



A Bayesian semiparametric approach for inference on the population partly conditional mean from longitudinal data with dropout

MARIA JOSEFSSON*

Department of Statistics, USBE, Umeå University, Sweden and Centre for Demographic & Ageing Research, Umeå University, Sweden

MICHAEL J. DANIELS

Department of Statistics, University of Florida, USA

SARA PUDAS

Department of Integrative Medical Biology, Umeå University, Sweden and Center for Functional Brain Imaging, Umeå University, Sweden

maria.josefsson@umu.se

SUMMARY

Studies of memory trajectories using longitudinal data often result in highly nonrepresentative samples due to selective study enrollment and attrition. An additional bias comes from practice effects that result in improved or maintained performance due to familiarity with test content or context. These challenges may bias study findings and severely distort the ability to generalize to the target population. In this study, we propose an approach for estimating the finite population mean of a longitudinal outcome conditioning on being alive at a specific time point. We develop a flexible Bayesian semiparametric predictive estimator for population inference when longitudinal auxiliary information is known for the target population. We evaluate the sensitivity of the results to untestable assumptions and further compare our approach to other methods used for population inference in a simulation study. The proposed approach is motivated by 15-year longitudinal data from the Betula longitudinal cohort study. We apply our approach to estimate lifespan trajectories in episodic memory, with the aim to generalize findings to a target population.

Keywords: BART; Memory; MNAR; Nonignorable dropout; Population inference; Sensitivity analysis; Truncation by death.

1. INTRODUCTION

Studies of lifespan trajectories in memory using longitudinal data present numerous methodological challenges including highly nonrepresentative samples, due to selective study enrollment and attrition, and practice effects, which results in improved or maintained performance due to familiarity with the test

*To whom correspondence should be addressed.

(Weuve and others, 2015). These challenges may bias study findings and severely distort the ability to generalize to the target population, even in well-designed studies.

Improvements in test scores upon repeated assessment due to practice effects (PEs) are well documented in the cognitive aging literature (e.g., Salthouse, 2010). PEs arise from increased familiarity with the assessment tools and may result in underestimating memory decline in longitudinal studies. Common statistical approaches for handling PEs include: ignoring, specifying an indicator for the first assessment, or modeling a linear trend (Vivot and others, 2016). These approaches however remain controversial due to, for example, the strong assumption that the practice gains are the same across the age range or the difficulties in separating effects of within-person change from PEs (Weuve and others, 2015; Hoffman and others, 2011).

Attrition is inevitable especially if individuals in the studied population are followed over a long time period. Standard methods often rely on the assumption of missing at random (MAR) which is invalid if the missingness is missing not at random (MNAR) or due to death. MNAR missingness can be dealt with in any analysis of incomplete data by exploring sensitivity to deviations from the MAR assumption. This can be accomplished in a Bayesian setting, by introducing sensitivity parameters that incorporate prior beliefs about the differences between respondents and nonrespondents and assigning an appropriate prior distribution (Daniels and Hogan, 2008). However, attrition due to death must be treated differently if a participant's death during follow-up truncates the outcome process (Kurland and others, 2009). For example, in research on cognitive aging, cognition after death is not defined and not of interest to study. Joint models with additional postprocessing, that truncate participants outcomes after death, have been proposed to address the combination of dropout and death (Rizopoulos, 2012). Two approaches for this type of mortal cohort inference are partly conditional inference, where focus is on the subpopulation who are still alive at that time point (Kurland, 2005), and principal stratification, where focus is on the subpopulation that would survive irrespective of exposure (Frangakis and Rubin, 2002; Frangakis and others, 2007). The latter approach has been proposed to handle truncation by death when interest is in estimating the survivors average causal effect (e.g., Josefsson and others, 2016; McGuinness and others, 2019; Shardell and Ferrucci, 2018). In contrast, partly conditional inference has been proposed for estimating noncausal associations. For example, by using augmented inverse probability weighting for longitudinal cohort data with truncation by death and MAR dropout (Wen and others, 2018) as well as MNAR dropout (Wen and Seaman, 2018). In addition, Li and Su (2018) proposed an approach for semicompeting risk data with MNAR missingness and death. This type of inference can for example be of interest, when the studied outcome is related to the health of all individuals who are alive. For example, in studies of health promotion and/or disease prevention among the elderly. These methods, however, fail to account for the complex nature of the sampling design such as selective study enrollment. While death is not viewed as a source of bias in longitudinal studies, because the objective is to generalize to those who are alive, selection and nonresponse bias (and additionally practice effects) often are, and therefore need to be addressed. Post-stratification is an approach to adjust for discrepancies between the sample and the population by using auxiliary information on the finite population. A classical weighting estimator for finite population inference is the Horvitz–Thompson (HT) estimator (Horvitz and Thompson, 1952). Although the HT estimator is design unbiased, it can potentially be very inefficient. In contrast, a model-based (MB) approach specifies a model for the study outcome, usually a regression model, which is then used to make predictions for the population, and hence, finite population quantities. Predictions are calculated by plugging in the auxiliary variables for all units in the population in the working model. Model-based approaches will generally perform equally good or better as weighting approaches if the model is correctly specified (Little, 2004). However, the correct model specification can be difficult when there is a large set of regressors, the relationship is nonlinear and/or includes interaction terms, and there are multiple observation times. Several approaches for post-stratification adjustment in a cross-sectional setting have shown improved performance compared to the HT and the MB regression estimator. Multilevel

regression combined with post-stratification (MRP; e.g., Gelman and Little, 1997, Park and others, 2004) have shown improved performance compared to standard weighting and MB regression estimators and was used for producing accurate population estimates from a nonrepresentative sample (Wang and others, 2015). The general regression estimator (GREG; Deville and Särndal, 1992) is a dual-modeling strategy that combines prediction and weighting. The approach requires both a model for the outcome and the participation mechanism and is double robust in the sense that it remains consistent if either one of the models is correctly specified. Bisbee (2019) combined a semiparametric machine learning approach [Bayesian Additive Regression Trees (BART); Chipman and others, 2010] with post-stratification for predicting opinions using cross-sectional data. Kern and others (2016) showed in a simulation study that BART and inverse probability weighting using random forests performed better than approximately doubly robust estimators for estimating the target population average treatment effect. Although showing improved performance compared to MRP and GREG, previous semiparametric approaches for population inference are *not valid for longitudinal data with dropout and deaths*.

In this study, we propose an approach for estimating the finite population mean of a longitudinal continuous outcome conditioning on being alive at a specific time point, i.e., the population partly conditional mean (PPCM). Specifically, we develop a flexible Bayesian semiparametric predictive estimator, when longitudinal auxiliary information is known for all units in the target population. The approach is to specify observed data models using BART and then to use assumptions with embedded sensitivity parameters to identify and estimate the PPCM. We evaluate the sensitivity of the results to untestable assumptions on MNAR dropout and PEs and further compare our approach to other methods used for population inference in a simulation study.

We are motivated by the Betula study, a prospective cohort study on memory, health, and aging. The aim of the current article is to extend previous results on cognitive lifespan trajectories (e.g., Rönnlund and others, 2005, Gorbach and others, 2017) by considering population partly conditional inference with MNAR missingness and practice effects. By using longitudinal micro-data from Statistics Sweden and the National Board of Health and Welfare, for both the sample and the target population, we are able to adjust for potential discrepancies in auxiliary variables and thereby improve the generalizability of study findings.

The remainder of the article is organized as follows. In Section 2, we present a motivating example. In Section 3, we present an MB approach for estimating the PPCM using longitudinal data with dropout and deaths, and in Section 4, we describe a Bayesian semiparametric modeling approach. In Section 5, we provide results from a simulation study and in Section 6 results from the empirical example using the Betula data. Conclusions are given in Section 7.

