**Recursive Partitioning Imputation Versus Full Information Maximum Likelihood in Finite
Mixture Modeling**

Marcus R. Waldman¹ and Katherine E. Masyn²

¹Harvard Graduate School of Education

²Division of Epidemiology and Biostatistics, School of Public Health, Georgia State University

Author Note

Marcus R. Waldman <https://orcid.org/0000-0002-3288-4803>

Marcus R. Waldman is now at the Department of Health Promotion, College of Public Health, University of Nebraska Medical Center.

We have no known conflict of interest to disclose.

Correspondence concerning this article should be addressed to Marcus Waldman, College of Public Health, University of Nebraska Medical Center, 42nd and Emile, Omaha, NE 68198.

Abstract

It is well established that omitting important variables that are related to the propensity for missingness can lead to biased parameter estimates and invalid inference. Nevertheless, researchers conducting a person-centered analysis ubiquitously adopt a full information maximum likelihood (FIML) approach to treat missing data in a manner that assumes the missingness is only related to the observed indicators and is not related to any external variables. Such an assumption is generally considered overly restrictive in the behavioral sciences where the data are observational in nature. At the same time, previous research has discouraged the adoption of multiple imputation to treat missing data in person-centered analyses because traditional imputation models make a single-class assumption and do not reflect the multiple group structure of data with latent subpopulations (Enders & Gottschall, 2011). However, more modern imputation models that rely on recursive partitioning do not impose a single-class structure to the data. Focusing on latent profile analysis, we demonstrate in simulations that in samples of $N = 1,200$ or greater, recursive partitioning imputation algorithms can effectively incorporate external information from auxiliary variables to attenuate nonresponse bias better than FIML and multivariate normal imputation. Moreover, we find that recursive imputation models lead to confidence intervals with adequate coverage and they better recover posterior class probabilities than alternative missing data strategies. Taken together, our findings point to the promise and potential of multiple imputation in person-centered analyses once remaining methodological gaps around pooling and class enumeration are filled.

Keywords: missing data, inclusive strategy, auxiliary variables, finite mixture modeling, latent profile analysis, classification and regression trees, random forests

Recursive Partitioning as an Alternative to Full Information Maximum Likelihood in Latent Profile Analysis

Over the last few decades, methodologists have implored researchers in the behavioral sciences to adopt “state of the art” (Schafer & Graham, 2002) strategies to address missing data. Missing at random (MAR; Rubin, 1976) procedures are especially encouraged because they leverage observable information to improve the robustness of conclusions by mitigating nonresponse bias (Hancock, Stapleton, & Mueller, 2019; T. D. Little, Jorgensen, Lang, & Moore, 2014; Wilkinson, 1999). Although valid inference with MAR strategies is technically only guaranteed if the missing data mechanism is ignorable (i.e., the propensity for missingness is explained only by observables), MAR strategies can also protect against nonresponse bias when the mechanism is nonignorable if the missing values are well predicted by observed variables (Enders, 2008; van Buuren, 2018). The most common MAR strategies employed in psychology and the behavioral sciences include multiple imputation (MI; Rubin, 1987) and full information maximum likelihood (FIML) (Allison, 2002; Enders, 2010; R. J. Little & Rubin, 2002; Molenberghs & Kenward, 2007).¹

Currently, researchers conducting latent profile analysis (LPA) rely almost exclusively on FIML to treat missing data. To highlight this fact, we identified thirty highly cited LPA studies in the literature to inform the simulation study later presented in this paper. We found that FIML was reported as the missing data strategy used to handle missing data in all cases, provided that a

¹ Weighted adjustment using propensity scores is an alternative MAR strategy to FIML and MI, and it is used heavily in large-scale surveys and in the clinical sciences (Molenberghs & Kenward, 2007). Although some studies in psychology have advocated for a weighted adjustment approach with promising simulation results in attrition contexts (Hayes & McArdle, 2017; Hayes, Usami, Jacobucci, & McArdle, 2015; McArdle, 2013), weighted adjustment is generally a complete case analysis (R. J. Little & Rubin, 2002) and is limited in its ability to treat multivariate missing data patterns like those observed in LPA (Hayes, 2017). We, therefore, do not consider weighted adjustment approaches.

missing data strategy was reported at all. None of the authors reported specifying auxiliary variables (AVs) that are external to the LPA to address the missing data problem. This makes clear that researchers conducting LPA are relying heavily on the viability of the assumption that the missing data mechanism is ignorable absent any external information. Such an ignorability assumption is generally considered overly restrictive given the complicated and imperfectly controlled contexts in which behavioral science data are collected.

To mitigate the associated nonresponse bias in observational data, methodologists strongly encourage researchers to adopt an “inclusive analysis strategy” (Collins, Schafer, & Kam, 2001) and incorporate many AVs to treat missing data with MAR procedures (Enders, 2010; T. D. Little et al., 2014; Rubin, 1996; Schafer, 1997; Schafer & Graham, 2002). When analyzing data using covariance structure models (i.e., regression and path analysis, factor models, or structural equation models), an inclusive strategy under FIML can be accomplished using specialized specification rules (Graham, 2003) or two-step procedures (Savalei & Bentler, 2009; Yuan & Bentler, 2000). Both augment the analytic model to factor in information from AVs. However, as we elaborate later, it is not obvious how to appropriately augment the mixture models used in LPA to incorporate AVs, and, even if it is possible, any developed procedure would likely be plagued by several limitations in real-world applications. Indeed, even in covariance structure models, it is common to quickly experience convergence issues when incorporating AVs in a FIML approach (Enders, 2010).

In contrast, MI is generally considered more flexible and better able to incorporate AVs than FIML because the model used to treat missing data (i.e., the imputation model) can be distinct from the analytic model used to make inferences and answer substantive research questions (Meng, 1994; Rubin, 1996; van Buuren, 2018). Moreover, MI can accommodate the

inevitable complications that arise in the real world, including treating the missingness on AVs, which themselves are comprised of multiple data types (continuous, categorical, etc.). This is especially true if the so-called fully conditional specification (van Buuren, Brand, Groothuis-Oudshoorn, & Rubin, 2006) so that missing data are treated using multivariate imputation by chained equations (MICE) (Azur, Stuart, Constantine, & Leaf, 2011; van Buuren et al., 2006).

Some have cautioned against using MI (Sterba, 2016) because popular models used to generate the imputed datasets assume that the data exhibit a single group structure (Enders, 2010; Enders & Gottschall, 2011). In a joint modeling imputation framework, for example, it is typically assumed that continuous data are drawn from a (single) multivariate normal distribution (see King, Honaker, Joseph, & Scheve, 2001; Schafer, 1997). Equivalently, if using MICE, the linear regression models popular in imputing continuous data assume a single-class structure with imputations drawn from a univariate normal distribution. The single-class assumption is problematic in LPA because such an analysis is person-centered in that the explicit goal is to identify multiple hidden subpopulations of individuals. Thus, creating imputed datasets from single-class imputation models risks obfuscating the very subgroups that the researcher intends to illuminate.

However, imputation models need not assume a single-class structure. For example, recursive partitioning regression models utilize a form of clustering to calculate predicted values for each observation. The clustering process assures that the imputations are drawn from a multiple group structure. Although recursive partitioning regression models are available in software when generating imputations using MICE, the performance of these methods has yet to be evaluated for LPA. This indicates a gap in literature, as these methods may provide a better option than FIML in treating missing data.

The purpose of this study is to evaluate whether recursive partitioning imputation models outperform FIML in treating missing data when the MAR assumption is conditional on AVs. Through a simulation study, we show that imputations created from recursive partitioning algorithms can effectively attenuate nonresponse bias and lead to valid inference outside of settings with small class sizes. In this way, we show that recursive partitioning imputation allows for an inclusive missing data strategy and is more effective than FIML in attenuating nonresponse bias outside of small class sizes. Thus, this study opens the door for MI to be an option for applied researchers treating missing data in person-centered analysis.

This paper is organized as follows. We introduce finite mixture models and a brief review into missing data theory as it relates to LPA in the Background. We next describe our simulation study design and motivate the metrics we used to evaluate performance in treating missing data. We then present the results, detail implications to applied researchers, and suggest avenues of research for future studies.

Background

In this section, we provide a brief background on the finite mixture models fit in LPA, we highlight the challenges associated with conducting an inclusive strategy with FIML, and we introduce recursive partitioning as a viable imputation strategy for data containing latent subpopulations.

Finite Mixture Models and LPA

As typically implemented in LPA, individuals are clustered into homogenous subpopulations (or classes) by fitting finite mixture models with Gaussian component densities. It is assumed that the indicators are normally distributed conditional on class membership:

$$[\mathbf{y}_i | \kappa = k] \sim \text{MVN}(\mu_k, \Sigma_k) \quad (1)$$

where $\mathbf{y}_i$ is the vector of latent class indicator data and κ represents the categorical latent variable representing class membership, while μ_k and Σ_k represent the within-class mean vector and variance-covariance matrix for class k , respectively.

Finite mixture models require the researcher to specify the total number of component densities to be fit to the data, denoted K . Typically, this value is not known a priori, so it must be inferred empirically from the data. The number of components is determined by fitting a sequence of mixture models with increasing numbers of classes until convergence becomes intractable. This process is referred to as enumerating the classes. A model within the sequence is then selected. For our purposes of evaluating the performance of recursive partitioning imputation in simulations, we assume that the number of classes in the population is known.

In finite mixture models, the latent variable, κ , is categorical and follows a multinomial distribution:

$$\kappa \sim \text{Multinomial}(\pi_1, \dots, \pi_K) \quad (2)$$

where π_k represents the likelihood of a randomly selected individual belonging to class k , and $k = 1, \dots, K$. From (1) and (2), the corresponding likelihood of observing the complete data Y is given by:

$$\mathcal{L}(Y|\theta) = \prod_{i=1}^n \sum_{k=1}^K \pi_k \mathcal{N}_J(\mathbf{y}_i | \mu_k, \Sigma_k) \quad (3)$$

where $\theta = \{\{\pi_k, \mu_k, \Sigma_k\}: k = 1, \dots, K\}$.

