

FUNCTIONAL INFERENCE IN SEMIPARAMETRIC MODELS USING THE PIGGYBACK BOOTSTRAP

JOHN R. DIXON¹, MICHAEL R. KOSOROK² AND BEE LENG LEE³

¹*Department of Statistics, Florida State University, Tallahassee, FL 32306, U.S.A.*

²*Department of Statistics, University of Wisconsin-Madison, Madison, WI 53706, U.S.A.*

³*Department of Statistics and Applied Probability, National University of Singapore, Singapore 117546, Singapore*

(Received October 14, 2003; revised May 25, 2004)

Abstract. This paper introduces the “piggyback bootstrap.” Like the weighted bootstrap, this bootstrap procedure can be used to generate random draws that approximate the joint sampling distribution of the parametric and nonparametric maximum likelihood estimators in various semiparametric models, but the dimension of the maximization problem for each bootstrapped likelihood is smaller. This reduction results in significant computational savings in comparison to the weighted bootstrap. The procedure can be stated quite simply. First obtain a valid random draw for the parametric component of the model. Then take the draw for the nonparametric component to be the maximizer of the weighted bootstrap likelihood with the parametric component fixed at the parametric draw. We prove the procedure is valid for a class of semiparametric models that includes frailty regression models arising in survival analysis and biased sampling models that have application to vaccine efficacy trials. Bootstrap confidence sets from the piggyback and weighted bootstraps are compared for biased sampling data from simulated vaccine efficacy trials.

Key words and phrases: Biased sampling, bootstrap, censored data, confidence sets, empirical process, Monte Carlo inference, semiparametric efficiency, survival analysis.

1. Introduction

We propose a computationally quick alternative to the weighted bootstrap in a general class of semiparametric models, which we call the “piggyback bootstrap”. The semiparametric models for which the methodology is applicable include frailty regression models arising in survival analysis and biased sampling models that have application to vaccine efficacy trials. In the general set-up, the model consists of a parametric component θ and a nonparametric component A . More specifically, we consider likelihoods of the form $\prod_{i=1}^n \ell(\theta, A)(D_i)$. Here the contribution of the i -th subject to the likelihood, $\ell(\theta, A)(D_i)$, depends upon the data vector D_i corresponding to the i -th subject, a vector $\theta \in \mathbb{R}^d$, and a nonnegative function of bounded variation $A(t)$ defined for t in some finite interval $[0, \tau]$. In our survival analysis applications $A(t)$ is a cumulative hazard, and in our biased sampling application $A(t)$ is a cumulative distribution function (cdf). Our regularity conditions given in Section 5 will further constrain the likelihood.

Simulation methods such as the bootstrap provide a way to use information from a sample of size n to generate random draws which accurately approximate the Gaussian limit process of the maximum likelihood estimates (MLEs), $\hat{\theta}_n$ and \hat{A}_n , of θ and A . These

draws should be realizations of random variables θ_n and A_n that satisfy the following asymptotic property: $\sqrt{n}(\theta_n - \hat{\theta}_n, A_n - \hat{A}_n)$ converges weakly, given the sample data, to the same distribution that $\sqrt{n}(\hat{\theta}_n - \theta_0, \hat{A}_n - A_0)$ does unconditional on the sample data, as $n \rightarrow \infty$. We make this statement precise in Section 2.

Computationally simple simulation methods have been developed for the Cox model for right censored data. Kim and Lee (2003) propose an empirical Bayes method. The regression coefficients θ_n are drawn from a posterior distribution for θ given the data. For each draw so obtained, one samples A for from the posterior distribution of A given θ_n , which conveniently involves only generating gamma random variables. Unfortunately, this technique relies on the special structure of the Cox partial likelihood, a feature not shared by other semiparametric survival models. Shen (2002) gives conditions on appropriate priors for the parametric and nonparametric components in a more general semiparametric setting. However, no general method is given for finding these priors, if they even exist. And there is no guarantee that the posterior distributions will be computationally easy to sample from.

Lin *et al.* (1994) propose a Monte Carlo method for the Cox model. In their scheme the MLEs are computed, and then independent standard normal variables are plugged in to a simple expression involving the MLEs to obtain the random draws. Although this method is very simple computationally, it is unclear if an analogous approach can be devised for other semiparametric settings.

In contrast to these two simulation methods, the weighted bootstrap (Rubin (1981), Praestgaard and Wellner (1993)) is broadly applicable to many semiparametric models. Each term of the likelihood is weighted by a positive random variable satisfying specified moment conditions. The resulting bootstrap likelihood is maximized over θ and A to give a random draw. Unfortunately, this maximization is computationally intense in most semiparametric models, and must be repeated for each desired draw.

Tsodikov (2003) recommends using profile likelihood maximization in order to reduce the computational difficulty in obtaining the MLEs. Let $\ell_n(\theta, A) = n\mathbb{P}_n\ell(\theta, A)$ denote the log-likelihood function based on a sample of size n , and $\ell_n^\circ(A, \theta) = \mathbb{P}_n^\circ\ell(\theta, A) \equiv \mathbb{P}_n\eta\ell(\theta, A)/(\mathbb{P}_n\eta)$ denote a corresponding bootstrap log-likelihood. The profile bootstrap log-likelihood is defined by

$$pl_n^\circ(\theta) \equiv \sup_A \ell_n^\circ(\theta, A) = \ell_n^\circ(\theta, \hat{A}_\theta^\circ)$$

where $\hat{A}_\theta^\circ \equiv \arg \max_A \ell_n^\circ(\theta, A)$. Then the value of θ which maximizes $pl_n^\circ(\theta)$ coincides with the θ component of the joint maximizer of $\ell_n^\circ(\theta, A)$, and is thus the bootstrap MLE, which we will denote by $\hat{\theta}_n^\circ$. And for the full parameter $\psi \equiv (\theta, A)$, the bootstrap MLE is thus $\hat{\psi}_n^\circ \equiv (\hat{\theta}_n^\circ, \hat{A}_{\hat{\theta}_n^\circ}^\circ)$. Dropping the superscript “ \circ ”, the above applies to the original likelihood.

In Tsodikov’s scheme, one must use a search algorithm, such as Newton-Raphson or the Powell method (Press *et al.* (1994)) employed by Tsodikov, to find the value of θ which maximizes $pl_n^\circ(\theta)$. And for each candidate search value θ^* , one must compute $\hat{A}_{\theta^*}^\circ$, which we call a “profile computation”. In the Cox Proportional Hazards model, \hat{A}_θ and \hat{A}_θ° have explicit forms in terms of θ . For instance, \hat{A}_θ is Breslow’s estimator. In the other models we consider, the most explicit expression we have for \hat{A}_θ° is a self-consistency equation:

$$\hat{A}_\theta^\circ(t) = f_n(t; \theta, \hat{A}_\theta^\circ),$$

where $f_n(t; \theta, A)$ is based on the n observations. And so we must accomplish the profile computation by iterations of a fixed point algorithm, which we describe in detail in Section 2. Tsodikov (2003) presents ways to obtain the relevant self-consistency equation in a variety of semiparametric settings. He points out that this procedure can be viewed as an MM or EM algorithm. Since the fixed point algorithm often requires many iterations (see the simulation study in Section 4 for some specific average iterations), it ends up being the main source of computational cost in the weighted bootstrap.

The contribution of our paper is to take advantage of the profile structure elucidated by Tsodikov to perform computationally efficient inference for semiparametric models. In our proposed *piggyback bootstrap*, we assume that draws for the parametric component θ_n are readily available (we discuss methods of obtaining such draws in Section 2) and that $\sqrt{n}(\theta_n - \hat{\theta}_n)$, given the sample data, converges in distribution to the unconditional limiting distribution of $\sqrt{n}(\hat{\theta}_n - \theta_0)$. Then, for each parametric draw θ_n , the piggyback bootstrap draw is $A_n = \arg \max_A \ell_n^\circ(\theta_n, A)$, resulting in the pair (θ_n, A_n) .

Thus the piggyback bootstrap decreases the dimension of the maximization problem for each set of bootstrap weights. This method requires only one profile computation for each set of bootstrap weights. In contrast, the full weighted bootstrap requires a profile computation for each candidate value in the search for the parametric maximizer. Since the profile computations are accomplished using computationally costly iterations of a fixed point algorithm, we use the number of profile computations required by a procedure as a measure of its computational complexity. We show that the piggyback bootstrap is significantly less computationally complex than the weighted bootstrap.

Section 2 describes our method in more detail. Section 3 presents several example models for which the piggyback bootstrap works. In Section 4, we provide a numerical study comparing the weighted and piggyback bootstraps in a biased sampling analysis of simulated vaccine efficacy data. Sufficient regularity conditions are given in Section 5, and Section 6 contains a brief discussion. Details on the proofs of the results are given in Appendix A, and Appendix B discusses verification of the regularity conditions.

2. The piggyback bootstrap

In this section we introduce our piggyback bootstrap approach to obtaining appropriate random draws for semiparametric inference.

The main idea is to first obtain valid random draws for the parametric component of the model. Usually, it is possible to do this in a manner that is computationally much less intense than repeatedly maximizing profile likelihoods. The second step is to piggyback the draws for the nonparametric component onto the parametric draws, by plugging the parametric draws into a bootstrap likelihood and maximizing over the nonparametric component holding the parametric part fixed. That is, for each $\theta_n^{(k)}$ drawn, $k = 1, \dots, m$, we generate i.i.d. random bootstrap weights η_1, \dots, η_n , and compute $\hat{A}_{\theta_n^{(k)}}^\circ = \arg \max_A \ell_n^\circ(\theta_n^{(k)}, A)$, where ℓ_n° is the bootstrapped log-likelihood using the given bootstrap weights. We assume that these bootstrap weights are nonnegative, with mean and variance 1 and with $\int_0^\infty \sqrt{P[\eta_1 > x]} dx < \infty$. The bootstrap weights should be generated so that they are independent of the data. Further, the draws for θ_n should be i.i.d. and independent of the bootstrap weights.

Due to measurability issues which may arise in applying our theory to specific cases, we use empirical process results from van der Vaart and Wellner (1996) (hereafter abbre-

viated VW). The precise form of the bootstrap log-likelihood is $\ell_n^\circ(A, \theta) = \mathbb{P}_n^\circ \ell(\theta, A) \equiv \mathbb{P}_n \eta \ell(\theta, A) / (\mathbb{P}_n \eta)$. Hoffmann-Jørgensen (HJ) weak convergence is denoted by \rightsquigarrow .

Section 1.12 of VW gives a useful characterization of HJ weak convergence in a metric space \mathbb{D} to a tight limit. For a metric space \mathbb{D} , define BL_1 to be the set of all $f : \mathbb{D} \mapsto \mathbb{R}$ with $\|f\|_{\mathbb{D}} \equiv \sup_{x \in \mathbb{D}} |f(x)| \leq 1$ and $|f(x) - f(y)| \leq d(x, y)$ for every x, y , where d is the metric on \mathbb{D} . Then $X_n \rightsquigarrow X$, where X is tight if and only if $\sup_{f \in BL_1} |E^* f(X_n) - E f(X)| \rightarrow 0$.

Loosely speaking, we define $B_n \approx C_n$ to mean that B_n has a limit law conditional on the data equal to the limit law of C_n . In specific, if \mathbb{G} denotes the limit law of C_n , we require $\sup_{h \in BL_1} |E_\eta h(B_n) - E h(\mathbb{G})| \rightarrow 0$ in outer probability and that B_n is asymptotically measurable unconditionally. This has the form of the conclusion of Theorem 2.9.6 of VW, which we employ in Appendix A.

Now we state the main theorem. If we consider $\theta_n^{(1)}, \dots, \theta_n^{(m)}$ to be realizations of a random vector θ_n , then the following establishes the validity of the new approach:

THEOREM 2.1. *Under regularity conditions,*

$$\sqrt{n}(\theta_n - \hat{\theta}_n, \hat{A}_{\hat{\theta}_n}^\circ - \hat{A}_{\hat{\theta}_n}) \approx \sqrt{n}(\hat{\theta}_n - \theta_0, \hat{A}_{\hat{\theta}_n} - A_0).$$

The regularity conditions are given in Section 5, and Appendix A contains the proof. Appendix B provides an outline for verifying the regularity conditions in applications, with specific detail given for the biased sampling model introduced in Section 3.

