PSYCHOMETRIKA—VOL. 83, NO. 1, 255–271
MARCH 2018
https://doi.org/10.1007/s11336-017-9571-z

MULTILEVEL MULTIVARIATE META-ANALYSIS WITH APPLICATION TO CHOICE OVERLOAD

BLAKELEY B. MCSHANE AND ULF BÖCKENHOLT

NORTHWESTERN UNIVERSITY

We introduce multilevel multivariate meta-analysis methodology designed to account for the complexity of contemporary psychological research data. Our methodology directly models the observations from a set of studies in a manner that accounts for the variation and covariation induced by the facts that observations differ in their dependent measures and moderators and are nested within, for example, papers, studies, groups of subjects, and study conditions. Our methodology is motivated by data from papers and studies of the choice overload hypothesis. It more fully accounts for the complexity of choice overload data relative to two prior meta-analyses and thus provides richer insight. In particular, it shows that choice overload varies substantially as a function of the six dependent measures and four moderators examined in the domain and that there are potentially interesting and theoretically important interactions among them. It also shows that the various dependent measures have differing levels of variation and that levels up to and including the highest (i.e., the fifth, or paper, level) are necessary to capture the variation and covariation induced by the nesting structure. Our results have substantial implications for future studies of choice overload.

Key words: multilevel, multivariate, meta-analysis, choice, overload.

1. Introduction

Contemporary psychological research can be dizzying in its complexity, and this complexity results in patterns of variation and covariation among the observations from a set of papers and studies that requires careful treatment in meta-analysis. For example, individual studies in a given domain can vary considerably in terms of their dependent measures and moderators; examine multiple conditions that result from the experimental manipulation of those moderators and give rise to multiple dependent effects of interest (e.g., simple effects and interaction effects); employ a mix of study designs (e.g., unmoderated versus moderated, between-subjects versus within-subjects, univariate versus multivariate); and feature different contexts, treatment manipulations, and measurement scales. Further, individual papers feature multiple studies that, while different, are quite similar particularly in comparison to studies featured in other papers.

However, the meta-analytic techniques typically employed in practice introduce a host of simplifications that fail to account for this complexity. For example, a common approach involves collapsing the observations from multiple conditions of each study to form a single effect of interest; converting the effects to a common, standardized scale such as the Cohen's d scale; and modeling the standardized effects via a linear mixed model with one or two variance component parameters. If differences in dependent measures or moderators are accounted for, this is typically done only via fixed main effects. These simplifications can result in *inter alia* the neglect of differences in dependent measures and moderators and miscalibrated inference.

The authors thank Alex Chernev for his encouragement and support of this article. Supplementary materials for this article are available online at http://www.blakemcshane.com/research.

Correspondence should be made to Blakeley B. McShane, Kellogg School of Management, Northwestern University, Evanston, IL 60208, USA. Email: b-mcshane@kellogg.northwestern.edu

In this paper, we introduce multilevel multivariate meta-analysis methodology that better accounts for the complexity of contemporary psychological research data. In particular, our methodology directly models the observations from a set of studies in a manner that accounts for the variation and covariation induced by the facts that observations differ in their dependent measures and moderators and are nested within, for example, papers, studies, groups of subjects, and study conditions. We also introduce two theoretically interesting and extremely parsimonious special cases of our methodology.

Our methodology is motivated by data from papers and studies of the choice overload hypothesis, the conjecture that an increase in the number of options from which to choose can result in adverse consequences such as a decrease in the likelihood of making a choice or a decrease in the satisfaction with a choice. Choice overload has already been the subject of two prominent meta-analyses (Scheibehenne et al., 2010; Chernev et al., 2015). These meta-analyses employed different variations of the simplifications to the data and model discussed above and arrived at contradictory conclusions: Scheibehenne et al. (2010) "found a mean effect size of virtually zero," whereas Chernev et al. (2015) found that "the overall effect of assortment size on choice overload is significant."

To resolve this difference, we apply our methodology to the set of fifty-seven studies from twenty-one papers originally examined by Chernev et al. (2015). By avoiding the simplifications employed in the two prior meta-analyses, our methodology more fully accounts for the complexity of choice overload data and provides richer insight. In particular, it shows that choice overload varies substantially as a function of the six dependent measures and four moderators examined in the domain and that there are potentially interesting and theoretically important interactions among them. It also shows that the various dependent measures have differing levels of variation and that levels up to and including the highest (i.e., the fifth, or paper, level) are necessary to capture the variation and covariation induced by the nesting structure.

2. Methodology

In this section, we introduce our multilevel multivariate meta-analysis methodology. Our methodology belongs, broadly speaking, to the class of meta-analytic techniques known as multivariate meta-analysis models (Kalaian & Raudenbush, 1996; Berkey et al., 1998; Becker, 2000). Such models employ a variant of the linear mixed model (Harville, 1977; Robinson, 1991) that accounts for covariation among sampling errors as well as among true effects.

Our methodology generalizes prior multivariate meta-analysis models in three important respects, namely to simultaneously accommodate (i) not two dependent measures but an arbitrary number of dependent measures; (ii) not a single effect of interest (arising from, for example, a two condition study) but an arbitrary number of study conditions that result from the experimental manipulation of moderators and give rise to multiple dependent effects of interest; and (iii) not two levels but an arbitrary number of levels that account for the variation and covariation induced by the fact that the observations are nested (e.g., within papers, studies, groups of subjects, and study conditions). These extensions are important because they are motivated by and respectful of key features of contemporary psychological research data and thereby allow our methodology to better account for the variation and covariation induced by the facts that observations differ in their dependent measures and moderators and are nested.

As a precursor to our model specification, we assume without loss of generality that each paper p features S_p studies and that there are a total of D unique dependent measures and C unique conditions across all studies and papers. We also assume, for notational simplicity, that each study of each paper measures each of the D dependent measures in each of the C conditions; as our methodology is fully general, we later relax this assumption.

We let $\mathbf{Y}_{p,s}$ be the matrix of dimension $C \times D$ containing a statistic (e.g., the mean) that summarizes the individual-level observations of each dependent measure d in each condition c in study s of paper p and we let $\mathbf{Y}_p = [\mathbf{Y}_{p,1}^{\mathrm{T}} \cdots \mathbf{Y}_{p,S_p}^{\mathrm{T}}]^{\mathrm{T}}$ be the matrix of dimension $S_pC \times D$. Our model specification for $\mathbf{Y}_p$ is given by

$$\mathbf{Y}_p = (\mathbf{1}_{S_p} \otimes \mathbf{A}) + \mathbf{B}_p + \mathbf{E}_p$$

where $\mathbf{1}_k$ is a column vector of length k containing all ones; $\mathbf{A}$ is a matrix of dimension $C \times D$ whose entries give the meta-analytic summary parameter for each condition and dependent measure and are sometimes called the fixed effects; $\mathbf{B}_p$ is a matrix of dimension $S_pC \times D$ whose entries are sometimes called the random effects for paper p; and $\mathbf{E}_p$ is a matrix of dimension $S_pC \times D$ of random (i.e., sampling) errors for paper p.[1]

