma_i^2. \quad (12)$$

With $\bar{y}_i$ and $\bar{y}$ in hand, it is possible to test the hypothesis $H_0 : \mu_1 - \mu_2 = 0$ by taking the test statistic

$$Q = \frac{(\bar{y}_1 - \bar{y}_2)^2}{w_1^{-1} + w_2^{-1}}, \quad (13)$$

or in the case of testing $p > 2$ group means

$$Q = \sum_{i=1}^p w_i (\bar{y}_i - \bar{y})^2, \quad (14)$$

and referring it to a χ^2 distribution with $(p - 1)$ degrees of freedom. If the significance level of the test is α and c_α denotes the $(1 - \alpha)$ quantile of the central χ^2 distribution, then H_0

can be rejected if $Q > c_a$.

Knapp and Hartung Method

The Knapp and Hartung (KH) method builds on the omnibus Q -test by applying an adjustment that accounts for the uncertainty in estimating the between-study variance. This adjustment has been found to ameliorate the Type I error rate of the Q -test in mixed-effects models (Rubio-Aparicio et al., 2020), and results in the F statistic given by

$$F = \frac{Q/(p-1)}{Q_w/(k-p)}, \quad (15)$$

where Q_w is defined as

$$Q_w = \sum_{i=1}^p \sum_{j=1}^{k_i} w_{ij} (y_{ij} - \bar{y}_i)^2. \quad (16)$$

The F statistic can be used to test the hypothesis $H_0 : \mu_1 = \dots = \mu_p$ by referring it to an F distribution with $(p-1)$ and $(k-p)$ degrees of freedom. If the significance level of the test is α and c_α denotes the $(1-\alpha)$ quantile of the central F distribution, then H_0 can be rejected if $F > c_\alpha$.

There are two important details to observe here. First, the expressions used for the Q - and F -tests are in effect the same as those used in a t -test and ANOVA, but use *weighted* forms of the group and grand means. Thus, assuming unequal τ_i^2 is equivalent to assuming unequal variances in a t -test or ANOVA. Second, and relatedly, the weights w_{ij}^* and w_{ij}^{**} are fundamentally tied to the between-study variances. It follows that if groups in a moderator truly have different τ_i^2 , then the weights using a pooled value will be incorrect, and so too the test statistic. This holds implications for both the Type I error and power of moderator tests, both of which are investigated in the section [Simulation Studies](#).

Motivating Example

To highlight the impact of the between-study variance on the resulting test statistic, we calculated the Q and F values for a hypothetical set of $k = 40$ studies where each study belonged to one of $p = 2$ groups. Following Williams et al. (2021), we assumed all σ_{ij}^2 and τ_i^2 to be known, with the σ_{ij}^2 fixed at 0.04 (roughly the variance of a standardized mean difference with $n = 50$). The heterogeneity parameter for the first group τ_1^2 was held constant at 0.05, but varied for the second group from 0.05 to 0.25 in steps of 0.01. We further varied the proportion of total studies belonging to each group such that the proportion of the 40 studies in the first group ranged from 0.5 to 0.9 in 0.1 increments. Lastly, we calculated the Q and F statistics,³ using weights that used either the heterogeneous τ_i^2 values or a pooled τ^2 value. The denominator for the F statistic was set to $Q_w = 0.67(k - p)$, which corresponds to a “moderate” degree of within-group heterogeneity (Hedges & Pigott, 2005). Pooled values of τ^2 were obtained by taking the arithmetic mean of the τ_i^2 across all studies. The main idea here is that if the between-study variances influence the value of the test statistics, then the values for the Q and F statistics should differ depending on whether equal or unequal τ_i^2 are assumed.

Figure 1 includes the results for this example. The y -axes denote the ratio of the test statistic using the pooled value of τ^2 (Q_{Equal} and F_{Equal}) to the test statistic using the heterogeneous values for τ_i^2 (Q_{Unequal} and F_{Unequal}). Thus, a ratio of 1 indicates that the test statistics are equal, and positive values indicate that the statistic using equal τ_i^2 is greater than the one using unequal τ_i^2 . The x -axes display the between-study variance for the second group, τ_2^2 , and colors differentiate the number of studies in the second group. In line with the classical literature on the t -test and ANOVA, the resulting test statistics are invariant to heterogeneous between-study variances under balanced sample sizes. They are also invariant under imbalanced sample sizes when the between-study variances are equal and known, as in this example. The test statistics differed, however, when both the

³ To provide analytical results in this example we used the equations given in the section [Statistical Power](#).

Figure 1

Motivating example of how the Q and F statistics vary according to whether equal or unequal between-study variances are assumed (see main text for details). As the imbalances grow between groups in between-study variances and sample sizes, so too does the ratio of the test statistics.

between-study variances and sample sizes were heterogeneous across the two groups. In our setup, greater disparities in either factor led to the statistics using equal between-study variances being larger than the ones using heterogeneous between-study variances. In the worst case with $k_2 = 4$ and $\tau_2^2 = 0.25$, the Q_{Equal} and F_{Equal} statistics were 1.42 times larger than their unequal counterparts.

Despite the simplicity of this example, it is clear that when heterogeneous variances and sample sizes go unaccounted for when testing moderators, the resulting test statistic can lead to overconfidence in the strength of the effect. Parallel to how Welch's versions of the t -test and ANOVA have been proposed deal with such scenarios, we too propose that unequal (between-study) variances should be assumed by default when testing categorical moderators. The motivation being that assuming unequal between-study variances will

result in the appropriate test statistics regardless of whether they are indeed equal or not, but assuming them to be equal will only be appropriate if the population between-study variances are truly equal. In the next section, we study the behavior of the MELSM when testing moderators under various realistic scenarios to ensure its suitability as a default method.

Simulation Studies

Because we are proposing that unequal between-study variances should be assumed by default when testing categorical moderators, it is of interest to examine how the mixed-effects location-scale model (MELSM) performs relative to the standard mixed-effects model (MEM) in situations where the variances are equal and when they are not. Accordingly, we conducted two Monte Carlo simulation studies comparing the performance of the two models in terms of Type I error and statistical power. Specifically, we compared these two quantities with respect to the Q statistic defined in (13) and the F statistic defined in (15). Unlike the preceding example, the use of Monte Carlo simulations allow us to examine the empirical performance of the Q - and F -tests in a more realistic setting because in practice the between- and within-study variances are not fixed and known in advance. We chose standardized mean differences as the effect size measure and a two-level categorical variable as a moderator.

Method

Following standard procedure in psychology, we assumed that each simulated study compared an experimental group to a control group on an arbitrary quantitative outcome (e.g., López-López et al., 2014; Sánchez-Meca & Marín-Martínez, 2008; Viechtbauer et al., 2015). Under a mixed-effects model the population standardized mean difference (SMD) for each study was defined as

$$\theta_{ij} = \frac{E_{ij} - C_{ij}}{S_{ij}} \quad (17)$$

where E_{ij} and C_{ij} were the population means for the experimental and control groups, respectively, and $S_{ij} = 1$ was the pooled within-study standard deviation. The population distribution for the θ_{ij} was taken to be $\mathcal{N}(\mu_i, \tau_i^2)$, where the subscript denotes the i th group. The specific values for μ_i and τ_i^2 were systematically varied in each simulation study and are given below. For the experimental and control groups, we defined the population distributions as $\mathcal{N}(\theta_{ij}, 1)$ and $\mathcal{N}(0, 1)$, respectively. Using these definitions, a study was simulated by first drawing a θ_{ij} , and then drawing $n_{E_{ij}} = n_{C_{ij}}$ variates from the latter two distributions and estimating the SMD in the usual way (Hedges & Olkin, 1985).

To be clear, the sample sizes $n_{E_{ij}}$ and $n_{C_{ij}}$ respectively refer to the sample size of the experimental and control groups *within* studies, and are not to be confused with the number of studies in each subgroup of a categorical moderator variable, k_i . The within-study sample sizes were set to $n_{E_{ij}} = n_{C_{ij}} = g$, where g is a gamma variate divided by two and rounded to the nearest integer. Gamma variates were drawn from a gamma distribution with shape equal to 0.65 and scale equal to 165, truncated to values greater than or equal to thirty. Similar procedures have been used to generate within-study sample sizes consistent with those commonly observed in social-behavioral studies (Brannick et al., 2019; Dahlke & Wiernik, 2020).

Using the above procedure, we generated a collection of $k = k_1 + k_2$ studies for each simulation trial and conducted moderator tests using both the MEM and the MELSM. All models were fit using REML estimation. Both the study sample size ($k_1 = 20$) and between-study variance ($\tau_1^2 = 0.2$) were held constant for the first group, but varied for the second group. The value for τ_2^2 varied from 0.05 to 0.4 in steps of 0.05, and values of 5, 10, 20, 30, 40 were examined for k_2 . These values were chosen because they represent empirical

values of τ^2 observed in psychological meta-analyses (van Erp et al., 2017). Additionally, the classical literature examining tests with unequal group variances focuses not on the raw values of variances and sample sizes, but on their ratios (Bartlett, 1936; Murphy, 1967; Scheffé, 1959; Welch, 1938), and in line with this convention, the study sample sizes and variances we considered approximately correspond to the ratios: a) $k_1/k_2 = (\frac{1}{4}, \frac{1}{2}, 1, 1.5, 2)$ and b) $\tau_1^2/\tau_2^2 = (\frac{1}{4}, \dots, 2)$. Using ratios rather than raw values further lends intuition for how the group variances and number of studies affect Type I error and power in situations outside of those considered here.

