Meta-Analysis in Vocational Behavior: A Systematic Review and
Recommendations for Best Practices

Cort W. Rudolph

Saint Louis University

Kahea Chang

Saint Louis University

Rachel S. Rauvola

Saint Louis University

Hannes Zacher

Leipzig University

Note: This is a pre-print version of an in-press, accepted manuscript. Please cite as:

Rudolph, C.W., Chang, K., Rauvola, R.S., & Zacher, H. (2020, In Press). Meta-analysis in vocational behavior: A systematic review and recommendations for best practices. *Journal of Vocational Behavior*.

Author Note

Cort W. Rudolph, Kahea Chang, & Rachel S. Rauvola, Department of Psychology, Saint Louis University, St. Louis, MO (USA). Hannes Zacher, Institute of Psychology, Leipzig University, Leipzig, Germany.

Correspondence concerning this article may be addressed to Cort W. Rudolph, Department of Psychology, Saint Louis University, Saint Louis MO (USA), e-mail: cort.rudolph@health.slu.edu

Abstract

Meta-analysis is a powerful tool for the synthesis of quantitative empirical research. Overall, the field of vocational behavior has benefited from the results of meta-analyses. Yet, there is still quite a bit to learn about how we can improve the quality of meta-analyses reported in this field of inquiry. In this paper, we systematically review all meta-analyses published in the *Journal of Vocational Behavior* (JVB) to date. We do so to address two related goals: First, based on guidance from various sources (e.g., the American Psychological Association's meta-analysis reporting standards; MARS), we introduce ten facets of meta-analysis that have particular bearing on statistical conclusion validity. Second, we systematically review meta-analyses published in JVB over the past 32 years, with a particular focus on the methods employed; this review informs a discussion of 19 associated "best practices" for researchers who are considering conducting a meta-analysis in the field of vocational behavior (or in related fields). Thus, this work serves as an important benchmark, indicating where we "have been" and where we "should go," with respect to conducting and reporting meta-analyses on vocational behavior topics.

Keywords: Systematic Review; Meta-Analysis; Methodology; Best Practices

Highlights:

- We review all meta-analyses published in the *Journal of Vocational Behavior* to date.
- We evaluate studies using ten criteria related to statistical conclusion validity.
- We derive 19 "best practices" for future meta-analyses in vocational behavior.

Meta-Analysis in Vocational Behavior: A Systematic Review and Recommendations for Best Practices

1. Introduction

The idea that quantitative research results can be statistically aggregated is not new. Over 100 years ago, Pearson (1904) empirically combined the results of multiple clinical studies of typhoid inoculations. Likewise, Fisher (1925) proposed a method for pooling p -values across null hypothesis significance tests. The introduction of more modern conceptualizations of “meta-analysis” as a research synthesis method is often attributed to Glass (1976), who is also credited for popularizing these ideas for the synthesis of psychological research. The field of vocational behavior was an early adopter of meta-analysis. Indeed, over the past 32 years, the flagship outlet for such work, the *Journal of Vocational Behavior* (JVB), has published 68 meta-analyses (see Figure 1), on topics ranging from occupational wellbeing (Assouline & Meir, 1987) to vocational interest congruence (Nye, Su, Rounds, & Drasgow, 2017).

Meta-analyses have arguably been impactful to the field of vocational behavior as a whole. For example, as of the writing of this manuscript, four of the 25 “most downloaded” JVB articles in the past 90 days (16%), and three of the 25 “most cited” articles since 2016 (12%), have been meta-analyses. Moreover, according to Web of Science citation counts, a meta-analysis of organizational commitment by Meyer, Stanley, Herscovitch, and Topolnytsky (2002) has been cited just over 2,100 times; the median citation count across all 68 meta-analyses published in JVB is 55 ($M = 147.02$, $SD = 309.95$; see also Figure 2).

In this manuscript, we systematically review and synthesize the entire corpus of meta-analytic articles that have been published in JVB. Systematic reviews are typically undertaken to synthesize the findings of primary empirical studies (e.g., Gough, Oliver, & Thomas, 2017). Our approach to this systematic review is somewhat different. Instead of integrating the findings of

meta-analyses published in JVB in a general sense, our primary focus is on the methods employed to conduct meta-analyses and on the structure used to report these meta-analyses.

Our goals for this systematic review are twofold: Our primary goal is to quantify the state of meta-analytic methods and to trace the development of meta-analytic methods applied to the study of vocational behavior phenomena over time, as published in JVB. We also aim to ascertain “gaps” that exist in the design, conduct, and presentation of meta-analytic studies published therein to date. Informed by the results of this systematic review, our second goal is to outline a set of “best practices” that are organized around the ten facets of our review and that guide the conduct and review of future meta-analyses in JVB, and for the field of vocational behavior more broadly defined (see Table 1). Thus, two overarching research questions that guide our review are, “How are meta-analyses published in JVB ‘done’?,” and “Do meta-analyses published in JVB conform to ‘best practices’?”

To answer these questions, we organize our review around ten interrelated facets of the design and conduct of meta-analysis that have particular bearing on *statistical conclusion validity* (i.e., the extent to which the conclusions about meta-analyzed relationships are correct or “reasonable”; Shadish, Cook, & Campbell, 2002). These facets were derived from multiple sources. First, we consulted the American Psychological Association’s (APA) meta-analysis reporting standards—a comprehensive effort to establish criteria against which the scientific rigor of a meta-analysis can be judged (MARS; APA, 2008, 2010). Second, we considered more recent suggestions for applying MARS standards specifically to meta-analyses in the organizational sciences (Kepes, McDaniel, Brannick, & Banks, 2013). Third, we referenced recent “best practice” recommendations for the conduct of meta-analyses (Siddaway, Wood, & Hedges, 2019). Finally, we triangulated advice from each of these three sources against contemporary

reference books regarding the design and conduct of meta-analyses (Borenstein, Hedges, Higgins, & Rothstein, 2011; Cooper, Hedges, & Valentine, 2009; Schmidt & Hunter, 2015).

In our online appendix, we offer a “crosswalk,” tying common advice across these multiple sources to the ten facets of meta-analysis and the 19 best practices we derive therefrom: <https://osf.io/pgujx/>. Importantly, the primary focus of our review is on the statistical methods involved in the conduct of meta-analyses, and not on the supporting methods involved in such reviews (for a comprehensive review of literature search strategies that support systematic reviews and meta-analyses, see Harari, Parola, Hartwell, & Riegelman, 2020).

Of note, our focus on the ten facets of meta-analysis is not designed to represent an exhaustive methodological summary and critique of every meta-analysis published in JVB to date. Rather, we focus on those ten facets of the design and conduct of meta-analysis that, if adopted prescriptively, would have the most “influence” on the broader applicability and impact of meta-analytic findings to the field as a whole. Moreover, our focus is on those facets of the meta-analytic process that are most *actionable* (i.e., those which researchers have most control over in the design, conduct, and reporting of meta-analyses), and that can be readily translated into best practices. Table 1 summarizes these ten facets and the best practice recommendations that we offer as guidance for researchers seeking to conduct meta-analyses of vocational behavior topics, including relevant cautionary notes and related practical advice, and notes about additional readings and resources to guide such efforts. To begin our discussion, we next summarize two predominant traditions of meta-analysis (i.e., Hedges-Olkin & Schmidt-Hunter) and then introduce the ten facets of meta-analysis that guided our review.

2. Two Traditions of Meta-Analysis

The term “meta-analysis” refers to a process of systematically and quantitatively summarizing a body of literature. Generally speaking, meta-analyses are conducted to achieve a

set of common goals. The overarching goal of any meta-analysis is to estimate average effects that are representative of population effects (e.g., population correlations, ρ_{xy}) based upon the cumulation of multiple sample effects (e.g., correlations from individual primary studies, r_{xy}). Moreover, meta-analyses generally involve procedures for differentially weighting such sample effects to account for variability in the precision of such estimates (e.g., weighting each r_{xy} by its respective sample size, n). Finally, meta-analyses typically provide estimates of the variability of effects from study to study (i.e., estimates of the heterogeneity of r_{xy} from one study, to another); to the extent that appreciable heterogeneity is observed, it is possible to model study-level characteristics as moderators (e.g., study design features) that might explain this variability.

Throughout the following, we will refer to two traditions of meta-analysis (named after the researchers who pioneered these methods): Hedges-Olkin and Schmidt-Hunter. These traditions largely represent meta-analytic “schools of thought,” which are characterized mainly by philosophical differences. At the risk of overgeneralizing, we characterize these traditions as more or less discrete (as have others, see Schmidt & Hunter, 2003), but believe that drawing such a distinction is helpful as there are core differences in the logic and, in some cases, underlying statistical models that characterize these two approaches. Hedges-Olkin-style meta-analysis emerged primarily from educational research (Hedges & Olkin, 1985), whereas Schmidt-Hunter-style meta-analysis emerged from industrial and organizational psychology research (Schmidt & Hunter, 2015). Owing partially to differences in dominant research methods and designs across disciplines (e.g., education vs. industrial and organizational psychology) the Hedges-Olkin approach is often used to synthesize the results of studies adopting experimental/intervention methods, whereas the Schmidt-Hunter approach is often used to synthesize the results of studies adopting observational methods. It is important to emphasize that these are by no means hard-and-fast rules, and that Hedges-Olkin-style meta-analysis is likewise often applied to

observational research, just as Schmidt-Hunter-style meta-analysis is likewise often applied to experimental research; both approaches offer a flexible set of tools for cumulating the results of primary studies that adopt a wide variety of research designs.

Indeed, both the Hedges-Olkin and Schmidt-Hunter traditions are quite similar in their underlying statistical procedures (i.e., both traditions approach the question of synthesis through the application of weighted least squares regression models). Some practical differences that are evident between these traditions include the particular weights used in specifying these weighted least squares models, the means of estimating the random effects variance component, and questions about how to account for small sample “bias” in the cumulation of correlation coefficients (i.e., Fisher’s *r*-to-*z* transformation; see Hedges & Olkin, 1985; Schmidt & Hunter, 2015). Such differences aside, if one were to conduct a Hedges-Olkin-style weighted least squares meta-regression with specific weights (i.e., sample-size weighting for individual study effects) and a Schmidt-Hunter random effects variance estimator, it would yield the same parameter estimates as a Schmidt-Hunter “bare bones” meta-analysis (i.e., a model accounting only for sampling error).

Practical differences notwithstanding, additional differences emerge when considering the philosophical perspectives that underlie these two traditions. For example, the Schmidt-Hunter approach focuses on estimating average effects (and variability around such effects) for hypothetical constructs. To ensure the “best” estimates of such effects, the Schmidt-Hunter approach is especially concerned with accounting for various sources of artefactual variance (i.e., sampling error, but also bias due to other features, including measurement error, selection effects such as range restrictions, and the influence of artificial dichotomization). Rather than focusing on hypothetical constructs, however still focusing on estimates of average effects, the Hedges-Olkin approach focuses on cumulating results across studies and describing the variability

observed between studies (e.g., that due to differences in measures, samples, etc.). As such, the Hedges-Olkin approach is arguably more focused on the development of formalized statistical and theoretical models, rather than on the representation of construct-level relationships.

To illustrate these philosophical differences in more practical terms, consider that Hedges-Olkin-style meta-analysis is typically applied to account for sampling error (i.e., an estimate of the unrepresentativeness of a given sample) alone as a statistical artefact. In contrast, Schmidt-Hunter-style meta-analysis was developed for the purposes of accounting for statistical artefacts that include, but also go beyond, sampling error (e.g., unreliability of predictor and criterion variables, range restriction, artificial dichotomization). Interestingly, this distinction is by no means strictly required, as Hedges and Olkin (1985, p. 135) discuss methods for accounting for statistical artefacts other than sampling error in their original work (although these approaches are seldom applied among meta-analyses that adopt this particular framework).

Indeed, Schmidt-Hunter-style meta-analysis is the methodology that is typically adopted to argue for the concept of *validity generalization* (VG), which is the process of establishing that observed criterion-related validities of various personnel selection instruments generalize across different settings (Schmidt & Hunter, 1977). Of note, we use these labels (i.e., Hedges-Olkin and Schmidt-Hunter) rather broadly, and primarily to indicate how statistical artefacts are treated across traditions. For example, within the Schmidt-Hunter “VG” school of thought, there are very closely related approaches to establishing such evidence (e.g., Raju, Burke, Normand, & Langlois, 1991 present a very similar approach to psychometric meta-analysis that differs only in how sampling error artefacts are accounted for in estimating heterogeneity).

As suggested, we have organized our review around ten common facets of meta-analysis, derived from the MARS standards and related resources, that cut across these two broader traditions of meta-analysis. These ten facets were chosen, as they have bearing on the statistical

conclusion validity of all meta-analyses, regardless of their grounding in any particular meta-analytic tradition. Table 1 summarizes these facets, and we review each in detail, next.

3. Ten Facets of Meta-Analysis

3.1. Consider Methodological and Substantive Moderators. Meta-analysis can address numerous research questions. However, those that are typically of primary interest to the synthesis of any given literature are related to the main and conditional (i.e., moderated) relationships between (assumed) independent–dependent/predictor–criterion variables. Such relationships are typically indexed with standardized mean difference effect size estimates, such as d_{Cohen} or g_{Hedges} , or with covariance-based effect sized estimates, such as zero-order correlations, r_{xy} (for a comprehensive treatment of various effect size metrics, see Ellis, 2010). With respect to main-effect relationships, typically the synthesized effects from primary studies represent an overall, aggregate estimate of an effect size (e.g., an “average” sample-size weighted r_{xy} , sometimes denoted as $\bar{r}_{xy}$) that is weighted by some index of study precision (e.g., sample size, n). Moderators in meta-analysis are understood to exist at the study level of analysis, and moderators are necessarily construed as study-level effects. Note that meta-analysis of interaction effects are possible (see Aguinis, Gottfredson, & Wright, 2011). However, zero-order relationships between interaction terms are still seldom reported in primary studies; no studies identified in our review considered meta-analyses of interaction terms.

There are two approaches that are generally seen in the literature to meta-analytically address questions of moderation: subgroup analysis and meta-regression. The subgroup analysis approach splits studies into subgroups (e.g., “high” vs. “low” levels) at different levels of the moderator variable (i.e., either because categorical moderators are directly coded from primary studies, or through some artificial “splitting” procedure in which otherwise continuous variables

are polytomized in some way, e.g., through a median split). Separate meta-analyses are then carried out at each level of the moderator. The meta-regression approach is a regression-based approach that essentially treats the effect sizes derived from primary studies as outcomes, regressing them onto one or more continuous or categorical moderators (i.e., representing features of these primary studies, and thus characterizing the nature of the effect sizes derived therefrom). Such models can be specified with moderators entered simultaneously or one at a time. As we will discuss below, both approaches have notable advantages and disadvantages.

