Published in final edited form as:

Multivariate Behav Res. 2013 May 1; 48(3): 340–369. doi:10.1080/00273171.2013.784862.

A Bayesian Approach for Estimating Mediation Effects with Missing Data

Craig K. Enders

Department of Psychology, Arizona State University

Amanda J. Fairchild

Department of Psychology, University of South Carolina

David P. MacKinnon

Department of Psychology, Arizona State University

Abstract

Methodologists have developed mediation analysis techniques for a broad range of substantive applications, yet methods for estimating mediating mechanisms with missing data have been understudied. This study outlined a general Bayesian missing data handling approach that can accommodate mediation analyses with any number of manifest variables. Computer simulation studies showed that the Bayesian approach produced frequentist coverage rates and power estimates that were comparable to those of maximum likelihood with the bias-corrected bootstrap. We share a SAS macro that implements Bayesian estimation and use two data analysis examples to demonstrate its use.

Keywords

Mediation; indirect effects; missing data; Bayesian estimation; bias corrected bootstrap; Sobel test

The evaluation of mediating mechanisms has become a critical element of behavioral science research, not only to assess whether (and how) interventions achieve their effects, but also more broadly to understand the causes of behavioral change. The significance of mediation hypotheses is evident in a variety of substantive domains ranging from drug prevention (e.g., perceived social norms, a psychological mediator, transmits the impact of a prevention program on adolescent drug use; Liu, Flay, et al., 2009), to health promotion (e.g., knowledge of fruit and vegetable benefits, a behavioral mediator, transmits the effect of a health promotion program on weight loss; Elliot, Goldberg, Kuehl, Moe, Breger, & Pickering, 2007), to epidemiology (e.g., weight-related biomarkers, a biological mediator, transmit the effect of a diet and stress reduction intervention on tumor markers of prostate cancer; Saxe, Major, Westerberg, Khandrika, & Downs, 2008). The wide appeal of mediation analyses continues to drive the development of new analytic procedures for assessing mediation effects.

Methodologists have developed mediation analysis techniques for a broad range of substantive applications. However, methods for estimating mediating mechanisms with missing data have been understudied. The purpose of this study is to extend Yuan and MacKinnon's (2009) work on Bayesian mediation analyses to the missing data context.

Specifically, we outline a Bayesian estimation approach and provide a SAS macro that applies the method to mediation analyses with any number of manifest variables. Bayesian mediation analyses provide several advantages over conventional estimation methods. First, the Bayesian approach should yield more accurate inferences than frequentist significance testing approaches because it does not require distributional assumptions. This advantage is particularly important because the sampling distribution of the mediated effect can be quite nonnormal. Second, the Bayesian paradigm offers a more interpretable interval estimate than the frequentist approach, providing a credible interval in which the probability of the population parameter falling within the bounds is specified. Third, it is possible to improve the power to detect mediation effects with missing data by incorporating prior information about the indirect effect. Finally, the estimation procedure that we describe is straightforward to implement in SAS.

The organization of the manuscript is as follows. The paper begins with a review of both the single mediator model and Bayes' theorem. Next, we describe the use of a Markov chain Monte Carlo algorithm known as data augmentation to simulate the distribution of a mediation effect. Having outlined the procedural details, we use a data analysis example to demonstrate Bayesian estimation, and we use computer simulations to study its performance. Finally, we show how to perform a Bayesian analysis that incorporates prior knowledge about a mediation effect.

Mediation Analysis

A mediating variable (M) conveys the influence of an independent variable (X) on a dependent variable (Y), thereby illustrating the mechanism through which the two variables are related (Baron & Kenny, 1986; Judd & Kenny, 1981; MacKinnon, 2008). The basic mediation model posits that X influences M , which in turn influences Y , as follows

$$\widehat{M} = \alpha X \quad (1)$$

$$\widehat{Y} = \tau' X + \beta M \quad (2)$$

where α is the slope coefficient from the regression of M on X , β is the partial coefficient from the regression of Y on X controlling for M , and τ is the slope from the regression of Y on M controlling for X . For simplicity, we omit the intercepts from the previous regression equations because these parameters do not contribute to the mediation effect.

Methodologists have proposed a variety of mediation tests (e.g., causal steps, Baron & Kenny, 1986; the difference in coefficients – , MacKinnon, Warsi, & Dwyer, 1995), but contemporary work suggests that the product of coefficients is the most flexible estimator because it extends to complex scenarios that include multiple mediators or outcomes, multilevel data structures, and categorical outcomes (e.g., MacKinnon, Lockwood, Brown, Wang, & Hoffman, 2007). The logic of the product of the coefficients estimator derives from path analysis, where the total effect of X on Y is decomposed into indirect and direct components (Alwin & Hauser, 1975). This decomposition posits that the total effect of X on Y is comprised of a portion that is transmitted through the intervening variable (i.e., the indirect effect) and a portion that is not impacted by the intervening variable (i.e., the direct effect). This decomposition is given by

$$\tau = \tau' + \alpha\beta \quad (3)$$

where τ is the total effect of X on Y , τ' is the direct effect, and $\alpha\beta$ is the indirect effect. The indirect effect, or product of coefficients estimator, is the product of the regression slope that

relates the independent variable to the mediator (i.e., the γ path) and the partial regression coefficient that relates the mediator to the outcome (i.e., the β path). We henceforth refer to this estimator as $\beta\gamma$.

Researchers routinely use a normal-theory standard error derived from the multivariate delta method to evaluate the statistical significance of $\beta\gamma$ (e.g., Sobel, 1982). Dividing $\beta\gamma$ by a normal theory standard error yields a Wald z statistic, and referencing this test statistic to a standard normal distribution gives a probability value for the mediated effect. Recent work has shown, however, that normal-theory standard errors are inaccurate because the product of two normally distributed variables is not necessarily normally distributed; the sampling distribution of $\beta\gamma$ can be markedly asymmetric and kurtotic even when the β and γ regression coefficients have normal sampling distributions (MacKinnon, Fritz, Williams, & Lockwood, 2007; MacKinnon, Lockwood, & Williams, 2004; MacKinnon, Lockwood, Hoffman, West, & Sheets, 2002; Shrout & Bolger, 2002; Williams & MacKinnon, 2008). This distributional violation is especially problematic in small samples and with small effect sizes. The consequence of applying normal-theory significance tests is that the probability value from the Wald z test is conservative, limiting the power to detect mediation effects and increasing Type II error rates. The methodological literature currently favors asymmetric confidence limits based either on the theoretical sampling distribution of the product (MacKinnon et al., 2007; Meeker, Cornwell, & Aroian, 1981) or on empirical resampling methods (Efron & Tibshirani, 1993) because these approaches improve the accuracy of significance tests. The Bayesian approach that we outline in this manuscript naturally yields asymmetric intervals as a byproduct of estimation.

Mediation Analyses with Missing Data

The last decade has seen a noticeable shift to missing data techniques that assume a missing at random (MAR) mechanism, such that the probability of missing data on an incomplete variable is unrelated to the would-be values of that variable (e.g., other variables in the data predict missingness). Maximum likelihood estimation and multiple imputation are currently the principal MAR-based methods in social science applications, and the Bayesian approach that we outline is a third option. Maximum likelihood estimation largely accommodates familiar complete-data mediation procedures. For example, multivariate delta standard errors (i.e., the Sobel test) and asymmetric confidence limits based on the bootstrap are available in software packages such as Mplus.

Multiple imputation is arguably less flexible for mediation analyses. Standard imputation procedures (e.g., draw m imputed data sets, perform a mediation analysis m times, pool the estimates and standard errors) encourage researchers to implement and pool delta method standard errors. As noted previously, this approach performs poorly with complete data, and there is no reason to expect the situation to improve with missing data. Unfortunately, bootstrap procedures that methodologists currently favor do not translate well to multiple imputation. The most natural way to implement the bootstrap is to impute the data then apply the bootstrap to each imputed data set (i.e., impute-then-bootstrap). The procedural steps are as follows: (a) generate m imputed data sets, (b) draw b bootstrap samples from each data set, (c) form m empirical sampling distributions for $\beta\gamma$ by fitting the mediation model to the b bootstrap samples from each imputed data set, and (d) find the $\beta\gamma$ values that correspond to the .025 and .975 quantiles of each empirical sampling distribution, $q_{.025}^m$ and $q_{.975}^m$, respectively. Because each of the m sets of bootstrap samples originates from a complete parent data set, averaging these quantiles or otherwise basing inferences on the bootstrap samples is inappropriate because the empirical sampling distributions reflect the variation of complete-data $\beta\gamma$ estimates (i.e., the bootstrap sampling distributions are narrower than their missing-data counterparts).