Our expectations for these simulation studies were as follows. Like with traditional tests of mean differences, we expected that for each moderator test, the Type I error rate and power rate would be approximately equivalent between the MEM and the MELSM when either the sample sizes were balanced and/or the between-study variances were equal. When the group sizes were imbalanced, we expected two general trends to emerge as a result of assuming equal between-study variances: when the group with a larger sample size simultaneously had a smaller between-study variance, then assuming equal variances would lead to a higher Type I error rate, and conversely, when the larger group had a larger between-study variance then assuming equal variances would result in a conservative Type I error rate. These scenarios were also expected to yield invalid power rates as the rejection rate would be too high or too low. We expected, however, that assuming unequal between-study variances in these scenarios would result in the nominal error rate and valid statistical power (i.e., no increased/decreased power at the cost of an inflated/conservative rejection rate). In sum, we anticipated there to be potentially serious costs in terms of Type I error and power when using the MEM, but not the MELSM.

For each simulation study, the total number of conditions was 2 (Q - or F -test) $\times 2$ (MEM or MELSM) $\times 5$ (k_2) $\times 8$ (τ_2^2) = 160, each of which was simulated 10,000 times. All simulations were conducted in the R environment (R Core Team, 2021) using the **metafor** package to fit the meta-regression models (Viechtbauer, 2010). All code used in this article

is openly available on the Open Science Framework, at <https://osf.io/gS6uw/>.

Type I Error

To examine how well the MELSM controls Type I error of the Q - and F - tests relative to the MEM, we set $\beta_0 = \beta_1 = 0$, or in terms of group means, we set $\mu_1 - \mu_2 = 0$. For both moderator tests, the α level was set to 0.05. Further, we considered “stringent” and “liberal” criteria for assessing adequate control of Type I error that respectively corresponded to the intervals $[0.046, 0.054]$ and $[0.043, 0.057]$. These intervals were determined by computing the critical values of a two-tailed test that a population proportion is equal to 0.05 under significance levels of 0.05 and 0.001, and with a sample size of 10,000. In other words, for each condition, values outside of the intervals would result in rejecting the null hypothesis that the Type I error rate is equal to 0.05, at the given significance level.

The results are presented in Figure 2. Type I error rates for both the Q - and F -tests are plotted as a function of the between-study heterogeneity for the second group (τ_2^2), whether the variances were treated as equal (MEM) or unequal (MELSM), and the number of studies in the second group (k_2). The dashed white line denotes the nominal error rate $\alpha = 0.05$, the light gray area region represents the region between $[0.043, 0.057]$, and the dark gray shaded area represents the region between $[0.046, 0.054]$. The color of the solid lines distinguishes between the number of studies included in the second group.

The rejection rates were in line with our expectations and largely mirror those in the classical t -test and ANOVA literature. When the between-study variances were equal ($\tau_1^2 = \tau_2^2 = 0.2$), the Type I error was roughly the same between the MEM and MELSM, regardless of study sample size. This was also the case across all values for τ_2^2 when the study sample sizes were equal ($k_1 = k_2 = 20$). The only time the MEM clearly controlled the Type I error rate better than the MELSM was when $k_2 = 5$ and the between-study variances were equal. In other words, in all but one case, the MELSM was on par with the MEM in adequately controlling the Type I error even when the assumption of equal

Figure 2

Type I error rates for the Q-test (Panel A) and F-test (Panel B) under the MEM (equal τ^2 assumed) and the MELSM (unequal τ_i^2 assumed). The white dashed line denotes 0.05, the light gray area spans [0.043, 0.057], and the dark gray area spans [0.046, 0.054]. For both moderator tests, the MEM only resulted in a well-controlled error rate when either the τ_i^2 were truly equal or when the sample sizes were balanced. When there were unequal τ_i^2 and unbalanced sample sizes, but the MEM was used, the error rate was inflated, or excessively conservative. Using the MELSM, however, mostly resulted in a well-controlled Type I error rate regardless of whether the variances and sample sizes were equal (except when $k_2 = 5$). Although both moderator tests had better overall Type I error under the MELSM, only the F-test had error rates consistently inside of the [0.046, 0.054] bounds.

between-study variances was true in the population. The similar rejection rates between the MEM and the MELSM under balanced sample sizes and equal τ_i^2 lends strong support to our view that almost nothing is lost by assuming unequal between-study variances by default. On the other hand, there may be serious ramifications by assuming equal between-study variances when they are unequal and are coupled with imbalanced sample sizes.

For both the Q- and F-test, when the between-study variances and sample sizes

simultaneously differed between groups but equal τ_i^2 were assumed, the Type I error rate deviated from the nominal rate. In particular, when the larger group (e.g., $k_1 > k_2$) also had a smaller between-study variance (e.g., $\tau_1^2 < \tau_2^2$), the error rate *increased* above the nominal level. In contrast, when the larger group had a larger between-study variance (e.g., $\tau_1^2 > \tau_2^2$), the error rate *decreased* below 0.05. In the most extreme cases, the MEM had a Type I error rate of 0.02 ($k_2 = 5, \tau_2^2 = 0.05$) and 0.12 ($k_2 = 5, \tau_2^2 = 0.4$); well outside the bounds of the wider [0.043, 0.057] interval.

When unequal between-study variances with the MELSM were assumed, the Type I error rate was mostly well-maintained across all conditions and regardless of moderator test, relative to the MEM, which assumed equal between-study variances. When comparing against the liberal and stringent intervals, however, the Q -test was expectedly worse than the F -test in controlling the Type I error. For the Q -test, the error rates in the majority of the conditions fell outside of the wider [0.043, 0.057] bounds. Even so, the error rates were still predominantly closer to nominal than those of the Q -test using the MEM. Meanwhile, the F -test achieved a much better Type I error rate, with most conditions achieving an error rate within the wide interval of [0.043, 0.057] and many falling between the stricter interval of [0.046, 0.054]. As mentioned above, the only cases where the moderators tests assuming unequal τ_i^2 clearly performed worse than their equal between-study variance counterparts occurred when $k_2 = 5$, though it is worth pointing out that all methods did relatively poorly in such conditions. In all other conditions, the Type I error rates under the MELSM were as good or better than those obtained using the MEM, with the F -test achieving the nominal error rate. These findings fortify the idea that the resulting inferences of moderator tests can be seriously compromised when equal between-study variances are assumed, but inferential integrity can be preserved (at least in part) by assuming unequal τ_i^2 .

Figure 3

Empirical power rates for the Q -test (Panel A) and F -test (Panel B) for $\beta_1 = 0.5$ and $\tau_1^2 = 0.2$. Power rates for conditions that exceeded acceptable Type I error bounds are excluded. Solid lines denote the MEM with equal τ_i^2 and dashed lines denote the MELSM with unequal τ_i^2 . The power rates for both models were similar under balanced sample sizes and equal between-study variances, but there were differences otherwise. Particularly, the moderators tests under the MEM produced higher or lower power relative to the MELSM, but in conditions where the tests also had Type I error rates that were above or below the nominal 0.05. In contrast, the moderator tests under the MELSM produced valid power rates across more conditions insofar as they did not incur any costs in terms of Type I error.

Statistical Power

To assess the empirical statistical power rates of the MEM and the MELSM with respect to the two moderator tests, we followed nearly the exact same procedure as in the previous simulation study, but set $\beta_1 = 0.5$, or $\mu_2 - \mu_1 = 0.5^4$. Such a moderator effect can be regarded as a representing a “medium” effect size for a standardized mean difference

⁴ We additionally conducted simulations for $\beta_1 = 0.2$ and $\beta_1 = 0.8$, but the results were qualitatively the same as those presented below. Thus, they are omitted here but can be found in the Appendix (Tables A2 and A3).

(Cohen, 1992).

Figure 3 displays the results for $\tau_2^2 = (0.05, 0.2, 0.4)$, corresponding to ratios of $\tau_1^2/\tau_2^2 = (\frac{1}{4}, 1, 2)$. Power rates are plotted as function of studies in the second group (k_2), heterogeneity in the second group (τ_2^2), whether the variances were treated as equal, and the moderator test. The dashed horizontal gray line denotes a power rate of 0.8.

Additionally, we excluded conditions where the Type I error rate for the MEM was greater than the error rate for the MELSM because they could misleadingly be interpreted as achieving higher power. That is, we omitted conditions where the MEM produced higher power rates than the MELSM, but only achieved higher power because they also had a higher Type I error rate. The complete set of results are included in the Appendix (Table A1). Similar to Type I error, the power rates between the MEM and the MELSM were nearly identical with balanced sample sizes and equal between-study variances, regardless of the moderator test. The lack of differences between the MEM and the MELSM in this scenario again supports our claim that there are essentially no drawbacks to assuming unequal τ_i^2 by default.

With respect to situations where the sample sizes were unbalanced and the between-study variances were simultaneously unequal, power rates diverged between the MEM and the MELSM. Most notably, in conditions where the MELSM produced as good or better Type I error rates than the MEM, the MELSM also resulted in as good or better power rates, regardless of moderator test. For example, for the F -test, when $k_2 < 20$ and $\tau_2^2 = 0.4$, the MEM produced an inflated Type I error rate and thus did not produce valid power rates. In contrast, when $k_2 \geq 20$, the MEM had valid power rates, but they were either equal to or lower than the power rates for the MELSM. Similar conclusions can be drawn when considering different values for k_2 and τ_2^2 . Together, these findings add support to our argument there is little lost by assuming heterogeneous between-study variances when testing moderator effects in regards to statistical power, but there may be a lot to gain.