2. MOTIVATING EXAMPLE

The aim of the empirical study is to estimate lifespan trajectories in episodic memory, with the goal to generalize findings to a target population. Two separate sources of data are available for this study; a longitudinal cohort study and a longitudinal database covering the target population.

The Betula study is a population-based cohort study with the objective to study how memory functions change over time and to identify risk factors for dementia (Nilsson and others, 1997). The participants were randomly recruited, stratified by age, from the population registry in the Ume municipality of Sweden. We consider longitudinal data from the first sample (S1) and four waves of data collection (T1–T4). There were 5 years in between each wave and the first wave of data collection was initiated in 1988–1990. A total of $n = 1000$ participants were included, 100 participants from each of the 10 age-cohorts: 35, 40, . . . , 80. In order to obtain a total of 100 subjects in each of the 10 different cohorts, 1976 persons had to be contacted. Of the 976 that never entered the study, 259 could not be reached, 130 had an illness (including dementia) to the extent that they could not participate, and 481 declined to participate (Nilsson and others,

1997). Memory was assessed at each wave using a composite of five episodic memory tasks, range: 0–76, where a higher score indicates better memory (Josefsson and others, 2012).

The second data source is the Linneaus Database (Malmberg and others, 2010), a longitudinal database covering every Swedish resident. The database includes annual data from Statistics Sweden and the National Board of Health and Welfare (similar information as for the Betula sample). In this study, we consider microdata for every unit of the population in the Ume municipality who were alive and nondemented in 1990 as the target population ($N = 9203$). Although longitudinal data is available annually, we restrict data to the years: 1990, 1995, 2000, and 2005, approximately corresponding to the years of testing in the Betula study.

A set of continuous and categorical auxiliary variables, linked to both selective study enrolment and memory, are included in both data sources. From the cause of death register, we know the death year for each deceased individual. In the Betula sample, 29.1% died during the study period and 20.0% of the target population, and there were 128 and 806 dementia cases (12.8% and 8.8%) in the sample and population, respectively.

3. A MODEL-BASED APPROACH FOR POPULATION INFERENCE

3.1. Cross-sectional setting

Consider a finite population $U = \{1, 2, \dots, N\}$. For each individual $i \in U$, a $[1 \times K]$ vector of auxiliary information \mathbf{x}_i is observed. A probability sample c of size n is drawn from U . Let y_i be the continuous study variable which is observed if $i \in c$, and $p(y_i | \cdot)$ refers to a probability density function. Suppose interest is in studying the finite population mean, $\mu_U = N^{-1} \sum_{i \in U} y_i$. If we assume the participation mechanism is ignorable in the sense that y_i and $i \in c$ are conditionally independent given \mathbf{x}_i (Rubin, 1976), the joint distribution of the outcome, participation mechanism, and auxiliary information can be factored into three conditional distributions $\prod_{i \in U} p(y_i | \mathbf{x}_i) p(i \in c | \mathbf{x}_i) p(\mathbf{x}_i)$.

A model-based (MB) prediction estimator incorporates the relationship between \mathbf{x} and y into the estimation of the population mean by $m(\mathbf{x}_i) = E(Y_i | X_i = x_i)$ and considers the finite population values $\{y_i; i \in U\}$ to be realizations of the model. Conditional on the auxiliary variables \mathbf{x}_i , let $p(y_i | \mathbf{x}_i)$ be a model of the form $y_i = m(\mathbf{x}_i) + \varepsilon_i$, where ε_i are uncorrelated and $E(\varepsilon_i) = 0$ and $\text{Var}(\varepsilon_i) = \sigma^2$. The predictions for each individual in the population, \hat{y}_i , are calculated by plugging in their auxiliary information \mathbf{x}_i for all $i \in U$ in the working model $m(\mathbf{x}_i)$. A MB estimator of the population mean is obtained by $\hat{\mu}^{MB} = \frac{1}{N} \sum_{i \in U} \hat{m}(\mathbf{x}_i)$, where $\hat{m}(\mathbf{x}_i)$ is the estimated mean function. We consider a setting with two separate data sources. Thus, we can not separate out the sample participants in the data covering the target population. Predictions must therefore be made for all participants in the target population.

3.2. Longitudinal data with dropout and death

We now consider the problem of finite population inference in the context of longitudinal data with dropout and death. Death must be treated differently than nonresponse since postdeath outcomes are truncated (and do not exist). Here, we consider partly conditional inference and, as such, are interested in estimating the finite population mean given survival up to that time point, that is, the population partly conditional mean (PPCM).

First we need some additional notation. For individual i at time point $t = 0, 1, \dots, T$, denote the outcome variable and the vector of auxiliary information by y_{it} and \mathbf{x}_{it} , respectively. The history of the time-varying variables are denoted with an overbar. For example, the outcome history for individual i up to and including time point t is denoted by $\bar{y}_{it} = \{y_{i0}, y_{i1}, \dots, y_{it}\}$. Let s_{it} denote survival, where $s_{it} = 1$ if an individual is alive at time t and 0 otherwise. Let r_{it} be a response indicator, where $r_{it} = 1$ if individual i participates in the study at time t and 0 otherwise. We assume monotone missingness, so if $r_{it} = 0$, $r_{ik} = 0$ for $k > t$,

and of course similarly for s_{it} . Note that, s_{it} is observed for all $i \in U$ and r_{it} is observed if $i \in c$. Note that, the number of individuals in the population decreases over time due to deaths.

Initially, we further assume the nonresponse mechanism to be missing at random conditional on being alive at time t (MARS). That is, $p(y_{it} | \bar{y}_{it-1}, \bar{r}_{it-1}, r_{it} = 0, s_{it} = 1, \bar{\mathbf{x}}_{it}) = p(y_{it} | \bar{y}_{it-1}, \bar{r}_{it} = 1, s_{it} = 1, \bar{\mathbf{x}}_{it})$. However, unlike previous works using MAR conditional on being alive (e.g., [Wen and others, 2018](#)), here, time of death is known for all individuals in the population; hence, it does not need to be modeled.

The working model for making predictions under MARS becomes

$$\hat{y}_{it} = \int_{\bar{y}_{it-1}} \hat{m}_t(\bar{y}_{it-1}, s_{it} = 1, \bar{\mathbf{x}}_{it}) \times \prod_{k=0}^{t-1} p(y_{ik} | y_{ik-1}, s_{ik} = 1, \bar{\mathbf{x}}_{ik}) d\bar{y}_{it-1}. \quad (3.1)$$

Monte Carlo integration is used to integrate over the outcome history \bar{y}_{t-1} . To do this we factor the joint distribution of \bar{y}_{t-1} into a sequence of 1D conditional distributions. Population outcome data is then sequentially sampled from each of these, for all $i \in U$. Finally, predictions for \hat{y}_{it} are obtained by plugging in \bar{y}_{it-1} and $\bar{\mathbf{x}}_{it}$ given $s_{it} = 1$ in the estimated model $\hat{m}_t(\bar{y}_{it-1}, s_{it} = 1, \bar{\mathbf{x}}_{it})$.

The MB estimator of the PPCM at time t is given by $\text{PPCM}_t^{MB} = \frac{1}{\sum_{i \in U} s_{it}} \sum_{i \in U} \hat{y}_{it} s_{it}$. Note that, in context of the Betula study interest is not the PPCM_t at a specific test wave, but rather the age-specific PPCM aggregated over test waves. That is,

$$\text{PPCM}^{MB}(\text{age}) = \frac{\sum_{i \in U_{\text{age}_t}} s_{it}}{\sum_{t=1}^T \sum_{i \in U_{\text{age}_t}} s_{it}} \sum_{t=1}^T \text{PPCM}_t^{MB}(\text{age}),$$

where $i \in U_{\text{age}_t}$ is the subset of individuals in age-cohort age at test wave t , and $\text{PPCM}_t^{MB}(\text{age}) = \frac{1}{\sum_{i \in U_{\text{age}_t}} s_{it}} \sum_{i \in U_{\text{age}_t}} \hat{y}_{it} s_{it}$.

