

Multiple Imputation: A Review of Practical and Theoretical Findings

Jared S. Murray

Abstract. Multiple imputation is a straightforward method for handling missing data in a principled fashion. This paper presents an overview of multiple imputation, including important theoretical results and their practical implications for generating and using multiple imputations. A review of strategies for generating imputations follows, including recent developments in flexible joint modeling and sequential regression/chained equations/fully conditional specification approaches. Finally, we compare and contrast different methods for generating imputations on a range of criteria before identifying promising avenues for future research.

Key words and phrases: Missing data, proper imputation, congeniality, chained equations, fully conditional specification, sequential regression multivariate imputation.

1. INTRODUCTION

Multiple imputation (MI) (Rubin, 1987) is a simple but powerful method for dealing with missing data. MI as originally conceived proceeds in two stages: A data disseminator creates a small number of completed datasets by filling in the missing values with samples from an imputation model. Analysts compute their estimates in each completed dataset and combine them using simple rules to get pooled estimates and standard errors that incorporate the additional variability due to the missing data.

MI was originally developed for settings in which statistical agencies or other data disseminators provide multiply imputed databases to distinct end-users. There are a number of benefits to MI in this setting: The disseminator can support approximately valid inference for a wide range of potential analyses with a small set of imputations, and the burden of dealing with the missing data is on the imputer rather than the analyst. All analyses conducted on the publicly available files can be based on the same set of imputations, ensuring

that differences in results are not due to the handling of missing data.

With the introduction of easy-to-use software to generate imputations and combine estimates, it has become increasingly common for users to create their own imputations prior to analysis. The set of methods available to generate imputations has also grown substantially, from simple parametric models and resampling methods to iterative classification and regression tree-based algorithms and flexible Bayesian nonparametric models. There are several textbook treatments of multiple imputation (e.g., Rubin, 1987, Little and Rubin, 2002, Van Buuren, 2012, Carpenter and Kenward, 2013) but fewer recent reviews of the variety of methods available to create multiply imputed files.

This paper provides a review of MI, with a focus on methods for generating imputations and the theoretical results and empirical evidence available to guide the selection and critique of imputation procedures. We restrict attention to methods for imputing item missing data (imputing the subset of values that are missing for an incomplete observation) in settings with independent observations. Much of the discussion also applies to other data structures, and to problems other than item missing data where MI has proven useful (see Reiter and Raghunathan, 2007 for some examples of other uses for multiple imputation).

The paper proceeds as follows: Section 2 briefly reviews the mechanics of multiple imputation for a

Jared S. Murray is Assistant Professor of Statistics, Department of Information, Risk, and Operations Management, and Department of Statistics and Data Science, University of Texas at Austin, 12110 Speedway B6500, Austin, Texas 78712, USA (e-mail: jared.murray@mccombs.utexas.edu).

scalar estimand. Section 3 reviews the conditions under which the usual MI rules give valid inference. Section 4 summarizes the practical implications of the theoretical results, particularly for choosing a method for generating imputations. Section 5 reviews methods for imputing a single variable subject to missingness. Section 6 reviews methods for imputing several variables. Section 7 discusses some of the considerations for choosing an imputation model. Section 8 concludes with discussion and directions for future work.

2. MULTIPLE IMPUTATION: HOW DOES IT WORK?

Let $Y_i = (Y_{i1}, Y_{i2}, \dots, Y_{ip})$ denote a p -dimensional vector of values corresponding to the i th unit and $R_i = (R_{i1}, R_{i2}, \dots, R_{ip})$ be a vector of indicator variables representing the response pattern, where $R_{ij} = 1$ if Y_{ij} is observed and is zero otherwise. We will use lowercase letters to distinguish fixed values from random variables, and denote the realized values in a particular dataset with a tilde (e.g., R_i is a random vector, r_i is a particular value that might be taken by R_i and \tilde{r}_i is the observed response pattern for unit i observed in a particular dataset).

Let $R = \{R_i : 1 \leq i \leq n\}$ with r and \tilde{r} defined similarly. The observed and missing values from a dataset of size n with response pattern R are denoted $Y_{\text{obs}}(R) = \{Y_{ij} : r_{ij} = 1, 1 \leq j \leq p, 1 \leq i \leq n\}$ and $Y_{\text{mis}}(R) = \{Y_{ij} : r_{ij} = 0, 1 \leq j \leq p, 1 \leq i \leq n\}$, respectively. Where the explicit dependence on the response pattern is a distraction we will drop the functional notation and simply refer to Y_{mis} and Y_{obs} .

We assume throughout that the missing data are *missing at random (MAR)* (Rubin, 1987), that is,

$$(2.1) \quad \Pr(R = \tilde{r} \mid Y_{\text{obs}}(\tilde{r}) = \tilde{y}_{\text{obs}}, Y_{\text{mis}}(\tilde{r}) = y_{\text{mis}}, \phi)$$

takes the same value for all y_{mis} and ϕ , where ϕ parameterizes our model of the response mechanism [the distribution of $(R \mid Y)$]. Under MAR, we do not need to explicitly model the response process to impute the missing data (Rubin, 1987, Result 2.3). MI may be used for missing data that are not MAR provided we explicitly model the response mechanism or make other identifying assumptions (see Rubin, 2003a for related discussion and examples of MI for non-MAR missing data).

2.1 Multiple Imputation for a Scalar Estimand

Let Q be an estimand of interest, which may be a function of complete data in a finite population or a model parameter. Let $\hat{Q}(Y)$ be an estimator of Q with

sampling variance U estimated by $\hat{U}(Y)$; where there is no ambiguity we refer to these as \hat{Q} and \hat{U} . In order to fix ideas, we focus on scalar Q . Inference for vector Q is similar in spirit; see Rubin (1987), Chapter 3, also Schafer (1997), Chapter 4, Section 3, or the review in Reiter and Raghunathan (2007), Section 2.1.

Assume $Y_{\text{mis}}^{(1)}, Y_{\text{mis}}^{(2)}, \dots, Y_{\text{mis}}^{(M)}$ are M imputations for Y_{mis} . Define $\hat{Q}^{(m)} = \hat{Q}(Y_{\text{obs}}, Y_{\text{mis}}^{(m)})$, the estimator computed using the m th completed dataset (with $\hat{U}^{(m)}$ defined similarly), and

$$(2.2) \quad \begin{aligned} \bar{Q}_M &= \sum_{m=1}^M \frac{\hat{Q}^{(m)}}{M}, & \bar{U}_M &= \sum_{m=1}^M \frac{\hat{U}^{(m)}}{M}, \\ B_M &= \sum_{m=1}^M \frac{(\hat{Q}^{(m)} - \bar{Q}_M)^2}{M-1}. \end{aligned}$$

These statistics form the basis for inference under MI: \bar{Q}_M averages the estimate computed in each imputed dataset to obtain an estimate of Q . The variance estimator of \bar{Q}_M has an ANOVA style decomposition:

$$(2.3) \quad T_M = \bar{U}_M + \left(1 + \frac{1}{M}\right) B_M,$$

where \bar{U}_M is an estimate of the variance of \hat{Q} if we had the complete data (“within-imputation” variance), and B_M estimates the excess variance due to the missing values (“between-imputation” variance). The factor $(1 + 1/M)$ is a bias adjustment for small M , as explained in Rubin (1987), Chapter 3.3.

MI was originally derived under Bayesian considerations. The Bayesian derivation of MI begins with the identities

$$(2.4) \quad \begin{aligned} P(Q \mid Y_{\text{obs}}) &= \int P(Q \mid Y_{\text{mis}}, Y_{\text{obs}}) P(Y_{\text{mis}} \mid Y_{\text{obs}}) dY_{\text{mis}}, \end{aligned}$$

$$(2.5) \quad \begin{aligned} E(Q \mid Y_{\text{obs}}) &= E(E(Q \mid Y_{\text{mis}}, Y_{\text{obs}}) \mid Y_{\text{obs}}), \end{aligned}$$

$$(2.6) \quad \begin{aligned} \text{Var}(Q \mid Y_{\text{obs}}) &= E(\text{Var}(Q \mid Y_{\text{mis}}, Y_{\text{obs}}) \mid Y_{\text{obs}}) \\ &\quad + \text{Var}(E(Q \mid Y_{\text{mis}}, Y_{\text{obs}}) \mid Y_{\text{obs}}). \end{aligned}$$

When imputations are generated from $P(Y_{\text{mis}} \mid Y_{\text{obs}})$, the MI statistics are Monte Carlo estimates of the relevant quantities:

$$(2.7) \quad \begin{aligned} \bar{Q}_M &\approx E(E(Q \mid Y_{\text{mis}}, Y_{\text{obs}}) \mid Y_{\text{obs}}) \\ &= E(Q \mid Y_{\text{obs}}), \end{aligned}$$

$$(2.8) \quad \bar{U}_M \approx E(\text{Var}(Q | Y_{\text{mis}}, Y_{\text{obs}}) | Y_{\text{obs}}),$$

$$(2.9) \quad (1 + 1/M)B_M \\ \approx \text{Var}(E(Q | Y_{\text{mis}}, Y_{\text{obs}}) | Y_{\text{obs}}),$$

$$(2.10) \quad T_M \approx \text{Var}(E(Q | Y_{\text{obs}})).$$

Rubin (1987) proposed constructing confidence intervals for Q based on an asymptotic normal approximation to the posterior distribution (2.4): Taking M to infinity, $(\bar{Q}_\infty - Q) \sim N(0, T_\infty)$ approximately in large samples. In large samples with finite M interval estimation for Q proceeds using a reference t -distribution for \bar{Q}_M : $(\bar{Q}_M - Q) \sim t_{\nu_M}(0, T_M)$. Rubin (1987) computed an approximate value for ν_M using a moment matching argument, obtaining $\nu_M = (M - 1)(1 + 1/r_M)^2$ where $r_M = (1 + 1/M)B_M/\bar{U}_M$ is a measure of the relative increase in variance due to non-response. Barnard and Rubin (1999) proposed an alternative degrees of freedom estimate with better behavior in moderate samples, suggesting it for general use. See Reiter and Raghunathan (2007) for a review of combining rules for more general estimands.

3. MULTIPLE IMPUTATION: WHEN DOES IT WORK?

In this section, we give a high-level review of some of the justifications for using MI and the estimators given above. Special consideration is given to results that can inform the selection of an imputation model.

