

Sandwich Boosting for Accurate Estimation in Partially Linear Models for Grouped Data

Elliot H. Young ^{*1} and Rajen D. Shah ^{†1}

¹Statistical Laboratory, University of Cambridge, UK

March 19, 2024

Abstract

We study partially linear models in settings where observations are arranged in independent groups but may exhibit within-group dependence. Existing approaches estimate linear model parameters through weighted least squares, with optimal weights (given by the inverse covariance of the response, conditional on the covariates) typically estimated by maximising a (restricted) likelihood from random effects modelling or by using generalised estimating equations. We introduce a new ‘sandwich loss’ whose population minimiser coincides with the weights of these approaches when the parametric forms for the conditional covariance are well-specified, but can yield arbitrarily large improvements in linear parameter estimation accuracy when they are not. Under relatively mild conditions, our estimated coefficients are asymptotically Gaussian and enjoy minimal variance among estimators with weights restricted to a given class of functions, when user-chosen regression methods are used to estimate nuisance functions. We further expand the class of functional forms for the weights that may be fitted beyond parametric models by leveraging the flexibility of modern machine learning methods within a new gradient boosting scheme for minimising the sandwich loss. We demonstrate the effectiveness of both the sandwich loss and what we call ‘sandwich boosting’ in a variety of settings with simulated and real-world data.

1 Introduction

Grouped data are commonplace in many scientific, econometric and sociological disciplines. Prime examples include: repeated measures data (e.g. multiple readings of patient data), longitudinal data (e.g. where weekly sales are recorded across multiple stores), and hierarchical data (e.g. educational datasets clustered by school and potentially further sub-clustered by classroom).

To fix ideas, consider a regression setting where we have available grouped data $(Y_i, D_i, X_i) \in \mathbb{R}^{n_i} \times \mathbb{R}^{n_i} \times \mathbb{R}^{n_i \times d}$ for $i = 1, \dots, I$, with Y_i a vector of responses and predictors (D_i, X_i) separated out into a covariate D_i , whose contribution to the response we are particularly interested in, and remaining covariates X_i ; for instance D_i may be a treatment whose effect we wish to assess after controlling for additional covariates X_i . In total therefore we have $N := \sum_{i=1}^I n_i$ observations,

^{*}EHY was supported by the EPSRC Doctoral Training Partnership.

[†]RDS was supported by an EPSRC Programme Grant EP/N031938/1. The authors have no conflicts of interest.

though typically not all independent. A simple but popular approach to modelling this data in practice is via a linear model of the form

$$Y_i = \beta D_i + X_i \gamma + \varepsilon_i. \quad (1)$$

Here $\varepsilon_i \in \mathbb{R}^{n_i}$ is a vector of errors such that $\mathbb{E}[\varepsilon_i | D_i, X_i] = 0$ and $(\beta, \gamma) \in \mathbb{R} \times \mathbb{R}^d$ are regression coefficients to be estimated, with β our primary target of inference.

A challenge in such settings is properly accounting for potential correlations between components of ε_i in order to obtain accurate estimates of the parameters. This may be achieved through a weighted least squares regression yielding estimates

$$\begin{pmatrix} \hat{\beta} \\ \hat{\gamma} \end{pmatrix} := M^{-1} \begin{pmatrix} \sum_i D_i^\top \hat{W}_i Y_i \\ \sum_i X_i^\top \hat{W}_i Y_i \end{pmatrix}, \quad \text{where} \quad M := \begin{pmatrix} \sum_i D_i^\top \hat{W}_i D_i & \sum_i D_i^\top \hat{W}_i X_i \\ \sum_i X_i^\top \hat{W}_i D_i & \sum_i X_i^\top \hat{W}_i X_i \end{pmatrix}, \quad (2)$$

in terms of weight matrices $\hat{W}_i \in \mathbb{R}^{n_i \times n_i}$ to be chosen. The optimal choice $\hat{W}_i \propto \text{Cov}[Y_i | D_i, X_i]^{-1}$ results in semiparametric efficient estimation of (β, γ) . A variety of approaches have been proposed for constructing the \hat{W}_i . Among the most popular are multilevel models (also known as random or mixed effects models) [Pinheiro and Bates, 2000, Fahrmeir and Tutz, 2001], which additionally make distributional assumptions on the errors ε_i , typically that of Gaussianity, and implicitly specify a particular parametrisation of $\text{Cov}[Y_i | D_i, X_i]$ in terms of the covariates through the introduction of latent random coefficients. Parameters are typically estimated through (restricted) maximum likelihood estimation [Hartley and Rao, 1967, Corbeil and Searle, 1976, Pinheiro and Bates, 2000]. An alternative is the marginal models framework [Heagerty and Zeger, 2000, Diggle et al., 2013, Fahrmeir and Tutz, 2001], which directly models the conditional covariance through a parametric form often estimated via generalised estimating equations [Liang and Zeger, 1986, Hardin and Hilbe, 2003, Ziegler, 2011]. Provided the forms of the conditional covariance are well-specified, any of these approaches will result in efficient estimates for β and γ .

It is however well-known that all models are wrong [Box, 1976], and it is of interest to understand, under misspecification, which approaches remain useful. Below we discuss the consequences of the two potential sources of misspecification, that of the conditional covariance $\text{Cov}[Y_i | D_i, X_i]$, and the conditional mean $\mathbb{E}[Y_i | D_i, X_i]$.

1.1 Conditional covariance misspecification

Misspecification of the conditional covariance has been given a good deal of attention in the literature. The generalised estimating equation approach that has come to be known as GEE1 [Liang et al., 1992] explicitly recognises the possibility of misspecification, and instead specifies what is referred to as a working model for the conditional covariance, with which to construct the weights. Valid inference is guaranteed even with arbitrary (fixed) weights as the estimator (2) is unbiased and standard errors may be based on a sandwich estimate of the variance of $\hat{\beta}$ [Huber, 1967, Gourieroux et al., 1984, Royall, 1986, Liang and Zeger, 1986],

$$\left\{ M^{-1} \begin{pmatrix} \sum_i D_i^\top \hat{W}_i \hat{R}_i \hat{R}_i^\top \hat{W}_i D_i & \sum_i D_i^\top \hat{W}_i \hat{R}_i \hat{R}_i^\top \hat{W}_i X_i \\ \sum_i X_i^\top \hat{W}_i \hat{R}_i \hat{R}_i^\top \hat{W}_i D_i & \sum_i X_i^\top \hat{W}_i \hat{R}_i \hat{R}_i^\top \hat{W}_i X_i \end{pmatrix} M^{-1} \right\}_{1,1}, \quad \hat{R}_i := Y_i - \hat{\beta} D_i - X_i \hat{\gamma} \in \mathbb{R}^{n_i}. \quad (3)$$

The idea behind the working covariance model however is to approximate the ground truth sufficiently well such that the resulting $\hat{\beta}$ has reasonably low variance; various estimation methods have been proposed for this purpose [Prentice and Zhao, 1991, Crowder, 1995, Lumley, 1996, Halekoh et al., 2006].




Figure 1: The (scaled) objective functions of the parameter ρ in the AR(1) working model for quasi-pseudo-Gaussian maximum likelihood (ML), GEE1 (GEE) and the corresponding asymptotic MSE relative to the minimum asymptotic MSE, in a settings with a ground truth given by an ARMA(2, 1) model of Example 1; coloured intervals at the bottom indicate ranges of ρ where the (global) minimisers corresponding to ML and GEE would be outperformed in terms of MSE.

While this intuition is basically well-founded, perhaps surprisingly, for a given model for the covariance, the success of these approaches depends crucially on the method of estimation, as we now demonstrate with two simple examples; specific details for these are given in Appendix A.2.

Example 1 (Conditional correlation misspecification). Consider a version of (1) with X_i omitted for simplicity, $n_i \equiv n$ and each error vector $\varepsilon_i \in \mathbb{R}^n$ given by the first n realisations of an ARMA(2, 1) model. Suppose the weights are estimated in the parametric class consisting of inverses of the covariance matrices of AR(1) processes, indexed by a single autoregressive parameter ρ ; note that the scale of the weights does not affect the resulting $\hat{\beta}$ so we do not consider this as a parameter. We consider two specific settings within this general setup. Figure 1 plots the objective functions that are minimised for two well-established methods for constructing weights based on estimates $\hat{\rho}$ of ρ .

The so-called quasi-pseudo-Gaussian maximum likelihood approach (ML) [Gourieroux and Monfort, 1993, McCullagh and Nelder, 1989, Ziegler, 2011] treats the errors as if they were normally distributed with correlation matrix given by the AR(1) process for a given ρ , and proceeds to maximise jointly over the unknown β, ρ and variance σ^2 , what would be the likelihood were this model to hold.

Motivated by the moment equation $\mathbb{E}[\varepsilon_{ij}\varepsilon_{ik}] = \sigma^2\rho^{|j-k|}$, a second approach (GEE) falling within the GEE1 framework estimates ρ by the minimiser of

$$\sum_{i=1}^I \sum_{j,k=1}^n (\hat{\varepsilon}_{ij}\hat{\varepsilon}_{ik} - \hat{\sigma}^2\rho^{|j-k|})^2, \quad \text{with} \quad \hat{\sigma} := \underset{\sigma>0}{\operatorname{argmin}} \sum_{i=1}^I \sum_{j=1}^n (\hat{\varepsilon}_{ij}^2 - \sigma^2)^2; \quad (4)$$

here the $\hat{\varepsilon}_i$ are the residuals from an initial unweighted least squares regression of Y_i on D_i .




Figure 2: Asymptotic MSEs of β -estimators relative to an unweighted estimator, using weights in a misspecified model class (5) estimated by each of ML and GEE in the settings of Example 2 parametrised by (λ, μ) .

We also plot in orange the asymptotic variance (equivalent to the mean squared error), i.e. the population equivalent of (3), of the β -estimator weighted by the inverse of an AR(1) working correlation matrix for a given value of ρ (the nomenclature ‘SL’ in the legend is explained in Section 1.3). In Setting (a), we see that optimising either of these objectives can result in suboptimal weights in terms of the resulting mean squared error, and any choice of $\rho \in [0.1, 0.8]$ would result in improved estimation of β . Setting (b) tells a similar story, but also illustrates issues that can arise due to local minima of the objective functions, which, in particular are typically not guaranteed to be convex. Attempting to optimise the GEE objective by initialising at $\rho = 0$ results in gradient descent converging to the highly suboptimal local optimum on the left as the derivative of the objective at $\rho = 0$ is slightly positive (see also Table 4 in Appendix A.1, which in particular demonstrates this issue also persists for Newton optimisation schemes). The resulting asymptotic variance of this final $\hat{\beta}$ is substantially worse than even that of the unweighted choice corresponding to $\rho = 0$.

Example 2 (Conditional variance misspecification). Consider an instance of model (1) with $n_i \equiv 1$, so the data are ungrouped, and $d = 1$. Suppose the distribution of the data is such that the true conditional variance $\text{Var}(\varepsilon_i | D_i, X_i = x) = \text{Var}(\varepsilon_i | X_i = x) =: \sigma_0^2(x)$ with $\sigma_0(x) = 2 + \tanh(\lambda(x - \mu))$; we shall consider different settings for the pair of parameters (λ, μ) . Consider using a misspecified class of functions of the form

$$\sigma(x; \eta) = 1 + 2\mathbb{1}_{[\eta, \infty)}(x), \quad (5)$$

where $\eta \in \mathbb{R}$, so the smooth curve of the tanh function is to be approximated by a step function.

Figure 2 plots the relative asymptotic mean squared errors of the weighted least squares estimates of β , given by the first component of (2), for quasi-pseudo-Gaussian maximum likelihood (ML) and GEE1-based approaches (GEE) for constructing weights $\hat{W}_i = 1/\{\sigma(X_i; \hat{\eta})\}^2$ based on estimates $\hat{\eta}$ of η . The former involves treating the errors as if they were normally distributed with standard deviation $\sigma(x; \eta)$ for some η while the latter estimates η by the minimiser of




Figure 3: Schematic of how the ML ‘projections’ $W_{ML,1}^*$ and $W_{ML,2}^*$ of the optimal weight function $W_0 \propto \text{Cov}(Y | D, X)^{-1}$ onto the nested weight classes $\mathcal{W}_1 \subset \mathcal{W}_2$ may be such that $W_{ML,1}^*$ in fact improves on $W_{ML,2}^*$ when the contours of the negative likelihood loss (dotted blue ellipses) do not align with those of the asymptotic MSE; the contours of the latter are here to be thought of as circles centred at W_0 (not shown), so that the distances, $V_1 < V_2$, are the variances of $\hat{\beta}$ constructed using weights $W_{ML,1}^*$ and $W_{ML,2}^*$ respectively. The projections $W_{SL,1}^*$ and $W_{SL,2}^*$ corresponding to the loss given by the asymptotic MSE itself (‘sandwich loss’) do not suffer from any such paradoxes.

the sum of the squared differences between the squared residuals from an initial least squares regression of Y_i on (D_i, X_i) , and $(\sigma(X_i; \eta))^2$; see also Robinson [1987], Carroll [1982], Tsiatis [2006], You et al. [2007] for examples of estimation of the conditional variance through a similar least squares approach for improving estimation of β in (partially) linear models.

We compare these strategies to a naive unweighted estimator, that is (2) with \hat{W}_i constant, which makes no attempt to take advantage of the heteroscedasticity in the errors to improve estimation. Note that such weights are permitted in the model class (5) used by ML and GEE in this example by taking η large, so (5) for some η necessarily gives a better approximation to the ground truth compared to the unweighted approach.

A first interesting observation is the quite different behaviour of the ML and GEE approaches here, with none appearing to be uniformly preferable to the other across all parameter settings. Perhaps more surprising however is the fact that for certain values of (μ, λ) , the performances of these more sophisticated approaches lead to an inflation of the variance over an unweighted estimator (of up to almost 80%). This worrying behaviour can obfuscate model selection via AIC or BIC, as even at the population level they can favour models that result in poorer estimation of the parameter of interest β .

To resolve this apparent paradox, notice that although the model classes used by ML and GEE here are richer, the optimal ‘projections’ of the nuisance function σ_0 (and corresponding optimal weight function $W_0 := \sigma_0^{-2}$) relating to their respective losses do not necessarily coincide with the optimal projection in the sense of the mean squared error of $\hat{\beta}$; in fact in general there is no reason for them to do so, as illustrated schematically in Figure 3 (see also Proposition 1 and Theorem 2 for a formalisation of this phenomenon). We return to these issues in Section 1.3, but first turn our attention to misspecification of the conditional mean $\mathbb{E}[Y_i | D_i, X_i]$.

1.2 Conditional mean misspecification

When the conditional expectation of the response given (D_i, X_i) is not necessarily expected to be linear, a popular model to consider is the partially linear model

$$\begin{aligned} Y_i &= \beta D_i + g_0(X_i) + \varepsilon_i, \\ D_i &= m_0(X_i) + \xi_i. \end{aligned} \tag{6}$$

Here $(Y_i, D_i, X_i, \varepsilon_i)$ are as in (1), so in particular $\mathbb{E}[\varepsilon_i | D_i, X_i] = 0$; g_0 and m_0 are potentially nonlinear row-wise functions, that is e.g. $g_0 : \mathbb{R}^d \rightarrow \mathbb{R}$, and writing X_{ij} for the j th row of matrix X_i , with a slight abuse of notation $g_0 : \mathbb{R}^{n_i \times d} \rightarrow \mathbb{R}^{n_i}$ is then defined via $(g_0(X_i))_j := g_0(X_{ij})$; and error $\xi_i \in \mathbb{R}^{n_i}$ in the D_i on X_i regression satisfies $\mathbb{E}[\xi_i | X_i] = 0$. Note that the model entails the conditional mean independence assumptions that

$$\mathbb{E}(Y_{ij} | D_i, X_i) = \mathbb{E}(Y_{ij} | D_{ij}, X_{ij}) \quad \text{and} \quad \mathbb{E}(D_{ij} | X_i) = \mathbb{E}(D_{ij} | X_{ij}).$$

Nevertheless, the model is flexible enough to well-approximate a wide variety of data generating processes, yet still permits easy interpretation of the contribution of D_i to the response. The second equation serves to model confounding due to X_i . In the ungrouped setting, i.e. where $n_i \equiv 1$, estimating m_0 in addition to g_0 forms a key part of the double / debiased machine learning (DML) framework [Chernozhukov et al., 2018] for inference about β , which in recent years has emerged as the dominant approach for estimation in partially linear models. The popularity of this paradigm is due to the fact that it accommodates the use of arbitrary machine learning methods for estimating the nuisance functions m_0 and g_0 , and requires only a relatively slow rate of $1/N$ for the product of the corresponding mean squared errors in order to yield estimates of β that converge at the parametric $1/\sqrt{N}$ rate.

Emmenegger and Bühlmann [2023] recently extended this approach to the grouped data setting with $n_i > 1$, assuming a parametric form of the covariance $\text{Cov}(\varepsilon_i | D_i, X_i)$ governed by a random effects model. To estimate β , they considered regressing each of Y_i and D_i on X_i using some independent auxiliary data and, with the resulting estimated regression functions, formed corresponding residuals \hat{R}_i^Y and \hat{R}_i^D to give

$$\hat{\beta} = \left(\sum_i \hat{R}_i^{D\top} \hat{W}_i \hat{R}_i^D \right)^{-1} \left(\sum_i \hat{R}_i^{D\top} \hat{W}_i \hat{R}_i^Y \right). \tag{7}$$

Here the weight matrices \hat{W}_i are formed as the inverse conditional covariances estimated using (restricted) maximum likelihood. In practice, sample-splitting and cross-fitting are used in place of auxiliary data, permitting semiparametric efficient estimates, provided the model is well-specified.

However, the flexibility in modelling the conditional mean afforded by DML comes with implications for potential misspecification of the conditional covariance. One key requirement of the approach above is that, in addition to assuming a parametric form for the conditional covariance, it should also not depend on D_i , i.e. we must have $\text{Cov}(\varepsilon_i | D_i, X_i) = \text{Cov}(\varepsilon_i | X_i)$. This restriction is in fact a fundamental limitation of approaches based on DML. It comes as a consequence of requiring Neyman orthogonality, a certain first-order insensitivity to plugging in potentially biased machine learning estimators. Writing $l_0(X_i) := \mathbb{E}(Y_i | X_i)$ and \hat{l} for the corresponding regression estimate, in our case, this entails

$$\sqrt{I} \mathbb{E} \left[\left(\hat{R}_i^Y - \beta \hat{R}_i^D \right)^\top \hat{W}_i \xi_i \right] = \sqrt{I} \mathbb{E} \left[\left(l_0(X_i) - \hat{l}(X_i) \right)^\top \mathbb{E} \left[\hat{W}_i \xi_i | X_i \right] \right] \tag{8}$$

Objective function	Asymptotic MSE of $\hat{\beta}$	Asymptotic integrated MSE of $\hat{\sigma}$
GEE	4.47	0.63
ML	6.82	0.87
Sandwich loss	4.10	1.10

Table 1: Quality of the estimated weights from GEE, ML and minimising the sandwich loss in Example 2 in terms of the derived $\hat{\beta}$ and estimate $\hat{\sigma}$ of σ_0 given by the square root of the inverse of the weights.

being approximately zero, which may not hold unless \hat{W}_i is a function of X_i alone. Thus misspecification of the conditional covariance, and the worrying consequences this may bring, deserve even greater attention when modelling the conditional mean in a flexible fashion through a partially linear model.

1.3 An overview of our contributions

To address the difficulties resulting from (inevitable) misspecification of the conditional covariance, we introduce a new approach for determining weight matrices in weighted least squares estimators (2) or the DML estimator in a partially linear model setting (7). Our proposal is to minimise a sandwich estimate of the variance of $\hat{\beta}$ (i.e. (3) or the equivalent for the DML estimate) over a given parametrisation of weight matrices, thereby directly targeting the primary objective of interest: the estimation performance of $\hat{\beta}$. We thus treat this sandwich estimate of the variance as a loss function—the ‘sandwich loss’ (SL)—by which we determine the weights.