Before utilizing this result, it is necessary to obtain draws $\theta_n^{(k)}$, $k = 1, \dots, m$, that have the right conditional distribution. The regularity condition that $\hat{\theta}_n$ is efficient implies that $\sqrt{n}(\hat{\theta}_n - \theta_0)$ is asymptotically mean zero normal with variance \tilde{I}_0^{-1} , where \tilde{I}_0 is the efficient Fisher information for θ . Thus one way to obtain the desired draws is to estimate \tilde{I}_0^{-1} with a consistent estimator \hat{V}_0 , and then let $\theta_n^{(k)} = \hat{\theta}_n + n^{-1/2} \hat{V}_0^{1/2} Z^{(k)}$, $k = 1, \dots, m$, where the $Z^{(k)}$ are independent standard normal vectors of length p , where p is the dimension of θ .

In some settings, such as the Cox model for right-censored data, a consistent estimator of \hat{V}_0 is not difficult to construct. Corollary 3 of Murphy and van der Vaart (2000) can be used to consistently estimate this covariance matrix in a number of other semiparametric settings. For the biased sampling example, we verify in Appendix B that the hypotheses of Corollary 3 of Murphy and van der Vaart hold, justifying our estimate described in Section 4.

An alternative Monte Carlo approach is presented in Lee (2000). This method builds on the results of Murphy and van der Vaart (2000). The proposed MCMC scheme generates θ_n from the density proportional to $\exp\{p\ell_n(\theta)\}$. Provided the draws stay sufficiently close to $\hat{\theta}_n$, we have $\sqrt{n}(\theta_n - \hat{\theta}_n) \approx \sqrt{n}(\hat{\theta}_n - \theta_0)$. See Lee (2000) for numerical studies of this algorithm, and Dixon (2003) for an implementation of this method in applying the piggyback bootstrap to an odds-rate regression on a Non-Hodgkin's Lymphoma data set.

In order to piggyback the draws for A onto the draws for θ we use the fixed point algorithm discussed in the Introduction. The same fixed point algorithm is used in implementing Tsodikov's (2003) method for computing the MLEs in the weighted bootstrap, as described in the Introduction. Before giving details on the algorithm, we mention that in the models we consider, we take the estimator for the nonparametric piece to be

a step function with jumps only at observation points. In the biased sampling models A is a cdf, and Vardi (1985) shows that the maximizer puts mass only at the observation points. In the survival analysis settings, the jumps are at observed failure times only. Murphy (1994) and Parner (1998) provide nice discussions of this issue. We will denote the jump at an observation t by $\Delta \hat{A}_\theta(t)$. Here and in what follows, for a right continuous function f we define $\Delta f(t) \equiv f(t) - f(t-)$.

Then the fixed point algorithm for evaluating $\Delta \hat{A}_\theta(t)$ for a fixed θ at each observation t is as follows:

- Step 1: Set $\Delta \hat{A}_\theta(t)^{(0)} \equiv g(t)$ for all observations t , where g is some initial guess function.

- Step 2: Compute

$$\Delta \hat{A}_\theta(t)^{(r+1)} \equiv f_n(t; \theta, \hat{A}_\theta^{(r)})$$

at each observation t .

- Step 3: If $\sup_t |\Delta \hat{A}_\theta^{(r+1)}(t) - \Delta \hat{A}_\theta^{(r)}(t)| < \epsilon$, stop and set $\hat{A}_\theta \equiv \hat{A}_\theta^{(r+1)}$. Otherwise repeat step 2 with r replaced by $r + 1$. Where ϵ is some tolerance level. (We took $\epsilon = .0001$ in our simulation study in Section 4.)

In our simulation study the algorithm described above always converged. And for the biased sampling model we consider in that study, Vardi (1985) has shown the self-consistency equation has a unique solution. Thus we know the resulting step function is the nonparametric MLE.

The next section describes several models that satisfy the regularity conditions of Theorem 2.1.

3. Examples

In this section we describe several models for which the piggyback bootstrap is applicable. The first example is the Cox model for right censored data. While computationally efficient inference for this model can be done with empirical Bayes (Kim and Lee (2003)) or the Monte Carlo approach of Lin *et al.* (1994), the example provides a straightforward illustration of the methodology. The next two examples are the odds-rate regression for right-censored data, and the correlated (and shared) gamma frailty model for right-censored data. A final example is the biased sampling model. The regularity conditions for this last example are established in Appendix B. Technical restrictions which simplify or are necessary in the verification of the regularity conditions are given for the biased sampling model in Appendix B, and also for the odds-rate model in Appendix B of Dixon, Kosorok, and Lee (2004) (hereafter abbreviated DKL). See Dixon (2003) for technical restrictions on the other models, as well as a detailed verification of the regularity conditions for these models.

3.1 Cox proportional hazards model

Efficient inference for θ in the Cox model for right censored data is based on the profile likelihood which is equivalent to the partial likelihood. The data consists of n i.i.d. realizations of (X, δ, Z) , where $X \equiv T \wedge C$ is the smaller of a failure time T and right censoring time C , $Z(\cdot)$ is a d -dimensional caglad (its components are left-continuous with right-hand limits) covariate process, and $\delta \equiv \mathbf{I}\{T \leq C\}$ is the censoring indicator, where $\mathbf{I}\{\cdot\}$ is the indicator function. The survival function for this model is given by $S(t | Z) \equiv P(T > t | Z) = \exp\{-\int_0^t e^{\theta' Z(s)} dA(s)\}$. Here θ is a d -vector of regression

coefficients, and $A(t)$, the cumulative baseline hazard, is a cadlag (right-continuous with left-hand limits) function of bounded variation on some finite interval $[0, \tau]$. Let $\hat{\theta}_n$ be the partial likelihood estimator for θ_0 and let \hat{V}_0 be the corresponding estimator of the variance of $\sqrt{n}(\hat{\theta}_n - \theta_0)$. The baseline hazard function can be estimated with Breslow's estimator $\hat{A}_{\hat{\theta}_n}$, where

$$\hat{A}_{\theta}(t) = \int_0^t \frac{\mathbb{P}_n dN(s)}{\mathbb{P}_n Y(s) e^{\theta' Z(s)}},$$

$N(t) = I\{X \leq t, \delta = 1\}$, $Y(t) = I\{X \geq t\}$, and \mathbb{P}_n is the empirical measure based on the data. A piggyback bootstrap is accomplished by drawing $\theta_n = \hat{\theta}_n + n^{-1/2} \hat{V}_0^{-1/2} Z$, where Z is a p -variate standard normal deviate. Then A_n is obtained by computing Breslow's estimate with bootstrap weights $\mathbb{P}_n^\circ(\cdot) \equiv \mathbb{P}_n \eta(\cdot) / (\mathbb{P}_n \eta)$ replacing the empirical weights $\mathbb{P}_n(\cdot)$.

3.2 Odds-rate model

A flexible regression model for right-censored data, which includes the Cox proportional hazards model and the proportional odds model as special cases, is the odds-rate regression model considered in Dabrowska and Doksum (1988), Scharfstein *et al.* (1998), and Lee (2000). The survival function for this model is defined as $S(t | Z) \equiv P(T > t | Z) = E[\exp\{-W \int_0^t e^{\beta' Z(s)} dA(s)\} | Z]$, where β is a d -vector of regression coefficients, W is an unobserved nonnegative gamma frailty with mean 1 and nonnegative variance γ , and $A(t)$, the cumulative baseline hazard, is a cadlag function of bounded variation on some finite interval $[0, \tau]$. After integrating over W , the survival function simplifies to $S(t | Z) = g_\gamma(\int_0^t e^{\beta' Z(s)} dA(s))$, where $g_\gamma(u) = (1 + \gamma u)^{-1/\gamma}$ for $\gamma \neq 0$ and $g_0(u) \equiv \lim_{\gamma \rightarrow 0} g_\gamma(u) = e^{-u}$. By multiplying the baseline hazard $e^{\beta' Z(t)} dA(t)$ by the random variable W , we are making an adjustment for misspecified or omitted covariates. Setting $\gamma = 0$ results in the Cox model, while setting $\gamma = 1$ yields the proportional odds model. We will focus on the case where γ is unknown.

Letting $\theta \equiv (\gamma, \beta)$ and assuming that censoring is independent of T given Z , the log-likelihood function for $\psi \equiv (\theta, A)$ in the odds-rate model is given by

$$\ell_n(\psi) = n \mathbb{P}_n \left\{ \int_0^\tau [\gamma \log g_\gamma(H_\psi^s(Y)) + \beta' Z(s) + \log a(s)] dN(s) + \log g_\gamma(H_\psi^\tau(Y)) \right\},$$

where $H_\psi^t(f) \equiv \int_0^t f(s) e^{\beta' Z(s)} dA(s)$ and $a \equiv dA/dt$.

As discussed in the Introduction, we consider A constant except for jumps at the observed failure times, replacing $a(s)$ with $\Delta A(s) (\equiv A(s) - A(s-))$. Then the bootstrap MLE may be computed in two steps via profile likelihood, as discussed in the introduction. For fixed θ , the self-consistency equation for \hat{A}_θ° is given by

$$\hat{A}_\theta^\circ(t) = \int_0^t W_n^\circ(u; \theta, \hat{A}_\theta^\circ)^{-1} dG_n^\circ(u),$$

where $G_n^\circ(u) = \mathbb{P}_n \eta \delta I\{X \leq u\}$ and $W_n^\circ(u; \theta, A) = \mathbb{P}_n [\eta (1 + \delta \gamma) e^{\beta' Z(u)} Y(u) / (1 + \gamma H_\psi^X(1))]$. Taking the bootstrap weights to be equal to 1 with probability 1, we have a self-consistency equation for the ordinary MLE.

DKL (2004) shows that under some additional technical restrictions, the piggyback bootstrap is applicable to this model. As discussed in Section 2, the MCMC approach

of Lee (2000) can be used to generate the parametric draws. See Dixon (2003) for an implementation of this method in applying the piggyback bootstrap to an odds-rate regression on a Non-Hodgkin’s Lymphoma data set.