3.1 Bayesian (In)validity Under MI

Since the MI estimators were derived under Bayesian arguments we might hope that MI yields valid Bayesian inference. In general, it does not. Suppose the analyst has specified a Bayesian model as $P_A(Y, Q) = P_A(Y | Q)P_A(Q)$. The analyst's inference is based on the posterior distribution

$$(3.1) \quad P_A(Q | Y_{\text{obs}}) \\ = \int P_A(Q | Y_{\text{mis}}, Y_{\text{obs}})P_A(Y_{\text{mis}} | Y_{\text{obs}})dY_{\text{mis}}.$$

Now suppose the imputer has generated imputations according to $Y_{\text{mis}}^{(m)} \sim P_I(Y_{\text{mis}} | Y_{\text{obs}})$. On computing $\hat{Q}(Y_{\text{obs}}, Y_{\text{mis}}^{(m)})$, the analyst has a draw from the hybrid model

$$(3.2) \quad P_H(Q | Y_{\text{obs}}) \\ = \int P_A(Q | Y_{\text{mis}}, Y_{\text{obs}})P_I(Y_{\text{mis}} | Y_{\text{obs}})dY_{\text{mis}}.$$

If $P_A(Y_{\text{mis}} | Y_{\text{obs}}) = P_I(Y_{\text{mis}} | Y_{\text{obs}})$, then MI delivers the analyst's posterior inference in the sense that $\hat{Q}^{(m)}$

is a draw from (3.1). If the posterior distribution for Q is approximately normal and M is not too small, the MI statistics will give a reasonable approximation to the posterior.

However, in practice the imputer and the analyst will likely have different models for $(Y_{\text{mis}} | Y_{\text{obs}})$. Even if one analyst should happen to share the same model as the imputer, the next analyst may have a different set of beliefs encoded in their model, resulting in $P_{A'}(Y_{\text{mis}} | Y_{\text{obs}}) \neq P_A(Y_{\text{mis}} | Y_{\text{obs}})$. In this case, the imputer cannot deliver valid Bayesian inference to both analysts with a single set of imputations. Since Bayesian validity is generally unattainable (and good repeated sampling behavior is desirable in its own right), MI is usually evaluated based on its frequentist properties. The remaining subsections explore conditions under which MI yields valid frequentist inference.

3.2 Frequentist Validity: Conditions on Complete Data Inference

We will follow Rubin (1996) and assume that the complete data inference is at least *confidence valid*, meaning that a nominal $100(1 - \alpha)\%$ confidence interval has actual coverage at least $100(1 - \alpha)\%$. (The stronger condition of *randomization validity* requires that the nominal and actual coverage rates agree.) We also assume that the sampling distribution of \hat{Q} is normal, so that valid confidence intervals can be obtained from \hat{Q} and \hat{U} . In this case, confidence validity requires that

$$(3.3) \quad E(\hat{Q}) = Q,$$

$$(3.4) \quad E(\hat{U}) \geq \text{Var}(\hat{Q}),$$

where the expectation and variance are over repeated sampling. Randomization validity obtains when $E(\hat{U}) = \text{Var}(\hat{Q})$. We depart slightly from Rubin (1996, 1987) in omitting any conditioning on fixed values in a finite population.

In practice normality and (3.3)–(3.4) may only hold asymptotically, or when particular modeling assumptions are correct. Whether this is plausible for a particular analysis will depend on the nature of \hat{Q} . For our purposes, we will assume that any necessary conditions for confidence validity with completely observed data are satisfied, since our primary consideration is the impact of missingness and imputation. Of course, if the complete data inference is not valid it would be unreasonable to expect MI or any other missing data procedure to remedy the issue.

3.3 Proper Imputation for Valid Inference

Chapter 4, Section 4.2 in Rubin (1987) outlines conditions under which MI inferences are randomization or confidence valid when $M = \infty$. Imputations satisfying these conditions for a particular estimand Q and posited response mechanism are known as *proper* imputations. Proper imputation coupled with valid complete data inference yields valid MI inference (Rubin, 1987, Result 4.1). It is important to remember that imputations are only proper with respect to a particular estimand Q and a posited response mechanism.

We focus on three essential conditions necessary for an imputation procedure to be proper for an estimand Q . (The other conditions are somewhat technical and generally not the source of improper imputations and invalid inference in practice.)

3.3.1 *Three essential conditions for proper imputation.* Rubin (1996) distilled the formal definition of proper imputation given in Rubin (1987), Section 4.2, into three conditions that generally ensure imputations are proper. They concern the behavior of the MI statistics under repeated realizations of the response mechanism, holding the sample values Y fixed [i.e., under repeated sampling from $P(R | Y)$]. The first two conditions require that \bar{Q}_∞ and \bar{U}_∞ be approximately unbiased for \hat{Q} and \hat{U} :

$$(3.5) \quad E(\bar{Q}_\infty | Y) \approx \hat{Q}(Y),$$

$$(3.6) \quad E(\bar{U}_\infty | Y) \approx \hat{U}(Y),$$

where the expectations are with respect to $P(R | Y)$.

Naturally, (3.5)–(3.6) will hold if $P(Y_{\text{mis}} | Y_{\text{obs}})$ is correctly specified by the imputer. However, imputations made under misspecified models can still satisfy (3.5)–(3.6) so long as they broadly capture the features of the predictive distribution that are relevant for computing Q and U and the proportion of missing data is not extreme. To see this more clearly, we can write

$$(3.7) \quad \begin{aligned} E(\bar{Q}_\infty | Y) &= \sum_{m=1}^{\infty} E(\hat{Q}(Y_{\text{obs}}(R), Y_{\text{mis}}^{(m)}(R)) | Y). \end{aligned}$$

With no missing data, the expectations inside the sum are all $\hat{Q}(Y)$. With modest amounts of missing data, the imputed values need to be sufficiently poor to overwhelm the influence of the observed data in computing Q . (What constitutes “sufficiently poor” naturally depends on Q .) Similar logic applies to \bar{U}_∞ .

The third condition for proper imputation is more subtle. It requires that the between-imputation variability B_∞ be approximately unbiased for the variance of \bar{Q}_∞ :

$$(3.8) \quad E(B_\infty | Y) \approx \text{Var}(\bar{Q}_\infty | Y).$$

Satisfying this condition generally requires that we account for uncertainty *in the imputation model itself* (or equivalently uncertainty in the parameters indexing a model class), since the observed data used to estimate the model, $Y_{\text{obs}}(R)$, varies over samples from the response mechanism. [Recall that the variance in (3.8) is with respect to $P(R | Y)$.]

Many seemingly reasonable stochastic imputation procedures fail to be proper because they do not satisfy (3.8); these include imputing from a model by plugging in the MLE or drawing imputations from the empirical distribution of observed cases (Rubin, 1987, Chapter 4). Accounting for uncertainty in the imputation model can be achieved (or approximated) in a variety of ways, such as sampling the parameters indexing a particular model class from their posterior under a Bayesian model or through small adjustments to the bootstrap (as described in Section 5.2). See Section 4.1 for further discussion.

3.4 Congeniality and Confidence Validity

It is well known that the MI estimate T_∞ can be inconsistent for certain choices of Q (Wang and Robins, 1998, Robins and Wang, 2000, Kim, 2002, Nielsen, 2003, Kim et al., 2006). The bias is typically positive and tends to have limited influence on coverage rates for common estimands when the amount of missingness is not extreme (Rubin, 2003a). Rubin (1996) reviewed early examples of inconsistency and gave sufficient conditions for MI inference to be confidence proper [i.e., for T_∞ to conservatively estimate $\text{Var}(\bar{Q}_\infty)$]; they are similar to the conditions in Section 3.3.1, averaged over repeated sampling of Y in addition to the response mechanism.

Meng (1994) introduced the concept of *congeniality* for understanding the inconsistency of the MI variance estimate. Roughly, an analysis procedure is congenial to an imputation model $P_I(Y_{\text{mis}} | Y_{\text{obs}})$ if we can take the complete data analysis and embed it into a Bayesian model $P_A(Y | Q)P_A(Q)$ such that:

1. Its posterior $P_A(Q | Y)$ recapitulates the desired analysis in the sense that

$$(3.9) \quad E_A(Q | Y) = \hat{Q}(Y), \quad \text{Var}_A(Q | Y) = \hat{U}(Y).$$

2. It matches the imputation model, that is,

$$(3.10) \quad P_A(Y_{\text{mis}} | Y_{\text{obs}}) = P_I(Y_{\text{mis}} | Y_{\text{obs}}).$$

Under congeniality, MI delivers samples from $P_A(Q | Y_{\text{obs}})$ (Section 3.1), which we have constructed to yield confidence valid inference. Unless the analyst is the imputer, congeniality is less a condition we should try to satisfy than one we should try to fail gracefully—uncongeniality is generally “the rule not the exception” (Xie and Meng, 2017), for the same reasons discussed in Section 3.1.

Xie and Meng (2017) revisited the behavior of MI inferences under uncongeniality and provided a host of new results. At a high level, their findings affirm and generalize common rules of thumb originating with Meng (1994): Even if the “true” model is nested within the imputer’s and the analyst’s models (e.g., if the imputation model includes both relevant and irrelevant covariates in an otherwise correctly specified regression model for the missing data), standard MI inference may be invalid. However, if the analyst’s procedure is *self-efficient* [meaning essentially that their estimator cannot be improved by ignoring relevant data (Meng, 1994, Meng and Romero, 2003)], then:

1. When the imputer’s model is more saturated than the analyst’s, the usual MI inference is confidence valid and generally robust.

2. When the imputer’s model is less saturated than the analyst’s, confidence validity is not guaranteed.

It is generally safer to conduct an uncongenial analyses under (1) than under (2), since conservative inferences will obtain. Xie and Meng (2017) also provide remarkably simple and broadly applicable (if somewhat exacting) alternative variance estimates that are valid under uncongeniality: Use $T_M^* = 2T_M$ for a vector Q , or sum and square the standard errors for a univariate Q : $T_M^* = (\sqrt{U_M} + \sqrt{B_M})^2 + (1/M)B_M$.

Like most strong theoretical results, Xie and Meng’s (2017) results depend on a number of assumptions. One of these assumptions is that the true model (“God’s model”) is nested within the imputation model class. In his discussion of the paper, Reiter (2017) notes that “[I]n my experience, very low coverage rates in MI confidence intervals arise more often from the imputation procedure generating bias in $[\bar{Q}_\infty]$ than from bias in the MI variance estimator,” often due to rote application of default imputation procedures. This has been in part a shared experience (Murray and Reiter, 2016), motivating the focus of this review on the specification of imputation models.

4. PRACTICAL IMPLICATIONS OF THEORETICAL RESULTS FOR IMPUTATION MODELING

The theoretical results summarized above suggest a number of practical considerations for generating imputations. These are reviewed below; for more detailed discussion and examples, see Rubin (1987, 1996), Little (1988), Van Buuren (2012). Throughout this section and the rest of the paper, we will continue to refer to procedures that generate imputations as “imputation models,” regardless of whether they are completely specified probability models.

4.1 Imputations Should Reflect Uncertainty About Missing Values and About the Imputation Model

The goal in multiple imputation is to account for uncertainty due to the missing values in subsequent inference. This is a different objective than estimating or predicting the missing values, which could generally be achieved via simpler means. The situation in MI is similar to the more familiar task of constructing valid predictive intervals with a regression model, where we need to account for uncertainty in the unobserved response as well as uncertainty in the regression fit.