We further assume that $\mathrm{vec}(\mathbf{B}_p) \overset{\mathrm{iid}}{\sim} \mathbf{N}(\mathbf{0}_{S_pCD}, \mathbf{D}_p)$ for all p; $\mathrm{vec}(\mathbf{E}_p) \overset{\mathrm{iid}}{\sim} \mathbf{N}(\mathbf{0}_{S_pCD}, \mathbf{R}_p)$ for all p; and $\mathrm{cov}(\mathrm{vec}(\mathbf{B}_p), \mathrm{vec}(\mathbf{E}_p)) = \mathbf{0}_{S_pCD, S_pCD}$ for all p where $\mathbf{0}_k$ is a column vector of length k containing all zeroes and $\mathbf{0}_{k_1,k_2}$ is a matrix of dimension $k_1 \times k_2$ containing all zeroes. Thus, the $\mathbf{R}_p$ give the level one variance-covariance matrices that model sampling error and the $\mathbf{D}_p$ give the level two and higher variance-covariance matrices that model the variation and covariation among the true values of the observations induced by the fact that the observations are nested. Given this specification, the $\mathbf{Y}_p$ have marginal distribution

$$\mathrm{vec}(\mathbf{Y}_p) \overset{\mathrm{iid}}{\sim} \mathbf{N}(\mathrm{vec}(\mathbf{1}_{S_p} \otimes \mathbf{A}), \mathbf{V}_p)$$

where $\mathbf{V}_p = \mathbf{D}_p + \mathbf{R}_p$. As is standard in meta-analysis, we assume that the $\mathbf{R}_p$ are known and that the $\mathbf{D}_p$ are a function of a parameter vector $\boldsymbol{\theta}$ (though for notational simplicity we suppress this dependence) and we seek to estimate $\mathbf{A}$ and $\boldsymbol{\theta}$.

Our specification for the $\mathbf{D}_p$ begins by noting that the $\mathbf{D}_p$ can be partitioned into blocks $\mathbf{D}_{p,d_1,d_2}$ for $d_1, d_2 \in \{1, \ldots, D\}$ where $\mathbf{D}_{p,d_1,d_2}$ is the $S_pC \times S_pC$ variance-covariance matrix of columns d_1 and d_2 of $\mathbf{B}_p$ (and where $\mathbf{D}_{p,d_1,d_2} = \mathbf{D}_{p,d_2,d_1}^{\mathrm{T}}$ due to the symmetry of $\mathbf{D}_p$). Given this, our specification for the $\mathbf{D}_{p,d_1,d_2}$ is given by

$$\mathbf{D}_{p,d_1,d_2} = \sum_{k=2}^{K} \sigma_{k,d_1,d_2} \mathbf{M}_{k,p}$$

where K gives the number of levels in the nesting structure; the σ_{k,d_1,d_2} give the degree of variation and covariation among the true values of the observations of dependent measures d_1 and d_2 induced by level k of the nesting structure; and the $\mathbf{M}_{k,p}$ are binary matrices with entries equal to one if the corresponding entries of $\mathrm{vec}(\mathbf{Y}_p)$ are nested in the same group at level k and zero otherwise.[2] We note that we start k, our index for the level, at two because the $\mathbf{D}_p$ give the level two

[1] A somewhat more general version of our model specification is given by $\mathbf{Y}_p = \mathbf{X}_p\mathbf{A} + \mathbf{Z}_p\mathbf{B}_p + \mathbf{E}_p$ where $\mathbf{Y}_p$ and $\mathbf{E}_p$ are as in the main text; $\mathbf{X}_p$ and $\mathbf{Z}_p$ are "regressor" matrices of dimension $S_pC \times q_1$ and $S_pC \times S_pq_2$, respectively; $\mathbf{A}$ is a matrix of dimension $q_1 \times D$; and $\mathbf{B}_p$ is a matrix of dimension $S_pq_2 \times D$. The specification in the main text sets $q_1 = q_2 = C$; $\mathbf{X}_p = \mathbf{1}_{S_p} \otimes \mathbf{I}_C$ where $\mathbf{I}_k$ is the $k \times k$ identity matrix; and $\mathbf{Z}_p = \mathbf{I}_{S_p} \otimes \mathbf{I}_C = \mathbf{I}_{S_pC}$. The more general specification requires that the $\mathbf{X}_p$ are known but allows the $\mathbf{D}_p$, $\mathbf{R}_p$, and possibly even the $\mathbf{Z}_p$ to be functions of a parameter vector $\boldsymbol{\theta}$. Provided the specification for the $\mathbf{Z}_p$ remains $\mathbf{Z}_p = \mathbf{I}_{S_pC}$ as in the main text, the specification for the $\mathbf{D}_p$ given there can be used under the more general specification; however, alternative specifications for the $\mathbf{Z}_p$ necessitate alternative specifications for the $\mathbf{D}_p$.

[2] Given the $\mathbf{Y}_p$ notation used in our model specification, the K^{th} level is here the paper level (and thus all entries of the $\mathbf{M}_{K,p}$ are one). However, this is a mere artifact of the notation: our methodology accommodates an arbitrary number of arbitrarily defined levels and the paper level either need not be one of or may be any one of these levels.

and higher variance–covariance matrices, whereas the $\mathbf{R}_p$ give the level one variance–covariance matrices. Given this specification, $\boldsymbol{\theta}$ consists of the variance component parameters $\{\sigma_{k,d_1,d_2}\}$ for $k \in \{2, \ldots, K\}$ and $d_1, d_2 \in \{1, \ldots, D\}$.

We term this specification for the $\mathbf{D}_p$ multilevel multivariate compound symmetry (MMCS) and note that it is quite parsimonious: it consists of only $(K-1)D(D+1)/2$ variance component parameters for a space of dimension $\sum_p S_p CD(S_p CD+1)/2$. Nonetheless, in addition to our general MMCS specification for the $\mathbf{D}_p$, we also consider two theoretically interesting and extremely parsimonious special cases of MMCS.

The first, which we term equal allocation multilevel multivariate compound symmetry (EAMMCS), constrains the fractional allocation of the variation and covariation induced by the components of the nesting structure to be equal across all dependent measure pairs and is specified as follows. Let $\sigma_{d_1,d_2} = \sum_{k=2}^{K} \sigma_{k,d_1,d_2}$ and $\pi_{k,d_1,d_2} = \sigma_{k,d_1,d_2}/\sigma_{d_1,d_2}$; in this case, the π_{k,d_1,d_2} give the fractional allocation of the degree of variation and covariation among the true values of the observations of dependent measures d_1 and d_2 induced by level k of the nesting structure. EAMMCS holds when we constrain this fractional allocation to be equal across all dependent measure pairs, namely by constraining $\pi_{k,d_1,d_2} = \pi_{k,d_3,d_4}$ for all d_1, d_2, d_3, d_4. Under this restricted specification, the parameter space consists of only $D(D+1)/2+(K-2)$ parameters (i.e., $\boldsymbol{\theta} = (\{\sigma_{d_1,d_2}\}, \{\pi_k\})$ for $d_1, d_2 \in \{1, \ldots, D\}$ and $k \in \{2, \ldots, K\}$ where $\sum \pi_k = 1$).

The second, which we term single correlation equal allocation multilevel multivariate compound symmetry (SCEAMMCS), further constrains the correlation across the dependent measures induced by the nesting structure to be equal across all dependent measure pairs—a hypothesis that is consistent with and may follow from their measuring the same construct—and is specified as follows. Let $\rho_{d_1,d_2} = \sigma_{d_1,d_2}/\sqrt{\sigma_{d_1,d_1}\sigma_{d_2,d_2}}$; in this case, the ρ_{d_1,d_2} give the correlation between dependent measures d_1 and d_2 induced by the nesting structure. SCEAMMCS holds when we constrain this correlation to be equal across all dependent measure pairs, namely by constraining $\rho_{d_1,d_2} = \rho_{d_3,d_4}$ for all d_1, d_2, d_3, d_4. Under this restricted specification, the parameter space consists of only $D + (K-1)$ parameters (i.e., $\boldsymbol{\theta} = (\{\sigma_{d,d}\}, \rho, \{\pi_k\})$ for $d \in \{1, \ldots, D\}$ and $k \in \{2, \ldots, K\}$ where $\sum \pi_k = 1$).[3]

Finally, our specification for the $\mathbf{R}_p$ is straightforward. We let the $\mathbf{R}_p$ be given by the observed variance–covariance matrix of the $\mathbf{Y}_p$; these can be easily calculated from the individual-level observations that the $\mathbf{Y}_p$ summarize. As the $\mathbf{R}_p$ give the level one variance–covariance matrices, they contain the analogue of the σ_{1,d_1,d_2} implicitly suggested by the notation used in our specification for the $\mathbf{D}_p$.