Substantive moderators refer to those that are theoretically relevant, meaning that they are derived from theory and have some bearing on the strength and/or direction of a relationship, because they are predicted to on the basis of said theory (e.g., job complexity as a moderator of the association between general mental ability and job performance; Schmidt & Hunter, 1998). Accordingly, the exact nature of substantive moderators investigated is highly dependent on the literature being synthesized and, therefore, the theoretical traditions being reviewed. Moreover, substantive moderators can be difficult to justify in meta-analysis, for example, because few theories specify conditional effects *at the study level*. This leads to potentially awkward hypothesizing and the possibility of mismatched inferences regarding the supposed mechanism tested against the level of analysis considered, particularly when interpreting such effects (e.g., inferences based upon “ecological relationships” and the commission of the “ecological fallacy,” wherein one makes an inference regarding individual-level phenomena on the basis of group-level relationships, see Robinson, 1950). Similar sentiments are echoed by Schmidt (2017, p. 474), who suggests, “As the saying goes, you cannot get blood out of a turnip...The accurate detection and calibration of moderators is extremely complex, difficult, and data demanding.”

Methodological moderators refer to those that are relevant to the way in which primary studies have been conducted, meaning that when considered across studies included in any meta-

analysis, they have some hypothesized bearing on the strength or direction of a relationship (e.g., subjective vs. objective indicators of job performance as a criterion variable). That is to say, methodological moderators matter, because the way in which the studies subjected to meta-analysis are conducted changes the nature of their effect sizes in a systematic way. Whereas the consideration of substantive moderators serves as a means to approximate formal tests of theory, the consideration of methodological moderators in meta-analysis is more so a concern for statistical conclusion validity. Indeed, the primary goal of meta-analysis is to make claims about the nature of effect sizes in the population (e.g., their direction and strength). If the direction and/or strength of meta-analytic estimates varies systematically as a function of certain methodological features, and these features are ignored in the computation of an overall effect size estimate, this otherwise unobserved heterogeneity would serve to occlude the proper estimation of such parameters. As with substantive moderators, the exact nature of methodological moderators is dependent on the type of literature being reviewed and the primary methods that this literature employs. For example, the types of methodological concerns would vary substantially between a literature that primarily relies on experimental methods versus one that relies on observational methods; even within studies using experimental versus observational methods, characteristics of study design can vary dramatically. Of particular importance, tests of both methodological and substantive moderators (where theoretically justifiable) should be considered in any meta-analytic effort; both are important to the synthesis of any literature.

Optimal strategies for testing moderators in meta-analysis are the subject of some debate. Some methodologists offer that simple subgroup analyses or hierarchically organized subgroup analyses are a preferred strategy (e.g., Schmidt, 2017), whereas others argue that meta-regression should be favored, particularly because it allows for testing multiple moderators simultaneously and thereby affords one some degree of “statistical control” over the process (e.g., Lipsey &

Wilson, 2001). With respect to the former, the hierarchical approach to subgroup analysis has several advantages, such as allowing for different estimates of heterogeneity across moderator levels. Moreover, because the number of studies included in such analyses is made explicit (i.e., because it is clear how many studies represent each moderator, at each level of the hierarchical analysis), it is arguably easier to understand the degree of uncertainty that underlies conclusions drawn from such procedures (e.g., owing to subgroupings that are represented by “low k ,” which is to say those based upon fewer studies, relatively speaking). With respect to the latter, the meta-regression approach has the advantage of not requiring the explicit polytomization of otherwise continuous moderator variables, thus allowing one to draw upon information from studies at all moderator levels to improve moderator precision.

Considering this point further, it should be clear that the subgroup approach to moderator analysis is simply a restricted form of the more general meta-regression approach, which stratifies on predictor levels (i.e., by running separate sub-group meta-analyses for each level of a moderator), rather than statistically conditioning on them (i.e., by directly regressing effect sizes onto moderators). Thus, like any regression-based model with covariates, a measure of caution is warranted when specifying and interpreting meta-regression models with multiple moderators considered simultaneously (see Becker et al., 2016, for general guidance about such cautions). Regardless of the methodology employed, we would argue that it is more important that methodological moderators *are* tested than by what means. The choice of “how” such effects are modeled will largely depend on the type of moderators being considered (e.g., those with an assumed underlying categorical versus continuous interval/ratio level distribution), the goals of such models (e.g., meta-regression can be used to approximate non-linear study-level relationships; see Sturman, 2003), and the number of effect sizes being considered (i.e., meta-regression is a “high k ” procedure).

To the latter point, regardless of the approach (i.e., subgroup analysis or meta-regression), it is important to keep in mind that testing moderation in meta-analysis is a statistical modeling approach like any other regression-based approach; like all such statistical models, these approaches can only benefit from large sample sizes (i.e., “high k ”) and variability on the variables of interest (i.e., having multiple studies representing all possible levels across one’s different moderators). Indeed, when choosing moderators, be they methodological or substantive in nature, it is very important to ensure there is adequate representation of effect sizes across levels thereof (i.e., that they represent “high k ” cases). Moreover, the management of multiple, often correlated moderators is a challenging task for the meta-analyst. For example, consider two methodological moderators: sample type (i.e., student samples vs. working samples) and study type (i.e., laboratory studies vs. field studies). Considering both of these moderators may seem prudent, however the information they provide may be largely redundant if, for example, a majority of the studies that consider “student samples” also happen to be “laboratory studies.” This issue would be exacerbated in cases where only one of these two variables is considered (e.g., as one could not unambiguously identify the source of variance being modeled, therefore running the risk of incorrectly attributing moderation to an unmeasured confound).

Ancillary to questions about conditional effects, but perhaps more important for the substantive conclusions drawn, meta-analysis often seeks to index the relative consistency or homogeneity of effect size estimates observed across primary studies. Evidence for heterogeneity of effect size estimates is often taken as a necessary precondition for the investigation of moderator variables (Hedges & Olkin, 1985). However, indexing such variability is in and of itself important to the broader conclusions gleaned from meta-analysis. While relevant, there is some confusion in the literature about “how” one should index heterogeneity. For example, one common metric of residual heterogeneity is the Q -statistic. However, it has been shown that the

statistical significance test associated with this statistic is influenced by the number of studies considered. Despite this, it is a widely reported index of heterogeneity. Another common approach to indexing heterogeneity is to compute credibility intervals, which describe a probable range of effect sizes that would be expected in the population (discussed further, below).

3.2. Use Random Effects Estimators. Perhaps no two concepts represented within the statistical vernacular are as misunderstood as “fixed effect” and “random effects” (e.g., Gelman, 2005). For meta-analysis, the distinction between fixed effect and random effects estimators is equally vexing. For our purposes (and without getting into mathematical technicalities), the difference between fixed effect and random effects estimators comes down to how the meta-analytic model treats between-study variability among effect size estimates when computing an overall meta-analytic effect. More complete treatments of these ideas can be found in Borenstein et al. (Borenstein, Hedges, Higgins, & Rothstein, 2010), Schmidt, Hunter, and colleagues (e.g., Hunter & Schmidt, 2000; Schmidt, Oh, & Hayes, 2009), and Hedges and Olkin (1985, p. 190).

When using a fixed effect estimator for a model considering an overall population effect (i.e., a so-called “intercept only” model, in which only an overall estimate of the effect size population parameter is being estimated, absent moderators), we assume that there is a single, “true effect” that is reflected in the observed effects that have been included in the estimate of the overall effect. In doing so, we assume that any differences in observed effects are due to sampling error alone. Generalizing a bit, a fixed effect model with moderators would assume that the specified moderator variable(s) account for some proportion of the non-artefactual heterogeneity observed in this population estimate. In practice, these are strict assumptions that are likely untenable in the vast majority of cases. Indeed, it is difficult to come up with a reasonable rationale for using a fixed effect estimator. Its justification would have to follow from a situation where a researcher was only interested in obtaining an overall effect estimate under

the assumption that they had included all possible studies from the population of studies in their analysis (which itself is largely untenable, and gets muddled by one's definition of "the population"). As suggested by Hunter and Schmidt (2000):

"The major problem with this assumption is that it is difficult (and perhaps impossible) to conceive of a situation in which a researcher would be interested only in the specific studies included in the meta-analysis and would not be interested in the broader task of estimation of the population effect sizes for the research domain as a whole" (p. 277).

In contrast, when using a random effects estimator for a model considering an overall population effect, we assume that this true effect *could* vary from study to study. Much like fixed effect models, random effects models assume that some amount of this variability is due to sampling error. Unlike fixed effect models, however, random effects models also assume that there is variability in this estimate that is attributable to the fact that observed effects themselves are sampled from a larger population of all possible observed effects. Thus, in a random effects model, the estimated variance among observed effects is decomposed into two component parts, which represent within-study and between-study variability, respectively. Scholars who are familiar with within- and between-unit variability distinctions in multilevel research may see parallels in this argument (see also Raudenbush & Bryk, 1985, for a classic treatment of this distinction). Indeed, a fixed effect model is a restricted case of a more general random effects model, where the between-study variability is *a priori* assumed to be zero.

There are various indices of between-study variability that have been proposed, and we will focus on two here: tau-squared (τ^2) and the standard deviation of rho (SD_ρ). First, tau-squared (τ^2) can be interpreted as the variance of the observed effects across the population of all possible studies, and thus reflects an estimate of the variance of the "true effect" parameter. Consistent with Raudenbush (2009), a credibility interval can likewise be computed around the true effect, $\hat{\mu}$, as a function of τ :

$$\hat{\mu} \pm z_{1-\frac{\alpha}{2}} \tau$$

where $z_{1-\frac{\alpha}{2}}$ is the $100 \times (1 - \frac{\alpha}{2})^{th}$ percentile of a standard normal distribution (e.g., 1.28 for $\alpha = .20$, corresponding to an 80% credibility interval). If, for example, an 80% credibility interval were computed, the range between the upper and lower bounds of this interval would represent an estimate of where 80% of the true effects would fall in the hypothetical population of studies.

Confusing matters somewhat, “credibility intervals” are also sometimes conflated with the concept of “prediction intervals” in the meta-analytic literature, although the former is based solely on estimates of estimated random-effects variance, τ , as shown above, whereas the latter additionally considers estimates of the sampling variance (i.e., thus accounting for both random effects and estimated individual-study sampling error). Prediction intervals serve to index the plausible range of population parameter values to be expected *if a new study were to be conducted* (see Higgins, Thompson, & Spiegelhalter, 2009; Inhout, Ioannidis, Rovers, & Goeman, 2016). Thus, if one’s goal is to generalize meta-analytic findings back to the population rather than to a hypothetical “new study,” there is generally more value in reporting the credibility interval than the prediction interval. Complicating matters further, the popular ‘metafor’ packages for the R statistical computing environment (Viechtbauer, 2010) uses the terms “credibility” and “prediction” intervals interchangeably. Thus, it is important to be clear which interval is being reported, and to interpret its width correctly.

When conducting Hedges-Olkin-style meta-analysis, one has to make a choice between adopting fixed effect versus random effects estimators. As alluded to, there are very few cases in which it *does not* make sense to choose a random effects estimator. In doing so, estimates of τ^2 are typically derived in Hedges-Olkin-style meta-analysis using either restricted maximum likelihood (REML) or DerSimonian-Laird estimators (DerSimonian & Laird, 1986). Of the two,

REML estimators perform more favorably, as they provide better estimates of between-study variance in various simulation studies (e.g., Langan et al., 2019; Veroniki et al., 2018).

Following the strict advice of Schmidt-Hunter-style meta-analysis, one does not have a choice between fixed effect versus random effects estimators, as the Schmidt-Hunter model is inherently a random effects model (see Schmidt et al., 2009, p. 100). Importantly, moving beyond such strict advice, hybridized approaches that blend Schmidt-Hunter and Hedges-Olkin estimation procedures are indeed possible, and in some cases have been advocated (e.g., Brannick, Potter, Benitez, & Morris, 2019). Estimates of between-study variability in the Schmidt-Hunter tradition typically adopt a different notation, standard deviation of rho (SD_ρ), to represent the non-artefactual variance in effect size estimates across studies. The standard deviation of rho (SD_ρ) is an estimate of the standard deviation of the distribution of “true” effect sizes (i.e., rho, ρ). Thus, SD_ρ is conceptually the same as $\sqrt{\tau^2}$, and can be used to estimate credibility intervals around the true effect (see also Schmidt & Hunter, 2015, p. 228).

Regardless of the tradition adopted, some caution must be applied when interpreting estimates of τ/SD_ρ . Specifically, in cases where the empirical distributions of effect sizes are not symmetrical, observed variability is more meaningful at certain ranges of such effects, and in particular when it occurs in the “center” of the empirical distribution of effects (see Wiernik, Kostal, Wilmot, Dilchert, & Ones, 2017). In this regard, Paterson et al. (2016) report that the distribution of correlations observed in the management literature is positively skewed, with a majority falling between $\rho = .15$ and $\rho = .40$ (Paterson, Harms, Steel, & Credé, 2016). Thus, SD_ρ must be understood in tandem with estimates of ρ and their associated credibility interval endpoints against a known empirical distribution of effect sizes. Given their correspondence, the same argument could be levied against estimates of τ .

3.3. Conduct Sensitivity Analyses. One of the many appeals of meta-analysis is that it provides a systematic methodology for synthesizing and summarizing research literatures. That said, it is important to acknowledge that researcher decisions play a distinct role in the process of conducting meta-analyses (e.g., Aguinis, Dalton, Bosco, Pierce, & Dalton, 2011; Geyskens, Krishnan, Steenkamp, & Cunha, 2009). Some have even argued that, owing to the degree of subjectivity surrounding such decisions, the various exercises of statistical analysis give little more than the “illusion of objectivity” (Berger & Berry, 1988). Salient examples of how researchers’ decisions can influence the outcomes of meta-analyses can be gleaned from case studies that compare the results of independent meta-analyses conducted within a common research domain (e.g., Nieminen, Nicklin, McClure, & Chakrabarti, 2011).

Despite these criticisms, the good news is that, in many cases, the influence of researchers’ decisions can be evaluated; this is the fundamental idea behind sensitivity analyses. As suggested by Cooper (2017), “Statistical sensitivity analyses are used to determine whether and how the conclusions of your analyses might differ if they were conducted using different statistical procedures or different assumptions about the data” (p. 265). Thus, for example, a meta-analyst might question whether the decision to combine effect sizes derived from studies variously adopting the long-form (e.g., 30 items) versus the short-form (e.g., 15 items) of a particular criterion measure. Comparing the meta-analytic effect sizes from studies adopting one form versus the other can be considered a form of sensitivity analysis. One could separately compute such effects (e.g., using a subgroup analysis; by treating “scale length” as a categorical moderator) along with their associated confidence intervals to answer a question like, “Do I reach different conclusions when I consider studies adopting different forms of the criterion measure?”

The goal of sensitivity analyses is to understand the robustness of meta-analytic findings with regard to different sets of statistical assumptions. Extending our example, if the answer to

the question posed above is “no,” which is to say that the conclusions drawn about the average meta-analytic effect do not seem to differ when considering different forms of the criterion measure, then a higher degree of confidence can be placed in this conclusion. If, however, the results are found to differ under different assumptions, then some degree of measured caution, or perhaps a different interpretation of one’s findings altogether, may be warranted.