Suppose we have a single variable subject to missingness, to be imputed using a regression model. If we were only concerned with reconstructing the missing values, we would just impute the fitted values. This would clearly lead to invalid MI inferences. Instead, MI propagates the intrinsic uncertainty about the missing values via some stochastic mechanism, for example, by adding a randomly generated residual to the regression prediction. However, to achieve at least approximately proper imputations we also need to account for uncertainty about the imputation model itself, that is, uncertainty in the fitted values of the regression model. Methods that do not appropriately reflect both sources of uncertainty tend to violate (3.8) and underestimate the between-imputation variance, yielding standard errors that are too small and anticonservative inferences (Rubin, 1987, 1996).

Bayesian imputation procedures provide a natural mechanism to account for model uncertainty. Imputations are generated from

$$(4.1) \quad \begin{aligned} P(Y_{\text{mis}} | Y_{\text{obs}}) \\ = \int P(Y_{\text{mis}} | \theta, Y_{\text{obs}}) P(\theta | Y_{\text{obs}}) d\theta, \end{aligned}$$

where θ is a parameter indexing a model for Y (or a model for Y_{mis} given Y_{obs}). To see how model uncertainty propagates, observe that imputations can be

sampled compositionally: For $1 \leq m \leq M$, first draw a value $\theta^{(m)} \sim P(\theta | Y_{\text{obs}})$ and then sample $Y_{\text{mis}}^{(m)} \sim P(Y_{\text{mis}} | \theta^{(m)}, Y_{\text{obs}})$. Model uncertainty is represented by $P(\theta | Y_{\text{obs}})$, and the intrinsic uncertainty about the missing values is represented by $P(Y_{\text{mis}} | \theta, Y_{\text{obs}})$. Approximations to full Bayesian inference have also proven useful: Rubin and Schenker's (1986) approximate Bayesian bootstrap for proper hot deck imputation is one early example (Section 5.2). Chapter 10 of Little and Rubin (2002) reviews several others.

Of course, Bayesian modeling is not magic—if θ indexes a class of misspecified models then we should expect our imputations and inferences to suffer, at least for estimands that are sensitive to this misspecification. For example, when Y_{mis} contains variables with significant skew a multivariate normal imputation model would likely yield approximately valid inference for marginal means but invalid inference for some marginal quantiles, since (3.5) can be violated when Q is an extreme quantile.

From a coverage perspective, model misspecification becomes increasingly consequential in large samples where the complete data standard errors are small and $P(\theta | Y_{\text{obs}})$ will tend to concentrate on the parameters of the “best” misspecified model. Even small biases due to misspecification in the imputation model can become large relative to the pooled standard errors. Enlarging the imputation model class $P(Y | \theta)$ via non- and semiparametric Bayesian modeling can guard against misspecification and also mitigate the artificial certainty implied by fixing a regular parametric model and only considering uncertainty in its parameters. Section 6.1.1 explores recent promising developments in this area.

4.2 Imputation Models Should Generally Include as Many Variables as Possible

There are multiple reasons for entertaining the largest possible imputation model: The missing at random assumption tends to be more tenable as more completely-observed variables are added to the imputation model. In addition, if variables predictive of the missing values are left out of the imputation model but used to compute Q or U , then the imputations will be improper—the imputed values will be incorrectly independent of the omitted variables, leading to bias over repeated imputations [violations of (3.5) or (3.6)] (Rubin, 1996). In this case, the analysis and imputation models are uncongenial in the “wrong” way—the imputer's model is less-saturated than the analysis

model. In sum, the cost of excluding a relevant variable (invalid inference) is often greater than the cost of including an irrelevant variable (roughly, additional variance). This is particularly relevant when the analyst and imputer are not the same, and the imputations must support many unspecified analyses. Even when the imputer and the analyst are the same it would be useful to generate one set of imputations that can support the usual process of iterative model building and refinement, rather than generating a new set of imputations for each analysis model that is considered. See Collins, Schafer and Kam (2001) and Schafer (2003) for further discussion of the tradeoffs involved.

These points are particularly relevant for design variables in complex surveys. Design-based estimators will typically use stratum and cluster information to compute U . Reiter, Raghunathan and Kinney (2006) show empirically that failing to account for an informative sampling design can lead to invalid inference. They suggest including indicator variables for strata and cluster membership in the imputation model, or including stratum fixed effects and cluster random effects in imputation models. It may be useful to include estimated response propensities or final adjusted survey weights (sampling weights with, e.g., calibration and post-stratification adjustments) as well, especially if complete design information is not available to the imputer (Rubin, 1996).

4.3 Imputation Models Should Be as Flexible as Possible

Finally, imputation models should try to “track the data” (Rubin, 1996) by modeling relevant features of the joint distribution of the missing values. Loosely, a feature of the joint distribution is relevant if it is a possible target of inference itself, or more generally if it yields a more accurate predictive distribution for the missing data. Interactions, nonlinearities and non-standard distributional forms are all potentially relevant features.

As Meng (1994) succinctly put it, “Sensible imputation models should not only use all available information to increase predictive power, but should also be as general and objective as practical in order to accommodate a potentially large number of different data analyses.” We would add that where possible, imputation models should have some capacity to *adapt* to unanticipated features of the data (such as interactions, nonlinearities and complex distributions), especially when the imputer has limited time and resources to spend on iteratively improving the imputation model.

5. GENERATING IMPUTATIONS FOR A SINGLE VARIABLE

We begin by cataloging some of the more common approaches to generating imputations for a single variable subject to missingness, conditional on other fully observed variables. In the next section, we consider how these can be extended to generate imputations for several variables.

5.1 Regression Modeling

Imputation by sampling from univariate regression models is conceptually straightforward. Generalized linear models and extensions to deal with complications such as zero-inflation and truncation are popular options; these are not reviewed in depth here but see, for example, Van Buuren and Oudshoorn (1999), Raghunathan et al. (2001), Su et al. (2011) or Van Buuren (2012) (Chapter 3). These methods are quite common in practice, but since most readers will be familiar and they are well-reviewed elsewhere we will not enumerate them here.

To generate proper imputations, some method should be used to account for parameter uncertainty—simple strategies like sampling from the regression model with parameters fixed at the observed data MLE are generally improper. Posterior sampling under a non- or weakly informative prior tends to be proper when the model fits well. Prior distributions can also ease problems like separation in logistic regression and apply helpful regularization in conditional models with many variables in the conditioning set (Su et al., 2011).

5.2 Hot Deck/Nearest Neighbor Methods

The hot deck and other nearest-neighbor methods (Chen and Shao, 2000, Andridge and Little, 2010) begin by defining a distance metric between cases in terms of the observed covariates. Imputations for a missing value are borrowed from a nearby completely observed case (the “donor”). These methods tend to be simpler to implement than fully specified regression models and often make fewer assumptions. However, these methods are far from assumption-free—the choice of distance metric, the definition of the donor pool, and how to sample from the donor pool all influence the quality of imputations.

The hot deck (Andridge and Little, 2010) defines distance via cross-classifications of fully observed variables which determine adjustment cells. Missing values are imputed by sampling with replacement from the pool of donors within the same cell. This strategy ensures that all imputations are plausible values,

which is an appealing feature relative to regression imputation. Complications arise when there are many fully observed variables to incorporate into the cross-classification or when the sample size is low, leading to many small or empty adjustment cells.

MI with the hot deck is also known to be improper for simple estimands like a population mean (Rubin and Schenker, 1986). The hot deck effectively assumes that the distribution of missing values within an adjustment cell is *exactly* the empirical distribution of the observed values within that cell, which leads to B having downward bias (due to ignoring uncertainty in the implicit imputation model). Rubin and Schenker (1986) propose a simple modification that makes the hot deck proper, based on an approximation to the Bayesian bootstrap (Rubin, 1981). Instead of sampling the n_m missing values from the empirical distribution of the n_o observed values within an adjustment cell, the approximate Bayesian bootstrap (ABB) first samples a set of n_o values with replacement from the observed data and then samples n_m imputed values with replacement from this set. This simple adjustment yields proper imputations for the population mean of the adjustment cell (Rubin and Schenker, 1986). (See also Kim, 2002 for a more accurate variance estimate in small samples.)

Predictive mean matching (PMM) (Little, 1988) instead measures the distance between cases by the distance between their predicted means for the variable subject to missingness (traditionally estimated using a linear regression, although in principle any method could be used to make the prediction). PMM generalizes the hot deck, which is a special case of PMM using saturated models with categorical predictors. By avoiding the discretization and making some assumptions about the relationships between the predictors and the response (such as linearity) PMM can handle more variables than the hot deck, but may be sensitive to the predictive model specification.

To define the donor pool, Heitjan and Little (1991) proposed sampling from a window of k nearby potential donors in PMM in the hope of making the method approximately proper. The donor’s value may be imputed, or its residual can be added to the predicted mean of the missing value to generate an imputation. Schenker and Taylor (1996) found these two approaches to perform similarly in simulations; the former will always impute a previously realized value, which may be desirable. See Vink et al. (2014) for an approach to semi-continuous variables. Morris, White and Royston (2014) compared newer developments

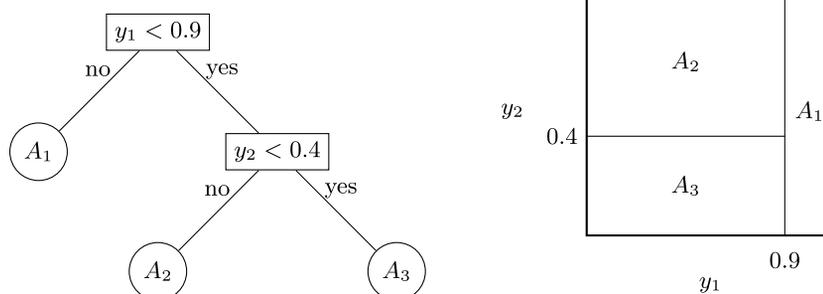


FIG. 1. (Left) An example CART tree, with internal nodes labeled by their splitting rules and terminal nodes given labels A_n . (Right) The corresponding partition of (Y_1, Y_2) .

and current implementations of these techniques, cautioning in particular against the imputation of a single nearest neighbor (which appears to be common in software implementations of PMM) as it is improper.

PMM and the hot deck can be made more adaptive using recursive partitioning. Reiter (2005) and Burgette and Reiter (2010) proposed imputation via classification and regression trees (CART, Breiman et al., 1984). A tree is grown using fully observed data to predict the variable subject to missingness. Then each incomplete case is assigned to its corresponding leaf, and an imputation is sampled from donors within in the same leaf. The imputer can control the size of the donor pool by growing the tree down to a specified minimum leaf size. This is a special case of PMM using CART to generate predictions; we could also think of it as an adaptive hot deck that leverages the most predictive variables and balances the size of the adjustment cells. Figure 1 shows an example tree grown on two variables (Y_1, Y_2) to impute a third (Y_3) , along with the corresponding partition which forms the adjustment cells.