Analytical Solution to Power

Figure 4

Analytical power for a moderator effect of 0.5 for the Q -test and F -test under fixed assumptions (see main text for details) for the MELSM. The lines corresponding to F_1 , F_2 , and F_3 denote the F -test with small, moderate, and large amounts of within group heterogeneity, respectively. The Q -test and F_3 -test most closely matched the observed power rates.

Beyond simulation, the power rates of these tests can be assessed analytically, albeit at the price of making a few more assumptions. Because simulation-based power analyses can be computationally demanding and time consuming, an analytic solution can be particularly useful if researchers wish to conduct a power analysis prior to carrying out a moderator test. To calculate statistical power for the omnibus Q -test, values must be assumed for k and p , and for each μ_i , τ_i^2 , σ_{ij}^2 , and k_i . Once these values have been chosen, power can be calculated as (Hedges & Pigott, 2005)

$$1 - \chi^2(c_\alpha | p - 1; \lambda), \quad (18)$$

where $\chi^2(\cdot)$ denotes the cumulative distribution function (CDF) of a χ^2 distribution with $(p - 1)$ degrees of freedom and non-centrality parameter λ . Recall that in our formulation,

p refers to the number of groups in the categorical moderator. The term c_α is the critical value for the central χ^2 distribution used for the Q -test in (13). The noncentrality parameter λ is given by

$$\lambda = \sum_{i=1}^p w_i (\mu_i - \mu)^2 \quad (19)$$

$$\mu = \frac{\sum_{i=1}^p w_i \mu_i}{\sum_{i=1}^p w_i}. \quad (20)$$

The expression to calculate power for the Knapp and Hartung method is similar to that of the Q -test, and can be computed as (Hartung et al., 2001)

$$1 - F(c_\alpha | p - 1, k - p; \lambda), \quad (21)$$

where $F(\cdot|\cdot)$ denotes the CDF of an F distribution with $(p - 1)$ and $(k - p)$ degrees of freedom, and non-centrality parameter λ . The term c_α is the critical value for the central F distribution used for the F -test in (15). The noncentrality parameter for the F -test is also similar to that of the Q -test and is defined as

$$\lambda = \frac{\sum_{i=1}^p w_i (\mu_i - \mu)^2}{\sum_{i=1}^p \sum_{j=1}^{k_i} w_{ij} (\theta_{ij} - \mu_i)^2}, \quad (22)$$

where θ_{ij} is the population effect size for the i th study in the j th group. Here, along with the values that must be assumed for the Q -test, values must also be inserted for θ_{ij} .

Because choosing values for the θ_{ij} can be particularly difficult, it has been suggested that values of $0.33(k - p)$, $0.67(k - p)$, and $(k - p)$ can be used for the denominator in (22), which correspond to small, moderate, and large degrees of within-group heterogeneity (Hedges & Pigott, 2001).

To aid researchers in sample size planning, we provide the analytical solutions to

power for a moderator effect of 0.5 under the same conditions as the power simulation, but varied k_2 in steps of 1 from 5 to 40. For both tests, we set the within-study variances σ_{ij}^2 to 0.2, which is approximately the sampling variance of a standardized mean difference with $n = 50$. For the F -test, we computed power for various degrees of within-group heterogeneity. In particular, F_1 , F_2 , and F_3 coincide with the F -test under a small, moderate, and large degree of within-group heterogeneity, respectively.

The power rates for the MELSM are plotted in Figure 4. Here, power is plotted as a function of studies in the second group (k_2), the between-study heterogeneity in the second group (τ_2^2), and moderator test. The pattern of power rates between the MEM and the MELSM was similar to the one above, and so power rates for the former are omitted. Regarding the between-study variance, τ_2^2 , power is higher across all conditions with smaller values. A similar statement can be made for the degree of within-group heterogeneity with respect to the F -test. Perhaps intuitively, the F -tests with small and moderate within-group heterogeneity (i.e., F_1 and F_2) were more powerful than the Q - and F_3 -tests across all conditions.

Mean Squared Error and Standard Error

We conducted two supplementary comparisons between the MEM and the MELSM. In the above empirical power simulation, we further collected: 1) the mean squared error (MSE) for the moderator effect (β_1) in each condition and 2) the average standard error of the moderator effect in each condition. This was done to investigate discrepancies in the quality of the estimates for the effect itself and the respective standard error between the MEM and the MELSM.

The results can be view in Figure 5. Panel A presents the MSE in each condition and Panel B displays the respective (average) standard errors. The gray lines denote τ_2^2 values spanning from 0.05 to 0.4 in increments of 0.05 (dark to light). Solid and dashed lines differentiate between the MEM and the MELSM, respectively. Although the standard

Figure 5

Panel A presents the mean-squared error for the moderator effect (β_1) under the MEM and the MELSM, and Panel B shows the respective average standard error. The gray lines denote τ_2^2 values spanning from 0.05 to 0.4 in increments of 0.05 (dark to light). As can be clearly seen by the overlap of lines, there is no difference in estimation accuracy by assuming equal or unequal between-study variances. There are discrepancies in standard errors, however. Under imbalanced group sizes and unequal τ_i^2 , the estimated standard errors obtained with the MEM were either too large or too little, depending on the condition. Consequently, the standard errors were responsible for too high and too low Type I error rates and power rates.

errors differ depending on whether the Q -test or F -test⁵ is used for the moderator effect, their results were so similar that they are only shown for the former. The MSE was practically identical regardless of which model was used to estimate the moderator effect. Point estimates of the moderator effect were evidently unaffected by whether equal or unequal τ_i^2 were assumed, whereas this choice affected the average standard error. When the number of studies was unequal and so were the between-study variances, the standard

⁵ Technically, the standard errors are for the corresponding z -test (Hedges & Pigott, 2005) and KH method (Knapp & Hartung, 2003), but this does not affect our arguments here given their equivalence with two groups (see Section [Tests for Categorical moderators](#))

errors were either too large or too little, depending on the condition. Although the between-study variances indeed play a role in estimating the moderator effect, it seems that it is their effect on the standard errors that is responsible for the discrepancies between the MEM and the MELSM observed in the above simulations studies.

Extension to 3 Groups

Naturally, one may wonder if the simulation study results above generalize to cases where there are more than two levels (groups) in a categorical moderator. To probe this idea, we included a third level in the categorical moderator and repeated the simulation studies examining Type I error rate and statistical power. The results were qualitatively the same as those with two groups. In cases where the population had equal τ_i^2 and had equal study sample sizes, the Type I error rate was relatively well controlled regardless of whether a MEM or MELSM was used. Again, the exceptions to this were cases where the groups had few studies. When there were imbalances in either τ_i^2 or study sample sizes, then using a standard MEM lead to inflated or conservative Type I error rates. Further, the power rates for the MELSM were as good or better than the MEM in all conditions where the latter had a valid (i.e., non-inflated) Type I error rate. The complete results for these simulations can be found in the appendix. Note that although not directly studied here, our simulation results for three groups qualitatively extend to testing multiple categorical moderators simultaneously (i.e., > 3 groups). This is because tests of moderator effects do not directly depend on the number of moderator variables, but the number of studies k and number of groups p (see Equations 13 – 16). The number of groups p will of course increase with higher number of moderators, thus reducing power if p becomes large relative to k .

Illustrative Example

In this section we use an empirical dataset to demonstrate the MELSM and contrast it against the standard MEM in an applied setting. The dataset was first used to synthesize the findings of $k = 46$ studies investigating the relationship between depression and specificity of future thinking (Gamble et al., 2019), and is available in the

psymetadata R package (Rodriguez & Williams, 2022). That is, whether higher levels of depression are correlated with vagueness in future thoughts. The original dataset contained 89 effect sizes (Pearson’s or biserial correlations), but because some of these effect sizes were collected from the same study, they did not conform with the standard assumption of independent effect sizes in meta-regression. This issue was resolved by averaging all effect sizes within a study⁶; a procedure that is routinely applied in meta-analyses to alleviate non-independence of effect sizes (Tipton et al., 2019).

One of the variables coded as a (trichotomous) moderator was the type of cue each study used to elicit participants’ thoughts about a future event: a single word (e.g., “laughing”), event cues (e.g., “New Years Eve”), or open (e.g., “I can imagine that, shortly, I...”). We thus fitted a MEM with cue type as a moderator for the location in addition to a MELSM with cue type as a moderator for both location and scale. The top of Figure 6 shows a forest plot (Lewis & Clarke, 2001) of the individual effect sizes, and the bottom displays the resulting moderator effect estimates from the MELSM.

Table 1 further shows, for the standard MEM and the MELSM, the estimates of the regression coefficients, their standard errors, and the resulting p -values. The KH method (equivalent to the F -test) was used to test the location coefficients in both models because it better maintains the nominal Type I error rate and valid power rates relative to the z -test (equivalent to the Q -test). Because the type of cue had three categories, it was coded as two dummy variables. Accordingly, the intercept β_0 captures the effect for the event cue type, β_1 captures the difference between the event cue type and the open cue type, and β_2 captures the difference between the event cue type and the single word cue type. As in the simulation studies, the location estimates were similar between the two models, but the MELSM had smaller standard errors for all three coefficients.

Consequently, all p -values for the moderator effects were smaller under the MELSM, and of

⁶ Although a three-level MELSM would accommodate dependent effect sizes, such models are not yet available in standard meta-analytic software.