Beyond checks on researcher decisions, sensitivity analyses can serve as checks on the robustness of one’s findings to various statistical concerns. One such analysis, sometimes termed “influence analysis,” is the consideration of statistical outliers (e.g., via so-called “leave-one-out” procedures; see Viechtbauer & Cheung, 2010). A related concern, the assessment of selection effects (e.g., publication bias), is an important sub-classification of sensitivity analyses, and will be discussed separately in some detail, below. Another common sensitivity analysis that has been prescribed is a consideration of study quality (e.g., Cooper, 2017).

Although various analytic tools exist to aide in the conduct of such sensitivity analyses in a piecemeal fashion (i.e., often built into statistical packages for the conduct of meta-analyses, e.g., the `metafor::influence()` function), a recent standalone toolkit for the conduct of “comprehensive” sensitivity analyses has been proposed by Field and colleagues (2018), and bears noting here (Field, Bosco, Kepes, McDaniel, & List, 2018). Regardless of the means by which they are conducted, when considering sensitivity analyses, it is important to note that any concern about model sensitivity (e.g., outliers, publication bias) could apply within or across levels of any given moderator. Despite this, such analyses are rarely considered, perhaps owing to dwindling sample sizes when considering moderators (i.e., as suggested, moderator analyses often represent “low k ” cases).

Meta-analysis has long been bemoaned as a “garbage in, garbage out” procedure, with arguments that poor-quality primary studies lead to poor quality meta-analytic conclusions. Over

time, various metrics of study quality have emerged (e.g., an early review by Moher et al., 1995, found over 30 different quality scales/checklists that had been applied to meta-analyses of randomized control trials alone). Absent glaring statistical flaws (e.g., clearly misreported effect sizes), we would strongly caution against outright excluding studies on the basis of study quality indices. However, it is perfectly acceptable, and in some cases warranted, to conduct sensitivity analyses to understand how certain features of study quality (e.g., completeness of statistical reporting) would affect the conclusions one draws from meta-analytic results. To some degree, this suggestion overlaps with our recommendation to consider methodological moderators.

3.4. Apply Appropriate Corrections for Statistical Artefacts. All meta-analyses, regardless of tradition, account for sampling error as a statistical artefact. For example, in the Hedges-Olkin fixed effect tradition, this is typically achieved by weighting each observed effect by the inverse of its sampling variance. The logic underlying this approach is that the sampling variance for any given observed effect is a reasonable proxy for the precision of that estimate. Somewhat differently, the Schmidt-Hunter tradition achieves the sampling error corrections via a more straightforward sample-size weighting procedure. In many respects, this is a preferred method. Indeed, the problem with the inverse variance weighting procedure is that sampling variance is an estimate, and these estimates themselves are affected by sampling error. In contrast, the sample size of any given observed effect is typically a known value (e.g., Schmidt et al., 2009). To this end, recent simulation studies demonstrate that sample-size weights outperform sample-estimated inverse variance weights (e.g., Bakbergenuly, Hoaglin, & Kulinskaya, 2019a, 2019b). That said, owing largely to the correspondence between observed effect sampling variance and sample size (i.e., all else being equal, *larger* “N” studies will result in observed effect estimates with *smaller* sampling variances), in many cases these two methods yield very similar results (Marín-Martínez & Sánchez-Meca, 2010).

Additionally, meta-analyses may consider one or more statistical artefact corrections beyond corrections for sampling error. When considering artefact corrections beyond sampling error, it is common and indeed encouraged to additionally consider measurement error (i.e., artefacts associated with the (un)reliability of the predictor and/or the criterion variable), range restriction, enhancement, other selection biases (i.e., artefacts associated with various of processes that artificially, yet systematically, truncate or augment the observed variance of the predictor and/or the criterion variable; see Dahlke & Wiernik, 2019), and artificial dichotomization (e.g., artefacts associated with splitting otherwise continuous variables into “high” and “low” groups, such as the distinction between clinically “normal” and “disordered” behaviors). For a broader and more complete discussion of these and related issues concerning statistical artefact corrections, please refer to Wiernik and Dahlke (2019).

As suggested, Hedges and Olkin (1985, p. 135) give advice for managing statistical artefacts, however corrections beyond sampling error are less common when adopting this tradition of meta-analysis. Schmidt and Hunter (2015) likewise offer advice on how to manage statistical artefacts that are present across all effect size estimates (i.e., by applying individual corrections to study-level effect sizes) or when artefacts are only sparsely reported (i.e., by applying artefact distributions). In our experience, it is almost always necessary to apply artefact distribution procedures to make such corrections; it is all but guaranteed that there will be missing artefact information across the observed effects that are provided in empirical studies. How such corrections are made (e.g., correcting each effect individually versus using artefact distributions) is an important choice. Despite the fact that guidance is offered by Hedges and Olkin (1985, p. 135) regarding corrections for artefacts beyond measurement error, we rarely see such corrections carried out in practice. We would argue that, if one of the goals of meta-analysis

is to comprehensively account for statistical artefacts (and that perhaps it should be), it is more appropriate to choose Schmidt-Hunter-style artefact distribution procedures than other methods.

As suggested, meta-analytic methods in the Schmidt-Hunter tradition are well geared for synthesizing construct-level relationships, while concurrently accounting for statistical artefacts to accurately represent population parameters. That being said, the Hedges-Olkin tradition arguably provides a more flexible framework that can, for example, more readily account for different types of effect sizes than the Schmidt-Hunter tradition (e.g., odds ratios; note that Schmidt & Hunter, 2015, explicitly eschew the consideration of odd ratios, suggesting that they are “rarely appropriate” and “rarely used in social science research” [p. 249], despite the fact that such indices are commonly aggregated in meta-analyses of epidemiological research, for example). That being said, nearly every methodological and statistical advance in that has emerged from the Hedges-Olkin tradition could be implemented in the Schmidt-Hunter tradition, and vice-versa. The only limitation to such an integration as of now is the implementation in statistical software, however with recent advances in open source software, this gap is quickly closing (e.g., Dahlke & Wiernik’s, 2018 ‘psychmeta’ package for R, incorporates multiple predictor variables into a meta-regression model, with sample-size weighting, the Schmidt-Hunter variance estimator, and sampling error variances computed using the average effect size).

Given this, we do not think that it is necessary to favor one tradition of meta-analysis over the other (i.e., as both have their relative strengths and weaknesses; both may be more or less appropriately applied to answer different types of research questions). Indeed, the choice of meta-analytic tradition largely boils down to the type of research question one wishes to address. That said, corrections for statistical artefacts are most typically applied to the synthesis of observational (i.e., correlational) studies. However, Schmidt and Hunter (2015) offer clear guidance for applying such corrections to experimental and/or interventions studies as well.

Indeed, the reliability of such manipulations is an equally important concern; however, this issue has largely remained untouched in the literature, and especially so within meta-analyses of experimental/intervention studies (Schmidt & Hunter, 2015, p. 261; Wiernik & Dahlke, 2019).

3.5. Acknowledge and Remediate Non-Independence Among Effect Size Estimates.

Like more general instances of linear models, meta-analytic models place a strict requirement on the independence of observations that are associated with any given hypothesis test (i.e., any given calculation of an overall meta-analytic effect). Many circumstances can give rise to non-independence among observed effects. If the same subjects are used in multiple tests of the same hypothesis across different measures of the same criterion, then these tests are not independent from one another. Non-independence can be extrapolated in numerous additional ways as well. Examples include the nesting of common effects within studies (i.e., but across otherwise independent samples), or the common representation of effects across researchers (e.g., who commonly use particular methodologies, procedures, or scales to study any given phenomena; see discussions of so-called “hierarchical effects” within discussions about multilevel meta-analytic procedures; e.g., Cheung, 2014; Tanner-Smith, Tipton, & Polanin, 2016).

In more practical terms, the assumption of independence means that either only one effect size representing the combination of a given predictor and criterion or independent and dependent variable can be extracted from each study or that some provision to account for the nesting of effect sizes within studies must be built into one’s model. Over time, this has been achieved in a variety of ways (e.g., ignoring non-independence, coding only one effect size per study, averaging/compositing across dependent effect sizes; Cooper’s, 2017 “shifting unit of analysis model”; for a discussion, see Cheung, 2014).

In our view, if one is interested in representing construct-level relationships as “overall” effects (e.g., as is often the case in meta-analyses of nomological networks), then the approach

described below to form composites to represent hypothesis tests for different combinations of observed effects should be favored. However, meta-analysts often face situations where decisions about how to represent dependent effects are not so clear cut. For example, consider a primary study where the predictor variable is assessed at “baseline” (i.e., time one) and that the criterion variable is measured four times over the course of one year (i.e., once at baseline, and then every three months thereafter). Let us also assume that this is a situation where time lags are not of relevance to the research question being asked, for example, because few longitudinal studies exist within a research domain (see Card, 2019, for an example of how to integrate questions about time into meta-analyses). How should one represent these measures as a single observed effect? One option might be to consider only concurrent measures of the criterion measure; thus, only “time one” observed effects would be coded and included in the analysis. Another option would be to combine (e.g., through appropriate compositing procedures) the observed effects across all four time points. We explore similar ideas further below, when discussing the representation of between- versus within-person effects.

A relatively new approach to dealing with such non-independence is the application of multivariate meta-regression models (Hedges, Tipton, & Johnson, 2010a, 2010b), with cluster robust variance estimators to account for non-independent effect size estimates (Tipton & Pustejovsky, 2015). Recent simulation studies have shown that this approach performs “better” relative to other means of remediating non-independence, especially effect size averaging (see Moeyaert et al., 2016). Despite this, research has yet to consider how these procedures perform in comparison to the compositing procedures that are suggested by Schmidt and Hunter for representing construct-level relationships, and future research should address this gap.

Although promising, two points about multivariate meta-regression models with cluster robust variance estimators bear some consideration: First, this approach assumes that the

covariance structure underlying dependent effects can be approximated or assumed. That said, emerging evidence from simulation studies suggests that such approximations/assumptions are possible, while maintaining appropriate type-I error rates, even in small samples (e.g., Tipton & Pustejovsky, 2015). Second, such procedures have yet to be widely implemented for the Schmidt-Hunter random effects procedure. Although possible to implement, Schmidt and Hunter (2015, p. 389) have cautioned against the application of multilevel/multivariate models, citing that distortions due to non-independence are often of less import than those tied to statistical artefacts such as unreliability. Thus, there remains a philosophical disconnect between these two traditions on the point of how best to account for observed effect non-independence.

Philosophical differences aside, there is no reason why artefact corrections—either individually applied or based on artefact distributions—could not be incorporated into such multilevel/multivariate models. As alluded to before, like many choices among statistical procedures, one has to make decision about “what” sources of error variance (e.g., non-independence; measurement error) are most important to control for when representing the intended effect. Given that unreliability and other statistical artefacts represent knowable upper boundaries on validity, we would argue that in most cases it makes sense to prioritize accounting for statistical artefacts beyond sampling error alone, while additionally considering how best to incorporate strategies to account for non-independence.

3.6. Apply Appropriate Methods for the Computation of Composite Variables.

Often, researchers are interested in representing construct-level relationships in meta-analysis. For example, an occupational health psychologist may be interested in representing observed effects represented as correlations with the construct “employee burnout” (see Maslach & Jackson, 1984). Sometimes, primary studies represent the effects of burnout as an overall index, but in other cases, effects representing individual dimensions of burnout are reported (e.g., effects

representing relationships with emotional exhaustion and depersonalization, separately).

Complicating matters more, the researcher might also be interested in elaborating meta-analytic relationship for *both* the overall construct (i.e., burnout) and its sub-dimensions (i.e., emotional exhaustion and depersonalization).

A common strategy to represent construct-level relationships is to create some estimate of the overall construct through the process of composite formation. Composites can be created in a number of ways, for example by pooling correlation coefficients in observational studies, or by pooling treatment effects in experimental/intervention studies by weighting outcome means and pooling their standard deviations prior to computing observed effect size estimates (see Lipsey & Wilson, 2001). Here we focus on the composition of correlation coefficients, as this represents a common strategy for representing construct-level relationships in Schmidt-Hunter meta-analysis.

If there are, for example, two correlations representing some predictor correlating with sub-dimension measures of burnout, emotional exhaustion ($r_{xy} = .20$) and depersonalization ($r_{xy} = .30$), it might seem prudent to take the arithmetic average of these two relationships and conclude that the composite relationship between the predictor and burnout is simply:

$$r_{xY} = \frac{(.20 + .30)}{2} = .25$$

However, this logic ignores the fact that the two sub-dimensions of burnout are correlated with one-another to some extent, but not perfectly. So, let us assume that emotional exhaustion and depersonalization correlate with one-another at $r_{xy} = .55$. If we apply formula 10.6 from Schmidt and Hunter (2015, p. 442), then the proper composite correlation is:

$$r_{xY} = \frac{(.20 + .30)}{\sqrt{2 + 2(1).55}} = .28$$

Following the logic of this formula, only if the correlation between emotional exhaustion and depersonalization was $r_{xy} = 1.00$ would the arithmetic average be an appropriate estimate of the composite. While the difference between $r_{xY} = .25$ and $r_{xY} = .28$ may seem trivial, as long as

the variables that comprise a composite are less than perfectly correlated with one another, failing to properly composite construct-level relationships would yield systematic under-estimates of such construct-level relationships. Additionally, along with the application of proper compositing procedures, it is important to consider the appropriate reliability estimate to accompany the construct-level composites. Owing to the fact that the Spearman-Brown formula (Brown, 1910; Spearman, 1910) makes strong assumptions about relationships between component parts of the composite, Schmidt and Hunter (2015) recommend Mosier's (1943) composite reliability formula instead. Likewise, for meta-analyses considering range restriction artefacts, similar compositing procedures are required for U-ratios (see Schmidt & Hunter, 2015). The 'psychmeta' package for the R statistical computing environment (Dahlke & Wiernik, 2018) provides functions for computing proper composite effect sizes (i.e., r_{xy} , d) to represent construct-level relationships, appropriate composite estimates reliabilities (r_{xx} , r_{yy}), and composite U-ratios.

Finally, it is worth noting that primary studies often collect just one sub-dimension of an otherwise multidimensional scale (e.g., measures of emotional exhaustion *or* depersonalization, but not both). If such scales are to be included into an overall representation of the construct along with studies that consider both sub-dimensions simultaneously, it would be good practice to conduct a sensitivity analysis to “check” whether this researcher decision has any bearing on the conclusions drawn regarding this construct-level relationship. For example, this could be done to address a question like, “Does the overall representation of burnout depend on either the number or type of sub-dimensions of burnout that are represented in the composite?”

3.7. Explore the Possibility of Selection Effects (i.e., Publication Bias). In primary studies, the term “selection effect” is used to describe a mechanism that systematically precludes certain people from a broader population of interest from being representatively sampled. Range restriction is one familiar form of selection bias in such studies. Given that selection has the

possibility of impacting the broader generalizability of a given study's findings, a great deal of statistical and econometric research has been dedicated to understanding and remediating such effects (e.g., in the year 2000, James Heckman won a Nobel Prize in economics for developing a method for estimating and correcting self-selection effects; see Heckman, n.d.).