Reiter (2005) and Burgette and Reiter (2010) drew ABB samples from within the leaves in an effort to generate proper imputations. Van Buuren (2012) (Algorithm 3.6) suggested also accounting for uncertainty in the tree itself by growing it on a different bootstrap sample for each imputed dataset. Doove, Van Buuren and Dusseldorp (2014) proposed imputation by growing a random forest (an ensemble of trees) (Breiman, 2001) of size k by bootstrapping the complete cases and (optionally) subsampling the variables, as in traditional applications of random forests. An imputed value is generated by sampling from the k trees and then following the procedure to generate a CART imputation. Shah et al. (2014) proposed fitting a random

forest, estimating its predictive error variance and generating imputations as the random forest prediction plus a normally distributed residual.

Limited results exist comparing these different recursive partitioning methods, and there is similarly limited guidance as to how they should be tuned. But they can be fast and effective imputation engines, particularly for large sets of categorical variables that take a relatively limited set of levels (see, e.g., Akande, Li and Reiter, 2017).

6. GENERATING IMPUTATIONS FOR MULTIPLE VARIABLES

There are two basic strategies for imputing multivariate missing data: Jointly modeling the variables subject to missingness, or specifying a collection of univariate conditional imputation models that condition on all the other variables [this approach goes under various names including sequential regression multivariate imputation (Raghunathan et al., 2001) and multiple imputation by chained equations (Van Buuren and Oudshoorn, 1999), but we will use “fully conditional specification” (FCS) as in Van Buuren et al., 2006]. Joint models can be further classified into “simultaneous” approaches that define a multivariate distribution $f(Y)$ directly or “sequential” approaches that build up a multivariate distribution using a ladder of conditional distributions, where the model for each variable conditions only on those earlier in the sequence. The Appendix has pointers to software implementations of many methods described in this section.

To describe the different approaches, we need some new notation: Let $Y_{j,\text{obs}}$ and $Y_{j,\text{mis}}$ denote the set of observed and missing values for the j th variable. Let Y_{imp} denote an imputed dataset, and $Y_{j,\text{imp}}$ denote a set of imputations for $Y_{j,\text{mis}}$. We will use the subscript $(-j)$ to denote the same quantities for all but the j th variable.

6.1 Joint Specification: Simultaneous Approaches

Early simultaneous joint modeling approaches were based on the multivariate normal (MVN) or t distribution; these are reviewed in Schafer (1997) and Little and Rubin (2002). For high dimensional continuous observations, low-rank structure can be imposed on the covariance matrix (Audigier, Husson and Josse, 2016). Various authors have proposed imputing categorical data under a misspecified MVN model, either leaving the continuous imputations for discrete variables as-is or rounding them based on some thresholds (Horton, Lipsitz and Parzen, 2003, Bernaards, Belin and Schafer, 2007). This is naturally more complicated when the discrete variables are not ordinal, particularly if they take many levels. Additionally, end users may not trust imputations from a data disseminator if the imputed data appear invalid. Therefore, it is often preferable to use models that are appropriate for the types of variables at hand.

For small numbers of strictly discrete variables, a simple multinomial model may be feasible. However, with a large number of discrete variables it is impossible to fit saturated multinomial models and further restrictions are necessary. Options include log-linear models (Schafer, 1997), latent class models (Vermunt et al., 2008, Gebregziabher and DeSantis, 2010, Vidotto, Vermunt and Kaptein, 2015) or multiple correspondence analysis (Audigier, Husson and Josse, 2017) [which is closely related to a certain class of multivariate logit models (Fithian and Josse, 2017)].

Joint models for mixed continuous and categorical data are also available. For the remainder of Section 6.1, suppose we have collected the continuous variables into a vector Y and the discrete variables into another vector X . The general location model (GLOM) (Olkin and Tate, 1961, Little and Schluchter, 1985, Schafer, 1997) assumes that $(Y | X = x) \sim N(\mu_x, \Sigma_x)$ and $X \sim \pi$. [Liu and Rubin, 1998 generalized the $(Y | X)$ model to the larger class of elliptically symmetric distributions.] The number of parameters in this saturated model grows rapidly with the sample space of X , so imputers typically impose further constraints. Examples include common covariance structure ($\Sigma_x \equiv \Sigma$ for all x), removing higher-order effects from the conditional means by specifying $\mu_x = D(x)B$ for a matrix of regression coefficients B and design vector $D(x)$, and imposing log-linear constraints on π to rule out higher-order interactions in the marginal model for X .

6.1.1 *Mixtures and nonparametric Bayesian models.* Even without additional parameter constraints, most parametric joint models make restrictive assumptions. Mixture models provide a simple and expressive way to enrich a parametric model class. For example, latent class models for categorical data are mixtures of independence models (log-linear models with only main effects) which have proven useful in multiple imputation (e.g., Vermunt et al., 2008, Gebregziabher and DeSantis, 2010). Mixtures of multivariate normal distributions can model complex features of joint continuous distributions (Böhning et al., 2007, Elliott and Stettler, 2007).

Several Bayesian nonparametric models have recently been proposed for multiple imputation. Most of these are based on infinite mixture models or their truncated approximations (but see Paddock, 2002 for an early exception based on Polya trees, and also the sequential regression approach in Xu, Daniels and Winterstein, 2016). Relative to parametric Bayesian approaches these models are appealing for their ability to grow in complexity with increasing sample size. Under some circumstances, this can allow the model to capture unanticipated structure like interactions and non-linear relationships or nonstandard distributions, reflecting these in the imputed values.

Recall that we have separated the data into vectors of categorical variables X and continuous variables Y . For imputing multivariate categorical data, Si and Reiter (2013) adopt a truncated version of the Dirichlet process mixture of product multinomials (DP-MPMN) proposed by Dunson and Xing (2009). This is a latent class model with a large number of classes (say k^X) and a particular prior over the class distribution.

Suppose the j th categorical variable takes (possibly unordered) values indexed by $1, 2, \dots, d_j$ and let $H_i^X \in \{1, \dots, k^X\}$ be a latent mixture component index for observation i . Let $\Pr(X_{ij} = x_{ij} | H_i^X = s) = \psi_{sx_{ij}}^{(j)}$. The DP-MPMN model assumes that

$$(6.1) \quad \Pr(H_i^X = s) = \phi_s^X,$$

$$(6.2) \quad \Pr(X_i = x_i | H_i^X = s, \Psi) = \prod_{j=1}^p \psi_{sx_{ij}}^{(j)},$$

so that the elements of X are conditionally independent given the latent class membership. The prior on ϕ is a truncated version of the stick-breaking construction for the Dirichlet process (DP) (Sethuraman, 1994), introduced in Ishwaran and James (2001) to simplify Gibbs

sampling in DP mixture models:

$$(6.3) \quad \begin{aligned} \phi_s^{\mathcal{X}} &= \xi_s \prod_{l < s} (1 - \xi_l), \\ \{\xi_s\}_{s=1}^{k^{\mathcal{X}}} &\stackrel{iid}{\sim} \text{Beta}(1, \alpha), \quad \xi_{k^{\mathcal{X}}} \equiv 1. \end{aligned}$$

The model is completed with prior distributions on Ψ and α (see Si and Reiter, 2013, for a complete specification). Manrique-Vallier and Reiter (2014a, 2014b) extended this model to assign zero probability to impossible values of X , such as cells that are logically impossible (pregnant men or children collecting retirement benefits) or necessarily empty due to skip patterns. Manrique-Vallier and Reiter (2016) introduced a variant of this model for edit-imputation that simultaneously accounts for missing values and observed values that are logically impossible but present due to measurement error. Hu, Reiter and Wang (2017) extended this model to nested data structures (i.e., hierarchical structures like individuals nested within households) in the presence of structural zeros.

For imputing continuous data, Kim et al. (2014) suggested a truncated DP mixture of multivariate normal distributions. Let $H_i^{\mathcal{Y}}$ be the mixture component index for record i . This model assumes that

$$(6.4) \quad \Pr(H_i^{\mathcal{Y}} = r) = \phi_r^{\mathcal{Y}},$$

$$(6.5) \quad (Y_i | H_i^{\mathcal{Y}} = r, -) \sim N(\mu_r, \Sigma_r),$$

with a prior on $\phi_r^{\mathcal{Y}}$ defined via a stick-breaking process similar to (6.3). Kim et al. (2014) modified the model in (6.5) to constrain the support of Y to a set \mathcal{A} with bounds determined by a set of linear inequalities, so that $\Pr(Y \notin \mathcal{A}) = 0$ under the prior. Kim et al. (2015) extended this approach to simultaneous edit-imputation, generating imputed values for observations outside of \mathcal{A} via a measurement error model.

Murray and Reiter (2016) built a hierarchical mixture model for mixed continuous and categorical observations by combining the models in (6.1)–(6.2) and (6.4)–(6.5), with two important adjustments. First, (6.5) is modified to include a regression on X with component-specific coefficients:

$$(6.6) \quad (Y_i | X_i = x_i, H_i^{\mathcal{Y}} = r, -) \sim N(D(x_i)B_r, \Sigma_r).$$

By default, the design matrix $D(x_i)$ encodes main effects. Allowing the component means to depend on X greatly reduces the number of mixture components necessary to capture $X - Y$ relationships. Second, the mixture component indices in each model are given a

hierarchical prior introduced by Banerjee, Murray and Dunson (2013):

$$(6.7) \quad \Pr(H_i^{\mathcal{X}} = s, H_i^{\mathcal{Y}} = r | Z_i = z) = \phi_{zs}^{\mathcal{X}} \phi_{zr}^{\mathcal{Y}},$$

$$(6.8) \quad \Pr(Z_i = z) = \lambda_z.$$

Here, λ_z is assigned a stick-breaking prior. Each pair $\phi_z^{\mathcal{X}} = (\phi_{z1}^{\mathcal{X}}, \dots, \phi_{zk^{\mathcal{X}}}^{\mathcal{X}})'$ and $\phi_z^{\mathcal{Y}} = (\phi_{z1}^{\mathcal{Y}}, \dots, \phi_{zk^{\mathcal{Y}}}^{\mathcal{Y}})'$ are probability vectors also assigned independent truncated stick breaking priors. This is a “mixture of mixtures” model; marginalizing over the latent variables the joint density is

$$(6.9) \quad \begin{aligned} f(X_i, Y_i) &= \sum_{z=1}^{k^{\mathcal{Z}}} \lambda_z \left(\sum_{r=1}^{k^{\mathcal{Y}}} \phi_{zr}^{\mathcal{Y}} N(Y_i; D(X_i)B_r, \Sigma_r) \right) \\ &\cdot \left(\sum_{s=1}^{k^{\mathcal{X}}} \phi_{zs}^{\mathcal{X}} \prod_{j=1}^p \psi_{sX_{ij}}^{(j)} \right). \end{aligned}$$

Each mixture component is itself composed of two mixture models, one for $(Y | X)$ and one for X . These lower-level mixtures share some parameters (B, Σ and Ψ), enforcing a degree of parsimony.