The asymptotic variance plotted in Figure 1 is precisely the population version of this sandwich loss, and minimising this will, by the very definition of the loss, deliver an estimator of β of minimal asymptotic variance among those considered. Returning to Example 2, Table 1 demonstrates that although ML and GEE are to be preferred in terms of estimating the true variance function σ_0^2 , they are worse when estimating β compared to our choice of weights tailored specifically for this purpose. While in these simple examples, the performances of the different methods are noticeably, but not radically different, we show later in Theorem 2 that there exist data generating distributions for which for a large class of misspecified working covariance models, the ratio of the variance for the β -estimators between either GEE or ML and our optimal weighting scheme, can be arbitrarily large.

One key message of our paper therefore is that particularly when there is a high risk of misspecification of the conditional covariance, the sandwich loss may be preferable over existing criteria when inference on the (partially) linear model parameter β is of primary interest. In fact, even in the case that the conditional covariance is well-specified, there is a danger that ML or GEE approaches could converge to weights that are only locally optimal for their respective losses. Since it is only the global optima of these losses that correspond to weights with favourable properties for estimation of β , there is no guarantee that the weights obtained are even locally optimal in terms of the resulting asymptotic variances. In this sense GEE and ML approaches may be more vulnerable to the consequences of local optima in their objective functions, compared to the sandwich loss; see Section 5.1.1.

A second main aim of our work is to introduce a new modelling strategy for working conditional covariances that can harness the power and flexibility of machine learning methods, similarly to how DML uses machine learning to accurately estimate nuisance functions. A challenge however is that standard regression methods cannot be directly deployed to construct

weight matrices. In order to make use of these, we first decompose the inverse of the weights into a working conditional variance of each entry of ε_i , and a working correlation that we model parametrically. We introduce a new gradient boosting approach for estimating these two components through minimising our sandwich loss; this takes as input a user-chosen regression method that is used within the boosting procedure to estimate the conditional variances. We demonstrate the favourable performance of our resulting ‘sandwich boosting’ method in a variety of numerical experiments.

In Section 2 we introduce the sandwich loss and compare its population version to ML and GEE-based equivalents. In Section 2.3 we verify that despite the unusual form of the sandwich loss, under relatively mild conditions, we can expect a minimiser of the sample version to converge to its population counterpart (Theorem 3). We introduce our general cross-fitted weighted estimation approach in Section 3.1 before describing our proposed sandwich boosting scheme in Section 3.2. Section 4 presents theory showing that our resulting estimator for the partially linear model coefficient β is asymptotically Gaussian under relatively mild conditions on the predictive ability of nuisance function estimators, permitting the construction of honest confidence intervals for β . In contrast to existing results in this context, our theory permits the group sizes n_i to grow with the number of groups I , and importantly accommodates misspecification of the conditional covariance. We present the results of a variety of numerical experiments on simulated and real-world data in Section 5 that further explore the themes hinted at in Examples 1 and 2, and demonstrate the effectiveness of our sandwich boosting approach. We conclude with a discussion in Section 6 outlining avenues for further work, including a sketch of an extension to estimating a coefficient function in a version of (6) with β replaced by a function $\beta(X_i)$ that is a linear combination of known basis functions, using a generalisation of the sandwich loss. The supplementary material contains the proofs of all results presented in the main text, additional theoretical results, further details on the examples and numerical experiments, and a detailed computational analysis of the sandwich boosting methodology. Sandwich boosting is implemented in the R package `sandwich.boost`¹. We also provide code to reproduce all simulations, results and figures presented in the manuscript². Below we briefly review some related work not necessarily covered elsewhere in the introduction, and collect together some notation used throughout the paper.

1.4 Other related literature

As indicated in the previous sections, our work connects to a vast literature on mixed effects models, also known as multilevel models or hierarchical models, and generalised estimating equations. Some recent developments in this area have looked at such models in high-dimensional contexts. In particular Li et al. [2022] considers using a particular proxy conditional covariance parametrised by a single parameter for computational simplicity. Li et al. [2018] considers a flexible conditional covariance specification through selecting from high-dimensional random effects via regularising terms in the Cholesky decomposition of the covariance matrix of the random effects.

Most closely related to our setup here however is the work of Emmenegger and Bühlmann [2023] who consider partially linear mixed effect models [Zeger and Diggle, 1994] in the double machine learning (DML) framework, popularised by Chernozhukov et al. [2018]; see also Kennedy [2022] for a recent review of this broad topic. Earlier work considered specific non-parametric estimators for g_0 in the (grouped) partially linear model framework, for example Huang et al. [2007] use regression splines to estimate g_0 and a GEE approach for estimating

¹<https://github.com/elliott-young/sandwich.boost/>

²https://github.com/elliott-young/sandwichboost_simulations

weights. Within the DML area, work related to the setting of the (ungrouped) partially linear model includes [Vansteelandt and Dukes \[2022\]](#) who propose new targets of inference and DML estimation strategies in potentially misspecified generalised partial linear models, and [Emmenegger and Bühlmann \[2021\]](#) who look at estimation in partially linear models with unobserved confounding in an instrumental variables setting using a DML approach and additional regularisation to reduce variance.

Boosting [[Schapire, 1990](#), [Freund and Schapire, 1996](#)], on which our sandwich boosting proposal is built, has received a lot of interest in recent years due to its success on modern datasets of interests. A long line of work (see for example [Breiman \[1999\]](#), [Mason et al. \[1999\]](#), [Friedman et al. \[2000\]](#), and [Bühlmann and Hothorn \[2007\]](#) for a review) in machine learning has resulted in the functional gradient descent perspective of boosting, which we make use of in developing our sandwich boosting proposal.

In general terms, our use of the sandwich loss involves selecting among estimators (in our case determined by weight functions) based on estimates of their quality (in our case, their MSEs). In this sense it is related to a number of statistical approaches, including, for example cross-validation. Of particular note is the recent work of [Park and Kang \[2021\]](#) who look at average treatment effect estimation in multilevel studies and pick from among a family of estimators based on augmented inverse propensity weighting [[Robins et al., 1994](#), [Robins and Rotnitzky, 1995](#)], one minimising an estimate of their variance.

1.5 Notation

We denote by \mathbb{S}^m the set of symmetric $m \times m$ matrices, and for $M \in \mathbb{S}^m$, we write $\Lambda_{\max}(M)$ and $\Lambda_{\min}(M)$ for its maximum and minimum eigenvalues respectively. We write $\mathbb{S}_{++}^m \subset \mathbb{S}^m$ for the set of positive definite matrices. Let Φ denote the cumulative distribution function of a standard Gaussian distribution. We will also use the shorthand $[I] := \{1, \dots, I\}$. For the uniform convergence results we will present, it will be helpful to write, for a law P governing the distribution of a random vector $U \in \mathbb{R}^d$, $\mathbb{E}_P U$ for its expectation and $\mathbb{P}_P(U \in B) =: \mathbb{E} \mathbb{1}_B(U)$ for any measurable $B \subseteq \mathbb{R}^d$. Further, given a sequence of families of probability distributions $(\mathcal{P}_I)_{I \in \mathbb{N}}$, and for a sequence of families of real-valued random variables $(A_{P,I})_{P \in \mathcal{P}_I, I \in \mathbb{N}}$ (which we note are permitted to depend on $P \in \mathcal{P}_I$), we write $A_{P,I} = o_{\mathcal{P}}(1)$ if $\lim_{I \rightarrow \infty} \sup_{P \in \mathcal{P}_I} \mathbb{P}_P(|A_{P,I}| > \epsilon) = 0$ for all $\epsilon > 0$, $A_{P,I} = o_{\mathcal{P}}(g(I))$ for a given function $g : (0, \infty) \rightarrow (0, \infty)$ if $g(I)^{-1} A_{P,I} = o_{\mathcal{P}}(1)$, and $A_{P,I} = O_{\mathcal{P}}(1)$ if for any $\epsilon > 0$ there exist $M_\epsilon, I_\epsilon > 0$ such that $\sup_{I \geq I_\epsilon} \sup_{P \in \mathcal{P}_I} \mathbb{P}_P(|A_{P,I}| > M_\epsilon) < \epsilon$.

2 The sandwich loss

In this section, we first outline a general weighted least squares framework for estimating β , within which we formally introduce the notion of the sandwich loss. In [Section 2.2](#) we study properties of the sandwich loss at the population level, compared to the ML and GEE-based approaches described in [Section 1.1](#). In [Section 2.3](#) we then study the behaviour of the sample version of the sandwich loss and show that under mild conditions, a minimiser will converge in probability to the population minimiser. For the remainder of this paper, we will work in the setting of the partially linear model [\(6\)](#), which includes as a special case, the linear model [\(1\)](#).

2.1 Weighted estimation

Here we outline a general strategy for weighted estimation of β , with which we will introduce our proposed sandwich loss. For simplicity, we describe our approach in terms of an auxiliary dataset, independent of our main data. In practice however, we use sample splitting and cross-fitting to construct our estimator, as described in [Section 3.1](#). As in the approach of [Emmenegger](#)

and Bühlmann [2023], we begin by regressing each of Y and D onto X using our main data to give estimates \hat{l} and \hat{m} of the row-wise conditional expectations $l_0(X_i) = \mathbb{E}(Y_i|X_i)$ and $m_0(X_i) = \mathbb{E}(D_i|X_i)$. With these we form respective vectors of residuals \tilde{R}^Y and \tilde{R}^D using which we find an initial estimate $\tilde{\beta}$ of β using (7) with the weights \hat{W}_i set to identity matrices:

$$\tilde{\beta} = \left(\sum_i \tilde{R}_i^{D\top} \tilde{R}_i^D \right)^{-1} \left(\sum_i \tilde{R}_i^{D\top} \tilde{R}_i^Y \right).$$

We then form ‘estimates’ of the errors ξ_i and ε_i given by $\tilde{\xi}_i := \tilde{R}_i^D$ and $\tilde{\varepsilon}_i := \tilde{R}_i^Y - \tilde{\beta} \tilde{R}_i^D$ to obtain a sandwich estimate of (N times) the variance of a β -estimator utilising given weight matrices $W(X_i)$:

$$\hat{L}_{\text{SL}}(W) := \left(\frac{1}{N} \sum_i \tilde{\xi}_i^\top W(X_i) \tilde{\xi}_i \right)^{-2} \left(\frac{1}{N} \sum_i \left(\tilde{\xi}_i^\top W(X_i) \tilde{\varepsilon}_i \right)^2 \right). \quad (9)$$

In the above, W is a function that given $X \in \mathbb{R}^{n \times d}$ for any $n \in \mathbb{N}$, outputs a matrix $W(X) \in \mathbb{R}^{n \times n}$. The sandwich loss \hat{L}_{SL} views the variance estimate as a function of W , and specifically as a measure of the quality of the weights W , somewhat analogously to how the likelihood views the density of the data as a function of parameters to be determined. Given a class \mathcal{W} of functions W , we then propose to find an (approximate) minimiser \hat{W} of \hat{L}_{SL} over \mathcal{W} (see Section 3.2 for our sandwich boosting approach for carrying this out). While the sandwich loss is thus a rather trivial (re-)definition, as we try to demonstrate in the present work, this shift in perspective from variance estimate to a loss function to be minimised, can lead to non-trivial improvements in terms of estimating β . Note that the sandwich loss is unaffected by any positive scaling of W : this is to be expected since the resulting $\hat{\beta}$ is equally invariant.

On the auxiliary data, we then set $\hat{W}_i := \hat{W}(X_i)$, and form a weighted β -estimate of the form (7). Our \hat{W} should then deliver a low variance estimate of β since this is precisely the way in which it was constructed. Indeed, we show in Section 4 that the variance estimate (9) consistently estimates the true variance (times N).

Note that although we have introduced the sandwich loss via a specific construction of the error estimates $\tilde{\varepsilon}_i$ and $\tilde{\xi}_i$, it is applicable more broadly to other such estimates. For example, in the simpler linear model setting, applying the Woodbury matrix identity shows that the sandwich variance estimate (3) also takes the form given in (9), for certain error estimates formed through weighted least squares regressions.

In the following section, we compare the sandwich loss to the ML and GEE-based approaches outlined in Section 1.1 theoretically by studying population versions of these.

2.2 Population level analysis

The examples in Section 1.1 hint at potential issues with the ML and GEE-based approaches, which unlike the sandwich loss, are not explicitly geared towards minimising the variance of the resulting $\hat{\beta}$: whilst when the conditional covariance is well-specified, their goals coincide with those of the sandwich loss, in the case of misspecification, this need not be the case. Recall that since the weight matrices are restricted to be functions of X_i alone due to the requirement of Neyman orthogonality, it is plausible that some form of misspecification is unavoidable.

To study the properties of the approaches under potential misspecification, we work for simplicity in a setting where we observe i.i.d. instances of the partially linear model (6) with

fixed finite group size $n_i \equiv n$, and consider population versions of the respective losses:

$$\begin{aligned} L_{\text{ML}}(W) &:= \mathbb{E} \left[-\log \det W(X) + \varepsilon^\top W(X) \varepsilon \right], \\ L_{\text{GEE}}(W) &:= \mathbb{E} \left[\|W(X)^{-1} - \varepsilon \varepsilon^\top\|^2 \right], \\ L_{\text{SL}}(W) &:= \left(\mathbb{E} \left[\xi^\top W(X) \xi \right] \right)^{-2} \left(\mathbb{E} \left[\left(\xi^\top W(X) \varepsilon \right)^2 \right] \right). \end{aligned}$$

Here ε, ξ and X are to be understood as generic versions of their counterparts with subscript i satisfying $\Lambda_{\min} \mathbb{E}(\varepsilon \varepsilon^\top | D, X) > \epsilon$ and $\Lambda_{\min} \mathbb{E}(\xi \xi^\top | X) > \epsilon$ almost surely for some $\epsilon > 0$, and $\mathbb{E}[\{\Lambda_{\max} \mathbb{E}(\xi \xi^\top | X)\}^2] < \infty$. Note that the GEE loss L_{GEE} is defined here with respect to an arbitrary matrix norm $\|\cdot\|$ derived from an inner product, such as the Frobenius norm.

In practice we would minimise empirical versions of these loss functions over a restricted class \mathcal{W} of weight functions W . The following family of such classes will facilitate our theoretical comparison of the loss functions. Let $q : \mathbb{R}^{n \times d} \rightarrow \mathbb{R}^{n \times d}$ be a measurable function representing some potential ‘coarsening’ of its input data $x \in \mathbb{R}^{n \times d}$. For example, q may be based on a partition of $\mathbb{R}^{n \times d}$ into disjoint regions, with $x \mapsto q(x)$ reducing the initial data $x \in \mathbb{R}^{n \times d}$ to a representative of the particular region into which x falls. Given some such q , we consider the class of all weight functions given by a measurable function $\tilde{W} : \mathbb{R}^{n \times d} \rightarrow \mathbb{S}^n$ composed with q such that the final outputs are invertible almost surely and obey mild integrability conditions:

$$\mathcal{W} := \left\{ \tilde{W} \circ q \text{ for } \tilde{W} : \mathbb{R}^{n \times d} \rightarrow \mathbb{S}^n : \mathbb{E}[\{\Lambda_{\max}(\tilde{W} \circ q(X))\}^2], \mathbb{E}[\{\Lambda_{\min}(\tilde{W} \circ q(X))\}^{-2}] < \infty \right\}. \quad (10)$$

The above setup includes as a special case, the setting where \mathcal{W} includes all appropriately integrable weight matrices $W : \mathbb{R}^{n \times d} \rightarrow \mathbb{S}^n; x \mapsto W(x)$ so in particular \mathcal{W} would then contain $x \mapsto \mathbb{E}(\varepsilon \varepsilon^\top | X = x)^{-1}$. On the other hand, when the image of q is finite, the resulting \mathcal{W} is essentially a parametric class.

Let us write $V_{\text{ML/GEE}} := L_{\text{SL}}(\mathbb{E}(\varepsilon \varepsilon^\top | q(X))^{-1})$ and V_{SL} for the infimum of L_{SL} over \mathcal{W} . Proposition 1 below shows that $V_{\text{ML/GEE}}$ is the asymptotic variance of both the ML and GEE losses, and gives a condition under which this coincides with V_{SL} .

Proposition 1. (a) L_{ML} and L_{GEE} are both minimised over \mathcal{W} (10) by $\mathbb{E}(\varepsilon \varepsilon^\top | q(X))^{-1}$.

(b) The asymptotic variance V_{SL} of the sandwich loss satisfies $V_{\text{SL}} \leq V_{\text{ML/GEE}}$, with equality if and only if for some constant $c > 0$,

$$\mathbb{E} \left[\text{tr} [\text{Cov}(\varepsilon | q(X))^{-1} \text{Cov}(\varepsilon | D, X)] \xi \xi^\top \mid q(X) \right] = c \mathbb{E} [\xi \xi^\top | q(X)]. \quad (11)$$

Note the condition (11) holds in the instance that $\text{Cov}(\varepsilon | D, X) = \text{Cov}(\varepsilon | q(X))$; but given variable D is considered here to be important enough for its associated parameter to be the target of our inference, it is not inconceivable that the errors depend on them after conditioning on $q(X)$. Moreover, we can expect that $\text{Cov}(\varepsilon | X) \neq \text{Cov}(\varepsilon | q(X))$ unless \mathcal{W} is a sufficiently rich class of functions; in the former case (11) may fail to hold even in the favourable case where $\text{Cov}(\varepsilon | D, X) = \text{Cov}(\varepsilon | X)$. In settings where (11) fails, it is possible for the ratio $V_{\text{ML/GEE}}/V_{\text{SL}}$ to be arbitrarily large, as Theorem 2 below shows.

Theorem 2. Suppose $q(X)$ is not almost surely constant. Then for all $M \geq 1$, and for all pairs $\Sigma, \Omega : \mathbb{R}^{n \times d} \rightarrow \mathbb{S}_{++}^n$ of positive definite matrices that are functions of $q(X)$, there exists a law on (ε, ξ) satisfying the conditions of the model (6) with $\mathbb{E}[\varepsilon \varepsilon^\top | q(X)] = \Sigma(q(X))$, $\mathbb{E}[\xi \xi^\top | q(X)] = \Omega(q(X))$ and

$$\frac{V_{\text{ML/GEE}}}{V_{\text{SL}}} \geq M.$$

While the discrepancy between $V_{\text{ML/GEE}}$ and V_{SL} is not always expected to be very large, it is nevertheless potentially a cause for concern that this can happen even when Σ and Ω are identity matrices, for example.

2.3 Sample level considerations

The previous section illustrated some of the advantages of the sandwich loss at the population level. The sandwich loss \hat{L}_{SL} (9) we propose however is unusual in the sense that it is not composed of a sum of independent terms typical of the objective functions of M-estimators. A further complication is that \hat{L}_{SL} involves estimates of the errors ξ_i and ε_i , rather than these errors themselves. The classical theory of M-estimation [van der Vaart, 1998, Chap. 5] is therefore not immediately applicable here, and it is not clear whether the useful advantages of the population level sandwich loss L_{SL} transfer over to its empirical counterpart. The result below however shows that under relatively mild conditions, minimisation of \hat{L}_{SL} over a parametric class \mathcal{W} (12) yields convergence in probability to the minimiser of L_{SL} . Note that since \hat{L}_{SL} is invariant to any positive rescaling of its argument, we are permitted to fix the scale of any weight function in \mathcal{W} ; we do this by asking for the maximum eigenvalue of any output weight matrix to be 1. We continue to work in the setup of the previous section, where our data consist of I i.i.d. groups of finite size n following the partially linear model (6).