Figure 6

Forest plot of the effect sizes in the Gamble et al. (2019) and the results from the MELSM with cue type as a moderator in the location and scale.

particular note, β_2 was not statistically significant for the MEM at the $\alpha = 0.05$ level whereas it was for the MELSM. The discrepancy in p -values between the two models

reinforces the simulation study findings that inferences drawn about the location coefficients may differ depending on whether a MEM or a MELSM is fit to the data. The differences in standard errors was due to the differences in the heterogeneity estimates between the models. The (pooled) estimate for heterogeneity under the MEM was $\hat{\tau}^2 = 0.06$, but varied across groups for the MELSM. For the event, open, and single cue types, the heterogeneity estimates were $\hat{\tau}_E^2 = 0.006$, $\hat{\tau}_O^2 = 0.09$, and $\hat{\tau}_S^2 = 0.067$, respectively. That is, under the MELSM, the groups were estimated to have 0.1, 1.5, and 1.1 as much heterogeneity than they would have been estimated to have under the MEM.

Table 1
Regression Estimates for Example Data

	MEM			MELSM		
	Est.	SE	<i>p</i>	Est.	SE	<i>p</i>
β_0	-0.05	0.07	0.52	-0.04	0.04	0.32
β_1	-0.05	0.09	0.57	-0.06	0.08	0.43
β_2	-0.19	0.11	0.09	-0.20	0.09	0.03

Note. Estimates of location coefficients under the standard mixed-effects model and the MELSM.

In applied settings, a commonly estimated metric in mixed-effects meta-regressions is pseudo- R^2 — a measure of how much heterogeneity is explained by the moderators (López-López et al., 2014; Raudenbush, 1994). Under the standard MEM, this metric is calculated as

$$R^2 = 1 - \frac{\hat{\tau}_{\text{full}}^2}{\hat{\tau}_{\text{null}}^2}, \quad (23)$$

where $\hat{\tau}_{\text{null}}^2$ corresponds to the heterogeneity parameter obtained via a null model (i.e., one without predictors) and $\hat{\tau}_{\text{full}}^2$ corresponds to the heterogeneity parameter obtained via the model containing the moderators. Because the MELSM instead yields τ_i^2 , a separate R^2 is can be calculated for each group

$$R_i^2 = 1 - \frac{\hat{\tau}_{i,\text{full}}^2}{\hat{\tau}_{i,\text{null}}^2} \quad (24)$$

where $\hat{\tau}_{i,\text{null}}^2$ is the heterogeneity parameter estimate for the i th group under a null location submodel. That is, the $\hat{\tau}_{i,\text{null}}^2$ estimates are obtained by fitting a MELSM with the categorical moderator in the scale component, but not the location. Further, $\hat{\tau}_{i,\text{full}}^2$ is obtained by fitting the MELSM with the moderator in both the location and scale components of the model. Having separate R_i^2 's enables researchers to assess the differential explanatory power of a moderator.

In this example, pseudo- R^2 was estimated to be 0 under the MEM, indicating that cue type did not have any explanatory power for the heterogeneity in effect sizes. Under the MELSM, the pseudo- R_i^2 metrics were found to be $R_E^2 = 0.29$, $R_O^2 = 0$, $R_S^2 = 0.22$ for the event, open, and single cue types, respectively. In other words, the results under the MELSM indicate that the event and single cue types explain more than 20% of the heterogeneity in their respective groups. Thus, one might conclude that the heterogeneity in the correlation between depression and the specificity of future thought is may be well-explained by event and single cue types, but not in the open cue type. We refer interested readers to Williams et al. (2021) for more information on computing and interpreting pseudo- R_i^2 in MELSMs.

Lastly, due to the MEM being nested within the MELSM, we compared the two models using a likelihood ratio test. Recall that the MEM can be recast as MELSM with only an intercept included in the scale component. Because the models did not differ in terms of their location coefficients, they did not need to be refit using ML instead of REML. The likelihood ratio test indicated that the MELSM provided an improved fit over the MEM $\chi^2(2) = 7.8$, $p = .02$.

Discussion

In this work we set out to establish that in testing moderator effects with a mixed-effects meta-regression, researchers should assume unequal between-study variances by default. We illustrated how meta-analytic moderator tests are analogous to the t -test and ANOVA, and thus, assuming unequal variances in moderator tests is no different than assuming unequal variances in classical test of mean differences — the latter of which is already common practice in many disciplines.

As evidenced by two simulation studies, there are few costs in terms of Type I error (or statistical power) by assuming unequal between-study variances and potentially serious drawbacks to assuming them equal. When the population between-study variances differ between groups in a categorical moderator, statistical tests of moderator effects using the standard mixed-effects model can result in either grossly inflated or overly conservative Type I error rates. In turn, these error rates result in power rates that are too low or misleadingly high. Meanwhile, when moderator tests are carried out using the mixed-effects location-scale model (MELSM), the Type I error is well-controlled regardless of equal variances or balanced sample sizes (except when there are few studies), and maintains valid power rates. These results substantiate the notion that researchers should assume unequal between-study variances as a default strategy.

As demonstrated in the Illustrative Example, MELSMs can be used to differentially assess the explanatory power of a categorical moderator. This is possible because the pseudo- R_i^2 statistic is based on the group-specific heterogeneity estimates, τ_i^2 . Consistent with this idea, other statistics that are traditionally based on τ^2 can be obtained at the group level by instead using τ_i^2 . For instance, Williams et al. (2021) use a MELSM to compute country-specific I^2 (Higgins & Thompson, 2002) values in a meta-analysis.

Future Directions and Limitations

The present work focused on categorical moderators due to their popularity in meta-analysis. However, a benefit of meta-regression techniques is the ability to

accommodate both categorical and continuous predictors. Accordingly, it is possible to include a continuous variable as a moderator for the between-study variance in a MELSM. Indeed, it has been suggested that effect size heterogeneity may be inversely proportional to study sample size (Bowater & Escarela, 2013). Given that there is currently a dearth of literature on continuous moderators of between-study heterogeneity, investigating the extent to which a MELSM can be considered by default with a continuous moderator is an interesting avenue for future exploration.

We demonstrated that the MELSM generally does well in controlling the Type I error rate, but we did so based on the implicit assumption that the included moderator was fixed in advance (e.g., through preregistration of hypotheses). That is, the observed error rates hold only when the moderator being tested is chosen a priori and not as a result of a data-driven process (see e.g., Berk et al., 2013). When the final meta-analytic model is chosen as a result of which moderators are (non-)significant, a potential consequence is an inflated Type I error rate. Although model selection techniques are relatively common in practice (Tipton et al., 2019), their use may contribute to recent concerns on the reproducibility of meta-analyses (Lakens et al., 2016).

In recent years, meta-analytic recommendations have expressed a preference towards methods that reflect complex dependencies among effect sizes (e.g., multiple effect sizes nested within a study). To deal with such complexities in the variance structure, researchers can employ multivariate meta-analysis (Kalaian & Raudenbush, 1996; Raudenbush et al., 1988), robust variance estimation techniques (Hedges et al., 2010; Tipton, 2015), or three-level random-effects models (Konstantopoulos, 2011; Van den Noortgate et al., 2013). The MELSM described in this work is a two-level model and allows only a single effect size per study. It is thus limited in reflecting the dependencies amongst effect sizes. However, we note that averaging effect sizes within studies or randomly selecting a single effect size per study remain the dominant forms of eliminating within-study dependencies (Tipton et al., 2019). In these cases, the MELSM can still be

fruitfully applied. Future research should aim to understand whether the arguments made throughout this paper similarly apply to three-level MELSMs.

Finally, it has been recommended that at least five studies should be included per group when testing a categorical moderator in order to obtain trustworthy estimates of τ_i^2 (Borenstein et al., 2009). Others have recommended that when the total number of studies is less than 20, the $\hat{\tau}_i^2$ should be pooled in order to more accurately estimate the between-study heterogeneity (Rubio-Aparicio et al., 2020). We agree with this suggestion when the focal interest lies in quantifying the between-study heterogeneity, but as suggested by our simulation studies, when the focal interest is in testing the effect of a moderator, assuming equal τ_i^2 when they are truly unequal and when there is an imbalance in study sample size may still lead to inadequate error rates.

Conclusion

Presently, mixed-effects location-scale models can be easily fit in the R statistical computing environment with the R packages **metafor** (Viechtbauer, 2021) and **blsmeta** (Williams et al., 2021). Our hope is that this work provides an impetus for the wide-scale adoption of location-scale models in meta-analytic software and by applied researchers for testing categorical moderators.