Finite mixture models are a probabilistic clustering algorithm in that each observation is assigned a nonzero probability value for belonging in each of the K classes. We denote the posterior class probability—i.e., the probability that is observation belongs to the k th class given

a set of parameter values θ —as $\tau_k(\mathbf{y}_i|\theta)$. This value can be obtained by simple application of Bayes rule and is given by

$$\tau_k(\mathbf{y}_i|\theta) = \frac{\pi_k \mathcal{N}_j(\mathbf{y}_i|\mu_k, \Sigma_k)}{\sum_{k=1}^K \pi_k \mathcal{N}_j(\mathbf{y}_i|\mu_k, \Sigma_k)}. \quad (4)$$

Inclusive Missing Data Strategy: Limitations with FIML

All missing data strategies can be justified by manipulating and studying the properties of the joint distribution formed by the data and the missing data mechanism. If profile indicator data are missing, then the complete indicator data Y is composed of both the missing and observed elements so that $Y = \{Y^{\text{obs}}, Y^{\text{mis}}\}$. The joint distribution of the complete data and the missing data mechanism represented by the missing data patterns R can then be written as:

$$\Pr(Y, R; \theta, \psi) = \Pr(Y^{\text{obs}}, Y^{\text{mis}}, R; \theta, \psi) \quad (4)$$

where Y represents the complete data with parameters θ , and ψ represents the parameters associated with the missing data mechanism. All MAR missing data strategies assume that the missing data mechanism is ignorable when estimating the θ parameters, resulting in the inferences being drawn from the observed data likelihood. The primary requirement for ignoring R and ψ during estimation is that the data must be MAR: that is, the missing profile indicator data (Y^{mis}) is independent of the missing data patterns (R) conditional on the observed data (Y^{obs}), as given by

$$Y^{\text{mis}} \perp R \mid Y^{\text{obs}} \quad (5)$$

(R. J. Little & Rubin, 2002, pp. 117–119; Rubin, 1976)².

² An additional requirement is distinctiveness of θ and ψ (R. J. Little & Rubin, 2002, pp. 119–120). Because distinctiveness is not considered as important (Schafer, 1997, p. 11), we do not elaborate.

Violations of the MAR assumption in (5) can lead to biased θ estimates and invalid inference. In real-world LPA settings where the data are often observational in nature so that the causes of missing data are beyond the control of the researcher, the MAR assumption is likely untenable. A more tenable assumption is that MAR holds conditional on Y^{obs} and other observables. In the latter setting, valid inference can theoretically be obtained by conditioning on additional variables:

$$Y^{\text{mis}} \perp R \mid (Y^{\text{obs}}, X) \quad (6)$$

where X are the additional variables that are incorporated into the model as AVs. Furthermore, it has been shown that highly predictive AVs can substantially reduce bias, even in settings where data remain missing not at random (Enders, 2010; van Buuren, 2018). For these reasons, methodologists recommend that applied researchers incorporate an inclusive strategy and incorporate many AVs when addressing missing data (Collins et al., 2001; Enders, 2010; Rubin, 1996; Schafer & Graham, 2002).

Outside of mixture modeling, researchers have numerous options to include AVs and adopt an inclusive strategy. If treating missing data with MI, the researcher can specify AVs as part of the imputation model and remove these external variables when fitting the analytic model to the imputed datasets. If the AVs themselves contain missing data, then this missingness can be treated in a straightforward manner under a MICE imputation procedure; the researcher simply includes the AVs as part of the larger chained equation.

Likewise, researchers have several options to include AVs if employing FIML estimation to treat missing data when fitting the covariance structure models associated with regression analysis, factor analysis, and structural equation modeling. An inclusive strategy can be implemented by fitting a “saturated correlates” model (Graham, 2003). A saturated correlates

model augments the analytic model by “working the AVs into the analysis without altering the substantive interpretation of the parameters” (Enders, 2010, p. 133). In this way, the MAR assumption is made more tenable by the inclusion of AVs without making parameter interpretations conditional on these variables, as would occur, for example, if the AVs were simply added as predictors in a regression analysis.

The saturated correlates model can be fit to covariance structure models using either a two-stage approach or specifying models according to Graham’s (2003) rules. The two-stage analysis mimics a method-of-moments estimation procedure. In the first stage, one estimates the sufficient statistics of the model (i.e., the variance-covariance matrix) in such a way that the estimates are informed by the AVs. With the variance-covariance matrix estimated, one can then fit the covariance structure model of interest (Savalei & Bentler, 2009; Yuan & Bentler, 2000). Alternatively, a saturated correlates model can be fit by correlating the AVs to any disturbance or uniqueness terms according to Graham’s (2003) rules.³ By correlating the AVs with these error terms, Graham’s (2003) rules “[transmit] the information [needed for the MAR assumption to be tenable] from the auxiliary variables to the analysis model variables without affecting the interpretation of the parameter estimates” (Enders, 2010, p. 135).

However, saturated correlates models have several practical limitations and is less flexible than MI. First, if the AVs themselves contain missing values, then the researcher either must adopt a listwise deletion strategy or specify the AVs as endogenous variables. If the latter approach is taken, then strong parametric assumptions about the distribution of the AVs must be made. For example, it is not uncommon to assume that binary AVs are treated as continuous so

³ We refer the reader to Enders (2010, Chapter 5) for illustrative examples of how to implement Graham’s (2003) rules with different types of covariance structure models.

as to specify the necessary correlations required as part of Graham's (2003) rules. Additionally, convergence issues are common with a saturated correlates approach, and convergence becomes increasingly more difficult to achieve as the number of AVs increases. Between the distributional assumptions implied by endogeneity and the convergence issues experienced with many AVs, researchers often need to make compromises in deciding AVs to exclude (Enders, 2010). Unlike saturated correlates models, such compromises are generally not necessary when conducting MI. This is because MICE is highly flexible in treating varied data types and is more robust against nonconvergence (van Buuren, 2018).

In addition to these practical limitations, it remains unclear how to work the AVs into the analysis without causing the definitions of the classes to change and individuals to switch classes in an LPA. In the Appendix, we show that specifying a mixture model according to Graham's (2003) rules can alter resulting class-specific mean estimates, even when the data are complete. Class means are often used to define the classes when naming the classes. Therefore, specifying a mixture model according to Graham's (2003) rules risks substantively changing the very definitions of the classes. In addition, there is evidence that individuals switched classes when the mixture was specified according to Graham's (2003) rules. This is shown by nonequivalent marginal probability estimates across the classes between a mixture model fit per the usual specification as opposed to a mixture model specified according to Graham's (2003) rules.

Why do Graham's (2003) rules alter key parameter interpretations when the latent variable is categorical (i.e., mixture models) and not when the latent variable is continuous (e.g., factor analysis)? In brief, implementing Graham's (2003) rules with categorical latent variables results in the AVs effectively becoming additional class indicators; the corresponding fitted mixture model results in parameter estimates that best explain differences in relations between

the AVs and the class indicators. However, AVs are external variables, so content validity considerations demand that they should not influence the definitions of the classes. What is clear is that by altering the parameters that define the classes, Graham's (2003) rules do not accomplish their fundamental goal of preserving interpretations in mixture models.

If Graham's (2003) rules fail to preserve key interpretations, it seems reasonable to question whether there is a two-stage procedure that could be developed to accomplish this goal, as is available with covariance structure models. Unfortunately, we do not consider a two-step approach possible. Mixture models require person-level data and cannot be fit to a covariance matrix, unlike the covariance structure models in regression analysis, factor analysis, or structural equation modeling. This is because mixture models estimate higher order moments in the data than simply the means and covariance matrices modeled in covariance structure models (Masyn, 2013; McLachlan & Peel, 2004). Consequently, there is no obvious way to translate the first step from the two-step approach by estimating the sufficient statistics.

In summary, it remains unclear how to conduct an inclusive missing data strategy when fitting a mixture model using FIML without changing the definitions of the classes. Even if a clever set of rules akin to Graham's (2003) rules could be identified for mixtures that accomplish this goal, there are still practical limitations that make an inclusive strategy with FIML unpalatable. These include increased convergence issues or invoking strong parametric assumptions to treat missing values in the AVs. MI is more flexible in treating a variety of data types and is less prone to practical problems like convergence issues, making the method more amenable to an inclusive strategy.

Theoretical Advantages and Current Limitations of MI in LPA

As an alternative to FIML, MI is better equipped to incorporate AVs (Enders, 2010; van Buuren, 2018). Unlike FIML, the multiple phase approach under MI separates and distinguishes the problem of treating missing data (i.e., the imputation phase of the analysis) from how substantive questions would be answered had the data been complete using the intended analytic model (i.e., the subsequent the analysis phase). Specifically, the model used to treat the missing data through generating imputations can be “uncongenial” (Meng, 1994) with the finite mixture model that would ultimately be fit to make inferences in LPA. This explicit separation of the missing data problem from the analysis problem protects the AVs from unduly influencing class membership above and beyond their role in treating the missing data. Additionally, saturated correlates models exhibit convergence difficulties when fit with a large number of AVs, a problem that is not observed with MI (Enders, 2010). Finally, unlike a saturated correlates approach, which places strict parametric assumptions on the AVs if these variables themselves contain missingness, missingness in the AVs can be addressed using nonparametric MI algorithms that are also more robust to misspecification issues (van Buuren, 2018). Thus, from a theoretical perspective, MI is often a more attractive alternative for dealing with missing data than FIML. Nevertheless, nuanced issues remain that threaten the validity of inferences drawn from mixture models when MI is used to treat missing data that contain latent subpopulations.