Theorem 3. *Let $\Psi \subset \mathbb{R}^d$ be some compact set and*

$$\mathcal{W} := \left\{ W(\psi) : \mathbb{R}^{n \times d} \rightarrow \mathbb{R}^{n \times n} \text{ with } \sup_{x \in \mathbb{R}^{n \times d}} \Lambda_{\max}(W(\psi)(x)) = 1 \text{ for } \psi \in \Psi \right\}. \quad (12)$$

Suppose $\psi^ \in \Psi$ is such that for all $\epsilon > 0$, $\inf_{\psi \in \Psi: \|\psi - \psi^*\| \geq \epsilon} L_{\text{SL}}(\psi) > L_{\text{SL}}(\psi^*)$. Assume the regularity conditions set out in Appendix B.3, and additionally that for estimates $(\tilde{\xi}_i, \tilde{\varepsilon}_i)$ of the error terms (ξ_i, ε_i) either*

- (a) $\mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\tilde{\xi}_i - \xi_i\|_2^4 \right] \vee \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\tilde{\varepsilon}_i - \varepsilon_i\|_2^4 \right] = o(1)$, or
- (b) $\mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\tilde{\xi}_i - \xi_i\|_2^2 \right] \vee \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\tilde{\varepsilon}_i - \varepsilon_i\|_2^2 \right] = o(N^{-\frac{1}{2}})$.

Then any sequence of approximate minimisers $\hat{\psi}_I$ with $\hat{L}_{\text{SL}}(\hat{\psi}_I) \leq \hat{L}_{\text{SL}}(\psi^) + o_P(1)$ satisfies*

$$\hat{\psi}_I = \psi^* + o_P(1).$$

3 Methodology

In this section we present our sandwich boosting weighted regression procedure. We first describe the basic outline of our approach for generic weight matrices employing cross-fitting in Section 3.1 and then in Section 3.2 present our boosting strategy to approximately optimise the sandwich loss within a flexible class of working covariance models.

3.1 Cross-fitting

In Section 2.1, we outlined a simplified approach for estimating β in the partially linear model (6) involving first obtaining ‘estimates’ $\tilde{\xi}_i$ and $\tilde{\varepsilon}_i$ of the errors ξ_i and ε_i with which to determine a weight function \hat{W} through minimising the sandwich loss (9). The second stage involved forming on an independent dataset, an estimate $\hat{\beta}$ through (7), so in particular, conditioning on the initial dataset, \hat{W} would be fixed. In practice, only a single dataset would be available,

and we construct the two independent datasets through sample splitting, employing a K -fold cross-fitting scheme to recover the loss in efficiency in using only part of the data to construct the final estimator.

Cross-fitting is a popular approach in semiparametric problems for ensuring the independence of nuisance parameter estimates from the data on which the final estimate of the target parameter is formed. This independence means that certain empirical process terms can be controlled straightforwardly even when arbitrary nuisance parameter estimators are used. Chernozhukov et al. [2018] and Emmenegger and Bühlmann [2023] use such cross-fitting in the regular and grouped partially linear models respectively, where the nuisance function estimates in question are \hat{l} and \hat{m} . In our case, cross-fitting additionally serves to guarantee independence of the weight function estimates.

Algorithm 1 details our method, with observation groups indexed by \mathcal{I}_k^c playing the roles of the initial datasets for $k = 1, \dots, K$, and those indexed by \mathcal{I}_k involved in the construction of the final estimator $\hat{\beta}$. Note that rather than forming separate estimates of β corresponding to each \mathcal{I}_k , we instead form sets of residuals \hat{R}_i^D and \hat{R}_i^Y , finally constructing $\hat{\beta}$ using these via (7), an approach known as DML2 [Chernozhukov et al., 2018]. A typical choice of K would be 5 or 10 and averaging over estimators $\hat{\beta}$ constructed using different random partitions $\mathcal{I}_1, \dots, \mathcal{I}_K$ can improve the performance in finite samples; see Section E.4 in the supplementary material for further discussion. As well as obtaining an estimate $\hat{\beta}$, we also calculate a sandwich estimate \hat{V} of the variance, with which an approximate $(1 - \alpha)$ -level confidence interval $\hat{C}(\alpha)$ may be constructed.

3.2 Sandwich boosting

Algorithm 1 introduced a generic approach for incorporating weight functions learnt from the data into an estimator for β via approximately minimising the sandwich loss over some class of functions \mathcal{W} . We now introduce an approach for performing this approximate minimisation over a class \mathcal{W} defined implicitly through a user-chosen regression method.

To introduce our approach, it is helpful to consider a class of proxy conditional covariances parametrised as

$$D_\sigma(\cdot) C_\theta(\cdot) D_\sigma(\cdot), \quad (13)$$

where, given an input $X_i \in \mathbb{R}^{n_i \times d}$, the functions D_σ and C_θ output

$$\begin{aligned} D_\sigma(X_i) &:= \text{diag}(\sigma(X_{i1}), \dots, \sigma(X_{in_i})) \\ C_\theta(X_i) &:= \left(\mathbb{1}_{\{j=k\}} + \mathbb{1}_{\{j \neq k\}} \rho_\theta(X_{ij}, X_{ik}) \right)_{(j,k) \in [n_i]^2}. \end{aligned}$$

Here, $\sigma : \mathbb{R}^d \rightarrow (0, \infty)$ and $\rho_\theta : \mathbb{R}^d \times \mathbb{R}^d \rightarrow [0, 1]$ for $\theta \in \Theta$ (where Θ is some closed convex set) are proxy conditional standard deviation and correlation functions that are to be modelled nonparametrically and parametrically respectively. Note that the working covariances (13) have the property that the jk th entry depends only on X_{ij} and X_{ik} ; this need not be the case for $\{\text{Cov}(Y_i | X_i)\}_{jk}$ for example, but is nevertheless a reasonable simplification.

Redefining $s \equiv 1/\sigma$, the corresponding weight class consists of functions of the form

$$W(s, \theta)(\cdot) = D_s(\cdot) C_\theta^{-1}(\cdot) D_s(\cdot), \quad (14)$$

understanding, $C_\theta^{-1}(X_i) := \{C_\theta(X_i)\}^{-1}$ up to an arbitrary positive scale factor. As an example, an equicorrelated working correlation may be parametrised as

$$C_\theta^{-1}(X_i) = \left(\mathbb{1}_{\{j=k\}} - \frac{\theta}{1 + \theta n_i} \right)_{(j,k) \in [n_i]^2}$$

Algorithm 1: Construction of β -estimator

Input: Data $(Y_i, D_i, X_i)_{i \in [I]}$; number of folds for cross-fitting K ; choices of regression methods for estimating (l_0, m_0) ; method for finding weight estimates; significance level α .

Partition $[I]$ randomly into K disjoint sets $\mathcal{I}_1, \dots, \mathcal{I}_K$ approximately equal in size.

for $k \in [K]$ **do**

Using data indexed by \mathcal{I}_k^c , fit estimates $\hat{l}^{(k)}$ and $\hat{m}^{(k)}$ of the functions l_0 and m_0 by regressing Y and D against X respectively.

for $i \in \mathcal{I}_k^c$ **do**

Calculate residuals $\tilde{R}_i^Y := Y_i - \hat{l}^{(k)}(X_i)$ and $\tilde{R}_i^D := D_i - \hat{m}^{(k)}(X_i)$.

Calculate $\tilde{\beta}_k := \left(\sum_{i \in \mathcal{I}_k^c} \tilde{R}_i^{D^\top} \tilde{R}_i^D \right)^{-1} \left(\sum_{i \in \mathcal{I}_k^c} \tilde{R}_i^{D^\top} \tilde{R}_i^Y \right)$.

Calculate $\tilde{\xi}_i := \tilde{R}_i^D$ and $\tilde{\varepsilon}_i := \tilde{R}_i^Y - \tilde{\beta}_k \tilde{R}_i^D$.

end

Find an (approximate) minimiser $\hat{W}^{(k)}$ of

$$\hat{L}_{\text{SL}}^{(k)}(W) := \left(\sum_{i \in \mathcal{I}_k^c} \tilde{\xi}_i^\top W(X_i) \tilde{\xi}_i \right)^{-2} \left(\sum_{i \in \mathcal{I}_k^c} \left(\tilde{\xi}_i^\top W(X_i) \tilde{\varepsilon}_i \right)^2 \right),$$

over some class of functions \mathcal{W} corresponding to a working covariance structure (e.g. using sandwich boosting, see Section 3.2).

for $i \in \mathcal{I}_k$ **do**

Set $\hat{R}_i^Y := Y_i - \hat{l}^{(k)}(X_i)$ and $\hat{R}_i^D := D_i - \hat{m}^{(k)}(X_i)$ and $\hat{W}_i := \hat{W}^{(k)}(X_i)$.

Define $\hat{\xi}_i := \hat{R}_i^D$ and $\hat{\varepsilon}_i := \hat{R}_i^Y - \tilde{\beta}_k \hat{R}_i^D$.

end

end

Calculate $\hat{\beta} := \left(\sum_{k=1}^K \sum_{i \in \mathcal{I}_k} \hat{R}_i^{D^\top} \hat{W}_i \hat{R}_i^D \right)^{-1} \left(\sum_{k=1}^K \sum_{i \in \mathcal{I}_k} \hat{R}_i^{D^\top} \hat{W}_i \hat{R}_i^Y \right)$.

Calculate $\hat{V} := N \left(\sum_{k=1}^K \sum_{i \in \mathcal{I}_k} \hat{\xi}_i^\top \hat{W}_i \hat{\xi}_i \right)^{-2} \left(\sum_{k=1}^K \sum_{i \in \mathcal{I}_k} \left(\hat{\xi}_i^\top \hat{W}_i \hat{\varepsilon}_i \right)^2 \right)$.

Calculate $\hat{C}(\alpha) := \left[\hat{\beta} - N^{-\frac{1}{2}} \hat{V}^{\frac{1}{2}} \Phi^{-1} \left(1 - \frac{\alpha}{2} \right), \hat{\beta} + N^{-\frac{1}{2}} \hat{V}^{\frac{1}{2}} \Phi^{-1} \left(1 - \frac{\alpha}{2} \right) \right]$.

Output: Estimator $\hat{\beta}$ for β , an estimator of its asymptotic variance \hat{V} and $1 - \alpha$ level confidence interval $\hat{C}(\alpha)$ for β .

for $\theta \in [0, \infty)$, with corresponding correlation ρ_θ given by $\theta/(1 + \theta)$. We also consider a version for nested group structures permitting two constant correlations and an autoregressive form suitable for longitudinal data where $\{C_\theta(X_i)\}_{jk} = \theta^{|j-k|}$; see Appendix C. Such inverse working correlations are among the classes of weights considered in the GEE1 framework [Liang and Zeger, 1986, Zeger and Liang, 1986, Ziegler, 2011]. A key difference here however is the greater flexibility afforded by learning the working inverse standard deviation s function through a boosting scheme, as we now explain. We also outline in Section 3.2.1 how our boosting scheme may be initialised at estimated weight functions derived using existing ML or GEE-based methods, for example, thereby increasing the flexibility of the functional forms considered.

Boosting has emerged as one of the most successful learning methods, with the XGBoost implementation [Chen and Guestrin, 2016] in particular dominating machine learning competitions such as those hosted on Kaggle [Bojer and Meldgaard, 2021]. Since its introduction in the work of Schapire [1990], it has been generalised and reinterpreted as a form of functional gradient descent of an objective function \hat{L} based on the data [Friedman et al., 2000, Mason et al., 1999, Bühlmann and Yu, 2003, Bühlmann and Hothorn, 2007]. For an objective func-

tion $f \mapsto \hat{L}(f) \in \mathbb{R}$ applied to function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, an individual boosting iteration involves perturbing the f -function by a step in the ‘direction’ of the f -score

$$U^{(f)}(f)(x) := \partial_f \hat{L}(f)(x) := \left. \frac{\partial}{\partial \alpha} \hat{L}(f + \alpha \delta_x) \right|_{\alpha=0}, \quad (15)$$

where $\delta_x : \mathbb{R}^d \rightarrow \mathbb{R}$ is the indicator function at $x \in \mathbb{R}^d$. Procedurally, the f -score $U^{(f)}(f)$ is evaluated at the data points and the regressed onto the data using a user-chosen ‘base learner’.

Typically \hat{L} takes the form of an empirical risk, so $\hat{L}(f) = \sum_i \ell(y_i, f(x_i))$ for some loss function ℓ and predictor–response pairs (x_i, y_i) . The corresponding f -score evaluated at the data point x_i then takes the simple form $\left. \frac{\partial}{\partial \alpha} \ell(y_i, \alpha) \right|_{\alpha=f(x_i)}$, a function of the i th observation alone. This allows for f -score calculation in linear time, as well as the possibility of parallelising computations for large data sets as exploited by XGBoost. In our case however where we wish to minimise the sandwich loss $(s, \theta) \mapsto \hat{L}_{\text{SL}}(W(s, \theta))$ (which recall is defined in terms of estimates $\tilde{\varepsilon}_i$ and $\tilde{\xi}_i$ of the errors) over weight functions parametrised by (s, θ) (14), we obtain as the s -score

$$U_{\text{SL}}^{(s)}(s, \theta)(X_{ij}) = -2b^{-3} \left\{ \left(\tilde{\xi}_i^\top A_{ij} \tilde{\xi}_i \right) \sum_{i'} c_{i'}^2 - bc_i \left(\tilde{\xi}_i^\top A_{ij} \tilde{\varepsilon}_i \right) \right\}, \quad (16)$$

where $W_i(s, \theta) := W(s, \theta)(X_i)$ and

$$A_{ij} := \left(\left\{ C_\theta^{-1}(X_i) \right\}_{kj} s(X_{ik}) \mathbb{1}_{\{l=j\}} + \left\{ C_\theta^{-1}(X_i) \right\}_{jl} s(X_{il}) \mathbb{1}_{\{k=j\}} \right)_{k,l \in [n_i]^2},$$

$$b := \sum_{i=1}^I \tilde{\xi}_i^\top W_i(s, \theta) \tilde{\xi}_i, \quad c_i := \tilde{\xi}_i^\top W_i(s, \theta) \tilde{\varepsilon}_i.$$

Thus the s -score at X_{ij} is a function of all the data points. Nevertheless, it may be computed for all X_{ij} at a cost of $O(\sum_i n_i^3)$. However, as we show in Appendix C, for the equicorrelated, nested and autoregressive working correlation structures, this cost is reduced to $O(N)$ and may be parallelised similarly to the standard setting of minimising an empirical risk. The critical factors in allowing this are: (a) that computing the matrix inverse $C_\theta^{-1}(X_i)$ present in $W_i(s, \theta)$, which for an arbitrary correlation ρ_θ may take $O(n_i^3)$ time, has a simple closed form; and (b) computation of the terms $\tilde{\xi}_i^\top A_{ij} \tilde{\xi}_i$ and $\tilde{\xi}_i^\top A_{ij} \tilde{\varepsilon}_i$ involving the sparse matrix A_{ij} can be arranged to be $O(1)$ by precomputing other terms appropriately.

Along with updating s by regressing the s -score above onto the X_{ij} and taking a step in the direction of the negative of this fitted regression function, we may also perform a regular projected gradient descent update for θ using the θ -score vector

$$U_{\text{SL}}^{(\theta)}(s, \theta) = -2b^{-3} \left\{ \left(\sum_i c_i^2 \right) \left(\sum_i \tilde{\xi}_i^\top \partial_\theta W_i(s, \theta) \tilde{\xi}_i \right) - b \sum_i c_i \left(\tilde{\xi}_i^\top \partial_\theta W_i(s, \theta) \tilde{\varepsilon}_i \right) \right\}, \quad (17)$$

which may be computed at no greater cost than the s -score above.

With these scores, our sandwich boosting algorithm is summarised in Algorithm 2; note π_Θ denotes projection onto the set Θ . Recall that in our cross-fitting scheme (Algorithm 1), we envisage applying boosting to approximately minimise a version of the sandwich loss corresponding to subsets of the observation indices. As is standard in boosting, the algorithm requires a choice of initialisers (in our case \hat{s}_1 and $\hat{\theta}_1$) and a base learner. In all of our numerical experiments, we take $\hat{s}_1 \equiv 1$, $\hat{\theta}_1 = 0$ and use additive penalised cubic regression splines implemented in the R package `mgcv` [Wood, 2017]. We select the number of boosting iterations m_{stop} by cross-validation, as recommended by [Bühlmann and Hothorn, 2007], though using our sandwich loss as the evaluation criterion. Note that the algorithm is stated for fixed step sizes $\lambda^{(\theta)}$ and $\lambda^{(s)}$ for simplicity; in Appendix E, we describe the specific choices and variable step size schemes used in our numerical results.

3.2.1 Initialising from other weighting schemes

The classes of weight functions that may be fitting using our computationally efficient boosting schemes with equicorrelated, autoregressive or nested correlations can be rather rich when used in conjunction with a flexible base learner. However, these classes would not encompass all those available using classical mixed effects modelling, for example. In order to further broaden the classes of weight functions that may be considered, one can start with an initial estimated weight or conditional covariance function $\hat{\Sigma}_{\text{init}}$ estimated through GEE or ML-based approaches, and fit a weight function of the form

$$\{\hat{\Sigma}_{\text{init}}(\cdot)\}^{-1/2} W(\cdot) \{\hat{\Sigma}_{\text{init}}(\cdot)\}^{-1/2},$$

where $W(\cdot)$ is of the form given by (14), using sandwich boosting. The boosting algorithm then serves to push the initial $\hat{\Sigma}_{\text{init}}$ in a better direction for the purposes of estimating β . This may be carried out easily by running Algorithm 2 on transformed error estimates $\tilde{\varepsilon}_i \mapsto \{\hat{\Sigma}_{\text{init}}(X_i)\}^{-1/2} \tilde{\varepsilon}_i$ and similarly for $\tilde{\xi}_i$. In fact, one can use multiple initialisers in this way, and pick among the best sandwich-boosted versions via cross-validation with the sandwich loss as the quality criterion.

Algorithm 2: Sandwich boosting algorithm

Input: Index set \mathcal{I} for boosting training set $(X_i, \tilde{\varepsilon}_i, \tilde{\xi}_i)_{i \in \mathcal{I}}$; boosting initialisers $\hat{s}_1 : \mathbb{R}^d \rightarrow (0, \infty)$ and $\hat{\theta}_1 \in \mathbb{R}$; θ -descent and s -descent step sizes $\lambda^{(\theta)}, \lambda^{(s)} > 0$; total number of boosting iterations m_{stop} ; base learner for regressing scores onto the data.

for $m = 1, \dots, m_{\text{stop}}$ **do**

 Calculate the N scores $U_{\text{SL}}^{(s)}(\hat{s}_m, \hat{\theta}_m)(X_{ij})$ (16) and θ -score $U_{\text{SL}}^{(\theta)}(\hat{s}_m, \hat{\theta}_m)$ (17) using Algorithm 3.

 Regress $\left(U_{\text{SL}}^{(s)}(\hat{s}_m, \hat{\theta}_m)(X_{ij}) \right)_{i \in \mathcal{I}, j \in [n_i]}$ onto $(X_{ij})_{i \in \mathcal{I}, j \in [n_i]}$ using the base learner to obtain the fitted regression function $x \mapsto \hat{u}_m(x)$.

$\hat{s}_{m+1} := \hat{s}_m - \lambda^{(s)} \hat{u}_m$.

$\hat{\theta}_{m+1} := \pi_{\Theta}(\hat{\theta}_m - \lambda^{(\theta)} U_{\text{SL}}^{(\theta)}(\hat{s}_m, \hat{\theta}_m))$.

end

Output: Weight function $\hat{W} := W(\hat{s}_{m_{\text{stop}}+1}, \hat{\theta}_{m_{\text{stop}}+1})$.

4 Theory

In this section we present results on asymptotic normality of the β -estimator of Algorithm 1, and coverage guarantees for the confidence interval construction therein. Recalling the setup of Section 3.2, we consider the case where the estimated weight functions are such that $\{\hat{W}^{(k)}(\cdot)\}^{-1}$ takes the form (13); these may be obtained using our sandwich boosting approach (Algorithm 2), but this is not required for our theoretical results. Let us define $\hat{\sigma}^{(k)}$ and $\hat{\theta}^{(k)}$ by

$$D_{\hat{\sigma}^{(k)}}(\cdot) C_{\hat{\theta}^{(k)}}(\cdot) D_{\hat{\sigma}^{(k)}}(\cdot) := \{\hat{W}^{(k)}(\cdot)\}^{-1},$$

and write $\hat{\rho}^{(k)} := \rho_{\hat{\theta}^{(k)}}$.