A second way to implement the bootstrap with imputation is to draw b samples from the incomplete data set and impute each bootstrap sample (i.e., bootstrap-then-impute). The procedural steps are as follows: (a) draw b bootstrap samples from the incomplete parent data set, (b) generate $m - 1$ imputed data set per sample, (c) estimate from each imputed data set, and (d) form an empirical sampling distribution of the estimates (or of the average of the m estimates if $m > 1$). Again, the .025 and .975 quantiles of this empirical sampling distribution define a 95% confidence interval that is the basis for statistical inferences. Unlike impute-then-bootstrap, this procedure should yield an appropriate empirical sampling distribution because an incomplete parent data set generates the bootstrap samples. However, bootstrap-then-impute is computationally intensive and difficult to implement in existing software packages.

The previous discussion highlights the fact that combining imputation and the bootstrap is, at best, a cumbersome solution for mediation analyses with missing data. We believe that a Bayesian framework provides a natural solution for mediation analyses because it yields a 95% confidence interval (technically, a 95% credible interval) based on the quantiles of an empirical distribution of estimates. This is precisely what researchers are striving for with the bootstrap. Further, the simple SAS macro that we present later in the manuscript provides a convenient method for estimating a wide range of mediation models (e.g., single- or multiple-mediator models with any number of predictors or outcomes). Coupled with the ability to incorporate prior information (e.g., mediation estimates from pilot data or a published study), we believe that Bayesian mediation is an important tool for researchers.

Bayesian Estimation

This section briefly outlines Bayesian estimation. A number of resources provide accessible introductions to the Bayesian paradigm (Bolstad, 2007; Enders, 2010; Yuan & MacKinnon, 2009), and Bayesian texts provide additional technical details (Gelman, Carlin, Stern, & Rubin, 2009). The classical frequentist approach that predominates the social sciences defines a parameter as a fixed but unknown value in the population. Repeated sampling yields a distribution of estimates that vary around the true population value. In contrast, the Bayesian paradigm views a parameter as a random variable that has a distribution (i.e., there is no single true parameter value in the population). The basic goal of a Bayesian analysis is to use a priori information and the data to describe the shape of this parameter distribution (the posterior). The researcher specifies a probability distribution that summarizes prior knowledge about the parameter, and a likelihood function quantifies the data's evidence about the parameter. Collectively, these two data sources define a posterior distribution that describes the relative probability of different parameter values.

Bayes' theorem provides the mathematical machinery behind Bayesian estimation. The theorem is

$$p(\theta|Y) = \frac{p(\theta) p(Y|\theta)}{p(Y)} \quad (4)$$

where $p(\theta|Y)$ is the posterior distribution of a parameter θ , Y is the data, $p(\theta)$ is the prior distribution of θ , $p(Y|\theta)$ is the likelihood function, and $p(Y)$ is the marginal distribution of the data. Because the denominator of the theorem is an unnecessary scaling factor that makes the area under the posterior integrate (i.e., sum) to one, the key part of the theorem reduces to

$$p(\theta|Y) \propto p(\theta) p(Y|\theta). \quad (5)$$

In words, Equation 5 says that the relative probability of a particular parameter value is proportional to its prior distribution times the likelihood function. Conceptually, the posterior distribution is a weighted likelihood function, where the prior probabilities increase or decrease the height of the likelihood. For example, if the prior assigns a high probability to a particular value of β , the corresponding point on the posterior is more elevated than the likelihood function. Conversely, if the prior specifies a low probability to a particular value of β , the posterior distribution relies more heavily on the likelihood function.

Applying Bayesian Estimation to a Covariance Matrix

There are at least two ways to apply the previous ideas to a mediation analysis. In the complete-data context, Yuan and MacKinnon (2009) applied Bayes' theorem to the β and Σ parameters that define the product of coefficients estimator. Although this approach is also applicable to missing data analyses, we propose an alternate procedure that instead applies Bayesian principles to a covariance matrix. Working with a covariance matrix is convenient because it provides a straightforward mechanism for incorporating prior information about a mediation effect. As we show below, specifying a prior distribution requires an estimate of the covariance matrix (or alternatively, a correlation matrix) and a degrees of freedom value. Because an estimate is often available from pilot data, a published study, or a meta-analysis, we believe that working with a covariance matrix is often easier than working directly with β and Σ . The remainder of this section describes the application of Bayesian estimation to a covariance matrix, and interested readers can consult Yuan and MacKinnon (2009) for details on applying Bayesian methods to regression coefficients.

The first step of a Bayesian analysis is to specify a prior probability distribution for the covariance matrix. Assuming normally distributed population data, this distribution belongs to the inverse Wishart family, a multivariate generalization of the chi-square distribution. More specifically, the prior distribution of a covariance matrix is

$$p(\Sigma) \propto W^{-1}(df_p, \Lambda_p^{-1}) \quad (6)$$

where W^{-1} denotes the inverse Wishart, df_p is the degrees of freedom, and Λ_p is a sum of squares and cross products matrix. In words, Equation 6 describes a researcher's a priori beliefs about the relative probability of different covariance matrices arising from normally distributed population data. The parameters df_p and Λ_p are so-called hyperparameters that define the center (i.e., the location) and the spread (i.e., the scale) of the probability distribution, respectively.

As noted previously, a historical data source can determine the prior distribution's hyperparameters. For example, the sum of squares and cross products matrix from a pilot data set with identical measures of X , M , and Y can provide values for Λ_p . Similarly, a covariance matrix from a published study or a meta-analysis can provide the necessary information via a simple conversion

$$\Lambda_p = (N_p - 1) \Sigma_p \quad (7)$$

where Σ_p and N_p are the covariance matrix and sample size from the prior data source, respectively. The analogous conversion for a correlation matrix is

$$\Lambda_p = (N_p - 1) \widehat{S} R_p \widehat{S} \quad (8)$$

where $\mathbf{R}_p$ is the prior correlation matrix, and $\hat{\mathbf{S}}$ is a diagonal matrix that contains the sample standard deviations (e.g., MAR-based maximum likelihood estimates). Pre- and post-multiplying $\mathbf{R}_p$ by the sample standard deviations expresses the prior correlations as a covariance matrix with diagonal elements equal to the sample variances.

As explained previously, a Bayesian analysis is essentially a weighted estimation problem that combines information from the prior and the data. The df_p value in Equation 6 determines the influence of the prior on the final results. For example, specifying $df_p = 20$ is akin to saying that the prior distribution contributes 20 data points worth of information to the estimation process. Larger degrees of freedom values are defensible when a researcher has a great deal of faith in the prior data source, and smaller values are appropriate when a researcher lacks confidence in the prior (e.g., because the historical data source used a slightly different population of participants, employed less reliable measures, etc.). In the absence of prior information, a researcher can specify a non-informative (i.e., diffuse) prior distribution that effectively contributes no information to estimation. In this situation, the data determine the final analysis results. Incorporating prior information into a mediation analysis can enhance precision and narrow confidence intervals (Yuan & MacKinnon, 2009). Although this advantage is lost with a non-informative prior, Bayesian estimation is still quite useful because it provides a mechanism for evaluating mediation effects without invoking distributional assumptions for . We demonstrate this point later in the manuscript.

In the Bayesian framework, the prior and the data combine to produce an updated distribution (a posterior) that describes the relative probability of different parameter values. The posterior distribution of a covariance matrix is also an inverse Wishart

$$p(\Sigma|Y) \propto W^{-1}(df, \Lambda^{-1}) \quad (9)$$

where df is a degrees of freedom and sum of squares and cross products matrix again define the center and the spread of the distribution, respectively. The degrees of freedom parameter is the sum of the degrees of freedom from the prior and the data, as follows

$$df = df_p + df_D = df_p + (N_D - 1) \quad (10)$$

and similarly combines two sources of information¹.

$$\Lambda = \Lambda_p + \Lambda_D \quad (11)$$

Mediation Analyses via Data Augmentation

Deriving the posterior distribution in Equation 9 is difficult or impossible with missing data, and the observed-data distribution of is even more daunting. Markov chain Monte Carlo (MCMC) provides a simulation-based approach to estimating this distribution. Conceptually, an MCMC analysis is analogous to the bootstrap in the sense that it generates an empirical distribution for each parameter, but it does so by drawing a large number of parameter values from their respective posterior distributions; in the context of a missing data analysis, Schafer (1997) refers to this process as parameter simulation. The data augmentation algorithm for multiple imputation (Schafer, 1997; Tanner & Wong, 1987) happens to be well suited for estimating the distribution of with incomplete data. As

¹The spread of the posterior distribution also depends on mean vector (e.g., see Schafer, 1997, p. 152). For simplicity, we omit these terms from Equation 11 because they vanish with a non-informative prior for the mean vector.

described below, data augmentation uses regression equations to impute the missing values, and it subsequently draws a new mean vector and covariance matrix from their posterior distributions. In the context of multiple imputation, researchers are primarily interested in the imputed values and view the parameter draws as a superfluous byproduct of data augmentation. In contrast, we are interested in the parameter draws and view the imputations as a byproduct of estimation.