References

- Bartlett, M. S. (1936). The information available in small samples. *Mathematical Proceedings of the Cambridge Philosophical Society*, 32(4), 560–566.
<https://doi.org/10.1017/S0305004100019290>
- Berk, R., Brown, L., Buja, A., Zhang, K., & Zhao, L. (2013). Valid post-selection inference. *Annals of Statistics*, 41(2), 802–837. <https://doi.org/10.1214/12-AOS1077>
- Borenstein, M., Hedges, L. V., Higgins, J. P., & Rothstein, H. R. (2009). *Introduction to meta-analysis*. John Wiley & Sons.
- Borenstein, M., & Higgins, J. P. T. (2013). Meta-Analysis and Subgroups. *Prevention Science*, 14(2), 134–143. <https://doi.org/10.1007/s11121-013-0377-7>
- Bowater, R. J., & Escarela, G. (2013). Heterogeneity and study size in random-effects meta-analysis. *Journal of Applied Statistics*, 40(1), 2–16.
<https://doi.org/10.1080/02664763.2012.700448>
- Brannick, M. T., Potter, S. M., Benitez, B., & Morris, S. B. (2019). Bias and Precision of Alternate Estimators in Meta-Analysis: Benefits of Blending Schmidt-Hunter and Hedges Approaches. *Organizational Research Methods*, 22(2), 490–514.
<https://doi.org/10.1177/1094428117741966>
- Bürkner, P.-C. (2018). Advanced Bayesian Multilevel Modeling with the R Package brms. *The R Journal*, 10(1), 395. <https://doi.org/10.32614/RJ-2018-017>
- Cohen, J. (1992). A power primer. *Psychological Bulletin*, 112(1), 155–159.
<https://doi.org/http://dx.doi.org/10.1037/0033-2909.112.1.155>
- Dahlke, J. A., & Wiernik, B. M. (2020). Not Restricted to Selection Research: Accounting for Indirect Range Restriction in Organizational Research. *Organizational Research Methods*, 23(4), 717–749. <https://doi.org/10.1177/1094428119859398>
- DeGeest, D. S., & Schmidt, F. L. (2010). The impact of research synthesis methods on industrial–organizational psychology: The road from pessimism to optimism about cumulative knowledge. *Research Synthesis Methods*, 1(3-4), 185–197.

<https://doi.org/10.1002/jrsm.22>

_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/jrsm.22>

- Delacre, M., Lakens, D., & Leys, C. (2017). Why Psychologists Should by Default Use Welch's t -test Instead of Student's t -test. *International Review of Social Psychology*, 30(1), 92–101. <https://doi.org/10.5334/irsp.82>
- DerSimonian, R., & Laird, N. (1986). Meta-analysis in clinical trials. *Controlled Clinical Trials*.
- Gamble, B., Moreau, D., Tippet, L. J., & Addis, D. R. (2019). Specificity of Future Thinking in Depression: A Meta-Analysis. *Perspectives on Psychological Science*, 14(5), 816–834. <https://doi.org/10.1177/1745691619851784>
- Hagger, M. S., Chatzisarantis, N. L. D., Alberts, H., Anggono, C. O., Batailler, C., Birt, A. R., Brand, R., Brandt, M. J., Brewer, G., Bruyneel, S., Calvillo, D. P., Campbell, W. K., Cannon, P. R., Carlucci, M., Carruth, N. P., Cheung, T., Crowell, A., De Ridder, D. T. D., Dewitte, S., ... Zwieneberg, M. (2016). A Multilab Preregistered Replication of the Ego-Depletion Effect. *Perspectives on Psychological Science*, 11(4), 546–573. <https://doi.org/10.1177/1745691616652873>
- Hartung, J., Makambi, K. H., & Argac, D. (2001). An Extended ANOVA F-Test with Applications to the Heterogeneity Problem in Meta-Analysis. *Biometrical Journal*, 43(2), 135–146. [https://doi.org/10.1002/1521-4036\(200105\)43:2<135::AID-BIMJ135>3.0.CO;2-H](https://doi.org/10.1002/1521-4036(200105)43:2<135::AID-BIMJ135>3.0.CO;2-H)
_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/1521-4036%28200105%2943%3A2%3C135%3A%3AAID-BIMJ135%3E3.0.CO%3B2-H>
- Hedeker, D., Mermelstein, R. J., & Demirtas, H. (2008). An Application of a Mixed-Effects Location Scale Model for Analysis of Ecological Momentary Assessment (EMA) Data. *Biometrics*, 64(2), 627–634. <https://doi.org/10.1111/j.1541-0420.2007.00924.x>
- Hedeker, D., Mermelstein, R. J., & Demirtas, H. (2012). Modeling between-subject and within-subject variances in ecological momentary assessment data using

- mixed-effects location scale models. *Statistics in Medicine*, 31(27), 3328–3336.
<https://doi.org/10.1002/sim.5338>
- Hedges, L. V. (1992). Meta-Analysis. *Journal of Educational Statistics*, 17(4), 279–296.
<https://doi.org/10.3102/10769986017004279>
- Hedges, L. V., & Olkin, I. (1985). *Statistical Methods for Meta-Analysis*. Academic Press.
- Hedges, L. V., & Pigott, T. D. (2001). The power of statistical tests in meta-analysis. *Psychological Methods*, 6(3), 203–217.
- Hedges, L. V., & Pigott, T. (2005). The Power of Statistical Tests for Moderators in Meta-Analysis. *Psychological methods*, 9, 426–45.
<https://doi.org/10.1037/1082-989X.9.4.426>
- Hedges, L. V., Tipton, E., & Johnson, M. C. (2010). Robust variance estimation in meta-regression with dependent effect size estimates. *Research Synthesis Methods*, 1(1), 39–65. <https://doi.org/10.1002/jrsm.5>
_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/jrsm.5>
- Hedges, L. V., & Vevea, J. L. (1998). Fixed- and random-effects models in meta-analysis. *Psychological Methods*, 3(4), 486–504. <https://doi.org/10.1037/1082-989X.3.4.486>
- Higgins, J. P. T., & Thompson, S. G. (2002). Quantifying heterogeneity in a meta-analysis. *Statistics in Medicine*, 21(11), 1539–1558. <https://doi.org/10.1002/sim.1186>
_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/sim.1186>
- Higgins, J. P. T., Thompson, S. G., & Spiegelhalter, D. J. (2009). A re-evaluation of random-effects meta-analysis. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 172(1), 137–159.
<https://doi.org/10.1111/j.1467-985X.2008.00552.x>
_eprint: <https://rss.onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-985X.2008.00552.x>

- Ioannidis, J. P. A., Patsopoulos, N. A., & Evangelou, E. (2007). Uncertainty in heterogeneity estimates in meta-analyses. *BMJ*, *335*(7626), 914–916.
<https://doi.org/10.1136/bmj.39343.408449.80>
- Kalaian, H. A., & Raudenbush, S. W. (1996). A multivariate mixed linear model for meta-analysis. *Psychological Methods*, *1*(3), 227–235.
<https://doi.org/10.1037/1082-989X.1.3.227>
- Klein, R. A., Vianello, M., Hasselman, F., Adams, B. G., Adams, R. B., Alper, S., Aveyard, M., Axt, J. R., Babalola, M. T., Bahník, Š., Batra, R., Berkics, M., Bernstein, M. J., Berry, D. R., Bialobrzeska, O., Binan, E. D., Bocian, K., Brandt, M. J., Busching, R., . . . Nosek, B. A. (2018). Many Labs 2: Investigating Variation in Replicability Across Samples and Settings. *Advances in Methods and Practices in Psychological Science*, *1*(4), 443–490.
<https://doi.org/10.1177/2515245918810225>
- Knapp, G., & Hartung, J. (2003). Improved tests for a random effects meta-regression with a single covariate. *Statistics in Medicine*, *22*(17), 2693–2710.
<https://doi.org/10.1002/sim.1482>
_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/sim.1482>
- Konstantopoulos, S. (2011). Fixed effects and variance components estimation in three-level meta-analysis. *Research Synthesis Methods*, *2*(1), 61–76.
<https://doi.org/10.1002/jrsm.35>
_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/jrsm.35>
- Lakens, D., Hilgard, J., & Staaks, J. (2016). On the reproducibility of meta-analyses: Six practical recommendations. *BMC Psychology*, *4*(1), 24.
<https://doi.org/10.1186/s40359-016-0126-3>
- Langan, D., Higgins, J. P. T., & Simmonds, M. (2017). Comparative performance of heterogeneity variance estimators in meta-analysis: A review of simulation studies.

- Research Synthesis Methods*, 8(2), 181–198. <https://doi.org/10.1002/jrsm.1198>
_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/jrsm.1198>
- Lee, Y., & Nelder, J. A. (2006). Double hierarchical generalized linear models (with discussion). *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 55(2), 139–185. <https://doi.org/10.1111/j.1467-9876.2006.00538.x>
_eprint: <https://rss.onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-9876.2006.00538.x>
- Lewis, S., & Clarke, M. (2001). Forest plots: Trying to see the wood and the trees. *BMJ : British Medical Journal*, 322(7300), 1479–1480.
- López-López, J. A., Marín-Martínez, F., Sánchez-Meca, J., Van den Noortgate, W., & Viechtbauer, W. (2014). Estimation of the predictive power of the model in mixed-effects meta-regression: A simulation study. *British Journal of Mathematical and Statistical Psychology*, 67(1), 30–48. <https://doi.org/10.1111/bmsp.12002>
_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/bmsp.12002>
- Murphy, B. P. (1967). Some two-sample tests when the variances are unequal: A simulation study. *Biometrika*, 54(3-4), 679–683. <https://doi.org/10.1093/biomet/54.3-4.679>
- Patterson, H. D., & Thompson, R. (1971). Recovery of inter-block information when block sizes are unequal. *Biometrika*, 58(3), 545–554.
- R Core Team. (2021). *R: A language and environment for statistical computing*. Manual. R Foundation for Statistical Computing. Vienna, Austria.
- Raudenbush, S. W. (1994). Random effects models. In *The handbook of research synthesis* (pp. 301–321). Russell Sage Foundation.
- Raudenbush, S. W., Becker, B. J., & Kalaian, H. (1988). Modeling multivariate effect sizes. *Psychological Bulletin*, 103(1), 111–120.
<https://doi.org/10.1037/0033-2909.103.1.111>