In meta-analysis, the concept of selection effects is understood similarly, in that the term is used to describe a mechanism that systematically precludes certain studies from a broader population of studies from being representatively sampled (e.g., located through exhaustive literature search efforts). Typically, the type of selection effect that one is most concerned about in meta-analysis is referred to as *publication bias*, or more classically as “the file drawer problem” (Rosenthal, 1979). The classic example of publication bias is tied to logic of null hypothesis significance testing: Studies that fail to reject the null hypothesis (i.e., those that fail to find a hypothesized effect at an *a priori* level of statistical significance, say $p < .05$) are less likely to be published than studies that do. What this means is that studies with affirmative conclusions about any given hypothesized effect are more likely to be present in a literature than those that do not find the expected effects. This is a problem for the meta-analyst, as their conclusions based upon published work alone would be more likely to reflect these affirmative conclusions. This issue is of particular concern for experimental/intervention studies, where the success or failure of any given study may rest on the efficacy of a manipulation/treatment, construed as a main effect. This issue is also a concern for observational studies; a salient example of such publication bias is outlined by McDaniel, Rothstein, and Whetzel (2006).

Given its implications for meta-analytic conclusions, a wide variety of methods have been developed to address publication bias (for a recent review and quantitative evaluation of such methods, see Carter, Schönbrodt, Gervais, & Hilgard, 2019). Each of these methods can be broadly classified into either detection methods or correction methods. Publication bias detection

methods aim to address the possibility that publication bias is present across the distribution of effect sizes considered for any given meta-analytic effect. An early example of a publication bias detection method is the so-called “failsafe N ” (Rosenthal, 1979). Failsafe N indexes the number of “missing studies” with a zero effect that, when added to the distribution of observed studies, would render the overall effect size not statistically significant. This method has been criticized for a number of reasons, including its overreliance of null hypothesis logic (e.g., Orwin, 1983). Moreover, across different methods, its application can result in varying estimates of the number of additional studies needed to nullify one’s effect (Becker, 2005). Accordingly, the general advice we would offer is that one should never use “failsafe N ” as an index of publication bias.

More contemporary publication bias detection methods hinge on observations and/or tests of funnel plot symmetry. Funnel plots are scatter plots that represent the joint distribution of effect sizes (on the X-axis) and an associated metric of their precision, such as their standard error (on the Y-axis). Such representations are termed “funnel plots,” because under the assumption that the precision of an estimated effect will increase as the size of the study (i.e., “ n ”) increases, effect estimates from “smaller n ” studies will scatter more widely at the bottom of the plot, with a narrowing of “larger n ” studies. Moreover, if there is no publication bias to speak of, this representation should approximate the shape of an (inverted) funnel. If, however, publication bias is present (e.g., because “smaller n ” studies with non-significant effects are precluded from the published literature), an asymmetrical funnel shape (e.g., one with a “bite” out of one of its lower corners, depending on the direction of the hypothesized effect) would be observed. Caution should be applied when applying “eyeball” tests to funnel plots, however, as it can be difficult to visually interpret such representations for either the presence or absence of publication bias (see Terrin, Schmid, & Lau, 2005).

Given shortcomings of “eyeball” tests of funnel plots, regression-based tests have also been offered to characterize the symmetry of funnel plots. For example, Egger, Smith, Schneider, and Minder (1997) proposed a regression-based test to gauge the symmetry of funnel plots. When doing “Egger’s test,” one performs an inverse variance-weighted least squares regression model, where effect size estimates are regressed onto their associated standard errors. The null hypothesis associated with the resulting slope parameter assumes no publication bias, which would manifest as a null relationship between effect size estimates and their standard errors. This hypothesis would be rejected if there was an appreciable association between effect size estimates and their standard errors (e.g., observing that “larger n ” studies have systematically smaller standard errors). Egger’s test can be criticized for assuming that the mechanism leading to publication bias is a function of sample size, that a symmetric funnel plot is to be expected (i.e., which may not hold true), and that deviations therefrom are indicative of publication bias.

Publication bias correction methods aim to address the possibility that publication bias is present, and to remediate its influence so as to arrive at meta-analytic estimates that are “corrected” for publication bias. One popular method for approximating this is “trim-and-fill” (Duval & Tweedie, 2000a, 2000b). Trim-and-fill is an approach that entails correcting for observed funnel plot asymmetry and proceeds in three iterative steps: First, “smaller n ” studies that lead to funnel plot asymmetry are “trimmed” from the overall distribution of effects. Second, the resulting trimmed effect size distribution is used to re-center the funnel plot. Third, the studies omitted in “step one” are used to “fill in” the missing effect sizes around the re-centered funnel plot. This procedure estimates both the number of missing studies (i.e., akin to the “failsafe N ”) and provides an adjusted overall effect size that is based on the distribution of effect sizes with the “missing” studies accounted for (i.e., imputed). Despite its promise, the trim-and-fill approach can be criticized on a number of grounds, similar to Egger’s test (e.g., assuming

publication bias is a function of “small sample” bias; that a symmetric funnel plot is to be expected). Moreover, this approach is known to perform poorly in the presence of effect size heterogeneity (see Peters et al., 2007; Terrin, Schmid, Lau, & Olkin, 2003).

We recommend that various assessments of selection effects be treated as sensitivity analyses. The idea is that one should triangulate evidence across multiple methods to arrive at a more holistically informed judgement about the possibility of selection effects. To this end, one promising method for detecting publication bias comes from a novel application of cumulative meta-analysis offered by McDaniel (2009). Cumulative meta-analysis is an iterative sensitivity analysis procedure, wherein studies are added one-by-one to a meta-analytic model to address the degree of “shift” in conclusions about the overall meta-analytic effect to particular studies. McDaniel (2009) proposes that effect sizes be sorted by their precision (e.g., high-to-low by their sample sizes) prior to entering them into a cumulative meta-analysis model. The cumulative estimates can then be iteratively examined for evidence of “drift” with the addition effect sizes.

The logic of this approach is straightforward: the meta-analytic effect resulting from studies entered into this model “earlier” represents more precise estimates of the population effect (e.g., because they represent “larger n ” studies). The meta-analytic effect observed “later” represents this effect with the addition of less precise estimates to the distribution that already captures more precise estimates. If less precise estimates (e.g., “smaller n ” studies) are being suppressed from the literature (i.e., a common theory of selection that would underscore the presence of publication bias), then one might observe that the cumulative effect “drifts” in a more positive direction as less precise estimates of the effect are added to the cumulative meta-analysis. This reflects the fact that those lower precision studies available in the literature tend to have higher estimates of the population effect than the higher precision estimates.

Given the uncertainty inherent in publication bias correction methods, we cannot universally recommend one method over another. However, in particular for experimental or intervention studies, where a strong theory of selection effects can be established *a priori*, Vevea and Hedges' (1995) weighted selection models may be considered, and are favorable in that they can be estimated even in the presence of heterogeneity—a noted shortcoming of the “trim-and-fill” method. Vevea and Hedges' (1995) model builds adjustments into the estimate of a meta-analytic effect that are derived from weights that represent pre-specified *p*-value intervals of interest. Thus, for example, one could test an assumption about whether studies with *p*-values greater than some threshold (e.g., $p < .05$) are differentially represented in the observed distribution of effect sizes. Such a model would particularly speak to a theory of selection based upon the logic of null hypothesis significance testing. This approach also makes corrections to the estimate of the meta-analytic effect on the basis of such weights and allows for a likelihood-ratio test comparing the fit of the corrected model to an uncorrected model. Thus, if the likelihood-ratio test is significant, this suggests that the corrected model is a better “fit” to the data and indicates that publication bias may be a concern.

3.8. Recognize Limits on Generalizability. Limits on the generalizability of inference are a topic of longstanding interest among methodologists and statisticians (e.g., Shadish et al., 2006). Concerns about threats to valid inference are likewise shared by meta-analysts (e.g., Matt & Cook, 2009). Of particular concern is the idea that threats to validity associated with any given primary study manifest in the aggregate when considered meta-analytically. For example, the generalizability of conclusions from primary studies is limited by a variety of factors, not least of which is the representativeness of the sample under consideration. Psychological research has long been criticized for an over-reliance on convenience samples that bear little correspondence to the broader population as a whole (Henrich, Heine, & Norenzayan, 2010; McNemar, 1946).

One of the advantages of meta-analytic synthesis is the ability to collapse across effects representing multiple different types of samples. The resulting estimate of the population effect is more broadly representative of the population as a whole than any given study considered in isolation. Although meta-analysis can help to bolster generalizability in this way, some important caveats around this idea must be understood.

Most notable among these caveats is a recognition that the types of samples included in a meta-analysis represent the sampling frame to which meta-analytic conclusions can be generalized. All else being equal, meta-analytic effects that are based on a relatively higher number of independent effects representing a diverse array of samples should better approximate the population parameter than those based on a relatively lower number of more homogenous samples. To understand the generalizability of a meta-analytic effect, it is therefore important for meta-analytic reports to provide summary statistics of study-level characteristics that typify relevant qualities of primary study participants. For example, Melchior, Bours, Schmitz, and Wittich (1997) present a small-scale meta-analysis ($k = 9$) of the correlates of burnout among psychiatric nurses. Given the narrow sampling frame for this effort, one would not necessarily expect these relationships to reflect the etiology of burnout, more broadly defined. There is still value in a narrow conceptualization like this; however, it is important to emphasize that the results of such efforts must be understood against such a conceptualization.

Considering the issue of generalizability further, the above discussion suggests that it is important to consider how many observed effects are “necessary” to conduct a meta-analytic review. This is an issue of statistical power, and one which has both statistical and logical underpinnings. Strictly speaking, assuming a fixed effect framework, the answer is $k = 2$. From a statistical perspective, the combination of $k = 2$ effects yields a more precise estimate of the true effect size than either effect size considered in isolation (e.g., Borenstein et al., 2011, p. 363).

Practically speaking, absent formal synthesis, researchers will attempt to apply some form of “cognitive algebra” to make sense of the two effects relative to one-another (Valentine, Piggot, & Rothstein, 2010). These ideas are more complex if we are to assume a random effects framework, in that when working with a low number of observed effects, it can be difficult to accurately represent the amount of between-study variance. Such estimates are integral to estimation of the overall effect and the inferences we make about it (i.e., as they have bearing on the estimation of the standard error for the overall effect). Therefore, both the point estimate and the confidence we put in that estimate are likely to reflect a higher-than-ideal amount of error (for extended discussions of statistical power in meta-analysis, see Hedges & Pigott, 2001; Pigott, 2012).

Finally, even with a “large k ” representing a diverse array of people, the accuracy of one’s meta-analytic database represents an important consideration and a distinct threat the validity of meta-analytic conclusions. As suggested by Shadish et al. (2002), sources of database inaccuracy can take multiple forms (e.g., coder “drift,” wherein a variety of factors lead coders to change their criteria for codes, such as practice effects, fatigue, etc.) but manifest as unreliability (i.e., inconsistency) in coding. Strategies for combatting these effects include employing multiple, independent coders, ensuring a clear and well-developed coding protocol, and having regular calibration and agreement meetings. For a more complete and general treatment of issues associated with interrater reliability and agreement, please refer to LeBreton and Senter (2008).

3.9. Properly Account for Between- and Within-Person Effect Sizes. Facets of the research designs employed within primary studies can have an important influence on the conduct of meta-analysis. One primary distinction in research designs that is often drawn is to distinguish between-person (e.g., “between subjects”, “interindividual”) from within-person (e.g., “within subjects”, “intraindividual”) research designs. Such designs can be employed for both experimental and observational methodologies. In experimental research designs, between-person

studies assign participants to one of perhaps many possible conditions, at random; within-person studies task participants with involvement in all conditions of a given study. In observational research designs, between-person studies collect measures of particular attributes once; within-person designs collect measures of such attributes more than once. While these fundamental distinctions may seem trivial, they have an important bearing on the conduct of meta-analyses, because the type of effect size estimates that are derived from each have important differences in their computation and interpretation that must be carefully considered.

Take, for example, a simple experimental design with a single independent variable with two levels (e.g., two levels of exposure to some treatment, “level one” and “level two”). It would be common practice to compute an effect size metric based upon a standardized mean difference between these two groups (e.g., d_{Cohen} or g_{Hedges}) using standard formulae (see Ellis, 2010 for such formulae). For example, d_{Cohen} can be simply defined as:

$$d_{Cohen} = \frac{M_1 - M_2}{\sigma}$$

where M_1 and M_2 define the means of the two “levels” described above, and σ is an estimate of the standard deviation, typically taken as the pooled standard deviation (i.e., assuming equal sample sizes, this would be the square root of the average variance noted across conditions).

This makes sense for the between-person case, however the standard formulae for computing standardized mean differences do not apply in the within-person case. Indeed, there are numerous ways to represent the denominator in the formula above, meaning that any number of different conceptualizations of within-person effect sizes are possible and yield results that are not directly comparable (please refer to Westfall, 2016, for computational examples).

Complicating matters more and depending on the nature of the research synthesis effort, the combination of effect sizes from between-person and within-person studies is challenging because they are not necessarily interpretable on the same metric (see Morris & DeShon, 2002).

Indeed, for repeated measures designs (e.g., pre-test/post-test/control group designs), the computation of a proper effect size metric of change is challenging, because the correlation between the measures over time (e.g., the correlations between “pre-test” and “post-test” scores) must be accounted for (see Morris, 2008). The calculation of an appropriate effect size therefrom is relatively straightforward; however, again complicating matters, it is often the case that such correlations are not provided and must be estimated. Overall, the consideration of within-person effect sizes is a complex issue marked both by flexibility on the part of the meta-analyst, and a lack of definite agreement about “what” is the proper index of such effects.

For observational studies, the effect size often taken is a correlation coefficient, representing some predictor–criterion relationship. In a between-person study, this correlation is straightforward to understand and code as long as it is reported in the primary research report. However, for within-person observational studies, it is not immediately clear what effect size to represent. For example, consider a hypothetical repeated measures observational study that adopts a complete two-wave fully crossed and time-lagged design to understand the relationship between a predictor, x , and criterion, y , variable. This study results in a 4x4 correlation matrix:

$$r_{xy} = \begin{bmatrix} r_{x1x1} & r_{x2x1} & r_{y1x1} & r_{y2x1} \\ r_{x1x2} & r_{x2x2} & r_{y1x2} & r_{y2x2} \\ r_{x1y1} & r_{x2y1} & r_{y1y1} & r_{y2y1} \\ r_{x1y2} & r_{x2y2} & r_{y1y2} & r_{y2y2} \end{bmatrix}$$

where, for example, r_{x1y2} defines the relationship between the predictor at time one, and the criterion at time two. Within the correlation matrix, there are four different effect size estimates of the predictor–criterion relationship (i.e., r_{x1y1} , r_{x2y1} , r_{x1y2} , r_{x2y2}). The question that must be addressed is, which of these relationships is the “best” estimate of the relationship of interest? By virtue of logic, one might preclude correlations where the outcome precedes the (theoretically assumed) antecedent in time (e.g., r_{x2y1}), yet that still leaves three possible relationships (i.e., two cross-sectional effects; one time-lagged effect) between the predictor and the criterion measures.