3.3 Correlated (and shared) gamma frailty model

Much of the theory for the odds-rate model is derived in a manner similar to the results of Parner (1998). Parner considers the correlated gamma frailty model for clustered data. The data is right censored with the observations for each of n clusters (e.g., family, pair of twins, group) consisting of m realizations of $(X_i, \delta_i, Z_i(u))$, $i = 1, \dots, n$. We associate with the ij -th individual, $i = 1, \dots, n$, $j = 1, \dots, m$, a frailty $W_i^{(j)} \equiv W_{i0} + W_{ij}$. The frailty components $W_{i0}, W_{i1}, \dots, W_{im}$ are assumed to be independent, unobservable, gamma-distributed random variables with parameters $(\nu, \eta), (\nu^*, \eta), \dots, (\nu^*, \eta)$, respectively. We adopt the parameterization used by Parner by considering the frailty parameters of the model to be $\theta \equiv \text{var}(W_0) = \nu/\eta^2$ and $\theta^* \equiv \text{var}(W_j) = \nu^*/\eta^2$. When θ^* is set equal to 0 we have the shared frailty model. The full parameter for the model is $\psi \equiv (\theta, \theta^*, \beta, A)$. The log-likelihood function is given by

$$\sum_{i=1}^n \left\{ \sum_{j=1}^m [\delta_{ij}(\beta' Z_{ij}(X_{ij}) + \log a(X_{ij}))] + \log \sum_{\mathbf{k} \in K_i} \alpha_i(\mathbf{k}, \psi) \right\},$$

where $\mathbf{k} = (k_1, \dots, k_m)$ and $K_i = \{\mathbf{k} \mid k_j \in \{0, \delta_{ij}\}, j = 1, \dots, m\}$. And the $\alpha_i(\mathbf{k}, \psi)$ are given by

$$\frac{\theta^{\delta_{i\cdot}} \Gamma(\theta\theta^{-2} + \delta_{i\cdot} - k_{\cdot})}{\Gamma(\theta\theta^{-2})(1 + \theta \cdot \Lambda_i(\psi))^{\theta\theta^{-2} + \delta_{i\cdot} - k_{\cdot}}} \prod_{j=1}^m \left\{ \frac{\Gamma(\theta^*\theta^{-2} + k_j)}{\Gamma(\theta^*\theta^{-2})} \frac{1}{(1 + \theta \cdot \Lambda_{ij}(\psi))^{\theta^*\theta^{-2} + k_j}} \right\}$$

where $\Lambda_{ij}(\psi) \equiv \int_0^T Y_{ij}(s) \exp(\beta' Z_{ij}(s)) da(s)$, $Y_{ij}(t) \equiv \mathbf{I}\{X_{ij} \geq t\}$, $\theta_{\cdot} \equiv \theta + \theta^*$, and otherwise a subscript “ \cdot ” denotes summation over the corresponding index, e.g. $\delta_{i\cdot} = \sum_{j=1}^m \delta_{ij}$. Using the bootstrap which assigns weight to the i -th cluster, $i = 1, \dots, n$, the bootstrap self-consistency equation for A is given by

$$\hat{A}_n^{\circ}(t) = \int_0^t \left(\frac{1}{n} \sum_{i=1}^n \eta_i \sum_{j=1}^m \hat{Z}_i^{(j)}(\theta, \theta^*, \beta, \hat{A}_n^{\circ}) Y_{ij}(s) \exp(\beta' X_{ij}(s)) \right)^{-1} d\bar{N}_{\cdot}^{\circ}(s).$$

Here $\hat{Z}_i^{(j)}(\psi) \equiv \hat{Z}_{i0}(\psi) + \hat{Z}_{ij}(\psi)$ with

$$\hat{Z}_{i0}(\psi) \equiv \frac{\sum_{\mathbf{k} \in K_i} \alpha_i(\mathbf{k}, \psi) b_{i0}(\mathbf{k}, \psi)}{\sum_{\mathbf{k} \in K_i} \alpha_i(\mathbf{k}, \psi)} \quad \text{and} \quad \hat{Z}_{ij}(\psi) \equiv \frac{\sum_{\mathbf{k} \in K_i} \alpha_i(\mathbf{k}, \psi) b_{ij}(\mathbf{k}, \psi)}{\sum_{\mathbf{k} \in K_i} \alpha_i(\mathbf{k}, \psi)},$$

where

$$b_{i0}(\mathbf{k}, \psi) = \frac{\theta\theta^{-1} + \theta \cdot (\delta_{i\cdot} - k_{\cdot})}{1 + \theta \cdot \Lambda_i(\psi)} \quad \text{and} \quad b_{ij}(\mathbf{k}, \psi) = \frac{\theta^*\theta^{-1} + \theta \cdot k_j}{1 + \theta \cdot \Lambda_{ij}(\psi)},$$

$\bar{N}_{\cdot}^{\circ} = \sum_{i=1}^n \eta_i \sum_{j=1}^m N_{ij}$, and (analogously to the odds-rate model setting) $N_{ij}(s) \equiv \mathbf{I}\{X_{ij} \leq s, \delta_{ij} = 1\}$. Taking $\eta_i = 1$ with probability 1, we have the self consistency equation for \hat{A}_n previously revealed by Parner (1998).

We omit verification of the conditions here, but with some additional technical restrictions, the piggyback bootstrap can be shown to be readily applicable to the correlated gamma frailty example. See Dixon (2003) for details.

3.4 *Biased sampling model*

As a final application, we consider a class of biased sampling models (Gilbert (2000)). Here the data consists of n i.i.d. realizations of $X = (I, Y)$. Here $I \in \{1, \dots, s\}$ is a random stratum, taking on the value i with selection probability $\lambda_i > 0$, $i = 1, \dots, s$, with $\sum_{i=1}^s \lambda_i = 1$. Given $I = i$, $Y \in [0, \tau]$ has distribution F_i defined on a sigma field of subsets \mathcal{B} of $[0, \tau]$ by $F_i(B, \theta, A) \equiv W_i^{-1}(\theta, A) \int_B w_i(u, \theta) dA(u)$ for $B \in \mathcal{B}$. The w_i , $i = 1, \dots, s$, are nonnegative (measurable) stratum weight functions assumed to be known up to the finite dimensional parameter $\theta \in \Theta$. $W_i(\theta, A) \equiv \int_0^\tau w_i(u, \theta) dA(u)$ is assumed to be finite for all $\theta \in \Theta$. The probability measure A is the unknown infinite dimensional parameter of interest. The goal is to estimate it based on information from samples from the F_i distributions, $i = 1, \dots, s$. Thus the semiparametric log-likelihood is given by

$$n\mathbb{P}_n \left\{ \lambda_I \frac{w_I(Y, \theta)}{W_I(\theta, A)} dA\{Y\} \right\}.$$

Denote the number of observations belonging to the i -th stratum by n_i , the total sample size by $n (= \sum_{i=1}^s n_i)$, and the i -th sampling fraction by $\lambda_{ni} = n_i/n$. Furthermore, let t_j , $j = 1, \dots, h$, denote the distinct observed realizations of Y , each with multiplicity r_j , and let n_{ij} be the number of observations from the i -th stratum with value t_j . The semiparametric likelihood is then proportional to

$$L_n(\psi \mid \text{data}) = \prod_{i=1}^s \prod_{j=1}^h \left[\frac{w_i(t_j, \theta) dA\{t_j\}}{W_i(\theta, A)} \right]^{n_{ij}}.$$

As demonstrated in Vardi (1985), we may instead maximize the partial likelihood defined by

$$L_{n1}(\theta, \tilde{B} \mid \text{data}) = \prod_{i=1}^s \prod_{j=1}^h \left[\frac{w_{ij}(\theta) B_i^{-1}}{\sum_{k=1}^s \lambda_{nk} w_{kj}(\theta) B_k^{-1}} \right]^{n_{ij}},$$

where $\tilde{B} = \{B_1, \dots, B_s\}$, subject to $B_i > 0$, $i = 1, \dots, s - 1$ and $B_s = 1$, and $w_{ij}(\theta) \equiv w_i(t_j, \theta)$. Again the MLE is computed in two steps via the notion of profile (partial) likelihood. The profile partial log-likelihood is given by $p\ell_n(\theta) \equiv \sup_{\tilde{B}} \log L_{n1}(\theta, \tilde{B}) \equiv \log L_{n1}(\theta, \hat{B}_\theta)$ where $\hat{B}_\theta \equiv \arg \max_{\tilde{B}} \log L_{n1}(\theta, \tilde{B})$. The MLE for θ coincides with the profile-partial-MLE obtained by maximizing $p\ell_n(\theta)$ with respect to θ . Then we take $\hat{\psi}_n \equiv (\hat{\theta}_n, \hat{A}_{\hat{\theta}_n})$.

As demonstrated in Vardi (1985), we may compute \hat{A}_θ by the following procedure. First find \hat{B} defined as the solution to the self-consistency equation

$$B_i(\theta) = \sum_{j=1}^h \frac{r_j w_{ij}(\theta)}{\sum_{k=1}^s n_k w_{kj}(\theta) B_k^{-1}(\theta)}$$

$i = 1, \dots, s$ subject to the constraints $B_i > 0$, $i = 1, \dots, s - 1$ and $B_s = 1$. Then we take

$$d\hat{A}_\theta(t_j) = \frac{r_j [\sum_{k=1}^s \lambda_{nk} w_{kj}(\hat{\theta}_n) \hat{B}_{nk}^{-1}(\hat{\theta}_n)]^{-1}}{\sum_{j=1}^h r_j [\sum_{k=1}^s \lambda_{nk} w_{kj}(\hat{\theta}_n) \hat{B}_{nk}^{-1}(\hat{\theta}_n)]^{-1}}.$$

In the weighted bootstrap case, we replace n_{ij} with the sum n_{ij}° of n_{ij} bootstrap weights in the above equations, and the other terms are modified accordingly, i.e., $n_i^\circ \equiv \sum_{j=1}^h n_{ij}^\circ$,

$n^\circ \equiv \sum_{i=1}^s n_i^\circ = n$ (since all bootstrap weights add up to n by construction), $r_j^\circ \equiv \sum_{i=1}^s n_{ij}^\circ$, and $\lambda_{ni}^\circ \equiv n_i^\circ/n$. Note that the use of the partial likelihood is for computational purposes only. It affects none of the theory because the nonparametric partial likelihood maximizer coincides with the nonparametric maximizer of the full likelihood.

The weighted bootstrap can be used to obtain draws for the parametric and nonparametric components of the model, as justified in Gilbert (1996). But, again, maximizing over both the parametric and nonparametric pieces in this approach is computationally intense. However, our piggyback approach applies in this situation, simplifying computations. To obtain draws for the parametric component, θ , Murphy and van der Vaart’s (2000) Corollary 3 can be used to consistently estimate \hat{I}_0 . Then the draws for the nonparametric piece are piggybacked on these draws as described in Section 2. The next section provides a simulation study using this approach, and verification of the regularity conditions is given in Appendix B.

4. Simulation study

We simulated vaccine efficacy trials in order to evaluate the coverage probabilities of confidence bands constructed using draws from the weighted and piggyback bootstraps. The simulated response represents percent divergence in the V3 loop amino acid sequence between the strain of HIV found in an infected subject versus the prototype strain used in the vaccine. These simulated data sets are modeled after ones found in Gilbert (1996).

We generated 200 data sets consisting of 400 independent observations. Half of the 400 observations in each data set come from a placebo group, and the other half from a vaccine group. The placebo group response was uniformly distributed on the interval $[0, 35]$, and the stratum bias functions are given by $w_1(y, \theta) = \exp(y\theta/35)$ and $w_2(y, \theta) = 1$. Here group 1 is the vaccine group and group 2 is the placebo group. Thus the cdf for the vaccine group is given by

$$F(y) = \frac{\int_0^y e^{s\theta_0/35} I_{[0,35]}(s) ds}{\int_0^\infty e^{s\theta_0/35} I_{[0,35]}(s) ds} = \frac{e^{y\theta_0/35} - 1}{e^{\theta_0} - 1},$$

and the cdf for the placebo group is the Uniform $[0, 35]$ cdf.

The choice of the Uniform $[0, 35]$ distribution is inspired by the use of this distribution in the simulation in Gilbert (1996). The value 35 corresponds to the suspected maximum possible percent divergence, and the true value of θ was taken to be $\theta_0 = 7.89$, arbitrarily. See Gilbert *et al.* (1998) for a scientific interpretation of the parameter θ .

For the parametric draws in our piggyback approach, we estimated the variance of the MLE $\hat{\theta}_n$ and then drew standard normal deviates with this variance and added them to $\hat{\theta}_n$ for each data set. To estimate the variance of $\hat{\theta}_n$, we applied Corollary 3 of Murphy and van der Vaart (2000). That is, we calculated

$$\lim_{h \rightarrow 0} -2 \frac{\log p\ell_n(\hat{\theta}_n + h) - \log p\ell_n(\hat{\theta}_n)}{400h^2} \equiv \hat{I}_0,$$

and took the variances of the normal deviates to be $(400\hat{I}_0)^{-1}$. See DKL for details on how this limit was computed in the simulations.

In order to maximize over the parametric component of the model in the weighted bootstrap (and in finding $\hat{\theta}_n$ to apply the method described in the previous paragraph for the piggyback bootstrap), we used a search algorithm that returns a value which is the

parametric MLE plus or minus .01. This algorithm is a modified grid search algorithm and does not compute any derivatives.