Data augmentation is a two-step algorithm that repeatedly cycles between an imputation step (I-step) and a posterior step (P-step). The I-step fills in the missing values with draws from a conditional distribution that depends on the observed data (i.e., the non-missing values) and the current estimates of μ and Σ . More formally, the I-step draws plausible score values from

$$\dot{Y}_t \tilde{p}(Y_{\text{mis}} | Y_{\text{obs}}, \dot{\theta}_{t-1}) \quad (12)$$

where $\dot{Y}_t$ represents an imputed value from iteration t (throughout the manuscript, we use a dot to denote a draw), Y_{mis} is the missing portion of the incomplete variable, Y_{obs} represents the observed data, and $\dot{\theta}_{t-1}$ contains the mean vector and covariance matrix from the previous iteration. Procedurally, data augmentation uses the elements in μ and Σ to construct regression equations that predict the incomplete variables from the complete variables, and it then draws an imputation for each missing value from a normal distribution that is centered a predicted score and has a variance equal to the residual variance from the regression model.

Having filled in the data, the P-step draws a new covariance matrix and mean vector from their respective complete-data posterior distributions. To begin, data augmentation uses the filled-in data set from the preceding I-step to estimate the sum of squares and cross products matrix, $\widehat{\Lambda}_t$. Combined with the prior, this matrix defines the posterior distribution of the covariance matrix, such that $\widehat{\Lambda}_t$ replaces Σ_D in Equation 9. The algorithm then uses Monte Carlo computer simulation to draw a new covariance matrix from the posterior distribution in Equation 9. This Monte Carlo procedure yields a simulation-based estimate of Σ that randomly differs from the covariance matrix that produced the imputation regression coefficients at the preceding I-step. We denote this draw as $\dot{\Sigma}_t$. The algorithm similarly draws a new mean vector from a multivariate normal distribution, the center and spread of which depends on mean estimates from the filled-in data and $\dot{\Sigma}_t$, respectively. We denote this draw as $\dot{\mu}_t$. Having completed a single computational cycle, data augmentation forwards the new parameter values to the next I-step, where the elements in $\dot{\mu}_t$ and $\dot{\Sigma}_t$ define the imputation regression equations for iteration $t + 1$.

The data augmentation algorithm repeatedly cycles between the I- and P-steps, often for thousands of iterations. The collection of $\dot{\mu}_t$ and $\dot{\Sigma}_t$ values forms an empirical posterior distribution for each element in the mean vector and the covariance matrix. Although the covariance matrix is not the substantive focus of a mediation analysis, the elements in $\dot{\Sigma}$ define the mediation model parameters. For example, the components of the product of coefficients estimator for a single-mediator model are as follows

$$\dot{\alpha} = \dot{\alpha}_{XM} / \dot{\alpha}_X^2 \quad (13)$$

$$\dot{\beta} = \left(\dot{\sigma}_{MY} \dot{\sigma}_X^2 - \dot{\sigma}_{XM} \dot{\sigma}_{XY} \right) / \left(\dot{\sigma}_X^2 \dot{\sigma}_M^2 - \dot{\sigma}_{XM}^2 \right) \quad (14)$$

and the direct effect is given by

$$\dot{\tau}' = \dot{\sigma}_{MY} \dot{\sigma}_X^2 - \dot{\sigma}_{XM} \dot{\sigma}_{XY} / \dot{\sigma}_X^2 \dot{\sigma}_M^2 - \dot{\sigma}_{XM}^2. \quad (15)$$

Although they may or may not be of interest, the covariance matrix elements also define the residual variances from the mediation model.

$$\dot{\sigma}_{\varepsilon_M}^2 = \dot{\sigma}_M^2 - \dot{\sigma}_{XM}^2 / \dot{\sigma}_X^2 \quad (16)$$

$$\dot{\sigma}_{\varepsilon_Y}^2 = \left(\dot{\sigma}_X^2 \dot{\sigma}_M^2 \dot{\sigma}_Y^2 - \dot{\sigma}_X^2 \dot{\sigma}_{MY} - \dot{\sigma}_M^2 \dot{\sigma}_{XY} - \dot{\sigma}_Y^2 \dot{\sigma}_{XM} + 2 \dot{\sigma}_{XM} \dot{\sigma}_{XY} \dot{\sigma}_{MY} \right) / \left(\dot{\sigma}_X^2 \dot{\sigma}_M^2 - \dot{\sigma}_{XM}^2 \right) \quad (17)$$

Analogous matrix expressions are applicable to multiple-mediator models.

Because each P-step yields values for the terms on the right side of the previous equations, data augmentation provides a simple mechanism for simulating draws from the posterior distribution of $\dot{\alpha}$. The steps are as follows: (a) generate a data augmentation chain comprised of several thousand (e.g., 10,000) computational cycles, (b) save the values in $\dot{\Sigma}_t$ at each P-step, and (c) use the elements in $\dot{\Sigma}_t$ to compute $\dot{\alpha} \beta_t$ from each cycle. The collection of $\dot{\alpha} \beta$ values forms a simulation-based posterior distribution for the mediation effect. Importantly, our procedure does not use the filled-in data sets to produce an inference about $\dot{\alpha} \beta$. Rather, the goal is to use familiar summary statistics to describe the center and spread of the posterior. For example, the mean or median $\dot{\alpha} \beta$ value describes the center of the posterior distribution and effectively serves the same purpose as a frequentist point estimate. Similarly, the posterior standard deviation is comparable to a frequentist standard error.

In addition to describing the center and spread of the posterior, the $\dot{\alpha} \beta$ values that correspond to the .025 and .975 quantiles of the posterior distribution, $q_{0.025}$ and $q_{0.975}$, define a credible interval that is analogous to a 95% confidence interval in the frequentist paradigm. The credible interval provides a straightforward mechanism for hypothesis testing. Consistent with standard practice, the mediation effect is statistically significant if this interval does not contain zero. Although the 95% credible interval functions like a frequentist confidence interval, it is important to note that its interpretation does not rely on a hypothetical process of drawing repeated samples from the population. Rather, the credible interval defines the range in which 95% of the parameter values fall. Importantly, inferences do not rely on the (often) untenable assumption that the estimates follow a normal distribution because the 95% credible interval is based on the empirical quantiles of the posterior distribution. This is an important advantage over maximum likelihood and multiple imputation procedures that utilize multivariate delta method standard errors (MacKinnon et al., 2007; MacKinnon, Lockwood, & Williams, 2004; MacKinnon et al., 2002).

Mediation via Bayesian Regression

Yuan and MacKinnon (2009) applied Bayesian estimation principles to the $\dot{\alpha}$ and regression coefficients, whereas we estimate a covariance matrix and subsequently solve for $\dot{\alpha}$ and $\dot{\beta}$. Because researchers can implement Yuan and MacKinnon's method in software packages such as Mplus (Muthén, 2010), it is important to consider how their approach differs from ours. With complete data, we would expect the two procedures to produce nearly identical point estimates, given the one-to-one relationship between a linear regression model and a saturated covariance matrix for multivariate normal outcomes. However, algorithmic differences cause these approaches to diverge with missing data. In

particular, the regression-based estimator can accommodate missing data on M or Y but requires complete data for X and covariates. Data augmentation can handle any missing data pattern. We briefly outline this issue below, and additional technical details are available elsewhere in the literature (Gelman et al., 2009; Schafer, 2007).