This issue becomes even more complicated with multi-wave designs (e.g., three or more waves of data collection), and is made worse by the fact that simple effects derived from such studies conflate between-person (i.e., stable, trait-like) and within-person (i.e., transient, state-like) variance (please refer to Hamaker, Kuiper, & Grasman, 2015 for an extended treatment of various issues associated with the conflation of between- and within-person sources of variance).

There are various ways to deal with this issue, and they each afford the analyst some flexibility in how to consider such relationships. For example, one might decide *a priori* to only code initial cross-sectional relationships (i.e., r_{xly1}) in such cases. Alternatively, if the goal of the meta-analysis itself is to quantify an over-time relationship (e.g., Nohe, Meier, Sonntag, & Michel, 2015), one might decide to only code the time-lagged effect (i.e., r_{xly2}). Still another way to deal with this is to code all relationships, and then collapse them either using a composite estimate (i.e., as discussed above; in doing so, one could also treat “lag” as a methodological moderator) or by adopting a multilevel model that accounts for the resulting non-independence among these observed effects. Multilevel models can also use robust variance estimates to correct the resulting standard errors for issues of clustering.

It also bears noting here that issues associated with the proper accounting of within- and between- effect size estimates can be abstracted beyond the “person” level of analysis (e.g., to within- vs. between-organization estimates; individual- vs. group-, organization-, or even nation-level relationships). For a consideration how such effects can be considered meta-analytically, please refer to Ostroff and Harrison (1999).

3.10. Recognize Limits on Causal Inferences. Meta-analysis is a powerful tool for research synthesis, but it is not a panacea. Because it relies on primary research studies to inform estimates of population effects, meta-analysis is beholden to the qualities of such research studies. Thus, the limits that various primary study designs place on one’s capacity to make

causal inferences also apply to meta-analytic summaries of such studies. A comprehensive discussion of the methodologies and logics of causal inference are well beyond the scope of this work, however please refer to Shadish et al. (2002) and Pearl (2009), respectively, for comprehensive details.

Here, we have made a rough distinction between experimental/intervention studies and observational studies, the former being better geared toward causal inference than the latter. To be clear, the goal of research, primary or otherwise, is not always to demonstrate causality, and it is by no means required for one to do so. However, for the meta-analyst, a recognition of the limitations of primary study research designs on one's meta-analytic conclusions bears additional consideration. For example, all else being equal, a meta-analysis of exclusively randomized control trials (RCTs) of a given phenomenon would instill more confidence regarding causality than the synthesis of quasi-experimental or observational studies of the same phenomenon.

Concerns about causal inference particularly plague the meta-analysis of observational studies, wherein it is typical to consider primarily cross-sectional (i.e., single time point, single source) correlations as effect size estimates. Additional confidence in the temporal precedence of (assumed) predictor–criterion relationships may be garnered by the consideration of prospective or time-lagged research designs (e.g., by treating timing of measurement as a methodological moderator). However, it should be recognized that temporal precedence is a necessary but insufficient condition for the inference of causality (e.g., Mackie, 1974).

4. Evaluating Meta-Analyses Published in JVB

With a clear sense of these ten important facets of meta-analysis, we next describe the methods used to obtain all meta-analyses published in JVB to-date and discuss the development of our coding protocol. Then, we report the results of our systematic review, looking at each of these ten facets of meta-analysis. Informed by these results, we additionally offer our “best

practices” recommendations. We then conclude with a discussion of these results, and ultimately a call for future research to adopt this prescriptive advice.

4.1 Method

4.1.1. Literature Search. To obtain the entire corpus of meta-analyses published in JVB, we focused our literature search exclusively on Elsevier’s ScienceDirect database. This was an ideal search strategy, because ScienceDirect comprehensively indexes all works published in the nearly 50-year history of the Journal. On March 26, 2019 we queried JVB abstracts in ScienceDirect using combinations of keywords that describe meta-analysis or related approaches to research synthesis (i.e., “meta-analysis” OR “meta analysis” OR “meta” OR “meta*” OR “validity generalization”) to obtain full texts of all studies. From this, we obtained $k = 85$ studies, spanning 32 years (1987 to 2019). Importantly, this search strategy also captured JVB “in press” and online-first publications. A full listing of all $k = 85$ studies considered in our systematic review is available in our online appendix: <https://osf.io/pgujx/>.

4.1.2. Inclusion and Exclusion Criteria. After obtaining full texts, two independent raters screened each article in full. Our primary inclusion criterion was that studies must use meta-analytic methods to synthesize primary empirical research. Thus, studies that do not use meta-analytic techniques (broadly defined) were excluded. For example, studies that mention “meta-analysis” in their abstracts, but do not employ meta-analytic synthesis in their methods, or studies using the phrase “meta” in another way (e.g., “meta-capacities”; Coetzee & Harry, 2014) were excluded. Likewise, we excluded studies that synthesize research using qualitative-systematic (as opposed to quantitative-systematic) methods (e.g., Eva et al., 2019). We also excluded studies that adopt meta-synthesis methods (i.e., the qualitative synthesis of qualitative research studies; Lazazzara, Tims, & de Gennaro, 2019). The application of these exclusion criteria led to a final sample of $k = 68$ meta-analytic studies (noted separated in our reference

section). Once this final database of all JVB meta-analyses was compiled, we systematically coded features of these studies, using an *a priori* developed coding protocol, described next.

4.1.3. Coding Protocol. As discussed above, our coding protocol was developed around ten facets of meta-analyses derived from the MARS standards and in consultation with related works that dictate contemporary “best practices” for the conduct of meta-analysis that have particular bearing on statistical conclusion validity. These ten facets are therefore based upon those methodological “best practices” as defined by MARS that are likewise reflected in recent recommendations for the conduct of meta-analyses, more generally (Kepes et al., 2013; Siddaway et al., 2019), as well as contemporary reference works that guide the design and conduct of meta-analyses (Borenstein et al., 2011; Cooper et al., 2009; Schmidt & Hunter, 2015).

To develop our coding protocol, we translated elements of the MARS and these other sources into more explicit criteria that we used to code each meta-analysis in our database. As suggested, these criteria were chosen explicitly to be the a) most *impactful*, inasmuch as they are likely to bolster the broader impact of any given meta-analysis, and b) most *actionable*, inasmuch as they can be readily implemented as “best practices” at the design, analysis, and/or reporting stages of any given meta-analysis. Please refer to Table 1 for a summary of these ten facets, along with related advice and additional resources.

With our protocol established, two independent coders (i.e., the second and third authors; Ph.D. students with experience conducting and critiquing meta-analyses) independently coded each article identified for inclusion from our database. These coders held regular meetings to ensure calibration and agreement; a third coder (i.e., the first author; a Ph.D.-level academic with 10+ years of experience conducting meta-analyses) helped rectify any disagreements encountered through discussions with the primary coding team. Our complete coding protocol, all data, and R code to replicate our analyses can be found on our project website: <https://osf.io/pgujx/>.

Regarding the consistency among coders, we calculated an index of the average percent agreement across all quantitative categories that were coded. Initially, the overall average percent agreement across these categories was 98.91%. All disagreements were discussed as a coding team in a formal consensus meeting, resulting in ultimately 100.00% agreement.

5. Results and Best Practices Recommendations

In the following, we present the results of our systematic literature review, organized around the ten facets of meta-analysis that guided this effort. Integrating these facets with the results of our review informs a set of 19 “best practices” that, if adopted, would serve to bolster valid inferences from meta-analytic data. Table 2 summarizes correlations among those coded methodological characteristics that correspond to all 68 meta-analyses included in our systematic review, whereas Table 3 provides a broader summary of the results of our review.

5.1. Consider Methodological and Substantive Moderators. The inclusion of methodological moderators is an important consideration for supporting the conclusions of meta-analyses. Testing the extent to which methodological variations across primary studies accounts for systematic variability in observed meta-analytic effects can help support the conclusions one draws. Our review found that 45.58% ($n = 31/68$) of meta-analyses published in JVB considered substantive moderators, whereas 47.06% ($n = 32/68$) considered methodological moderators. 50% of these studies reported more than one methodological moderator (i.e., 50%, $n = 16/32$). In all but one case (93.75%, $n = 15/16$) these multiple moderators were considered separately, as subgroup analyses; in one case, such moderators were entered simultaneously into a meta-regression model (6.26%, $n = 1/16$). Of those methodological moderators tested, the most common were related to study design (e.g., comparing cross-sectional versus longitudinal relationships), the types of predictor/criterion measures used, and publication status (e.g., published versus unpublished). To this end, the following best practice should be followed:

Meta-Analytic Best Practice #1: Anticipate and test methodological moderators.

5.2. Use Random Effects Estimators. Choices of estimators make a large difference in the interpretation of meta-analytic models of “true effects.” In nearly every case, random effects estimators should be favored. Additionally, it is important to make clear the amount of between-study heterogeneity that exists, and to quantify it appropriately. Our review found that 61.76% of meta-analyses ($n = 42/68$) adopted the tradition of Schmidt and Hunter, whereas 19.12% ($n = 13/68$) adopted the similar approach recommended by Raju, Burke, Normand, and Langlois, (1991). Thus, the majority (i.e., 80.79%) of meta-analyses published in JVB adopt some form of psychometric meta-analysis. Moreover, 19.12% ($n = 13/68$) adopted the tradition of Hedges and Olkin. Of the $n = 13/68$ (19.12%) meta-analyses that followed the Hedges and Olkin tradition, 61.54% (8/13) reported using a fixed effect estimator, whereas 38.46% (5/13) did not report what type of estimator was used.

Considering heterogeneity, our review found that 83.32% of meta-analyses ($n = 57/68$) published in JVB reported some estimate of the heterogeneity of effect size estimates. The most commonly reported indices of heterogeneity were $SD\rho$ ($n = 39$), %Var ($n = 23$), and Q_b ($n = 22$); of note, most studies reported more than one index of heterogeneity. Additionally, 58.88% ($n = 38/68$) of meta-analyses reported credibility intervals. To this end, the following best practices should be followed:

Meta-Analytic Best Practice #2: Use random effects estimators.

Meta-Analytic Best Practice #3: Report point estimates of between-study heterogeneity and credibility intervals.

Meta-Analytic Best Practice #4: Interpret estimates of heterogeneity cautiously, and in concert with true effect estimates, associated credibility intervals, and assumptions about the distribution of effects in the population.

5.3. Conduct Sensitivity Analyses. Sensitivity analyses can provide evidence about the tenability of researcher decisions and the robustness of one's meta-analytic conclusions to statistical concerns. Whenever possible, researchers should make their decisions clear, and if possible consider their potential influence *a priori*. Absent this, and recognizing that many such decisions occur during the coding process, researchers should conduct sensitivity analyses whenever possible to support their choices. Our review found that 10.20% ($n = 7/68$) of meta-analyses in JVB conducted explicit sensitivity analyses (*n.b.* we coded separately for selection effects/publication bias analyses). The most commonly reported sensitivity analysis was of influential/outlier cases. To this end, the following best practice should be followed:

Meta-Analytic Best Practice #5: Conduct sensitivity analyses to demonstrate the robustness of meta-analytic conclusions to researcher decisions.

5.4. Apply Appropriate Corrections for Statistical Artefacts. The choice to consider corrections for statistical artefacts above-and-beyond sampling error is an important consideration in meta-analyses of both observational and experimental/intervention studies. That said, we cannot readily think of a reason why one would *not* want to make such corrections, even if information about such artefacts was only sparsely available. Our review found that 80.88% ($n = 55/68$) of meta-analyses in JVB corrected for statistical artefacts beyond sampling error. The most common artefacts corrected for were predictor ($n = 49$) and/or criterion ($n = 43$) unreliability. Despite this, a wide range of potential statistical artefacts may be accounted for meta-analytically (e.g., range restriction, dichotomization of variables). Fewer studies found in our review corrected for other such artefacts (e.g., only $n = 4$ studies accounted for range restriction; $n = 1$ study accounted for dichotomization). When studies corrected for artefacts, 61.81% ($n = 34/55$) corrected effect sizes individually, whereas 29.09% ($n = 16/55$) reported using artefact distributions. Additionally, 9.09% of studies ($n = 5/55$) did not clearly report the

method used to make corrections. We also noted that 63.64% ($n = 35/55$) of studies reported both corrected and uncorrected effect size estimates. With these observations, we therefore make the following best practice recommendations:

Meta-Analytic Best Practice #6: Consider appropriate corrections for statistical artefacts.

Meta-Analytic Best Practice #7: Report corrected and uncorrected meta-analytic estimates.

5.5. Acknowledge and Remediate Non-Independence Among Effect Size Estimates.

Non-independence of observed effect size estimates is a longstanding concern in meta-analysis, regardless of the particular meta-analytic tradition followed. The particular way in which non-independence is addressed represents a choice that researchers make when conducting meta-analyses. Because there are a number of ways of managing non-independence, the particular strategy adopted should be made explicit in reporting one's meta-analytic methodology. In considering choices among such approaches, it is important to understand their relative strengths and limitations, and to be cognizant of the sources of error that one is most interested in managing. Our review found that 57.35% ($n = 39/68$) of meta-analyses explicitly acknowledged the issue of non-independence of effect size estimates. To remediate this concern, 92.31% ($n = 36/39$) of studies suggested that they coded only one effect/study. No studies identified through our review adopted a multilevel/multivariate model with robust variance estimation to account for nesting/non-independence. Thus, we offer the following best practice recommendation:

Meta-Analytic Best Practice #8: Recognize issues related to, and describe strategies for remediating, non-independence of effect size estimates.

5.6. Apply Appropriate Methods for the Computation of Composite Variables. One common goal of meta-analyses is to represent construct-level relationships. If this is the case, then it is important to use appropriate formulae for the construction of composite effects. Failure to do so will result in misestimates of the intended parameter. Our review found that 20.59% ($n =$

14/68) of meta-analyses report some efforts to create composite construct-level variables. 57.14% ($n = 8/14$) of meta-analyses that report composite relationships use formulae from Schmidt and Hunter, whereas 21.43% ($n = 3/14$) report simply averaging across effect size estimates. Of these studies, 7.14% ($n = 1/14$) report creating composite reliability estimates to inform artefact corrections. From this study, it was not clear whether formulae from Spearman-Brown (Brown, 1910; Spearman, 1910) or Mosier (1943) were applied. Thus, we offer the following best practice recommendation:

Meta-Analytic Best Practice #9: Use appropriate compositing procedures to represent construct-level relationships and associated reliabilities.

5.7. Explore the Possibility of Selection Effects (i.e., Publication Bias). Selection effects such as publication bias can unduly influence the conclusions of meta-analyses and should be estimated and explored. Although no method is best *per se*, treating tests of publication bias as sensitivity analyses can bolster confidence in one's conclusions. Importantly, beyond publication bias, other types of selection effects have an influence on meta-analysis. For example, meta-analyses that set reporting language inclusion criteria (e.g., that all studies must be reported in English) are automatically building selection effects into their conclusions. Moreover, as will be discussed next, such criteria put limits on the broader generalizability of one's findings. Our review found that 30.88% ($n = 21/68$) of meta-analyses report some analysis of selection effects tied to publication bias. Of those, the most common analyses conducted were failsafe N and funnel plot asymmetry/regression-based methods. The more proactive course of action to remediate publication bias is to take all possible steps to identify unpublished literature. Our review found that 88.24% ($n = 60/68$) of meta-analyses report some attempts to uncover such literature. Most commonly, these attempts come from including dissertations/theses ($n = 49$),

author contacts ($n = 20$), and searching conference programs ($n = 13$). Accordingly, we offer the following best practice recommendations:

Meta-Analytic Best Practice #10: Develop a theory of selection effects and test it.