For simplicity of the exposition, similarly to Sections 2.2, here we consider the case where our data are i.i.d. copies of the group of n observations $(Y, D, X) \in \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R}^{n \times d}$ following the partially linear model (6). Recall that X is a matrix whose rows, denoted $X_j \in \mathbb{R}^d$, are not

necessarily independent or identically distributed. We also relax the i.i.d. assumption at the group level to permit non-identically distributed groups of unequal size in Appendix D.

Our results here are based in part on Emmenegger and Bühlmann [2023], but build on them in two key ways. Firstly, we permit the conditional covariance $\text{Cov}(Y | D, X)$ to be misspecified, i.e., for the (likely) possibility that the probability limit of the $\hat{W}^{(k)}(X)$ is not some multiple of $\text{Cov}(Y | D, X)^{-1}$. Secondly, we consider asymptotic regimes that allow the group size n to diverge with the total number of observations $N = nI$ at rates we will formalise later. Throughout, we assume that the number of folds K used in cross-fitting is finite.

We state our results as uniform convergence results over the sequence of classes of distributions $(\mathcal{P}_I)_{I \in \mathbb{N}}$ such that for all I sufficiently large and for all $P \in \mathcal{P}_I$, the following hold. Note that in the below, $\delta, \mu_\Sigma, \mu_\Omega, \gamma, \alpha$ and κ are to be thought of as constants, not depending on P . The values of these are not relevant in the case where the group size n is finite, but play a role in the rate of growth permitted when it is diverging. Moreover $a \lesssim b$ denotes $a \leq cb$ for constant $c > 0$ not depending on P . We have however suppressed the dependence on P in l_0, σ^* etc.

Assumption A1 (Moment assumptions).

$$(A1.1) \quad \text{There exists } \delta > 0 \text{ such that } \mathbb{E}_P \left[\|\varepsilon\|_2^{4+\delta} \right]^{\frac{1}{4+\delta}} \lesssim \sqrt{n} \text{ and } \mathbb{E}_P \left[\|\xi\|_2^{4+\delta} \right]^{\frac{1}{4+\delta}} \lesssim \sqrt{n}.$$

$$(A1.2) \quad \text{The covariance matrices } \Sigma(D, X) := \mathbb{E}_P [\varepsilon \varepsilon^\top | D, X] \text{ and } \Omega(X) := \mathbb{E}_P [\xi \xi^\top | X] \text{ satisfy } \Lambda_{\min}(\Sigma) \gtrsim 1 \text{ and } \Lambda_{\min}(\Omega) \gtrsim 1 \text{ almost surely. Further, } \Lambda_{\max}(\Sigma) \lesssim n^{\mu_\Sigma} \text{ and } \Lambda_{\max}(\Omega) \lesssim n^{\mu_\Omega} \text{ almost surely for some } \mu_\Sigma, \mu_\Omega \in [0, 1].$$

Lower values of μ_Σ and μ_Ω will permit faster rates of divergence of n (see Appendix D). Note that when Σ is close to the equal correlation working covariance of Section 3.2, we can expect $\mu_\Sigma = 1$. For our simplified result in Corollary 5 we set $\mu_\Omega = 1$. In the below, whenever we condition on estimated regression functions, this is to be understood as conditioning on the data used to train these.

Assumption A2 (Accuracy of regression function estimators). Define the maximum within group estimation errors of regression functions l_0 and m_0 :

$$\mathcal{R}_l := \max_{j \in [n]} \mathbb{E}_P \left[\left(\hat{l}^{(1)}(X_j) - l_0(X_j) \right)^2 \middle| \hat{l}^{(1)} \right], \quad \mathcal{R}_m := \max_{j \in [n]} \mathbb{E}_P \left[\left(\hat{m}^{(1)}(X_j) - m_0(X_j) \right)^2 \middle| \hat{m}^{(1)} \right].$$

Then the errors of these nuisance function estimators satisfy:

$$(A2.1) \quad \mathcal{R}_m (\mathcal{R}_l \vee \mathcal{R}_m) = o_{\mathcal{P}}(N^{-1}),$$

$$(A2.2) \quad \mathcal{R}_m \vee \mathcal{R}_l = o_{\mathcal{P}}(1),$$

$$(A2.3) \quad \max_{j \in [n]} \mathbb{E}_P \left[\left| \hat{l}^{(1)}(X_j) - l_0(X_j) \right|^{4+\delta} \middle| \hat{l}^{(1)} \right] = O_{\mathcal{P}}(1) \text{ and} \\ \max_{j \in [n]} \mathbb{E}_P \left[\left| \hat{m}^{(1)}(X_j) - m_0(X_j) \right|^{4+\delta} \middle| \hat{m}^{(1)} \right] = O_{\mathcal{P}}(1).$$

The assumptions on the regression function estimates are relatively weak and identical to those in Emmenegger and Bühlmann [2023], with what is typically the strongest requirement (A2.1) permitting nonparametric rates of $o_{\mathcal{P}}(N^{-1/2})$ for each of \mathcal{R}_m and \mathcal{R}_l . Faster rates than this however weaken conditions on how n may diverge; see Corollary 5 below.

Assumption A3 (Stability of weight function estimates). Suppose there also exists deterministic functions $\sigma^* : \mathbb{R}^d \rightarrow \mathbb{R}$ and $\rho^* : \mathbb{R}^d \times \mathbb{R}^d \rightarrow (0, \infty)$ whose estimators satisfy:

$$(A3.1) \quad \mathcal{R}_\sigma := \max_{j \in [n]} \mathbb{E}_P \left[(c^* \hat{\sigma}^{(1)}(X_j) - \sigma^*(X_j))^2 \mid \hat{\sigma}^{(1)} \right] = o_{\mathcal{P}}(1)$$

where $c^* := \operatorname{arginf}_{c > 0} \max_{j \in [n]} \mathbb{E}_P \left[(c \hat{\sigma}^{(1)}(X_j) - \sigma^*(X_j))^2 \right]$,

$$(A3.2) \quad \mathcal{R}_\rho := \max_{j \neq j'} \mathbb{E}_P \left[(\hat{\rho}^{(1)}(X_j, X_{j'}) - \rho^*(X_j, X_{j'}))^2 \mid \hat{\rho}^{(1)} \right] = o_{\mathcal{P}}(1).$$

Further suppose the associated weights $W^* := W(\sigma^*, \rho^*)(X)$ satisfy

$$(A3.3) \quad \Lambda_{\max}(W^*) \lesssim 1 \text{ and } \Lambda_{\min}(W^*) \gtrsim n^{-\gamma} \text{ for some } \gamma \in [0, \mu_\Sigma] \text{ almost surely,}$$

$$(A3.4) \quad \Lambda_{\min}\left(W^{*\frac{1}{2}} \Sigma W^{*\frac{1}{2}}\right) \gtrsim n^{-\kappa} \text{ and } \Lambda_{\max}\left(W^{*\frac{1}{2}} \Sigma W^{*\frac{1}{2}}\right) \lesssim n^\alpha \text{ for some } \kappa \in [0, \gamma] \text{ and } \alpha \in [0, \mu_\Sigma] \text{ almost surely.}$$

Assumptions [A3.1](#) and [A3.2](#) require a probabilistic limit for our estimates of the weight function, but this need not correspond closely to the inverse of Σ . The eigenvalue assumption [A3.4](#) however does loosely quantify the discrepancy between these, and κ and α impact the permitted divergence rate of n . The reason for introducing the c^* is that the estimated weights need not be on the same scale as Σ^{-1} (recall that the sandwich loss is invariant to positive scaling of its argument).

Theorem 4. *Consider Algorithm 1. Let the sequence of distribution families $(\mathcal{P}_I)_{I \in \mathbb{N}}$ for (Y, D, X) be such that for all I sufficiently large, and for all $P \in \mathcal{P}_I$, Assumptions [A1](#), [A2](#) and [A3](#) are satisfied. Further, suppose that the group size n is either finite, or diverges at a rate satisfying Assumption [A4](#) in Appendix D. Then defining*

$$V := \left(\mathbb{E}_P \left[\xi^\top W^* \xi \right] \right)^{-2} \left(\mathbb{E}_P \left[\left(\varepsilon^\top W^* \xi \right)^2 \right] \right),$$

we have that $\hat{\beta}$ is uniformly asymptotically Gaussian

$$\lim_{I \rightarrow \infty} \sup_{P \in \mathcal{P}_I} \sup_{t \in \mathbb{R}} \left| \mathbb{P}_P \left(\sqrt{N/V} (\hat{\beta} - \beta) \leq t \right) - \Phi(t) \right| = 0,$$

and moreover the above holds with V replaced by \hat{V} .

The result shows in particular that $\hat{C}(\alpha)$ constructed in Algorithm 1 is an asymptotic $(1-\alpha)$ -level honest confidence interval, under the given assumptions. Corollary 5 below specialises a version of Theorem 4 for diverging group sizes for two cases of interest where relatively simple forms of (conservative) rate requirements on n are available.

Corollary 5. *Adopt the setup and notation of Theorem 4 but suppose $\delta \geq 4$ in [A1.1](#) and [A2.3](#), and additionally that $\mathcal{R}_\rho = O_{\mathcal{P}}(N^{-1})$. Suppose the estimated weight functions $\hat{W}^{(k)}$ are constructed to fall within classes \mathcal{W} (see Section 3.2) corresponding to one of the following two settings:*

(i) *Equicorrelated working correlation, but where the true conditional correlation $\operatorname{Corr}(Y \mid X, D)$ may be arbitrary;*

(ii) *Autoregressive AR(1) working correlation and when $\mu_\Sigma = 0$ (see [A1.2](#)).*

Then the conclusions of Theorem 4 holds for diverging group sizes at the following rates:

$$\begin{aligned} \text{Equicorrelated (i):} \quad & n = o \left(N^{\frac{1}{3+2\kappa}} \wedge (N \mathcal{R}_m (\mathcal{R}_l \vee \mathcal{R}_m))^{-\frac{1}{\kappa}} \wedge \mathcal{R}_\sigma^{-\frac{1}{2(\kappa+2)}} \right), \\ \text{Autoregressive (ii):} \quad & n = o \left(N^{\frac{1}{3}} \wedge \mathcal{R}_l^{-1} \wedge \mathcal{R}_\sigma^{-1} \right). \end{aligned}$$

We discuss each of the cases (i) and (ii) in turn. Case (i) places few restrictions on the true conditional covariance and so results in a more stringent requirement on the growth rate of n . Recall that the middle term $N \mathcal{R}_m(\mathcal{R}_l \vee \mathcal{R}_m)$ is required to be $o_P(1)$ by [A2.1](#) and small values of the parameter $\kappa \in [0, 1]$ indicate better approximation of Σ in [A3.4](#). Case (ii) enforces that $\mu_\Sigma = 0$: this occurs for instance when the true correlation function is upper bounded by an exponentially decaying function with separation (satisfied e.g. for ARMA processes). As such, this condition may be appropriate for longitudinal data. The rate requirement on n is relatively weak: both \mathcal{R}_l and \mathcal{R}_σ need only satisfy a rate requirement weaker than that on \mathcal{R}_m entailed by [A2.1](#) in order for n to be permitted to grow at any rate $o(N^{1/3})$.

5 Numerical experiments

In this section we explore the empirical properties of the sandwich boosting estimator on a number of simulated and real-world datasets. In all cases where covariates X_i are available in addition to our covariate of interest D_i , we fit partially linear models using the approach of [Algorithm 1](#), estimating nuisance regression functions m_0 and l_0 using cubic regression splines implemented in the `mgcv` package [[Wood, 2017](#)]; however in addition to using weight functions $\hat{W}^{(k)}$ selected using sandwich boosting, we also compare to versions with these selected using quasi-pseudo-Gaussian maximum likelihood (ML) and GEE1 (GEE) based methods; see [Section E.1](#) of the Appendix for further details. Note that the use of [Algorithm 1](#) with ML is essentially the approach of [Emmenegger and Bühlmann \[2023\]](#), as in the `dmlalg` package [[Emmenegger, 2021](#)], but using a robust sandwich estimate of the variance.

[Section 5.1](#) explores four simulated settings with varying degrees of misspecification of the conditional covariance. [Section 5.2](#) looks at two datasets: the first, on orange juice sales grouped by store, highlights the benefits of flexible variance modelling via the nonparametric $s \equiv 1/\sigma$ component in our sandwich boosting scheme ([Section 3.2](#)); and the second, a longitudinal study on women’s wages, provides a real-life example of the phenomena seen in [Examples 1 and 2](#) where more complex conditional covariance models can lead to poorer β -estimation when weights are selected using ML or GEE-based approaches, and where the minimiser of the sandwich loss is rather different to those corresponding to the ML and GEE objectives.

5.1 Simulated data

We look at four simulated scenarios, and in each case consider different weight classes \mathcal{W} : we describe these classes below in terms of their implied working covariances i.e. in terms of the inverses of the weight matrices. For all the approaches we use the cross-fitting scheme ([Algorithm 1](#)) with $K = 2$ folds.

- **Homoscedastic:** Depending on the setting, this consists of either equicorrelated or autoregressive AR(1) working correlations scaled by a constant variance, with the single parameter estimated either by maximising a Gaussian likelihood (ML), an approach of the form given in [\(4\)](#) (GEE) or minimising the sandwich loss through projected gradient descent (only used in the setting of [Section 5.1.3](#)).
- **Heteroscedastic:** This uses the same working correlations as in the homoscedastic case, but allows for more flexibility in the working conditional variance function with specifics depending on the estimation method used.
 - **ML:** We model the logarithm of the conditional variance function with a polynomial basis in the covariate X_i with the number of basis functions (restricted to at most

- 4 to avoid numerical instabilities) determined by cross-validation using Gaussian log-likelihood loss. This is carried out using the `nlme` package [Pinheiro et al., 2022].
- **GEE:** We perform a cubic penalised spline regression of the squared residuals $\tilde{\varepsilon}_{ij}^2$ onto X_i using the `mgcv` package [Wood, 2017]) to obtain an estimate of the working conditional variance function with the parameter of the correlation subsequently estimated using the `geeglm` package [Halekoh et al., 2006].
 - **Sandwich loss:** We use sandwich boosting as described in Section 3.2.

Section 5.1.1 considers a well-specified setting where the optimal true conditional covariance weights can in principle be replicated by the heteroscedastic weight estimators; Section 5.1.2 looks at a misspecified setting where the optimal weights W_0 depend on D and varies the degree of misspecification; and Sections 5.1.3 and 5.1.4 explore the effect of varying group sizes in a settings with mildly misspecified conditional correlation and variance respectively. In all cases, we simulate data with equal independent and identically distributed groups of equal size n , which we vary across the settings. The mean squared errors of β -estimators corresponding to each method and setting pair, averaged over 500 repetitions, are shown in Figure 4. Coverage of associated confidence intervals is reported in Section F of the Appendix; in particular the coverage probabilities of nominal 95% confidence intervals based on our sandwich boosting approach varies from 0.938–0.984.

5.1.1 Increasing model complexity

Consider $I = 2000$ i.i.d. instances of $(Y, D, X) \in \mathbb{R}^{10} \times \mathbb{R}^{10} \times \mathbb{R}^{10}$ following the partially linear model (6) with the following properties: the group size $n = 10$ where $\beta = 1$ is the target parameter of interest, X is componentwise i.i.d. uniform $U[-5, 5]$, $m_0(x) = \cos x$, $g_0(x) = \tanh x$ (with the \cos and \tanh functions applied component-wise), $\xi | X \sim N_{10}(\mathbf{0}, \Omega(X))$ and $\varepsilon | (D, X) \sim N_{10}(\mathbf{0}, \Sigma(X))$ with covariance matrices given by

$$\Sigma(X) = \left((\mathbb{1}_{\{j=k\}} + 0.2 \cdot \mathbb{1}_{\{j \neq k\}}) \sigma_0(X_j) \sigma_0(X_k) \right)_{(j,k) \in [10]^2}, \quad \sigma_0(x) = 2 + \cos(\lambda x),$$

$$\Omega(X) = (\mathbb{1}_{\{j=k\}} + 0.1 \cdot \mathbb{1}_{\{j \neq k\}})_{(j,k) \in [10]^2}.$$

Here $\lambda \geq 0$ is the Lipschitz constant (complexity parameter) of the conditional variance function σ_0 , that we will vary. We use homoscedastic and heteroscedastic working covariance classes with equicorrelated working correlation (noting that the true correlation is also constant here). As $\varepsilon | (D, X) \stackrel{d}{=} \varepsilon | X$, by Proposition 1 the population minimiser of the sandwich loss and those corresponding to the ML and GEE approaches should all coincide, and so from this perspective, the former should have no clear advantage in terms of the performance of the resulting β -estimator. One might therefore expect all heteroscedasticity-accommodating methods to perform similarly here since they need only model the conditional variance σ_0 sufficiently well to yield the semiparametrically optimal MSE (referred to as the oracle MSE).

The top left panel of Figure 4 shows that this appears to be the case when $\lambda \leq 0.75$, but for larger values of the complexity parameter, both the ML and GEE approaches appear to struggle. The latter displays a somewhat erratic trajectory, peaking at an MSE 9 times that of the oracle, and even greatly exceeding that of its homoscedastic counterpart (note that the curves for the homoscedastic GEE and ML estimators almost coincide). This behaviour seems likely to be due to the GEE approach finding local optima, which recall, need not be local optima for the asymptotic MSE objective, i.e., the sandwich loss; see also Example 1 Setting (b) and Appendix A.1 for a similar phenomenon. The sandwich boosted estimator in comparison remains relatively robust to this increase in model complexity, maintaining performance comparable to the oracle estimator.

5.1.2 Increasing covariance misspecification

Here we consider $I = 10^4$ instances of the partially linear model (6) with $n = 4$, $\beta = 1$, $g_0(x) = \tanh x$ and $m_0(x) = \cos x$. Errors (ε, ξ) are generated by introducing an unobserved confounder ζ between ε and ξ inspired by the proof of Theorem 2 as follows:

$$X \sim N_4(\mathbf{0}, 0.9\mathbf{1}\mathbf{1}^\top + 0.1I_4), \quad B | X \sim \text{Ber}(p(X)), \quad p(x) := \mathbb{1}_{[0,\infty)}(\bar{x}) + \eta^{-1}\mathbb{1}_{(-\infty,0)}(\bar{x}),$$

$$\zeta := p(X)^{-1}B, \quad \xi | (X, \zeta) := N_4(\mathbf{0}, \zeta I_4), \quad \varepsilon | (D, X, \zeta) := N_4(\mathbf{0}, \zeta \Sigma), \quad \Sigma = (0.2^{|j-k|})_{(j,k) \in [4]^2}.$$

Here $\eta \geq 1$ acts as ‘misspecification parameter’, with larger values indicating greater confounding and violation of the condition (11) for equivalence of the ML, GEE and sandwich losses. Note that \bar{x} denotes the mean of the entries $x \in \mathbb{R}^n$.

As we see in the top right panel of Figure 4, the performances of the heteroscedastic ML and GEE estimators deteriorate with increasing η and yield worse MSEs than even an unweighted least squares estimator as the extent of covariance misspecification increases. In contrast, despite being equally restricted to use a misspecified class of weights that are a function of X alone, the sandwich boosted $\hat{\beta}$ has a substantially smaller MSE compared to the approaches considered, with its advantage increasing with increasing η .




Figure 4: Mean squared error of β -estimators for the four simulated experiments in Section 5.1 over 500 independent repeats. The top left plot gives the MSE relative to the oracle, while the others are relative to the unweighted estimator.

5.1.3 Mild conditional correlation misspecification

We now consider the simple setting of (grouped) linear regression:

$$Y = \beta D + \varepsilon, \quad D \sim N_n \left(\mathbf{0}, \frac{1}{8} \mathbf{1}\mathbf{1}^\top + \frac{7}{8} I_n \right),$$

$$\varepsilon \sim \text{ARMA}(2, 1), \quad \phi = (0.3, 0.6), \quad \vartheta = -0.5.$$

Here ϕ and ϑ are the autoregressive and moving average parameters respectively. We take $\beta = 1$ and i.i.d. Gaussian innovations for the ARMA process. We consider settings with $I = 2^{15}/n$ and group sizes $n \in \{2^r : r = 1, 2, \dots, 8\}$. To demonstrate the effect of correlation misspecification, we use a constant working variance and fit all models with an AR(1) working correlation; this yields a misspecified correlation for group sizes $n \geq 3$. Example 1, Setting (a) corresponds to this setting with $n = 100$. The bottom left panel of Figure 4 shows that the sandwich loss here outperforms the competing approaches, with the GEE approach leading to an inflation in MSE over an unweighted approach for moderate group sizes.