- Rodriguez, J. E., & Williams, D. R. (2022). Psymetadata: An R Package Containing Open Datasets from Meta-Analyses in Psychology. *Journal of Open Psychology Data*, 10(1), 8. <https://doi.org/10.5334/jopd.61>
- Rubio-Aparicio, M., López-López, J. A., Viechtbauer, W., Marín-Martínez, F., Botella, J., & Sánchez-Meca, J. (2020). Testing Categorical Moderators in Mixed-Effects Meta-analysis in the Presence of Heteroscedasticity. *The Journal of Experimental Education*, 88(2), 288–310. <https://doi.org/10.1080/00220973.2018.1561404>
- Rubio-Aparicio, M., Sánchez-Meca, J., López-López, J. A., Botella, J., & Marín-Martínez, F. (2017). Analysis of categorical moderators in mixed-effects meta-analysis: Consequences of using pooled versus separate estimates of the residual between-studies variances. *British Journal of Mathematical and Statistical Psychology*, 70(3), 439–456. <https://doi.org/10.1111/bmsp.12092>
_eprint: <https://bpspsychub.onlinelibrary.wiley.com/doi/pdf/10.1111/bmsp.12092>
- Sánchez-Meca, J., & Marín-Martínez, F. (2008). Confidence intervals for the overall effect size in random-effects meta-analysis. *Psychological Methods*, 13(1), 31–48.
<https://doi.org/10.1037/1082-989X.13.1.31>
- Scheffé, H. (1959). *The Analysis of Variance*. John Wiley & Sons.
- Schoemann, A. M. (2016). Using multiple group modeling to test moderators in meta-analysis. *Research Synthesis Methods*, 7(4), 387–401.
<https://doi.org/10.1002/jrsm.1200>
_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/jrsm.1200>
- Thompson, S. G., & Sharp, S. J. (1999). Explaining heterogeneity in meta-analysis: A comparison of methods. *Statistics in Medicine*, 18(20), 2693–2708.
[https://doi.org/10.1002/\(SICI\)1097-0258\(19991030\)18:20<2693::AID-SIM235>3.0.CO;2-V](https://doi.org/10.1002/(SICI)1097-0258(19991030)18:20<2693::AID-SIM235>3.0.CO;2-V)
_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/%28SICI%291097-0258%2819991030%2918%3A20%3C2693%3A%3AAID-SIM235%3E3.0.CO%3B2-V>

- Tipton, E. (2015). Small sample adjustments for robust variance estimation with meta-regression. *Psychological Methods*, 20(3), 375–393.
<https://doi.org/10.1037/met0000011>
- Tipton, E., Pustejovsky, J. E., & Ahmadi, H. (2019). Current practices in meta-regression in psychology, education, and medicine. *Research Synthesis Methods*, 10(2), 180–194. <https://doi.org/10.1002/jrsm.1339>
_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/jrsm.1339>
- van Agteren, J., Iasiello, M., Lo, L., Bartholomaeus, J., Kopsaftis, Z., Carey, M., & Kyrios, M. (2021). A systematic review and meta-analysis of psychological interventions to improve mental wellbeing. *Nature Human Behaviour*, 5(5), 631–652. <https://doi.org/10.1038/s41562-021-01093-w>
- Van den Noortgate, W., López-López, J. A., Marín-Martínez, F., & Sánchez-Meca, J. (2013). Three-level meta-analysis of dependent effect sizes. *Behavior Research Methods*, 45(2), 576–594. <https://doi.org/10.3758/s13428-012-0261-6>
- van Erp, S., Verhagen, J., Grasman, R. P. P. P., & Wagenmakers, E.-J. (2017). Estimates of Between-Study Heterogeneity for 705 Meta-Analyses Reported in *Psychological Bulletin* From 1990–2013. *Journal of Open Psychology Data*, 5(1), 4.
<https://doi.org/10.5334/jopd.33>
- Viechtbauer, W. (2010). Conducting Meta-Analyses in *R* with the **metafor** Package. *Journal of Statistical Software*, 36(3). <https://doi.org/10.18637/jss.v036.i03>
- Viechtbauer, W. (2021). *Metafor: A meta-analysis package for R*. Manual.
- Viechtbauer, W., & López-López, J. A. (2022). Location-scale models for meta-analysis. *Research Synthesis Methods*, n/a(n/a), 1–19. <https://doi.org/10.1002/jrsm.1562>
_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/jrsm.1562>
- Viechtbauer, W., López-López, J. A., Sánchez-Meca, J., & Marín-Martínez, F. (2015). A comparison of procedures to test for moderators in mixed-effects meta-regression

models. *Psychological Methods*, 20(3), 360–374.

<https://doi.org/http://dx.doi.org/10.1037/met0000023>

Welch, B. L. (1938). The Significance of the Difference Between Two Means when the Population Variances are Unequal. *Biometrika*, 29(3/4), 350–362.

<https://doi.org/10.2307/2332010>

Welch, B. L. (1947). The Generalization of ‘Student’s’ Problem when Several Different Population Variances are Involved. *Biometrika*, 34(1/2), 28.

<https://doi.org/10.2307/2332510>

Williams, D. R., Rodriguez, J. E., & Bürkner, P.-C. (2021). Putting Variation into Variance: Modeling Between-Study Heterogeneity in Meta-Analysis.

<https://doi.org/10.31234/osf.io/9vkqy>

Appendix A

Complete simulation results (two groups)

Power for $\beta_1 = 0.5$

Table A1

Power for $\beta_1 = 0.5$

τ_2^2	k_2	Q -test		F -test	
		MEM	MELSM	MEM	MELSM
0.05	5	0.614	0.813	0.544	0.782
	10	0.867	0.926	0.839	0.908
	20	0.970	0.969	0.964	0.964
	30	0.989	0.980	0.986	0.976
	40	0.993	0.982	0.992	0.98
0.10	5	0.576	0.701	0.527	0.664
	10	0.828	0.864	0.794	0.842
	20	0.945	0.944	0.939	0.939
	30	0.978	0.971	0.972	0.962
	40	0.984	0.976	0.983	0.973
0.15	5	0.547	0.618	0.501	0.58
	10	0.780	0.799	0.754	0.775
	20	0.920	0.920	0.913	0.913
	30	0.957	0.950	0.956	0.951
	40	0.974	0.967	0.973	0.966
0.20	5	0.522	0.540	0.491	0.514
	10	0.741	0.739	0.712	0.713
	20	0.888	0.887	0.878	0.877
	30	0.937	0.937	0.926	0.924

	40	0.958	0.956	0.955	0.953
0.25	5	0.513	0.503	0.465	0.462
	10	0.705	0.688	0.683	0.661
	20	0.856	0.858	0.839	0.839
	30	0.913	0.920	0.895	0.902
	40	0.938	0.945	0.933	0.941
0.30	5	0.501	0.463	0.47	0.422
	10	0.680	0.642	0.644	0.6
	20	0.823	0.824	0.81	0.811
	30	0.885	0.896	0.874	0.887
	40	0.918	0.933	0.904	0.922
0.35	5	0.495	0.430	0.454	0.393
	10	0.656	0.599	0.619	0.566
	20	0.793	0.793	0.771	0.773
	30	0.859	0.881	0.847	0.868
	40	0.890	0.915	0.875	0.906
0.40	5	0.486	0.404	0.438	0.361
	10	0.627	0.555	0.593	0.526
	20	0.761	0.762	0.746	0.747
	30	0.828	0.856	0.807	0.839
	40	0.864	0.902	0.852	0.895

Note. “MEM” indicates power rates observed under the standard mixed-effects model and “MELSM” indicates power rates observed under the MELSM with the moderator included in the scale submodel.

Power for $\beta_1 = 0.2$

Table A2*Power for $\beta_1 = 0.2$*

τ_2^2	k_2	Q -test		F -test	
		MEM	MELSM	MEM	MELSM
0.05	5	0.09	0.25	0.07	0.22
	10	0.21	0.30	0.18	0.28
	20	0.35	0.35	0.33	0.33
	30	0.45	0.38	0.43	0.37
	40	0.50	0.39	0.50	0.39
0.10	5	0.11	0.21	0.09	0.19
	10	0.20	0.26	0.17	0.22
	20	0.31	0.31	0.29	0.29
	30	0.38	0.35	0.37	0.33
	40	0.43	0.37	0.41	0.35
0.15	5	0.13	0.19	0.11	0.17
	10	0.21	0.23	0.17	0.20
	20	0.28	0.28	0.25	0.25
	30	0.33	0.32	0.31	0.30
	40	0.37	0.35	0.34	0.32
0.20	5	0.15	0.18	0.12	0.16
	10	0.20	0.21	0.18	0.19
	20	0.25	0.25	0.24	0.24
	30	0.30	0.31	0.28	0.28
	40	0.32	0.32	0.31	0.32
0.25	5	0.15	0.17	0.14	0.15
	10	0.20	0.19	0.17	0.16

	20	0.24	0.24	0.21	0.21
	30	0.27	0.28	0.24	0.26
	40	0.28	0.31	0.27	0.29
0.30	5	0.16	0.16	0.14	0.15
	10	0.20	0.18	0.17	0.15
	20	0.22	0.22	0.20	0.20
	30	0.23	0.26	0.23	0.25
	40	0.26	0.29	0.24	0.28
0.35	5	0.18	0.17	0.15	0.14
	10	0.19	0.16	0.16	0.14
	20	0.21	0.21	0.19	0.19
	30	0.21	0.24	0.20	0.23
	40	0.23	0.28	0.21	0.26
0.40	5	0.19	0.15	0.16	0.14
	10	0.19	0.16	0.17	0.14
	20	0.21	0.21	0.18	0.18
	30	0.20	0.24	0.18	0.22
	40	0.21	0.27	0.19	0.25

Note. $k_1 = 20$; $\tau_1^2 = 0.2$. “MEM” indicates power rates observed under the standard mixed-effects model and “MELSM” indicates power rates observed under the MELSM with the moderator included in the scale submodel.