Yuan and MacKinnon (2009) employ a classic Gibbs sampler for Bayesian regression that repeatedly samples regressions coefficients and a residual variance from their respective posterior distributions (Gelman et al., 2009, Ch. 14). Importantly, the sampler follows standard regression procedures by fixing exogenous variables at their observed values. Defining explanatory variables as fixed effectively requires complete data for the predictors because the regression model does not incorporate a probability distribution for these variables (MAR-based imputation schemes necessarily require a distribution for incomplete variables). However, because the model assumes normally distributed residuals for M and Y , the sampler can accommodate missing values for the outcomes by adding a computational step that draws imputations from a normal distribution that conditions on the regression model parameters. However, implementing Yuan and MacKinnon's procedure in a path analytic framework (e.g., with Mplus 7) requires that either M or Y is complete because the likelihood is not defined for cases with missing data on both outcomes.

By modeling the population data as a multivariate normal distribution with an unstructured covariance matrix, Schafer's (1997) data augmentation algorithm effectively treats all variables as outcomes, regardless of their role in the analysis model. This parameterization is flexible for mediation analyses because the algorithm can generate imputations for any variable in the model. However, it is worth noting data augmentation model is more complex because it incorporates a probability distribution (and thus additional parameters) for the exogenous variables. This added model complexity can increase the number of iterations required to achieve convergence. All things being equal, the Gibbs sampler is likely to converge more quickly because it treats the mean vector and covariance matrix of the predictors as known constants across iterations, whereas data augmentation introduces stochastic variation for these parameters. Fortunately, algorithmic differences that result from model complexity are probably negligible in most cases because data augmentation often converges very quickly (Schafer, 1997).

As an aside, Yuan and MacKinnon's (2009) approach can be modified to accommodate missing exogenous variables by treating each incomplete predictor as the sole indicator of a pseudo-latent variable. By fixing the loading to one and the residual variance to zero, this approach effectively converts an explanatory variable to an outcome without altering its exogenous status in the model. Enders (2010, pp. 116-188) gives additional details on the latent variable approach for incomplete predictors.

Data Analysis Example 1

To illustrate Bayesian mediation via data augmentation, we analyzed data from a study that sought to reduce the risk for cardiovascular disease in firefighters (Moe et al., 2002; Ranby, MacKinnon, Fairchild, Elliot, Kuehl, & Goldberg, 2011). Due to the extreme physical demands of their job, firefighters have a markedly elevated risk for cardiovascular disease. Accordingly, the PHLAME ("Promoting healthy lifestyles: Alternative models' effects") intervention aimed to decrease risk factors in this group by promoting healthy eating and exercise habits. For this example, we consider the impact of a team-based treatment program (X) on intent to eat fruit and vegetables (Y) through two putative mediators: the extent to which coworkers eat fruit and vegetables (M_1) and knowledge of healthy eating benefits (M_2). Additionally, we used two auxiliary variables in the imputation model, friends' and partner's consumption of fruit and vegetables and partner's consumption of fruit and

vegetables. Figure 1 shows a path diagram of the mediation model. To avoid clutter, the residual covariance between the mediators is omitted, as are the auxiliary variable correlations. Table 1 gives maximum likelihood descriptive statistics for the analysis variables.

Although multiple imputation programs typically output the parameter draws required for computing mediation model parameters, the computations require considerable data manipulation, particularly for models with multiple mediators or multiple outcomes. To simplify the process, we created a SAS macro program that fully automates the Bayesian analysis. The macro uses the MI procedure in SAS to implement data augmentation and subsequently uses the IML procedure to compute the mediation model parameters. To illustrate the macro, Figure 2 shows the inputs for the data analysis example, with bold typeface denoting user-supplied values. The first line (%include) calls a SAS file named “Bayesian Mediation Macro.sas” that contains the main macro program. The remaining lines specify the input variables for the macro. To implement the macro, the user simply needs to execute the program in Figure 2 after altering the elements in bold typeface. The macro program and the data analysis programs are available for download at www.appliedmissingdata.com.

The macro requires four sets of inputs. The first set of inputs consists of the file path for the raw data, the input variable list, and a numeric missing value code. We assume that the input text file is in free format and with a single placeholder value for missing data. Note that the input data may contain variables that are not used in the analysis. The second set of inputs contains information about the prior distribution. Leaving these entries blank, as they are in the figure, invokes a standard non-informative prior. Later in the manuscript we demonstrate how to specify an informative prior distribution. The third set of inputs specifies the analysis variables. Variables in the mediation model are specified as X , M , or Y variables. Although our example is a relatively simple two-mediator model, the macro is flexible and can accommodate single- or multiple-mediator models within any number of X , M , and Y variables. Auxiliary variables (e.g., correlates of incomplete variables or predictors of missingness) contribute to the imputation process but are not part of the mediation model. Specifying which variables are complete or missing is unnecessary, but it is important to note that the linear imputation model in PROC MI assumes that the incomplete variables are multivariate normal. At a minimum, this assumption precludes the use of incomplete categorical variables with more than two categories; if complete, such variables can be used as X variables or as auxiliary variables, but they must be dummy or effect coded. The final set of inputs specifies the burn-in period, the number of iterations, the desired point estimate (median or mean), and a random number seed (the seed can be left blank, but specifying a value ensures that the program will return the same results each time it executes). The specifications in our example instruct the macro to discard the initial set of 500 burn-in iterations and form the posterior distributions from the ensuing 10,000 iterations. Note that our decision to use 10,000 iterations was somewhat arbitrary and had no material impact on the analysis results (point estimates and 95% interval limits from a 1,000-iteration chain were identical to the third decimal). Although recommendations vary, using 2,000 or fewer iterations is often sufficient (e.g., Gelman & Shirley, 2011; Gelman et al., 2009).

To facilitate convergence assessments, the macro produces trace plots of the mediation parameters from the burn-in period. A trace plot is a line graph that displays the iterations on the horizontal axis and the simulated parameter values on the vertical axis. To illustrate, Figure 3 shows a plot of the values for the indirect effect of the intervention on intentions via coworker eating habits. The absence of long-term upward or downward trends suggests that a stable posterior distribution generated the parameter values. For interested readers, a

number of resources further address the use of trace plots as a convergence diagnostic (Enders, 2010; Schafer, 1997; Schafer & Olsen, 1998).

Turning to the posterior distribution, the macro produces a histogram and a kernel density plot for each mediated effect. To illustrate, Figure 4 shows a plot of the values for the indirect effect of the intervention on intentions via coworker eating habits. In our example, this distribution was based on draws from the 10,000 iterations following the burn-in period. As seen in the figure, the posterior distribution was nonnormal, with skewness of $S = .46$ and excess kurtosis of $K = .59$. Again, the shape of the posterior underscores the problem with applying significance tests that assume a normal distribution (and thus a symmetric confidence interval) for (e.g., tests based on the multivariate delta method).

In a Bayesian analysis, researchers can use straightforward summary statistics to describe and evaluate mediation effects. For example, the mean and the median describe the center of the posterior distribution in a manner that is analogous to a frequentist point estimate. Although both values are legitimate estimates of the mediation effect, our computer simulations suggest that the median has a lower mean squared error, particularly when the sample size or the effect size is small (situations where the posterior distribution of is nonnormal). Our macro allows users to specify the mean or the median as a point estimate. In lieu of a formal significance test, the parameter values that correspond to the .025 and .975 quantiles of the posterior distribution form a credible interval that is analogous to a 95% confidence interval in the frequentist paradigm. The macro produces a table that includes the posterior median and the 95% credible interval bounds for each mediated effect.

Figure 5 shows the macro output from the example analysis. Considering the indirect effect of the intervention on intentions via coworker eating habits, the posterior median was $Mdn = .083$, and the 95% credible interval limits were $q_{0.025} = .0010$ and $q_{0.975} = .193$. Because this interval does not contain zero, we can conclude that the mediation effect is statistically significant. The corresponding absolute proportion mediated effect size was .264 (i.e., the absolute value of divided by the absolute value of the total effect; MacKinnon, 2009, pp. 82–83), indicating that approximately 26% of the effect of the team intervention on intent to eat fruit and vegetables was mediated through the social norm of coworkers eating fruit and vegetables. Turning to the indirect effect of the intervention on intentions via knowledge of healthy eating benefits, the posterior median was $Mdn = .063$, and the 95% credible interval limits were $q_{0.025} = -.00007$ and $q_{0.975} = .162$, and the corresponding absolute proportion mediated was .203. From a frequentist perspective, this indirect effect is non-significant because the credible interval contained zero. Importantly, both credible intervals are asymmetric around the posterior median, which owes to the positive skew of the posterior distributions.