5.1.4 Mild conditional variance misspecification

We now consider $I = 2^{15}/n$ i.i.d. instances of the the partially linear model (6) via

$$X_j \stackrel{\text{i.i.d.}}{\sim} U[-2, 2], \quad m_0(x) = -6e^{-x}, \quad g_0(x) = \tanh(x), \quad \xi | X \sim N_n(\mathbf{0}, 9I_n)$$

$$\varepsilon | (D, X) \sim N_n(\mathbf{0}, \Sigma(D, X)), \quad \Sigma_{jk}(D, X) = \begin{cases} \sigma_0^2(D_j, X_j) & \text{if } j = k \\ 0.2\sigma_0(D_j, X_j)\sigma_0(D_k, X_k) & \text{if } j \neq k \end{cases},$$

$$\sigma_0(d, x) = 2 + \tanh(d - 3x),$$

for $n \in \{2^r : r = 1, 2, \dots, 8\}$, taking $\beta = 1$. Note therefore that as $\text{Cov}(\varepsilon | D, X) \neq \text{Cov}(\varepsilon | X)$, the optimal weights that depend also on D are not in the weight classes considered. All methods are fitted using an equicorrelated working correlation. We again see in the bottom right Figure 4 superior performance of sandwich boosting over the competitors here.

5.2 Real-world data analyses

Here we present analyses of two datasets. We fit GEE and (heteroscedastic) sandwich boosting approaches as in the previous section, but use mixed effects models (MEM) to give a family of (working) conditional covariance functions in the ML framework by taking certain covariates as random effects. We continue to use these within Algorithm 1, using the `lme4` package [Bates et al., 2015] to obtain the weights in the MEM case as in Emmenegger and Bühlmann [2023], but reporting robust sandwich estimates of the variance of the $\hat{\beta}$ constructed. We use $K = 5$ folds for cross-fitting. To mitigate the randomness of the resulting estimators on the sample splits themselves we aggregate the β and variance estimators obtained over 50 random independent sample splits using the approach of Chernozhukov et al. [2018], Emmenegger and Bühlmann [2023]; see Appendix E for details.

5.2.1 Orange juice price elasticity

We analyse historical data on orange juice sales, available from the James M. Kilts Center, University of Chicago Booth School of Business [James M. Kilts Center, Accessed: 2022]. The dataset is composed of grouped store-level scanner price and sales data over a 121 week period from 83 Dominick's Finer Foods stores and consists of $N = 9649$ observations.

Our goal is to estimate the price elasticity of a brand of orange juice (Tropicana) during this time period. We do this via a partially linear model (6) of the logarithm of the quantity of sales ($Y = \log(\text{Sales})$) on the logarithm of the price ($D = \log(\text{Price})$) accounting for confounding by events in time ($X = \text{week number}$), the coefficient β of D giving the price elasticity.

Method	$\hat{\beta}$	\hat{V}	Reduction in \hat{V} relative to homoscedastic GEE estimator (%)
Sandwich Boosting	-2.97	30.9	26.0
Homoscedastic GEE	-3.18	41.7	0
Heteroscedastic GEE	-3.19	44.6	-7.0
Intercept only MEM	-3.17	41.8	-0.17
Intercept + Time MEM	-3.18	42.5	-2.0

Table 2: Estimates for the price elasticity of orange juice and corresponding variance estimates.

Table 2 shows the price elasticity estimators and associated sandwich variances estimates \hat{V} ; note GEE and sandwich boosting approaches here use an equicorrelated working correlation. We see that our sandwich boosting estimator has a 26.0% reduction in variance compared to the second best homoscedastic GEE estimator. Note that in the cases of both the mixed effects models (MEM) and GEE estimators, as a broader weight class is used, we observe poorer performance in estimating the price elasticity, illustrating the phenomenon described in Figure 3. Also note that the sandwich boosted and heteroscedastic GEE estimators model weights in the same class. However, whilst the heteroscedastic GEE estimator does not usefully model any heteroscedasticity in the data (and in fact exhibits the worst performance out of all the estimators considered), sandwich boosting successfully estimates a helpful weighting scheme. Figure 5a demonstrates an s -function output from sandwich boosting, effectively capturing the general and seasonal trends in volatility that are not learnt by the other estimators.

5.2.2 National longitudinal survey of young working women


Here we consider a dataset from the National Longitudinal Survey of Young Women [Bureau of Labor Statistics, 2004] containing the wages of 4711 young working women, each measured at approximately 6 time points per woman, totalling $N = 28\,534$ observations. We measure the effect of work experience in their current related sector ($D = \text{work experience}$) on the logarithm of wages ($Y = \log(\text{wage})$), controlling for age and tenure ($X = (\text{age}, \text{tenure})$), with weights a function of X , using an equicorrelated working correlation for sandwich boosting and GEE approaches.

Table 3 gives the β -estimators and associated variances \hat{V} for the approaches considered. Similarly to the orange juice price elasticity analysis, we see that the broader model classes of the mixed effects model and heteroscedastic GEE estimators yield larger variances than their respective homoscedastic counterparts. The sandwich boosting estimator gives a modest 7.4% reduction in variance over the second smallest homoscedastic GEE variance. Interestingly however, this improvement is entirely due to the sandwich loss rather than the potentially richer model class used by sandwich boosting: the s function output by sandwich boosting is almost constant, and so it effectively uses a constant working variance and estimates a single working correlation parameter. The sandwich loss objective for this correlation (parametrised in terms of θ ; see Section 3.2) is plotted in Figure 5b alongside the objectives corresponding to the homoscedastic GEE and intercept only MEM approaches. We see that the respective


Method	$\hat{\beta}$	\hat{V}	Reduction in \hat{V} relative to homoscedastic GEE estimator (%)
Sandwich Boosting	0.040	0.083	7.4
Homoscedastic GEE	0.040	0.089	0
Heteroscedastic GEE	0.039	0.091	-1.6
Intercept only MEM	0.040	0.092	-3.5
Intercept + Age + Tenure MEM	0.042	0.120	-34.9

Table 3: Estimates of β and associated variance estimates \hat{V} relating to the effect of work experience on wages from the National Longitudinal Survey of Young Women.

θ -minimisers differ substantially, with the asymptotic variance (the sandwich loss) evaluated at the minimisers of the GEE and MEM approaches being larger than the asymptotic variance evaluated at any $\theta \in [0.14, 1.09]$, corresponding to any working correlation in the range $\rho \in [0.12, 0.51]$.



(a) Sandwich boosted \hat{s} -function output for the orange juice price elasticity dataset.



(b) The (scaled) objective functions used to determine θ for each of the intercept only mixed effects model (ML), homoscedastic GEE estimator (GEE) and sandwich loss with equicorrelated working correlation.

Figure 5: Outputs relating to data analyses in Section 5.2.

6 Discussion

In this work we have highlighted and clarified the shortcomings of some popular classical methods in the estimation of weights for weighted least squares-type estimators in partially linear models when the conditional covariance is misspecified. We instead advocate for choosing weights to minimise a sandwich estimate of the variance, what we call the sandwich loss in this context. A main contribution of ours, in the spirit of the trend towards using machine learning

methods for the purposes of statistical inference, is a practical gradient boosting scheme for approximately minimising this loss over a potentially flexible family of functions defined implicitly through a user-chosen base-learner. Despite the unusual form of our loss that does not decompose as a sum over data points as with the standard case of the empirical risk, we show that for certain versions of our algorithm, the boosting updates can be performed in linear time.

Our work offers a number of directions for future research. On the computational side, it would be useful to investigate broader classes of working correlations that could be accommodated within sandwich boosting to yield linear time updates. It could also be fruitful to consider the use of the sandwich loss in other classes of models, for example it would be of interest to develop these ideas in the context of generalised (partially) linear marginal models, and beyond.

Thus far we have only considered estimators of a single scalar quantity. In other situations, one may be interested in estimating several parameters simultaneously, and in such cases there are several modifications of the basic sandwich loss that may be helpful to explore. For example, consider the following generalisation of the partially linear model (6):

$$\begin{aligned} Y &= \beta(X) \circ D + g_0(X) + \varepsilon, \\ D &= m_0(X) + \xi. \end{aligned} \tag{18}$$

Here \circ denotes the Hadamard product, $\beta(X)$ is a row-wise function of X , and all other terms are as before; thus the response of the j th observation within the i th group satisfies (with a slight abuse of notation) the mean function relationship $\mathbb{E}[Y_j | D, X] = \beta(X_j)D_j + g_0(X_j)$. We suppose $\beta(X)$ admits the basis expansion

$$\beta(X) = \sum_{l=1}^L \phi_l \varphi_l(X), \tag{19}$$

for some L known basis functions $(\varphi_l)_{l \in [L]}$ (also row-wise functions of X), and unknown vector of parameters $\phi := (\phi_1, \dots, \phi_L) \in \mathbb{R}^L$ to be estimated. For example, the model where $\beta(X) = \phi_1 + \phi_2 X$ corresponds in the classical linear model setting to fitting an ‘interaction term’ between D and X . Given a consistent estimator $\hat{\phi}$ of ϕ , the mean squared error of the resulting plug-in β function estimator $\hat{\beta}(X) = \sum_{l=1}^L \hat{\phi}_l \varphi_l(X)$ satisfies

$$\mathbb{E} \left[\left(\hat{\beta}(X) - \beta(X) \right)^2 \right] \approx \text{tr}(\Phi V_\phi), \quad \Phi := \left(\mathbb{E}[\varphi_l(X) \varphi_{l'}(X)] \right)_{l \in [L], l' \in [L]} \quad \text{and} \quad V_\phi := \text{Var} \hat{\phi}. \tag{20}$$

This suggests the following approach. Consider a class of weighted ϕ -estimators $\hat{\phi} = \hat{\phi}(W)$ where W is a weight function among a class of \mathcal{W} , similarly to our previous framework in Section 2.1 of weighted β -estimators. Then given estimates $\hat{\Phi}$ of Φ and $\hat{V}_\phi(W)$ of $\text{Var} \hat{\phi}(W)$, we can consider a generalised sandwich loss of the form

$$\hat{L}_{\text{MSE}(\hat{\beta})}(W) := \text{tr}(\hat{\Phi} \hat{V}_\phi(W)), \tag{21}$$

which we may attempt to minimise using a sandwich boosting approach; further details are given in Appendix G.

Our sandwich loss \hat{L}_{SL} is defined with respect to estimated errors $\tilde{\xi}_i$ and $\tilde{\varepsilon}_i$ derived from initial regressions, which in particular, take no advantage of the dependence structure in the data, unlike the final estimate $\hat{\beta}$. Clearly initial weighted regressions could deliver improved estimates of the errors, in turn giving an improved estimate of $\hat{\beta}$. This suggests a scheme with weights and residuals being updated iteratively, analogous to iterative generalised least squares [Goldstein, 1986, 1989]. In the simple linear model setting, where the residuals are derived from linear regressions, a generalised sandwich loss of the form (21) may be appropriate for delivering accurate estimates of the errors. How to do this for a general regression is less clear but would certainly be worthy of further investigation.

References

- D. Bates, M. Mächler, B. Bolker, and S. Walker. Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1):1–48, 2015. doi: 10.18637/jss.v067.i01.
- C. S. Bojer and J. P. Meldgaard. Kaggle forecasting competitions: An overlooked learning opportunity. *International Journal of Forecasting*, 37(2):587–603, 2021. ISSN 0169-2070. doi: <https://doi.org/10.1016/j.ijforecast.2020.07.007>. URL <https://www.sciencedirect.com/science/article/pii/S0169207020301114>.
- G. E. P. Box. Science and statistics. *Journal of the American Statistical Association*, 71(356): 791–799, 1976. doi: 10.1080/01621459.1976.10480949.
- L. Breiman. Prediction Games and Arcing Algorithms. *Neural Computation*, 11(7):1493–1517, 10 1999. doi: 10.1162/089976699300016106.
- P. Bühlmann and T. Hothorn. Boosting Algorithms: Regularization, Prediction and Model Fitting. *Statistical Science*, 22(4):477 – 505, 2007. doi: 10.1214/07-STS242.
- P. Bühlmann and B. Yu. Boosting with the l2 loss. *Journal of the American Statistical Association*, 98(462):324–339, 2003. doi: 10.1198/016214503000125.
- U. D. o. L. Bureau of Labor Statistics. National longitudinal survey of young women, 1968-1988 (rounds 1-15). Produced and distributed by the Center for Human Resource Research (CHRR), The Ohio State University, 2004. URL <https://www.stata-press.com/data/r10/nlswork.dta>.
- R. J. Carroll. Adapting for Heteroscedasticity in Linear Models. *The Annals of Statistics*, 10 (4):1224 – 1233, 1982. doi: 10.1214/aos/1176345987.
- T. Chen and C. Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 785–794. Association for Computing Machinery, 2016. doi: 10.1145/2939672.2939785.
- V. Chernozhukov, D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, and J. Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 2018. doi: 10.1111/ectj.12097.
- R. R. Corbeil and S. R. Searle. Restricted maximum likelihood (reml) estimation of variance components in the mixed model. *Technometrics*, 18(1):31–38, 1976.
- M. Crowder. On the use of a working correlation matrix in using generalised linear models for repeated measures. *Biometrika*, 82(2):407–410, 1995.
- P. Diggle, P. Heagerty, K.-Y. Liang, and S. Zeger. *Analysis of Longitudinal Data*, volume Second edition of *Oxford Statistical Science Series*. OUP Oxford, Oxford, 2013.
- C. Emmenegger. *dmlalg: Double machine learning algorithms*, 2021. URL <https://CRAN.R-project.org/package=dmlalg>. R-package available on CRAN.
- C. Emmenegger and P. Bühlmann. Regularizing double machine learning in partially linear endogenous models. *Electronic Journal of Statistics*, 15(2):6461–6543, 2021.
- C. Emmenegger and P. Bühlmann. Plug-in machine learning for partially linear mixed-effects models with repeated measurements. *Scandinavian Journal of Statistics*, 2023. doi: 10.1111/sjos.12639.

- L. Fahrmeir and G. Tutz. *Multivariate Statistical Modelling Based on Generalized Linear Models*. Springer Series in Statistics. Springer, 2 edition, 2001.
- Y. Freund and R. E. Schapire. Experiments with a new boosting algorithm. In *International Conference on Machine Learning*, pages 148–156, 1996.
- J. Friedman, T. Hastie, and R. Tibshirani. Additive logistic regression: a statistical view of boosting (with discussion and a rejoinder by the authors). *The Annals of Statistics*, 28(2): 337 – 407, 2000. doi: 10.1214/aos/1016218223.
- H. Goldstein. Multilevel mixed linear model analysis using iterative generalized least squares. *Biometrika*, 73(1):43–56, 1986.
- H. Goldstein. Restricted unbiased iterative generalized least-squares estimation. *Biometrika*, 76(3):622–623, 1989.
- C. Gourieroux and A. Monfort. Pseudo-likelihood methods. In *Econometrics*, volume 11 of *Handbook of Statistics*, pages 335–362. Elsevier, 1993.
- C. Gourieroux, A. Monfort, and A. Trognon. Pseudo maximum likelihood methods: Theory. *Econometrica*, 52(3):681–700, 1984.
- U. Halekoh, S. Højsgaard, and J. Yan. The r package geepack for generalized estimating equations. *Journal of Statistical Software*, 15/2:1–11, 2006.
- J. W. Hardin and J. M. Hilbe. *Generalized estimating equations*. Chapman and Hall, 2003.
- H. O. Hartley and J. N. K. Rao. Maximum-likelihood estimation for the mixed analysis of variance model. *Biometrika*, 54(1/2):93–108, 1967.
- P. J. Heagerty and S. L. Zeger. Marginalized multilevel models and likelihood inference. *Statistical Science*, 15(1):1–19, 2000.
- J. Z. Huang, L. Zhang, and L. Zhou. Efficient estimation in marginal partially linear models for longitudinal/clustered data using splines. *Scandinavian Journal of Statistics*, 34(3):451–477, 2007.
- P. J. Huber. The behaviour of maximum likelihood estimates under nonstandard conditions. *Proceedings of the Fifth Berkeley Symposium*, pages 221–223, 1967.
- U. o. C. B. S. o. B. James M. Kilts Center. Dominick’s finer foods dataset. Available from Chicago Booth Research Data Center: <https://www.chicagobooth.edu/research/kilts/research-data/dominicks>, Accessed: 2022.
- E. H. Kennedy. Semiparametric doubly robust targeted double machine learning: a review. *arXiv preprint arXiv:2203.06469*, 2022.
- S. Li, T. T. Cai, and H. Li. Inference for high-dimensional linear mixed-effects models: A quasi-likelihood approach. *Journal of the American Statistical Association*, 117(540):1835–1846, 2022.
- Y. Li, S. Wang, P. X.-K. Song, N. Wang, L. Zhou, and J. Zhu. Doubly regularized estimation and selection in linear mixed-effects models for high-dimensional longitudinal data. *Statistics and its Interface*, 11(4):721, 2018.

- K.-Y. Liang and S. L. Zeger. Longitudinal data analysis using generalized linear models. *Biometrika*, 73(1):13–22, 1986.
- K.-Y. Liang, S. L. Zeger, and B. Qaqish. Multivariate regression analyses for categorical data. *Journal of the Royal Statistical Society. Series B (Methodological)*, 54(1):3–40, 1992.
- T. Lumley. Generalized estimating equations for ordinal data: A note on working correlation structures. *Biometrics*, 52(1):354–361, 1996.
- L. Mason, J. Baxter, P. Bartlett, and M. Frean. Boosting algorithms as gradient descent. In *Advances in Neural Information Processing Systems*, volume 12. MIT Press, 1999.
- P. McCullagh and J. A. Nelder. *Generalized linear models*. Monographs on statistics and applied probability (Series) ; 37. Chapman and Hall, London, 2 edition, 1989.
- C. Park and H. Kang. A more efficient, doubly robust, nonparametric estimator of treatment effects in multilevel studies. *arXiv preprint arXiv:2110.07740*, 2021.
- J. Pinheiro, D. Bates, and R Core Team. *nlme: Linear and Nonlinear Mixed Effects Models*, 2022. URL <https://CRAN.R-project.org/package=nlme>. R package version 3.1-161.
- J. C. Pinheiro and D. M. Bates. *Mixed-Effects Models in S and S-PLUS*, volume 1 of *Springer Statistics and Computing*. Springer, New York, 2000.
- R. L. Prentice and L. P. Zhao. Estimating equations for parameters in means and covariances of multivariate discrete and continuous responses. *Biometrics*, 47(3):825–839, 1991.
- J. M. Robins and A. Rotnitzky. Semiparametric efficiency in multivariate regression models with missing data. *Journal of the American Statistical Association*, 90(429):122–129, 1995.
- J. M. Robins, A. Rotnitzky, and L. P. Zhao. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American statistical Association*, 89(427):846–866, 1994.
- P. Robinson. Asymptotically efficient estimation in the presence of heteroskedasticity of unknown form. *Econometrica*, 55(4):875–891, 1987.
- R. M. Royall. Model robust confidence intervals using maximum likelihood estimators. *International Statistical Review / Revue Internationale de Statistique*, 54(2):221–226, 1986.
- R. E. Schapire. The strength of weak learnability. *Machine Learning*, 5(2):197–227, 1990. doi:10.1007/BF00116037.
- A. A. Tsiatis. *Semiparametric theory and missing data*, volume 1 of *Springer series in statistics*. Springer, New York, 2006.
- A. W. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 1998.
- S. Vansteelandt and O. Dukes. Assumption-lean inference for generalised linear model parameters. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(3):657–685, 2022.
- S. Wood. *Generalized Additive Models: An Introduction with R*. Chapman and Hall/CRC, 2 edition, 2017.