3.3. Nonignorable dropout among survivors and practice effects

We introduce a set of sensitivity parameters to assess the impact of violations to the MARS assumption for the missingness mechanism. We additionally introduce a sensitivity parameter that accounts for practice effects (PEs). The general strategy is to model the observed data distribution and to use priors on the sensitivity parameters to identify the full-data model.

Previous studies of the Betula data suggests that individuals who drop out have lower performance and steeper decline in memory ([Josefsson and others, 2012](#)). Thus, we expect dropout to be missing not at random conditioning on survival (MNARS). To allow for deviations from the MARS assumption, we introduce a parameter γ_{it} to identify the expected outcome for dropouts among survivors. That is, for all $t > 0$, $E(Y_{it} | \bar{y}_{it-1}, \bar{r}_{it-1}, r_{it} = 0, s_{it} = 1, \bar{\mathbf{x}}_{it}) = E(Y_{it} - \gamma_{it} | \bar{y}_{it-1}, \bar{r}_{it} = 1, s_{it} = 1, \bar{\mathbf{x}}_{it})$. Note that, after we condition on \bar{r}_{it-1} we implicitly assume y_{it} is conditionally independent of the wave at which the drop-out occurred. If $\gamma_{it} = 0$ this implies a MARS assumption on the expectation for the outcome, and if $\gamma_{it} > 0$ this implies a negative location shift in the outcome at the unobserved test wave.

Practice effects (PEs) arise from increased familiarity with the assessment tools and may result in underestimating the decline in memory in longitudinal studies. Let y_{it}^* denote the observed memory score (in contrast to y_{it} which denotes an individual's actual memory function) and δ_{it} denote a sensitivity parameter. Then, for $t > 0$, $E(Y_{it} | \bar{y}_{it-1}, \bar{r}_{it} = 1, s_{it} = 1, \bar{\mathbf{x}}_{it}) = E(Y_{it}^* - \delta_{it} | \bar{y}_{it-1}, \bar{r}_{it} = 1, s_{it} = 1, \bar{\mathbf{x}}_{it})$, where $\delta_{it} > 0$ implies an overestimated memory performance due to practice effects. Note that we assume no PEs at the initial testing.

Our approach allows to explore sensitivity to the unverifiable assumptions by specifying informative priors for the sensitivity parameters $\gamma_{i1}, \dots, \gamma_{iT}$, and $\delta_{i1}, \dots, \delta_{iT}$. We specify triangular distributed priors

conditioning on auxiliary variables, $\gamma_{it} \sim \text{Tri}(A_t(\bar{\mathbf{x}}_{it}), 0, A_t(\bar{\mathbf{x}}_{it}))$ and $\delta_{it} \sim \text{Tri}(0, B_t(\bar{\mathbf{x}}_{it}), B_t(\bar{\mathbf{x}}_{it}))$, where the three parameters of the Triangular distribution are the minimum, the mode and the maximum. We restrict the parameters of the Triangular distribution to a plausible range of values, reflecting the analysts' beliefs. Note that, using this distribution we are able to place more prior weight at one of the endpoints while still only needing to specify a lower- or an upper bound (and not an explicit variance parameter). In the Analysis of the Betula data, we specify values of $A_t(\bar{\mathbf{x}}_{it})$ and $B_t(\bar{\mathbf{x}}_{it})$ in context of the study.

With the sensitivity parameters, identification of the PPCM based on the working model in (3.1), uses

$$\begin{aligned} \hat{y}_{it} = & \int_{\bar{y}_{t-1}^*} \sum_{\bar{r}_{it}} \left\{ [m_t(\bar{y}_{it-1}^*, \bar{r}_{it}, s_{it} = 1, \bar{\mathbf{x}}_{it}) - \gamma_{it} I_{(r_{it}=0)} - \delta_{it}] \times \right. \\ & \prod_{k=0}^{t-1} p(y_{ik}^* - \gamma_{ik} I_{(r_{ik}=0)} \mid y_{ik-1}^*, \bar{r}_{ik}, s_{ik} = 1, \bar{\mathbf{x}}_{ik}) \times \\ & \left. p(r_{ik} \mid \bar{y}_{ik-1}^*, \bar{r}_{ik-1}, \bar{\mathbf{x}}_{ik}, s_{ik} = 1) \right\} d\bar{y}_{it-1}^*. \end{aligned} \quad (3.2)$$

An implication of (3.2) is that Monte Carlo integration is implemented over both the outcome history, \bar{y}_{t-1} , and the response history, \bar{r}_t . The sensitivity parameters, γ_t and δ_t , must also be sampled for all $i \in U$. If $r_{ik} = 0$, \hat{y}_{ik} is shifted downwards by γ_{ik} for $k = 1, \dots, t$. Predictions for \hat{y}_{it} are similarly obtained by plugging in \bar{y}_{it-1} and $\bar{\mathbf{x}}_{it}$ in the estimated model given $s_{it} = 1$. Note that, for $t \geq 1$, \hat{y}_{it} is further shifted downwards by δ_{it} , thereby adjusting for PEs.

4. A SEMIPARAMETRIC APPROACH FOR ESTIMATING THE PPCM

We propose a Bayesian semiparametric modeling approach based on Bayesian Additive Regression Trees (BART; Chipman and others, 2010) for the working model in (3.2) using the observed data and the sensitivity parameters.

4.1. Semiparametric estimation of the outcomes and dropout

We specify BART models for the conditional distributions of the time-varying variables y_{it} and r_{it} . The distribution of the continuous outcome, $p(y_{it} \mid \bar{y}_{it-1}, s_{it})$ is specified as normal, $y_{it} \sim N(m_t(\bar{y}_{it-1}, \bar{\mathbf{x}}_{it}), \sigma_t^2)$, for the subset that satisfies $\bar{r}_{it} = 1$ and $s_{it} = 1$. The mean function at wave t , $m_t(\bar{y}_{it-1}, \bar{\mathbf{x}}_{it})$, is given by the sum of $k = 1, \dots, K_{y_t}$ regression trees, denoted by $g_{y_t}^k(\bar{y}_{it-1}, \bar{\mathbf{x}}_{it}; T_{y_t}^k, M_{y_t}^k)$. Each regression tree consists of a set of decision rules (a tree structure), denoted by $T_{y_t}^k$, leading down to b^k terminal node parameters, denoted by $M_{y_t}^k = (\rho_{y_t}^k(1), \dots, \rho_{y_t}^k(b^k))$. Hence, an individual's values for $(\bar{y}_{it-1}, \bar{\mathbf{x}}_{it})$ is linked to a single terminal node by following the decision rules for each tree and is assigned the associated terminal node parameter. Similarly, the BART models for the binary response indicators r_{it} are specified as probit models, $\pi_{it}(\bar{y}_{it-1}, \bar{\mathbf{x}}_{it}) = \Phi\left(\sum_{k=1}^{K_{r_t}} g_{r_t}^k(\bar{y}_{it-1}, \bar{\mathbf{x}}_{it}; T_{r_t}^k, M_{r_t}^k)\right)$, where Φ denotes the cumulative density function of the standard normal distribution and $\pi_{it}(\bar{y}_{it-1}, \bar{\mathbf{x}}_{it})$ is the probability of being observed at wave t given $(\bar{y}_{it-1}, \bar{\mathbf{x}}_{it})$ for the subset that satisfies $\bar{r}_{it-1} = 1$ and $s_{it} = 1$. Note that, $r_{i0} = 1$ and $s_{i0} = 1$ for all individuals, and that $\pi_{it} = 0$ if $r_{it-1} = 0$.