Power for $\beta_2 = 0.8$

Table A3

Power for $\beta_1 = 0.8$

Q -test	F -test
-----------	-----------

τ_2^2	k_2	MEM	MELSM	MEM	MELSM
0.05	5	0.97	0.99	0.95	0.99
	10	1.00	1.00	1.00	1.00
	20	1.00	1.00	1.00	1.00
	30	1.00	1.00	1.00	1.00
	40	1.00	1.00	1.00	1.00
0.10	5	0.94	0.97	0.93	0.96
	10	1.00	1.00	0.99	1.00
	20	1.00	1.00	1.00	1.00
	30	1.00	1.00	1.00	1.00
	40	1.00	1.00	1.00	1.00
0.15	5	0.92	0.93	0.90	0.91
	10	0.99	0.99	0.99	0.99
	20	1.00	1.00	1.00	1.00
	30	1.00	1.00	1.00	1.00
	40	1.00	1.00	1.00	1.00
0.20	5	0.89	0.88	0.87	0.85
	10	0.98	0.98	0.98	0.98
	20	1.00	1.00	1.00	1.00
	30	1.00	1.00	1.00	1.00
	40	1.00	1.00	1.00	1.00
0.25	5	0.87	0.83	0.85	0.81
	10	0.97	0.97	0.97	0.96
	20	1.00	1.00	1.00	1.00
	30	1.00	1.00	1.00	1.00
	40	1.00	1.00	1.00	1.00
0.30	5	0.85	0.80	0.83	0.77

	10	0.96	0.95	0.96	0.94
	20	1.00	1.00	0.99	0.99
	30	1.00	1.00	1.00	1.00
	40	1.00	1.00	1.00	1.00
0.35	5	0.82	0.75	0.80	0.71
	10	0.96	0.94	0.94	0.92
	20	0.99	0.99	0.99	0.99
	30	1.00	1.00	1.00	1.00
	40	1.00	1.00	1.00	1.00
0.40	5	0.81	0.71	0.78	0.67
	10	0.94	0.90	0.92	0.89
	20	0.99	0.99	0.99	0.99
	30	1.00	1.00	1.00	1.00
	40	1.00	1.00	1.00	1.00

Note. $k_1 = 20$; $\tau_1^2 = 0.2$. “MEM” indicates power rates observed under the standard mixed-effects model and “MELSM” indicates power rates observed under the MELSM with the moderator included in the scale submodel.

Appendix B

Simulation Results for Three Groups

In order to examine whether the results in Section [Simulation Studies](#) generalized to situations involving the comparison of more than two groups, we conducted supplemental simulation studies. These followed the same procedures described in [Simulation Studies](#), but included a third level in the categorical moderator. Thus, in addition to k_2 and τ_2^2 , we also systematically varied the between-study variance and study sample size for the third group, k_3 and τ_3^2 . These quantities were varied according the values in Table [B1](#). Within the columns for k_i and τ_i^2 , each row represents a condition. For example, in the first condition for study sample size, $k_2 = k_3 = 5$, and for the between-study variance, $\tau_2^2 = \tau_3^2 = 0.2$ in the first condition. In total, the total number of conditions was 2 (Q - or F -test) $\times 2$ (mixed-effects model or MELSM) $\times 5$ (k_i) $\times 9$ (τ_i^2) = 180 conditions. In each condition, the second and third groups were compared to the first group, whose values remained fixed at $\tau_1^2 = 0.2$ and $k_1 = 20$.

To study how well the MELSM controls the Type I error of the Q - and F -tests relative to the mixed-effects model when there are more than two groups, we set $\beta_0 = \beta_1 = \beta_2 = 0$, or in terms of group means, $\mu_1 = \mu_2 = \mu_3 = 0$. Both moderators tests were conducted for β_1 and for β_2 . The α level for each test was set to 0.05. The results for β_1 can be seen in Table [B2](#) and the results for β_2 can be seen in Table [B3](#).

To study the empirical statistical power rates of the moderator tests with respect to the two moderator tests, we kept everything the same as in the Type I error rate simulation, with the exception that we change the values of the moderator. For the power simulation study, the first moderator effect was set to $\beta_1 = 0.2$ or $\mu_2 - \mu_1 = 0.2$, and the second moderator effect was set to $\beta_1 = 0.8$ or $\mu_2 - \mu_1 = 0.8$. The power rates for the moderator effects can be found in Tables [B4](#) and [B5](#), respectively.

The patterns observed for the Type I error and power rates for the mixed-effects model and the MELSM qualitatively mirror those in the section [Simulation Studies](#) and are

briefly described in the [Discussion](#).

Table B1
Simulation Values

k_i		τ_i^2	
k_2	k_3	τ_2^2	τ_3^2
5	5	0.2	0.2
10	10	0.1	0.2
20	20	0.2	0.1
40	40	0.1	0.1
10	40	0.1	0.4
40	10	0.4	0.1
		0.4	0.4
		0.2	0.4
		0.4	0.2

Note. $k_1 = 20$, $\tau_1^2 = 0.2$

Table B2
Type I Error Rate for β_2

				Q -test		F -test	
τ_2^2	τ_3^2	k_2	k_3	MEM	MELSM	MEM	MELSM
0.2	0.2	5	5	0.05	0.09	0.06	0.09
		10	10	0.05	0.06	0.06	0.07
		20	20	0.05	0.05	0.05	0.05
		40	40	0.05	0.05	0.05	0.05
		10	40	0.05	0.06	0.06	0.06
		40	10	0.05	0.05	0.05	0.05
0.1	0.2	5	5	0.02	0.08	0.03	0.08
		10	10	0.03	0.06	0.03	0.07
		20	20	0.04	0.05	0.04	0.05
		40	40	0.05	0.06	0.05	0.06

			10	40	0.02	0.06	0.03	0.06
			40	10	0.07	0.06	0.07	0.06
0.2	0.1		5	5	0.05	0.09	0.07	0.10
			10	10	0.06	0.06	0.07	0.07
			20	20	0.07	0.06	0.08	0.06
			40	40	0.07	0.05	0.08	0.06
			10	40	0.09	0.06	0.09	0.07
			40	10	0.06	0.05	0.06	0.06
0.1	0.1		5	5	0.03	0.08	0.03	0.07
			10	10	0.04	0.06	0.05	0.07
			20	20	0.06	0.05	0.07	0.06
			40	40	0.08	0.06	0.08	0.05
			10	40	0.05	0.06	0.06	0.06
			40	10	0.07	0.05	0.08	0.06
0.1	0.4		5	5	0.01	0.08	0.02	0.08
			10	10	0.02	0.06	0.02	0.06
			20	20	0.02	0.05	0.02	0.06
			40	40	0.03	0.06	0.03	0.06
			10	40	0.01	0.06	0.01	0.06
			40	10	0.05	0.05	0.06	0.06
0.4	0.1		5	5	0.11	0.10	0.13	0.10
			10	10	0.10	0.07	0.11	0.07
			20	20	0.08	0.06	0.09	0.06
			40	40	0.06	0.05	0.06	0.05
			10	40	0.15	0.07	0.14	0.07
			40	10	0.04	0.05	0.04	0.06
0.4	0.4		5	5	0.09	0.10	0.11	0.10

			10	10	0.06	0.06	0.08	0.07
			20	20	0.04	0.06	0.05	0.06
			40	40	0.02	0.05	0.03	0.05
			10	40	0.05	0.06	0.05	0.07
			40	10	0.03	0.05	0.03	0.06
0.2	0.4		5	5	0.04	0.08	0.05	0.09
			10	10	0.03	0.06	0.04	0.07
			20	20	0.03	0.05	0.03	0.06
			40	40	0.03	0.06	0.03	0.06
			10	40	0.02	0.06	0.02	0.07
			40	10	0.04	0.05	0.04	0.06
0.4	0.2		5	5	0.10	0.10	0.12	0.11
			10	10	0.08	0.06	0.09	0.07
			20	20	0.06	0.06	0.07	0.06
			40	40	0.04	0.05	0.05	0.06
			10	40	0.09	0.06	0.10	0.07
			40	10	0.03	0.05	0.04	0.06

Note. $k_1 = 20$; $\tau_1^2 = 0.2$. “MEM” indicates power rates observed under the standard mixed-effects model and “MELSM” indicates power rates observed under the MELSM with the moderator included in the scale submodel.