To understand these nuanced issues, it is important to consider some of the theoretical underpinnings that justify MI as a missing data strategy. MI, like FIML, requires that the missing data mechanism be ignorable. Unlike FIML, however, MI takes a Bayesian perspective in treating the missing data and relies on the Bayesian Central Limit theorem for valid frequentist

inference. In particular, Little and Rubin (2002, p. 210) show that the posterior density for parameters θ given the observed data is given by

$$\Pr(\theta|Y^{\text{obs}}) \propto \int_{Y^{\text{mis}}} \Pr(\theta|Y^{\text{mis}}, Y^{\text{obs}}) \Pr(Y^{\text{mis}}|Y^{\text{obs}}, \theta) dY^{\text{mis}} \quad (7)$$

where Y^{obs} and Y^{mis} represent observed and missing data, respectively. AVs can be incorporated in (7) to ensure that the ignorability requirement is met or to mitigate the effect of violations to the MAR assumption if the missing data mechanism is nonignorable:

$$\Pr(\theta|Y^{\text{obs}}) \propto \int_{Y^{\text{mis}}} \Pr(\theta|Y^{\text{mis}}, Y^{\text{obs}}) \Pr(Y^{\text{mis}}|Y^{\text{obs}}, X, \theta) dY^{\text{mis}} \quad (8)$$

where X is the set of AVs. The integral in (8) implies that valid inferences about θ can be accomplished through a marginalizing process over all plausible values for Y^{mis} under the assumption that true posterior predictive distribution for the missing data (i.e., $[Y^{\text{mis}}|Y^{\text{obs}}, X, \theta]$) can be sampled independently (R. J. Little & Rubin, 2002; Rubin, 1987). Sampling from the true posterior predictive distribution $[Y^{\text{mis}}|Y^{\text{obs}}, X, \theta]$ is guaranteed to lead to “Bayesianly proper” (Schafer, 1997) imputed datasets and valid Bayesian inference. In large samples and when the regularity conditions are met, the Bayesian Central Limit theorem ensures that Bayesian inference is consistent with frequentist inference (Freedman, 1963, 1965; Le Cam, 1986). However, imputations sampled from a distribution other than the true posterior predictive distribution are not theoretically guaranteed to lead to valid Bayesian inference, but, in many cases, they can perform well in real-world situations (van Buuren, 2018). Unfortunately, LPA is not one of those situations.

In fact, multiple studies have shown that if missingness depends on group membership (such as subpopulations), then bias can result if group membership information is not

incorporated in the imputation model (Collins et al., 2001; Enders & Gottschall, 2011). This is because the imputed datasets may not reflect key features of the data's structure if group membership information is not incorporated into the imputation models. As part of an inclusive strategy, this means that the researcher either specifies many interaction terms with group membership information, or the researcher fits separate imputation models across the groups entirely (Allison, 2002; Enders, 2010).

However, in a person-centered analysis, an individual's group (i.e., their latent subpopulation) is not directly observable. This precludes the adoption of inclusive analytic strategies that explicitly incorporate group information. The fact that group membership cannot be directly observed has negative consequences when single-class imputation models are used to construct plausible values. This is because the imputations generated from a single-class model fail to reproduce the multiple group structure of the data, resulting in imputations that are not Bayesianly proper. This is true even if the data are missing completely at random (Enders, 2010; Enders & Gottschall, 2011). In the supplemental materials, we provide an intuitive explanation to help the visualize why single-class imputation models are problematic even if the data are missing completely at random.

Conditional Mean Imputation by Recursive Partitioning Prediction Models

Recursive partitioning refers to a stratification algorithm by which partitions are formed by successively splitting up the data in a hierarchical fashion, ultimately leading to rectangular clusters of observations (Breiman, 1984). Specifically, an automated procedure identifies a splitting variable and then selects the optimal cut point on that variable in a manner that results in the maximum reduction on a so-called impurity metric. This variable selection and splitting process repeats itself until a termination condition is met: either the change on the impurity

measure no longer sufficiently decreases, or the number of observations in a resulting partition reaches some minimum value (Strobl, Malley, & Tutz, 2009). The final partition (i.e., terminal node) for any given observation can be determined by a set of decision rules that follow a tree-like structure, often referred to as a decision tree.

In regression contexts, Breiman's (1984) classification and regression tree algorithm (CART) is a highly popular application of recursive partitioning for prediction. Specifically, upon termination of the recursive partitioning process, a summary statistic (e.g., mean, median, etc.) using only outcome values from observations in the terminal node are utilized (see Strobl, Malley, & Tutz, 2009 for more details). Because the same variable is chosen as the splitting variable multiple times, CARTs can accommodate nonlinear and even discontinuous functions. Moreover, CARTs automatically search for the presence of high-order interaction terms (Burgette & Reiter, 2010; Doove, van Buuren, & Dusseldorp, 2014) in the recursive partitioning process.

An example of a decision tree from CART with a minimum terminal node size of 200 observations is illustrated in Panel A in Figure 1. Here, Y_3 is an outcome variable with two predictors, Y_1 and Y_2 . A predicted Y_3 value can be obtained by first identifying the terminal node of an observation using that observation's Y_1 and Y_2 value and then calculating a summary statistic using only the Y_3 values for observations clustered in the identified terminal node. The corresponding rectangular partitions that result from the decision tree are shown in Panel B of Figure 1.

The machine learning community has since extended Breiman's (1984) CART algorithm with post hoc procedures to improve predictions coming from recursive partitioning regression algorithms. The random forest (RF; Breiman, 2001) algorithm represents one important

extension. RFs are known to produce more accurate predictions that better generalize to out-of-sample data than predictions from CART with pruning alone (Hastie, Tibshirani, & Friedman, 2009).

Naturally, the machine learning community was also first to employ CART and RF predictions for treating missing values as part of a conditional mean imputation procedure (Bárcena & Tusell, 2000; Conversano & Cappelli, 2002; Conversano & Siciliano, 2009; Creel & Krotki, 2006; D'Ambrosio, Aria, & Siciliano, 2012; Ishwaran, Kogalur, Blackstone, & Lauer, 2008; Stekhoven & Bühlmann, 2011). Although the predicted value may best approximate the true value on average, it is well known that conditional mean imputation does not lead to valid statistical inference because it does not account for sampling variability or other important sources of variability (van Buuren, 2018). In particular, conditional mean imputation leads to speciously narrow confidence intervals, inflated Type-I error rates, and, critically in multivariate application, positively biased point estimates among association parameters that define the relations among the variables (R. J. Little & Rubin, 2002, Chapter 4). One may expect that biased association parameters would translate to bias across parameters of substantive interest, such as the class-specific means of the profile indicators.

CART and RF Imputation: A Multiple Group Missing Data Approach

To resolve the problems associated with conditional mean imputation, several researchers have extended recursive partitioning imputation algorithms to reflect the important sources of variability around predicted values (Burgette & Reiter, 2010; Doove et al., 2014; Shah, Bartlett, Carpenter, Nicholas, & Hemingway, 2014; Sovilj et al., 2016). In order to capture important sources of variability, values are randomly drawn from a donor pool that includes all values that would have otherwise been used to inform the aggregate summary statistic to construct a

predicted value. Randomly sampling from the donor pool implies that the plausible values are constructed by sampling from an empirical distribution. As applied to the example in Figure 1, there are six possible donor pools to impute at Y_3 value from Y_1 and Y_2 because each donor pool corresponds directly to the terminal nodes. Thus, the donor pool size is equivalent to the terminal node size. The empirical distribution of the donor pools is shown in Panel C. In practice a minimum donor pool size of 200 observations is generally far too coarse for imputation to perform well and software typically defaults to donor pool sizes of as few as 5 observations.

Valid inference is only guaranteed if the empirical distribution sufficiently approximates the true posterior predictive distribution of the missing values, $[Y^{\text{mis}}|Y^{\text{obs}}, X, \theta]$ in (8). However, previous empirical studies from imputation methods that utilize donor pools to sample from empirical distributions have shown strong performance so long as the empirical distribution approximates the true distribution well (van Buuren, 2018).

Imputation algorithms that incorporate random partitioning through CART or RF models work in much the same way as when the goal is prediction. Specifically, CART and RF regression models can be embedded in a MICE procedure in multivariate missing data settings. MICE proceeds by sweeping across the data and sequentially fitting univariate CART or RF regression models to construct a donor pool for the missing values. At each sequential step, imputations are then drawn from the donor pool and substituted for the missing values. Sweeping across the dataset in this way repeats multiple times so that a Markov chain Monte Carlo (MCMC) is formed (van Buuren et al., 2006). Multiple imputed datasets are then constructed by sampling from the corresponding MCMC chain. We refer to the corresponding CART and RF MICE procedures as MICE-CART or MICE-RF, respectively.

In the R (R Core Team, 2020) computing environment, researchers can implement MICE-CART using the `mice` (van Buuren & Groothuis-Oudshoorn, 2010) package. MICE-RF can be implemented using either the `mice` or `CALIBERrfimpute` (Shah, Bartlett, Hemingway, Nicholas, & Hingorani, 2014) packages. For convenience, only the `mice` package is employed in this study. It is important to note that neither implementation prunes to increase parsimony in the decision trees used to construct the donor pools. The advantages and disadvantages of not pruning are detailed later in the Discussion section.

MICE-CART and MICE-RF are promising alternatives to imputing with single-class models. Doove et al. (2014) and Shah et al. (2014) showed that generating imputed data from MICE-CART and MICE-RF models substantially reduces bias when high-order interactions in multiple regression contexts are present but not modeled, even compared to alternative nonparametric methods, such as predictive mean matching. These findings are consistent with other literature demonstrating the utility of recursive partitioning algorithms to detect interactions (Jacobucci, Grimm, & McArdle, 2017; Morgan & Sonquist, 1963; Strobl et al., 2009).

The favorable performance of MICE-CART and MICE-RF with unmodeled interaction terms has direct implications for LPA. In the hypothetical event that class membership was known, one could specify appropriate interactions with dummy variables to generate proper imputations. Thus, the challenge of imputing class indicators without knowing subpopulation membership is highly related to the challenge of imputing data when interactions are present but are left unmodeled, a scenario where CARTs and RFs excel.

Another reason why MICE-CART and MICE-RF are expected to perform well is because they incorporate clustering when constructing imputations (van Buuren, 2018). Specifically, the

rectangular partitions are themselves clusters that represent homogenous subgroups and these clusters are used to form the donor pool. The clustering process in the first step guarantees that the imputations have some multiple group structure. The structure implied by imputed data generated from MICE-CART and MICE-RF may be qualitatively different than the true structure of data with subpopulations in LPA. Visually, this can be seen in Panel B of Figure 1 where the clusters implied by the true latent subpopulations are circular in orientation and not rectangularly shaped as would be implied by the rectangular partitions. Nevertheless, by clustering observations, recursive partitioning imputation models clearly avoids the problems of imputing from a single-class model.

Finally, an important application of mixture modeling separate from latent profile analysis is semiparametric density estimation where a multivariate density is approximated using finite mixture models. There, the goal is to approximate the joint distribution of the data in a flexible and nonparametric manner. MICE-CART and MICE-RF are both nonparametric methods, implying that the imputation models are more aligned with the goal of semiparametric density estimation than traditional, parametric imputation models.