J. You, G. Chen, and Y. Zhou. Statistical inference of partially linear regression models with heteroscedastic errors. *Journal of Multivariate Analysis*, 98(8):1539–1557, 2007. doi: 10.1016/j.jmva.2007.06.011.

S. L. Zeger and P. J. Diggle. Semiparametric models for longitudinal data with application to cd4 cell numbers in hiv seroconverters. *Biometrics*, pages 689–699, 1994.

S. L. Zeger and K.-Y. Liang. Longitudinal data analysis for discrete and continuous outcomes. *Biometrics*, 42(1):121–130, 1986.

A. Ziegler. *Generalized estimating equations*. Lecture notes in statistics (Springer-Verlag). Springer, New York, 2011.

List of Figures

1	The (scaled) objective functions of the parameter ρ in the AR(1) working model for quasi-pseudo-Gaussian maximum likelihood (ML), GEE1 (GEE) and the corresponding asymptotic MSE relative to the minimum asymptotic MSE, in a settings with a ground truth given by an ARMA(2, 1) model of Example 1; coloured intervals at the bottom indicate ranges of ρ where the (global) minimisers corresponding to ML and GEE would be outperformed in terms of MSE.	3
2	Asymptotic MSEs of β -estimators relative to an unweighted estimator, using weights in a misspecified model class (5) estimated by each of ML and GEE in the settings of Example 2 parametrised by (λ, μ)	4
3	Schematic of how the ML ‘projections’ $W_{ML,1}^*$ and $W_{ML,2}^*$ of the optimal weight function $W_0 \propto \text{Cov}(Y D, X)^{-1}$ onto the nested weight classes $\mathcal{W}_1 \subset \mathcal{W}_2$ may be such that $W_{ML,1}^*$ in fact improves on $W_{ML,2}^*$ when the contours of the negative likelihood loss (dotted blue ellipses) do not align with those of the asymptotic MSE; the contours of the latter are here to be thought of as circles centred at W_0 (not shown), so that the distances, $V_1 < V_2$, are the variances of $\hat{\beta}$ constructed using weights $W_{ML,1}^*$ and $W_{ML,2}^*$ respectively. The projections $W_{SL,1}^*$ and $W_{SL,2}^*$ corresponding to the loss given by the asymptotic MSE itself (‘sandwich loss’) do not suffer from any such paradoxes.	5
4	Mean squared error of β -estimators for the four simulated experiments in Section 5.1 over 500 independent repeats. The top left plot gives the MSE relative to the oracle, while the others are relative to the unweighted estimator.	21
5	Outputs relating to data analyses in Section 5.2.	24

Supplementary material relating to ‘Sandwich Boosting for Accurate Estimation in Partially Linear Models for Grouped Data’, by Elliot H. Young and Rajen D. Shah

A Supplementary material relating to Section 1

A.1 Further details for Example 1

Recall that in Example 1 we consider the grouped linear model

$$Y_i = \beta D_i + \varepsilon_i,$$

where $(Y_i, D_i) \in \mathbb{R}^n \times \mathbb{R}^n$ are iid; note that since we plot the population level objective functions of the ML and GEE approaches, along with the asymptotic variance of the resulting $\hat{\beta}$ for each given value of ρ , there is no need to specify the total number of groups I . In Setting (a), we take $n = 100$ and ε_i as the first n realisations of an ARMA(2, 1) model with unit variance and autoregressive parameters $\phi = (0.3, 0.6)$ and moving average parameter $\vartheta = -0.5$. Further, we take $D_i \sim N_n(\mathbf{0}, \frac{1}{8}\mathbf{1}\mathbf{1}^\top + \frac{7}{8}I_n)$.

In Setting (b), we take $n = 30$ and the ARMA(2, 1) model parameters $\phi = (0.1, 0.85)$, $\vartheta = -0.4$. Compared to the Setting (a), the AR(1) working correlation model is less severely misspecified. In this setting, the $\hat{\rho}$ obtained by minimising the ML objective results in a $\hat{\beta}$ that has the same asymptotic mean squared error (MSE) as the unweighted estimator (obtained by setting $\rho = 0$ in this working correlation model, i.e. not attempting to model the correlation at all). The GEE objective is also non-convex in ρ , and in fact a gradient descent procedure minimising the GEE objective initialised at $\rho = 0$ (or a Newton optimisation approach as used in the `geepack` package [Halekoh et al., 2006]) will incorrectly estimate $\hat{\rho}$ at the local minimum at $\hat{\rho} = -0.71$. Using this estimate for ρ results in a $\hat{\beta}$ with MSE 2.4 times that of the unweighted estimator (and 3.4 times that of the $\hat{\beta}$ resulting from minimising the sandwich loss); see Table 4.

Objective function	$\hat{\rho}$	Asymptotic MSE of $\hat{\beta}$ relative to sandwich loss
Unweighted	0.00	1.4
GEE (<code>geepack</code>)	-0.71	3.4
ML (<code>gls</code>)	0.00	1.4
Sandwich loss (<code>sandwich.boosting</code>)	0.30	1.0

Table 4: Mean squared errors of the weighted β -estimators weighted using ML (implemented in `gls`), GEE1 (`geepack`) and gradient descent of the sandwich loss (`sandwich.boosting`) for the simulation example in Section A.1.

A.2 Further details on Example 2

We give further details on misspecified conditional variance motivating example in Section 1.1. Note specifically this example compares the three loss functions considered at the population level (and not the sample level) therefore exploring asymptotic properties of each loss function. We consider the ungrouped ($n_i \equiv 1$) setting where the random variables $(Y, D, X) \in \mathbb{R} \times \mathbb{R} \times \mathbb{R}$ satisfy

$$X \sim U[0, 1], \quad D = \xi, \quad Y = D\beta + \epsilon,$$

for some (arbitrary) $\beta \in \mathbb{R}$ and where $\mathbb{E}[\varepsilon | D, X] = \mathbb{E}[\xi | X] = 0$, $\mathbb{E}[\xi^2 | X] = 1$ and $\sigma_0^2(X) := \text{Var}(\varepsilon | D, X) = 2 + \tanh(\lambda(X - \mu))$ for constants $\lambda \geq 0$ and $\mu \in [0, 1]$.

We consider using a misspecified class of functions of the form

$$\sigma(x; \eta) = 1 + 2\mathbb{1}_{[\eta, \infty)}(x),$$

where $\eta \in \mathbb{R}$, with estimates $\hat{\eta}$ of an optimal η minimising one of the ML, GEE or SL losses (each method discussed in Section 1). At the population level for this example these losses can be expressed as

$$L_{\text{ML}}(\eta) = \mathbb{E} \left[\log \left(\sigma^2(X; \eta) + \frac{\sigma_0^2(X)}{\sigma^2(X; \eta)} \right) \right], \quad L_{\text{GEE}}(\eta) = \mathbb{E} \left[(\sigma_0^2(X) - \sigma^2(X; \eta))^2 \right],$$

$$\text{and } L_{\text{SL}}(\eta) = \left(\mathbb{E} \left[\frac{1}{\sigma^2(X; \eta)} \right] \right)^{-2} \left(\mathbb{E} \left[\frac{\sigma_0^2(X)}{\sigma^4(X; \eta)} \right] \right).$$

The model class considered is therefore parametrised by a single parameter η , with each loss function minimised either analytically or via a line search.

Given a minimising $\hat{\eta}$ of one the three losses, the (scaled) mean squared errors of the σ_0 -estimator and β -estimator (weighted using the σ_0 -estimator) can be calculated analytically as

$$\mathbb{E} \left[(\sigma(X; \hat{\eta}) - \sigma_0(X))^2 \right], \text{ and } \left(\mathbb{E} \left[\frac{1}{\sigma^2(X; \hat{\eta})} \right] \right)^{-2} \left(\mathbb{E} \left[\frac{\sigma_0^2(X)}{\sigma^4(X; \hat{\eta})} \right] \right),$$

respectively, with results as given in Table 1

B Proof of results in Section 3

B.1 Proof of Proposition 1

We first show that L_{GEE} is minimised over \mathcal{W} by $\mathbb{E}[\varepsilon\varepsilon^\top | q(X)]^{-1}$; this follows immediately from the decomposition of $L_{\text{GEE}}(W)$ as

$$\begin{aligned} \mathbb{E} \left[\|W(X)^{-1} - \varepsilon\varepsilon^\top\|^2 \right] &= \mathbb{E} \left[\|W(X)^{-1} - \mathbb{E}[\varepsilon\varepsilon^\top | q(X)]\|^2 \right] + \mathbb{E} \left[\|\varepsilon\varepsilon^\top\|^2 - \|\mathbb{E}[\varepsilon\varepsilon^\top | q(X)]\|^2 \right] \\ &\quad - 2\mathbb{E} \left[\mathbb{E} \left[\langle \varepsilon\varepsilon^\top - \mathbb{E}[\varepsilon\varepsilon^\top | q(X)], W(X)^{-1} \rangle \mid q(X) \right] \right], \\ &= \mathbb{E} \left[\|W(X)^{-1} - \mathbb{E}[\varepsilon\varepsilon^\top | q(X)]\|^2 \right] + \mathbb{E} \left[\|\varepsilon\varepsilon^\top\|^2 - \|\mathbb{E}[\varepsilon\varepsilon^\top | q(X)]\|^2 \right]. \end{aligned}$$

Note this result holds for any matrix norm derived from an inner product.

We now show that L_{ML} is minimised over \mathcal{W} by $\mathbb{E}[\varepsilon\varepsilon^\top | q(X)]^{-1}$. Noting that $L_{\text{ML}}(W) = \mathbb{E}[\ell_{\text{ML}}(W(X))]$ where

$$\ell_{\text{ML}}(W(X)) := -\log \det W(X) + \text{tr} \left(W(X) \mathbb{E}[\varepsilon\varepsilon^\top | q(X)] \right),$$

we aim to minimise the function $x \mapsto \ell_{\text{ML}}(W(x))$ for arbitrary $x \in \mathbb{R}^{n \times d}$. We define $\bar{W} \in \mathcal{W}$ by $\bar{W}(x) := \mathbb{E}[\varepsilon\varepsilon^\top | q(X) = q(x)]^{-1}$. Then given some fixed $x \in \mathbb{R}^{n \times d}$ and $W \in \mathcal{W}$ we define the function $g(t) = \ell_{\text{ML}}((1-t)\bar{W}(x) + tW(x))$ for $t \in [0, 1]$. The second order Taylor expansion of g then gives

$$\ell_{\text{ML}}(W(x)) = \ell_{\text{ML}}(\bar{W}(x)) + g'(0) + \frac{g''(\tilde{t})}{2},$$

for some $\tilde{t} \in [0, 1]$. Now, by Jacobi's formula,

$$g'(t) = -\text{tr} \left(\left[\{(1-t)\bar{W}(x) + tW(x)\}^{-1} - \mathbb{E}[\varepsilon\varepsilon^\top | q(X) = q(x)] \right] \{W(x) - \bar{W}(x)\} \right),$$

and so

$$g'(0) = -\text{tr} \left(\left\{ \tilde{W}(x)^{-1} - \mathbb{E}[\varepsilon\varepsilon^\top | q(X) = q(x)] \right\} \{W(x) - \bar{W}(x)\} \right) = 0.$$

Further, by direct calculation

$$\begin{aligned} g''(\tilde{t}) &= \text{tr} \left(\left(\{(1-\tilde{t})\bar{W}(x) + \tilde{t}W(x)\} \{W(x) - \bar{W}(x)\} \right)^2 \right) \\ &= \left\| \{(1-\tilde{t})\bar{W}(x) + \tilde{t}W(x)\}^{\frac{1}{2}} \{W(x) - \bar{W}(x)\} \{(1-\tilde{t})\bar{W}(x) + \tilde{t}W(x)\}^{\frac{1}{2}} \right\|_{\text{F}}^2 \geq 0, \end{aligned}$$

where $\|\cdot\|_{\text{F}}$ denotes the Frobenius norm, and with equality holding if and only if $W(x) = \bar{W}(x)$. Therefore whenever $W(x) \neq \bar{W}(x)$,

$$\ell_{\text{ML}}(W(x)) > \ell_{\text{ML}}(\tilde{W}(x)),$$

and so L_{ML} is minimised over \mathcal{W} by $\mathbb{E}[\varepsilon\varepsilon^\top | q(X)]^{-1}$.

Finally, we will show that

$$\begin{aligned} V_{\text{SL}} &:= \inf_{W \in \mathcal{W}} L_{\text{SL}}(W) \\ &= \left(\mathbb{E} \left[\mathbb{E}[\text{vec}(\xi\xi^\top) | q(X)] \mathbb{E}[\text{vec}(\xi\varepsilon^\top)\text{vec}(\varepsilon\xi^\top)^\top | q(X)]^{-1} \mathbb{E}[\text{vec}(\xi\xi^\top) | q(X)] \right] \right)^{-1}. \end{aligned}$$

This can be shown by noting that for $W \in \mathcal{W}$

$$\begin{aligned} L_{\text{SL}}(W) &= \frac{\mathbb{E}[\xi^\top W \varepsilon \varepsilon^\top W \xi]}{(\mathbb{E}[\xi^\top W \xi])^2} = \frac{\mathbb{E}[\text{vec}(W)^\top \mathbb{E}[\text{vec}(\xi\varepsilon^\top)\text{vec}(\varepsilon\xi^\top)^\top | q(X)] \text{vec}(W)]}{\left(\mathbb{E}[\mathbb{E}[\text{vec}(\xi\xi^\top) | q(X)]^\top \text{vec}(W)] \right)^2} \\ &\geq \left(\mathbb{E} \left[\mathbb{E}[\text{vec}(\xi\xi^\top) | q(X)] \mathbb{E}[\text{vec}(\xi\varepsilon^\top)\text{vec}(\varepsilon\xi^\top)^\top | q(X)]^{-1} \mathbb{E}[\text{vec}(\xi\xi^\top) | q(X)] \right] \right)^{-1} \\ &= \left(\mathbb{E} \left[\mathbb{E}[\text{vec}(\xi\xi^\top) | q(X)] \mathbb{E}[(\varepsilon\xi^\top) \otimes (\xi\varepsilon^\top) | q(X)]^{-1} \mathbb{E}[\text{vec}(\xi\xi^\top) | q(X)] \right] \right)^{-1}, \end{aligned}$$

by the Cauchy–Schwarz inequality, and by using the identities $a^\top Ab = \text{vec}(ab^\top)^\top \text{vec}(A)$ and $\text{vec}(ab^\top)\text{vec}(cd^\top)^\top = (bd^\top) \otimes (ac^\top)$ (for vectors $a, b, c, d \in \mathbb{R}^n$ and matrix $A \in \mathbb{R}^{n \times n}$), where \otimes denotes the Kronecker product. Equality holds in the above if and only if

$$W \in \left\{ W \in \mathcal{W} : \text{vec}(W(X)) = c \mathbb{E}[(\varepsilon\xi^\top) \otimes (\xi\varepsilon^\top) | q(X)]^{-1} \mathbb{E}[\text{vec}(\xi\xi^\top) | q(X)] \text{ for some } c > 0 \right\}.$$

Therefore $V_{\text{SL}} \leq V_{\text{ML}/\text{GEE}}$ with equality if and only if the condition

$$\mathbb{E}[(\varepsilon\xi^\top) \otimes (\xi\varepsilon^\top) | q(X)] \text{vec} \left(\mathbb{E}[\varepsilon\varepsilon^\top | q(X)]^{-1} \right) = c \text{vec} \left(\mathbb{E}[\xi\xi^\top | q(X)] \right),$$

holds for some $c > 0$. Using the identity $\text{vec}(ABC) = (C^\top \otimes A)\text{vec}(B)$ (for matrices $A, B, C \in \mathbb{R}^{n \times n}$) this condition is equivalent to

$$\mathbb{E}[(\xi\varepsilon^\top) \mathbb{E}[\varepsilon\varepsilon^\top | q(X)]^{-1} (\varepsilon\xi^\top) | q(X)] = c \mathbb{E}[\xi\xi^\top | q(X)].$$

Further,

$$\begin{aligned}\mathbb{E} \left[(\xi \xi^\top) \mathbb{E} [\varepsilon \varepsilon^\top | q(X)]^{-1} (\varepsilon \xi^\top) | q(X) \right] &= \mathbb{E} \left[\xi \operatorname{tr} [\operatorname{Cov}(\varepsilon | q(X))^{-1} \operatorname{Cov}(\varepsilon | D, X)] \xi^\top | q(X) \right] \\ &= \mathbb{E} \left[\operatorname{tr} [\operatorname{Cov}(\varepsilon | q(X))^{-1} \operatorname{Cov}(\varepsilon | D, X)] \xi \xi^\top | q(X) \right],\end{aligned}$$

and so this condition is equivalent to

$$\mathbb{E} \left[\operatorname{tr} [\operatorname{Cov}(\varepsilon | q(X))^{-1} \operatorname{Cov}(\varepsilon | D, X)] \xi \xi^\top | q(X) \right] = c \mathbb{E} [\xi \xi^\top | q(X)]. \quad (22)$$

In particular when $q(X) = X$ and $\operatorname{Cov}(\varepsilon | D, X) = \operatorname{Cov}(\varepsilon | X)$ the condition (22) holds.

B.2 Proof of Theorem 2

We will adopt the shorthand $\Sigma_q = \Sigma(q(X))$ and $\Omega_q = \Omega(q(X))$. As $q(X)$ is almost surely non-constant there exists measurable $\Gamma \subset \mathbb{R}^{n \times d}$ such that $\mathbb{P}(q(X) \in \Gamma) \in (0, 1)$. With this, we define the random variables $Q := \eta^{-1} \mathbb{1}_{\{q(X) \in \Gamma\}} + \mathbb{1}_{\{q(X) \in \Gamma^c\}}$ for some $\eta > 1$ to be determined later, and $\zeta := Q^{-1}B - 1$ where $B \sim \operatorname{Ber}(Q)$ with $B \perp\!\!\!\perp X | Q$. Further, define $\xi := \sqrt{1 + \zeta} R_\xi$ and $\varepsilon := \sqrt{1 + \zeta} R_\varepsilon$ where $R_\xi | (X, B) \sim N(0, \Omega_q)$ and $R_\varepsilon | (X, B, R_\xi) \sim N(0, \Sigma_q)$.