4.2. Algorithm

The estimator of the PPCM as described in Section 3 can be computed using the algorithm in Table 1. In practice, draws from the posterior distribution of the BART models are generated using Markov chain

Table 1. *Algorithm for estimation of the PPCM as described in Sections 3 and 4.*

-
1. *Models for the outcomes and response mechanisms:*
For $t = 1, \dots, T$, sample from the observed data posteriors for the parameters of the conditional distribution of y_{it} and r_{it} using DART.
 2. *Sensitivity parameters:*
For all $i \in U$ and $t = 1, \dots, T$, sample one set from the prior distributions for γ_{it} and δ_{it} .
 3. *Compute predicted means:*
For all $i \in U$ and $t = 1, \dots, T$, compute \hat{y}_{it} in [3.2] as follows:
Sequentially (in t) sample response and outcome data from $r_{it}^* \sim \text{Ber}(\pi_{it})$ and $y_{it-1}^* = \hat{m}_{t-1}(\bar{y}_{it-2}^*, \bar{r}_{it-1}^*, s_{it-1} = 1, \bar{\mathbf{x}}_{it-1}) + \gamma_{t-1}I(\hat{r}_{it-1}=0) + \varepsilon_{it}^*$. The conditional probability π_{it} conditions on $\bar{\mathbf{x}}_{it}, \bar{y}_{it-1}^*, \bar{r}_{it-1}^*$. For all $s_{it} = 1$, compute $\hat{y}_{it} = \hat{m}_t(\bar{y}_{it-1}^*, \bar{r}_{it}^*, s_{it} = 1, \bar{\mathbf{x}}_{it}) - \gamma_{it}I(r_{it}^*=0) - \delta_{it}$.
 4. *Compute the PPCM:*
Compute one posterior sample of PPCM = $\frac{1}{\sum_{t=1}^T \sum_{i \in U} s_{it}} \sum_{t=1}^T \sum_{i \in U} \hat{y}_{it} s_{it}$.
 5. Repeat steps 2–4 for each posterior sample.
-

Monte Carlo (MCMC). The parameters of the models for y_{it} and r_{it} are assumed independent and thus their posteriors can be sampled simultaneously. We use the sparse Dirichlet splitting rule prior for BART [the sparse Dirichlet splitting rule prior for Bayesian Additive Regression Trees (DART); Linero, 2018] to encourage parsimony, implemented in the R package *BART* for continuous and binary responses. To simplify notation we denote the conditional probability $\pi_{it}(\bar{y}_{it-1}, \bar{\mathbf{x}}_{it})$ by π_{it} and the conditional expectation for the outcome at time t , $\mu_{it}(\bar{y}_{it-1}, \bar{\mathbf{x}}_{it})$ by μ_{it} .

5. SIMULATION STUDY

In this simulation study, we compare five estimators for population inference. The sample and population size, the response rate, and the strength of association between \hat{y}_{it} and y_{it} (i.e., the R^2 estimated using BART) are chosen to mirror the Betula study.

5.1. Data generating process

We consider a finite population of size $N = 10\,000$, a sample of size $n = 1000$, and 1000 simulated data sets were generated. The auxiliary variables are generated independently as follows, $x_1, x_2 \sim \text{Bernoulli}(0.5)$ and $x_3, x_4, \dots, x_8 \sim \text{Uniform}(-1, 1)$, where $x_5 - x_8$ are uncorrelated with the outcome. We consider two time points ($t = 0, 1$). The response rate for y_{i1} was set to approximately 75%. Here, interest is in estimating the population mean at $t = 1$, $\mu_U = N^{-1} \sum_{i \in U} y_{i1}$. The approaches are compared in terms of bias, standard deviation (SD), mean squared error (MSE), and coverage of 95% credible intervals. We consider five true outcome models:

1. *Generalized linear additive models for the sample selection, response mechanism, and outcomes.*
Given the auxiliary variables $x_1 - x_8$, the outcome values for $i \in U$ were generated as $y_{0i} = -1 - x_{1i} + x_{2i} + x_{3i} + x_{4i} + \varepsilon_{i1}$ and $y_{1i} = -1 - x_{1i} + x_{2i} + x_{3i} + x_{4i} - 0.3y_{0i} + \varepsilon_{i2}$, where $\varepsilon_i \sim N(0, 1)$. The sample selection was generated from the following model $\text{logit}(\pi_i^c) = -2.67 - 0.4x_{1i} + 0.4x_{2i} + 0.4x_{3i} + 0.4x_{3i}$. In each selected sample, for $i \in c$ nonresponse to the study variable y_1 was generated from the following model $\text{logit}(\pi_i^r) = -2.7 + 1.2x_{1i} + 1.2x_{2i} + 1.2x_{3i} + 1.2x_{4i} - 1.2y_{0i}$.

2. *Interaction and nonlinear dependencies for the response mechanism.* y_{0i} , y_{1i} , and π_i^c were generated as for Scenario 1. The response mechanism was generated from the following model $\text{logit}(\pi_i^r) = -2.7 - x_{1i} + x_{2i} + x_{3i} + x_{4i} + y_{0i} + x_{3i}x_{4i} + x_{3i}x_{1i} + y_{0i}x_{1i}$.
3. *Interactions, nonlinear dependencies and skew normal error terms.* y_{0i} , π_i^r , and π_i^c were generated as for Scenario 2. y_{1i} was generated according to $y_{1i} = -0.87 - 0.4x_{3i} + 0.8x_{3i}^2 + 0.8x_{3i}^3 + 0.4x_{4i} + 0.8x_{1i} + 0.8x_{2i} + 0.4y_{0i} - 0.4x_{1i}y_{0i} + \varepsilon_{it}$, where ε_{it} were generated from the skew normal distribution, such that $\varepsilon_{it} \sim SN(-1.6 \times \frac{5}{\sqrt{1+5^2}} \times \sqrt{\frac{2}{\pi}}, 1.6, 5)$ for $t=0,1$, i.e. a right skewed variable with 0 mean, a variance of 1 and a skewness of 1.3.
4. *Interactions, nonlinear dependencies and practice effects.* y_{0i} , π_i^r and π_i^c were generated as for Scenario 3. For the outcome y_{1i} , we added a moderate PE of 0.1 to Scenario 3.
5. *Interactions, nonlinear dependencies, and deaths.* While y_{0i} , y_{1i} , π_i^r , and π_i^c were generated as for Scenario 3, survival at $t = 1$, s_{1i} , was generated from $\text{logit}(\pi_i^s) = 1.7 + 0.35x_{1i} + 0.35x_{2i} + 0.35x_{3i} + 0.35x_{4i}$, for a overall death rate of 12%.

5.2. Estimators for population inference

We compare five estimators for population inference. Our semiparametric model-based approach (MB-sp), was implemented as described in Section 4 and Table 1, but fixing the sensitivity parameters to 0. However, for the fourth scenario, we used the PE sensitivity parameter and specified a triangular prior reflecting practice effects, $\delta_{it} \sim \text{Tri}(0, b, b)$ with $b = 0.05, 0.1, 0.15$. Note that for the other approaches the sensitivity parameters were set to 0 for all scenarios. Details of the other four estimators (MB-lm, HT, MRP, and GREG) for longitudinal data are given in Appendix A of the [Supplementary materials](#). Briefly, MB-lm, is a parametric version of the MB-sp, specifying the working models as Bayesian additive linear regression models instead of using DART. Default uninformative priors were used. HT is an extension of the classic Horvitz–Thompson weighting estimator, where the inclusion weights were replaced by longitudinal *probability of participation weights*. The general regression estimator (GREG), combines prediction and longitudinal weighting. For GREG, the weights were computed as for HT and the prediction models were estimated using additive linear regression models. MRP is an extension of multilevel regression and post-stratification ([Gelman and Little, 1997](#)), where the binary covariates and the outcome at $t = 0$ were added to the model as fixed effects and the six continuous variables were first categorized into quartiles and added to the model as random effects.