DeYoreo, Reiter and Hillygus (2016) used a similar hierarchical mixture model constructed based on different considerations, splitting the variables into sets based on their type (ordinal or nominal) and high or low rates of missing values. An expressive model class is specified for the variables with high rates of missing values, and a simpler model class is utilized for variables with low rates of missingness. Ordinal variables are explicitly modeled as such by thresholding mixtures similar to (6.6).

Further extensions, combinations and enhancements of these models are possible. Despite their complexity, all of these models have been shown to perform well for MI with real, complicated data and little or no tuning.

6.2 Fully Conditional Specification

FCS avoids explicit joint probability models by specifying a collection of univariate conditional imputation models instead (Van Buuren and Oudshoorn, 1999, Raghunathan et al., 2001). Each univariate model typically conditions on all the remaining variables. In FCS, the missing values are imputed by iteratively sampling from these conditional models:

1. Begin by filling in Y_{mis} with plausible values to generate an initial completed dataset, stored in Y_{imp} .

2. For $1 \leq j \leq p$, use a univariate imputation method to sample new imputed values for $Y_{j,\text{mis}}$ from a distribution $P(Y_{j,\text{mis}} | Y_{j,\text{obs}}, Y_{(-j),\text{imp}})$, and store them in $Y_{j,\text{imp}}$.

3. Iterate the previous step until apparent convergence and return the final value of Y_{imp} .

This process is repeated M times, saving the returned value as one of the M imputations. Any of the univariate imputation methods in the previous section could be used. This lends FCS some flexibility relative to the joint-simultaneous approaches described above.

But this flexibility comes at a cost: Even if each g_j is a completely specified probability model, taken together they often do not correspond to a proper joint distribution for Y (Arnold and Press, 1989, Arnold, Castillo and Sarabia, 2001). A set of full conditional distributions that do not correspond to any joint distribution is said to be *incompatible*. Simple adjustments like adding polynomial terms or interactions to univariate regression models can induce incompatibility (Liu et al., 2014).

While the algorithm above looks like a standard Gibbs sampler, if the conditional models are incompatible the behavior of the FCS imputation algorithm is unclear: The imputations from the FCS algorithm given above may converge to a unique limiting distribution, or fail to converge to any unique limiting distribution, or converge to different distributions depending on the initial values and/or order of the updates. Li, Yu and Rubin (2012) give examples of incompatible FCS models with fixed parameters whose imputations either diverge or converge to different stationary distributions depending on the order of their updates. This phenomenon seems to be rare in real data, and Zhu and Raghunathan (2015) note that estimating rather than fixing parameters ameliorates at least some of the problems in Li, Yu and Rubin's (2012) examples.

There are some limited convergence results available when the fully conditional specification comprises univariate Bayesian regression models. Liu et al. (2014) study an iterative FCS imputation procedure that uses a set of Bayesian regression models $g_j(Y_{ij} | Y_{(-j)}, \theta_j)$ with prior distributions $\pi_j(\theta_j)$. With a slight abuse of notation, define

$$(6.10) \quad \begin{aligned} &g_j(Y_{j,\text{obs}} | Y_{(-j),\text{imp}}, \theta_j) \\ &= \prod_{i=1}^n g_j(Y_{ij} | Y_{j,\text{imp}}, \theta_j)^{R_{ij}}, \end{aligned}$$

Algorithm 1 Iterative FCS sampler from Liu et al. (2014)

For $1 \leq j \leq p$,

1. Sample $\theta_j \sim \pi_j(\theta_j | Y_{j,\text{obs}}, Y_{(-j),\text{imp}}) \propto g_j(Y_{j,\text{obs}} | Y_{(-j),\text{imp}}, \theta_j) \pi_j(\theta_j)$
 2. Sample $Y_{j,\text{imp}} \sim g_j(Y_{j,\text{imp}} | Y_{j,\text{obs}}, Y_{(-j),\text{imp}}, \theta_j)$
-

$$(6.11) \quad \begin{aligned} &g_j(Y_{j,\text{imp}} | Y_{j,\text{obs}}, Y_{(-j),\text{imp}}, \theta_j) \\ &= \prod_{i=1}^n g_j(Y_{ij} | Y_{(-j),\text{imp}}, \theta_j)^{1-R_{ij}}. \end{aligned}$$

Algorithm 1 gives one iteration of an iterative FCS sampler under these models.

We can compare this approach to a proper MCMC algorithm under a joint model. Specifically, we consider a collapsed Gibbs sampler (Liu, 1994) that targets $P(Y_{\text{mis}} | Y_{\text{obs}}) = \int P(Y_{\text{mis}}, \theta | Y_{\text{obs}}) d\theta$ directly, by jointly sampling $(Y_{j,\text{mis}}, \theta | Y_{j,\text{obs}}, Y_{(-j),\text{imp}})$ at each step. It is impractical to use directly, but it is helpful to make comparisons with Algorithm 1.

Let the joint model be given by $f(Y_i | \theta)$, with full conditionals $f_j(Y_{ij} | Y_{(-j)}, \theta)$ and joint prior distribution $\pi(\theta)$ [where $\theta = (\theta_1, \theta_2, \dots, \theta_p)$]. Define $f_j(Y_{j,\text{obs}} | Y_{(-j),\text{imp}}, \theta)$ and $f_j(Y_{j,\text{imp}} | Y_{j,\text{obs}}, Y_{(-j),\text{imp}}, \theta)$ as in equations (6.10)–(6.11). Algorithm 2 gives one iteration of the collapsed Gibbs sampler.

Under some regularity conditions, the two algorithms are equivalent in finite samples if we can write $\pi(\theta) = \pi_j(\theta_j) \pi_{(-j)}(\theta_1, \theta_2, \dots, \theta_{j-1}, \theta_{j+1}, \dots, \theta_p)$ for any j and the set of g_j 's are compatible and correspond to the full conditionals of f (Hughes et al., 2014). This is sufficient to ensure that the conditional distributions in both steps of each algorithm agree.

If $\pi(\theta) \neq \pi_j(\theta_j) \pi_{(-j)}(\theta_1, \theta_2, \dots, \theta_{j-1}, \theta_{j+1}, \dots, \theta_p)$ for some j but the conditional models are compatible and correspond to the full conditionals of f , the two algorithms agree as $n \rightarrow \infty$ provided the FCS algorithm has a unique stationary distribution (Liu et al., 2014). Intuitively, in this case the data in $Y^{(-j)}$ influence θ_j indirectly through the other parameters, but

Algorithm 2 Collapsed Gibbs sampler for a joint model

For $1 \leq j \leq p$,

1. Sample $\theta \sim \pi(\theta | Y_{j,\text{obs}}, Y_{(-j),\text{imp}}) \propto f_j(Y_{j,\text{obs}} | Y_{(-j),\text{imp}}, \theta) \pi(\theta)$
 2. Sample $Y_{j,\text{imp}} \sim f(Y_{j,\text{imp}} | Y_{j,\text{obs}}, Y_{(-j),\text{imp}}, \theta)$
-

the FCS algorithm ignores this information. Asymptotically, the priors become irrelevant in regular parametric models, but in finite samples inference based on the FCS imputations may be inefficient in this regime (Seaman and Hughes, 2016).

Finally, Liu et al. (2014) show that if the FCS algorithm uses an inconsistent set of models but has a unique stationary distribution then MI estimates computed using imputations from Algorithm 1 are consistent provided that the following conditions hold:

1. The collection of conditional models are incompatible, but become compatible with a joint model f after constraining θ .
2. The model class defined by f contains the true distribution that generated the data.

These are rather restrictive; verifying a unique stationary distribution is challenging, as is checking condition 1 above. It also seems unlikely that condition 2 will hold exactly for the simple parametric models in common use. Zhu and Raghunathan (2015) provide some further convergence results for FCS algorithms where each observation is missing at most one value, but without assuming a unique stationary distribution for the FCS chain.

6.3 Joint Specifications: Sequential Approach

Sequential approaches to imputation modeling fix a permutation of $1, 2, \dots, p$ and build up a joint distribution from a series of univariate models. For example, if the variables are already in the desired order we would

have

$$(6.12) \quad f(Y) = f_1(Y_1) f_2(Y_2 | Y_1) f_3(Y_3 | Y_2, Y_1) \cdots \cdot f_p(Y_p | Y_{p-1}, \dots, Y_1).$$

Examples of this approach include (Lipsitz and Ibrahim, 1996, Ibrahim, Lipsitz and Chen, 1999, Ibrahim et al., 2005, Lee and Mitra, 2016, Xu, Daniels and Winterstein, 2016), among others.

Provided that each f_j is a proper univariate probability model, a sequential specification always defines a coherent joint model, unlike FCS approaches. However, different orderings will generally lead to different joint distributions and potentially different fits. Heuristics have been proposed for selecting the order, for example, ordering variables by their types (e.g., Ibrahim, Lipsitz and Chen, 1999) or percentage of missing values (e.g., Rubin and Schafer, 1990). The latter is particularly well motivated when the missing data are monotone (when there is an ordering such that $R_{ij} = 0 \Rightarrow R_{i,j'} = 0$ for $j' > j$). If the missing data are not exactly monotone, one can identify a permutation that is nearly monotone and use FCS or delete observed values to “monotonize” the missing data pattern, so that proper sequential techniques can be used for the majority of missing values (as in Rubin, 2003b and extended in Li et al., 2014).

Another consideration in joint-sequential modeling is that variables early in the sequence may have complex distributions because they are marginalized over many related covariates. For example, Figure 2 shows the joint distribution of householder earnings and age,

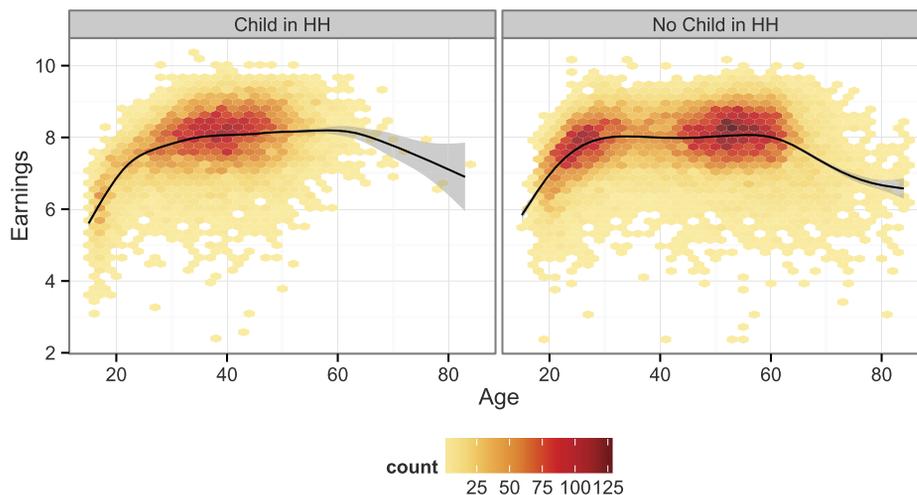


FIG. 2. Joint distribution of householder age and log total earnings, stratified on whether the household includes one of the householder’s own children, using the population Murray and Reiter (2016) constructed from complete cases in the first wave of the Survey of Income and Program Participation’s 2008 panel.

conditional on whether the householder has any children living in the same household (the data are from complete cases in wave one of the Survey of Income and Program Participation's 2008 panel). The distributions are quite complicated, and it would be difficult to capture them well with simple parametric regression models in any order.