By construction we then have $\mathbb{E}[\zeta | X] = 0$ and hence $\mathbb{E}[\xi | X] = \mathbb{E}[\varepsilon | D, X] = 0$, $\mathbb{E}[\xi \xi^\top | q(X)] = \Omega_q$ and $\mathbb{E}[\varepsilon \varepsilon^\top | q(X)] = \Sigma_q$. Further, we can show that

$$\mathbb{E}[\zeta^2 | q(X)] = \frac{1 - Q}{Q},$$

and thus

$$\begin{aligned}\mathbb{E} \left[(\varepsilon \xi^\top) \otimes (\xi \varepsilon^\top) | q(X) \right] &= \mathbb{E} \left[(1 + \zeta)^2 (R_\varepsilon R_\xi^\top) \otimes (R_\xi R_\varepsilon^\top) | q(X) \right] \\ &= (1 + \mathbb{E}[\zeta^2 | q(X)]) \mathbb{E} \left[(R_\varepsilon R_\xi^\top) \otimes (R_\xi R_\varepsilon^\top) | q(X) \right] \\ &= Q^{-1} \mathbb{E} \left[(R_\varepsilon R_\xi^\top) \otimes (R_\xi R_\varepsilon^\top) | q(X) \right].\end{aligned}$$

Therefore, using $V_{\text{ML}}, V_{\text{GEE}}$ and V_{SL} as in the proof of Proposition 1 (also noting that $V_{\text{ML}} = V_{\text{GEE}}$),

$$\begin{aligned}\frac{V_{\text{ML/GEE}}}{V_{\text{SL}}} &= \left(\mathbb{E} \left[\operatorname{vec} \left(\mathbb{E}[\xi \xi^\top | q(X)] \right)^\top \operatorname{vec} \left(\mathbb{E}[\varepsilon \varepsilon^\top | q(X)]^{-1} \right) \right] \right)^{-2} \\ &\quad \cdot \mathbb{E} \left[\operatorname{vec} \left(\mathbb{E}[\varepsilon \varepsilon^\top | q(X)]^{-1} \right)^\top \mathbb{E} \left[(\varepsilon \xi^\top) \otimes (\xi \varepsilon^\top) | q(X) \right] \operatorname{vec} \left(\mathbb{E}[\varepsilon \varepsilon^\top | q(X)]^{-1} \right) \right] \\ &\quad \cdot \mathbb{E} \left[\operatorname{vec} \left(\mathbb{E}[\xi \xi^\top | q(X)] \right)^\top \mathbb{E} \left[(\varepsilon \xi^\top) \otimes (\xi \varepsilon^\top) | q(X) \right]^{-1} \operatorname{vec} \left(\mathbb{E}[\xi \xi^\top | q(X)] \right) \right] \\ &= \left(\mathbb{E} \left[\operatorname{vec}(\Omega_q)^\top \operatorname{vec}(\Sigma_q^{-1}) \right] \right)^{-2} \\ &\quad \cdot \underbrace{\mathbb{E} \left[Q^{-1} \operatorname{vec}(\Sigma_q^{-1})^\top \mathbb{E} \left[(R_\varepsilon R_\xi^\top) \otimes (R_\xi R_\varepsilon^\top) | q(X) \right] \operatorname{vec}(\Sigma_q^{-1}) \right]}_{=: A_1} \\ &\quad \cdot \underbrace{\mathbb{E} \left[Q \operatorname{vec}(\Omega_q)^\top \mathbb{E} \left[(R_\varepsilon R_\xi^\top) \otimes (R_\xi R_\varepsilon^\top) | q(X) \right]^{-1} \operatorname{vec}(\Omega_q) \right]}_{=: A_2}.\end{aligned}$$

Then,

$$\begin{aligned}
A_1 &= \mathbb{E} \left[Q^{-1} \text{vec}(\Sigma_q^{-1})^\top \mathbb{E} \left[(R_\varepsilon R_\xi^\top) \otimes (R_\xi R_\varepsilon^\top) \mid q(X) \right] \text{vec}(\Sigma_q^{-1}) \right] \\
&= \mathbb{E} \left[Q^{-1} \mathbb{E} \left[\text{vec}(\Sigma_q^{-1})^\top \{ (R_\varepsilon R_\xi^\top) \otimes (R_\xi R_\varepsilon^\top) \} \text{vec}(\Sigma_q^{-1}) \mid q(X) \right] \right] \\
&= \mathbb{E} \left[Q^{-1} \mathbb{E} \left[\text{vec}(\Sigma_q^{-1})^\top \text{vec} \left(\mathbb{E} \left[R_\xi R_\varepsilon^\top \Sigma_q^{-1} R_\xi R_\varepsilon^\top \mid q(X) \right] \right) \mid q(X) \right] \right] \\
&= \mathbb{E} \left[Q^{-1} \mathbb{E} \left[\text{tr} \left(\Sigma_q^{-1} R_\xi R_\xi \Sigma_q^{-1} R_\varepsilon R_\varepsilon^\top \right) \mid q(X) \right] \right] \\
&= \mathbb{E} \left[Q^{-1} \text{tr}(\Sigma_q^{-1} \Omega_q) \right],
\end{aligned}$$

and

$$\begin{aligned}
A_2 &= \mathbb{E} \left[Q \text{vec}(\Omega_q)^\top \mathbb{E} \left[(R_\varepsilon R_\xi^\top) \otimes (R_\xi R_\varepsilon^\top) \mid q(X) \right]^{-1} \text{vec}(\Omega_q) \right] \\
&= \mathbb{E} \left[Q \text{vec}(\Omega_q)^\top \text{vec}(\Sigma_q^{-1}) \right] \\
&= \mathbb{E} \left[Q \text{tr}(\Sigma_q^{-1} \Omega_q) \right],
\end{aligned}$$

where we use

$$\mathbb{E} \left[(R_\varepsilon R_\xi^\top) \otimes (R_\xi R_\varepsilon^\top) \mid q(X) \right]^{-1} \text{vec}(\Omega_q) = \text{vec}(\Sigma_q^{-1}),$$

which follows as

$$\begin{aligned}
\Omega_q &= \Omega_q \Sigma_q^{-1} \Sigma_q \\
&= \mathbb{E} \left[R_\xi R_\xi^\top \Sigma_q^{-1} R_\varepsilon R_\varepsilon^\top \mid q(X) \right] \\
&= \mathbb{E} \left[R_\xi R_\varepsilon^\top \Sigma_q^{-1} R_\xi R_\varepsilon^\top \mid q(X) \right] \\
&= \mathbb{E} \left[(R_\varepsilon R_\xi^\top) \otimes (R_\xi R_\varepsilon^\top) \mid q(X) \right] \\
\Rightarrow \text{vec}(\Omega_q) &= \mathbb{E} \left[(R_\varepsilon R_\xi^\top) \otimes (R_\xi R_\varepsilon^\top) \mid q(X) \right] \text{vec}(\Sigma_q^{-1}).
\end{aligned}$$

Define

$$t(X) := \text{tr}(\Sigma_q^{-1} \Omega_q).$$

Then

$$\begin{aligned}
\frac{V_{\text{ML/GEE}}}{V_{\text{SL}}} &= \frac{\mathbb{E} [Qt(X)] \mathbb{E} [Q^{-1}t(X)]}{(\mathbb{E} [t(X)])^2} \\
&= (\mathbb{E} [t(X)])^{-2} \\
&\quad \cdot \{ \eta^{-1} \mathbb{E} [t(X) \mid q(X) \in \Gamma] \mathbb{P}(q(X) \in \Gamma) + \mathbb{E} [t(X) \mid q(X) \in \Gamma^c] \mathbb{P}(q(X) \in \Gamma^c) \} \\
&\quad \cdot \{ \eta \mathbb{E} [t(X) \mid q(X) \in \Gamma] \mathbb{P}(q(X) \in \Gamma) + \mathbb{E} [t(X) \mid q(X) \in \Gamma^c] \mathbb{P}(q(X) \in \Gamma^c) \} \\
&\geq \eta \frac{\mathbb{E} [t(X) \mid q(X) \in \Gamma] \mathbb{E} [t(X) \mid q(X) \in \Gamma^c]}{(\mathbb{E} [t(X)])^2} \cdot \mathbb{P}(q(X) \in \Gamma) \mathbb{P}(q(X) \in \Gamma^c) \geq M,
\end{aligned}$$

taking

$$\eta = \left(M^{-1} \frac{\mathbb{E} [t(X) \mid q(X) \in \Gamma] \mathbb{E} [t(X) \mid q(X) \in \Gamma^c]}{(\mathbb{E} [t(X)])^2} \cdot \mathbb{P}(q(X) \in \Gamma) \mathbb{P}(q(X) \in \Gamma^c) \right)^{-1} \vee 2.$$

B.3 Regularity conditions for Theorem 3

For the statement of Theorem 3 we also impose the following regularity conditions:

- (i) The map $\psi \mapsto W(\psi)(x)$ is continuous for all $x \in \mathbb{R}^{n \times d}$.
- (ii) $\inf_{\psi \in \Psi} \inf_{x \in \mathbb{R}^{n \times d}} \Lambda_{\min}(W(\psi)(x)) > 0$.
- (iii) The errors satisfy the moment bounds $\mathbb{E} [\|\varepsilon\|_2^4] \leq C_\varepsilon$ and $\mathbb{E} [\|\xi\|_2^4] \leq C_\xi$ for finite constants $C_\varepsilon, C_\xi > 0$.

B.4 Proof of Theorem 3

Using Theorem 5.7 of [van der Vaart \[1998\]](#) it suffices to show

$$\sup_{\psi \in \Psi} \left| \hat{L}_{\text{SL}, I}(\psi) - L_{\text{SL}}(\psi) \right| = o_P(1). \quad (23)$$

Define the error terms

$$\mathcal{E}_i^{(\xi)} := \tilde{\xi}_i - \xi_i, \quad \mathcal{E}_i^{(\varepsilon)} := \tilde{\varepsilon}_i - \varepsilon_i,$$

and subsequently the terms

$$\begin{aligned} a_i^{(1)}(\psi) &:= \mathcal{E}_i^{(\xi)T} W_i(\psi) \mathcal{E}_i^{(\varepsilon)}, \\ a_i^{(2)}(\psi) &:= \mathcal{E}_i^{(\xi)T} W_i(\psi) \varepsilon_i + \mathcal{E}_i^{(\varepsilon)T} W_i(\psi) \xi_i \\ \mathcal{A}_I(\psi) &:= \frac{1}{N} \sum_{i=1}^I \left(\xi_i^\top W(\psi)(X_i) \varepsilon_i \right)^2 - \frac{1}{n} \mathbb{E} \left[\left(\xi^\top W(\psi)(X) \varepsilon \right)^2 \right], \\ \mathcal{A}_I(\psi) &:= \underbrace{\mathcal{A}_I(\psi) + \frac{2}{N} \sum_{i=1}^I \left(\xi_i^\top W_i(\psi) \varepsilon_i \right) \left(a_i^{(1)}(\psi) + a_i^{(2)}(\psi) \right) + \frac{1}{N} \sum_{i=1}^I \left(a_i^{(1)}(\psi) + a_i^{(2)}(\psi) \right)^2}_{=: G_I(\psi)}, \\ b_i(\psi) &:= 2\mathcal{E}_i^{(\xi)T} W_i(\psi) \xi_i + \mathcal{E}_i^{(\varepsilon)T} W_i(\psi) \mathcal{E}_i^{(\xi)}, \\ \mathcal{B}_I(\psi) &:= \frac{1}{N} \sum_{i=1}^I \xi_i^\top W(\psi)(X_i) \xi_i - \frac{1}{n} \mathbb{E} \left[\xi^\top W(\psi)(X) \xi \right], \\ B_I(\psi) &:= \mathcal{B}_I(\psi) + \frac{1}{N} \sum_{i=1}^I b_i(\psi). \end{aligned}$$

Then we can decompose

$$\begin{aligned} \sup_{\psi \in \Psi} \left| \hat{L}_{\text{SL}, I}(\psi) - L_{\text{SL}}(\psi) \right| &= \sup_{\psi \in \Psi} \left| \frac{N^{-1} \sum_{i=1}^I \left(\hat{\xi}_i^\top W(\psi)(X_i) \hat{\varepsilon}_i \right)^2}{\left(N^{-1} \sum_{i=1}^I \hat{\xi}_i^\top W(\psi)(X_i) \hat{\xi}_i \right)^2} - \frac{n^{-1} \mathbb{E} \left[\left(\xi^\top W(\psi)(X) \varepsilon \right)^2 \right]}{\left(n^{-1} \mathbb{E} \left[\xi^\top W(\psi)(X) \xi \right] \right)^2} \right| \\ &= \sup_{\psi \in \Psi} \left| \frac{n^{-1} \mathbb{E} \left[\left(\xi^\top W(\psi)(X) \varepsilon \right)^2 \right] + \mathcal{A}_I(\psi)}{\left(n^{-1} \mathbb{E} \left[\xi^\top W(\psi)(X) \xi \right] \right)^2 + B_I(\psi)} - \frac{n^{-1} \mathbb{E} \left[\left(\xi^\top W(\psi)(X) \varepsilon \right)^2 \right]}{\left(n^{-1} \mathbb{E} \left[\xi^\top W(\psi)(X) \xi \right] \right)^2} \right|. \end{aligned}$$

First we claim that $\sup_{\psi \in \Psi} |\mathcal{A}_I(\psi)| = o_P(1)$. This follows by the uniform law of large numbers [\[Jennrich, 1969\]](#), which may be applied as $\mathbb{E} \left[\sup_{\psi \in \Psi} \left(\xi^\top W(\psi)(X) \varepsilon \right)^2 \right] \leq C_\varepsilon^{\frac{1}{2}} C_\xi^{\frac{1}{2}} < \infty$ due to

Assumption (iii). It can similarly be shown that $\sup_{\psi \in \Psi} |\mathcal{B}_I(\psi)| = o_P(1)$. Now,

$$\begin{aligned}
& \mathbb{E} \left[\sup_{\psi \in \Psi} \left| \frac{1}{N} \sum_{i=1}^I b_i(\psi) \right| \right] \\
& \leq \frac{1}{N} \sum_{i=1}^I \left(2\mathbb{E} \left[\sup_{\psi \in \Psi} \left| \mathcal{E}_i^{(\xi)T} W_i(\psi) \xi_i \right| \right] + \mathbb{E} \left[\sup_{\psi \in \Psi} \mathcal{E}_i^{(\xi)T} W_i(\psi) \mathcal{E}_i^{(\xi)} \right] \right) \\
& \leq 2\mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\mathcal{E}_i^{(\xi)}\|_2 \|\xi_i\|_2 \right] + \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\mathcal{E}_i^{(\xi)}\|_2^2 \right] \\
& \leq 2\mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\mathcal{E}_i^{(\xi)}\|_2^2 \right]^{\frac{1}{2}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\xi_i\|_2^2 \right]^{\frac{1}{2}} + \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\mathcal{E}_i^{(\xi)}\|_2^2 \right] \\
& = o(1),
\end{aligned}$$

and so by Markov's inequality, $\sup_{\psi \in \Psi} \left| N^{-1} \sum_{i=1}^I b_i(\psi) \right| = o_P(1)$. Further,

$$\begin{aligned}
& \sup_{\psi \in \Psi} |G_I(\psi)| \\
& \leq \sup_{\psi \in \Psi} \frac{2}{N} \sum_{i=1}^I \left| \xi_i^\top W_i(\psi) \varepsilon_i \right| \cdot \left| a_i^{(1)}(\psi) \right| + \sup_{\psi \in \Psi} \frac{2}{N} \sum_{i=1}^I \left| \xi_i^\top W_i(\psi) \varepsilon_i \right| \cdot \left| a_i^{(2)}(\psi) \right| \\
& \quad + \sup_{\psi \in \Psi} \frac{2}{N} \sum_{i=1}^I \left(a_i^{(1)}(\psi) \right)^2 + \sup_{\psi \in \Psi} \frac{2}{N} \sum_{i=1}^I \left(a_i^{(2)}(\psi) \right)^2 \\
& \leq 2 \cdot \underbrace{\frac{1}{N} \sum_{i=1}^I \|\xi_i\| \|\varepsilon_i\| \|\mathcal{E}_i^{(\xi)}\| \|\mathcal{E}_i^{(\varepsilon)}\|}_{=: G_I^{(I)}} + 2 \cdot \underbrace{\frac{1}{N} \sum_{i=1}^I \left(\|\xi_i\| \|\varepsilon_i\| \left\{ \|\mathcal{E}_i^{(\xi)}\| \|\varepsilon_i\| + \|\mathcal{E}_i^{(\varepsilon)}\| \|\xi_i\| \right\} \right)}_{=: G_I^{(II)}} \\
& \quad + 2 \cdot \underbrace{\frac{1}{N} \sum_{i=1}^I \|\mathcal{E}_i^{(\xi)}\|^2 \|\mathcal{E}_i^{(\varepsilon)}\|^2}_{=: G_I^{(III)}} + 4 \cdot \underbrace{\frac{1}{N} \sum_{i=1}^I \left(\|\mathcal{E}_i^{(\xi)}\|^2 \|\varepsilon_i\|^2 + \|\mathcal{E}_i^{(\varepsilon)}\|^2 \|\xi_i\|^2 \right)}_{=: G_I^{(IV)}}
\end{aligned}$$

by the triangle equality. We proceed by showing that each term on the right hand side is $o_P(1)$. In the case of Assumption (a) we have

$$\begin{aligned}
\mathbb{E} \left[G_I^{(I)} \right] & \leq \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\xi_i\|_2^2 \|\varepsilon_i\|_2^2 \right]^{\frac{1}{2}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\mathcal{E}_i^{(\xi)}\|_2^2 \|\mathcal{E}_i^{(\varepsilon)}\|_2^2 \right]^{\frac{1}{2}} \\
& \leq C_\varepsilon^{\frac{1}{4}} C_\xi^{\frac{1}{4}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\mathcal{E}_i^{(\xi)}\|_2^4 \right]^{\frac{1}{4}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\mathcal{E}_i^{(\varepsilon)}\|_2^4 \right]^{\frac{1}{4}} = o(1),
\end{aligned}$$

by the Cauchy–Schwarz inequality (twice). We can similarly show (again by multiple uses of the Cauchy–Schwarz inequality) that $\mathbb{E}[G_I^{(II)}], \mathbb{E}[G_I^{(III)}], \mathbb{E}[G_I^{(IV)}] = o(1)$. It then follows by Markov's inequality (alongside Lemma 12) that $G_I^{(I)}, G_I^{(II)}, G_I^{(III)}, G_I^{(IV)} = o_P(1)$, and so $\sup_{\psi \in \Psi} |G_I(\psi)| = o_P(1)$. If alternatively Assumption (b) holds then we consider each term separately:

Term $G_I^{(I)}$:

$$\begin{aligned}
\mathbb{E} \left[(G_I^{(I)})^{\frac{2}{3}} \right] &\leq 2^{\frac{2}{3}} N^{\frac{1}{3}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\xi_i\|_2^{\frac{2}{3}} \|\varepsilon_i\|_2^{\frac{2}{3}} \|\mathcal{E}_i^{(\xi)}\|_2^{\frac{2}{3}} \|\mathcal{E}_i^{(\varepsilon)}\|_2^{\frac{2}{3}} \right] \\
&\leq N^{\frac{1}{3}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\varepsilon_i\|_2^2 \|\xi_i\|_2^2 \right]^{\frac{1}{3}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\mathcal{E}_i^{(\varepsilon)}\|_2 \|\mathcal{E}_i^{(\xi)}\|_2 \right]^{\frac{2}{3}} \\
&\leq N^{\frac{1}{3}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\mathcal{E}_i^{(\varepsilon)}\|_2^2 \right]^{\frac{1}{3}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\mathcal{E}_i^{(\xi)}\|_2^2 \right]^{\frac{1}{3}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\varepsilon_i\|_2^4 \right]^{\frac{1}{6}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\xi_i\|_2^4 \right]^{\frac{1}{6}} \\
&= o(1),
\end{aligned}$$

the first inequality follows as $(\sum_i x_i)^\eta \leq \sum_i x_i^\eta$ for any $\eta \in (0, 1]$, $x_i \geq 0$, the second follows by Hölder's inequality, and the third by the Cauchy–Schwarz inequality.

Term $G_I^{(II)}$:

$$\begin{aligned}
\mathbb{E} \left[(G_I^{(II)})^{\frac{4}{5}} \right] &\leq N^{\frac{1}{5}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\xi_i\|_2^{\frac{4}{5}} \|\varepsilon_i\|_2^{\frac{8}{5}} \|\mathcal{E}_i^{(\xi)}\|_2^{\frac{4}{5}} \right] + N^{\frac{1}{5}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\xi_i\|_2^{\frac{8}{5}} \|\varepsilon_i\|_2^{\frac{4}{5}} \|\mathcal{E}_i^{(\varepsilon)}\|_2^{\frac{4}{5}} \right] \\
&\leq N^{\frac{1}{5}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\xi_i\|_2^{\frac{4}{3}} \|\varepsilon_i\|_2^{\frac{8}{3}} \right]^{\frac{3}{5}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\mathcal{E}_i^{(\xi)}\|_2^2 \right]^{\frac{2}{5}} \\
&\quad + N^{\frac{1}{5}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\xi_i\|_2^{\frac{8}{3}} \|\varepsilon_i\|_2^{\frac{4}{3}} \right]^{\frac{3}{5}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\mathcal{E}_i^{(\varepsilon)}\|_2^2 \right]^{\frac{2}{5}} \\
&\leq N^{\frac{1}{5}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\mathcal{E}_i^{(\xi)}\|_2^2 \right]^{\frac{2}{5}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\mathcal{E}_i^{(\xi)}\|_2^2 \right]^{\frac{1}{3}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\xi_i\|_2^4 \right]^{\frac{1}{5}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\varepsilon_i\|_2^4 \right]^{\frac{2}{5}} \\
&\quad + N^{\frac{1}{5}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\mathcal{E}_i^{(\varepsilon)}\|_2^2 \right]^{\frac{2}{5}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\mathcal{E}_i^{(\xi)}\|_2^2 \right]^{\frac{1}{3}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\xi_i\|_2^4 \right]^{\frac{2}{5}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\varepsilon_i\|_2^4 \right]^{\frac{1}{5}} \\
&= o(1),
\end{aligned}$$

where the first inequality follows as $(\sum_i x_i)^\eta \leq \sum_i x_i^\eta$ for any $\eta \in (0, 1]$, $x_i \geq 0$, the second follows by the Cauchy–Schwarz inequality, and the third by Hölder's inequality.