The interpretation of component path coefficients in the mediation model may be of particular interest in program evaluation work where a researcher is interested in the action theory or conceptual theory underlying a program (Chen, 1990). The same summary statistics detailed above can describe the posterior distributions of $\dot{\alpha}$, $\dot{\beta}$, and $\dot{\tau}$ individually. As seen in Figure 5, the tabular output also includes numeric summaries for the component paths as well as for the total effect of X on Y .

Computer Simulations

We performed a series of Monte Carlo simulations to study the performance of Bayesian mediation. With one exception, the data-generating model was consistent with full mediation and did not include a direct effect of X on Y . We generated 5000 artificial data sets within each cell of a fully crossed design that manipulated three experimental

conditions: the missing data mechanism (missing completely at random and missing at random), the sample size ($N = 50, 100, 300$, and 1000), and the magnitude of the mediation effect. The mediation factor was comprised of five combinations of α and β values: $\alpha = 0$ and $\beta = 0$, $\alpha = .10$ and $\beta = .10$, $\alpha = .30$ and $\beta = .30$, $\alpha = .50$ and $\beta = .50$, and $\alpha = .30$ and $\beta = .10$. Because the simulation variables were standard normal, the α and β values are identical to Pearson correlations. Consequently, the mediation model parameters correspond to Cohen's (1988) effect size benchmarks (i.e., zero/zero, small/small, medium/medium, large/large, medium/small). Finally, note that the $\alpha = 0$ and $\beta = 0$ condition had a direct effect that was moderate in magnitude (i.e., $\gamma = .30$). The direct effect was zero in all other conditions.

We used the SAS IML procedure to generate three standard normal variables and subsequently used Cholesky decomposition to impose the desired correlation structure. Next, we imposed a 20% missing data rate on both M and Y . In the missing completely at random (MCAR) simulation, we used uniform random numbers to independently impose missing data, such that cases with $u_i < .20$ had a missing value on M or Y . In the MAR simulation, the value of X determined missingness on M and Y . Beginning with the highest X value and working in descending order, we independently deleted M and Y scores with a .75 probability until 20% of the values on each variable were missing. This produced a situation where cases with high scores on X had higher missing data rates on M and Y . Note that we chose not to manipulate the amount of missing data because this design characteristic tends to produce predictable and uninteresting findings (e.g., as the missing data rate increases, power decreases). Rather, we chose a constant 20% missing data rate because we felt that it was extreme enough to expose any problems.

We used the MI procedure in SAS version 9.2 to implement data augmentation. Using the standard non-informative prior distributions for μ and Σ , we generated 10,500 cycles of data augmentation and saved the simulated parameter estimates from each P-step. Next, we used the formulas from Equations 13 and 14 to convert the elements in $\hat{\Sigma}_t$ into mediation model parameters. For each of the 200,000 replications (2 missing data mechanisms by 4 sample size conditions by 5 effect size conditions by 5000 iterations), we discarded the parameter draws from the initial 500 burn-in cycles and formed an empirical distribution of 10,000 parameter draws. To be complete, we also used Mplus 7 to implement the regression-based Gibbs sampler from Yuan and MacKinnon (2009). There is no reason to expect differences between data augmentation and the Gibbs sampler because these are simply two algorithmic approaches to the same estimation problem. It is nevertheless useful to calibrate our method to an existing software package.

To assess the relative performance of the Bayesian approach, we also used maximum likelihood missing data handling and multiple imputation to estimate the mediation effect. These are the principle MAR-based analysis approaches that enjoy widespread use in substantive applications. To implement maximum likelihood estimation, we input the raw data files to Mplus 6.1 and used the MODEL INDIRECT command to estimate γ . Mplus offers a variety of significance testing options for mediation effects, and we examined three different possibilities: normal-theory maximum likelihood with multivariate delta method (i.e., Sobel) standard errors, the bias corrected bootstrap, and the percentile bootstrap. Although the methodological literature has discounted significance tests based on normal-theory standard errors, we chose this option because it likely reflects the analytic practice of many substantive researchers, especially those who use statistical software packages such as SPSS or SAS.

We used the MI procedure in SAS version 9.2 to implement multiple imputation. Using the standard non-informative prior distributions for μ and Σ , we generated 100 imputed data sets from a data augmentation chain with 200 burn-in and 200 between-imputation cycles; we

chose the relatively large number of imputed data sets in order to maximize power of the subsequent significance tests (Graham, Olchowski, & Gilreath, 2007). We then used the REG procedure to estimate the component regression coefficients (see Equations 1 and 2) and subsequently computed β and its normal-theory (i.e., Sobel) standard error from each of the 100 sets of imputations. Finally, we used Rubin's (1987) pooling rules to generate a point estimate, standard error, and confidence interval for each of the 200,000 artificial data sets.

The outcome variables for the simulations were relative bias, power, and confidence interval coverage. Because the Bayesian framework defines a parameter as a random variable rather than a fixed quantity, frequentist definitions of power and confidence interval coverage do not apply. Nevertheless, we felt that it was important to evaluate the Bayesian procedure from a frequentist perspective in order to make comparisons with other MAR-based procedures. To quantify power for the Bayesian analyses, we obtained the .025 and the .975 quantiles from each posterior distribution and defined power as the proportion of replications where the 95% credible interval did not include zero. We applied the same procedure to the maximum likelihood bootstrap procedures. For maximum likelihood and multiple imputation with normal-theory test statistics, the proportion of replications within each design cell that produced a statistically significant β coefficient served as an empirical estimate of power.

For Bayesian (and bootstrap) confidence interval coverage, we used the .025 and the .975 quantiles from each posterior (sampling) distribution to compute the proportion of the credible intervals within each design cell that included the population parameter from the data-generating model. For maximum likelihood and multiple imputation with normal-theory standard errors, coverage was the proportion of replications where the 95% confidence interval included the true population value. In the frequentist paradigm, confidence interval coverage value directly relates to the Type I error rate, such that 95% coverage corresponds to the nominal 5% Type I error rate, a 90% coverage rate corresponds to a twofold increase in Type I errors, and so on.

Simulation Results

As expected with MCAR and MAR mechanisms, all missing data handling approaches produced relatively accurate point estimates with no appreciable bias. Consequently, we limit our subsequent presentation to power and confidence interval coverage. Further, because the results from the MCAR and MAR simulations were virtually identical, we limit the presentation to the MAR condition.

Table 2 displays the coverage values for each design cell in the MAR simulation. The table illustrates that the effect size and the sample size influenced coverage rates and that the sample size moderated the impact of effect size, such that larger sample sizes yielded more accurate coverage. The $\beta = 0$ and $\gamma = 0$ condition produced coverage rates that exceeded the nominal .95 rate at every sample size, indicating conservative inferences. The same was true for the condition where both coefficients had a small effect size (i.e., $\beta = .10$ and $\gamma = .10$), although the coverage rates approached the nominal level as the sample size increased. Finally, when either β or γ (or both) had at least a medium effect size, Bayesian estimation and maximum likelihood estimation with bootstrap standard errors produced accurate coverage rates, particularly with sample sizes of 100 or larger. In contrast, coverage rates based on multivariate delta (i.e., Sobel) standard errors were not as accurate at smaller sample sizes and only approached the nominal level at larger N s. In the complete-data context, Yuan and MacKinnon (2009) compared Bayesian and frequentist coverage values and found a very similar pattern of results.

Table 3 gives empirical power estimates for the five methods. For Bayesian estimation and maximum likelihood estimation with bootstrap standard errors, these estimates reflect the proportion of replications where the credible (or confidence) interval did not contain zero, whereas the power values for the normal-theory test statistics reflect the proportion of replications that produced a statistically significant test statistic. As seen in the table, the bias corrected bootstrap produced the highest power values, followed by Bayesian estimation and the percentile bootstrap, respectively. Power values for the normal-theory (i.e., Sobel) test were noticeably lower, except in situations where the sample size or the effect size was large.

Collectively, the results in Tables 2 and 3 suggest that Bayesian estimation yields coverage and power values that are comparable to the best maximum likelihood approach (maximum likelihood estimation with the bias corrected bootstrap). As expected, Bayesian estimation was superior to normal-theory significance tests based on multivariate delta method standard errors. It is important to reiterate that we implemented standard non-informative prior distributions in the simulations. Bayesian estimation can produce a substantial power advantage (i.e., narrower confidence intervals) when using informative prior distributions (Yuan & MacKinnon, 2009). We address this topic in the next section. Finally, note that data augmentation and the Gibbs sampler (two algorithmic approaches for implementing Bayesian estimation) were virtually identical, as expected.