7. CHOOSING AND ASSESSING AN IMPUTATION STRATEGY

7.1 Comparing FCS and Joint Approaches

FCS and joint approaches have competing strengths. FCS models are relatively simple to implement and widely available in software, especially compared to joint-sequential approaches. Joint-simultaneous models including the multivariate normal, log-linear models and the GLOM are also easy to set up and widely available, but inflexible in practice even relative to simple FCS procedures (e.g., Van Buuren, 2007, Stuart et al., 2009, He et al., 2010, Drechsler, 2010, Kropko et al., 2014).

More sophisticated joint models can be challenging to implement, although this is changing—many of the nonparametric Bayesian methods have publicly available implementations (the Appendix). However, even with a good implementation the nonparametric Bayesian models are generally more computationally expensive than simpler joint models (especially those based on low-rank methods, for example, Audigier, Husson and Josse, 2016, 2017) or FCS methods. Joint-sequential approaches currently take more effort to set up, but they inherit many of the positive features of FCS and joint-simultaneous approaches (univariate models that are readily assessed and modified but also consistent with joint models).

The convergence properties of FCS in general settings is still mostly an open question. The behavior of FCS algorithms under non- or quasi-Bayesian imputation procedures like PMM is entirely an open question. While the lack of a coherent joint distribution does undermine the theoretical justifications for MI inference detailed in Rubin (1987), experience with FCS in simulations and real applications does not seem to suggest that either lack of convergence or compatibility with a joint model are necessarily overriding concerns.

In fact, under the current theoretical results ensuring that the imputations generated by FCS converge to the imputations under a proper joint model requires using restrictive (implicit) joint models and there is strong empirical evidence that these joint models can be too

simple to perform well with realistic data (e.g., Murray and Reiter, 2016, Akande, Li and Reiter, 2017). Therefore, at this point it would probably be a mistake to choose the models in an FCS imputation routine to try to ensure convergence; it seems much more important to use flexible, adaptive imputation models wherever possible, whether using a joint or FCS imputation strategy.

Imputers who do choose to use FCS should use flexible univariate models wherever possible and take care to assess apparent convergence of the algorithm, for example, by computing traces of pooled estimates or other statistics and using standard MCMC diagnostics (Gelman et al., 2013, Chapter 11). It may also be helpful to examine the results of many independent runs of the algorithm with different initializations and to use random scans over the p variables to try to identify any convergence issues and mitigate possible order dependence.

7.2 Practical Considerations Derived from MI Theory

We can also compare methods on the practical considerations derived from theoretical results as summarized in Section 4.

7.2.1 Accounting for uncertainty. Most of the methods reviewed above include some mechanism for reflecting imputation model uncertainty. Bayesian or approximately Bayesian methods (including the approximate Bayesian bootstrap) do this naturally, whether part of a joint modeling or FCS imputation routine. Their behavior is not well understood in the FCS setting, however. Tree-based methods seem promising for some applications, but more work is required to find parameter settings and resampling strategies that make them reliably proper.

7.2.2 Include as many variables as possible. Joint-sequential models may be easier to fit than FCS with many covariates, since all but one univariate model will include fewer than p predictors. Simultaneous joint models somewhat lag behind sequential and FCS approaches here. This is particularly true with mixed data types and many fully observed covariates—most of these models are not easily adapted to condition on additional covariates, so fully observed variables must be included as additional variables in the joint model. Modeling fully observed variables instead of conditioning on them can waste “degrees of freedom” and lead to poorer model fit for the conditional distribution of the missing data. Carefully constructed models can help (DeYoreo, Reiter and Hillygus, 2016), but seem to only go so far.

7.2.3 Use flexible imputation models. Non- and semiparametric methods (Bayesian and otherwise, such as sequential tree-based methods) are flexible in their ability to capture certain unanticipated features of the data. Empirically, these methods can outperform existing default MI procedures in simulations, particularly when the simulations are not built around simple parametric models themselves. More of these realistic evaluations are needed, as discussed in Sections 7.3 and 8.

However, with flexible imputation models it can be challenging to manually adjust the imputation model to incorporate prior information or address model misfit. Incorporating meaningful prior information into non-parametric Bayesian imputation models is challenging but not impossible; see, for example, [Schifeling and Reiter \(2016\)](#) for a strategy to include prior information in DP-MPMN models. While iterative imputation model refinement and assessment is ideal, it is not always possible. Empirical evidence suggests that flexible imputation models are much better as defaults than simple parametric models or PMM using linear models.

7.3 Empirical Comparisons Between Methods

Empirical comparisons of several different imputation models on realistic datasets are relatively rare. Most papers introducing a new imputation model evaluate it using synthetic data generated from a researcher-specified multivariate probability model. The new imputation model is typically compared to a small number of competitors. These simulation studies can be informative, for example, both [Burgette and Reiter \(2010\)](#) and [Doove, Van Buuren and Dusseldorp \(2014\)](#) found evidence that imputations for continuous values generated via recursive partitioning can preserve interactions but underestimate main effects. However, models that are easy to simulate from and present in a paper will naturally be gross simplifications of the distribution of data in real populations.

Simulations based on repeated sampling from realistic populations can be more informative. In these studies, a population is compiled from existing data. Random samples are taken from these populations and values are “blanked out” via a known stochastic non-response mechanism. Each of the resulting incomplete datasets are multiply imputed and used to compute a range of estimates and confidence intervals, assessing the bias, coverage and efficiency of the MI estimates under the imputation model. Since the missing values

are known, these can all be compared against the frequentist operating characteristics of the complete data procedure without appeal to asymptotic theory or other approximations. While the results are specific to a particular population and a set of estimands, this framework is much closer to reality than fully synthetic examples.

There are several recent examples of this kind of evaluation: [Akande, Li and Reiter \(2017\)](#) compared FCS with CART, the DP-MPMN model described in 6.1.1, and a default application of FCS with main effects multinomial logistic regression in a large repeated-sampling study of imputation using categorical data from the American Community Survey. The DP-MPMN imputations tended to yield better coverage than FCS-CART overall, but had much worse coverage for a small number of estimands. [Manrique-Vallier and Reiter \(2014b\)](#) also demonstrated the utility of accounting for structural zeros in this model with a population constructed from publicly available data from the U.S. Census. A default version of [Murray and Reiter’s \(2016\)](#) joint model for mixed data types outperformed FCS using the default settings in R’s `mice` package ([Van Buuren and Groothuis-Oudshoorn, 2011](#)) in a large repeated-sampling study with data from the Survey of Income and Program Participation. Evidence suggested that misspecification bias was primarily to blame for FCS’s poor performance.

7.4 Imputation Model Diagnostics

A more obvious way to choose between imputation models is by fitting multiple and choosing the one that appears to fit the data best. Checking the fit of imputation models is challenging, but some approaches have been proposed. For methods that employ univariate regressions, imputers can examine standard diagnostics for those models ([Abayomi, Gelman and Levy, 2008](#), [Su et al., 2011](#)). [Abayomi, Gelman and Levy \(2008\)](#) suggested other diagnostic plots comparing imputed and observed values, primarily comparing marginal and bivariate distributions. Under MAR, the distribution of missing values may be different than the distribution of observed values; [Bondarenko and Raghunathan \(2016\)](#) used estimated response propensities to adjust for this and make diagnostic plots more comparable. [He and Zaslavsky \(2012\)](#) proposed posterior predictive checks, comparing the distribution of estimands computed on the multiply imputed datasets to the distribution of those estimands computed on entirely synthetic datasets generated by the imputation method (see

also Nguyen, Lee and Carlin, 2015). These checks require the imputer to choose relevant estimands and generate many samples from posterior predictive distributions, which can be computationally expensive.

8. CONCLUSION

Over thirty years after Rubin's extensive treatment of MI (Rubin, 1987), experience with the method has cemented its reputation as a principled and practical solution to missing data problems. MI remains an active and fertile research area. While the behavior of the MI estimates have been the subject of intense scrutiny, relatively little is known about the comparative merits of various imputation models that have been proposed in recent years. Considerations based on theoretical findings suggest the use of more flexible imputation models where possible. Empirical evidence also suggests that simple defaults (MVN/log-linear models, or default FCS imputation using simple imputation models such as PMM with linear mean functions or regression models including only main effects) should be avoided, or at least carefully scrutinized.

Nonparametric Bayesian methods for generating imputations have recently emerged as a promising technique for generating imputations. In addition to new model development, more work is needed on scalable posterior computation with these models. In addition, the heuristic justification for why Bayesian MI "tends to be proper" is based on the asymptotic behavior of parametric Bayesian models (Rubin, 1987). It would be interesting to revisit this argument from the perspective of Bayesian nonparametric models, where the asymptotics are more involved (see Rousseau, 2016 for a recent review). For example, can semiparametric Bernstein-von-Mises results be derived for likely targets of MI inference under Bayesian nonparametric models used for imputation?

Joint-sequential approaches appear understudied and underutilized in the literature, perhaps because they currently require more intervention to set up. More research is needed on the implications of choosing different permutations of the variables in joint-sequential approaches. Further development of algorithmic approaches for selecting good joint-sequential variable orderings in the same vein as Li et al. (2014) would also be welcome. There remains considerable work to be done in characterizing the behavior of FCS approaches to generating imputations; while some theoretical results exist, they are limited in scope and do not address some of the most effective variants of these algorithms (including PMM and CART).

More empirical comparisons of imputation methods and models are also needed. The field would benefit greatly from a repository of ready-to-use synthetic populations constructed from real data files. A common set of samples from these populations complete with missing values already generated would allow for easy comparisons across methods. A forward-thinking statistical agency could kickstart this repository, providing a public good (and possibly improving the state of their own missing data imputation routines) by sponsoring an imputation challenge in the spirit of a Kaggle competition.

The applications of MI have grown far beyond imputing item missing data in public use files: MI is used with synthetic data for disclosure limitation (Rubin, 1993, 2002, Raghunathan, Reiter and Rubin, 2003), to adjust for measurement error (Cole, Chu and Greenland, 2006, Blackwell, Honaker and King, 2015) and to perform statistical matching/data fusion (Rässler, 2004, Reiter, 2012, Fosdick, DeYoreo and Reiter, 2016). In these new settings, the amount of missing data can be much greater than typical applications of MI for item missing data, and imputation model development, selection and assessment is even more consequential. We expect that new models and methods for multiple imputation will be an active research area for the foreseeable future.