We now return to our simplifying assumption, namely that each study of each paper measures each of the D dependent measures in each of the C conditions. When dependent measure d in condition c is not observed in a given study of a given paper, we simply (i) remove the corresponding row of $\text{vec}(\mathbf{Y}_p)$, $\text{vec}(\mathbf{1}_{S_p} \otimes \mathbf{A})$, $\text{vec}(\mathbf{B}_p)$, and $\text{vec}(\mathbf{E}_p)$ and (ii) remove the corresponding row and column of $\mathbf{D}_p$ and $\mathbf{R}_p$ when forming the likelihood for paper p. This relaxes our simplifying assumption and is fully general.

Estimation of our model is as follows. We first estimate the $\mathbf{R}_p$ using conventional estimators and, as is standard in meta-analysis, assume they are known. Next, we estimate $\boldsymbol{\theta}$ (and thus the $\mathbf{D}_p$) using restricted (or residual or reduced) maximum likelihood (REML) (Harville, 1977; Robinson, 1991) conditional on the estimates of the $\mathbf{R}_p$. Finally, we estimate $\mathbf{A}$ and its variance–covariance matrix using generalized least squares conditional on the estimates of the $\mathbf{R}_p$ and $\boldsymbol{\theta}$.

[3] While a single correlation multilevel multivariate compound symmetry model specification that constrains the correlation across the dependent measures induced by the nesting structure to be equal across all dependent measure pairs without also constraining the fractional allocation to be equal (and that thus consists of only $(K-1)D + (K-1)D(D-1)/2 + 1$ parameters) is possible, we find such a specification implausible and thus do not consider it here.

3. Application to Choice Overload

In this section, we present our application to choice overload. We begin by briefly reviewing choice overload theory and two prior meta-analyses of choice overload. We then discuss our data, model specifications, and model estimates. We conclude by comparing our results to those of the two prior meta-analyses of choice overload and discussing implications for future studies of choice overload.

3.1. Theory

Common belief holds that better outcomes result when individuals have a greater number of options from which to choose. This is also "the standard line among social scientists who study choice. If we're rational, they tell us, added options can only make us better off" (Schwartz, 2004). However, a large number of studies conducted over the last fifteen years suggest this is not necessarily the case: choice sets that contain many options can trigger choice overload that makes individuals worse off. For example, individuals facing large choice sets may be less likely to make a choice or be less satisfied with a choice.

Given the counterintuitive and wide-reaching implications of this finding, the choice overload hypothesis has attracted a considerable amount of attention. Researchers have studied choice overload in a host of product categories including inexpensive everyday goods such as jam and toothpaste (Iyengar & Lepper, 2000; Chernev, 2005), expensive luxury goods such as hotel resorts and vacation packages (Chernev, 2006; Goodman & Malkoc, 2012), expensive durable goods such as mobile phones and laptops (Fasolo et al., 2009; Sela et al., 2009), and socially important goods such as charities and mutual funds (Scheibehenne et al., 2009; Morrin et al., 2012).

Researchers have also examined the effect of small versus large choice sets on several dependent measures. In particular, they have investigated behavioral outcomes such as choice deferral (i.e., the likelihood of postponing or choosing not to make a choice; Shah & Wolford, 2007; Townsend & Kahn, 2014), switching likelihood (i.e., the likelihood of switching to an alternative option; Chernev, 2003a), assortment choice (i.e., the likelihood of selecting a small versus large choice set; Chernev, 2006; Chernev & Hamilton, 2009), and option selection (i.e., the likelihood of selecting some "target" option; Gourville & Soman, 2005; Sela et al., 2009). They have also explored the effect of choice overload on subjective states such as choice satisfaction (i.e., a subjective self-assessment of satisfaction with the chosen option; Iyengar & Lepper, 2000; Diehl & Poynor, 2010), decision regret (i.e, a subjective self-assessment of regret for the chosen option; Lin & Wu, 2006; Haynes, 2009), and decision confidence (i.e, a subjective self-assessment of confidence in the chosen option; Chernev, 2003a, 2003b).

In addition, researchers have examined moderators of choice overload, that is, factors that exacerbate, attenuate, nullify, or reverse it. These moderators can be divided into two broad classes: extrinsic (or objective) factors and intrinsic (or subjective) factors. Among the former are choice set complexity (i.e., aspects of the decision task that influence the values of the particular choice options without necessarily influencing the structural aspects of the decision problem at hand; Chernev, 2005; Gourville & Soman, 2005) and decision task difficulty (i.e., general structural characteristics of the decision problem that do not influence the values of the particular choice options; Greifeneder et al., 2010; Inbar et al., 2011) while among the latter are preference uncertainty (i.e., the degree to which individuals have articulated preferences with respect to the decision at hand; Chernev, 2003b; Mogilner et al., 2008) and the decision goal (i.e., the degree to which individuals aim to minimize the cognitive effort involved in making a choice; Oppewal & Koelemeijer, 2005; Lin & Wu, 2006).

Chernev et al. (2015) present a conceptual model of choice overload based on Payne et al. (1993) that we depict graphically in Fig. 1. We use this conceptual model to inform how

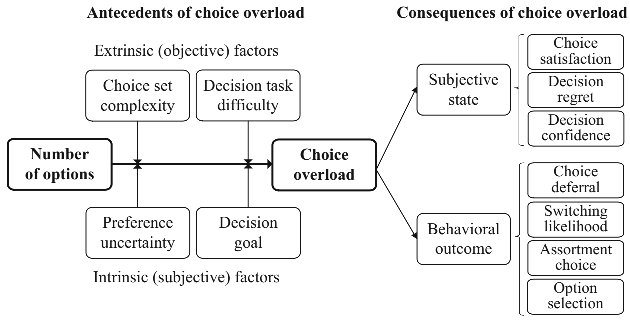

FIGURE 1.
Conceptual model of the impact of choice set size on choice overload. The four antecedents of choice overload are operationalized as follows: (i) the complexity of the choice set describes the aspects of the decision set associated with the particular values of the choice options: the presence of a dominant option in the choice set, the overall attractiveness of the options in the choice set, and the relationship between individual options in the decision set (alignability and complementarity); (ii) the difficulty of the decision task refers to the general structural characteristics of the decision problem: time constraints, decision accountability, and number of attributes describing each option; (iii) preference uncertainty refers to the degree to which individuals have articulated preferences with respect to the decision at hand and has been operationalized by two factors: the level of product-specific expertise and the availability of an articulated ideal point; and (iv) the decision goal reflects the degree to which individuals aim to minimize the cognitive effort involved in making a choice among the options contained in the available choice sets and is operationalized by two measures: decision intent (buying vs. browsing) and decision focus (choosing a choice set vs. choosing a particular option). In this context, we expect higher levels of decision task difficulty, greater choice set complexity, higher preference uncertainty, and a more prominent, effort-minimizing goal to produce greater choice overload. Source: Chernev et al. (2015), Figure 1.

our methodology accounts for the fact that observations of choice overload data differ in their dependent measures and moderators.