Simulation Design

To study the performance of MICE-CART and MICE-RF in addressing missing data when the missingness depends on external variables, we conducted a simulation study. We used the Web of Science database to identify 30 frequently cited articles employing LPA that were published between 2008-2018 in developmental and educational psychology journals in order to ensure decisions regarding parameter values reflect those observed in applied literature. From each study, we recorded the sample size of the analytic dataset, the number of classes selected, the number of indicators used to construct the profiles, the proportion of observations in the

smallest profile, and the observed entropy value (Table 1). We also noted typical rates of missingness and how missing data were treated for each study. We refer the reader to the supplemental material for more information regarding the inclusion and exclusion criteria. There, we also provide the data that we extracted when reviewing each of the studies.

Missing Data Rates Observed in Applied Research

The 30 LPA studies provided limited information regarding missing data rates, when such information was provided at all. Indeed, only about one-third (11 studies) reported missingness information. Of those reporting indicator missingness information, the reporting range of missing value rates across individual indicators was the most common reporting method (seven of 11 studies), followed by reporting the proportion of observations with complete data (three of 11 studies). A single study reported covariance coverage rates, which is only appropriate in covariance structure models.

The reporting observations with complete cases would have been ideal to inform the simulation design. However, predicting this value from either the range of missing data rates or the covariance coverage rates proved difficult. For example, one typical study included five indicators and reported missingness rates ranging from 10.7% to 25%. The percentage of observations with complete data in this study could range anywhere between 0-75%. In the simulation study, we fixed the missing data rate such that 50% of observations are missing at least one indicator value. With the chosen missing data mechanism (discussed below) this corresponds to marginal missingness rates for each variable of approximately 25%.

Number of Indicators and Classes Observed in Applied Research

The modal number of profiles in the selected LPA articles was found to be $J = 4$ indicators. Correspondingly, four classes were the most frequently reported number of classes in

the 30 selected LPA articles. However, we chose to simulate data from a three-class model ($K = 3$) to make the simulation study time-feasible and to minimize convergence issues encountered with fitting mixture models with more components.

Class Separation Values Observed in Applied Research

The 25th and 75th quantiles of the reported entropy values were .74 and .88, respectively. To achieve entropy values that reflect these typical values, the Mahalanobis distances (MD) between the class means were manipulated so that the entropy values from the simulation roughly matched the 25th and 75th quantile values from the LPA studies, on average across replications. The results from this manipulation exercise are shown in Figure 2. In conducting this manipulation exercise, the unequal mixing condition was chosen because that condition was highly represented in the LPA studies. We found that the corresponding MD values associated with the 25th and 75th entropy quantiles were $MD = 2.84$ and $MD = 3.70$, respectively.

Manipulated Simulation Conditions

We hypothesized that the relative performance of recursive partitioning imputation would be dependent on the sample size, the mixing proportions, and the class separation. We manipulated the sample size because we believe sample size influences the quality of the donor pools. This is because the donor pool represents an empirical distribution that approximates a posterior distribution. As sample size increases, the empirical distribution is expected to converge to the true posterior distribution for Y^{obs} , leading to improved imputations. Thus, we manipulated the sample size between $N = 300$ (small sample) and $N = 1,200$ (large sample), according to the interquartile range that is typically observed in applied research as indicated by the reviewed studies.

We also manipulated the mixing proportions because we anticipated that the quality of the donor pool would degrade in small class sizes. We expected that the donor pool would be a particularly poor reflection of the true posterior predictive distribution for Y^{mis} with unequal mixing proportions and small class sizes. To test this hypothesis, we manipulated the mixing proportions between the condition with equal mixing and a condition with unequal mixing, such that the smallest class represented about 10% of the sample. This value is slightly less than would be predicted when regressing the minimum class proportion on the total number of indicators from the data extracted from the 30 applied studies, but it is certainly consistent with proportions observed in literature.

Finally, we manipulated the separation between the classes. In the weakly and strongly separated condition, all classes were separated $MD = 2.86$ and $MD = 3.70$, respectively, as determined was necessary to reflect the interquartile range of observed entropy values from the literature review ($IQR = .74-.88$). Although we anticipated that recursive partitioning algorithms would outperform imputation procedures that assume a single class and do not embed any clustering in this setting, we note that hard clustering (like that formed by the rectangular stratum in recursive partitioning) is less appropriate when classes are weakly separated. As a result, performance was expected to decline in weakly separated settings because recursive partitioning algorithms do not account for uncertainty in classification.

In addition to the above primary manipulated conditions, we also identified a secondary set of conditions that we hypothesized would influence the amount of nonresponse bias observed in FIML estimates and the degree to which multiple imputation decreases this nonresponse bias. However, the relative performance of recursive partitioning imputation compared to alternatives was not expected to be sensitive to these manipulated conditions. The secondary conditions

include: (1) whether there are mean differences by classes as would occur when the auxiliary variable is a distal outcome, and (2) whether profile membership moderates the correlation between the auxiliary variable and profile indicators. Specifically, in the mean differences condition, the mean value between the reference class and two nonreference classes will be separated by $\delta = 0.5$ SD, while the nonreference classes will be separated by $\delta = 1$ SD. Finally, in the moderation condition, the correlation between the AV and the profile indicators is set to $r = .4$ for the first class, a zero correlation for the second class, and $r = -.4$ for the last class.

We fully crossed the manipulated conditions for a total of 32 simulation conditions. We conducted 500 replications across each condition. For each replication and for each condition, we constructed $M = 100$ imputed datasets.

Data Generating Mechanism and Missing Data Mechanism

We generated complete indicator data and data for a single AV using a three-class Gaussian mixture model as a template (Panel A of Figure 3). Specifically, data were generated from this template by first randomly drawing profile memberships for each observation from a multinomial distribution and then subsequently drawing values for the profile indicators and AV from a multivariate normal distribution with a profile-specific population mean vector and a unitary variance-covariance matrix. Drawing values for the AV jointly with the profile indicators separately across profiles allowed for mean differences across profiles, and it also allowed for different covariance specifications with the indicators across the profiles. Both of these specifications were required for the secondary manipulation conditions.

Profile indicator data that were set to missing in a manner which ensured that the propensity for missingness was informed by the auxiliary variable so that the data are MAR only

if the AV. Observations could be missing up to three indicator values, resulting in 15 possible missing data patterns total (including the pattern corresponding to no missingness).

Implementing the missing data mechanism to assign missing values proceeded in four steps. First, observations were assigned to a missing data pattern using a latent response variable formulation. In particular, observations were first assigned a latent response variable, η_i , by adding random noise to that observation's AV value,

$$\eta_i^* = AV_{1i} + \epsilon_i$$

where $\epsilon_i \sim N(0, 0.1)$. Next, the observations were ranked based on the latent response variable value and observations in the bottom half of the distribution were then identified as missing data for at least one indicator datapoint. Observations were could then be assigned to missing data for one, two, or three indicators. The number of indicators missing data was determined by sampling from a multinomial distribution with three categories, all with equal probability. Finally, the specific indicators among the set of $\{Y_1, Y_2, Y_3, Y_4\}$ that would be set to missing sampled randomly without replacement using the results from the previous step. Under this missing data mechanism, missingness averaging approximately 25% on each indicator while half of the observations contained complete data.

Analytic Model

We fit the standard latent profile model diagrammed in Panel B in Figure 3 to either the observed dataset of profile indicators or to the imputed datasets using Mplus version 8.0 (Muthén & Muthén, 2017). Subsequently, results were exported to R using the `MplusAutomation` package (Hallquist & Wiley, 2018). We evaluated the performance of FIML in estimating parameters from the observed data likelihood function. We also evaluated the performance of a multivariate normal imputation procedure using the `Amelia` package (Honaker, King, &

Blackwell, 2011) in R. The multivariate normal imputation is denoted EMS-MVN because the expectation-maximization with sampling (EMS; King et al., 2001) imputation algorithm is employed by *Amelia*. We assessed the performance of several MICE imputation procedures, including MICE-CART and MICE-RF using the `mice` R package.

Evaluative Criteria

All pooling was conducted using Rubin's (1987) rules, which assume that the posterior distribution in (8) is approximated by a normal distribution, i.e.

$$[\theta | Y^{\text{obs}}, X] \sim \mathcal{N}_{q_K}(\bar{\theta}, \hat{T}), \quad (9)$$

$$\bar{\theta} = \frac{1}{M} \sum_{m=1}^M \hat{\theta}_m \quad (10)$$

$$\hat{T} = \bar{U} + \left(1 + \frac{1}{M}\right) B \quad (11)$$

where q_K is the number of parameters in θ and depends on the number of mixture components K fit to the data. The posterior distribution is centered at the mean of the maximum likelihood estimates $\hat{\theta}_m$ obtained from fitting the mixture model to each $m = 1, \dots, M$ imputed datasets. Finally, $\hat{T}$ is the observed-data posterior variance; it is a composite of the within imputation variance covariance matrix ($\bar{U}$) and the between imputation variance covariance matrix (B), each given by

$$\bar{U} = \sum_{m=1}^M I^{-1}(\hat{\theta}_m | Y_m^{\text{imp}}) \quad (12)$$

$$B = \frac{1}{M-1} (\hat{\theta}_m - \bar{\theta})(\hat{\theta}_m - \bar{\theta})^T \quad (13)$$

where $I^{-1}(\hat{\theta}_m | Y_m^{\text{imp}})$ is the observed information matrix obtained by fitting the model to the m th imputed dataset Y_m^{imp} .

Recovery of μ_k and π_k Parameters (i.e., Bias)

We analyzed the absolute and relative bias of the class-specific mean estimates as well as the marginal class probability estimates. We defined relative bias as the ratio:

$$\text{Relative Bias} = \frac{\text{Bias}(\hat{\theta}^{\text{imp}})}{\text{Bias}(\hat{\theta}^{\text{FIML}})} \quad (14)$$

where the “imp” superscript refers to one of the missing data procedures tested in this simulation study (e.g., EMS-MVN, MICE-CART, or MICE-RF). Relative bias values near zero indicate complete attenuation of the bias, whereas relative bias values greater than one or less than negative one imply that the bias is greater in magnitude for the alternative procedure relative to FIML. Finally, negative values are indicative that the direction of the bias is opposite from that observed in FIML.