The inclusion weights used in HT and GREG must be estimated using cell weight adjustment. To avoid sparse cells only $x_1 - x_4$ were considered for computing the weights, hence, the uncorrelated variables $x_5 - x_8$ were omitted. The continuous variables x_3 and x_4 were first categorized into tertiles. Every unique combination of the categorized variables constituted an adjustment cell. Sparse cells, $n_j < 20$, were combined with their nearest, nonsparse neighbor, i.e., the cell that has the most similar combination of auxiliary variables. Weights larger than 30 were trimmed.

The HT and GREG estimators and their 95% confidence intervals were estimated using the *mase* package in R, the working models for MB-lm and MRP were fitted using the *MCMCpack* package and *MCMCglmm* package in R ([R Core Team, 2018](#)). MCMC convergence and mixing were monitored using trace plots.

5.3. Simulation results

We present the bias, empirical standard deviation (SD; calculated as the standard deviation of the parameter estimates), mean-squared error (MSE), and coverage probabilities for the 95% credible intervals (CP) from the 1000 simulations in Table 2. The results show that inferences using the sample are highly biased, have

Table 2. Results from 1000 simulations with a finite population of size 10 000 and a sample size of 1000. Bias, empirical standard deviation (SD; calculated as the standard deviation of the parameter estimates), mean-squared error (MSE), and coverage probabilities in % for the 95% credible intervals (CP). Bias and SD are multiplied by 100 for ease of presentation. Sample: inference using the sample. MB-sp, a semiparametric model-based approach using DART; MB-lm, a parametric model-based approach using additive linear regression; MRP, multilevel regression and post-stratification; GREG, general regression estimator; HT, the Horvitz–Thompson estimator.

Method	Scenario 1				Scenario 2			
	Bias	SD	MSE	CP	Bias	SD	MSE	CP
Sample	18.8	5.3	3.8	3.7	18.8	5.1	3.8	3.1
MB-sp	0.5	4.2	0.2	93.6	0.4	4.3	0.2	94.2
MB-lm	0.2	3.9	0.2	95.2	−0.1	4.2	0.2	94.5
MRP	1.6	4.3	0.2	89.6	1.6	4.6	0.2	88.9
GREG	0.1	4.1	0.2	89.6	−0.1	4.2	0.2	89.5
HT	4.8	4.7	0.4	86.7	7.2	4.7	0.7	73.9
	Scenario 3				Scenario 4			
	Bias	SD	MSE	CP	Bias	SD	MSE	CP
Sample	41.1	5.6	17.2	0.0	51.0	5.5	26.3	0.0
MB-sp	1.5	4.9	0.3	91.6	11.8	4.6	2.6	23.3
MB-lm	11.8	5.1	1.7	33.0	21.7	5.0	5.0	0.9
MRP	19.7	5.4	4.2	1.9	29.6	5.7	9.1	0.1
GREG	9.1	5.4	1.1	45.5	18.6	5.2	3.7	2.1
HT	21.2	6.4	4.9	6.9	30.9	6.3	9.9	0.3

the highest MSEs and the lowest coverage probabilities under all four scenarios ($\leq 3.7\%$). The weighting approach (HT) performs poorly with large bias, SDs and MSEs, and the lowest CPs compared to the other four approaches across Scenario 1–4.

Scenarios 1 and 2 give similar findings for MB-sp, MB-lm, MRP, and GREG. That is, all approaches are unbiased, or nearly unbiased, and have similar SDs and MSEs. MB-sp and MB-lm have CPs above or close to 95%, while MRP and GREG have CPs just below 90%. Under Scenario 3 and 4, MB-lm, MRP and GREG are severely biased and their CPs are low (Scenario 3: 6.9–42.0% and Scenario 4: 0.1–2.1%). MRP has the highest biases and MSEs, and the lowest CPs. In contrast, under Scenario 3 our MB-sp approach is nearly unbiased, has the lowest MSE, and a CP of 91.6%. Under Scenario 4 (where the true PE was 0.1), our MB-sp approach revealed a bias of 0.118 (SD = 0.046) and a CP of 23.3%. Which was nevertheless best among all the approaches. For the analyses adjusting for practice effects, when the upper bound (and mode) of the Triangular prior was set to 0.05, the bias was 0.087 (SD = 0.048). If the upper bound was set to the true PE, 0.1, the bias was further reduced to 0.045 (SD = 0.043). Finally, if the upper bound was set to 0.15 (with a mean equal to the true PE) the bias was 0.019 (SD = 0.047) and the CP was 90.1%, which is comparable to Scenario 3. The CPs were 50.0% and 85.2% for the upper bounds of 0.05 and 0.1, respectively. These results demonstrate the importance of the PE adjustment. For Scenario 5, adjusting for truncation due to death, we compared, for simplicity of implementation, results for the three model-based approaches. MB-sp, MB-lm, and MRP revealed results comparable to Scenario 3, i.e., MB-sp: bias = 0.023 (SD = 0.047) and CP = 92.3%, MRP: bias = 0.195 (SD = 0.06) and CP = 2.4%, and MB-lm: bias = 0.115 (SD = 0.050) and CP = 32.6%.

6. ANALYSIS OF THE BETULA DATA

We applied the proposed semiparametric approach for estimating the PPCM(*age*) to the Betula data. Here, interest is in estimating the average memory performance across the adult lifespan among nondemented individuals given survival up to a specific age. Note, that both death and dementia are considered to be sources of truncation in the analysis. We consider population inference in the context of longitudinal cohort data in the presence of practice effects, nonignorable dropout, and death. Longitudinal auxiliary data for both the sample and target population is available for four waves of data collection. These include the baseline variables age, sex, having children (Y/N), and highest level of education, and additionally, several income variables, benefits received from the government, and marital status were treated as time varying. Details of baseline characteristics are found in Table 3. Details of the International Classification of Diseases codes are given in [Appendix B](#) of the [Supplementary material](#) available at *Biostatistics* online.

6.1. Sensitivity parameters

We explore sensitivity to default assumptions by specifying informative priors for the sensitivity parameters. Since there is generally little information about the distributional form for the SPs, the prior distributions are often chosen by the analyst reflecting prior beliefs about the departures from the default assumptions.

We assume triangular priors for δ_{it} , reflecting the improvement in memory performance that occur at repeated testings, i.e., practice effects (PEs). The priors were specified as $\delta_{i1} \sim \text{Tri}(0, U_{\delta_1}(a_{i1}), U_{\delta_1}(a_{i1}))$ and $\delta_{i2}, \delta_{i3} \sim \text{Tri}(0, U_{\delta_2}(a_{i2}), U_{\delta_2}(a_{i2}))$, where $A_{it} = a_{it}$ is the age of individual i at wave t . The upper bounds are given by; $U_{\delta_1}(a_{i1}) = 4.8 - 0.1 \times a_{i1} + 5.2 \times 10^{-4} \times a_{i1}^2$ and $U_{\delta_2}(a_{i2}) = 11.0 - 0.3 \times a_{i2} + 1.9 \times 10^{-3} \times a_{i2}^2$.