3.2. Prior Meta-analyses

Given the considerable attention the choice overload hypothesis has attracted, it may be unsurprising that choice overload has already been the subject of two prominent meta-analyses (Scheibehenne et al., 2010, Chernev et al., 2015). These meta-analyses arrived at contradictory conclusions: Scheibehenne et al. (2010) "found a mean effect size of virtually zero" and concluded that only "idiosyncratic moderators...explain when and why choice overload reliably occurs," whereas Chernev et al. (2015) found that "when moderating variables are taken into account the overall effect of assortment size on choice overload is significant."

Importantly, both of these meta-analyses employed the simplifications to the data and model discussed in Sect. 1 and thus failed to fully account for the complexity of choice overload data. In particular, they simplified the analysis along three lines. First, both collapsed the observations across study conditions and used the simple contrast between small and large choice sets as measured on the standardized Cohen's d scale as the response variable. This resulted in the distinction among the dependent measures either being largely ignored (Scheibehenne et al., 2010) or accounted for only via fixed main effects (Chernev et al., 2015). Second, they largely ignored (Scheibehenne et al., 2010) or only partially accounted for (Chernev et al., 2015) the four moderators studied in the domain. Third, they did not account for the variation and covariation in the observations induced by the fact that they differ in their dependent measures and are nested within papers, studies, groups of subjects, and study conditions. In particular, Scheibehenne et al. (2010) employed a two-level model with a single variance component parameter that ignores *inter alia* potentially differing levels of variation among the dependent measures and the covariation among observations from the same paper and study. Similarly, Chernev et al. (2015) employed a three-level model with two variance component parameters that ignores *inter alia* potentially differing levels of variation among the dependent measures and the covariation among observations from the same study beyond the covariation among observations from the same paper.

Given the differing simplifications employed, it is perhaps unsurprising that these two prior meta-analyses arrived at contradictory conclusions. To resolve this difference, we employ our methodology that more fully accounts for the complexity of choice overload data.

3.3. Data

For consistency with prior meta-analyses, we examine choice overload using the same set of fifty-seven studies from twenty-one papers examined by Chernev et al. (2015). This allows us to be sure any differences between their results and ours are due to differences in the meta-analysis model employed rather than the underlying data.

Although we examine the same studies and papers as Chernev et al. (2015), our datasets are entirely distinct. In particular, and as noted in the prior subsection, Chernev et al. (2015) collapsed the observations across study conditions and analyzed the simple contrast between small and large choice sets as measured on the standardized Cohen's d scale as their univariate response variable. In contrast, we use the summary statistic(s) from each study condition as our multivariate response variable, thereby better preserving the nature of the data. For binary dependent measures (i.e., assortment choice, choice deferral, option selection, and switching likelihood), we use the proportion as our summary statistic; however, we conduct our analysis on the log odds scale so all estimates remain bounded between zero and one when converted to proportions. For integer-scale dependent measures (i.e., confidence, regret, and satisfaction), we use the mean as our summary statistic. We convert all integer-scale dependent measures to the one-to-nine scale as this was the most commonly used scale in our choice overload data; our conversion is via a linear transformation,[4] and while more sophisticated conversions are possible, they can require moving beyond the linear mixed model.

Following Chernev et al. (2015), and again to be sure any differences between their results and ours are due to differences in the meta-analysis model employed, we treat satisfaction and confidence as a single dependent measure thus leaving six dependent measures and four moderators of choice overload. The data are coded such that higher values imply more positive outcomes for all dependent measures (i.e., choice deferral, switching likelihood, and regret are reverse-coded), and thus choice overload occurs when a dependent measure is lower for large versus small choice sets.

In total, our data consist of 172 observations from 154 conditions measured on 147 groups of subjects from fifty-seven studies from twenty-one papers. Consequently, the nesting structure consists of $K = 5$ levels with observations nested within papers (level five), studies (level four), groups of subjects (level three), and study conditions (level two).

Before proceeding, we make three comments about the paper and study designs employed in our choice overload data. First, we present the number of papers and studies measuring each dependent measure and manipulating each moderator in Table 1. As can be seen, two-fifths of the dependent measure/moderator combinations are entirely unexamined while two-thirds (over three-quarters) are examined by only one paper (three or fewer studies). This suggests many entries of $\mathbf{A}$ are inestimable while many others are only weakly estimable.

Second, only six papers and only six studies from five papers measure more than one dependent measure; all of these measure only two dependent measures with satisfaction as one of them. Consequently, the vast majority of the σ_{k,d_1,d_2}, $d_1 \neq d_2$ are inestimable while the remaining ones are only weakly estimable.

[4] In particular, for integer-scale dependent measure y with scale minimum m and scale maximum M, we use $1 + \frac{9-1}{M-m}(y - m)$ in place of y.

TABLE 1.
Dependent measures and moderators of choice overload.

Dependent measure	Moderator				
	No moderator	Choice set complexity	Decision goal	Decision task difficulty	Preference uncertainty
Assortment choice	0 (0)	1 (8)	2 (9)	0 (0)	0 (0)
Choice deferral	3 (6)	1 (1)	0 (0)	2 (2)	2 (2)
Option selection	1 (4)	2 (3)	0 (0)	1 (2)	1 (1)
Regret	1 (1)	0 (0)	1 (1)	2 (2)	0 (0)
Satisfaction	3 (7)	0 (0)	2 (2)	2 (3)	4 (5)
Switching likelihood	0 (0)	0 (0)	0 (0)	0 (0)	1 (4)

Each cell gives the number of papers (studies) that measure the dependent measure given by the row and manipulate the moderator given by the column. Papers that measure multiple dependent measures or manipulate multiple moderators (studies that measure multiple dependent measures) appear in more than one cell and thus the total in the table is greater than the number of unique papers (studies). Two-fifths of the dependent measure/moderator combinations are entirely unexamined while two-thirds (over three-quarters) are examined by only one paper (three or fewer studies).

Third, none of the eight studies from six papers that follow either within-subjects or multivariate study designs reported the correlation of the observations relevant for the $\mathbf{R}_p$—an unfortunate but common practice (Becker et al., 2004; Riley, 2009). Thus, in our principal analysis, we follow the recommendation of Ishak et al. (2008) and assume these correlations are zero. However, we also follow the recommendation of Riley (2009) and conduct a sensitivity analysis to assess the impact of the assumption of zero correlation; such an analysis reveals no important differences in results (see "Appendix A" for details).

3.4. *Model Specifications*

Before proceeding to model estimates, we first consider a sequence of nested simplifications to the MMCS and EAMMCS specifications for $\boldsymbol{\theta}$ presented in Sect. 2. The first and simplest model specification we consider is the so-called fixed effects (or one-level) model specification that sets the $\sigma_{k,d_1,d_2} = 0$ for all k, d_1, d_2 for the MMCS specification (or equivalently the $\sigma_{d_1,d_2} = 0$ for all d_1, d_2 for the EAMMCS specification).

The second model specification we consider sets the $\sigma_{k,d,d}$ equal to one another for all d and the $\sigma_{k,d_1,d_2} = 0$ for all $k, d_1 \neq d_2$ for the MMCS specification (or equivalently the $\sigma_{d,d}$ equal to one another for all d and the $\sigma_{d_1,d_2} = 0$ for all $d_1 \neq d_2$ for the EAMMCS specification). We term this the Equal Variance specification and note that this specification is analogous to those employed in prior meta-analyses of choice overload (Scheibehenne et al., 2010 employed a two-level version of this specification while Chernev et al., 2015 employed a three-level version of this specification; as noted, we employ a five-level version and also consider a richer specification for $\mathbf{A}$).