Density Estimation

Whereas the proceeding section focused on recovering select statistics of typical interest to applied researchers (e.g., μ_k and π_k) to answer substantive research questions, this section studies the performance of each MAR strategy in recovering the full joint distribution of the true data generating mechanism. We used the Kullback-Leibler (KL) divergence to quantify the degree to which the distribution implied by fitting a mixture model to the observed (FIML) or imputed (EMS-MVN, MICE-CART, etc.) data matched the joint distribution in the population. In this way, we were interested in the KL-divergence to evaluate how well the implied multivariate distribution estimated from one of the missing data strategies approximated the true data generating mechanism. Indeed, one important application of finite mixture models is semiparametric applications where a joint density is approximated (e.g., density estimation). The

KL-divergence is a metric for evaluating how well the implied density from the parameter estimates matches the true population density.

The KL-divergence between the model-estimated distribution and the population distribution is given by:

$$\text{KL}(\theta||\hat{\theta}) = \int_Y \text{Pr}(Y|\theta) \log \frac{\text{Pr}(Y|\theta)}{\text{Pr}(Y|\hat{\theta})} dY \quad (15)$$

where $\text{Pr}(Y|\theta)$ is the density for a $K = 3$ class Gaussian mixture model and $\hat{\theta}$ is the estimate of the sufficient statistics from the sample after employing one of the MAR strategies. We used Monte Carlo integration with 50,000 sampled nodes to approximate the integral in (15).

The KL-divergence is always greater than zero and takes on the value zero if the estimated joint distribution exactly equals the true data generating distribution. However, the value of KL-divergence is not on an interpretable scale. Therefore, we mapped the KL-divergence on an interpretable scale and calculated the percent reduction in the KL-divergence relative to the KL-divergence observed from the FIML parameter estimates. This percent reduction is given by:

$$\text{Percent Reduction KL}(\theta||\hat{\theta}) = 100 \times \left(1 - \frac{\text{KL}(\theta||\hat{\theta}^{\text{imp}})}{\text{KL}(\theta||\hat{\theta}^{\text{FIML}})} \right) \% \quad (16)$$

where $\hat{\theta}^{\text{FIML}}$ is the FIML estimate and $\hat{\theta}^{\text{imp}}$ is the estimate obtained from an imputation procedure. A near 100% reduction in the KL-divergence corresponds to estimates that approximate the true data generating mechanism extremely well and reflects a holistic attenuation of bias across all sufficient statistics estimated by fitting a mixture model to a sample. In contrast, a negative percent reduction corresponds to a set of estimates that, relative to FIML, fails to well-approximate the full joint distribution.

95% CI Coverage Rates

To assess coverage, we took the mean of the individual 95% confidence interval coverage rates across replications. We did not analyze the mean of the coverage rates of the profile variance estimates, as these quantities are generally not of substantive interest to applied researchers. Following Muthén & Muthén (2002), we considered coverage rates falling between .91 to .98 as acceptable because they neither unduly risked invalid inference through narrow intervals nor were they too conservative to be considered needlessly inefficient.

Imputation Efficiency

We analyzed the increase in variance attributable to the missingness. A typical statistic for this quantity is the average relative increase in variance (ARIV), which can be roughly thought of as the mean increase in the standard errors squared due to between-imputation uncertainty. The ARIV is calculated as

$$\text{ARIV} = \frac{\text{tr}(\bar{U}^{-1}\hat{T})}{M} \quad (17)$$

which mathematically translates to the mean of the proportional increase in the eigenvalues observed in the within- vs. total- imputation asymptotic covariance matrices. It is important to note that the ARIV results in an unbiased estimate of the (average) increase in variance attributable to missingness only if $\bar{U}$ is an unbiased estimate of the complete-data asymptotic variance-covariance matrix. Bayesianly proper imputations are sufficient to guarantee that this condition holds (Schafer, 1997), but in simulation studies, the complete-data asymptotic variance-covariance matrix can be obtained directly. Thus, we calculated the average relative increase in variance from this matrix directly using the formula:

$$\text{ARIV} = \frac{\text{tr}(I(\hat{\theta}|Y)\hat{T})}{M} \quad (18)$$

where $I(\hat{\theta}|Y)$ is the observed information matrix using the complete data, and $\hat{T}$ is the covariance matrix of the posterior distribution $[\theta|Y^{\text{obs}}, X]$ approximated by Rubin's (1987) rules during the pooling phase.

We note that imputation efficiency is not to be conflated with traditional notions of sampling variability. Imputation efficiency relates to the variability due to sampling different imputations (e.g., between-imputation variability). Problems with inference from this source of variability can be resolved by increasing the number of imputations in the imputed dataset. In contrast, sampling variability refers to both the within- and between- imputation variability, where the within-variability is what akin to traditional notions of sampling variability and is sensitive to sample size (van Buuren, 2018, p. 43).

Recovery of Posterior Class Probabilities

Finally, we assessed the recovery of the resulting posterior class probabilities estimated using FIML and MI. The estimated posterior class probabilities may be discrepant from the population values using the complete data because of either nonresponse bias or sampling error. Missing data will also lead to discrepancies. Imputation variability, for example, will also result in discrepant posterior class probabilities between the imputed dataset and the complete dataset, even if the true population parameter values were known.

Noting that the posterior class probabilities are themselves sufficient statistics for a multinomial distribution, we again leveraged the KL-divergence to measure the recovery between (a) the probability mass function implied by the complete data with parameter set to their population values (i.e., $\tau_k(\mathbf{y}_i|\theta)$) and (b) the probability mass function implied by the fitting a mixture model to either the observed data (i.e., $\tau_k(\mathbf{y}_{i,m}^{\text{obs}}|\hat{\theta}^{\text{FIML}})$) or the imputed data, (i.e., $\tau_k(\mathbf{y}_{i,m}^{\text{imp}}|\hat{\theta}^{\text{imp}})$). For FIML, we summarized the discrepancies in the probability mass

functions by averaging across observations in the datasets; the mean KL-divergence in the case of FIML is given by

$$\overline{\text{KL}}(\tau||\tau^{\text{FIML}}) = \frac{1}{N} \sum_{i=1}^N \sum_{k=1}^K \tau_k(\mathbf{y}_i|\theta) \log \frac{\tau_k(\mathbf{y}_i|\theta)}{\tau_k(\mathbf{y}_i^{\text{obs}}|\hat{\theta}^{\text{FIML}})}. \quad (19)$$

In the case of MI, we summarized the discrepancies by averaging across both imputations and observations,

$$\overline{\text{KL}}(\tau||\tau^{\text{imp}}) = \frac{1}{M} \sum_{m=1}^M \frac{1}{N} \sum_{i=1}^N \sum_{k=1}^K \tau_k(\mathbf{y}_i|\theta) \log \frac{\tau_k(\mathbf{y}_i|\theta)}{\tau_k(\mathbf{y}_{i,m}^{\text{imp}}|\hat{\theta}_m)}. \quad (20)$$

To place the results on an interpretable scale, we then calculated percent reduction in the mean KL divergence of an imputation method compared to FIML,

$$\text{Percent Reduction } \overline{\text{KL}}(\tau||\tau^{\text{imp}}) = 100 \times \left(1 - \frac{\overline{\text{KL}}(\tau||\tau^{\text{imp}})}{\overline{\text{KL}}(\tau||\tau^{\text{FIML}})} \right) \%. \quad (21)$$

The percent reduction in the mean KL divergence will be positive if the imputation procedure results a decreased discrepancy between estimated and true posterior class probabilities, on average, relative to FIML. On the other hand, if the imputation procedure increased the discrepancy on average relative to FIML, the percent reduction will take on a negative value.

Results

Recovery of μ_k and π_k Parameters

In recovering class means in large samples ($N = 1,200$), recursive partitioning imputation better attenuated nonresponse bias than FIML and multivariate normal imputation (see Panel A of Figure 4). We found superior recovery of all μ_k parameters when classes were set to be equally mixed. These results also mostly held for the unequal mixing conditions, with the bias for the mean of Y_4 in the smallest class (i.e., class $k = 3$) representing the one exception;

nevertheless, the increase in bias relative to FIML and multivariate normal imputation for this parameter was found to be substantively small in size ($\delta = 0.05$ SD). Between the two recursive partitioning methods evaluated, we found that MICE-CART best attenuated nonresponse bias in large samples.

In small samples, recursive partitioning only attenuated nonresponse bias relative to FIML and multivariate normal imputation when the classes were set to be equally mixed (see Panel B of Figure 4). We found the only exception to this rule was the MICE-RF estimates for the class mean of Y_1 in $k = 1$. In contrast, recursive partitioning imputation often failed to attenuate nonresponse bias when the classes were set to the unequal mixing conditions. For example, both MICE-CART and MICE-RF resulted in a sharp increase in bias for the class mean of Y_4 in $k = 3$ relative to FIML. Additionally, both recursive partitioning imputation methods failed to match the bias attenuation exhibited by multivariate normal imputation for class mean of Y_1 in $k = 3$ when the classes were to the unequal mixing condition and the classes were weakly separated. In sum, recursive partitioning imputation only consistently attenuates nonresponse bias in small samples when the classes are equally mixed.

With regard to the marginal class probabilities (i.e., π_k), recursive partitioning imputation also attenuated nonresponse bias for the marginal class probabilities in the large sample ($N = 1,200$) (see Panel A of Figure 5). Two exceptions include π_1 in the equal mixing condition as well as π_3 in the unequal mixing condition. In both cases the bias was within a probability value of .015 of the true parameter values. We consider this discrepancy small in substantive size as it implies that the class membership proportions are biased by only 1.5%. Lastly, compared to multivariate normal imputation, MICE-CART and MICE-RF resulted in less biased estimates in all settings.

The bias attenuation of marginal class probabilities exhibited by recursive partitioning imputation in large samples was not fully replicated when sample sizes were set to the small sample condition ($N = 300$; see Panel B of Figure 5). As with the case of the class means, the performance of recursive partitioning imputation in recovering marginal class probabilities was sensitive to the mixing and separation conditions. Indeed, both MICE-CART and MICE-RF resulted in increased bias in the unequal mixture setting when the classes were weakly separated ($MD = 2.86$, $entropy = .74$). This increase in bias was especially pronounced when the classes were weakly separated.