Data Analysis Example 2

Thus far, we have focused on non-informative prior distributions. In the context of complete-data mediation analyses, Yuan and MacKinnon (2009) showed that an informative prior could provide a substantial increase in statistical power. This is particularly useful for behavioral science applications where mediation effects are often modest in magnitude. In this section, we use our SAS macro to illustrate a mediation analysis with an informative prior distribution. To keep the example simple, we consider a single-mediator model involving the indirect effect of the team intervention on intentions via coworker eating habits. Further, we omit the auxiliary variables from the previous example.

As described previously, implementing an informative prior distribution requires an a priori estimate of the covariance or correlation matrix and a degrees of freedom value (e.g., from pilot data, a published study, or a meta-analysis). For the purposes of illustration, suppose that the following prior estimates of the covariance and mean vector are available from pilot data.

$$\Sigma_p = \begin{bmatrix} .250 & .188 & .075 \\ .188 & 2.250 & .450 \\ .075 & .450 & 2.250 \end{bmatrix}$$

$$\mu_p = \begin{bmatrix} .50 \\ 3.50 \\ 3.25 \end{bmatrix}$$

The MI procedure (and thus our macro) uses a `_TYPE_ = COV` data set containing estimates and μ to define the prior distributions (the program automatically converts the matrices into the necessary hyperparameters). Our macro program requires a file path for this data set (in text format) as well as a corresponding variable list. To illustrate, Figure 6 shows the macro inputs for the analysis example, and Figure 7 shows the contents of the `_TYPE_ =`

COV text data set containing the prior parameters. Importantly, the variable list for the prior data must have `_TYPE_` and `_NAME_` as the first two variables, and the format of the data file itself should not deviate from that in Figure 7. The `N` and `N_MEAN` rows of the input data set contain the degrees of freedom values for the covariance matrix and mean vector, respectively. Although the decision is subjective, we arbitrarily assigned 30 degrees of freedom to the prior distribution, which effectively meant that the prior contributed roughly 30% as much information as the data. Finally, note that the values in the `MEAN` and `N_MEAN` rows can be set at zero if mean estimates are not available. Such would be the case when using Equation 8 to convert a correlation matrix from a published study that uses measures with different metrics.

Consistent with the earlier example, we generated a 10,500-cycle data augmentation chain and discarded the initial set of 500 burn-in iterations. For comparison purposes, we estimated the single-mediator model with a standard non-informative prior. The analysis produced a posterior median of $Mdn = .079$ and a 95% credible interval bounded by values of $q_{0.025} = -.003$ and $q_{0.975} = .192$. Again, notice that this interval was asymmetric around the center of the distribution because the posterior was positively skewed. Importantly, implementing an informative prior reduced the variance of the posterior distribution and therefore narrowed the width of the credible interval (i.e., increased power). Specifically, the posterior median was $Mdn = .085$, and the 95% credible interval was bounded by values of $q_{0.025} = .011$ and $q_{0.975} = .188$. Notice that the mediation effect is statistically different from zero because the credible interval no longer contains zero. To visually illustrate the analysis results, Figure 8 shows kernel density plots from the single-mediator model. The plots clearly demonstrate that the posterior distribution based on an informative prior is narrower than that of the non-informative prior. In frequentist terms, the reduction in variance translates into greater power.

Discussion

Assessing mediating mechanisms is a critical component of behavioral science research. Methodologists have developed mediation analysis techniques for a broad range of substantive applications, but methods for missing data have thus far gone unstudied. Consequently, the purpose of this study was to outline a Bayesian approach for estimating and evaluating mediation effects with missing data. Our approach is closely related to that of Yuan and MacKinnon (2009) in the complete-data context but applies Bayesian machinery to a covariance matrix rather than to the γ and β coefficients themselves. In our view, specifying the mediation model in terms of Σ has important advantages. First, working with a covariance matrix simplifies the process of specifying a prior distribution because researchers need only have an a priori estimate of Σ or $\mathbf{R}$. Suitable estimates are often available from pilot data, published research studies, or meta-analyses. Having a convenient mechanism for specifying prior distributions is particularly important because it can partially compensate for the reduced precision inherent with missing data. Second, and perhaps most importantly, applying data augmentation to a covariance matrix accommodates missing values on any variable in the mediation model. As noted previously, the classic Gibbs sampler for Bayesian regression (Yuan & MacKinnon, 2009) is more restrictive in this regard.

Even with a non-informative prior, our simulation results suggest that Bayesian estimation yields frequentist coverage values and power estimates that are comparable to maximum likelihood estimation with the bias corrected bootstrap; based on the complete-data literature, we expected the bias corrected bootstrap to be the gold standard against which to compare the Bayesian method. As expected, Bayesian estimation was superior to normal-theory significance tests, particularly when the sample size or the effect size of either the

or paths was small. These findings align with those of Yuan and MacKinnon (2009) in the complete-data context. Although we did not pursue simulations with informative prior distributions, past research and statistical theory predict that Bayesian estimation would produce narrower confidence intervals and greater precision than frequentist alternatives, including the bootstrap. Our second analysis example supports this assertion.

Our study suggests a number of avenues for future research. First, the lack of existing research on this topic prompted us to limit the simulations to a single-mediator model. Although this model is exceedingly common in the behavioral sciences, future studies should investigate models with multiple mediators and outcomes. Our SAS macro is general and can accommodate these additional complexities. Second, our approach assumes that the incomplete variables are multivariate normal. The literature suggests that data augmentation and the bootstrap perform well with nonnormal incomplete variables (Demirtas, Freels, & Yucel, 2008; Enders, 2001; Graham & Schafer, 1999). Further, flexible MCMC methods are now available for mixtures of categorical and continuous variables (Enders, 2012; Goldstein, Carpenter, Kenward, & Levin, 2009; van Buuren, 2007), and these could easily be adapted to accommodate mediation models. Investigating the issue of nonnormal data would be a fruitful avenue for future research, particularly given the recent interest in mediation models with categorical outcomes (MacKinnon et al., 2007). Fourth, recent research suggests that the bias corrected bootstrap yields elevated Type I errors under particular sample size and effect size configurations (Fritz, Taylor, & MacKinnon, 2012). Although we did not observe between-method differences in Type I errors (for brevity, we did not report the Type I error rates from the $\alpha = 0$ and $\beta = 0$ design cells), we did not examine the combination of conditions that produce elevated Type I errors. Future studies should explore the differences between Bayesian estimation and the bias corrected bootstrap in these problematic scenarios. Finally, developing missing data methods for multilevel mediation is an important area of future research. Although our approach could be extended to multilevel data by sampling separate covariance matrices at level-1 and level-2, it would necessarily be limited to certain classes of models (e.g., models with only random intercepts). Consistent with its single-level counterpart, the Gibbs sampler for multilevel models (e.g., Yuan and MacKinnon, 2009) is an imperfect solution because it cannot accommodate incomplete predictor variables or missing data patterns where both M and Y are incomplete. Developing a Bayesian approach to multilevel mediation that could handle general missing data patterns would be an important contribution to the literature.

In sum, Bayesian estimation provides a straightforward mechanism for estimating and evaluating mediation effects with missing data. Our simulations suggest that the Bayesian approach should yield comparable results to the best possible maximum likelihood analysis (maximum likelihood with bias corrected bootstrap). When researchers have access to a priori information about a mediation effect (e.g., from pilot data, a published study, or a meta-analysis), Bayesian estimation can provide a substantial increase in power. We believe that specifying the prior information in terms of a covariance or a correlation matrix will allow researchers to take advantage of this important benefit.

Acknowledgments

This article was supported in part by National Institute on Drug Abuse grant DA09757 (David MacKinnon, PI). Amanda Fairchild is supported by National Institute on Drug Abuse grant 1R01DA030349-01A1 (Amanda Fairchild, PI).

References

- Alwin DF, Hauser RM. The decomposition of effects in path analysis. *American Sociological Review*. 1975; 40:37–47.