The third model specification we consider sets (i) the $\sigma_{k,d_1,d_2} = 0$ for all $k, d_1 \neq d_2$ for the MMCS specification and (ii) the $\sigma_{d_1,d_2} = 0$ for all $d_1 \neq d_2$ for the EAMMCS specification. We term this the Unequal Variance specification and note that this specification is designed to relax the Equal Variance assumption employed in prior meta-analyses of choice overload.

The fourth model specification we consider imposes no constraints on (i) the σ_{k,d_1,d_2} for the MMCS specification and (ii) the σ_{d_1,d_2} and the π_k for the EAMMCS specification other than that imposed by the paper and study designs (i.e., inestimable parameters are set to zero as is common in practice). We note that the two-level version of this specification is analogous to prior

TABLE 2.
Model specification results.

Model specification	MMCS		EAMMCS
Fixed effects		−450.61 (0)	
Equal Variance		−111.55 (4)	
Unequal Variance	−100.46 (19)		−102.83 (9)
No Constraints	−97.44 (26)		−101.74 (13)

The table gives the REML log likelihood of (number of $\boldsymbol{\theta}$ parameters estimated by) each model specification. The EAMMCS Unequal Variance specification seems to best balance model fit against the number of estimated parameters.

multivariate meta-analysis models (Kalaian & Raudenbush, 1996; Berkey et al., 1998; Becker, 2000).

We present the REML log likelihood of and the number of $\boldsymbol{\theta}$ parameters estimated by each of these model specifications in Table 2. For comparison with prior meta-analysis models, we note that the REML log likelihood of (number of $\boldsymbol{\theta}$ parameters estimated by) the two-level version of the Equal Variance specification is −138.82 (1) and the two-level version of the No Constraints specification is −123.28 (10). From this table and these results, it is clear that (i) our five-level model specifications improve upon prior meta-analysis models thus demonstrating their benefit in our application to choice overload; (ii) the MMCS specifications do not improve substantially on the EAMMCS specifications thus reflecting that the fractional allocation of the variation and covariation induced by the four components of the nesting structure is equal (or at least relatively similar) across the dependent measures; (iii) the Unequal Variance specification improves substantially upon the Equal Variance specification thus casting doubt on the Equal Variance assumption employed in prior meta-analyses of choice overload; and (iv) the No Constraints specification does not improve substantially on the Unequal Variance specification thus reflecting weak covariation among the various dependent measure pairs or the fact that few papers and studies measure more than one dependent measure.

We note we also considered the SCEAMMCS specification for $\boldsymbol{\theta}$ presented in Sect. 2. As this specification has no corresponding MMCS counterpart specification, we do not present it in Table 2. However, the REML log likelihood of (number of $\boldsymbol{\theta}$ parameters estimated by) this specification is −102.79 (10) thus further reflecting weak covariation among the various dependent measure pairs or the fact that few papers and studies measure more than one dependent measure.

In sum, it seems the EAMMCS Unequal Variance specification best balances model fit against the number of estimated parameters (e.g., it is the preferred model specification when using the Akaike information criterion; Akaike, 1974). Given this, we proceed by analyzing the EAMMCS Unequal Variance specification in greater depth.

3.5. Model Estimates

In this subsection, we discuss our model estimates of $\mathbf{A}$ and $\boldsymbol{\theta}$ beginning with the former. We present our estimates of the entries of $\mathbf{A}$, the meta-analytic summary parameter for each condition and dependent measure, along with estimates from individual studies in Fig. 2. The large points represent the estimates of the entries of $\mathbf{A}$ while the vertical lines represent estimates of plus and minus one standard error; the small points represent estimates from individual studies. Estimates for the binary dependent measures are presented on the probability scale while estimates for the integer-scale dependent measures are presented on the one-to-nine scale. For all dependent measures except assortment choice, choice overload occurs when a dependent measure is lower

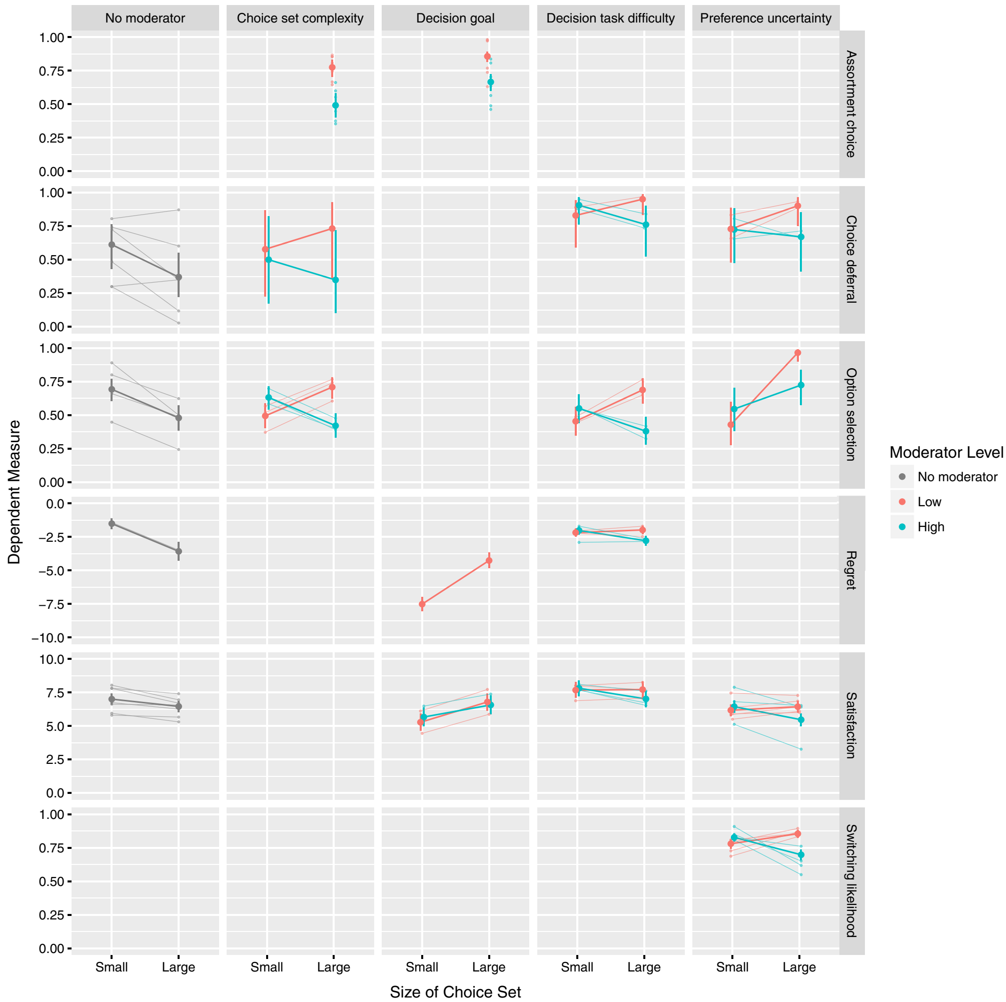

FIGURE 2.
Choice overload results. The large points represent the estimates of the entries of **A**, the meta-analytic summary parameter for each condition and dependent measure, while the vertical lines represent estimates of plus and minus one standard error. The small points represent estimates from individual studies. Estimates for the binary dependent measures are presented on the probability scale while estimates for the integer-scale dependent measures are presented on the one-to-nine scale. Choice overload occurs when a dependent measure is lower for large versus small choice sets (or for lower values when the dependent measure is assortment choice). Choice overload varies substantially as a function of the dependent measure and moderator. These estimates and estimated standard errors along with z-statistics are available in "Appendix B".

for large versus small choice sets; as assortment choice measures the choice of a large versus a small choice set, choice overload occurs when the proportion choosing the large choice set is lower. These estimates and estimated standard errors are available in "Appendix B"; because the estimates for large and small choice sets conditional on the dependent measure, moderator, and moderator level tend to be highly correlated, we also provide z-statistics, which give the difference in the estimates for large versus small choice sets divided by the estimated standard error of this difference, in "Appendix B." The empty subplots and the subplots with few small points that represent estimates from individual studies highlight the respective facts that no or few papers and studies measure the corresponding dependent measure/moderator combinations (see also Table 1).