We performed 2000 bootstrap repetitions for both the piggyback and weighted bootstraps for each of the 200 simulated data sets. For reasons discussed in the Introduction, we use the number of profile computations required by each bootstrap procedure as a measure of computational complexity. In the weighted bootstrap, we profiled (maximized) over the nonparametric component of the model on average 58189.755 times per data set. The mean number of iterations of the fixed point algorithm for computing these A values was 35.758 with standard deviation 12.375. In order to obtain the variance estimates for the parametric component for use in the piggyback bootstrap, we profiled over the nonparametric component of the model on average 31.995 times per data set. (Note that this includes the profile computations needed to find the MLE for each data set.) The mean number of iterations for computing each A was 36.909 with standard deviation 12.468. In piggybacking the draws for A on the 2000 values of θ so drawn for each of the data sets, the mean number of iterations was 43.468 with standard deviation 9.867. Thus in the piggyback bootstrap we profiled over the nonparametric component of the model an average of 2031.995 times per data set compared to the average of 58189.755 times for the weighted bootstrap, a ratio of 1 to 29. It might be possible to streamline our maximization algorithm to decrease the number of times we must profile over the nonparametric component in the weighted bootstrap. But to beat the piggyback bootstrap in this sense, we would need to profile over the nonparametric component on average less than $2031.995/2000$ times for each of the 2000 bootstrap likelihoods. This would require that the first candidate value we try in each bootstrap likelihood search happens to be the MLE (plus or minus some small error) almost all the time, which is clearly beyond even the best search algorithms.

In constructing confidence bands for the cdfs of the vaccine and placebo groups, we considered only the middle 75 percent of the observations. The reason for this is that the bands narrow towards the ends of the data set due to the nature of the estimate. Indeed, the estimated value of the cdf for the largest observation will always be 1, and so the confidence band at that point will never contain the true value. To construct a level α pointwise confidence interval at an observation y , we find the middle α percent of the 2000 cdf estimates induced by the draws at y . A 95 percent simultaneous confidence band is constructed by finding the smallest value of α so that 95 percent of the cdfs induced by the draws are contained within the level α pointwise confidence intervals for all observations y in the middle 75 percent of the observations.

See Fig. 1 for the resulting confidence intervals and bands for one of the simulated data sets. Since the piggyback and weighted bootstraps only specify the values of the intervals and bands at the observation points, we interpolated to create these figures. We tried two interpolation schemes, which we call leftpoint and midpoint. In these schemes, the estimates were extended to the range of the data by creating a step function, with the observations in the data set taken as either the left endpoint or the midpoint, respectively, of each step. See DKL for a precise description.

Table 1 presents the coverage of these bands so defined for the piggyback approach and the weighted bootstrap for three interpolation schemes: none, leftpoint, and midpoint. The values for “none” are obtained by only considering whether or not a confidence band contains the true cdf at each observation in the data set, and not between observations as is the case for interpolation. Note that we prove in this paper that the two methods are asymptotically equivalent. Even for this moderate sample size, we see that




Fig. 1. MLE cdfs for placebo and vaccine groups with 95% confidence intervals and confidence bands, for a simulated biased sampling data set.

Table 1. Coverage of 95% confidence bands generated with the piggyback bootstrap (pb) and weighted bootstrap (wb) approaches for three different interpolation schemes.

| | pb | wb | pb | wb |
|---------------|---------|---------|---------|---------|
| Interpolation | placebo | placebo | vaccine | vaccine |
| none | 92.5% | 93.5% | 93% | 92.5% |
| leftpoint | 91% | 92% | 91.5% | 90.5% |
| midpoint | 92% | 93% | 93% | 92.5% |

the two methods perform close to the same. Although the piggyback bootstrap does slightly better than the weighted bootstrap for the vaccine group, and slightly worse than the weighted bootstrap for the placebo group.

We also constructed confidence intervals for the parametric component of the model. For each data set, the .975 and .025 percentiles of the draws for the parametric component were taken as the ends of the corresponding confidence interval. The coverage of these intervals was 94.5 percent using the weighted bootstrap, and 93 percent using the piggyback bootstrap, both of which are within 2 Monte Carlo standard deviations of 0.95 ($0.03 = 2\sqrt{0.95 \times .05/200}$).

5. Regularity conditions

In this section, we present sufficient technical conditions for the piggyback bootstrap to have the desired asymptotic property given in Theorem 2.1. This section may be skipped at first reading. Conditions (PB1) through (PB3) are conditions on the structure and smoothness of the information operator. Condition (PB4) involves asymptotic conditions on the Monte Carlo parameter estimates. Condition (PB5) requires sufficient smoothness of the score operator for the infinite dimensional parameter. Condition (PB6) requires the estimators to be efficient.

As noted in the Introduction, we consider likelihoods of the form $\prod_{i=1}^n \ell(\theta, A)(D_i)$. Here the contribution of the i -th subject to the likelihood, $\ell(\theta, A)(D_i)$, depends upon

the data vector D_i corresponding to the i -th subject, a vector $\theta \in \mathbb{R}^d$, and a nonnegative function of bounded variation $A(t)$ defined for t in some finite interval $[0, \tau]$. Restrictions on the parameter spaces of θ and A will be model specific. See Appendix B for restrictions on the biased sampling model, Appendix B of DKL for restrictions on the odds-rate model, and Dixon (2003) for restrictions on the other example models.

Let $h = (h_1, h_2)$, where $h_1 \in \mathbb{R}^d$ and h_2 is a nonnegative function of bounded variation on $[0, \tau]$. Let $\|\cdot\|_i$ denote the Euclidean norm for $i = 1$, the total variation norm for $i = 2$, and $\|\cdot\|_1 + \|\cdot\|_2$ for $i = 3$ (i.e. $\|\psi\|_3 = \|\theta\|_1 + \|A\|_2$). Let \mathcal{S} be a set. We use $o_p^{\mathcal{S}}(R_n)(s)$ to denote a term $F_n(s) = R_n Q_n(s)$ such that $\sup_{s \in \mathcal{S}} \|Q_n(s)\|$ converges in outer probability to 0 as $n \rightarrow \infty$ (here $\|\cdot\| = \|\cdot\|_i$ for $i = 1, 2$, or 3 as appropriate). We assume the score operator for the model,

$$\begin{aligned} U^n[\theta, A](h) &\equiv \frac{\partial}{\partial t} \ell_n \left(\theta + th_1, A + t \int_0^\tau h_2(u) dA(u) \right) \Big|_{t=0} \\ &= \frac{\partial}{\partial t} \ell_n(\theta + th_1, A) \Big|_{t=0} + \frac{\partial}{\partial t} \ell_n \left(\theta, A + t \int_0^\tau h_2(u) dA(u) \right) \Big|_{t=0}, \end{aligned}$$

exists. Denote the first term in the last line $U_{n,1}[\theta, A](h_1) \equiv n\mathbb{P}_n U_1[\theta, A](h_1)$ and the second $U_{n,2}[\theta, A](h_2) \equiv n\mathbb{P}_n U_2[\theta, A](h_2)$. Let $\bar{h}_1 \in \mathbb{R}^d$, and \bar{h}_2 be a nonnegative function of bounded variation on $[0, \tau]$. We assume the following derivative exists

$$(5.1) \quad \begin{aligned} &\frac{\partial}{\partial s} U_i[\theta + s\bar{h}_1, A + s\bar{h}_2](h_i) \Big|_{s=0} \\ &= \frac{\partial}{\partial s} U_i[\theta + s\bar{h}_1, A](h_i) \Big|_{s=0} + \frac{\partial}{\partial s} U_i[\theta, A + s\bar{h}_2](h_i) \Big|_{s=0}, \quad i = 1, 2. \end{aligned}$$

The following are the regularity conditions:

(PB1) The negative of the first term of the right hand side of (5.1) is of the form $\bar{h}'_1 \hat{\sigma}_{i,1}[\theta, A](h_i)$, and the negative of the second term is of the form $\int_0^\tau \hat{\sigma}_{i,2}[\theta, A](h_i)(u) dh_2(u)$ for $i = 1, 2$, where $\hat{\sigma}_{i,1}[\theta, A](h_i)$ is a random d -vector depending upon an observation, and $\hat{\sigma}_{i,2}[\theta, A](h_i)(u)$ is a random function of bounded variation on $[0, \tau]$ depending upon an observation for $i = 1, 2$. Define $\hat{\sigma}_{n,i,j}[\theta, A](h_i) \equiv \mathbb{P}_n \hat{\sigma}_{i,j}[\theta, A](h_i)$, $i, j = 1, 2$. Then $\hat{\sigma}_{n,i,j}[\theta, A](h_i)$ has total variation bounded by some M over all n .

(PB2) Denote $\sigma_{i,j}[\theta, A](h_i) \equiv P_0 \hat{\sigma}_{i,j}[\theta, A](h_i)$, $i, j = 1, 2$, where P_0 is the expectation under the true model. Let $\mathcal{H} = \mathcal{H}_1 \times \mathcal{H}_2$, where $\mathcal{H}_1 = \mathbb{R}^d$ and \mathcal{H}_2 is a set of functions which includes all $h(s) = \sigma_{2,2}[\theta_0, A_0]^{-1}(\mathbb{I}\{s \leq t\})$ with $t \in [0, \tau]$. And define $\|\cdot\|_j^{\mathcal{H}_2} \equiv \sup_{h \in \mathcal{H}_2} \|\cdot\|_j$, $j = 1, 2$. Then $\lim_{\theta \rightarrow \theta_0, A \rightarrow A_0} \|\sigma_{2,j}[\theta, A](h) - \sigma_{2,j}[\theta_0, A_0](h)\|_j^{\mathcal{H}_2} = 0$, $j = 1, 2$.

(PB3) For some $c > 0$, $\{\hat{\sigma}_{2,1}[\theta, A](h)_i, \|\theta - \theta_0\|_1 \leq c, \|A - A_0\|_2 \leq c, h \in \mathcal{H}_2\}$, $i = 1, \dots, d$, and $\{\hat{\sigma}_{2,2}[\theta, A](h)(u), \|\theta - \theta_0\|_1 \leq c, \|A - A_0\|_2 \leq c, h \in \mathcal{H}_2, u \in [0, \tau]\}$ are Donsker and bounded.

(PB4) (i) $\sqrt{n}(\hat{\theta}_n - \theta_n) \approx \sqrt{n}(\hat{\theta}_n - \theta_0)$ and the common limiting distribution is tight.

(ii) $\sqrt{n}(\hat{A}_{\theta_0}^\circ - \hat{A}_{\theta_0}) \approx \sqrt{n}(\hat{A}_{\theta_0} - A_0)$ and the common limiting distribution is tight and cadlag.

(iii) $\sqrt{n} \|\hat{A}_{\theta_n}^\circ - A_0\|_2 \leq \tilde{O}_p(1) + o_p^{\mathcal{H}_2}(1 + \sqrt{n} \|\theta_n - \theta_0\|_1)$. Here and in what follows, $\tilde{O}_p(1)$ denotes a random sequence which converges weakly to a tight, cadlag limit.

(PB5) $P_0 U_2(\theta_0, A_0)(\cdot) = 0$ and $\{U_2(\theta_0, A_0)(h), h \in \mathcal{H}_2\}$ is P_0 -Donsker.

(PB6) $\sqrt{n}(\hat{\theta}_n - \theta_0, \hat{A}_{\hat{\theta}_n} - A_0)'$ is asymptotically linear, regular, and efficient. And the information operator

$$\begin{pmatrix} \sigma_{1,1} & \sigma_{2,1} \\ \sigma_{1,2} & \sigma_{2,2} \end{pmatrix}$$

is continuously invertible and onto. Here, and in other cases where the full notation is not necessary for clarity, we abbreviate $\sigma_{i,j}[\theta_0, A_0] \equiv \sigma_{i,j}$, $i, j = 1, 2$.