- Baron RM, Kenny DA. The moderator-mediator variable distinction in social psychological research: conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology*. 1986; 51:1173–1182. [PubMed: 3806354]
- Bolstad, WM. Introduction to Bayesian statistics. 2nd Ed. Wiley; New York: 2007.
- Chen, HT. Theory-driven evaluations. Sage; Newbury Park, CA: 1990.
- Cohen, J. Statistical power analysis for the behavioral sciences. 2nd Ed. Erlbaum; Hillsdale, NJ: 1988.
- Demirtas H, Freels SA, Yucel RM. Plausibility of multivariate normality assumption when multiple imputing non-Gaussian continuous outcomes: A simulation assessment. *Journal of Statistical Computation and Simulation*. 2008; 78:69–84.
- Efron, B.; Tibshirani, RJ. An Introduction to the Bootstrap. Chapman & Hall; New York: 1993.
- Elliot DL, Goldberg L, Kuehl KS, Moe EL, Breger RKR, Pickering MA. The PHLAME (Promoting Healthy Lifestyles: Alternative Models' Effects) firefighter study: Outcomes of two models of behavior change. *Journal of Occupational and Environmental Medicine*. 2007; 49(2):204–213. [PubMed: 17293760]
- Enders CK. A chained equations approach for imputing single and multilevel data with categorical and continuous variables. 2012 Manuscript in preparation.
- Enders, CK. Applied missing data analysis. Guilford Press; New York: 2010.
- Enders CK. The impact of nonnormality on full information maximum likelihood estimation for structural equation models with missing data. *Psychological Methods*. 2001; 6:352–370. [PubMed: 11778677]
- Fritz MS, Taylor AB, MacKinnon DP. Explanation of two anomalous results in statistical mediation analysis. *Multivariate Behavioral Research*. 2012; 47:61–87.
- Gelman, A.; Carlin, JB.; Stern, HS.; Rubin, DB. Bayesian data analysis. 2nd Ed. Chapman & Hall; Boca Raton, FL: 2009.
- Gelman, A.; Shirley, K. Inference from simulations and monitoring convergence. In: Brooks, S.; Gelman, A.; Jones, G.; Meng, XL., editors. *Handbook of Markov Chain Monte Carlo*. CRC Press; Boca Raton, FL: 2011.
- Goldstein H, Carpenter J, Kenward MG, Levin KA. Multilevel models with multivariate mixed response types. *Statistical Modelling*. 2009; 9:173–197.
- Graham JW, Olchowski AE, Gilreath TD. How many imputations are really needed? Some practical clarifications of multiple imputation theory. *Prevention Science*. 2007; 8:206–213. [PubMed: 17549635]
- Graham, JW.; Schafer, JL. On the performance of multiple imputation for multivariate data with small sample size. In: Hoyle, R., editor. *Statistical strategies for small sample research*. Sage; Thousand Oaks, CA: 1999. p. 1-29.
- Judd CM, Kenny DA. Process analysis: estimating mediation in treatment evaluations. *Evaluation Review*. 1981; 5:602–619.
- Liu LC, Flay BR, et al. Evaluating Mediation in Longitudinal Multivariate Data: Mediation effects for the Aban Aya Youth Project drug prevention program. *Prevention Science*. 2009; 10:197–207. [PubMed: 19288196]
- MacKinnon, DP. Introduction to statistical mediation analysis. Erlbaum; Mahwah, NJ: 2008.
- MacKinnon DP, Fritz MS, Williams J, Lockwood CM. Distribution of the product confidence limits for the indirect effect: Program PRODCLIN. *Behavior Research Methods*. 2007; 39:384–389. [PubMed: 17958149]
- MacKinnon DP, Lockwood CM, Brown, Wang W, C.H. Hoffman JM. The intermediate endpoint effect in logistic and probit regression. *Clinical Trials*. 2007; 4:499–513. [PubMed: 17942466]
- MacKinnon DP, Lockwood CM, Hoffman JM, West SG, Sheets V. A comparison of methods to test mediation and other intervening variable effects. *Psychological Methods*. 2002; 7:83–104. [PubMed: 11928892]
- MacKinnon DP, Lockwood CM, Williams J. Confidence limits for the indirect effect: Distribution of the product and resampling methods. *Multivariate Behavioral Research*. 2004; 39:99–128. [PubMed: 20157642]

- MacKinnon DP, Warsi G, Dwyer JH. A simulation study of mediated effect measures. *Multivariate Behavioral Research*. 1995; 30(1):41–62. [PubMed: 20157641]
- Meeker, WQ.; Cornwell, LW.; Aroian, LA. The product of two normally distributed random variables. In: Kenney, W.; Odeh, R., editors. *Selected tables in mathematical statistics*. Vol. Vol. VII. American Mathematical Society; Providence, IR: 1981. p. 1-256.
- Moe EL, Elliot DL, Goldberg L, et al. Promoting Healthy Lifestyles: Alternative Models' Effects (PHLAME). *Health Education Research*. 2002; 17:586–596. [PubMed: 12408203]
- Muthén, BO. Bayesian analysis in Mplus: A brief introduction. 2010. Retrieved from <http://statmodel.com/download/IntroBayesVersion%203.pdf>
- Ranby K, MacKinnon DP, Fairchild AJ, Elliot DL, Kuehl KS, Goldberg L. The PHLAME (Promoting Healthy Lifestyles: Alternative Models' Effects) Firefighter Study: Testing Mediating Mechanisms. *Journal of Occupational Health Psychology*. 2011; 16:501–513. [PubMed: 21728433]
- Rubin, DB. Multiple imputation for nonresponse in surveys. Wiley; Hoboken, NJ: 1987.
- Saxe GA, Major JM, Westerberg L, Khandrika S, Downs TM. Biological Mediators of Effect of Diet and Stress Reduction on Prostate Cancer. *Integrative Cancer Therapies*. 2008; 7:130–138. [PubMed: 18815144]
- Schafer, JL. Analysis of incomplete multivariate data. Chapman & Hall; Boca Raton, FL: 1997.
- Schafer JL, Olsen MK. Multiple imputation for multivariate missing-data problems: A data analyst's perspective. *Multivariate Behavioral Research*. 1998; 33:545–571.
- Shrout PE, Bolger N. Mediation in experimental and nonexperimental studies: New procedures and recommendations. *Psychological Methods*. 2002; 7(4):422–445. [PubMed: 12530702]
- Sobel ME. Asymptotic confidence intervals for indirect effects in structural equation models. *Sociological Methodology*. 1982; 13:290–312.
- Tanner MA, Wong WH. The calculation of posterior distributions by data augmentation. *Journal of the American Statistical Association*. 1987; 82:528–540.
- van Buuren S. Multiple imputation of discrete and continuous data by fully conditional specification. *Statistical Methods in Medical Research*. 2007; 16:219–242. [PubMed: 17621469]
- Williams J, MacKinnon DP. Resampling and distribution of the product methods for testing indirect effects in complex models. *Structural Equation Modeling*. 2008; 15:23–51. [PubMed: 20179778]
- Yuan Y, MacKinnon DP. Bayesian mediation analysis. *Psychological Methods*. 2009; 14:301–322. [PubMed: 19968395]

Figure 1.
Mediation model for Data Analysis Example 1.

```

%include 'c:\temp\bayesian mediation macro.sas';

* rawdata = file path of input text file;
* varlist = variable order in input text file;
* missingcode = missing value code;
* prior = file path of _TYPE_ = COV data;
* priorvarlist = variable order in _TYPE_ = COV data;
* x = x variables and covariates;
* m = mediators;
* y = outcome variables;
* auxiliary = missing data auxiliary variables;
* pointest = point estimate (median or mean);
* burnin = number of burn-in iterations;
* iterations = number of mcmc iterations;
* seed = random number seed;

%mediation(
rawdata = c:\temp\phlame.dat,
varlist = txgroup coworker knowledge intentions friends partner,
missingcode = -99,
prior = ,
priorvarlist = ,
x = txgroup,
m = coworker knowledge,
y = intentions,
auxiliary = friends partner,
pointest = median,
burnin = 500,
iterations = 10000,
seed = 20319);

```

Figure 2.

SAS macro input for Data Analysis Example 1. The values in bold typeface denote user-specified values.

Figure 3. Trace plot of the ab values from the first 500 cycles of data augmentation. The absence of long-term upward or downward trends suggests that a stable posterior distribution generated the simulated parameter values.

Figure 4. Kernel density plot of the posterior distribution from 10,000 cycles of data augmentation.