TABLE 3.
Estimates and estimated standard errors of $\sigma_{d,d}$.

Dependent measure	Estimate	SE
Assortment choice	0.52	0.13
Choice deferral	1.53	0.48
Option selection	0.53	0.21
Regret	0.37	0.31
Satisfaction	0.90	0.17
Switching likelihood	0.17	0.21

We present estimates on the standard deviation scale (i.e., estimates of $\sqrt{\sigma_{d,d}}$) rather than the variance scale so as to match the scale of the $\mathbf{Y}_p$ and $\mathbf{A}$. The various dependent measures have differing levels of variation.

Figure 2 shows that choice overload varies substantially as a function of the dependent measure and moderator. For example, choice overload occurs for the high level of the decision task difficulty moderator when option selection or satisfaction is the dependent measure; however, it is reversed for the low level of the decision task difficulty moderator when option selection is the dependent measure while it is nullified when satisfaction is the dependent measure. It also occurs when there are no moderators regardless of the dependent measure while it fails to occur when there are moderators—regardless of the moderator level—when choice deferral is the dependent measure.

In sum, while choice overload reliably occurs for some dependent measure/moderator combinations and it is reliably reversed for others, for still others the evidence is quite mixed. Indeed, it is clear there are interactions among the dependent measures and moderators and thus fully accounting for the various dependent measures and moderators is critical for modeling and understanding choice overload.

We present our estimates of the $\sigma_{d,d}$ parameters of $\boldsymbol{\theta}$ along with estimated standard errors in Table 3; we present estimates on the standard deviation scale (i.e., estimates of $\sqrt{\sigma_{d,d}}$) rather than the variance scale so as to match the scale of the $\mathbf{Y}_p$ and $\mathbf{A}$. The various dependent measures have differing levels of variation, with estimates ranging from close to zero for switching likelihood to substantial for choice deferral.

Finally, we present our estimates of the π_k parameters of $\boldsymbol{\theta}$; as $\sum \pi_k = 1$, these estimates reflect the fractional allocation of the variation and covariation induced by the four components of the nesting structure. We estimate the fractional allocation induced at the paper level is 0.36, at the study level is 0.53, at the group of subjects level is 0.01, and at the condition level is 0.10. The covariation among the observations from the same paper and study is quite large; however, because only four studies from three papers follow within-subjects study designs, the covariation at the group of subjects level and the condition level is difficult to dissociate in these data.

3.6. Comparisons to Prior Meta-analyses and Implications for Future Studies

We assess our results in light of prior meta-analyses of choice overload. We then discuss some implications of our results for future studies of choice overload.

Beginning with Scheibehenne et al. (2010), our results are inconsistent with their finding of "a mean effect size of virtually zero" and conclusion that only "idiosyncratic moderators...explain when and why choice overload reliably occurs" that they obtained by largely ignoring the distinction among the dependent measures and moderators. In contrast, Fig. 2 and "Appendix B" show that choice overload varies substantially as a function of the dependent measure and moderator and reliably occurs for many dependent measure/moderator combinations. Further, their two-level model with a single variance component parameter assumes that (i) the various dependent measures have the same level of variation and (ii) the variation and covariation induced by the four components of the nesting structure is induced only at the lowest level such that there is no covariation induced by the nesting structure. In contrast, Tables 2 and 3 show that the various dependent measures have differing levels of variation while our estimates of the π_k show that 0.90 of the fractional allocation of the variation and covariation induced by the four components of the nesting structure is induced by levels higher than the lowest such that there is substantial covariation induced by the nesting structure. Falsely assuming no covariation is induced by the nesting structure typically results in overly confident inference.

Moving onto Chernev et al. (2015), our results are largely consistent with their finding that "when moderating variables are taken into account the overall effect of assortment size on choice overload is significant." However, our results are rather more nuanced because they show interactions among the dependent measures and moderators. For example, their Table 2 indicates that choice overload occurs in studies with no moderators, is exacerbated (reversed) in studies with moderators when the moderator is set to the high (low) level, and does not much differ among the various dependent measures. In contrast, Fig. 2 and "Appendix B" show choice overload varies substantially as a function of the dependent measure and moderator and that there are interactions among them. Further, their three-level model with two variance component parameters assumes that (i) the various dependent measures have the same level of variation and (ii) there is no covariation among observations from the same study beyond the covariation among observations from the same paper. In contrast, Tables 2 and 3 show that the various dependent measures have differing levels of variation while our estimates of the π_k show that there is substantial additional covariation among observations from the same study beyond the covariation among observations from the same paper. In particular, our estimates of the π_k show that 0.53 of the variation and covariation induced by the four components of the nesting structure is due to covariation among observations from the same study beyond the covariation among observations from the same paper. Because their estimate of the fractional allocation of the variation and covariation induced by the nesting structure at the paper level is 0.72 and at the observation level is 0.28, their inference is likely to be overly conservative.

More broadly, our results have implications for future studies of choice overload. For example, our results show choice overload fails to occur when there are moderators—regardless of the moderator level—when choice deferral is the dependent measure. However, they also show that choice deferral is the dependent measure with by far the largest level of variation. As this is a prominent and important dependent measure in the choice overload literature (e.g., it was the primary dependent measure featured in the original studies of choice overload by Iyengar and Lepper 2000), this suggests future research should examine this dependent measure and its interaction with various moderators in greater depth. In addition, for some dependent measures (e.g., option selection), when the moderator is set to the low level, choice overload tends to be reversed, while for other dependent measures (e.g., satisfaction), it tends to be nullified. This suggests future research should develop richer theory and examine how various moderators exacerbate, attenuate, nullify, or reverse choice overload depending on the dependent measure. Further, the fact that choice overload occurs for all dependent measures when there is no moderator but there is a lack of agreement among the dependent measures when there is a moderator suggests that future research should investigate whether the various dependent measures are equally valid in assessing choice overload.

4. Discussion

We have introduced multilevel multivariate meta-analysis methodology that better accounts for the complexity of contemporary psychological research data, in particular the variation and covariation induced by the facts that observations differ in their dependent measures and moderators and are nested. Our methodology generalizes prior multivariate meta-analysis models in three important respects, namely to simultaneously accommodate an arbitrary number of dependent measures; an arbitrary number of study conditions that give rise to multiple dependent effects of interest; and an arbitrary number of levels that account for the fact that the observations are nested.

We have also introduced two theoretically interesting and extremely parsimonious special cases of MMCS, namely the EAMMCS specification that constrains the fractional allocation of the variation and covariation induced by the various components of the nesting structure to be equal across all dependent measures pairs as well as the SCEAMMCS specification that further constrains the correlation across the dependent measures induced by the nesting structure to be equal across all dependent measures pairs.