We assume the bootstrap weights η_1, \dots, η_n above are i.i.d., nonnegative, and with mean and variance 1. A technical condition on the bootstrap weights useful in establishing the validity of the weighted bootstrap and condition (PB4)(ii) is that $\int_0^\infty \sqrt{P[\eta_1 > x]} dx < \infty$. Note that we assume the bootstrap weights are generated so that they are independent of the data, and that the draws θ_n generated for the parametric component are i.i.d. and independent of the bootstrap weights.

In Appendix B we give an outline of arguments one can use in verifying the regularity conditions in models such as our example models. We discuss verification of the regularity conditions for the biased sampling example there. Appendix B of DKL discusses verification of the regularity conditions for the odds-rate model as well. See Dixon (2003) for a detailed verification of the regularity conditions for the other example models.

DKL discuss the relevance of the above conditions in proving the main theorem. Here we discuss condition (PB4) only. Condition (PB4)(i) simply states that we have valid draws for the parametric component of the model. The tightness of the limit distribution allows us to apply Slutsky's Theorem (Example 1.4.7 of VW) in various parts of the proof. In Section 2 we discussed methods of obtaining such draws in practice. Condition (PB4)(ii) requires that the bootstrap is valid when θ is fixed at θ_0 . This should hold whenever the bootstrap is valid for the full ψ . Note that we require the limit in (PB4)(ii) to be tight and cadlag. This is useful in various parts of the proofs in light of Lemma 5.1 below.

LEMMA 5.1. *Suppose we have an expression of the form*

$$(5.2) \quad Q_n(h) = \sqrt{n} \int_0^\tau f_n(h)(u) dg_n(u)$$

where $f_n(h) = o_p^{t_2}(1)(h)$ has total variation bounded by some M over all n , and $\sqrt{n}g_n$ converges weakly to a tight, cadlag limit. Then $Q_n(h) = o_p^{t_2}(1)(h)$.

See DKL (2004) for a proof, which employs arguments from the proof of Lemma A.3 of Biliias *et al.* (1997).

For example we use Lemma 5.1 in the last step of the proof of Lemma A.4, taking $g_n = \hat{A}_{\theta_0} - A_0$ and

$$f_n(h)(u) = \mathbb{P}_n \hat{\sigma}_{2,2}[\theta_0, A_0 + t_n(h)(\hat{A}_{\theta_0} - A_0)](h)(u) - \sigma_{2,2}[\theta_0, A_0](h)(u).$$

A similar lemma is useful in going from the second to last to the last line of the proof of Lemma A.2. Note that we first observe $\hat{A}_{\hat{\theta}_n}^\circ - \hat{A}_{\hat{\theta}_n} = (\hat{A}_{\hat{\theta}_n}^\circ - A_0) - (\hat{A}_{\hat{\theta}_n} - A_0)$, and the integral on the second to last line is broken up accordingly. The integral with respect to $d[\sqrt{n}(\hat{A}_{\hat{\theta}_n} - A_0)]$ that results can be dealt with using Lemma 5.1 above. But

we do not necessarily have that $\sqrt{n}(\hat{A}_{\hat{\theta}_n}^\circ - A_0)$ converges weakly to a tight, cadlag limit. Condition (PB4)(iii) is a weaker condition than this, which turns out to be easier to verify in our examples. In this case we use the following lemma, taking $g_n = \hat{A}_{\hat{\theta}_n}^\circ - A_0$ and

$$f_n(h)(u) = \mathbb{P}_n^\circ \{ \hat{\sigma}_{2,2}[\hat{\theta}_n + t_n(h)(\theta_n - \hat{\theta}_n), \hat{A}_{\hat{\theta}_n} + t_n(h)(\hat{A}_{\hat{\theta}_n}^\circ - \hat{A}_{\hat{\theta}_n})](h)(u) \} \\ - \sigma_{2,2}[\theta_0, A_0](h)(u).$$

LEMMA 5.2. *Suppose we have an expression of the form (5.2) where $f_n(h) = o_p^{\mathcal{H}^2}(1)(h)$ and $\|\sqrt{n}g_n\|_2 \leq \tilde{O}_p(1) + o_p^{\mathcal{H}^2}(1 + \sqrt{n}\|\theta_n - \theta_0\|_1)$. Then $Q_n(h) = o_p^{\mathcal{H}^2}(1)(h)$.*

See DKL for a proof of this lemma, which is similar to the proof of Lemma 5.1.

6. Discussion

We have demonstrated that in semiparametric models where both the finite dimensional parameter θ and the infinite dimensional parameter A are \sqrt{n} consistent, it is possible to significantly decrease the dimensionality of the maximization required for Monte Carlo inference compared to the weighted bootstrap.

The proposed piggyback bootstrap algorithm is easy to implement. First, draws for the parametric component θ_n are generated with the same sampling distribution as the MLE $\hat{\theta}_n$. Under the regularity conditions, this can often be accomplished by estimating the covariance matrix of $\hat{\theta}_n$ and adding normal deviates with this covariance to $\hat{\theta}_n$. In some models, such as the Cox model, a theoretical estimate of the covariance is readily available. In other models, such as the biased sampling model in the simulation study, one can use Corollary 3 of Murphy and van der Vaart (2000) to accomplish the variance estimation. This involves computing a limit involving the profile likelihood evaluated at various values of θ . We gave details on this method in Section 4. An alternative is Lee's (2000) MCMC approach, which is based on other results from Murphy and van der Vaart (2000). Next, for each parametric draw θ_n , we piggyback the draw for A onto the draw for θ by performing the profile computation $A_n = \arg \max_A \ell_n^\circ(\theta_n, A)$. This computation can be accomplished using the fixed point algorithm given in Section 2. Our main result is that the resulting pair (θ_n, A_n) has the correct limit distribution for joint semiparametric inference, as discussed in Section 2.

The regularity conditions necessary for the method, given in Section 5, are not difficult to check in practice. Appendix B provides an outline for verifying the regularity conditions for models such as our example models.

For right censored survival data, the method applies to the Cox proportional hazards model, the proportional odds model, and the odds-rate model. In the case of clustered survival data, the procedure applies to the shared and correlated gamma frailty models. An application not arising in survival analysis is to biased sampling models, which arise in vaccine efficacy trials.

In the odds-rate model, the frailty W is assumed to be gamma distributed with mean 1 and variance γ . We are confident the method can be proved valid for frailties following other distributions with unknown parameters. Asymptotic theory for some such models has been worked out by Kosorok *et al.* (2004), including the validity of the weighted bootstrap. In addition, the method should be applicable to the more general proportional

hazards random effects regression models for clustered survival data considered in Vaida and Xu (2000). However, in the more complex cases, asymptotic theory has not yet been worked out. Perhaps the most complex case worked out thus far is the correlated gamma frailty model of Parner (1998).

While the precise choice of the distribution of $\{\eta_i\}$ in the weighted bootstrap has no effect asymptotically, the rate of convergence may be affected. Newton and Raftery (1994) discuss different choices in the context of parametric maximum likelihood. They demonstrate that unit exponential weights, which are Dirichlet after standardizing, perform well. Our own experience is that exponential weights also work well for semiparametric inference.

Our main theoretical result on the validity of the piggyback bootstrap for semiparametric inference was confirmed in the simulation study in Section 4. We saw that for simulated vaccine trial data from a biased sampling model, the piggyback bootstrap and the weighted bootstrap performed close to the same in semiparametric inference, with the piggyback bootstrap providing a dramatic improvement in computational efficiency.

Appendix A: Proof of main theorem

In this appendix we prove the main theorem and the intermediate lemmas.

PROOF OF THEOREM 2.1. This theorem follows from the expansion

$$(A.1) \quad \sqrt{n}(\hat{A}_{\hat{\theta}_n}^\circ(t) - \hat{A}_{\hat{\theta}_n}(t)) = -\sqrt{n}(\theta_n - \hat{\theta}_n)' \sigma_{2,1}[\theta_0, A_0][\sigma_{2,2}[\theta_0, A_0]^{-1}(\mathbf{I}\{u \leq t\})] + \sqrt{n}(\hat{A}_{\theta_0}^\circ(t) - \hat{A}_{\theta_0}(t)) + o_p^{[0, \tau]}(1)(t),$$

which we prove below. Let U_n have the distribution of $\sqrt{n}(\hat{\theta}_n - \theta_0)$ and let V_n have the distribution of $\sqrt{n}(\hat{A}_{\theta_0}^\circ(t) - A_0(t))$ and let U_n and V_n be independent. Since the bootstrap weights and the draws θ_n are independent given the data, (PB4) implies

$$\begin{aligned} & -\sqrt{n}(\theta_n - \hat{\theta}_n)' \sigma_{2,1}[\sigma_{2,2}^{-1}(\mathbf{I}\{u \leq t\})] + \sqrt{n}(\hat{A}_{\theta_0}^\circ(t) - \hat{A}_{\theta_0}(t)) \\ & \approx -U_n' \sigma_{2,1}[\sigma_{2,2}^{-1}(\mathbf{I}\{u \leq t\})] + V_n \\ & \rightsquigarrow -Z_\theta' \sigma_{2,1}[\sigma_{2,2}^{-1}(\mathbf{I}\{u \leq t\})] + Z_A(t), \end{aligned}$$

where Z_θ is mean zero Gaussian with covariance $(\sigma_{1,1} - \sigma_{1,2}\sigma_{2,2}^{-1}\sigma_{2,1})^{-1}$ and Z_A is mean zero Gaussian with covariance $\sigma_{2,2}^{-1}$, and Z_θ is independent of Z_A . Thus

$$\sqrt{n}(\theta_n - \hat{\theta}_n, \hat{A}_{\hat{\theta}_n}^\circ - \hat{A}_{\hat{\theta}_n}) \approx Z,$$

where $Z' \equiv (Z_\theta, Z_A - \sigma_{2,2}^{-1}\sigma_{1,2}Z_\theta)'$ is mean zero Gaussian with covariance

$$\begin{bmatrix} (\sigma_{1,1} - \sigma_{2,1}\sigma_{2,2}^{-1}\sigma_{1,2})^{-1} & -(\sigma_{1,1} - \sigma_{2,1}\sigma_{2,2}^{-1}\sigma_{1,2})^{-1}\sigma_{2,1}\sigma_{2,2}^{-1} \\ -\sigma_{2,2}^{-1}\sigma_{1,2}(\sigma_{1,1} - \sigma_{2,1}\sigma_{2,2}^{-1}\sigma_{1,2})^{-1} & \sigma_{2,2}^{-1} + \sigma_{2,2}^{-1}\sigma_{1,2}(\sigma_{1,1} - \sigma_{2,1}\sigma_{2,2}^{-1}\sigma_{1,2})^{-1}\sigma_{2,1}\sigma_{2,2}^{-1} \end{bmatrix}$$

which is the inverse of the information operator. \square

Before verifying expansion (A.1), we need the following preliminary lemmas in addition to Lemmas 5.1 and 5.2 given in Section 5.

LEMMA A.1. Under conditions (PB2) and (PB3), if $\tilde{\theta}_n \rightarrow_p \theta_0$ and $\tilde{A}_n \rightarrow_p A_0$, then

$$\mathbb{P}_n \hat{\sigma}_{2,j}[\tilde{\theta}_n, \tilde{A}_n](h) - \sigma_{2,j}[\theta_0, A_0](h) = o_p^{\mathcal{H}_2}(1)(h), \quad j = 1, 2.$$

This result also holds with \mathbb{P}_n° replacing \mathbb{P}_n .