APPENDIX: SOFTWARE FOR MULTIPLE IMPUTATION

Pointers to many software implementations of MI methods are available at <http://www.stefvanbuuren.nl/mi/Software.html>, an updated version of Appendix A of Van Buuren (2012). As of December 2017, it is missing links to R packages for several nonparametric Bayesian joint models: These include the R packages MixedDataImpute (imputation for mixed continuous and categorical missing values using the model in Murray and Reiter, 2016), NPBayesImpute (imputation for multivariate categorical data, possibly with structural zeros, as presented in Si and Reiter, 2013, Manrique-Vallier and Reiter, 2014a, 2014b) and NestCategBayesImpute (imputation for multivariate categorical data with hierarchical data structures, as described in Hu, Reiter and Wang, 2017).

ACKNOWLEDGMENTS

The author gratefully acknowledges support from the National Science Foundation under grant numbers SES-1130706, SES-1631970 and DMS-1043903. Any

opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the funding agencies.

REFERENCES

- ABAYOMI, K., GELMAN, A. and LEVY, M. (2008). Diagnostics for multivariate imputations. *J. Roy. Statist. Soc. Ser. C* **57** 273–291. [MR2440009](#)
- AKANDE, O., LI, F. and REITER, J. (2017). An empirical comparison of multiple imputation methods for categorical data. *Amer. Statist.* **71** 162–170. [MR3668704](#)
- ANDRIDGE, R. R. and LITTLE, R. J. A. (2010). A review of hot deck imputation for survey non-response. *Int. Stat. Rev.* **78** 40–64.
- ARNOLD, B. C., CASTILLO, E. and SARABIA, J. M. (2001). Conditionally specified distributions: An introduction. *Statist. Sci.* **16** 249–274. [MR1874154](#)
- ARNOLD, B. C. and PRESS, J. S. (1989). Compatible conditional distributions. *J. Amer. Statist. Assoc.* **84** 152–156. [MR0999673](#)
- AUDIGIER, V., HUSSON, F. and JOSSE, J. (2016). Multiple imputation for continuous variables using a Bayesian principal component analysis. *J. Stat. Comput. Simul.* **86** 2140–2156. [MR3491013](#)
- AUDIGIER, V., HUSSON, F. and JOSSE, J. (2017). MIMCA: Multiple imputation for categorical variables with multiple correspondence analysis. *Stat. Comput.* **27** 501–518. [MR3599686](#)
- BANERJEE, A., MURRAY, J. and DUNSON, D. B. (2013). Bayesian learning of joint distributions of objects. In *Proceedings of the 16th International Conference on Artificial Intelligence and Statistics (AISTATS)*, Scottsdale, AZ.
- BARNARD, J. and RUBIN, D. B. (1999). Miscellanea. Small-sample degrees of freedom with multiple imputation. *Biometrika* **86** 948–955.
- BERNAARDS, C. A., BELIN, T. R. and SCHAFER, J. L. (2007). Robustness of a multivariate normal approximation for imputation of incomplete binary data. *Stat. Med.* **26** 1368–1382. [MR2345726](#)
- BLACKWELL, M., HONAKER, J. and KING, G. (2015). A unified approach to measurement error and missing data. *Sociol. Methods Res.* **46** 303–341.
- BÖHNING, D., SEIDEL, W., ALFÓ, M., GAREL, B., PATILEA, V., WALTHER, G., DI ZIO, M., GUARNERA, U. and LUZI, O. (2007). Imputation through finite Gaussian mixture models. *Comput. Statist. Data Anal.* **51** 5305–5316. [MR2370885](#)
- BONDARENKO, I. and RAGHUNATHAN, T. (2016). Graphical and numerical diagnostic tools to assess suitability of multiple imputations and imputation models. *Stat. Med.* **35** 3007–3020. [MR3528239](#)
- BREIMAN, L. (2001). Random forests. *Mach. Learn.* **45** 5–32.
- BREIMAN, L., FRIEDMAN, J., STONE, C. J. and OLSHEN, R. A. (1984). *Classification and Regression Trees*. Wadsworth Advanced Books and Software, Belmont, CA. [MR0726392](#)
- BURGETTE, L. F. and REITER, J. P. (2010). Multiple imputation for missing data via sequential regression trees. *Am. J. Epidemiol.* **172** 1070–1076.
- CARPENTER, J. and KENWARD, M. (2013). *Multiple Imputation and Its Application*, 1st ed. Wiley, New York.
- CHEN, J. and SHAO, J. (2000). Nearest neighbor imputation for survey data. *J. Off. Stat.* **16** 113–131.
- COLE, S. R., CHU, H. and GREENLAND, S. (2006). Multiple-imputation for measurement-error correction. *Int. J. Epidemiol.* **35** 1074–1081.
- COLLINS, L. M., SCHAFER, J. L. and KAM, C. M. (2001). A comparison of inclusive and restrictive strategies in modern missing data procedures. *Psychol. Methods* **6** 330–351.
- DEYOREO, M., REITER, J. P. and HILLYGUS, D. S. (2017). Bayesian mixture models with focused clustering for mixed ordinal and nominal data. *Bayesian Anal.* **12** 679–730. [MR3655872](#)
- DOOVE, L. L., VAN BUUREN, S. and DUSSELDORP, E. (2014). Recursive partitioning for missing data imputation in the presence of interaction effects. *Comput. Statist. Data Anal.* **72** 92–104. [MR3139350](#)
- DRECHSLER, J. (2010). Multiple imputation of missing values in the wave 2007 of the IAB Establishment Panel. IAB Discussion Paper.
- DUNSON, D. B. and XING, C. (2009). Nonparametric Bayes modeling of multivariate categorical data. *J. Amer. Statist. Assoc.* **104** 1042–1051. [MR2562004](#)
- ELLIOTT, M. R. and STETTLER, N. (2007). Using a mixture model for multiple imputation in the presence of outliers: The “healthy for life” project. *J. Roy. Statist. Soc. Ser. C* **56** 63–78. [MR2339163](#)
- FITHIAN, W. and JOSSE, J. (2017). Multiple correspondence analysis and the multilogit bilinear model. *J. Multivariate Anal.* **157** 87–102. [MR3641738](#)
- FOSDICK, B. K., DEYOREO, M. and REITER, J. P. (2016). Categorical data fusion using auxiliary information. *Ann. Appl. Stat.* **10** 1907–1929. [MR3592042](#)
- GEBREGZIABHER, M. and DESANTIS, S. M. (2010). Latent class based multiple imputation approach for missing categorical data. *J. Statist. Plann. Inference* **140** 3252–3262. [MR2659852](#)
- GELMAN, A., CARLIN, J. B., RUBIN, D. B., VEHTARI, A., DUNSON, D. B. and STERN, H. S. (2014). *Bayesian Data Analysis*, 3rd ed. CRC Press, Boca Raton, FL. [MR3235677](#)
- HE, Y. and ZASLAVSKY, A. M. (2012). Diagnosing imputation models by applying target analyses to posterior replicates of completed data. *Stat. Med.* **31** 1–18. [MR2868986](#)
- HE, Y., ZASLAVSKY, A. M., LANDRUM, M. B., HARRINGTON, D. P. and CATALANO, P. (2010). Multiple imputation in a large-scale complex survey: A practical guide. *Stat. Methods Med. Res.* **19** 653–670. [MR2744515](#)
- HEITJAN, D. F. and LITTLE, R. J. A. (1991). Multiple imputation for the fatal accident reporting system. *J. Roy. Statist. Soc. Ser. C* **40** 13–29.
- HORTON, N. J., LIPSITZ, S. R. and PARZEN, M. (2003). A potential for bias when rounding in multiple imputation. *Amer. Statist.* **57** 229–232. [MR2016255](#)
- HU, J., REITER, J. P. and WANG, Q. (2017). Dirichlet process mixture models for modeling and generating synthetic versions of nested categorical data. *Bayesian Anal.* **12** 679–703. [MR3737948](#)
- HUGHES, R. A., WHITE, I. R., SEAMAN, S. R., CARPENTER, J. R., TILLING, K. and STERNE, J. A. C. (2014). Joint modelling rationale for chained equations. *BMC Med. Res. Methodol.* **14** 28.