In our application to choice overload, our analysis provided richer insight compared to two prior meta-analyses. In particular, it showed that choice overload varies substantially as a function of the six dependent measures and four moderators examined in the domain and that there are potentially interesting and theoretically important interactions among them. It also showed that the various dependent measures have differing levels of variation and that levels up to and including the highest (i.e., the fifth, or paper, level) are necessary to capture the variation and covariation induced by the nesting structure.

While our SCEAMMCS specification did not outperform other model specifications in our application to choice overload, we believe this reflects the fact that few papers and studies of choice overload measure more than one dependent measure and consequently that allowing for covariation among the various dependent measure pairs simply does not much impact the model fit in this application. Nonetheless, we believe this highly parsimonious model specification will prove a powerful benchmark specification in future applications.

We note that underlying the SCEAMMCS specification is, in essence, a factor analytic model featuring a single factor. Future research may consider extending this specification to allow for multiple factors. This will allow for a richer pattern of covariation among the various dependent measure pairs while remaining parsimonious.

Another potential model extension involves accounting for authors as well as papers, studies, groups of subjects, and study conditions. If authors can be treated as random, our model can trivially accommodate them by treating them as a sixth level. Nonetheless, this requires some degree of special treatment: whereas it is quite clear what it means for observations to be nested in the same group for groups such as papers, studies, groups of subjects, and study conditions, it is less clear what this means for a group such as authors where, for example, partial overlap is possible. Thus, future research may consider extending the model to account for authors considering perhaps both fixed as well as random treatments.

In conclusion, we reiterate that the complexity of contemporary psychological research data results in patterns of variation and covariation among the observations from a set of studies that requires careful treatment in meta-analysis. In particular, individual studies vary in terms of their dependent measures and moderators; examine multiple conditions that give rise to multiple dependent effects of interest; and feature different contexts, treatment manipulations, and measurement scales while individual papers feature multiple studies that vary but are similar in comparison to studies featured in other papers. There is mounting evidence of and growing appreciation for these differences across papers and studies and the complexity that results from them as well as the need to address this complexity in sample size planning (McShane & Böckenholt, 2014), in meta-analysis (McShane & Böckenholt, 2017), and in adjusting for publication bias (McShane et al., 2016). The methodology presented here directly addresses this complexity by accommodating an arbitrary number of dependent measures, study conditions, and levels in the nesting structure. Because this methodology is motivated by and respectful of key features of contemporary psychological research data, it is quite general and widely applicable and we expect it to yield rich insight in future applications.

Appendix A: Sensitivity Analysis

None of the eight studies from six papers that follow either within-subjects or multivariate study designs reported the correlation of the observations relevant for the $\mathbf{R}_p$—an unfortunate but common practice (Becker et al., 2004; Riley, 2009). Thus, in our principal analysis, we followed the recommendation of Ishak et al. (2008) and assumed these correlations were zero. In this appendix, we follow the recommendation of Riley (2009) and present a sensitivity analysis that assesses the impact of the assumption of zero correlation on our results. As discussed below, our analysis reveals no important differences in results; this is perhaps unsurprising as (i) only eight studies from six papers follow either within-subjects or multivariate study designs and (ii) the influence of within-study correlation is small when, as is generally the case in our data, the variation among the true values of the observations is large relative to sampling error (Riley, 2009).

Across these eight studies from six papers, there were eight distinct correlation parameters that were unreported. Consequently, we set these correlation parameters in turn to 0.1, 0.3, and 0.5 (i.e., a small, medium, and large correlation; Cohen, 1992) and recompute the $\mathbf{R}_p$ under these assumptions. We then refit the EAMMCS Unequal Variance specification using each of the $3^8 = 6561$ versions of the $\mathbf{R}_p$.

Across the 6561 models, the REML log likelihood ranged from -103.37 to -102.15, a range of only 1.22. The REML log likelihood of the model presented in the main text that assumed zero correlation (i.e., diagonal $\mathbf{R}_p$) was -102.83, solidly in the middle of this range. In terms of the estimates of $\mathbf{A}$, 98.1% (73.3%) were within 0.10 (0.01) of the estimates of the model presented in the main text; further, 100% (100%) were included within the 95% (50%) interval estimates of the model presented in the main text. The RMSE of the estimates as compared to the estimates of the model presented in the main text was 0.03. In terms of the estimates of the $\sigma_{d,d}$ parameters of $\boldsymbol{\theta}$ (which we evaluate on the standard deviation scale (i.e., estimates of $\sqrt{\sigma_{d,d}}$) so as to match the scale of the $\mathbf{Y}_p$ and $\mathbf{A}$), 100.0% (51.3%) were within 0.10 (0.01) of the estimates of the model presented in the main text. The RMSE of the estimates as compared to the estimates of the model presented in the main text was 0.02. In terms of the estimates of the π_k parameters of $\boldsymbol{\theta}$, 100.0% (71.9%) were within 0.10 (0.01) of the estimates of the model presented in the main text. The RMSE of the estimates as compared to the estimates of the model presented in the main text was 0.01.

Appendix B: Model Estimates

Dependent measure	Moderator	Moderator level	Small choice set Estimate	Small choice set SE	Large choice set Estimate	Large choice set SE	z
Assortment choice	Choice set complexity	Low			1.24	0.38	
	Choice set complexity	High			−0.03	0.37	
	Decision goal	Low			1.79	0.31	
	Decision goal	High			0.69	0.28	
Choice deferral	No moderator	No moderator	0.46	0.73	−0.53	0.74	−2.75
	Choice set complexity	Low	0.31	1.56	1.01	1.56	0.82
	Choice set complexity	High	0.00	1.56	−0.62	1.56	−0.75
	Decision task difficulty	Low	1.58	1.21	2.97	1.36	1.27
	Decision task difficulty	High	2.27	1.10	1.16	1.06	−1.61
	Preference uncertainty	Low	0.99	1.07	2.22	1.11	1.74
	Preference uncertainty	High	0.97	1.07	0.71	1.06	−0.41
Option selection	No moderator	No moderator	0.81	0.39	−0.08	0.38	−3.48
	Choice set complexity	Low	−0.02	0.38	0.89	0.38	3.17
	Choice set complexity	High	0.54	0.38	−0.32	0.38	−3.05
	Decision task difficulty	Low	−0.18	0.44	0.79	0.45	3.33
	Decision task difficulty	High	0.20	0.44	−0.49	0.44	−2.41
	Preference uncertainty	Low	−0.29	0.69	3.33	1.15	3.19
	Preference uncertainty	High	0.18	0.68	0.97	0.67	1.21
Regret	No moderator	No moderator	−1.50	0.40	−3.57	0.70	−3.16
	Decision goal	Low	−7.51	0.53	−4.25	0.59	5.24
	Decision task difficulty	Low	−2.17	0.31	−1.97	0.31	0.74
	Decision task difficulty	High	−2.01	0.31	−2.78	0.35	−2.48

Dependent measure	Moderator	Moderator level	Small choice set		Large choice set		z
			Estimate	SE	Estimate	SE	
Satisfaction	No moderator	No moderator	7.00	0.42	6.45	0.42	−2.87
	Decision goal	Low	5.27	0.64	6.79	0.64	4.60
	Decision goal	High	5.67	0.70	6.57	0.70	1.97
	Decision task difficulty	Low	7.68	0.59	7.71	0.59	0.12
	Decision task difficulty	High	7.81	0.59	7.01	0.60	−2.40
	Preference uncertainty	Low	6.17	0.44	6.44	0.44	1.10
	Preference uncertainty	High	6.44	0.46	5.46	0.47	−3.08
Switching likelihood	Preference uncertainty	Low	1.27	0.21	1.78	0.23	2.04
	Preference uncertainty	High	1.57	0.22	0.84	0.19	−3.20

Estimates and estimated standard errors of the entries of **A**. Estimates for the binary dependent measures (i.e., assortment choice, choice deferral, option selection, and switching likelihood) are presented on the log odds scale, while estimates for the integer-scale dependent measures (i.e., satisfaction and regret) are presented on the one-to-nine scale. Assortment choice measures the choice of a large versus a small choice set; thus, estimates for the choice of the large choice set are presented. The z-statistic gives the difference in the estimates for large versus small choice sets divided by the estimated standard error of this difference. Choice overload occurs when a dependent measure is lower for large versus small choice sets (or for lower values when the dependent measure is assortment choice). These estimates are depicted graphically in Fig. 2. Estimates of $\boldsymbol{\theta}$ are available in the main text.