Meta-Analytic Best Practice #11: Triangulate the results of multiple tests of publication bias as sensitivity analyses.

Meta-Analytic Best Practice #12: Be cautious of correction-based methods for addressing publication bias.

5.8. Recognize Limits on Generalizability. There are a wide variety of issues that serve as threats to the validity of meta-analytic conclusions and consequently represent barriers to the broader generalizability of such findings. Such effects are bound, among others, by the type of primary studies under consideration, the number of studies considered, and how accurately relevant information from them is representing in a meta-analytic database. Our review found that 33.82% ($n = 23/68$) of meta-analyses report some information about the samples under investigation, with most studies most commonly reporting information about sample demographics ($n = 15$), countries ($n = 10$), and job/industry types ($n = 9$). Additionally, 73.91% ($n = 17/23$) of these meta-analyses reported using such sample characteristics as moderators of overall effect size estimates. Our review also found that, across meta-analyses reporting such figures, the median “minimum k ” (i.e., the minimum number of studies deemed appropriate to elaborate on a meta-analytic relationship) was $k = 3$ ($M = 9.31$, $SD = 13.80$). The median overall number of independent effects reported was $k = 341$ ($M = 520.61$, $SD = 586.42$), and the median overall sample size was $n = 31,586$ ($M = 89,399.81$, $SD = 156,597.20$). Finally, our review found that 60.29% ($n = 41/68$) of meta-analyses report using multiple coders, and of these 80.49% ($n = 34/41$) report some evidence of intercoder agreement. When coding disagreements were found, the most common strategy reported for reconciling such issues was through discussions among

the coding team (70.73%, $n = 29/41$). Accordingly, we offer the following best practice suggestions:

Meta-Analytic Best Practice #13: Report study-level characteristics and be careful not to overgeneralize meta-analytic findings beyond such characteristics.

Meta-Analytic Best Practice #14: Make explicit the minimum number of studies necessary to elaborate on a meta-analytic effect.

Meta-Analytic Best Practice #15: When possible, use multiple coders and check their agreement.

5.9. Properly Account for Between- and Within-Person Effect Sizes. It is important for meta-analysts to acknowledge the issues that are inherent in considering between-person and within-person estimates of effect size. That said, there is no single prescription for “how” to manage these various effects, and the “best” strategy to employ is likely to differ depending on the nature of one’s meta-analytic study. In preparing a meta-analysis, it is important to recognize issues involved with effect sizes derived from different types of research designs, and to be very clear with reporting how they are dealt with. The flexibility that can be applied to this end is a double-edged sword—on the one hand, there are a variety of means to deal with this issue; on the other hand, different conclusions could be reached depending on one’s choices in this matter. Thus, it would be best to establish *a priori* one’s intentions for how to manage this issue. Our review found that 92.65% ($n = 63/68$) of meta-analyses reported in JVB focus on observational studies, whereas 7.35% ($n = 5/68$) focus on (quasi)-experimental/intervention studies. Among those reporting the synthesis of (quasi)-experimental/intervention studies, we found that 20% ($n = 1/5$) provided an explicit consideration of the treatment of effect sizes from within-subjects designs versus between-subjects designs. Among those reporting the synthesis of observational studies we found that 19.04% ($n = 12/63$) provided an explicit consideration of the treatment of

effect sizes from within-subjects designs versus between-subjects designs. Accordingly, we make the following best practice recommendations:

Meta-Analytic Best Practice #16: Be cognizant of differences in the computation of between-person versus within-person effect sizes.

Meta-Analytic Best Practice #17: Make decisions about managing between-person versus within-person effect sizes transparent.

5.10. Recognize Limits on Causal Inferences. There is very little that a meta-analyst can do to change how any given phenomenon is investigated in primary studies; in many respects, they are beholden to the literature that is available to be synthesized. Meta-analysis has been criticized for its “chilling effect” on research literatures (i.e., that it can serve as a be-all, end-all summary of a body of work, which stifles future innovation in a given topic area; e.g., Judge, Thoresen, Bono, & Patton, 2001). On the contrary, if meta-analysts more readily recognize the limits of the inferences gleaned from a given synthesis, then meta-analysis becomes a powerful tool to call for “better” research. Whether one’s goal is to index causal relationships, or to provide a descriptive accounting of a literature, making such aims clear and explicit is important for the broader interpretation of any given meta-analysis. Our review found that 23.81% ($n = 15/62$) of meta-analyses of observational studies reported in JVB provided some explicit recognition of the limitations of causal inferences from the synthesis of observational studies. Moreover, across all meta-analyses, 50.00% ($n = 34/68$) recognized the limitations of cross-sectional research designs to this end. Finally, 20.59% ($n = 14/68$) of studies reported separate analyses of prospective studies (i.e., those that separate the measurement of the predictor and criterion variable in time). Accordingly, we make the following “best practice” recommendations:

Meta-Analytic Best Practice #18: Acknowledge the types of research designs available to be synthesized.

Meta-Analytic Best Practice #19: Recognize the limitations of primary studies for informing causal claims from meta-analytic synthesis.

6. Discussion

6.1. Summary of Findings. As suggested in the introduction, two research questions guided this effort. The answer to the first, “How are meta-analyses published in JVB ‘done’?,” can be gleaned from selected results of our systematic review. Indeed, the prototypical meta-analysis published in JVB tends to be based on a) cross-sectional, b) observational studies, and uses c) Schmidt-Hunter methods to make d) artefact corrections for predictor and criterion unreliability. The answer to the second question, “Do meta-analyses published in JVB conform to ‘best practices?’,” is somewhat more nuanced, but can also be gleaned from selected results of our systematic review. In some respects, meta-analyses published in JVB do generally conform to “best practices.” For example, the fact that a majority of papers identified in our review correct for statistical artefacts is promising. However, we also observed that those papers that do make such corrections tend to apply individual corrections rather than those based on artefact distributions. Moreover, the observation that most studies make efforts to locate unpublished studies is likewise promising. However, the fact that most studies do not otherwise assess for the presence of publication bias, and that those that do tend to use inferior methods, is concerning. The correlations in Table 2 suggest that methodological features of meta-analyses are not necessarily related to one-another. Moreover, such features are largely unrelated to impact, however this conclusion should be interpreted cautiously in light of the negative correlation between publication year and impact (see also Figure 2). Below, we provide some thoughts on the conduct of meta-analyses, structured around more general recommendations for meta-analysts to embrace open science practices in their work. Then, we discuss some of the limitations of our

review, and offer some thoughts on the “future” of meta-analytic work in the field of vocational behavior.

6.2. General Recommendations. Beyond the 19 “best practices” offered here, there are some more general recommendations for the conduct of meta-analysis that bear some consideration. Most notably, as we have discussed, meta-analysis is a methodology that is marked by a number of decision points—choices that the researcher must make about how to conduct their research synthesis. In general, meta-analysis proceeds in a systematic way, however such decisions open the process up to criticism, particularly if they are not well justified or it is not clear whether they were made *a priori*. Thus, our general advice to this end is for meta-analysts to embrace the open science zeitgeist, and pre-register such decisions to the extent that it is possible (Nosek et al., 2015; see Kleine, Rudolph, & Zacher, 2019, for an example). Indeed, many of the more general issues we raise here regarding the choices meta-analysts make could be addressed if such decisions were pre-registered, that is to say, if such decisions were codified up front, and before data were collected, coded, and analyzed. To this end, there are existing platforms that allow for the pre-registering of systematic reviews, including meta-analyses (e.g., Open Science Framework: osf.io; PROSPERO: crd.york.ac.uk/prospero). We would strongly encourage researchers to consider pre-registration of meta-analyses as a means of realizing our recommendations to make meta-analytic decisions more transparent.

Taking the concept of pre-registration further, editors and journals should consider adopting a registered reports format, to allow for meta-analytic protocols to be peer-reviewed and accepted “in principle.” Beyond pre-registration, emerging philosophies of open science have additional bearing on the process of conducting meta-analysis. For example, the sharing of analysis code and data (e.g., code sheets) would serve the literature well by allowing for fully

replicable analyses. We would also strongly encourage editors and journals to consider requiring such information be provided explicitly.

6.3. Limitations and Future Directions. No effort at systematically reviewing a literature is without some limitations. Based on the goals of our review, we have purposefully limited the scope of our efforts to only meta-analyses published in JVB. Given that the issues identified here are likely to have bearing on meta-analyses published elsewhere, we encourage researchers to adopt a similar perspective on meta-analyses within other literatures. Such efforts would allow for one to take stock of whether “best practices” for the conduct of such efforts are being adopted, and if not, to call for more attention to such issues. Another (somewhat ironic) critique of this manuscript is that we focused exclusively on published meta-analyses. Thus, our work here reflects some degree of systematic selection akin to survivorship bias; all papers considered here were vetted through peer review and/or editorial guidance. The question this raises is, what would non-published (or pre-review “first drafts”) meta-analyses look like if reviewed similarly? Finally, we developed our coding profile on the basis of established best practice recommendations and resources. We acknowledge, if faced with the same task and resources, that others may have chosen different areas of focus; however, we would also suggest that each of the areas we argue for here are of relevance to a wide variety of meta-analytic efforts.

6.4. Conclusions. By “taking stock” of the current state of meta-analyses published in JVB, this work serves as a benchmark for future research that considers applications of meta-analytic methods. Given that meta-analysis holds a top position in the “hierarchy of evidence” (e.g., Greenhalgh, 1997), we hope this work will serve as a guide for future researchers who are planning and conducting meta-analyses on important vocational behavior topics. Thus, we present this work as a formal call for future meta-analyses on such topics to pay closer attention to these “best practices.”

7.0 References

- Aguinis, H., Dalton, D. R., Bosco, F. A., Pierce, C. A., & Dalton, C. M. (2011). Meta-analytic choices and judgment calls: Implications for theory building and testing, obtained effect sizes, and scholarly impact. *Journal of Management*, 37, 5–38.
doi:10.1177/0149206310377113
- Aguinis, H., Gottfredson, R. K., & Wright, T. A. (2011). Best-practice recommendations for estimating interaction effects using meta-analysis. *Journal of Organizational Behavior*, 32, 1033–1043. doi:10.1002/job.719
- American Psychological Association. (2008). Reporting standards for research in psychology: Why do we need them? What might they be? *American Psychologist*, 63, 839–851.
doi:10.1037/0003-1066X.1063.1039.1839.
- American Psychological Association. (2010). *Publication manual of the American psychological association* (6th ed.). Washington, DC: American Psychological Association.
- Assouline, M., & Meir, E. I. (1987). Meta-analysis of the relationship between congruence and well-being measures. *Journal of Vocational Behavior*, 31, 319–332. doi:10.1016/0001-8791(87)90046-7
- Bakbergenuly, I., Hoaglin, D. C., & Kulinskaya, E. (2019a). Simulation study of estimating between-study variance and overall effect in meta-analysis of odds-ratios. arXiv preprint arXiv:1902.07154.
- Bakbergenuly, I., Hoaglin, D. C., & Kulinskaya, E. (2019b). Simulation study of estimating between-study variance and overall effect in meta-analysis of standardized mean difference. arXiv preprint arXiv:1903.01362.
- Becker, B. J. (2005). Failsafe N or file-drawer number. In H. R. Rothstein, A. J. Sutton, & M.

- Borenstein (Eds), *Publication bias in meta-analysis* (pp. 111–126). Chichester, UK: John Wiley & Sons.
- Becker, T. E., Atinc, G., Breugh, J. A., Carlson, K. D., Edwards, J. R., & Spector, P. E. (2016). Statistical control in correlational studies: 10 essential recommendations for organizational researchers. *Journal of Organizational Behavior*, 37, 157–167. doi:10.1002/job.2053
- Berger, J. O., & Berry, D. A. (1988). Statistical analysis and the illusion of objectivity. *American Scientist*, 76, 159–165.
- Borenstein, M., Hedges, L. V., Higgins, J. P. T., & Rothstein, H. R. (2010). A basic introduction to fixed effect and random effects models for meta-analysis. *Research Synthesis Methods*, 1, 97–111. doi:10.1002/jrsm.12
- Borenstein, M., Hedges, L. V., Higgins, J. P., & Rothstein, H. R. (2011). *Introduction to meta-analysis*. Hoboken, NJ: John Wiley & Sons.
- 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, 490–514. doi:10.1177/1094428117741966
- Brown, W. (1910). Some experimental results in the correlation of mental abilities. *British Journal of Psychology*, 3, 296–322. doi:10.1111/j.2044-8295.1910.tb00207.x
- Card, N. A. (2019). Lag as moderator meta-analysis: A methodological approach for synthesizing longitudinal data. *International Journal of Behavioral Development*, 43, 80–89. doi:10.1177/0165025418773461
- Carter, E. C., Schönbrodt, F. D., Gervais, W. M., & Hilgard, J. (2019). Correcting for bias in

- psychology: A comparison of meta-analytic methods. *Advances in Methods and Practices in Psychological Science*, 2, 115–144. doi:10.1177/2515245919847196
- Cheung, M. W.-L. (2014). Modeling dependent effect sizes with three-level meta-analyses: A structural equation modeling approach. *Psychological Methods*, 19, 211–229. doi:10.1037/a0032968
- Coetzee, M., & Harry, N. (2014). Emotional intelligence as a predictor of employees' career adaptability. *Journal of Vocational Behavior*, 84, 90–97. doi:10.1016/j.jvb.2013.09.001
- Cooper, H. (2017). *Research synthesis and meta-analysis: A step-by-step approach* (5th edition). Thousand Oaks, CA: SAGE Publications.
- Cooper, H., Hedges, L. V., & Valentine, J. C. (Eds.). (2009). *The handbook of research synthesis and meta-analysis*. New York, NY: Russell Sage Foundation.
- Dahlke, J. A., & Wiernik, B. M. (2018). psychmeta: An R package for psychometric meta-analysis. *Applied Psychological Measurement*. doi:10.1177/0146621618795933.
- Dahlke, J. A., & Wiernik, B. M. (2019). Not restricted to selection research: Accounting for indirect range restriction in organizational research. *Organizational Research Methods*. doi:10.1177/1094428119859398 [In Press Accepted Manuscript]
- DerSimonian R. & Laird N. (1986). Meta-analysis in clinical trials. *Controlled Clinical Trials*, 7, 177–188. doi:10.1016/0197-2456(86)90046-2
- Duval, S., & Tweedie, R. (2000a). A nonparametric "trim and fill" method of accounting for publication bias in meta-analysis. *Journal of the American Statistical Association*, 95, 89–98. doi:10.1080/01621459.2000.10473905
- Duval, S., & Tweedie, R. (2000b). Trim and fill: A simple funnel-plot-based method of testing and adjusting for publication bias in meta-analysis. *Biometrics*, 56, 455–463. doi:10.1111/j.0006-341x.2000.00455.x