Table 3. *Baseline descriptive statistics for the Betula sample and the target population. Mean (standard deviation) for continuous variables and N (%) for categorical variables. CVD, cardiovascular disease. * amounts in Swedish kronor rounded to the nearest thousand.*

	Sample	Population
N	1000	9203
Age	57.5 (14.4)	52.0 (13.3)
Male	470 (47.0)	4431 (48.1)
Having children	856 (85.6)	7952 (86.4)
≤ 9 years of education	308 (30.8)	2505 (27.2)
10–12 years of education	352 (35.4)	3288 (35.7)
> 12 years of education	240 (24.2)	2516 (27.3)
Education level unknown	100 (10.1)	894 (9.7)
Married	628 (63.2)	5960 (64.8)
Widow	149 (14.9)	715 (7.8)
Disposable income*	1021.3 (412.8)	1058.6 (490.5)
Earned income*	868.9 (936.2)	1045.2 (952.6)
Retirement income*	275.1 (404.3)	165.4 (335.0)
Early retirement income*	37.2 (172.6)	29.6 (151.1)
Unemployment benefits*	4.4 (48.3)	9.1 (64.5)
Health benefits*	68.7 (220.0)	86.7 (227.8)
Diagnosed with diabetes prior 1990	14 (1.4)	185 (2.0)
Diagnosed with CVD prior 1990	97 (9.7)	593 (6.4)

To obtain these we considered two comparable samples (S1 and S3) from the Betula study. The estimated practice effects were derived as the mean differences in memory performance between participants taking the test for the 2nd/3rd time (sample S1) versus those taking the test for the 1st/2nd time (sample S3), while adjusting for age and age squared. Hence, $U_{\delta_{t+1}} = E(Y_{it+1}|A_{it+1} = a_{it+1}, S1, R_{it+2} = 1) - E(Y_{it}|A_{it} = a_{it+1}, S3, R_{it+1} = 1)$, for $t = 0, 1$, and $j = 1, \dots, N_{S3}$ refers to individuals in sample S3. By conditioning on R_{it+1} we avoid nonresponse bias. Note that, practice effects were larger for younger participants and at $t = 1$ compared to $t = 0$. We similarly specify a triangular prior for $\gamma_{it} \sim \text{Tri}(L_{\gamma_{it}}(a_{it}), 0, L_{\gamma_{it}}(a_{it})),$ reflecting a decline in memory after dropping out of the study. The bounds are given by; $L_{\gamma_{it}}(a_{it}) = -8.0 - 0.3 \times a_{it} + 3.9 \times 10^{-3} \times a_{it}^2$, obtained from estimating, $E(Y_{i1} - Y_{i0}|A_{i1} = a_{i1}, S1, R_{i1} = 1)$, i.e., the change in memory performance between the first and second wave for responders while adjusting for age and age squared. Note that, this imply that γ_{it} is the same for all t and that older participants decline more quickly than younger participants after drop out.

6.2. Results

We estimated the PPCM using our proposed Bayesian semiparametric approach with the sensitivity parameters. For each chain, the first 1000 iterations were discarded as burn-in, and a total of 1000 posterior samples of the PPCM were obtained. Convergence of the posterior samples was monitored using trace plots.

In the main analysis, the priors for the SPs were specified as described in the previous section. For the sensitivity analysis, we compare the results for different values of the SPs. Specifically, the SPs were set to (i) $\gamma_{it} = 0$ while δ_{it} was specified as in the main analysis (MARS and PE adjustment), (ii) $\delta_{it} = 0$ for $t = 2, 3, 4$, while γ_{it} was specified as in the main analysis (MNARS and no PE adjustment), (iii) the SPs were specified as twice as large as in the main analysis ($2 \times \delta_{it}$ for $t = 2, 3, 4$, and similarly, $2 \times \gamma_{it}$), and (iv) $\gamma_{it} = 0$ and $\delta_{it} = 0$ for $t = 2, 3, 4$ (MARS and no PE adjustment). We moreover compared the results from our mortal-cohort analysis to an analysis assuming an immortal cohort (Wen and others, 2018), treating both death and dropout as MAR ($\gamma_{it} = 0$); hence, deceased participants are implicitly included after death. In this scenario, only baseline auxiliary variables were considered since x_{it} is missing when $s_{it} = 0$ and $\delta_{it} = 0$ for all t (no PE adjustment).

The results from the main analysis and the various sensitivity analyses are presented in Figure 1. The corresponding 95% credible intervals are also plotted for the main analysis, the immortal-cohort analysis, and for the analysis assuming MARS and no PE adjustment (Figure 1a). For the main analysis, the age-specific PPCM revealed an initial decline in memory performance between the ages 35 to 65 and accelerated decline after the age of 65. Assuming an immortal cohort, no PEs and MAR nonresponse (both for dropout and death) resulted in a significantly higher estimated memory performance across adulthood, with a slightly greater decline at younger ages and less decline in memory after the age of 65. An analysis assuming MARS missingness and not adjusting for practice effects resulted in a significantly higher estimated memory performance between the ages 40–90 compared to the main analysis.

It is apparent in Figure 1b and c that adjusting for both PEs and MNARS dropout resulted in lower estimated memory performance across adulthood, although the discrepancy varied in magnitude with respect to age. When the SPs were specified as twice as large as in the main analysis (Figure 1c), we see an increased reduction in memory performance, this is more pronounced at older ages. As expected, comparing a PE adjusted analysis to an analysis without PE adjustment (assuming MNARS dropout) revealed lower memory performance for the PE adjusted analysis, although the discrepancy was more pronounced at younger ages. Comparing results from an analysis assuming MNARS and MARS while adjusting for PEs, also revealed lower memory performance, and the difference was more pronounced at older ages.