In conclusion, recursive partitioning imputation performs well in recovering class means and marginal class probabilities when sample sizes are large ($N = 1,200$), and MICE-CART performs as well or better than MICE-RF in attenuating bias in such settings. However, when sample sizes were set to the small condition ($N = 300$), both recursive partitioning method failed to consistently attenuate nonresponse bias across all conditions. In the supplementary materials, we show that these conclusions are insensitive to whether to the secondary simulation conditions, including class mean differences in the AV and whether class membership moderates the relationship between the AV and the profile indicators.

Density Estimation, 95% CI Coverage, and Imputation Efficiency

We report the percent reduction in the KL divergence and coverage rates across the simulation conditions in Table 2. On these metrics, we once again find that recursive partitioning imputation outperforms FIML and multivariate normal imputation in large samples ($N = 1,200$). Both MICE-CART and MICE-RF lead to a full set of parameter estimates that better approximate the true data generating joint distribution and maximize the percent reduction in the $KL(\theta|\hat{\theta})$, on average across replications (EMS-MVN: Range: 14.3-57.5%; MICE-CART:

Range: 71.3-73.6%; MICE-RF: Range: 68.3-72.8%). Finally, of the two recursive partitioning imputation algorithms, we find that MICE-CART led to imputations that better recaptured to true joint distribution and maximized the percent reduction in $KL(\theta|\hat{\theta})$. In the supplemental materials we show that these results are insensitive to manipulation of the secondary simulation conditions.

Similar to our findings regarding the attenuation of nonresponse bias, we found that these positive results failed to translate to the small sample conditions ($N = 300$). In small sample settings, multivariate normal imputation most frequently outperformed recursive partitioning imputation and resulted in a greater percent reduction in the $KL(\theta|\hat{\theta})$, except when classes were equally mixed and strongly separated.

With regards to confidence intervals and imputation efficiency, recursive partitioning also demonstrated adequate performance. Average coverage rates across the parameters fell in the acceptable range of .91 to .98 (Muthén & Muthén, 2002), with the exception of the MICE-RF confidence interval for the class means when classes were equally mixed and sample sizes were large (weakly separated: coverage = .89, strongly separated: coverage = .86). Similarly, when sample sizes were small ARIV values were below or commensurate with those observed for multivariate normal imputation.

Recovery of Posterior Class Probabilities

Figure 6 illustrates the percent reduction in the $\overline{KL}(\tau|\hat{\tau})$, the measure we used to assess improvements in classification quality. In large samples ($N = 1,200$), we found that recursive partitioning improved classification quality compared to FIML and multivariate normal imputation. Notably, the substantive size of the improvement is generally small and on the order of a 10% reduction in $\overline{KL}(\tau|\hat{\tau})$.

At first glance, it may seem surprising that the reduction in $\overline{KL}(\tau|\hat{\tau})$ is small in substantive size when sample sizes are large. We note, however, that the reduction is confounded by two competing sources of variability that both influence how well the posterior class probabilities are recovered. Specifically, the posterior class probabilities from an imputed dataset will only equal the complete data values if (a) the imputed values exactly equal the true missing values (i.e., no imputation error) and (b) the corresponding parameter estimates from the imputed dataset exactly equal the complete data estimates (i.e., no sampling error). Because the imputed dataset are laden with imputation and sampling variability, the posterior class probabilities are not perfectly recoverable. In the asymptotic case that sample sizes are infinitely large, the degree of recovery mainly reflects imputation variability. Thus, the fact that we observed a small decrease in the percent reduction of $\overline{KL}(\tau|\hat{\tau})$ relative to FIML simply indicates that, relative to sampling error, imputation error is a large source of variability.

Discussion

This simulation study suggests that when the data are MAR conditional on AVs, recursive partitioning imputation models (and MICE-CART, in particular) can outperform FIML and single-class imputation alternatives in terms of mitigating nonresponse bias and recovering the joint distribution of the data generating mechanism. In particular, MICE-CART imputation performed as well or better as well as FIML and multivariate normal imputation in large samples ($N = 1,200$) with respect to attenuating nonresponse bias on key parameters of interest (i.e., μ_k and π_k), overall density estimation of the true data generating joint distribution, and the recovery of posterior class probabilities. We found acceptable 95% CI coverage rates and ARIV estimates. Indeed, the ARIV for MICE-CART was slightly larger but commensurate with the relative increase in standard errors of multivariate normal imputation. This implies that users would not

need to greatly increase the number of imputed datasets to achieve optimal power using MICE-CART as compared to traditional imputation approaches. Thus, although these results suggest that researchers may rely on the traditional recommendations of $M = 10$ to 20 datasets just as they would multivariate normal imputation (van Buuren, 2018), we note that computing power has advanced markedly over the decades to make analysis of $M = 50$ to 100 datasets at best a trivial inconvenience.

In small samples ($N = 300$), recursive partitioning imputation methods only demonstrated superior performance when classes were strongly separated and equally mixed. Because it is rather unlikely that applied researchers experience data conditions where classes are both equally mixed and strongly separated in practice, we conclude that neither MICE-CART nor MICE-RF demonstrated adequate performance in small sample settings in treating the missing data.

Why do the recursive partitioning methods perform poorly in small samples is implied when in the unequal mixing condition? This finding can be explained by the degradation of the donor pool used to sample the imputations in these settings. A high-quality donor pool is one whose empirical distribution well approximates the posterior predictive distribution of the missing values, $[Y^{\text{mis}}|Y^{\text{obs}}, X, \theta]$, and is effective at matching missing values with observations from the same latent class. As implemented in the `mice` R package, MICE-CART and MICE-RF create decision trees with stopping criteria determined by a minimum node size of five observations. One would expect that for a fixed node size, a decision tree would include more cut points in large samples than in small samples. Thus, in large samples, observations are more likely to be stratified into clusters so that the donor pools are more homogenous with respect to the joint distribution of the profile indicators and auxiliary variables and, therefore, are clustered together into the same latent class. The improved homogeneity implies that observations with

missing values are more finely matched to observations that are similar with respect to the joint distribution of the profile indicators and AVs. In contrast, decision trees constructed with small observations with a fixed terminating criterion lead to fewer cut points, decreased stratification, and clusters that are less homogenous. The matching process is, therefore, coarser, resulting in an empirical distribution that does not approximate the true posterior predictive distribution of the missing values, $[Y^{\text{mis}}|Y^{\text{obs}}, X, \theta]$, well.

As demonstrated by the simulations, high-quality donor pools for observations in a small class representing only 10% of the population are difficult to construct if the classes are weakly separated so that entropy values average near .74. This is not surprising, given that the number of observations belonging to the small class is only $N = 30$, on average. Compounded by greater overlap between the classes induced by weak separation, one would expect a heterogenous donor pools containing observations belonging to a mix of latent classes so to satisfy the fixed node size requirement.

Improving the quality of the donor pool in LPA likely requires researchers to rely on parametric assumptions to sample from a known distribution rather than an empirical distribution (as is done in forming donor pools). There are several avenues by which this can proceed. First, covariance structure models can be imbedded within a recursive partitioning framework so that, at each terminal node, a covariance and mean structure are estimated. This “SEM Trees” (Brandmaier, von Oertzen, McArdle, & Lindenberger, 2013) approach is currently an area of active research with promising results.

Alternatively, regression mixture imputation models could be developed to directly sample imputations from a latent class model. This is also an area of active research. A Bayesian nonparametric imputation model utilizing Dirichlet processes to model class membership in

order to fit an “infinite” mixture model has recently received attention in the machine learning community (Sovilj et al., 2016). Alternatively, EM with sampling imputation algorithms with finite mixtures have also been developed and utilized for imputation in the behavioral sciences (Vidotto, Vermunt, & Kaptein, 2015). A critical challenge with a finite mixture approach, however, is that the number of subpopulations is not known a priori, and the number of classes supported by the data may be sensitive to nonresponse bias. Regardless of the mixture modeling approach, it is necessary to make these algorithms flexible enough to incorporate many AVs which themselves contain missing data to reflect the multiple data types that appear in the real world.

Finally, we acknowledge that this study only focuses on the challenges associated with the imputation phase and leaves the pooling phase unaddressed. Before multiple imputation can become mainstream or be recommended over FIML, appropriate strategies for pooling model fit information must be identified.

In summary, we have shown that multiple imputation can outperform FIML in the more realistic situation where the MAR assumption is only tenable with the inclusion of AVs. As expected, FIML results in biased parameter estimates in real-world data conditions. Thus, applied researchers should test the MAR assumption before conducting an LPA. This can be done by evaluating if variables available to the researcher in the dataset explain missing data patterns above and beyond that explained by observed profile indicator values. If predictive variables are found, researchers employing FIML to treat missing data with sample sizes above $N = 1,200$ should conduct a sensitivity analysis using MICE-CART imputation. Such a sensitivity analysis can provide valuable evidence regarding whether conclusions about class

definitions and inferences regarding class proportion are sensitive to violations of the MAR assumption.

References

- Allison, P. D. (2002). *Missing data*. Thousand Oaks, CA: SAGE Publications.
- Azur, M. J., Stuart, E. A., Constantine, F., & Leaf, P. J. (2011). Multiple imputation by chained equations: what is it and how does it work? *International Journal of Methods in Psychiatric Research*, 20(1), 40–49. <https://doi.org/10.1002/mp.329>
- Bárcena, M. J., & Tusell, F. (2000). Tree-based algorithms for missing data imputation. *Proceedings in Computational Statistics*, 193–198.
- Brandmaier, A. M., von Oertzen, T., McArdle, J. J., & Lindenberger, U. (2013). Structural equation model trees. *Psychological Methods*, 18(1), 71–86. <https://doi.org/10.1037/a0030001>
- Breiman, L. (1984). *Classification and regression trees*. Belmont, CA, USA: Wadsworth International Group.
- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32. <https://doi.org/10.1023/A:1010933404324>
- Burgette, L. F., & Reiter, J. P. (2010). Multiple imputation for missing data via sequential regression trees. *American Journal of Epidemiology*, 172(9), 1070–1076.
- Collins, L. M., Schafer, J. L., & Kam, C.-M. (2001). A comparison of inclusive and restrictive strategies in modern missing data procedures. *Psychological Methods*, 6(4), 330–351. <https://doi.org/10.1037/1082-989X.6.4.330>
- Conversano, C., & Cappelli, C. (2002). Missing data incremental imputation through tree based methods. In *Compstat* (pp. 455–460). Springer.
- Conversano, C., & Siciliano, R. (2009). Incremental tree-based missing data imputation with lexicographic ordering. *Journal of Classification*, 26(3), 361–379.