PROOF. Fix $\epsilon > 0$ and observe that for $j = 1, 2$,

$$\begin{aligned} & \lim_{n \rightarrow \infty} P^*(\|\mathbb{P}_n \hat{\sigma}_{2,j}[\tilde{\theta}_n, \tilde{A}_n](h) - \sigma_{2,j}[\theta_0, A_0](h)\|_j^{\mathcal{H}_2} > \epsilon) \\ & \leq \lim_{m \rightarrow \infty} \lim_{n \rightarrow \infty} \left\{ P^* \left(\|\tilde{\theta}_n - \theta_0\|_1 > \frac{1}{m} \right) + P^* \left(\|\tilde{A}_n - A_0\|_2 > \frac{1}{m} \right) \right. \\ & \quad \left. + P^* \left(\|\mathbb{P}_n \hat{\sigma}_{2,j}[\tilde{\theta}_n, \tilde{A}_n](h) - \sigma_{2,j}[\theta_0, A_0](h)\|_j^{\mathcal{H}_2} > \epsilon, \right. \right. \\ & \quad \left. \left. \|\tilde{\theta}_n - \theta_0\|_1 \leq \frac{1}{m}, \|\tilde{A}_n - A_0\|_2 \leq \frac{1}{m} \right) \right\} \\ & \leq \lim_{m \rightarrow \infty} \lim_{n \rightarrow \infty} P^* \left(\|\mathbb{P}_n \hat{\sigma}_{2,j}[\tilde{\theta}_n, \tilde{A}_n](h) - \sigma_{2,j}[\tilde{\theta}_n, \tilde{A}_n](h)\|_j^{\mathcal{H}_2} \right. \\ & \quad \left. + \|\sigma_{2,j}[\tilde{\theta}_n, \tilde{A}_n](h) - \sigma_{2,j}[\theta_0, A_0](h)\|_j^{\mathcal{H}_2} > \epsilon, \right. \\ & \quad \left. \|\tilde{\theta}_n - \theta_0\|_1 \leq \frac{1}{m}, \|\tilde{A}_n - A_0\|_2 \leq \frac{1}{m} \right) \\ & \leq \lim_{m \rightarrow \infty} \lim_{n \rightarrow \infty} P^* \left(\|\mathbb{P}_n \hat{\sigma}_{2,j}[\tilde{\theta}_n, \tilde{A}_n](h) - \sigma_{2,j}[\tilde{\theta}_n, \tilde{A}_n](h)\|_j^{\mathcal{H}_2} > \tilde{\epsilon}, \right. \\ & \quad \left. \|\tilde{\theta}_n - \theta_0\|_1 \leq \frac{1}{m}, \|\tilde{A}_n - A_0\|_2 \leq \frac{1}{m} \right), \end{aligned}$$

for some $\tilde{\epsilon} > 0$ by condition (PB2). The last expression equals zero by condition (PB3) and the fact that a Donsker class is also Glivenko-Cantelli. For the case of \mathbb{P}_n° replacing \mathbb{P}_n , make the same replacement in the proof, and note that condition (PB3) implies $\{\eta \hat{\sigma}_{2,j}[\theta, A](h), \|\theta - \theta_0\| \leq c, \|A - A_0\| \leq c, h \in \mathcal{H}_2\}$ is Donsker and hence Glivenko-Cantelli for $j = 1, 2$. \square

LEMMA A.2.

$$\begin{aligned} \sqrt{n}(\mathbb{P}_n^\circ - \mathbb{P}_n)U_2(\hat{\theta}_n, \hat{A}_{\hat{\theta}_n})(h) &= \sqrt{n}(\theta_n - \hat{\theta}_n)' \sigma_{2,1}[\theta_0, A_0](h) \\ & \quad + \sqrt{n} \int_0^\tau \sigma_{2,2}[\theta_0, A_0](h)(u) d(\hat{A}_{\hat{\theta}_n}^\circ - \hat{A}_{\hat{\theta}_n})(u) \\ & \quad + o_p^{\mathcal{H}_2}(1)(h). \end{aligned}$$

PROOF. By definition of \hat{A}_θ° and $\hat{A}_{\hat{\theta}_n}$ we have

$$0 = (\mathbb{P}_n^\circ - \mathbb{P}_n)U_2(\hat{\theta}_n, \hat{A}_{\hat{\theta}_n})(h) + \mathbb{P}_n^\circ(U_2(\theta_n, \hat{A}_{\hat{\theta}_n}^\circ)(h) - U_2(\hat{\theta}_n, \hat{A}_{\hat{\theta}_n})(h)).$$

Using 2.4.8 of Abraham, Marsden, and Ratiu (1988) (hereafter abbreviated AMR) to expand the rightmost term of the preceding equation, under condition (PB1), we obtain

$$(\mathbb{P}_n^\circ - \mathbb{P}_n)U_2(\hat{\theta}_n, \hat{A}_{\hat{\theta}_n})(h)$$

$$\begin{aligned}
 &= -\mathbb{P}_n^\circ \left\{ \frac{\partial}{\partial s} U_2[\hat{\theta}_n + s(\theta_n - \hat{\theta}_n), \hat{A}_{\hat{\theta}_n} + s(\hat{A}_{\hat{\theta}_n}^\circ - \hat{A}_{\hat{\theta}_n})](h) \Big|_{s=t_n(h)} \right\} \\
 &= -\mathbb{P}_n^\circ \left\{ \frac{\partial}{\partial s} U_2[\hat{\theta}_n + t_n(h)(\theta_n - \hat{\theta}_n) + s(\theta_n - \hat{\theta}_n), \right. \\
 &\quad \left. \hat{A}_{\hat{\theta}_n} + t_n(h)(\hat{A}_{\hat{\theta}_n}^\circ - \hat{A}_{\hat{\theta}_n}) + s(\hat{A}_{\hat{\theta}_n}^\circ - \hat{A}_{\hat{\theta}_n})](h) \Big|_{s=0} \right\} \\
 &= \mathbb{P}_n^\circ \{ (\theta_n - \hat{\theta}_n)' \hat{\sigma}_{2,1}[\hat{\theta}_n + t_n(h)(\theta_n - \hat{\theta}_n), \hat{A}_{\hat{\theta}_n} + t_n(h)(\hat{A}_{\hat{\theta}_n}^\circ - \hat{A}_{\hat{\theta}_n})](h) \} \\
 &\quad + \mathbb{P}_n^\circ \left\{ \int_0^\tau \hat{\sigma}_{2,2}[\hat{\theta}_n + t_n(h)(\theta_n - \hat{\theta}_n), \right. \\
 &\quad \left. \hat{A}_{\hat{\theta}_n} + t_n(h)(\hat{A}_{\hat{\theta}_n}^\circ - \hat{A}_{\hat{\theta}_n})](h)(u) d(\hat{A}_{\hat{\theta}_n}^\circ - \hat{A}_{\hat{\theta}_n})(u) \right\},
 \end{aligned}$$

for some $t_n(h) \in [0, 1]$. Note that

$$\sigma_{2,j}[\theta_0, A_0](h) = \mathbb{P}_n^\circ \{ \hat{\sigma}_{2,j}[\hat{\theta}_n + t_n(h)(\theta_n - \hat{\theta}_n), \hat{A}_{\hat{\theta}_n} + t_n(h)(\hat{A}_{\hat{\theta}_n}^\circ - \hat{A}_{\hat{\theta}_n})](h) \} + o_p^{\mathcal{H}^2}(1)(h)$$

by Lemma A.1 and conditions (PB2), (PB3) and (PB4). This, condition (PB4), and Lemmas 5.1 and 5.2 imply,

$$\begin{aligned}
 &\sqrt{n}(\theta_n - \hat{\theta}_n)' \sigma_{2,1}(h) \\
 &= \sqrt{n} \mathbb{P}_n^\circ \{ (\theta_n - \hat{\theta}_n)' \hat{\sigma}_{2,1}[\hat{\theta}_n + t_n(h)(\theta_n - \hat{\theta}_n), \hat{A}_{\hat{\theta}_n} + t_n(h)(\hat{A}_{\hat{\theta}_n}^\circ - \hat{A}_{\hat{\theta}_n})](h) \} \\
 &\quad + o_p^{\mathcal{H}^2}(1)(h),
 \end{aligned}$$

and

$$\begin{aligned}
 &\sqrt{n} \mathbb{P}_n^\circ \left\{ \int_0^\tau \hat{\sigma}_{2,2}[\hat{\theta}_n + t_n(h)(\theta_n - \hat{\theta}_n), \right. \\
 &\quad \left. \hat{A}_{\hat{\theta}_n} + t_n(h)(\hat{A}_{\hat{\theta}_n}^\circ - \hat{A}_{\hat{\theta}_n})](h)(u) d(\hat{A}_{\hat{\theta}_n}^\circ - \hat{A}_{\hat{\theta}_n})(u) \right\} \\
 &= \int_0^\tau \mathbb{P}_n^\circ \{ \hat{\sigma}_{2,2}[\hat{\theta}_n + t_n(h)(\theta_n - \hat{\theta}_n), \\
 &\quad \hat{A}_{\hat{\theta}_n} + t_n(h)(\hat{A}_{\hat{\theta}_n}^\circ - \hat{A}_{\hat{\theta}_n})](h)(u) \} d[\sqrt{n}(\hat{A}_{\hat{\theta}_n}^\circ - \hat{A}_{\hat{\theta}_n})](u) \\
 &= \sqrt{n} \int_0^\tau \sigma_{2,2}[\theta_0, A_0](h)(u) d(\hat{A}_{\hat{\theta}_n}^\circ - \hat{A}_{\hat{\theta}_n})(u) + o_p^{\mathcal{H}^2}(1)(h). \quad \square
 \end{aligned}$$

LEMMA A.3.

$$\sqrt{n}(\mathbb{P}_n^\circ - \mathbb{P}_n)U_2(\hat{\theta}_n, \hat{A}_{\hat{\theta}_n})(h) \approx \sqrt{n}(\mathbb{P}_n - P_0)U_2(\theta_0, A_0)(h).$$

PROOF. If we show that

$$\sqrt{n}(\mathbb{P}_n^\circ - \mathbb{P}_n)U_2(\hat{\theta}_n, \hat{A}_{\hat{\theta}_n})(h) = \sqrt{n}(\mathbb{P}_n^\circ - \mathbb{P}_n)U_2(\theta_0, A_0)(h) + o_p^{\mathcal{H}^2}(1)(h),$$

then the conclusion follows as a consequence of Theorem 2.9.6 of VW. To see this, note that by Slutsky's Theorem (Example 1.4.7 of VW) and condition (PB5),

$$\sqrt{n}(\mathbb{P}_n^\circ - \mathbb{P}_n)U_2(\theta_0, A_0)(h) = \frac{\sqrt{n}\mathbb{P}_n\eta U_2(\theta_0, A_0)(h)}{\mathbb{P}_n\eta} - \sqrt{n}\mathbb{P}_n U_2(\theta_0, A_0)(h)$$

$$\begin{aligned}
 &= \sqrt{n}\mathbb{P}_n\eta U_2(\theta_0, A_0)(h) - \sqrt{n}\mathbb{P}_n U_2(\theta_0, A_0)(h) + o_p^{\mathcal{H}^2}(1)(h) \\
 &= \sqrt{n}\mathbb{P}_n(\eta - 1)(U_2(\theta_0, A_0)(h) - P_0 U_2(\theta_0, A_0)(h)) \\
 &\quad + o_p^{\mathcal{H}^2}(1)(h) \\
 &\approx \sqrt{n}(\mathbb{P}_n - P_0)U_2(\theta_0, A_0)(h).
 \end{aligned}$$

The last line follows by the aforementioned theorem from VW. Thus all that remains is to prove the first display of the proof.