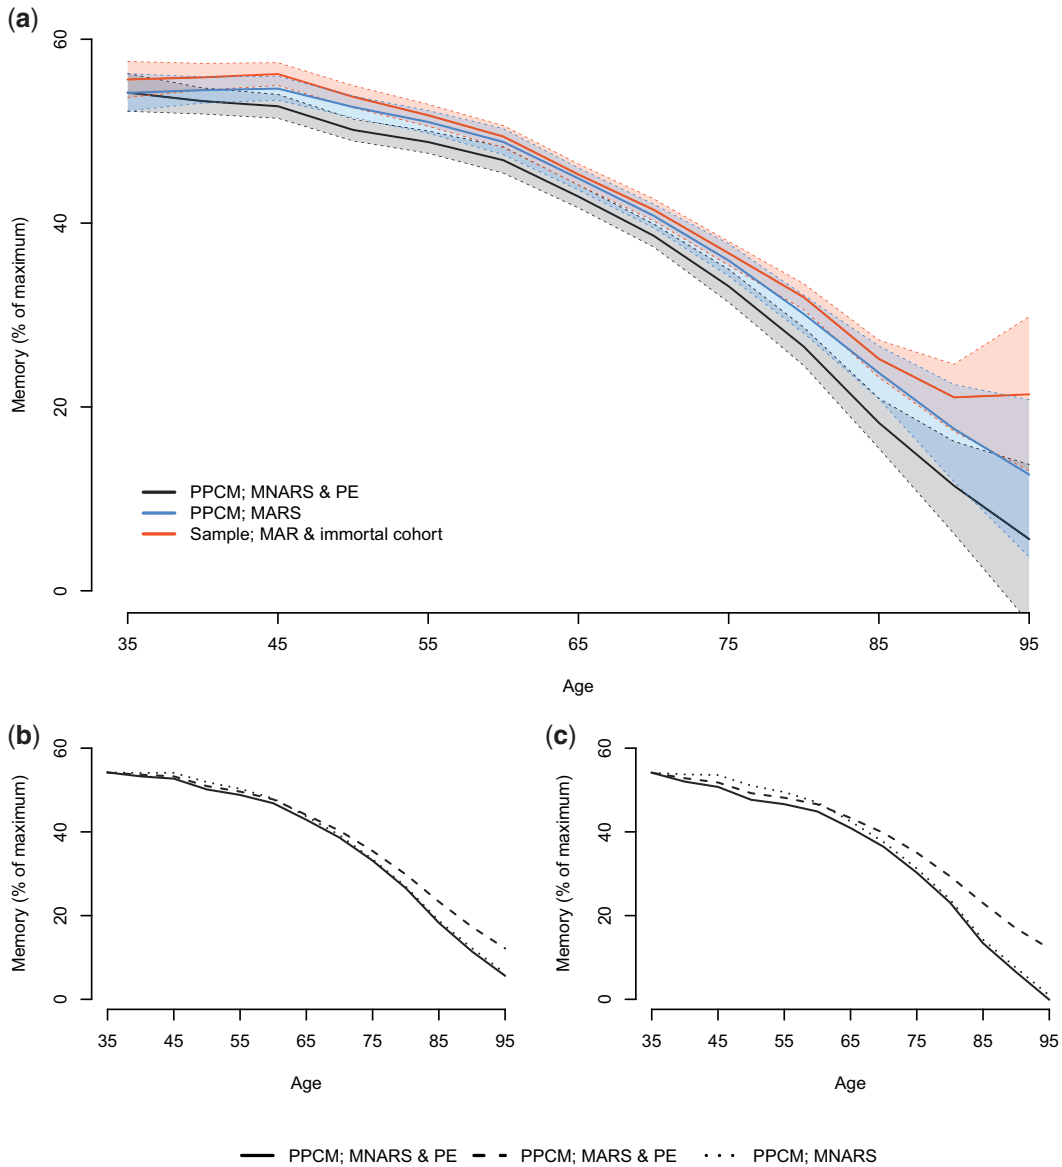


Fig. 1. Estimating lifespan trajectories of memory for the Betula data using our MB-sp approach. (a) *PPCM, MNARS, & PE*: the main analysis; *PPCM, MARS*: the sensitivity parameters were all set to 0; *Sample, MAR & immortal cohort*: sample inference assuming MAR missingness for both dropouts and deaths and no practice effects. Panels (b) and (c) show a sensitivity analysis for $k = 1$ and $k = 2$, respectively, comparing the main analysis (*PPCM, MNARS, & PE*) to an analysis, where $\gamma_i = 0$ and δ_{it} were specified as for the main analysis (*PPCM, MARS, & PE*), and to an analysis, where $\delta_{it} = 0$, while γ_i was specified as for the main analysis (*PPCM, MNARS*).

We compare the results for the Betula data using our approach (MB-sp) with the four other estimators for population inference, MB-lm, HT, MRP, and GREG. The sensitivity parameters, γ_{it} and δ_{it} for $t = 1, 2, 3$, were for simplicity all set to 0 when estimating the PPCM, i.e. assuming MARS missingness and no PEs.

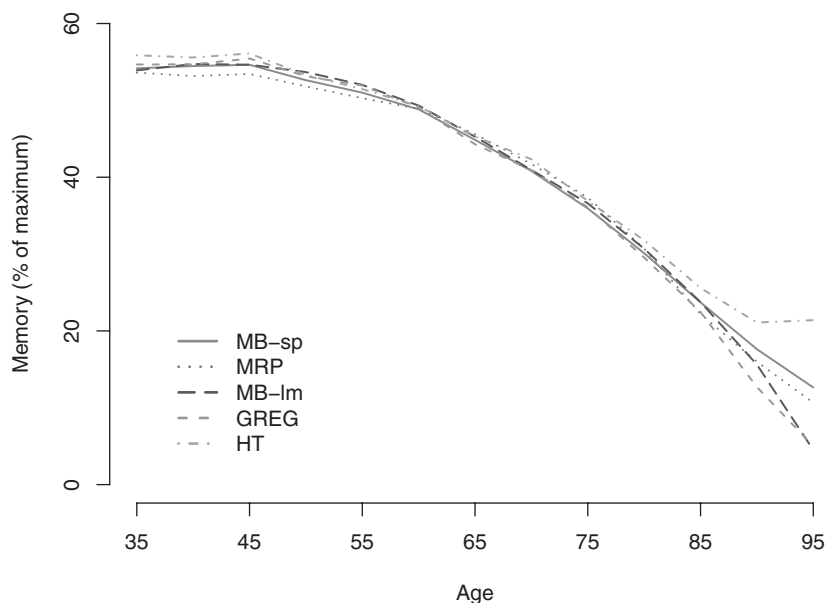


Fig. 2. Analysis of the Betula data. Results from estimating the population partly conditional mean using (1) MB-sp, our semiparametric model-based approach using DART; (2) MB-lm, a parametric model-based approach using additive linear regression; (3) MRP, multilevel regression and post-stratification, (4) GREG, general regression estimator; and (5) HT, the Horvitz–Thompson estimator. The sensitivity parameters γ_i , δ_{i2} , δ_{i3} , and δ_{i4} were all set to 0.

Details of the estimation procedures for MB-lm, HT, MRP, and GREG are described in Section 5.2 and in [Appendix A](#) of the [Supplementary material](#) available at *Biostatistics* online. However, some adjustments were made. For MRP, the continuous variables were first categorized into quartiles plus an additional category if the variable was 0 (added to the model as random effects). Due to excessive number of cells with a small or zero sample size we only used a subset of the baseline auxiliary variables for computing the cell adjustment weights π_i in HT and GREG. These were age, sex, level of education, widowhood, and history of cardiovascular disease.

Results are found in Figure 2. Compared to our semiparametric approach (MB-sp), the HT estimator overestimated memory performance across adulthood, and the three other parametric approaches revealed an accelerated decline in memory performance for individuals 70–95 years old. The large discrepancy in memory performance for HT, and for GREG at older ages, is likely a result from using only a subset of the auxiliary variables, dichotomizing continuous variables, and collapsing zero or spars cells, since this may introduce bias in the weights. Given the findings of the simulation study, we are inclined to believe the discrepancies for the parametric approaches (MB-lm and MRP) from our semiparametric approach are likely due to model-misspecification. Note that, MRP gives results most similar to our semiparametric approach.

7. CONCLUSIONS

This article proposes a Bayesian semiparametric predictive estimator for estimating the population partly conditional mean when a large set of longitudinal auxiliary variables is known for all units in the target population. A key feature is the flexible modeling approach that effectively addresses nonlinearity and complex interactions. Additionally, BART (using the sparse Dirichlet splitting rule prior) demonstrated

excellent predictive performance when irrelevant regressors were added, diminishing the need to carry out formal variable or model selection.

Our study is motivated by the fact that it is becoming increasingly difficult to recruit study participants, which may severely distort the ability to generalize study findings. The increased availability of microdata covering the population in many countries however, makes postsampling adjustments an attractive tool. Although weighting is the most popular technique, a large set of auxiliary (possibly continuous) variables makes cell weight adjustment difficult. In this setting model-based approaches are more attractive, but put stronger requirements on correct model specification. As expected, the results of the simulation study showed that the weighting approach (HT) performed poorly across a wide range of scenarios, despite a simplified scenario where uncorrelated variables were excluded. This is likely a result from collapsing sparse cells and dichotomizing the continuous auxiliary variables, thereby introducing bias in the weights. In contrast, the model-based approaches and GREG all performed well under the correct specification of the outcome model, although, our semiparametric method was the only approach that gave unbiased results for the more realistic scenario with unknown nonlinearity and interactions. Furthermore, under the scenario with practice effects our approach performed relatively well compared to the other approaches, showing the importance of adjusting for practice effects.

The goal of the empirical study was to estimate lifespan trajectories in memory for a target population. The results revealed an initial decline in memory performance between the ages 35 to 65, in contrast to previous studies showing stable performance up to the age of 60 (Rönnlund and others, 2005; Gorbach and others, 2017). Furthermore, the standard approach for estimation in previous literature that assumes an immortal cohort, no PEs and MAR nonresponse revealed significantly higher memory performance across the adult lifespan compared to our approach. This suggests that in previous studies the magnitude of memory performance across adulthood is likely overestimated while the rate of change is likely underestimated, especially at older ages. This is due to both selective study enrollment and attrition.

Our approach allows for Bayesian inference under MNAR missingness and truncation by death, as well as the ability to characterize uncertainty about practice effects. This was accomplished by introducing sensitivity parameters (SPs) that incorporated prior beliefs. A strength of the current approach is that inference with SPs and a mortal cohort is relatively easy to implement and communicate to nonstatisticians. However, specifying an appropriate prior distribution can sometimes be difficult; an alternative approach could be a tipping point analysis (Yan and others, 2009). In a tipping point analysis, subject matter experts can discuss whether the tipping point for the SPs are plausible, which may aid in making judgment based on study findings.

8. SOFTWARE

R code for the main analysis of the Betula data and the simulation study using our approach is available on https://github.com/m4ryjo/BSP_PS.git.