Table B3

Type I Error Rate for β_2

τ_2^2	τ_3^2	k_2	k_3	<i>Q</i> -test		<i>F</i> -test	
				MEM	MELSM	MEM	MELSM
0.2	0.2	5	5	0.05	0.09	0.06	0.09
		10	10	0.05	0.06	0.06	0.07

		20	20	0.05	0.05	0.06	0.06
		40	40	0.05	0.05	0.05	0.05
		10	40	0.05	0.05	0.05	0.06
		40	10	0.05	0.06	0.06	0.07
0.1	0.2	5	5	0.06	0.09	0.07	0.10
		10	10	0.06	0.06	0.07	0.07
		20	20	0.06	0.05	0.07	0.06
		40	40	0.07	0.05	0.08	0.06
		10	40	0.05	0.05	0.06	0.06
		40	10	0.09	0.06	0.09	0.06
0.2	0.1	5	5	0.02	0.08	0.03	0.08
		10	10	0.03	0.06	0.04	0.07
		20	20	0.04	0.05	0.04	0.06
		40	40	0.05	0.06	0.05	0.06
		10	40	0.07	0.06	0.07	0.06
		40	10	0.03	0.06	0.03	0.06
0.1	0.1	5	5	0.03	0.08	0.03	0.08
		10	10	0.04	0.06	0.05	0.06
		20	20	0.06	0.05	0.07	0.06
		40	40	0.08	0.06	0.09	0.06
		10	40	0.07	0.06	0.08	0.06
		40	10	0.06	0.06	0.06	0.07
0.1	0.4	5	5	0.11	0.09	0.13	0.10
		10	10	0.10	0.07	0.10	0.06
		20	20	0.08	0.05	0.09	0.06
		40	40	0.06	0.05	0.07	0.06
		10	40	0.04	0.05	0.04	0.06

		40	10	0.14	0.06	0.14	0.07
0.4	0.1	5	5	0.01	0.08	0.02	0.08
		10	10	0.02	0.06	0.02	0.07
		20	20	0.02	0.06	0.02	0.06
		40	40	0.02	0.06	0.02	0.06
		10	40	0.05	0.06	0.05	0.06
		40	10	0.01	0.06	0.01	0.06
0.4	0.4	5	5	0.08	0.09	0.10	0.10
		10	10	0.06	0.06	0.07	0.07
		20	20	0.04	0.05	0.05	0.06
		40	40	0.02	0.05	0.03	0.06
		10	40	0.03	0.05	0.03	0.06
		40	10	0.04	0.06	0.05	0.07
0.2	0.4	5	5	0.10	0.09	0.12	0.10
		10	10	0.08	0.06	0.09	0.07
		20	20	0.06	0.05	0.07	0.06
		40	40	0.05	0.06	0.05	0.05
		10	40	0.04	0.05	0.04	0.06
		40	10	0.09	0.06	0.10	0.07
0.4	0.2	5	5	0.04	0.08	0.05	0.10
		10	10	0.03	0.06	0.04	0.07
		20	20	0.03	0.06	0.03	0.06
		40	40	0.02	0.06	0.03	0.06
		10	40	0.04	0.05	0.04	0.05
		40	10	0.02	0.06	0.02	0.07

Note. $k_1 = 20$; $\tau_1^2 = 0.2$. “MEM” indicates power rates observed under the standard mixed-effects model and “MELSM” indicates power rates observed under the MELSM with the moderator included in the scale submodel.

Table B4

Power for $\beta_1 = 0.2$

τ_2^2	τ_3^2	k_2	k_3	<i>Q</i> -test		<i>F</i> -test	
				MEM	MELSM	MEM	MELSM
0.2	0.2	5	5	0.13	0.17	0.14	0.18
		10	10	0.17	0.19	0.19	0.20
		20	20	0.24	0.25	0.26	0.27
		40	40	0.31	0.32	0.32	0.33
		10	40	0.18	0.20	0.19	0.21
		40	10	0.30	0.31	0.32	0.33
0.1	0.2	5	5	0.09	0.21	0.11	0.20
		10	10	0.16	0.23	0.18	0.25
		20	20	0.26	0.29	0.28	0.31
		40	40	0.36	0.36	0.37	0.37
		10	40	0.15	0.24	0.17	0.26
		40	10	0.40	0.37	0.41	0.36
0.2	0.1	5	5	0.13	0.17	0.16	0.18
		10	10	0.21	0.19	0.22	0.20
		20	20	0.29	0.25	0.29	0.25
		40	40	0.37	0.32	0.38	0.32
		10	40	0.24	0.19	0.26	0.20
		40	10	0.32	0.32	0.34	0.32
0.1	0.1	5	5	0.10	0.20	0.12	0.21
		10	10	0.20	0.24	0.22	0.26

		20	20	0.32	0.30	0.35	0.32
		40	40	0.44	0.36	0.45	0.37
		10	40	0.23	0.24	0.24	0.25
		40	10	0.42	0.36	0.43	0.37
0.1	0.4	5	5	0.07	0.20	0.09	0.22
		10	10	0.11	0.24	0.13	0.25
		20	20	0.19	0.31	0.19	0.31
		40	40	0.25	0.37	0.26	0.37
		10	40	0.07	0.24	0.07	0.25
		40	10	0.35	0.36	0.37	0.37
0.4	0.1	5	5	0.18	0.14	0.20	0.16
		10	10	0.20	0.15	0.22	0.16
		20	20	0.24	0.19	0.25	0.20
		40	40	0.28	0.26	0.28	0.26
		10	40	0.27	0.15	0.28	0.16
		40	10	0.21	0.25	0.23	0.27
0.4	0.4	5	5	0.14	0.14	0.17	0.16
		10	10	0.15	0.15	0.17	0.16
		20	20	0.16	0.19	0.16	0.19
		40	40	0.17	0.25	0.18	0.27
		10	40	0.12	0.15	0.13	0.15
		40	10	0.19	0.26	0.20	0.27
0.2	0.4	5	5	0.10	0.17	0.12	0.18
		10	10	0.13	0.19	0.15	0.20
		20	20	0.17	0.25	0.19	0.26
		40	40	0.21	0.31	0.23	0.33
		10	40	0.09	0.20	0.10	0.20

			40	10	0.27	0.32	0.28	0.32
0.4	0.2	5	5		0.17	0.14	0.19	0.16
		10	10		0.19	0.15	0.20	0.16
		20	20		0.21	0.20	0.22	0.20
		40	40		0.23	0.26	0.25	0.26
		10	40		0.20	0.15	0.21	0.15
		40	10		0.20	0.25	0.21	0.25

Note. $k_1 = 20$; $\tau_1^2 = 0.2$. “MEM” indicates power rates observed under the standard mixed-effects model and “MELSM” indicates power rates observed under the MELSM with the moderator included in the scale submodel.

Table B5

Power for $\beta_2 = 0.8$

τ_2^2	τ_3^2	k_2	k_3	<i>Q</i> -test		<i>F</i> -test	
				MEM	MELSM	MEM	MELSM
0.2	0.2	5	5	0.87	0.87	0.90	0.89
		10	10	0.98	0.98	0.99	0.98
		20	20	1.00	1.00	1.00	1.00
		40	40	1.00	1.00	1.00	1.00
		10	40	1.00	1.00	1.00	1.00
		40	10	0.98	0.98	0.99	0.98
0.1	0.2	5	5.00	0.88	0.87	0.90	0.88
		10	10	0.99	0.98	0.99	0.98
		20	20	1.00	1.00	1.00	1.00
		40	40	1.00	1.00	1.00	1.00
		10	40	1.00	1.00	1.00	1.00
		40	10	0.99	0.98	0.99	0.98

0.2	0.1	5	5.00	0.93	0.96	0.94	0.96
		10	10	0.99	1.00	1.00	1.00
		20	20	1.00	1.00	1.00	1.00
		40	40	1.00	1.00	1.00	1.00
		10	40	1.00	1.00	1.00	1.00
		40	10	1.00	1.00	1.00	1.00
0.1	0.1	5	5	0.94	0.96	0.95	0.97
		10	10	1.00	1.00	1.00	1.00
		20	20	1.00	1.00	1.00	1.00
		40	40	1.00	1.00	1.00	1.00
		10	40	1.00	1.00	1.00	1.00
		40	10	1.00	1.00	1.00	1.00
0.1	0.4	5	5	0.80	0.69	0.82	0.71
		10	10	0.94	0.90	0.96	0.91
		20	20	0.99	0.99	0.99	0.99
		40	40	1.00	1.00	1.00	1.00
		10	40	1.00	1.00	1.00	1.00
		40	10	0.97	0.90	0.97	0.91
0.4	0.1	5	5	0.90	0.96	0.93	0.97
		10	10	0.99	1.00	0.99	1.00
		20	20	1.00	1.00	1.00	1.00
		40	40	1.00	1.00	1.00	1.00
		10	40	1.00	1.00	1.00	1.00
		40	10	0.98	1.00	0.98	1.00
0.4	0.4	5	5	0.77	0.70	0.79	0.71
		10	10	0.92	0.90	0.93	0.91
		20	20	0.98	0.99	0.98	0.99

			40	40	1.00	1.00	1.00	1.00
			10	40	1.00	1.00	1.00	1.00
			40	10	0.90	0.91	0.91	0.91
0.2	0.4		5	5	0.79	0.69	0.81	0.71
			10	10	0.93	0.90	0.95	0.91
			20	20	0.99	0.99	0.99	0.99
			40	40	1.00	1.00	1.00	1.00
			10	40	1.00	1.00	1.00	1.00
			40	10	0.95	0.91	0.95	0.91
0.4	0.2		5	5	0.85	0.87	0.87	0.88
			10	10	0.97	0.98	0.98	0.98
			20	20	1.00	1.00	1.00	1.00
			40	40	1.00	1.00	1.00	1.00
			10	40	1.00	1.00	1.00	1.00
			40	10	0.96	0.98	0.96	0.98

Note. $k_1 = 20$; $\tau_1^2 = 0.2$. “MEM” indicates power rates observed under the standard mixed-effects model and “MELSM” indicates power rates observed under the MELSM with the moderator included in the scale submodel.