References

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, *19*(6), 716–723.

Becker, B. J. (2000). Multivariate meta-analysis. In H. E. A. Tinsley & S. D. Brown (Eds.), *Handbook of applied multivariate statistics and mathematical modeling*. San Diego: Academic Press.

Becker, B. J., Hedges, L. V., & Pigott, T. D. (2004). *Statistical analysis policy brief*. Oslo: The Campbell Collaboration.

Berkey, C . S., Hoaglin, D . C., Antczak-Bouckoms, A., Mosteller, F., & Colditz, G . A. (1998). Meta-analysis of multiple outcomes by regression with random effects. *Statistics in Medicine*, *17*(22), 2537–2550. ISSN 1097-0258.

Chernev, A. (2003a). When more is less and less is more: The role of ideal point availability and assortment in consumer choice. *Journal of Consumer Research*, *30*(2), 170–183.

Chernev, A. (2003b). Product assortment and individual decision processes. *Journal of Personality and Social Psychology*, *85*, 151–162.

Chernev, A. (2005). Feature complementarity and assortment in choice. *Journal of Consumer Research*, *31*(4), 748–759.

Chernev, A. (2006). Decision focus and consumer choice among assortments. *Journal of Consumer Research*, *33*(1), 50–59.

Chernev, A., & Hamilton, R. (2009). Assortment size and option attractiveness in consumer choice among retailers. *Journal of Marketing Research*, *46*(3), 410–420.

Chernev, A., Böckenholt, U., & Goodman, J. (2015). Choice overload: A conceptual review and meta-analysis. *Journal of Consumer Psychology*, *25*(2), 333–358.

Cohen, J. (1992). A power primer. *Psychological Bulletin*, *112*, 155–159.

Diehl, K., & Poynor, C. (2010). Great expectations?! Assortment size, expectations and satisfaction. *Journal of Marketing Research*, *47*(2), 312–322.

Fasolo, B., Carmeci, F. A., & Misuraca, R. (2009). The effect of choice complexity on perception of time spent choosing: When choice takes longer but feels shorter. *Psychology and Marketing*, *26*(3), 213–228.

Goodman, J. K., & Malkoc, S. A. (2012). Choosing here and now versus there and later: The moderating role of psychological distance on assortment size preferences. *Journal of Consumer Research*, *39*(4), 751–768.
Gourville, J . T., & Soman, D. (2005). Overchoice and assortment type: When and why variety backfires. *Marketing Science*, *24*(3), 382–395.
Greifeneder, R., Scheibehenne, B., & Kleber, N. (2010). Less may be more when choosing is difficult: Choice complexity and too much choice. *Acta Psychologica*, *133*(1), 45–50.
Harville, D. A. (1977). Maximum likelihood approaches to variance component estimation and to related problems. *Journal of the American Statistical Association*, *72*(358), 320–338.
Haynes, G. A. (2009). Testing the boundaries of the choice overload phenomenon: The effect of number of options and time pressure on decision difficulty and satisfaction. *Psychology & Marketing*, *26*(3), 204–212.
Inbar, Y., Botti, S., & Hanko, K. (2011). Decision speed and choice regret: When haste feels like waste. *Journal of Experimental Social Psychology*, *47*(3), 533–540.
Ishak, K. J., Platt, R. W., Joseph, L., & Hanley, J. A. (2008). Impact of approximating or ignoring within-study covariances in multivariate meta-analyses. *Statistics in Medicine*, *27*, 670–686.
Iyengar, S. S., & Lepper, M. R. (2000). When choice is demotivating: Can one desire too much of a good thing? *Journal of Personality and Social Psychology*, *79*(6), 996–1006.
Kalaian, H. A., & Raudenbush, S. W. (1996). A multivariate mixed linear model for meta-analysis. *Psychological Methods*, *1*(3), 227–235.
Lin, C.-H., & Wu, P.-H. (2006). The effect of variety on consumer preferences: The role of need for cognition and recommended alternatives. *Social Behavior and Personality: An International Journal*, *34*(7), 865–875.
McShane, B. B., & Böckenholt, U. (2014). You cannot step into the same river twice: When power analyses are optimistic. *Perspectives on Psychological Science*, *9*(6), 612–625.
McShane, B. B., & Böckenholt, U. (2017). Single paper meta-analysis: Benefits for study summary, theory-testing, and replicability. *Journal of Consumer Research*, *43*, 1048–1063.
McShane, B. B., Böckenholt, U., & Hansen, K. T. (2016). Adjusting for publication bias in meta-analysis: An evaluation of selection methods and some cautionary notes. *Perspectives on Psychological Science*, *11*(5), 730–749.
Mogilner, C., Rudnick, T., & Iyengar, S. S. (2008). The mere categorization effect: How the presence of categories increases choosersí perceptions of assortment variety and outcome satisfaction. *Journal of Consumer Research*, *35*(2), 202–215.
Morrin, M., Broniarczyk, S . M., & Inman, J . J. (2012). Plan format and participation in 401 (k) plans: The moderating role of investor knowledge. *Journal of Public Policy & Marketing*, *31*(2), 254–268. Fall.
Oppewal, H., & Koelemeijer, K. (2005). More choice is better: Effects of assortment size and composition on assortment evaluation. *International Journal of Research in Marketing*, *22*(1), 45–60.
Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge: Cambridge University Press.
Riley, R. D. (2009). Multivariate meta-analysis: The effect of ignoring within-study correlation. *Journal of the Royal Statistical Society, A*, *172*, 789–811.
Robinson, G. K. (1991). That blup is a good thing: The estimation of random effects. *Statistical Science*, *6*(1), 15–32.
Scheibehenne, B., Greifeneder, R., & Todd, P. M. (2009). What moderates the too-much-choice effect? *Psychology and Marketing*, *26*(3), 229–253.
Scheibehenne, B., Greifeneder, R., & Todd, P. M. (2010). Can there ever be too many options? A meta-analytic review of choice overload. *Journal of Consumer Research*, *37*(3), 409–425.
Schwartz, B. (2004). *The paradox of choice: Why more is less*. New York: Ecco.
Sela, A., Berger, J., & Liu, W. (2009). Variety, vice, and virtue: How assortment size influences option choice. *Journal of Consumer Research*, *35*(6), 941–951.
Shah, A. M., & Wolford, G. (2007). Buying behavior as a function of parametric variation of number of choices. *Psychological Science*, *18*(5), 369–370.
Townsend, C., & Kahn, B. E. (2014). The 'visual preference heuristic': The influence of visual versus verbal depiction on assortment processing, perceived variety, and choice overload. *Journal of Consumer Research*, *40*(5), 993–1015.

Manuscript Received: 31 AUG 2016
Final Version Received: 3 APR 2017
Published Online Date: 19 MAY 2017