<https://doi.org/10.1007/s00357-009-9038-8>

- Creel, D. V., & Krotki, K. (2006). Creating imputation classes using classification tree methodology. *American Statistical Association (Hg.): Proceedings of the Section on Survey Research Methods*, 2884–2887.
- D'Ambrosio, A., Aria, M., & Siciliano, R. (2012). Accurate tree-based missing data imputation and data fusion within the statistical learning paradigm. *Journal of Classification*, 1–32.
- Doove, L. L., van Buuren, S., & Dusseldorp, E. (2014). Recursive partitioning for missing data imputation in the presence of interaction effects. *Computational Statistics & Data Analysis*, 72, 92–104.
- Enders, C. K. (2008). A note on the use of missing auxiliary variables in full information maximum likelihood-based structural equation models. *Structural Equation Modeling*, 15(3), 434–448.
- Enders, C. K. (2010). *Applied missing data analysis*. New York, NY: The Guilford Press.
- Enders, C. K., & Gottschall, A. C. (2011). Multiple imputation strategies for multiple group structural equation models. *Structural Equation Modeling*, 18(1), 35–54.
- Freedman, D. A. (1963). On the asymptotic behavior of Bayes' estimates in the discrete case. *The Annals of Mathematical Statistics*, 34(4), 1386–1403.
- Freedman, D. A. (1965). On the asymptotic behavior of Bayes estimates in the discrete case II. *The Annals of Mathematical Statistics*, 36(2), 454–456.
- Graham, J. W. (2003). Adding missing-data-relevant variables to FIML-based structural equation models. *Structural Equation Modeling*, 10(1), 80–100.
- Hallquist, M. N., & Wiley, J. F. (2018). MplusAutomation: An R package for facilitating large-scale latent variable analyses in Mplus. *Structural Equation Modeling: A Multidisciplinary*

Journal, 25(4), 621–638.

Hancock, G. R., Stapleton, L. M., & Mueller, R. O. (2019). *The reviewer's guide to quantitative methods in the social sciences* (2nd ed.). New York, NY: Routledge.

Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: data mining, inference, and prediction* (2nd ed.). New York, NY: Springer.

Hayes, T. (2017). *Using classification and regression trees (CART) and random forests to address missing data*. [Doctoral dissertation, University of Southern California]. ProQuest Dissertations and Theses. Retrieved from <http://search.proquest.com.ezp-prod1.hul.harvard.edu/docview/2071284125?accountid=11311>

Hayes, T., & McArdle, J. J. (2017). Evaluating the performance of CART-based missing data methods under a missing not at random mechanism. *Multivariate Behavioral Research*, 52(1), 113–114. <https://doi.org/10.1080/00273171.2016.1264287>

Hayes, T., Usami, S., Jacobucci, R., & McArdle, J. J. (2015). Using classification and regression trees (CART) and random forests to analyze attrition: Results from two simulations. *Psychology and Aging*, 30(4), 911–929. <https://doi.org/10.1037/pag0000046>

Honaker, J., King, G., & Blackwell, M. (2011). Amelia II: A program for missing data. *Journal of Statistical Software*, 45(7), 1–47.

Ishwaran, H., Kogalur, U. B., Blackstone, E. H., & Lauer, M. S. (2008). Random survival forests. *The Annals of Applied Statistics*, 841–860.

Jacobucci, R., Grimm, K. J., & McArdle, J. J. (2017). A comparison of methods for uncovering sample heterogeneity: structural equation model trees and finite mixture models. *Structural Equation Modeling: A Multidisciplinary Journal*, 24(2), 270–282. <https://doi.org/10.1080/10705511.2016.1250637>

- King, G., Honaker, J., Joseph, A., & Scheve, K. (2001). Analyzing incomplete political science data: An alternative algorithm for multiple imputation. *American Political Science Review*, 95(1), 49–69.
- Le Cam, L. (1986). *Asymptotic methods in statistical decision theory*. New York, NY: Springer New York.
- Little, R. J., & Rubin, D. B. (2002). *Statistical analysis with missing data* (2nd ed.). Hoboken, N.J.: John Wiley & Sons.
- Little, T. D., Jorgensen, T. D., Lang, K. M., & Moore, E. W. G. (2014). On the joys of missing data. *Journal of Pediatric Psychology*, 39(2), 151–162.
<https://doi.org/10.1093/jpepsy/jst048>
- Masyn, K. E. (2013). Latent class analysis and finite mixture modeling. In T. D. Little (Ed.), *The oxford handbook of quantitative methods* (pp. 551–611). New York, NY: Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780199934898.013.0025>
- McArdle, J. J. (2013). Dealing with longitudinal attrition using logistic regression and decision tree analyses. *Contemporary Issues in Exploratory Data Mining in the Behavioral Sciences*, 282–311.
- McLachlan, G., & Peel, D. (2004). *Finite mixture models*. New York, NY: John Wiley & Sons.
- Meng, X.-L. (1994). Multiple-imputation inferences with uncongenial sources of input. *Statistical Science*, 9(4), 538–558. <https://doi.org/10.1214/ss/1177010269>
- Molenberghs, G., & Kenward, M. (2007). *Missing data in clinical studies* (Vol. 61). New York, NY: John Wiley & Sons.
- Morgan, J. N., & Sonquist, J. A. (1963). Problems in the analysis of survey data, and a proposal. *Journal of the American Statistical Association*, 58(302), 415–434.

<https://doi.org/10.1080/01621459.1963.10500855>

Muthén, L. K., & Muthén, B. O. (2002). How to use a Monte Carlo study to decide on sample size and determine power. Teacher's Corner. *Structural Equation Modeling*, 9(4), 599–620.

https://doi.org/10.1207/S15328007SEM0904_8

Muthén, L. K., & Muthén, B. O. (2017). Mplus User's Guide. Eighth Edition. Los Angeles,

California: Muthén & Muthén. Retrieved from https://www.statmodel.com/html_ug.shtml

R Core Team. (2020). R: A language and environment for statistical computing. Vienna, Austria.

Retrieved from <https://www.r-project.org/>

Rubin, D. B. (1976). Inference and missing data. *Biometrika*, 63(3), 581–592.

Rubin, D. B. (1987). *Multiple imputation for nonresponse in surveys*. New York, NY: Johnson

Wiley & Sons. <https://doi.org/10.1002/9780470316696>

Rubin, D. B. (1996). Multiple imputation after 18+ years. *Journal of the American Statistical*

Association, 91(434), 473–489. <https://doi.org/10.1080/01621459.1996.10476908>

Savalei, V., & Bentler, P. M. (2009). A two-stage approach to missing data: Theory and

application to auxiliary variables. *Structural Equation Modeling: A Multidisciplinary*

Journal, 16(3), 477–497.

Schafer, J. L. (1997). *Analysis of incomplete multivariate data*. Boca Raton, FL: CRC Press.

Schafer, J. L., & Graham, J. W. (2002). Missing data: our view of the state of the art.

Psychological Methods, 7(2), 147.

Shah, A. D., Bartlett, J., Hemingway, H., Nicholas, O., & Hingorani, H. (2014).

CALIBERrfimpute: Imputation in MICE using Random Forest.

Shah, A. D., Bartlett, J. W., Carpenter, J., Nicholas, O., & Hemingway, H. (2014). Comparison

of random forest and parametric imputation models for imputing missing data using MICE:

a CALIBER study. *American Journal of Epidemiology*, 179(6), 764–774.

<https://doi.org/10.1093/aje/kwt312>

Sovilj, D., Eirola, E., Miche, Y., Björk, K.-M., Nian, R., Akusok, A., & Lendasse, A. (2016).

Extreme learning machine for missing data using multiple imputations. *Neurocomputing*, 174(PA), 220–231. <https://doi.org/10.1016/j.neucom.2015.03.108>

Stekhoven, D. J., & Bühlmann, P. (2011). MissForest—non-parametric missing value imputation for mixed-type data. *Bioinformatics*, 28(1), 112–118.

Sterba, S. K. (2016). Cautions on the use of multiple imputation when selecting between latent categorical versus continuous models for psychological constructs. *Journal of Clinical Child & Adolescent Psychology*, 45(2), 167–175.

Strobl, C., Malley, J., & Tutz, G. (2009). An introduction to recursive partitioning: Rationale, application, and characteristics of classification and regression trees, bagging, and random forests. *Psychological Methods*, 14(4), 323–348. <https://doi.org/10.1037/a0016973>

van Buuren, S. (2018). *Flexible imputation of missing data* (2nd ed.). Boca Raton, FL: CRC Press.

van Buuren, S., Brand, J. P. L., Groothuis-Oudshoorn, C. G. M., & Rubin, D. B. (2006). Fully conditional specification in multivariate imputation. *Journal of Statistical Computation and Simulation*, 76(12), 1049–1064.

van Buuren, S., & Groothuis-Oudshoorn, K. (2010). mice: Multivariate imputation by chained equations in R. *Journal of Statistical Software*, 1–68.

Vidotto, D., Vermunt, J. K., & Kaptein, M. C. (2015). Multiple imputation of missing categorical data using latent class models: State of the art. *Psychological Test and Assessment Modeling*, 57(4), 542–576.

Wilkinson, L. (1999). Statistical methods in psychology journals: Guidelines and explanations.

American Psychologist, 54(8), 594.

Yuan, K., & Bentler, P. M. (2000). Three likelihood-based methods for mean and covariance

structure analysis with nonnormal missing data. *Sociological Methodology*, 30(1), 165–200.