SUMMARY OF POSTERIOR DISTRIBUTIONS			
Total Effect of X on Y (Tau Paths)			
Estimate	Median	LCL	UCL
txgroup -> intentions	0.323235	-.054293	0.696910
Direct Effect of X on M (A Paths)			
Estimate	Median	LCL	UCL
txgroup -> coworker	0.533389	0.282380	0.785818
txgroup -> knowledge	0.197340	0.000523	0.394753
Direct Effect of M on Y Controlling for X (B Paths)			
Estimate	Median	LCL	UCL
coworker -> intentions	0.163396	0.001897	0.323125
knowledge -> intentions	0.344401	0.143104	0.542487
Direct Effect of X on Y Controlling for M (Tau Prime Paths)			
Estimate	Median	LCL	UCL
txgroup -> intentions	0.166932	-.216748	0.552133
Indirect Effect of X on Y via M (AB Paths)			
Estimate	Median	LCL	UCL
txgroup -> coworker -> intentions	0.083065	0.000997	0.192888
txgroup -> knowledge -> intentions	0.063327	-.000073	0.161574
Proportion Mediated Effect Sizes			
Estimate	Median	LCL	UCL
txgroup -> coworker -> intentions	0.264484	0.024072	3.770156
txgroup -> knowledge -> intentions	0.202989	0.019962	2.500836

Figure 5.
Macro output from Data Analysis Example 1.

```

%include 'c:\temp\bayesian mediation macro.sas';

* rawdata = file path of input text file;
* varlist = variable order in input text file;
* missingcode = missing value code;
* prior = file path of _TYPE_ = COV data;
* priorvarlist = variable order in _TYPE_ = COV data;
* x = x variables and covariates;
* m = mediators;
* y = outcome variables;
* auxiliary = missing data auxiliary variables;
* pointest = point estimate (median or mean);
* burnin = number of burn-in iterations;
* iterations = number of mcmc iterations;
* seed = random number seed;

%mediation(
rawdata = c:\temp\phlame.dat,
varlist = txgroup coworker knowledge intentions friends partner,
missingcode = -99,
prior = c:\temp\ex2priorcov.dat,
priorvarlist = _type_ $ _name_ $ txgroup coworker intentions,
x = txgroup,
m = coworker,
y = intentions,
auxiliary = ,
pointest = median,
burnin = 500,
iterations = 10000,
seed = 20319);

```

Figure 6.

SAS macro input for Data Analysis Example 2. The values in bold typeface denote user-specified values.

COV	txgroup	0.2500	0.1875	0.0750
COV	coworker	0.1875	2.2500	0.4500
COV	intentions	0.0750	0.4500	2.2500
N	.	30.0000	30.0000	30.0000
MEAN	.	0.5000	3.5000	3.2500
N_MEAN	.	30.0000	30.0000	30.0000

Figure 7.

Contents of the _TYPE_ = COV text data set containing the prior parameters for Data Analysis Example 2. The macro in Figure 6 references this data set as c:\temp\ex2priorcov.dat.

Figure 8.

Kernel density graph of the posterior distribution of $\alpha\beta$ from 10,000 cycles of data augmentation. The dashed distribution reflects a non-informative prior, and the solid distribution is based on an informative prior with 30 degrees of freedom.

Table 1

Maximum Likelihood Descriptives from Data Analysis Example 1

	1.	2.	3.	4.	5.	6.
1. Friends' healthy eating habits	1.000	0.488	-0.009	0.333	-0.059	0.299
2. Partner's healthy eating habits	0.243	1.000	0.078	0.818	0.057	0.571
3. Treatment program (X)	-0.016	0.093	1.000	0.129	0.048	0.079
4. Coworker eating habits (M_1)	0.234	0.404	0.220	1.000	0.024	0.265
5. Knowledge of health benefits (M_2)	-0.055	0.037	0.108	0.022	1.000	0.298
6. Healthy eating intentions (Y)	0.169	0.227	0.108	0.150	0.221	1.000
Mean	3.329	3.858	0.569	3.642	6.237	2.951
Std. Dev.	1.188	1.693	0.491	1.196	0.910	1.484
% Missing	0.456	22.096	0.000	21.185	21.640	43.052

Note. Covariances are on the upper diagonal and correlations are on the lower diagonal.

Table 2

Confidence Interval Coverage from the MAR Simulation

N		Skew.	Kurt.	Bayes (DA)	Bayes (Gibbs)	MI (Sobel)	ML (Sobel)	ML (BCB)	ML (NB)
= 0, = 0									
50	0	-0.009	4.265	0.999	0.999	0.999	1.000	0.993	0.997
100	0	0.012	4.079	0.998	0.998	1.000	1.000	0.992	0.997
300	0	0.009	3.970	0.999	0.999	1.000	1.000	0.993	0.998
1000	0	0.035	3.964	1.000	1.000	1.000	1.000	0.986	0.997
= .10, = .10									
50	0.01	0.119	3.847	0.996	0.997	0.999	0.986	0.981	0.994
100	0.01	0.257	3.290	0.995	0.995	0.999	0.970	0.966	0.991
300	0.01	0.522	2.224	0.978	0.980	0.965	0.909	0.914	0.963
1000	0.01	0.612	0.894	0.950	0.952	0.928	0.920	0.954	0.939
= .30, = .30									
50	0.09	0.559	1.854	0.952	0.961	0.934	0.900	0.932	0.943
100	0.09	0.587	0.993	0.947	0.948	0.918	0.909	0.948	0.939
300	0.09	0.437	0.322	0.951	0.952	0.936	0.936	0.954	0.947
1000	0.09	0.255	0.093	0.953	0.954	0.950	0.951	0.950	0.948
= .50, = .50									
50	0.25	0.516	0.717	0.952	0.961	0.935	0.922	0.949	0.938
100	0.25	0.414	0.328	0.945	0.947	0.936	0.930	0.944	0.941
300	0.25	0.257	0.102	0.955	0.955	0.952	0.947	0.952	0.949
1000	0.25	0.144	0.029	0.950	0.950	0.946	0.949	0.947	0.946
= .30, = .10									
50	0.03	0.280	2.699	0.980	0.968	0.996	0.959	0.930	0.970
100	0.03	0.377	1.747	0.968	0.972	0.978	0.947	0.917	0.956
300	0.03	0.351	0.658	0.951	0.951	0.952	0.945	0.935	0.944
1000	0.03	0.220	0.193	0.948	0.950	0.952	0.949	0.941	0.942

Note. DA = data augmentation, Gibbs = Gibbs sampler (Mplus), MI = multiple imputation, ML = maximum likelihood, Sobel = Sobel z test, BCB = bias corrected bootstrap, NB = naive bootstrap.

Table 3

Empirical Power Estimates from the MAR Simulation

<i>N</i>		Skew.	Kurt.	Bayes (DA)	Bayes (Gibbs)	MI (Sobel)	ML (Sobel)	ML (BCB)	ML (NB)
= .10, = .10									
50	0.01	0.119	3.847	0.005	0.003	0.001	0.001	0.018	0.007
100	0.01	0.257	3.290	0.012	0.010	0.001	0.002	0.034	0.015
300	0.01	0.522	2.224	0.066	0.062	0.016	0.018	0.111	0.059
1000	0.01	0.612	0.894	0.497	0.496	0.273	0.291	0.539	0.429
= .30, = .30									
50	0.09	0.559	1.854	0.123	0.095	0.057	0.065	0.213	0.134
100	0.09	0.587	0.993	0.427	0.409	0.260	0.271	0.542	0.425
300	0.09	0.437	0.322	0.974	0.974	0.951	0.953	0.981	0.973
1000	0.09	0.255	0.093	1.000	1.000	1.000	1.000	1.000	1.000
= .50, = .50									
50	0.25	0.516	0.717	0.665	0.618	0.595	0.563	0.744	0.645
100	0.25	0.414	0.328	0.970	0.968	0.961	0.953	0.978	0.968
300	0.25	0.257	0.102	1.000	1.000	1.000	1.000	1.000	1.000
1000	0.25	0.144	0.029	1.000	1.000	1.000	1.000	1.000	1.000
= .30, = .10									
50	0.03	0.280	2.699	0.029	0.021	0.007	0.014	0.063	0.031
100	0.03	0.377	1.747	0.067	0.060	0.023	0.025	0.121	0.070
300	0.03	0.351	0.658	0.276	0.269	0.194	0.188	0.336	0.264
1000	0.03	0.220	0.193	0.744	0.740	0.745	0.720	0.749	0.728

Note. DA = data augmentation, Gibbs = Gibbs sampler (Mplus), MI = multiple imputation, ML = maximum likelihood, Sobel = Sobel z test, BCB = bias corrected bootstrap, NB = naive bootstrap.