To establish the first display of the proof, note that 2.4.8 of AMR reveals

$$\begin{aligned}
 &\sqrt{n}(\mathbb{P}_n^\circ - \mathbb{P}_n)U_2(\hat{\theta}_n, \hat{A}_{\hat{\theta}_n})(h) - \sqrt{n}(\mathbb{P}_n^\circ - \mathbb{P}_n)U_2(\theta_0, A_0)(h) \\
 &= -\sqrt{n}(\mathbb{P}_n^\circ - \mathbb{P}_n)\{(\hat{\theta}_n - \theta_0)' \hat{\sigma}_{2,1}[\theta_0 + t_n(h)(\hat{\theta}_n - \theta_0), A_0 + t_n(h)(\hat{A}_{\hat{\theta}_n} - A_0)](h)\} \\
 &\quad - \sqrt{n}(\mathbb{P}_n^\circ - \mathbb{P}_n)\left\{\int_0^\tau \hat{\sigma}_{2,2}[\theta_0 + t_n(h)(\hat{\theta}_n - \theta_0), \right. \\
 &\qquad\qquad\qquad \left. A_0 + t_n(h)(\hat{A}_{\hat{\theta}_n} - A_0)](h)(u)d(\hat{A}_{\hat{\theta}_n} - A_0)(u)\right\},
 \end{aligned}$$

for some $t_n(h) \in [0, 1]$. Note that $(\mathbb{P}_n^\circ - \mathbb{P}_n)\hat{\sigma}_{2,j}[\theta_0 + t_n(h)(\hat{\theta}_n - \theta_0), A_0 + t_n(h)(\hat{A}_{\hat{\theta}_n} - A_0)](h) = o_p^{\mathcal{H}^2}(1)(h)$ for $j = 1, 2$ by conditions (PB2), (PB3) and (PB4), and Lemma A.1. This, Lemma 5.1, and condition (PB4) completes the proof of the first display, and hence, the lemma. \square

LEMMA A.4.

$$\sqrt{n}(\mathbb{P}_n - P_0)U_2(\theta_0, A_0)(h) = \sqrt{n} \int_0^\tau \sigma_{2,2}[\theta_0, A_0](h)(u)d(\hat{A}_{\theta_0} - A_0)(u) + o_p^{\mathcal{H}^2}(1)(h).$$

PROOF. By condition (PB5), the definition of \hat{A}_{θ_0} , and 2.4.8 of AMR,

$$\begin{aligned}
 &\sqrt{n}(\mathbb{P}_n - P_0)U_2(\theta_0, A_0)(h) \\
 &= \sqrt{n}\mathbb{P}_n(U_2(\theta_0, A_0)(h) - U_2(\theta_0, \hat{A}_{\theta_0})(h)) \\
 &= \sqrt{n}\mathbb{P}_n\left\{\int_0^\tau \hat{\sigma}_{2,2}[\theta_0, A_0 + t_n(h)(\hat{A}_{\theta_0} - A_0)](h)(u)d(\hat{A}_{\theta_0} - A_0)(u)\right\} + o_p^{\mathcal{H}^2}(1)(h),
 \end{aligned}$$

for some $t_n(h) \in [0, 1]$. Note that

$$\mathbb{P}_n\hat{\sigma}_{2,2}[\theta_0, A_0 + t_n(h)(\hat{A}_{\theta_0} - A_0)](h) = \sigma_{2,2}[\theta_0, A_0](h) + o_p^{\mathcal{H}^2}(1)(h)$$

by conditions (PB2), (PB3) and (PB4), and Lemma A.1. This, Lemma 5.1, and condition (PB4) completes the proof. \square

PROOF OF (A.1). Note that by condition (PB4),

$$\sqrt{n} \int_0^\tau \sigma_{2,2}[\theta_0, A_0](h)(u)d(\hat{A}_{\theta_0}^\circ - \hat{A}_{\theta_0})(u) \approx \sqrt{n} \int_0^\tau \sigma_{2,2}[\theta_0, A_0](h)(u)d(\hat{A}_{\theta_0} - A_0)(u).$$

This and the conclusions of Lemmas A.2, A.3, and A.4 give

$$\begin{aligned} & \sqrt{n} \int_0^\tau \sigma_{2,2}[\theta_0, A_0](h)(u) d(\hat{A}_{\theta_0}^\circ - \hat{A}_{\theta_0})(u) \\ &= \sqrt{n}(\theta_n - \hat{\theta}_n)' \sigma_{2,1}[\theta_0, A_0](h) \\ & \quad + \sqrt{n} \int_0^\tau \sigma_{2,2}[\theta_0, A_0](h)(u) d(\hat{A}_{\theta_n}^\circ - \hat{A}_{\hat{\theta}_n})(u) + o_p^{\mathcal{H}_2}(1)(h). \end{aligned}$$

Since \mathcal{H}_2 includes all functions of the form $h(u) = \sigma_{2,2}[\theta_0, A_0]^{-1}(\mathbf{I}\{u \leq t\})$, for all $t \in [0, \tau]$, we have established (A.1), and the proof is complete. \square

Appendix B: Verifying the regularity conditions for examples

In this appendix we verify that the regularity conditions hold for the biased sampling model, adding some technical restrictions as necessary. See DKL for the odds-rate model, and Dixon (2003) for a verification of the regularity conditions and the additional technical restrictions for the other example models. We start with a general outline that the arguments follow in all of the examples, and which should serve as a guide to verifying the regularity conditions in other applications.

As our approach is an alternative to the weighted bootstrap, we note that the validity of the weighted bootstrap for the Cox proportional hazards model, the proportional odds model, and the odds-rate model follows from Kosorok *et al.* (2004). For the shared and correlated gamma frailty models, the validity of the weighted bootstrap which assigns weight to clusters (as described in Section 3) follows by similar arguments. And Gilbert (1996) establishes the validity of the weighted bootstrap for the biased sampling example.

B.1 General outline

Verification of the regularity conditions follows the same outline in each of our examples. The same outline should be applicable to other semiparametric models.

Conditions (PB1) and (PB2) can be readily verified by inspection of the score and information operators, which we give for the biased sampling model below. Condition (PB6) can be verified in our examples using the Z-Estimator master theorem (Theorem 3.3.1 of VW) and Convolution Theorems (Theorems 5.1–5.3 of Bickel *et al.* (1993)). The fact that the information operator is continuously invertible and onto is established using a standard result from functional analysis (see Theorem 3.4 and Corollary 3.8 of Kress (1989)).

Verification of conditions (PB3) and (PB5) is routine in our examples using Donsker theorems from VW. As an example, consider the piece $H_\psi^X(f) \equiv \int_0^t f(s)e^{\theta'Z(s)}dA(s)$ which appears in many of our survival analysis examples. We have the following lemma.

LEMMA B.1. *For each fixed $p < \infty$ and each fixed $\psi \in \Theta_q \times BV_r$ with $q, r < \infty$, the class $\{H_\psi^T(f) : f \in BV_p\}$ is P_0 -Donsker. Here $\Theta_q \equiv \{\theta : \|\theta\| \leq q\}$.*

PROOF. Classes of cells and uniformly bounded classes of functions of bounded variation are standard examples of Donsker classes. Thus $Y(\cdot)(\equiv \mathbf{I}\{X \geq \cdot\})$, $\beta'Z(\cdot)$, and $f(\cdot)$ are uniformly bounded Donsker classes on $[0, \tau]$. The exponential function is Lipschitz on compact sets, therefore the Lipschitz Transformation Donsker Permanence Theorem (Theorem 2.10.6 of VW) gives that $\exp(\beta'Z(\cdot))$ is Donsker on $[0, \tau]$. Since

the product of two Lipschitz functions is Lipschitz on compact sets, another application of this theorem gives that $f(\cdot)e^{\beta'Z(\cdot)}Y(\cdot)$ is Donsker. Thus the claim follows from the continuous mapping theorem. \square

Condition (PB4)(ii) can be verified in our examples by the following argument, which is similar to the proof of the validity of the weighted bootstrap for proportional hazards frailty regression models given in Kosorok *et al.* (2004).

LEMMA B.2 *Condition (PB4)(ii) holds in our examples.*

Sketch of Proof. Define $\hat{\psi}_\theta^\circ \equiv (\theta, \hat{A}_\theta^\circ)$ and $\hat{\psi}_\theta \equiv (\theta, \hat{A}_\theta)$. Applying the Z-Estimator master theorem of VW gives that $\sqrt{n}(\hat{\psi}_{\theta_0}^\circ - \psi_0) = \sqrt{n}\mathbb{P}_n^\circ U^\tau(\psi_0)(\sigma^{-1}(\cdot)) + o_p(1)$ unconditionally, and $\sqrt{n}(\hat{\psi}_{\theta_0} - \psi_0) = \sqrt{n}\mathbb{P}_n U^\tau(\psi_0)(\sigma^{-1}(\cdot)) + o_p(1)$, where $o_p(1)$ denotes a quantity $\rightarrow 0$ uniformly in outer probability. Thus $\sqrt{n}(\hat{\psi}_{\theta_0}^\circ - \hat{\psi}_{\theta_0}) = \sqrt{n}(\mathbb{P}_n^\circ - \mathbb{P}_n)U^\tau(\psi_0)(\sigma^{-1}(\cdot)) + o_p(1)$ unconditionally. Finally, note that

$$\begin{aligned} \sqrt{n}(\mathbb{P}_n^\circ - \mathbb{P}_n)U^\tau(\psi_0)(h) &= \frac{\sqrt{n}\mathbb{P}_n^\circ \eta U^\tau(\psi_0)(h)}{\mathbb{P}_n^\circ \eta} - \sqrt{n}\mathbb{P}_n U^\tau(\psi_0)(h) \\ &= \sqrt{n}\mathbb{P}_n^\circ \eta U^\tau(\psi_0)(h) - \sqrt{n}\mathbb{P}_n U^\tau(\psi_0)(h) + o_p^{\mathcal{H}^2}(1)(h) \\ &= \sqrt{n}\mathbb{P}_n(\eta - 1)(U^\tau(\psi_0) - P_0 U^\tau(\psi_0))(h) + o_p^{\mathcal{H}^2}(1)(h) \\ &\approx \sqrt{n}(\mathbb{P}_n - P_0)U^\tau(\psi_0)(h), \end{aligned}$$

where the last relation follows by the Multiplier Central Limit Theorem (Theorem 2.9.6 of VW). \square

Condition (PB4)(iii) can be verified in our examples by a similar argument to one used in the proof of Theorem 3.4 of Lee (2000), which we give in the proof of the following lemma:

LEMMA B.3 *Condition (PB4)(iii) holds in our examples.*

Sketch of Proof. Define

$$\tilde{D}_n(A)(h) \equiv \mathbb{P}_n^\circ U^\tau \begin{pmatrix} \theta_n \\ A \end{pmatrix} \begin{pmatrix} 0 \\ h \end{pmatrix}, \quad D_n(A)(h) \equiv \mathbb{P}_n^\circ U^\tau \begin{pmatrix} \theta_0 \\ A \end{pmatrix} \begin{pmatrix} 0 \\ h \end{pmatrix},$$

and

$$D_0(A)(h) \equiv P_0 U^\tau \begin{pmatrix} \theta_0 \\ A \end{pmatrix} \begin{pmatrix} 0 \\ h \end{pmatrix}.$$

Note that $\tilde{D}_n(\hat{A}_{\theta_n}^\circ) = 0$ and $D_0(A_0) = 0$. We can use Lemma 3.3.5 of VW to obtain

$$\sqrt{n}(\tilde{D}_n - D_0)(\hat{A}_{\theta_n}^\circ) - \sqrt{n}(\tilde{D}_n - D_0)(A_0) = o_p^{\mathcal{H}^2}(1 + \sqrt{n}\|\hat{\psi}_n^\circ - \psi_0\|_3)(h),$$

and

$$\begin{aligned} \sqrt{n}(\tilde{D}_n - D_n)(A_0) &= \sqrt{n}(\tilde{D}_n - D_0)(A_0) - \sqrt{n}(D_n - D_0)(A_0) \\ &= o_p^{\mathcal{H}^2}(1 + \sqrt{n}\|\theta_n - \theta_0\|_1)(h), \end{aligned}$$

from which we have

$$\begin{aligned} \sqrt{n}(D_0(\hat{A}_{\theta_n}^\circ) - D_0(A_0)) &= \sqrt{n}(D_0(\hat{A}_{\theta_n}^\circ) - \tilde{D}_n(\hat{A}_{\theta_n}^\circ)) \\ &= -\sqrt{n}(\tilde{D}_n - D_0)(A_0) + o_p^{\mathcal{H}^2}(1 + \sqrt{n}\|\hat{\psi}_n^\circ - \psi_0\|_3)(h) \\ &= -\sqrt{n}(D_n - D_0)(A_0) + o_p^{\mathcal{H}^2}(1 + \sqrt{n}\|\hat{\psi}_n^\circ - \psi_0\|_3)(h) \\ &= \tilde{O}_p(1) + o_p^{\mathcal{H}^2}(1 + \sqrt{n}\|\hat{\psi}_n^\circ - \psi_0\|_3)(h). \end{aligned}$$

The last equality in the preceding display follows from the Z-Estimator master theorem of VW. Since σ is continuously invertible, we have for some $c > 0$ that

$$\|D_0(A) - D_0(A_0)\| \geq c\|A - A_0\|_2 + o(\|A - A_0\|_2).$$

Thus, $\sqrt{n}\|\hat{A}_{\theta_n}^\circ - A_0\|_2(c + o_p^{\mathcal{H}^2}(1)(h)) \leq \tilde{O}_p(1) + o_p^{\mathcal{H}^2}(1 + \sqrt{n}\|\theta_n - \theta_0\|_1)(h)$. Multiplying each side by $(c + o_p^{\mathcal{H}^2}(1)(h))^{-1}$ and applying Slutsky's Theorem (Example 1.4.7 of VW), we have the desired result. \square

Draws θ_n satisfying condition (PB4)(i) can be generated in our examples by applying results from Murphy and van der Vaart (2000).