Tables

Table 1

Descriptive Statistics of Study Characteristics in Reviewed Articles

	# of Studies	M	SD	Min	.25 Q-tile	Median	.75 Q-tile	Max
Sample size, N	30	932.5	886.8	137	292	640	1167	4417
# of indicators, J	30	5.2	2.44	2	4	5	6	12
# of components, K	30	3.9	1.37	2	3	4	4	8
Entropy	27	.82	.09	.63	.74	.85	.88	.98
Min. class prop., π_{min}	29	.11	0.1	.02	.06	.09	.13	.47

Table 2
Density Recovery, Imputation Efficiency, and 95% CI Coverage

	Small Sample ($N = 300$)				Large Sample ($N = 1,200$)			
	% Reduc. KL	ARIV	95% CI Cover. ^a		% Reduc. KL	ARIV	95% CI Cover. ^a	
			μ_k	π_k			μ_k	π_k
Unequal Mixing & Weakly Separated								
Complete Data	100.0	-	.87	.78	100.0	-	.94	.94
FIML	0.0	-	.76	.70	0.0	-	.68	.85
EMS-MVN	79.9	6.44	.96	.98	57.5	0.65	.88	.90
MICE-CART	36.0	6.71	.95	.93	73.3	0.99	.94	.96
MICE-RF	51.9	6.42	.96	.94	72.8	0.69	.92	.96
Unequal Mixing & Strongly Separated								
Complete Data	100	-	.93	.96	100	-	.95	.99
FIML	0.0	-	.82	.92	0.0	-	.61	.88
EMS-MVN	57.0	1.4	.95	.99	14.3	0.70	.83	.74
MICE-CART	50.9	1.60	.94	.98	73.6	0.73	.93	.96
MICE-RF	56.3	1.54	.95	.98	68.3	0.78	.90	.95
Equal Mixing & Weakly Separated								
Complete Data	100	-	.92	.93	100	-	.95	.99
FIML	0.0	-	.82	.87	0.0	-	.64	.82
EMS-MVN	74.6	1.09	.96	.98	42.0	0.60	.82	.81
MICE-CART	44.7	1.75	.95	.98	71.3	0.63	.93	.95
MICE-RF	57.5	1.47	.95	.98	68.9	0.52	.89	.91
Equal Mixing & Strongly Separated								
Complete Data	100	-	.94	1.00	100	-	.94	1.00
FIML	0	-	.82	.95	0	-	.54	.78
EMS-MVN	36.02	0.88	.93	.93	-8.95	0.77	.74	.70
MICE-CART	46.09	0.71	.93	.97	73.93	0.63	.91	.92
MICE-RF	48.11	0.78	.93	.97	66.38	0.59	.86	.87

Notes. $\pi_3 = .10$ in unequal mixing condition. Classes separated by MD = 2.86 (entropy $\approx .74$) under weak

separation and MD = 3.70 (entropy $\approx .88$) under strong separation.

^a Averaged across classes (π_k) and indicator variables (μ_k).

Figures

Figure 1

Example of Nodes, Rectangular Partitions, and Donor Pools from CART

Notes. Minimum node size = 200 observations. Grey = Class 1. Black = Class 2. (A) Decision tree with terminal nodes in regressing Y_3 on Y_1 and Y_2 . (B) Corresponding rectangular partitions. (C) Distribution of observed Y_3 values in each terminal node which can be used to sample imputations from a donor pool.

Figure 2

Manipulating Class Separation to Achieve Desired Entropy Values

Notes. $K = 3$ classes. $J = 4$ indicators. Classes separated by the same MD.

Figure 3

Template and Analytic LPA Models

Notes. Latent profile models reflecting (A) the data generating mechanism for simulating complete data, and (B) the analytic models fit to either the observed or the imputed data.

Figure 4

Parameter Recovery: Class Indicator Means, μ_k

Figure 5

Parameter Recovery: Marginal Class Probabilities, π_k

Notes. Tail of arrow represents absolute bias of FIML estimates. Head of arrow corresponds to absolute bias of the imputation strategy.

Figure 6

Large Sample ($N = 1,200$) Classification Quality: Percent Reduction in $\overline{KL}(\tau|\hat{\tau})$

Notes. These are the notes for figure 6.

Appendix

To highlight the difficulty in implementing a saturated correlates approach in a latent profile analysis, we utilized data from the Early Childhood Longitudinal Study, Kindergarten (ECLS-K) 1998 cohort. With empirical data we show that implementing Graham's rules (2003) results in individuals switching classes and it leads to undesired changes in class definitions. In other words, while Graham's rules (2003) preserves the interpretation of parameters in regression models, structural equation models, or other covariance structure models, these rules do not preserve the interpretation of the parameters that define the classes in latent profile analysis.

The ECLS-K is a nationally representative, longitudinal study following children enrolled in either full-time or part-time Kindergarten in the 1998-1999 school year to the end of eighth grade. The dataset contains a wide range of achievement, behavioral, psychological, and school-environment outcomes being collected up to two times a school year. We build on an introduction to LPA for applied researchers provided by Berlin, Williams, & Parra (2014)⁴ investigating distinct subpopulation of Black, non-Hispanic adolescents in eighth grade who differ according to dietary intake, physical activity, and sedentary behaviors.

Measures and Inclusion Criteria

The constructs of interest include weekly physical activity, sedentary behaviors, and dietary intake. Three items measured physical activities, including participation in school sports

⁴ Berlin, K. S., Williams, N. A., & Parra, G. R. (2014). An introduction to latent variable mixture modeling (part 1): Overview and cross-sectional latent class and latent profile analyses. *Journal of Pediatric Psychology*, 39(2), 174–187. Retrieved from <http://dx.doi.org/10.1093/jpepsy/jst084>

(three categories), participation in non-school sports (four categories), days exercised in the past seven days (0-7 days; 8 categories), and average days in PE per week (0-5 days; 6 categories). Six items measured sedentary behaviors by asking the number of hours per day (0-24 hours; 25 categories) spent watching TV, playing videogames, or using the internet. Nine items measured dietary intake by asking the number of days (0-7 days; 8 categories) specified foods (e.g., carrots, potatoes, fruit, fast food, etc.) or drinks (e.g., a glass of milk, a glass of juice, drank soda, etc.) were consumed.

Because LPA relies on the assumption that indicators are continuous, we parceled items to construct three profile indicators that measure overall physical activity, sedentary behaviors, and dietary intake. The ACTIVITY parcel was constructed by standardizing all items before aggregating. The SEDENTARY and DIETARY parcels were constructed using the items in their raw scales because all items were on the same scale (i.e., hours/day or days/week). The auxiliary variable used in this example is the student's predicted BMI in eighth grade. This value was predicted by fitting a latent growth curve model to BMI data collected in seven waves between Kindergarten and eighth grade.

Only data for individuals who identify as Black, non-Hispanic in eighth grade are analyzed in the illustrative example. We required that all class indicator data be present in order to be included in the study. Requiring complete data implies that any difference in the parameter estimates across the specified models is the result of one specification changing the parameter's interpretation, and not differences because missing data are being treated. Applying these inclusion criteria results in a sample size of $N = 608$ girls and boys.

Analytic Strategy

To identify if Graham's (2003) rules preserves the interpretation of the parameters of the latent classes, we fit two separate mixture models. We specify the first mixture model according to the regular specification. The second mixture model is specified according to Graham's (2003) rules. The total number of classes specified was $K = 3$, which was determined by minimizing the BIC.

If Graham's (2003) rules preserves the interpretation of the parameters that define the classes, then the class-specific means should be equivalent across the two fitted models. Similarly, if the two models result in equivalent classifications into the latent classes, then this should be reflected by the marginal class probabilities being equivalent.

Appendix Figure 1

Regular Specification

Mplus Code:

```
%OVERALL%
[ c#1*2.62697 ];
[ c#2*1.96629 ];
```

Appendix Figure 2

Graham's (2003) Rules Specification

Mplus Code:

```
%OVERALL%
[ c#1*2.01049 ];
[ c#2*1.56094 ];
```

```

%C#1%
activity WITH sedent*0.02210;
activity WITH diet*0.05300;
sedent WITH diet*-0.05011;

[ activity*-0.02476 ];
[ sedent*2.65952 ];
[ diet*2.26605 ];

activity*0.40506;
sedent*1.21632;
diet*0.26683;

%C#2%
activity WITH sedent*0.28524;
activity WITH diet*0.05147;
sedent WITH diet*-1.00020;

[ activity*0.11883 ];
[ sedent*5.16571 ];
[ diet*3.03186 ];

activity*0.33318;
sedent*5.12648;
diet*0.85182;

%C#3%
activity WITH sedent*1.82689;
activity WITH diet*0.03423;
sedent WITH diet*-0.14302;

[ activity*-0.50638 ];
[ sedent*10.96636 ];
[ diet*3.25351 ];

activity*0.21963;
sedent*28.22388;
diet*1.09756;

%C#1%
activity WITH sedent*0.02103;
activity WITH diet*0.05051;
activity WITH bmi7hat*-0.43547;
sedent WITH diet*-0.04691;
sedent WITH bmi7hat*0.12228;
diet WITH bmi7hat*-0.05482;

[ activity*-0.01788 ];
[ sedent*2.58428 ];
[ diet*2.22146 ];
[ bmi7hat*23.64010 ] (c1);

activity*0.40754;
sedent*1.07994;
diet*0.24012;
bmi7hat*26.41688;

%C#2%
activity WITH sedent*0.22583;
activity WITH diet*0.05782;
activity WITH bmi7hat*-0.09564;
sedent WITH diet*-0.83204;
sedent WITH bmi7hat*-1.07637;
diet WITH bmi7hat*-0.01836;

[ activity*0.07593 ];
[ sedent*4.89019 ];
[ diet*2.90937 ];
[ bmi7hat*23.64010 ] (c1);

activity*0.36193;
sedent*5.17748;
diet*0.66108;
bmi7hat*15.65743;

%C#3%
activity WITH sedent*0.03752;
activity WITH diet*0.05405;
activity WITH bmi7hat*1.35984;
sedent WITH diet*-0.05141;
sedent WITH bmi7hat*-41.36480;
diet WITH bmi7hat*-9.50095;

[ activity*-0.23295 ];
[ sedent*9.79538 ];
[ diet*3.80696 ];
[ bmi7hat*23.64010 ] (c1);

activity*0.32718;
sedent*32.94598;
diet*1.51608;
bmi7hat*188.48395;

```

Results

The marginal class probabilities and standardized class means across the indicator parcels for both specifications is shown in the figure below. We find that specifying according to Graham's (2003) rules does not result in equivalent parameter estimates in the class means,

suggesting that such a specification risks changing the definitions of the classes. Moreover, the two specifications lead to differences in the marginal class probabilities. For example, the third class was found to comprise only 4.6% of the population according to the regular specification, but the corresponding class comprised 7.6% of the population when the mixture model is specified according to Graham's (2003) rules. Thus, Graham's (2003) rules neither preserves the parameters that define the classes nor results in observations being classified into the same classes.

Appendix Figure 3

Profile Plots and Marginal Class Probabilities