- Egger, M., Smith, G. D., Schneider, M., & Minder, C. (1997). Bias in meta-analysis detected by a simple, graphical test. *BMJ*, 315, 629–634. doi:10.1136/bmj.315.7109.629
- Ellis, P. D. (2010). *The essential guide to effect sizes: Statistical power, meta-analysis, and the interpretation of research results*. New York, NY: Cambridge University Press.
- Eva, N., Newman, A., Jiang, Z., & Brouwer, M. (2019). Career optimism: A systematic review and agenda for future research. *Journal of Vocational Behavior*. [In Press Accepted Manuscript]. doi:10.1016/j.jvb.2019.02.011
- Field, J., Bosco, F. A., Kepes, S., McDaniel, M. A., & List, S. L. (2018). Introducing a comprehensive sensitivity analysis tool for meta-analytic reviews. Paper presented at the 78th Annual Conference of the Academy of Management, Chicago, IL.
- Fisher, R.A. (1925). *Statistical methods for research workers*. Edinburgh: Oliver and Boyd.
- Gelman, A. (2005, January 25). Why I don't use the term "fixed and random effects." [Blog post]. Retrieved from: https://statmodeling.stat.columbia.edu/2005/01/25/why_i_dont_use
- Geyskens, I., Krishnan, R., Steenkamp, J. B. E. M., & Cunha, P. V. (2009). A review and evaluation of meta-analysis practices in management research. *Journal of Management*, 35, 393–419. doi:10.1177/0149206308328501
- Glass G. V. (1976). Primary, secondary, and meta-analysis of research. *Educational Researcher*, 5(10), 3–8. doi:10.3102/0013189X005010003.
- Gough, D., Oliver, S., & Thomas, J. (Eds.). (2017). *An introduction to systematic reviews*. Thousand Oaks, CA: SAGE.
- Greenhalgh, T. (1997). How to read a paper: Getting your bearings (deciding what the paper is about). *British Medical Journal*, 315, 243–246. doi:10.1136/bmj.315.7102.243
- Hamaker, E. L., Kuiper, R. M., & Grasman, R. P. (2015). A critique of the cross-lagged panel

- model. *Psychological Methods*, 20, 102–116. doi:10.1037/a0038889
- Harari, M., Parola, H.R., Hartwell, C.J., Riegelman, A. (202). Literature searches in systematic reviews and meta-analyses: A review, evaluation, and recommendations. *Journal of Vocational Behavior*. [In Press Accepted Manuscript]
- Heckman, J. J. (n.d.). Retrieved from: <https://www.econlib.org/library/Enc/bios/Heckman.html>
- Hedges, L.V. & Olkin, I. (1985). *Statistical methods for meta-analysis*. Orlando, FL: Academic Press.
- Hedges, L. V., & Pigott, T. D. (2001). The power of statistical tests in meta-analysis. *Psychological Methods*, 6, 203–217. doi:10.1037//1082-989X.6.3.20
- Hedges, L. V., Tipton, E., & Johnson, M. C. (2010a). Robust variance estimation in meta-regression with dependent effect size estimates. *Research Synthesis Methods*, 1(1), 39-65. doi: 10.1002/jrsm.5
- Hedges, L. V., Tipton, E., & Johnson, M. C. (2010b). Erratum: Robust variance estimation in meta-regression with dependent effect size estimates. *Research Synthesis Methods*. doi.org/10.1002/jrsm.17
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). Most people are not WEIRD. *Nature*, 466, 29. doi:10.1038/466029a
- Higgins, J. P., 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, 137–159. doi:10.1111/j.1467-985X.2008.00552.x
- Hunter, J. E., & Schmidt, F. L. (2000). Fixed effects vs. random effects meta-analysis models:

- Implications for cumulative research knowledge. *International Journal of Selection and Assessment*, 8, 275–292. doi:10.1111/1468-2389.00156
- Inthout, J., Ioannidis, J. P., Rovers, M. M., & Goeman, J. J. (2016). Plea for routinely presenting prediction intervals in meta-analysis. *BMJ Open*, 6, e010247, 1–5. doi:10.1136/bmjopen-2015-010247
- Judge, T. A., Thoresen, C. J., Bono, J. E., & Patton, G. K. (2001). The job satisfaction–job performance relationship: A qualitative and quantitative review. *Psychological Bulletin*, 127, 376–407. doi:10.1037//0033-2909.127.3.376
- Kepes, S., McDaniel, M. A., Brannick, M. T., & Banks, G. C. (2013). Meta-analytic reviews in the organizational sciences: Two meta-analytic schools on the way to MARS (the Meta-Analytic Reporting Standards). *Journal of Business and Psychology*, 28, 123–143. doi:10.1007/s10869-013-9300-2
- Kleine, A. K., Rudolph, C. W., & Zacher, H. (2019). Thriving at work: A meta-analysis. *Journal of Organizational Behavior*, 40, 973–999. doi:10.1002/job.2375
- Knight, G. P., Fabes, R. A., & Higgins, D. A. (1996). Concerns about drawing causal inferences from meta-analyses: An example in the study of gender differences in aggression. *Psychological Bulletin*, 119, 410–421. doi:10.1037/0033-2909.119.3.410
- Langan, D., Higgins, J. P., Jackson, D., Bowden, J., Veroniki, A. A., Kontopantelis, E., Viechtbauer, W., & Simmonds, M. (2019). A comparison of heterogeneity variance estimators in simulated random-effects meta-analyses. *Research Synthesis Methods*, 10, 83–98. doi:10.1002/jrsm.1316
- Lazazzara, A., Tims, M., & de Gennaro, D. (2019). The process of reinventing a job: A meta–

- synthesis of qualitative job crafting research. *Journal of Vocational Behavior* [In Press Accepted Manuscript]. doi:10.1016/j.jvb.2019.01.001
- LeBreton, J. M., & Senter, J. L. (2008). Answers to 20 questions about interrater reliability and interrater agreement. *Organizational Research Methods, 11*, 815–852.
doi:10.1177/1094428106296642
- Lipsey, M.W. & Wilson, D.B. (2001). *Practical meta-analysis*. Thousand Oaks, CA: Sage.
- Mackie, J. L. (1974). *The cement of the universe: A study of causation*. Oxford, UK: Clarendon Press.
- Marín-Martínez, F., & Sánchez-Meca, J. (2010). Weighting by inverse variance or by sample size in random-effects meta-analysis. *Educational and Psychological Measurement, 70*, 56–73. doi:10.1177/0013164409344534
- Maslach, C. & Jackson, S. E. (1984). Burnout in organizational settings. *Applied Social Psychology Annual, 5*, 133–153.
- Matt, G.E., & T.D. Cook. (2009). Threats to the validity of generalized inferences. In H. Cooper, L.V. Hedges, and J.C. Valentine (Eds.) *The handbook of research synthesis and meta-analysis* (pp. 537–560). New York, NY: Russell Sage Foundation.
- McDaniel, M.A. (2009). Cumulative meta-analysis as a publication bias method. Paper presented at the 24th Annual Meeting of the Society for Industrial and Organizational Psychology. New Orleans, LA (USA).
- McDaniel, M. A., Rothstein, H. R., & Whetzel, D. L. (2006). Publication bias: A case study of four test vendors. *Personnel Psychology, 59*, 927–953. doi:10.1111/j.1744-6570.2006.00059.x
- McNemar, Q. (1946) Opinion-attitude methodology. *Psychological Bulletin, 43*, 289–374.
doi:10.1037/h0060985

- Melchior, M. E. W., Bours, G. J. J. W., Schmitz, P., & Wittich, Y. (1997). Burnout in psychiatric nursing: A meta-analysis of related variables. *Journal of Psychiatric and Mental Health Nursing*, 4, 193–201. doi:10.1046/j.1365-2850.1997.00057.x
- Meyer, J. P., Stanley, D. J., Herscovitch, L., & Topolnytsky, L. (2002). Affective, continuance, and normative commitment to the organization: A meta-analysis of antecedents, correlates, and consequences. *Journal of Vocational Behavior*, 61, 20–52. doi:10.1006/jvbe.2001.1842
- Moeyaert, M., Ugille, M., Natasha Beretvas, S., Ferron, J., Bunuan, R., & Van den Noortgate, W. (2017). Methods for dealing with multiple outcomes in meta-analysis: A comparison between averaging effect sizes, robust variance estimation and multilevel meta-analysis. *International Journal of Social Research Methodology*, 20, 559–572. doi:10.1080/13645579.2016.1252189
- Moher, D., Jadad, A. R., Nichol, G., Penman, M., Tugwell, P., & Walsh, S. (1995). Assessing the quality of randomized controlled trials: An annotated bibliography of scales and checklists. *Controlled Clinical Trials*, 16, 62–73. doi:10.1016/0197-2456(94)00031-W
- Morris, S. B. (2008). Estimating effect sizes from pretest-posttest-control group designs. *Organizational Research Methods*, 11, 364–386. doi:10.1177/1094428106291059
- Morris, S. B., & DeShon, R. P. (2002). Combining effect size estimates in meta-analysis with repeated measures and independent-groups designs. *Psychological Methods*, 7, 105–125. doi:10.1037//1082-989X.7.1.105
- Mosier, C. I. (1943). On the reliability of a weighted composite. *Psychometrika*, 8, 161–168. doi:10.1007/bf02288700

- Nieminen, L. R., Nicklin, J. M., McClure, T. K., & Chakrabarti, M. (2011). Meta-analytic decisions and reliability: A serendipitous case of three independent telecommuting meta-analyses. *Journal of Business and Psychology*, 26, 105–121. doi:10.1007/s10869-010-9185-2
- Nohe, C., Meier, L. L., Sonntag, K., & Michel, A. (2015). The chicken or the egg? A meta-analysis of panel studies of the relationship between work–family conflict and strain. *Journal of Applied Psychology*, 100, 522–536. doi:10.1037/a0038012
- Nosek, B. A., Alter, G., Banks, G. C., Borsboom, D., Bowman, S. D., Breckler, S. J., ... & Contestabile, M. (2015). Promoting an open research culture. *Science*, 348, 1422–1425. doi:10.1126/science.aab2374
- Nye, C. D., Su, R., Rounds, J., & Drasgow, F. (2017). Interest congruence and performance: Revisiting recent meta-analytic findings. *Journal of Vocational Behavior*, 98, 138–151. doi:10.1016/j.jvb.2016.11.002
- Orwin, R. G. (1983). A fail-safe N for effect size in meta-analysis. *Journal of Educational Statistics*, 8, 157–159. doi:10.2307/1164923
- Ostroff, C., & Harrison, D. A. (1999). Meta-analysis, level of analysis, and best estimates of population correlations: Cautions for interpreting meta-analytic results in organizational behavior. *Journal of Applied Psychology*, 84, 260–270. doi:10.1037/0021-9010.84.2.260
- Paterson, T. A., Harms, P. D., Steel, P., & Credé, M. (2016). An assessment of the magnitude of effect sizes: Evidence from 30 years of meta-analysis in management. *Journal of Leadership & Organizational Studies*, 23, 66–81. doi:10.1177/1548051815614321
- Pearl, J. (2009). *Causality: Models, reasoning, and inference*. Cambridge University Press.
- Pearson, K. (1904). Report on certain enteric fever inoculation statistics. *British Medical Journal*, 2, 1243–1246. doi:10.1136/bmj.2.2288.1243

- Peters, J. L., Sutton, A. J., Jones, D. R., Abrams, K. R., & Rushton, L. (2007). Performance of the trim and fill method in the presence of publication bias and between-study heterogeneity. *Statistics in Medicine*, 26, 4544–4562. doi:10.1002/sim.2889
- Pigott, T. (2012). *Advances in meta-analysis*. New York, NY: Springer.
- Raju, N. S., Burke, M. J., Normand, J., & Langlois, G. M. (1991). A new meta-analytic approach. *Journal of Applied Psychology*, 76, 432–446. doi:10.1037/0021-9010.76.3.432
- Raudenbush, S. W. (2009). Analyzing effect sizes: Random effects models. In H. Cooper, L. V. Hedges, & J. C. Valentine (Eds.), *The handbook of research synthesis and meta-analysis* (2nd ed., pp. 295–315). New York, NY: Russell Sage Foundation.
- Raudenbush, S. W., & Bryk, A. S. (1985). Empirical Bayes meta-analysis. *Journal of Educational Statistics*, 10, 75–98. doi:10.3102/10769986010002075
- Robinson, W. S. (1950). Ecological correlations and the behavior of individuals. *American Sociological Review*, 15, 351–357. doi:10.2307/2087176
- Rosenthal, R. (1979). The file drawer problem and tolerance for null results. *Psychological Bulletin*, 86, 638–664. doi:10.1037/0033-2909.86.3.638
- Schmidt, F. (2010). Detecting and correcting the lies that data tell. *Perspectives on Psychological Science*, 5, 233–242. doi:10.1177/1745691610369339
- Schmidt, F. L. (2017). Statistical and measurement pitfalls in the use of meta-regression in meta-analysis. *Career Development International*, 22, 469–476. doi:10.1108/CDI-08-2017-0136
- Schmidt, F. L., & Hunter, J. E. (1977). Development of a general solution to the problem of validity generalization. *Journal of Applied Psychology*, 62, 529–540. doi:10.1037/0021-9010.62.5.529

- Schmidt, F. L., & Hunter, J. E. (1998). The validity and utility of selection methods in personnel psychology: Practical and theoretical implications of 85 years of research findings. *Psychological Bulletin*, 124, 262–274. doi:10.1037/0033-2909.124.2.262
- Schmidt, F.L. & Hunter, J.E. (2003). Meta-analysis. In J. A. Shinka, & W. R. Velicer (Eds.), *Handbook of psychology, volume 2: Research methods in psychology* (pp. 533–554). Hoboken, NJ: John Wiley & Sons.
- Schmidt, F. L., & Hunter, J. E. (2015). *Methods of meta-analysis: Correcting error and bias in research findings* (3rd edition). Thousand Oaks, CA: Sage.
- Schmidt, F. L., Le, H., & Oh, I.-S. (2009). Correcting for the distorting effects of study artifacts in meta-analysis. In H. Cooper, L. V. Hedges, & J. C. Valentine (Eds.), *The handbook of research synthesis and meta-analysis* (pp. 317–333). New York, NY: Russell Sage Foundation.
- Schmidt, F. L., Oh, I. S., & Hayes, T. L. (2009). Fixed-versus random-effects models in meta-analysis: Model properties and an empirical comparison of differences in results. *British Journal of Mathematical and Statistical Psychology*, 62, 97–128. doi:10.1348/000711007X255327
- Shadish, W. R. (1996). Meta-analysis and the exploration of causal mediating processes: A primer of examples, methods, and issues. *Psychological Methods*, 1, 47–65. doi:10.1037/1082-989X.1.1.47
- Shadish, W., Cook, T. D., & Campbell, D. T. (2002). *Experimental and quasi-experimental designs for generalized causal inference*. New York, NY: Houghton Mifflin.
- Siddaway, A. P., Wood, A. M., & Hedges, L. V. (2019). How to do a systematic review: A best