LEMMA B.4. *The hypotheses of Corollary 3 of Murphy and van der Vaart (2000) hold, and so we can obtain draws for the parametric component of the model satisfying (PB4)(i).*

Sketch of Proof. Based on the discussion in Chapter 25 of van der Vaart (2000) an approximately least favorable submodel for estimating θ in the presence of A is

$$dA_t(\theta, A) = (1 + (\theta - t)^\top \nu_0)dA,$$

where $\nu_0 : \mathbb{R} \mapsto \mathbb{R}^d$ is the least favorable direction at (θ_0, A_0) defined by

$$h^\top \nu_0 = \sigma_{2,2}^{-1} \sigma_{1,2} h, \quad h \in \mathbb{R}^d.$$

Let $\tilde{\theta}_n \rightarrow \theta_0$ almost surely under P_0 . It follows from the definition of \hat{A}_θ that $L_n(\tilde{\theta}_n, \hat{A}_{\tilde{\theta}_n}) \geq L_n(\theta_n, \hat{A}_{\theta_n})$. Therefore, noting that $(\tilde{\theta}_n, \hat{A}_{\tilde{\theta}_n})$ is consistent for (θ_0, A_0) , the method for proving consistency of the MLE in Theorem 3.2 of Lee (2000) (see also Theorem 1 of Parner (1998) or Theorem 3 of Kosorok *et al.* (2004)) can be adapted in a straightforward manner to show that $\hat{A}_{\tilde{\theta}_n} \rightarrow A_0$ almost surely under P_0 . Upon substituting $\theta = t$ and $A = A_t(\theta, A)$ into the log-likelihood function, it is seen that the path is smooth in t and is continuously differentiable. Arguments from the proof of Lemma B.1 give that the Glivenko-Cantelli and Donsker conditions in Murphy and van der Vaart (2000) are satisfied. All that remains to be established is that their equation (11) holds. Murphy and van der Vaart note that this equation holds whenever

$$\|\hat{A}_{\tilde{\theta}_n} - A_0\| = O_P(\|\tilde{\theta}_n - \theta_0\|) + o_P(n^{-1/2}).$$

This is certainly implied by

$$\sqrt{n}\|\hat{A}_{\tilde{\theta}_n} - A_0\| \leq O_P(1) + o_P(1 + \sqrt{n}\|\tilde{\theta}_n - \theta_0\|),$$

which can be proven in a similar manner to Lemma B.3. \square

B.2 *A biased sampling model*

Observe that for the biased sampling models

$$\begin{aligned}
 U_1[\theta, A](h_1) &= h_1 \left(\frac{\dot{w}_i(y, \theta)}{w_i(y, \theta)} - \frac{\int \dot{w}_i(z, \theta) dA(z)}{\int w_i(z, \theta) dA(z)} \right), \\
 U_2[\theta, A](h_2) &= h_2(y) - \frac{\int h_2(z) w_i(z, \theta) dA(z)}{\int w_i(z, \theta) dA(z)}, \\
 \hat{\sigma}_{1,1}[\theta, A](h_1) &= -h_1 \left(\frac{\ddot{w}_i(y, \theta) w_i(y, \theta) - (\dot{w}_i(y, \theta))^2}{(w_i(y, \theta))^2} \right. \\
 &\quad \left. - \frac{\int \ddot{w}_i(z, \theta) dA(z) \int w_i(z, \theta) dA(z) - (\int \dot{w}_i(z, \theta) dA(z))^2}{(\int w_i(z, \theta) dA(z))^2} \right), \\
 \hat{\sigma}_{1,2}[\theta, A](h_1)(v) &= h_1 \left(\frac{\dot{w}_i(v, \theta) \int w_i(z, \theta) dA(z) - w_i(v, \theta) \int \dot{w}_i(z, \theta) dA(z)}{(\int w_i(z, \theta) dA(z))^2} \right), \\
 \hat{\sigma}_{2,1}[\theta, A](h_2) &= \frac{\int h_2(z) \dot{w}_i(z, \theta) dA(z) \int w_i(z, \theta) dA(z)}{(\int w_i(z, \theta) dA(z))^2} \\
 &\quad - \frac{\int h_2(z) w_i(z, \theta) dA(z) \int \dot{w}_i(z, \theta) dA(z)}{(\int w_i(z, \theta) dA(z))^2},
 \end{aligned}$$

and

$$\hat{\sigma}_{2,2}[\theta, A](h_2)(v) = \frac{h_2(v) w_i(v, \theta) \int w_i(z, \theta) dA(z) - w_i(v, \theta) \int h_2(z) w_i(z, \theta) dA(z)}{(\int w_i(z, \theta) dA(z))^2}.$$

For the biased sampling model examined in the numerical studies, we used $w_1(z, \theta) = e^{z\theta}$ and $w_2(z, \theta) = 1$, and we assume θ lies in a known compact subset of \mathbb{R} . Then conditions (PB1) and (PB2) are readily verified. Conditions (PB3) and (PB5) hold by arguments similar to the proof of Lemma B.1 and the proof of Theorem 5.8 of Gilbert (1996). The mean zero assertion of condition (PB5) follows upon taking expectations in the above expression for U_2 . Conditions (PB4)(i), (PB4)(ii) and (PB4)(iii) hold by the arguments of Lemmas B.4, B.2 and B.3. Here the necessary identifiability for using arguments from the proof of Theorem 3.2 of Lee (2000), as described in the sketch of the proof of Lemma B.4, follows from Theorem 4.4 of Gilbert (1996). Theorem 5.11 of Gilbert (1996) proves the validity of the bootstrap for bootstrap weights that satisfy our aforementioned conditions. Finally, condition (PB6) follows from Theorem 5.8 and the discussion on pp. 106–107 of Gilbert (1996). Note that Gilbert takes \mathcal{H}_2 to be the set of functions of bounded supremum norm. Since this contains the set of functions of bounded total variation norm, the proofs in Gilbert immediately imply that our regularity conditions hold, except one must establish that the information operator is continuously invertible and onto for this smaller space of functions. This holds by the same arguments Gilbert uses for the larger class of functions.

REFERENCES

- Abraham, R., Marsden, J. E. and Ratiu, T. (1988). *Manifolds, Tensor Analysis, and Applications*, Springer-Verlag, New York.
- Bickel, P. J., Klaassen, C. A. J., Ritov, Y. and Wellner, J. A. (1993). *Efficient and Adaptive Estimation for Semiparametric Models*, Johns Hopkins University Press, Baltimore, Maryland.
- Bilias, Y., Gu, M. and Ying, Z. (1997). Towards a general asymptotic theory for Cox model with staggered entry, *Annals of Statistics*, **25**, 662–682.

- Dabrowska, D. M. and Doksum, K. A. (1988). Estimation and testing in a two-sample generalized odds-rate model, *Journal of the American Statistical Association*, **83**, 744–749.
- Dixon, J. R. (2003). The piggyback bootstrap for functional inference in semiparametric models, Ph.D. Dissertation, Department of Statistics, University of Wisconsin-Madison.
- Dixon, J. R., Kosorok, M. R. and Lee, B. L. (2004). Technical report to accompany ‘Functional inference in semiparametric models using the piggyback bootstrap’, Technical Report, Department of Statistics, Florida State University.
- Gilbert, P. B. (1996). Sieve analysis: Statistical methods for assessing differential vaccine protection against HIV types, Ph.D. Dissertation, Department of Biostatistics, University of Washington.
- Gilbert, P. B. (2000). Large sample theory of maximum likelihood estimates in semiparametric biased sampling models, *Annals of Statistics*, **28**, 151–194.
- Gilbert, P. B., Self, S. G. and Ashby, M. A. (1998). Statistical methods for assessing differential vaccine protection against human immunodeficiency virus types, *Biometrics*, **54**, 799–814.
- Kim, Y. and Lee, J. (2003). Bayesian bootstrap for proportional hazards models, *Annals of Statistics*, **31**, 1905–1922.
- Kosorok, M. R., Lee, B. L. and Fine, J. P. (2004). Robust inference for proportional hazards univariate frailty regression models, *Annals of Statistics*, **32**, 1448–1491.
- Kress, R. (1989). *Linear Integral Equations*, Springer-Verlag, Berlin.
- Lee, B. L. (2000). Efficient semiparametric estimation using Markov chain Monte Carlo, Ph.D. Dissertation, Department of Statistics, University of Wisconsin-Madison.
- Lin, D. Y., Fleming, T. R. and Wei, L. J. (1994). Confidence bands for survival curves under the proportional hazards model, *Biometrika*, **81**, 73–81.
- Murphy, S. A. (1994). Consistency in a proportional hazards model incorporating a random effect, *Annals of Statistics*, **22**, 712–731.
- Murphy, S. A. and van der Vaart, A. W. (2000). On profile likelihood, *Journal of the American Statistical Association*, **95**, 449–465.
- Newton, M. A. and Raftery, A. E. (1994). Approximate Bayesian inference with the weighted likelihood bootstrap, *Journal of the Royal Statistical Society, Series B*, **56**, 3–26.
- Parner, E. (1998). Asymptotic theory for the correlated gamma-frailty model, *Annals of Statistics*, **26**, 183–214.
- Praetgaard, J. and Wellner, J. A. (1993). Exchangeably weighted bootstraps of the general empirical process, *Annals of Probability*, **21**, 2053–2086.
- Press, W., Flannery, B., Teukolsky, S. and Vetterling, W. (1994). *Numerical Recipes in Pascal: The Art of Scientific Computing*, Cambridge University Press, New York.
- Rubin, D. B. (1981). The Bayesian bootstrap, *Annals of Statistics*, **9**, 130–134.
- Scharfstein, D. O., Tsiatis, A. A. and Gilbert, P. B. (1998). Semiparametric efficient estimation in the generalized odds-rate class of regression models for right-censored time-to-event data, *Lifetime Data Analysis*, **4**, 355–391.
- Shen, X. (2002). Asymptotic normality of semiparametric and nonparametric posterior distributions, *Journal of the American Statistical Association*, **97**, 222–235.
- Tsodikov, A. (2003). Semiparametric models: A generalized self-consistency approach, *Journal of the Royal Statistical Society, Series B*, **65**, 759–774.
- van der Vaart, A. W. (2000). *Asymptotic Statistics*, Cambridge University Press, New York.
- van der Vaart, A. W. and Wellner, J. A. (1996). *Weak Convergence and Empirical Processes: With Applications to Statistics*, Springer, New York.
- Vaida, F. and Xu, R. (2000). Proportional hazards model with random effects, *Statistics in Medicine*, **19**, 3309–3324.
- Vardi, Y. (1985). Empirical distributions in selection bias models, *Annals of Statistics*, **13**, 178–203.