- IBRAHIM, J. G., LIPSITZ, S. R. and CHEN, M. H. (1999). Missing covariates in generalized linear models when the missing data mechanism is non-ignorable. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **61** 173–190. [MR1664045](#)
- IBRAHIM, J. G., CHEN, M. H., LIPSITZ, S. R. and HERRING, A. H. (2005). Missing data methods for generalized linear models: A comparative review. *J. Amer. Statist. Assoc.* **100** 332–346. [MR2166072](#)
- ISHWARAN, H. and JAMES, L. F. (2001). Gibbs sampling methods for stick-breaking priors. *J. Amer. Statist. Assoc.* **96** 161–173. [MR1952729](#)
- KIM, J. K. (2002). A note on approximate Bayesian bootstrap imputation. *Biometrika* **89** 470–477. [MR1913974](#) [MR1913974](#)
- KIM, J. K., BRICK, J. M., FULLER, W. A. and KALTON, G. (2006). On the bias of the multiple-imputation variance estimator in survey sampling. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **68** 509–521. [MR2278338](#)
- KIM, H. J., REITER, J. P., WANG, Q., COX, L. H. and KARR, A. F. (2014). Multiple imputation of missing or faulty values under linear constraints. *J. Bus. Econom. Statist.* **32** 375–386. [MR3238592](#)
- KIM, H. J., COX, L. H., KARR, A. F., REITER, J. P. and WANG, Q. (2015). Simultaneous edit-imputation for continuous microdata. *J. Amer. Statist. Assoc.* **110** 987–999. [MR3420678](#)
- KROPKO, J., GOODRICH, B., GELMAN, A. and HILL, J. (2014). Multiple imputation for continuous and categorical data: Comparing joint multivariate normal and conditional approaches. *Polit. Anal.* **22** 497–519.
- LEE, M. C. and MITRA, R. (2016). Multiply imputing missing values in data sets with mixed measurement scales using a sequence of generalised linear models. *Comput. Statist. Data Anal.* **95** 24–38. [MR3425936](#)
- LI, F., YU, Y. and RUBIN, D. B. (2012). Imputing missing data by fully conditional models: Some cautionary examples and guidelines. Dept. Statistical Science, Duke Univ., Durham, NC.
- LI, F., BACCINI, M., MEALLI, F., ZELL, E. R., FRANGAKIS, C. E. and RUBIN, D. B. (2014). Multiple imputation by ordered monotone blocks with application to the anthrax vaccine research program. *J. Comput. Graph. Statist.* **23** 877–892. [MR3224660](#)
- LIPSITZ, S. R. and IBRAHIM, J. G. (1996). A conditional model for incomplete covariates in parametric regression models. *Biometrika* **83** 916–922.
- LITTLE, R. J. A. (1988). Missing-data adjustments in large surveys. *J. Bus. Econom. Statist.* **6** 287–296.
- LITTLE, R. J. A. and RUBIN, D. B. (2002). *Statistical Analysis with Missing Data*, 2nd ed. Wiley-Interscience, Hoboken, NJ. [MR1925014](#)
- LITTLE, R. J. A. and SCHLUCHTER, M. D. (1985). Maximum likelihood estimation for mixed continuous and categorical data with missing values. *Biometrika* **72** 497–512. [MR0817564](#)
- LIU, J. S. (1994). The collapsed Gibbs sampler in Bayesian computations with applications to a gene regulation problem. *J. Amer. Statist. Assoc.* **89** 958–966. [MR1294740](#)
- LIU, C. and RUBIN, D. B. (1998). Ellipsoidally symmetric extensions of the general location model for mixed categorical and continuous data. *Biometrika* **85** 673–688. [MR1665830](#)
- LIU, J., GELMAN, A., HILL, J., SU, Y.-S. and KROPKO, J. (2014). On the stationary distribution of iterative imputations. *Biometrika* **101** 155–173. [MR3180663](#)
- MANRIQUE-VALLIER, D. and REITER, J. P. (2014a). Bayesian estimation of discrete multivariate latent structure models with structural zeros. *J. Comput. Graph. Statist.* **23** 1061–1079. [MR3270711](#)
- MANRIQUE-VALLIER, D. and REITER, J. P. (2014b). Bayesian multiple imputation for large-scale categorical data with structural zeros. *Surv. Methodol.* **40** 125–134.
- MANRIQUE-VALLIER, D. and REITER, J. P. (2016). Bayesian simultaneous edit and imputation for multivariate categorical data. *J. Amer. Statist. Assoc.* **112** 1708–1719. [MR3750893](#)
- MENG, X.-L. (1994). Multiple-imputation inferences with uncongenial sources of input. *Statist. Sci.* **9** 538–558.
- MENG, X.-L. and ROMERO, M. (2003). Discussion: Efficiency and self-efficiency with multiple imputation inference. *Int. Stat. Rev.* **71** 607–618.
- MORRIS, T. P., WHITE, I. R. and ROYSTON, P. (2014). Tuning multiple imputation by predictive mean matching and local residual draws. *BMC Med. Res. Methodol.* **14** 75.
- MURRAY, J. S. and REITER, J. P. (2016). Multiple imputation of missing categorical and continuous values via Bayesian mixture models with local dependence. *J. Amer. Statist. Assoc.* **111** 1466–1479. [MR3601702](#)
- NGUYEN, C. D., LEE, K. J. and CARLIN, J. B. (2015). Posterior predictive checking of multiple imputation models. *Biom. J.* **57** 676–694.
- NIELSEN, S. F. (2003). Proper and improper multiple imputation. *Int. Stat. Rev.* **71** 593–607.
- OLKIN, I. and TATE, R. F. (1961). Multivariate correlation models with mixed discrete and continuous variables. *Ann. Math. Stat.* **32** 448–465.
- PADDOCK, S. M. (2002). Bayesian nonparametric multiple imputation of partially observed data with ignorable nonresponse. *Biometrika* **89** 529–538. [MR1929160](#)
- RAGHUNATHAN, T. E., REITER, J. P. and RUBIN, D. B. (2003). Multiple imputation for statistical disclosure limitation. *J. Off. Stat.* **19** 1–16.
- RAGHUNATHAN, T. E., LEPKOWSKI, J. M., VAN HOEWYK, J. and SOLENBERGER, P. (2001). A multivariate technique for multiply imputing missing values using a sequence of regression models. *Surv. Methodol.* **27** 85–96.
- RÄSSLER, S. (2004). Data fusion: Identification problems, validity, and multiple imputation. *Aust. J. Stat.* **33** 153–171.
- REITER, J. P. (2002). Satisfying disclosure restrictions with synthetic data sets. *J. Off. Stat.* **18** 531.
- REITER, J. P. (2005). Using CART to generate partially synthetic public use microdata. *J. Off. Stat.* **21** 441.
- REITER, J. P. (2012). Bayesian finite population imputation for data fusion. *Statist. Sinica* **22** 795–811. [MR2954362](#)
- REITER, J. (2017). Discussion: Dissecting multiple imputation from a multi-phase inference perspective: What happens when God’s, imputer’s and analyst’s models are uncongenial? *Statist. Sinica*. **27** 1578–1583. [MR3701498](#)
- REITER, J. P. and RAGHUNATHAN, T. E. (2007). The multiple adaptations of multiple imputation. *J. Amer. Statist. Assoc.* **102** 1462–1471. [MR2372542](#)
- REITER, J. P., RAGHUNATHAN, T. E. and KINNEY, S. K. (2006). The importance of modeling the sampling design in multiple imputation for missing data. *Surv. Methodol.* **32** 143.
- ROBINS, J. M. and WANG, N. (2000). Inference for imputation estimators. *Biometrika* **87** 113–124. [MR1766832](#)

- ROUSSEAU, J. (2016). On the frequentist properties of Bayesian nonparametric methods. *Annual Review of Statistics and Its Application* **3** 211–231.
- RUBIN, D. B. (1981). The Bayesian bootstrap. *Ann. Statist.* **9** 130–134. [MR0600538](#)
- RUBIN, D. B. (1987). *Multiple Imputation for Nonresponse in Surveys*. Wiley, New York. [MR0899519](#)
- RUBIN, D. B. (1993). Discussion: Statistical disclosure limitation. *J. Off. Stat.* **9** 461–468.
- RUBIN, D. B. (1996). Multiple imputation after 18+ years. *J. Amer. Statist. Assoc.* **91** 473–489.
- RUBIN, D. B. (2003a). Discussion on multiple imputation. *Int. Stat. Rev.* **71** 619–625.
- RUBIN, D. B. (2003b). Nested multiple imputation of NMES via partially incompatible MCMC. *Stat. Neerl.* **57** 3–18. [MR2055518](#)
- RUBIN, D. B. and SCHAFER, J. L. (1990). Efficiently creating multiple imputations for incomplete multivariate normal data. In *Proc. Statistical Computing Section of the American Statistical Association* 83–88. Amer. Statist. Assoc., Alexandria, VA.
- RUBIN, D. B. and SCHENKER, N. (1986). Multiple imputation for interval estimation from simple random samples with ignorable nonresponse. *J. Amer. Statist. Assoc.* **81** 366–374. [MR0845877](#)
- SCHAFER, J. L. (1997). *Analysis of Incomplete Multivariate Data*. Chapman & Hall, London. [MR1692799](#)
- SCHAFER, J. L. (2003). Multiple imputation in multivariate problems when the imputation and analysis models differ. *Stat. Neerl.* **57** 19–35. [MR2055519](#)
- SCHENKER, N. and TAYLOR, J. M. G. (1996). Partially parametric techniques for multiple imputation. *Comput. Statist. Data Anal.* **22** 425–446.
- SCHIFELING, T. A. and REITER, J. P. (2016). Incorporating marginal prior information in latent class models. *Bayesian Anal.* **11** 499–518. [MR3472000](#)
- SEAMAN, S. R. and HUGHES, R. A. (2016). Relative efficiency of joint-model and full-conditional-specification multiple imputation when conditional models are compatible: The general location model. *Stat. Methods Med. Res.* DOI:10.1177/0962280216665872.
- SETHURAMAN, J. (1994). A constructive definition of Dirichlet priors. *Statist. Sinica* **4** 639–650. [MR1309433](#)
- SHAH, A. D., BARTLETT, J. W., CARPENTER, J., NICHOLAS, O. and HEMINGWAY, H. (2014). Comparison of random forest and parametric imputation models for imputing missing data using MICE: A CALIBER study. *Am. J. Epidemiol.* **179** 764–774.
- SI, Y. and REITER, J. P. (2013). Nonparametric Bayesian multiple imputation for incomplete categorical variables in large-scale assessment surveys. *J. Educ. Behav. Stat.* **38** 499–521.
- STUART, E. A., AZUR, M., FRANGAKIS, C. and LEAF, P. (2009). Multiple imputation with large data sets: A case study of the children’s mental health initiative. *Am. J. Epidemiol.* **169** 1133–1139.
- SU, Y.-S., GELMAN, A., HILL, J., YAJIMA, M. et al. (2011). Multiple imputation with diagnostics (mi) in R: Opening windows into the black box. *J. Stat. Softw.* **45** 1–31.
- VAN BUUREN, S. (2007). Multiple imputation of discrete and continuous data by fully conditional specification. *Stat. Methods Med. Res.* **16** 219–42.
- VAN BUUREN, S. (2012). *Flexible Imputation of Missing Data*. CRC Press, Boca Raton, FL.
- VAN BUUREN, S. and GROOTHUIS-OUDSHOORN, K. (2011). mice: Multivariate imputation by chained equations in R. *J. Stat. Softw.* **45** 1–67.
- VAN BUUREN, S. and OUDSHOORN, K. (1999). *Flexible Multivariate Imputation by MICE*. TNO Prevention Center, Leiden, The Netherlands.
- VAN BUUREN, S., BRAND, J. P. L., GROOTHUIS-OUDSHOORN, C. G. M. and RUBIN, D. B. (2006). Fully conditional specification in multivariate imputation. *J. Stat. Comput. Simul.* **76** 1049–1064. [MR2307507](#)
- VERMUNT, J. K., VAN GINKEL, J. R., VAN DER ARK, L. A. and SIJTSMA, K. (2008). Multiple imputation of incomplete categorical data using latent class analysis. *Sociol. Method.* **38** 369–397.
- VIDOTTO, D., VERMUNT, J. K. and KAPTEIN, M. C. (2015). Multiple imputation of missing categorical data using latent class models: State of art. *Psychol. Test Assess. Model.* **57** 542–576.
- VINK, G., FRANK, L. E., PANNEKOEK, J. and VAN BUUREN, S. (2014). Predictive mean matching imputation of semicontinuous variables. *Stat. Neerl.* **68** 61–90. DOI:10.1111/stan.12023. [MR3168318](#)
- WANG, N. and ROBINS, J. M. (1998). Large-sample theory for parametric multiple imputation procedures. *Biometrika* **85** 935–948. [MR1666715](#)
- XIE, X. and MENG, X.-L. (2017). Dissecting multiple imputation from a multi-phase inference perspective: What happens when God’s, imputer’s and analyst’s models are uncongenial? *Statist. Sinica.* **27** 1485–1545. [MR3701490](#)
- XU, D., DANIELS, M. J. and WINTERSTEIN, A. G. (2016). Sequential BART for imputation of missing covariates. *Biostatistics* **17** 589–602. [MR3603956](#)
- ZHU, J. and RAGHUNATHAN, T. E. (2015). Convergence properties of a sequential regression multiple imputation algorithm. *J. Amer. Statist. Assoc.* **110** 1112–1124. [MR3420688](#)