- practice guide for conducting and reporting narrative reviews, meta-analyses, and meta-syntheses. *Annual Review of Psychology*, 70, 747–770. doi:10.1146/annurev-psych-010418-102803
- Spearman, C. (1910). Correlation calculated from faulty data. *British Journal of Psychology*, 3, 171–195. doi:10.1111/j.2044-8295.1910.tb00206.x
- Sturman, M. C. (2003). Searching for the inverted U-shaped relationship between time and performance: Meta-analyses of the experience/performance, tenure/performance, and age/performance relationships. *Journal of Management*, 29, 609–640. doi:10.1016/S0149-2063(03)00028-X
- Tanner-Smith, E., Tipton, E., & Polanin, J. (2016) Handling complex meta-analytic data structures using robust variance estimates: A tutorial in R. *Journal of Developmental and Life-Course Criminology*, 2, 85–112. doi:10.1007/s40865-016-0026-5
- Terrin, N., Schmid, C. H., & Lau, J. (2005). In an empirical evaluation of the funnel plot, researchers could not visually identify publication bias. *Journal of Clinical Epidemiology*, 58, 894–901. doi:10.1016/j.jclinepi.2005.01.006
- Terrin, N., Schmid, C. H., Lau, J., & Olkin, I. (2003). Adjusting for publication bias in the presence of heterogeneity. *Statistics in Medicine*, 22, 2113–2126. doi:10.1002/sim.1461
- Tipton, E., & Pustejovsky, J. E. (2015). Small-sample adjustments for tests of moderators and model fit using robust variance estimation in meta-regression. *Journal of Educational and Behavioral Statistics*, 40, 604–634. doi:10.3102/1076998615606099
- Valentine, J. C., Pigott, T. D., & Rothstein, H. R. (2010). How many studies do you need? A primer on statistical power for meta-analysis. *Journal of Educational and Behavioral Statistics*, 35, 215–247. doi:10.3102/1076998609346961
- Veroniki, A. A., Jackson, D., Viechtbauer, W., Bender, R., Bowden, J., Knapp, G., ... & Salanti,

- G. (2016). Methods to estimate the between-study variance and its uncertainty in meta-analysis. *Research Synthesis Methods*, 7, 55–79. doi:10.1002/jrsm.1164
- Vevea, J. L., & Hedges, L. V. (1995). A general linear model for estimating effect size in the presence of publication bias. *Psychometrika*, 60, 419–435. doi:10.1007/BF02294384
- Viechtbauer, W. (2010). Conducting meta-analyses in R with the metafor package. *Journal of Statistical Software*, 36(3), 1–48. doi:10.18637/jss.v036.i03
- Viechtbauer, W., & Cheung, M. W. L. (2010). Outlier and influence diagnostics for meta-analysis. *Research Synthesis Methods*, 1, 112–125. doi:10.1002/jrsm.11
- Westfall, J. (2016, March 25). Five different “Cohen’s d” statistics for within-subject designs [Blog post]. Retrieved from <http://jakewestfall.org/blog/index.php/2016/03/25/five-different-cohens-d-statistics-for-within-subject-designs/>
- Wiernik, B. M., & Dahlke, J. A. (2019). Obtaining unbiased results in meta-analysis: The importance of correcting for statistical artefacts. *Advances in Methods and Practices in Psychological Science*. [In Press Accepted Manuscript].
- Wiernik, B. M., Kostal, J. W., Wilmot, M. P., Dilchert, S., & Ones, D. S. (2017). Empirical benchmarks for interpreting effect size variability in meta-analysis. *Industrial and Organizational Psychology*, 10, 472–479. doi:10.1017/iop.2017.44

Table 1. Ten Facets of Meta-Analysis & Best Practice Recommendations

Ten Facets of Meta-Analysis	Recommended "Best Practices"	Cautionary Notes & Practical Advice	Additional Readings & Resources
1. Consider Methodological & Substantive Moderators	- #1: Anticipate and test methodological moderators	- Be mindful of "small k " situations that arise when subsetting larger analyses into tests of moderators - Consider multiple moderators, and account for the possibility of correlations among moderators	Hedges & Piggot (2001); Valentine, Piggot, & Rothstein (2010) Piggot (2012, p. 21)
2. Use Random (Not Fixed) Effects Estimators	- #2: Use random effects estimators - #3: Report point estimates of between-study heterogeneity and credibility intervals - #4: Interpret estimates of heterogeneity cautiously, and in concert with true effect estimates, associated credibility intervals, and assumptions about the distribution of effects in the population	- Practically, there are very few situations in which fixed effects meta-analytic models "make sense" - Be careful not to confuse "credibility intervals" with "prediction intervals" (and interpret each appropriately) - Not all estimates of heterogeneity are created equal	Hedges and Olkin (1985, p. 190); Hunter & Schmidt (2000); Schmidt, Oh, & Hayes (2009) Higgins, Thompson, & Spiegelhalter (2009); Inhout, Ioannidis, Rovers, & Goeman, (2016) Wiernik, Kostal, Wilmot, Dilchert, & Ones (2017)
3. Conduct Sensitivity Analyses	- #5: Conduct sensitivity analyses to demonstrate the robustness of meta-analytic conclusions	- Carefully elucidate how various decisions may affect meta-analytic conclusions; pre-register such decisions whenever possible	Aguinis et al. (2011); Field et al. (2018); Geyskens, Krishnan, Steenkamp, & Cunha (2009)
4. Apply Appropriate Corrections for Statistical Artefacts	- #6: Consider appropriate corrections for statistical artefacts - #7: Report corrected and uncorrected meta-analytic estimates	- Corrections beyond sampling error (i.e., for unreliability; artificial dichotomization; selection effects) are generally warranted - Be wary of meta-analyses that report only uncorrected (or corrected) effects	Dahlke & Wiernik (2019); Schmidt, Le, & Oh, (2009) Schmidt (2010); Schmidt & Hunter (2015)
5. Acknowledge and Remediate Non-independence Among Effect Size Estimates	- #8: Recognize issues related to, and describe strategies for remediating, non-independence of effect size estimates	- Meta-analyses that average or "pick one" effect to account for non-independence should be cautiously interpreted	Hedges, Tipton, & Johnson (2010a, 2010b); Moeyaert et al., (2016)
6. Apply Appropriate Methods for the Computation of Composite Variables	- #9: Use appropriate compositing procedures to represent construct-level relationships and associated reliabilities	- Averaging effect size estimates is not a preferable strategy to represent construct-level relationships	Mossier (1943); Schmidt & Hunter (2015)
7. Explore the Possibility of Selection Effects (i.e., Publication Bias)	- #10: Develop a theory of selection effects and test it - #11: Triangulate the results of multiple tests of publication bias as sensitivity analyses	- The mechanisms that lead to selection effects are likely to vary as a function of research methodology (i.e., experimental vs. observational) - No single test will unequivocally indicate the presence or absence of publication bias	McDaniel (2009); Vevea & Hedges (1995) Carter, Schönbrodt, Gervais, & Hilgard (2019)

	- #12: Be cautious of correction-based methods for addressing publication bias	- There is no panacea for the post hoc remediation of publication bias	Becker (2005); Orwin (1983)
8. Recognize Limits on Generalizability	- #13: Report study-level characteristics and be careful not to overgeneralize meta-analytic findings beyond such characteristics - #14: Make explicit the minimum number of studies necessary to elaborate on a meta-analytic effect - #15: When possible, use multiple coders and check their agreement	- Meta-analyses only reflect characteristics of the sample of studies upon which they are based - Carefully consider the "minimum k " necessary to meaningfully represent an effect; pre-register such decisions whenever possible - Employ multiple coders, especially for "high inference" coding tasks	Matt & Cook (2009) Borenstein et al. (2011 p. 363); Valentine, Piggot, & Rothstein (2010) LeBreton & Senter (2008); Shadish et al. (2002)
9. Properly Account for Within and Between Person Effect Sizes	- #16: Be cognizant of differences in the computation of between-person versus within-person effect sizes - #17: Make decisions about managing between-person versus within-person effect sizes	- Carefully interpret the results of meta-analyses that combine estimates of within- and between person effects - Anticipate how to handle various effect sizes present in a literature <i>a priori</i> ; pre-register such decisions whenever possible	Morris (2008); Morris & DeShon (2002) Card (2019)
10. Recognize Limits on Causal Inferences	- #18: Acknowledge the types of research designs available to be synthesized - #19: Recognize the limitations of primary studies for informing causal claims from meta-analytic synthesis	- Observational studies are not well geared for demonstrating causality, meta-analytically or otherwise - Correlations are still correlations in the population	Pearl (2009); Shadish et al. (2002) Knight, Fabes, & Higgins (1996); Shadish (1996)

Table 2. Correlations Among Coded Study Characteristics

Criteria	1.	2.	3.	4.	5.	6.	7.	8.	9.	10.	11.	12.	13.	14.	15.	16.	17.
1. Year of Publication	1.00																
2. Web of Science Citations	-0.45	1.00															
3. Substantive Moderators? (1=Yes, 2=No) ^a	0.18	0.01	1.00														
4. Methodological Moderators (1=Yes, 2=No) ^b	0.16	-0.19	-0.15	1.00													
5. Heterogeneity? (1=Yes, 2=No) ^c	-0.14	-0.11	0.03	0.03	1.00												
6. Credibility Intervals? (1=Yes, 2=No) ^d	-0.02	-0.04	-0.32	0.29	0.42	1.00											
7. Sensitivity Analyses? (1=Yes, 2=No) ^e	-0.13	-0.02	-0.02	-0.02	-0.03	0.02	1.00										
8. Artefact Corrections? (1=Yes, 2=No) ^f	-0.06	-0.09	-0.29	0.27	0.19	0.39	-0.04	1.00									
9. Non-Independence (1=Yes, 2=No) ^g	-0.08	-0.03	-0.08	0.13	0.17	0.28	0.01	0.12	1.00								
10. Composite Effect Sizes? (1=Yes, 2=No) ^h	-0.17	-0.26	-0.03	-0.12	0.20	0.21	0.19	0.19	0.20	1.00							
11. Publication Bias (1=Yes, 2=No) ⁱ	-0.10	-0.06	0.20	-0.09	0.30	-0.06	0.18	-0.02	0.07	-0.02	1.00						
12. Unpublished Studies (1=Yes, 2=No) ^j	-0.30	0.05	-0.04	-0.04	0.00	0.16	0.10	-0.14	0.06	0.15	-0.01	1.00					
13. Sample Descriptives (1=Yes, 2=No) ^k	0.06	-0.01	0.60	0.01	0.00	-0.23	0.07	-0.03	-0.26	-0.10	0.10	-0.02	1.00				
14. Multiple Coders (1=Yes, 2=No) ^l	-0.14	0.14	0.00	0.00	0.11	0.07	0.11	-0.04	0.03	0.08	0.17	0.19	-0.16	1.00			
15. Design (Experimental = 1, Observational = 2) ^m	-0.05	0.07	0.17	-0.23	-0.06	-0.17	-0.08	-0.64	-0.05	-0.12	-0.04	0.09	-0.03	0.07	1.00		
16. Cross-Sectional Limitations (1=Yes, 2=No) ⁿ	-0.22	0.02	-0.20	-0.27	0.00	0.07	0.06	0.14	0.17	0.18	-0.04	-0.11	-0.15	0.03	-0.27	1.00	
17. Prospective Studies (1=Yes, 2=No) ^o	-0.20	0.09	0.02	0.11	-0.02	0.15	0.17	-0.04	0.23	0.25	0.05	0.16	0.07	0.02	0.22	0.13	1.00

Note. $n = 68$; Correlations $\geq |.26|$ are statistically significant at $p < .05$; Superscript letters cross-reference to criteria in Table 3.

Table 3. Summary of Systematic Review Results

Criteria	Category	Count	n	Percent
1. Does this study include substantive moderators? ^a	No	37	68	54.41%
	Yes	31	68	45.59%
2. Does this study include methodological moderators? ^b	No	36	68	52.94%
	Yes	32	68	47.06%
3. Are estimates of heterogeneity reported? ^c	No	11	68	16.18%
	Yes	57	68	83.82%
4. What "tradition" of meta-analysis was followed?	Hedges-Olkin	13	68	19.11%
	Schmidt-Hunter	42	68	61.76%
	Other	13	68	19.12%
5. If Hedges-Olkin approach to meta-analysis, was a fixed effect or random effects estimator used?	Random Effects	8	13	61.54%
	Not Stated	5	13	38.46%
6. Are credibility intervals reported? ^d	No	30	68	44.12%
	Yes	38	68	55.88%
7. Were explicit sensitivity analyses conducted? ^e	No	61	68	89.71%
	Yes	7	68	10.29%
8. Were corrections for statistical artefacts made? ^f	No	13	68	19.12%
	Yes	55	68	80.88%
9. How were corrections for statistical artefacts carried out?	Artefact Distribution	16	55	29.09%
	Individual Corrections	34	55	61.82%
	Unclear	5	55	9.09%
10. Were corrected and uncorrected effect size estimates reported?	No	20	55	36.36%
	Yes	35	55	63.64%
11. Is there an explicit acknowledgement of a strategy for dealing with non-independent effect size estimates? ^g	No	29	68	42.65%
	Yes	39	68	57.35%
12. Are composite variables/constructs used to represent effect size estimates? ^h	No	54	68	79.41%
	Yes	14	68	20.59%
13. Are selection effects (i.e., publication bias) explored/assessed? ⁱ	No	47	68	69.12%
	Yes	21	68	30.88%
14. Were attempts made to recover unpublished data? ^j	No	8	68	11.76%
	Yes	60	68	88.24%
15. Is information given about the types of samples under investigation? ^k	No	45	68	66.18%
	Yes	23	68	33.82%
16. Are multiple coders used? ^l	No	27	68	39.71%
	Yes	41	68	60.29%
17. Was evidence of intercoder agreement provided?	No	8	41	19.51%
	Yes	33	41	80.49%
18. Is this primary a meta-analysis of (quasi)-experimental or observational studies? ^m	Experimental	5	68	7.35%
	Observational	63	68	92.65%
19. Is there an explicit recognition of limits to causal inference for correlational studies?	No	48	63	76.19%
	Yes	15	63	23.81%
20. Is there an explicit recognition of limits to causal inference from cross-sectional research designs? ⁿ	No	34	68	50.00%
	Yes	34	68	50.00%
21. Are prospective studies considered? ^o	No	54	68	79.41%
	Yes	14	68	20.59%

Note. Superscript letters cross-reference to criteria in Table 2.

Figure 1. *Plot of the Number Meta-Analyses Published in JVB by Year of Publication*

Note. Dashed line corresponds to observed density function.

Figure 2. *Plot of Web of Science Citation Counts by Year of Publication*

Note. “WoS” = Web of Science. For ease of interpretation, “Year of Publication” reflects four, seven-year intervals. For the sake of simplifying this visualization, meta-analyses with citations > 2,000 have been trimmed from this representation.