SUPPLEMENTARY MATERIAL

Supplementary material is available online at <http://biostatistics.oxfordjournals.org>.

ACKNOWLEDGMENTS

We thank Prof. Rolf Adolfsson and project coordinator Annelie Nordin Adolfsson, Department of Clinical Sciences at Ume University, for providing data from the Betula project for analyses.

Conflict of Interest: None declared.

FUNDING

The Swedish Foundation for Humanities and Social Sciences (P17-0196:1 to M.J.), and National Institutes of Health (R01 CA 183854 to M.J.D. and R01 GM 112327 to M.J.D.). The Betula Project is supported by Knut and Alice Wallenberg foundation and the Swedish Research Council (K2010-61X-21446-01).

REFERENCES

- BISBEE, J. (2019). BARP: improving mister p using Bayesian additive regression trees. *American Political Science Review* **113**, 1060–1065.
- CHIPMAN, H. A., GEORGE, E. I., MCCULLOCH, R. E. and others. (2010). BART: Bayesian additive regression trees. *The Annals of Applied Statistics* **4**, 266–298.
- DANIELS, M. J. AND HOGAN, J. W. (2008). *Missing Data in Longitudinal Studies: Strategies for Bayesian Modeling and Sensitivity Analysis*. Boca Raton: Chapman and Hall/CRC.
- DEVILLE, J.-C. AND SÄRNDAL, C.-E. (1992). Calibration estimators in survey sampling. *Journal of the American Statistical Association* **87**, 376–382.
- FRANGAKIS, C. E. AND RUBIN, D. B. (2002). Principal stratification in causal inference. *Biometrics* **58**, 21–29.
- FRANGAKIS, C. E., RUBIN, D. B., AN, M.-W. AND MACKENZIE, E. (2007). Principal stratification designs to estimate input data missing due to death. *Biometrics* **63**, 641–649.
- GELMAN, A. AND LITTLE, T. C. (1997). Poststratification into many categories using hierarchical logistic regression. *Survey Methodology* **23**, 127–136.
- GORBACH, T., PUDAS, S., LUNDQUIST, A., ORÄDD, G., JOSEFSSON, M., SALAMI, A., DE LUNA, X. AND NYBERG, L. (2017). Longitudinal association between hippocampus atrophy and episodic-memory decline. *Neurobiology of aging* **51**, 167–176.
- HOFFMAN, L., HOFER, S. M. AND SLIWINSKI, M. J. (2011). On the confounds among retest gains and age-cohort differences in the estimation of within-person change in longitudinal studies: a simulation study. *Psychology and Aging* **26**, 778.
- HORVITZ, D. G. AND THOMPSON, D. J. (1952). A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association* **47**, 663–685.
- JOSEFSSON, M., DE LUNA, X., DANIELS, M. J. AND NYBERG, L. (2016). Causal inference with longitudinal outcomes and non-ignorable dropout: estimating the effect of living alone on cognitive decline. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **65**, 131–144.
- JOSEFSSON, M., DE LUNA, X., PUDAS, S., NILSSON, L.-G. AND NYBERG, L. (2012). Genetic and lifestyle predictors of 15-year longitudinal change in episodic memory. *Journal of the American Geriatrics Society* **60**, 2308–2312.
- KERN, H. L., STUART, E. A., HILL, J. AND GREEN, D. P. (2016). Assessing methods for generalizing experimental impact estimates to target populations. *Journal of Research on Educational Effectiveness* **9**, 103–127.
- KURLAND, B. F., JOHNSON, L. L., EGLESTON, B. L. AND DIEHR, P. H. (2009). Longitudinal data with follow-up truncated by death: match the analysis method to research aims. *Statistical Science: A Review Journal of the Institute of Mathematical Statistics* **24**, 211.
- LI, Q. AND SU, L. (2018). Accommodating informative dropout and death: a joint modelling approach for longitudinal and semicompeting risks data. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **67**, 145–163.
- LINERO, A. R. (2018). Bayesian regression trees for high-dimensional prediction and variable selection. *Journal of the American Statistical Association* **113**, 626–636.
- LITTLE, R. J. (2004). To model or not to model? competing modes of inference for finite population sampling. *Journal of the American Statistical Association* **99**, 546–556.

- MALMBERG, G., NILSSON, L.-G. AND WEINEHALL, L. (2010). Longitudinal data for interdisciplinary ageing research. design of the linnaeus database. *Scandinavian Journal of Public Health* **38**, 761–767.
- MCGUINNESS, M. B., KASZA, J., KARAHALIOS, A., GUYMER, R. H., FINGER, R. P. AND SIMPSON, J. A. (2019). A comparison of methods to estimate the survivor average causal effect in the presence of missing data: a simulation study. *BMC Medical Research Methodology* **19**, 223.
- NILSSON, L.-G., BÄCKMAN, L., ERNGRUND, K., NYBERG, L., ADOLFSSON, R., BUCHT, G., KARLSSON, S., WIDING, M. AND WINBLAD, B. (1997). The Betula prospective cohort study: memory, health, and aging. *Aging, Neuropsychology, and Cognition* **4**, 1–32.
- PARK, D. K., GELMAN, A. AND BAFUMI, J. (2004). Bayesian multilevel estimation with poststratification: state-level estimates from national polls. *Political Analysis* **12**, 375–385.
- R Core Team. (2018). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- RIZOPOULOS, D. (2012). *Joint Models for Longitudinal and Time-to-Event Data, with Applications* in R. Boca Raton: Chapman & Hall/CRC.
- RÖNNLUND, M., NYBERG, L., BÄCKMAN, L. AND NILSSON, L.-G. (2005). Stability, growth, and decline in adult life span development of declarative memory: cross-sectional and longitudinal data from a population-based study. *Psychology and Aging* **20**, 3–18.
- RUBIN, D. B. (1976). Inference and missing data. *Biometrika* **63**, 581–592.
- SALTHOUSE, T. A. (2010). Influence of age on practice effects in longitudinal neurocognitive change. *Neuropsychology* **24**, 563.
- SHARDELL, M. AND FERRUCCI, L. (2018). Joint mixed-effects models for causal inference with longitudinal data. *Statistics in Medicine* **37**, 829–846.
- VIVOT, A., POWER, M. C., GLYMOUR, M. M., MAYEDA, E. R., BENITEZ, A., SPIRO III, A., MANLY, J. J., PROUST-LIMA, C., DUFOUIL, C. AND GROSS, A. L. (2016). Jump, hop, or skip: modeling practice effects in studies of determinants of cognitive change in older adults. *American Journal of Epidemiology* **183**, 302–314.
- WANG, W., ROTHSCHILD, D., GOEL, S. AND GELMAN, A. (2015). Forecasting elections with non-representative polls. *International Journal of Forecasting* **31**, 980–991.
- WEN, L. AND SEAMAN, S. R. (2018). Semi-parametric methods of handling missing data in mortal cohorts under non-ignorable missingness. *Biometrics* **74**, 1427–1437.
- WEN, L., TERRERA, G. M. AND SEAMAN, S. R. (2018). Methods for handling longitudinal outcome processes truncated by dropout and death. *Biostatistics* **19**, 407–425.
- WEUVE, J., PROUST-LIMA, C., POWER, M. C., GROSS, A. L., HOFER, S. M., THIÉBAUT, R., CHÊNE, G., GLYMOUR, M. M., DUFOUIL, C., MELODEM Initiative and others. (2015). Guidelines for reporting methodological challenges and evaluating potential bias in dementia research. *Alzheimer's & Dementia* **11**, 1098–1109.
- YAN, X., LEE, S. AND LI, N. (2009). Missing data handling methods in medical device clinical trials. *Journal of Biopharmaceutical Statistics* **19**, 1085–1098.

[Received November 24, 2020; revised March 17, 2021; accepted for publication March 17, 2021]