Term $G_I^{(III)}$:

$$\mathbb{E} \left[(G_I^{(III)})^{\frac{1}{2}} \right] \leq N^{\frac{1}{2}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\mathcal{E}_i^{(\xi)}\|_2^2 \right]^{\frac{1}{2}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\mathcal{E}_i^{(\varepsilon)}\|_2^2 \right]^{\frac{1}{2}} = o(1),$$

by $(\sum_i x_i)^\eta \leq \sum_i x_i^\eta$ for any $\eta \in (0, 1]$, $x_i \geq 0$ and using the Cauchy–Schwarz inequality.

Term $G_I^{(IV)}$:

$$\begin{aligned} \mathbb{E} \left[(G_I^{(IV)})^{\frac{2}{3}} \right] &\leq N^{\frac{1}{3}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\mathcal{E}_i^{(\xi)}\|_2^{\frac{4}{3}} \|\varepsilon_i\|_2^{\frac{4}{3}} \right] + N^{\frac{1}{3}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\mathcal{E}_i^{(\varepsilon)}\|_2^{\frac{4}{3}} \|\xi_i\|_2^{\frac{4}{3}} \right] \\ &\leq N^{\frac{1}{3}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\mathcal{E}_i^{(\xi)}\|_2^2 \right]^{\frac{2}{3}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\varepsilon_i\|_2^4 \right]^{\frac{1}{3}} \\ &\quad + N^{\frac{1}{3}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\mathcal{E}_i^{(\varepsilon)}\|_2^2 \right]^{\frac{2}{3}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^I \|\xi_i\|_2^4 \right]^{\frac{1}{3}} \end{aligned}$$

By applying Markov's inequality to each of the above terms we have $G_I^I, G_I^{II}, G_I^{III}, G_I^{IV} = o_P(1)$, and therefore with Lemma 12 we have $\sup_{\psi \in \Psi} |G_I(\psi)| = o_P(1)$.

Therefore, combining the results above with the triangle inequality

$$\sup_{\psi \in \Psi} |A_I(\psi)| = o_P(1) \quad \text{and} \quad \sup_{\psi \in \Psi} |B_I(\psi)| = o_P(1). \quad (24)$$

Now, by the triangle inequality

$$\begin{aligned} \sup_{\psi \in \Psi} \left| \hat{L}_{SL,I}(\psi) - L_{SL}(\psi) \right| &\leq \left(\sup_{\psi \in \Psi} \left| \frac{1}{(\mathbb{E} [\xi^\top W(\psi)(X)\xi] + B_I(\psi))^2} - \frac{1}{(\mathbb{E} [\xi^\top W(\psi)(X)\xi])^2} \right| \right) \\ &\quad \cdot \left(\sup_{\psi \in \Psi} \mathbb{E} \left[(\xi^\top W(\psi)(X)\varepsilon)^2 \right] + \sup_{\psi \in \Psi} |A_I(\psi)| \right) + \frac{\sup_{\psi \in \Psi} |A_I(\psi)|}{\inf_{\psi \in \Psi} (\mathbb{E} [\xi^\top W(\psi)(X)\xi])^2}. \end{aligned}$$

We have already shown that $\mathbb{E} \left[\sup_{\psi \in \Psi} (\xi^\top W(\psi)(X)\varepsilon)^2 \right] < \infty$ and it is similarly straightforward to show that $\mathbb{E} \left[\inf_{\psi \in \Psi} \xi^\top W(\psi)(X)\xi \right] > 0$. Therefore, it remains to show that

$$\sup_{\psi \in \Psi} \left| \frac{1}{(\mathbb{E} [\xi^\top W(\psi)(X)\xi] + B_I(\psi))^2} - \frac{1}{(\mathbb{E} [\xi^\top W(\psi)(X)\xi])^2} \right| = o_P(1). \quad (25)$$

To show this, it will be helpful to define

$$\begin{aligned} c &:= \mathbb{E} \left[\inf_{\psi \in \Psi} \xi^\top W(\psi)(X)\xi \right] > 0, \\ C &:= \mathbb{E} \left[\sup_{\psi \in \Psi} \xi^\top W(\psi)(X)\xi \right] < \infty \end{aligned}$$

Then, further defining $B_I^{\text{sup}} := \sup_{\psi \in \Psi} |B_I(\psi)|$ it follows that

$$\begin{aligned} &\mathbb{P}_P \left(\sup_{\psi \in \Psi} \left| \frac{1}{(\mathbb{E} [\xi^\top W(\psi)(X)\xi] + B_I(\psi))^2} - \frac{1}{(\mathbb{E} [\xi^\top W(\psi)(X)\xi])^2} \right| > \epsilon \right) \\ &\leq \mathbb{P}_P \left(\frac{2cB_I^{\text{sup}} - (B_I^{\text{sup}})^2}{C^4 (1 + c^{-1}B_I^{\text{sup}})^2} > \epsilon \right) \\ &\leq \mathbb{P}_P \left(\frac{2cB_I^{\text{sup}} - (B_I^{\text{sup}})^2}{C^4 (1 + c^{-1}B_I^{\text{sup}})^2} > \tilde{\epsilon} \right) \quad \text{where } \tilde{\epsilon} := \min \left(\epsilon, \frac{1}{6}C^{-4}c^2 \right) \\ &\leq \mathbb{P}_P (B_I^{\text{sup}} > \lambda) \quad \text{where } \lambda := (1 + c^{-2}C^4\tilde{\epsilon})^{-1} \left(c - c^{-1}C^4\tilde{\epsilon} - \sqrt{c^2 - 3C^4\tilde{\epsilon}} \right) > 0 \\ &= \mathbb{P}_P \left(\sup_{\psi \in \Psi} |B_I(\psi)| > \lambda \right) = o(1) \quad \text{by (24)}. \end{aligned}$$

The penultimate step follows by solving the quadratic inequality in B_I^{sup} . This proves the claim (25) and hence (23), completing the proof.

C Further details on sandwich boosting

C.1 Working correlation structures that exhibit computational benefits

We outline three common working correlation structures that allow for computationally fast sandwich boosting, at computational order $O(N)$ as opposed to $O(\sum_i n_i^3)$.

Equicorrelated: The equicorrelated working correlation structure (and its scaled inverse) is given by

$$C_\theta(X_i) = (\mathbb{1}_{\{j=k\}} + \mathbb{1}_{\{j \neq k\}} \rho_\theta)_{(j,k) \in [n_i]^2},$$

$$C_\theta^{-1}(X_i) = \left(\mathbb{1}_{\{j=k\}} - \frac{\theta}{1 + \theta n_i} \right)_{(j,k) \in [n_i]^2},$$

where the constant correlation $\rho_\theta = \theta/(1 + \theta)$ is given in terms of the parameter $\theta \in [0, \infty)$. This working correlation is particularly popular for repeated measures data, with this class of weights in the homoscedastic case (setting $\sigma^* \equiv s^* \equiv 1$) the same class as used when fitting an intercept only mixed effects model.

Autoregressive, AR(1): The autoregressive structure is given by the working correlation matrix (and its scaled inverse) of an AR(1) process

$$C_\theta(X_i) = (\theta^{|j-k|})_{(j,k) \in [n_i]^2},$$

$$C_\theta^{-1}(X_i) = \begin{cases} (\mathbb{1}_{\{j=k\}} + \theta^2 \mathbb{1}_{\{1 < j=k < n_i\}} - \theta \mathbb{1}_{\{|j-k|=1\}})_{(j,k) \in [n_i]^2} & \text{if } n_i > 1 \\ (1 - \theta^2) & \text{if } n_i = 1 \end{cases}.$$

This working correlation is therefore of particular relevance for longitudinal datasets where the correlation between observations is bounded by a function that decays exponentially with separation.

Hierarchical / Nested: The nested working correlation models hierarchical groupings as follows. Consider the i th group (of size n_i) to be further partitioned into M sub-groups of sizes (n_{i1}, \dots, n_{iM}) (such that $\sum_{m=1}^M n_{im} = n_i$). The nested group structure is then given by the block matrix

$$C_\theta(X_i) = \begin{pmatrix} E_{i,11} & \cdots & E_{i,1M} \\ \vdots & \ddots & \vdots \\ E_{i,M1} & \cdots & E_{i,MM} \end{pmatrix}, \quad E_{i,mm'} = \begin{cases} \rho_1 J_{n_{im} \times n_{im}} + (1 - \rho_1) I_{n_{im}} & \text{if } m = m' \\ \rho_2 J_{n_{im} \times n_{im'}} & \text{if } m \neq m', \end{cases}$$

for the vector $\rho = (\rho_1, \rho_2) \in [0, 1) \times [0, 1)$ satisfying $\rho_2 \leq \rho_1$ (suppressing the dependence on θ in ρ_θ), where $J_{a \times b}$ denotes an $a \times b$ dimensional matrix of ones and I_a denotes the $a \times a$ dimensional identity matrix. This nested correlation can be reparametrised in terms of $\theta =$

$(\theta_1, \theta_2) := \left(\frac{\rho_1 - \rho_2}{1 - \rho_1}, \frac{\rho_2}{1 - \rho_1} \right) \in [0, \infty) \times [0, \infty)$ as

$$C_\theta^{-1}(X_i) = (1 + \theta_1 + \theta_2)I_{n_i} - \mathcal{B}_i(\theta),$$

$$\mathcal{B}_i(\theta) = \begin{pmatrix} \mathcal{A}_{i,11}(\theta)J_{n_{i1} \times n_{i1}} & \mathcal{A}_{i,12}(\theta)J_{n_{i1} \times n_{i2}} & \cdots & \mathcal{A}_{i,1M}(\theta)J_{n_{i1} \times n_{iM}} \\ \mathcal{A}_{i,21}(\theta)J_{n_{i2} \times n_{i1}} & \mathcal{A}_{i,22}(\theta)J_{n_{i2} \times n_{i2}} & \cdots & \mathcal{A}_{i,2M}(\theta)J_{n_{i2} \times n_{iM}} \\ \vdots & \vdots & \ddots & \vdots \\ \mathcal{A}_{i,M1}(\theta)J_{n_{iM} \times n_{i1}} & \mathcal{A}_{i,M2}(\theta)J_{n_{iM} \times n_{i2}} & \cdots & \mathcal{A}_{i,MM}(\theta)J_{n_{iM} \times n_{iM}} \end{pmatrix},$$

$$\mathcal{A}_{i,mm'}(\theta) = \frac{\theta_1}{1 + \theta_1 n_{im}} \mathbb{1}_{\{m=m'\}} + \frac{\theta_2}{1 + \theta_2 \sum_{l=1}^M \frac{n_{il}}{1 + \theta_1 n_{il}}} \cdot \frac{1}{(1 + \theta_1 n_{im})(1 + \theta_1 n_{im'})} \quad (m, m' \in [M]),$$

calculated by the Woodbury matrix identity.

C.2 Calculation of s-scores for computationally efficient sandwich boosting

Note that, as discussed in Section 2, the scores for sandwich boosting can be calculated in $O(\sum_i n_i^3)$ computations for a generic (inverse) working correlation $C_\theta^{-1}(\cdot)$, but the three working correlations introduced are special cases that can be calculated in $O(N)$ operations. We first introduce a general algorithm to calculate scores for a generic working correlation parametrised by θ , and then present this theory applied to the special cases of the working correlations introduced (where we have a reduction in computational order).

Therefore, the order of computation to calculate the score functions for sandwich boosting is $O(\sum_i n_i^3)$ for a general correlation matrix structure whose inverse $C_\theta^{-1}(\cdot)$ is not known (and so needs to be computed analytically). In the special case of Toeplitz correlation structures, or where the inverse correlation structure $C_\theta^{-1}(\cdot)$ is known, this order of computation can be reduced to $O(\sum_i n_i^2)$. Note that this computational order can be further reduced to $O(N)$ in the three special cases of correlation structures posed in this paper. We give an example of these, implementing Algorithm 3 to the equicorrelation structure, with scores calculated in $O(N)$ operations as in Algorithm 4. For notational simplicity for these algorithms we introduce the notation

$$\mathcal{C}_{ijk}(\theta) := \{C_\theta^{-1}(X_i)\}_{jk}, \quad (26)$$

for the inverse working correlation structure.

Similar algorithms exist for the autoregressive AR(1) and nested groups working correlations (and are also working correlation cases that achieve the faster $O(N)$ computational rates when calculating scores for performing sandwich boosting).

C.3 Theoretical results under working correlation structures

The uniform asymptotic convergence results of Section 4 works on the general working covariance model parametrised by (σ, ρ) functions and / or parameters. For computational speed and stability for sandwich boosting we consider the weight class reparametrised in terms of (s, θ) . For completeness we reformulate the assumptions on the deterministic limits (σ^*, ρ^*) of Assumption A2 in terms of the deterministic limits (s^*, θ^*) , with a particular focus on the three prototypical working correlation parametrisations considered in Section C.

Corollary 6. *Consider the setting of Theorem 4 where the weight class \mathcal{W} takes one of the three forms outlined in Section C in terms of (s, θ) . Further, suppose there exists the deterministic function $s^* : \mathbb{R}^d \rightarrow \mathbb{R}$ and vector θ^* whose estimators in place of Assumptions A3.1-A3.3 in Assumption A3 satisfy*

Algorithm 3: Computationally lean sandwich boosting score calculation for an arbitrary working correlation

Input: Index set \mathcal{I} for boosting training; Working correlation structure \mathcal{C}_i (26); Estimates of (grouped) errors $(\hat{\xi}_i, \hat{\varepsilon}_i, X_i)_{i \in \mathcal{I}}$; Estimates of (s, θ) at which scores to be calculated.

Calculate $A^{(1)}, A^{(2)}$:

for $i \in \mathcal{I}$ **do**

$$\begin{cases} A_i^{(1)} := \hat{\xi}_i^\top W_i \hat{\xi}_i = \sum_{k=1}^{n_i} \sum_{k'=1}^{n_i} \mathcal{C}_{ikk'}(\theta) s(X_{ik}) s(X_{ik'}) \hat{\xi}_{ik} \hat{\xi}_{ik'}, \\ A_i^{(2)} := \hat{\xi}_i^\top W_i \hat{\varepsilon}_i = \sum_{k=1}^{n_i} \sum_{k'=1}^{n_i} \mathcal{C}_{ikk'}(\theta) s(X_{ik}) s(X_{ik'}) \mathcal{C}_{ikk'}(\theta) \hat{\xi}_{ik} \hat{\varepsilon}_{ik'}. \end{cases}$$

end

Calculate $A^{(3)}, A^{(4)}$:

for $i \in \mathcal{I}$ **and** $j \in [n_i]$ **do**

$$\begin{cases} A_{ij}^{(3)} := 2 \sum_{k=1}^{n_i} \mathcal{C}_{ijk}(\theta) s(X_{ik}) \hat{\xi}_{ij} \hat{\xi}_{ik}, \\ A_{ij}^{(4)} := \sum_{k=1}^{n_i} \mathcal{C}_{ijk}(\theta) s(X_{ik}) \left(\hat{\xi}_{ij} \hat{\varepsilon}_{ik} + \hat{\xi}_{ik} \hat{\varepsilon}_{ij} \right). \end{cases}$$

end

Calculate $A_{\text{sum}}^{(1)} := \sum_{i \in \mathcal{I}} A_i^{(1)}$ and $A_{\text{sum.sq}}^{(2)} := \sum_{i \in \mathcal{I}} A_i^{(2)^2}$.

Calculate s -scores:

for $i \in \mathcal{I}$ **and** $j \in [n_i]$ **do**

$$U_{\text{SL}}^{(s)}(s, \theta)(X_{ij}) = -2 \left(A_{\text{sum}}^{(1)} \right)^{-3} \left[A_{\text{sum.sq}}^{(2)} A_{ij}^{(3)} - A_{\text{sum}}^{(1)} A_i^{(2)} A_{ij}^{(4)} \right].$$

end

Calculate the θ -score:

$$U_{\text{SL}}^{(\theta)}(s, \theta) = -2 \left(A_{\text{sum}}^{(1)} \right)^{-3} \left[A_{\text{sum.sq}}^{(2)} \left(\sum_{i \in \mathcal{I}} \sum_{k=1}^{n_i} \sum_{k'=1}^{n_i} \frac{\partial \mathcal{C}_{ikk'}(\theta)}{\partial \theta} s(X_{ik}) s(X_{ik'}) \hat{\xi}_{ik} \hat{\xi}_{ik'} \right) - A_{\text{sum}}^{(1)} \left(\sum_{i \in \mathcal{I}} A_i^{(2)} \sum_{k=1}^{n_i} \sum_{k'=1}^{n_i} \left(\frac{\partial \mathcal{C}_{ikk'}(\theta)}{\partial \theta} s(X_{ik}) s(X_{ik'}) \hat{\xi}_{ik} \hat{\varepsilon}_{ik'} \right) \right) \right].$$

Output: The set of N s -scores $\left(X_{ij}, U_{\text{SL}}^{(s)}(s, \theta)(X_{ij}) \right)_{i \in \mathcal{I}, j \in [n_i]}$ and θ -score $U_{\text{SL}}^{(\theta)}(s, \theta)$.

$$(A3.1') \quad \mathcal{R}_{s^*} := \max_{j \in [n]} \mathbb{E}_P \left[\left(c^* \hat{s}^{(1)}(X_j) - s^*(X_j) \right)^2 \mid \hat{s}^{(1)} \right] = o_{\mathcal{P}}(1)$$

$$\text{where } c^* := \operatorname{arg\,inf}_{c > 0} \max_{j \in [n]} \mathbb{E}_P \left[\left(c \hat{s}^{(1)}(X_j) - s^*(X_j) \right)^2 \right],$$

$$(A3.2') \quad \mathcal{R}_{\theta^*} := \max_{k \in [K]} \left\| \hat{\theta}^{(k)} - \theta^* \right\|_2 = o_{\mathcal{P}}(1),$$

(A3.3') For all $k \in [K]$, $\hat{s}^{(k)}$ are bounded above and below uniformly by some positive constants and, for the relevant working correlation:

- **Equicorrelation:** $\hat{\theta}^{(k)}$ is bounded above and below uniformly by constants in $[0, \infty)$,
- **Longitudinal:** $\hat{\theta}^{(k)}$ is bounded above and below uniformly by constants in $[0, 1)$,
- **Hierarchical / Nested:** $\hat{\theta}^{(k)}$ is component-wise bounded above and below uniformly by constants in $[0, \infty)$.

Then the results of Theorem 4 hold in this alternative setup.

A proof of this corollary is given in Section D.4.

Algorithm 4: Computationally lean sandwich boosting score calculation for the equicorrelated working correlation

Input: Index set \mathcal{I} for boosting training; Estimates of (grouped) errors $(\hat{\xi}_i, \hat{\varepsilon}_i, X_i)_{i \in \mathcal{I}}$;
 Estimates of (s, θ) at which scores to be calculated.

Calculate $A^{(1)}, A^{(2)}, B^{(1)}, B^{(2)}$:

for $i \in \mathcal{I}$ **do**

$$B_i^{(1)} := \sum_{k=1}^{n_i} s(X_{ik}) \hat{\xi}_{ik},$$

$$B_i^{(2)} := \sum_{k=1}^{n_i} s(X_{ik}) \hat{\varepsilon}_{ik},$$

$$A_i^{(1)} := \hat{\xi}_i^\top W_i \hat{\xi}_i = \sum_{k=1}^{n_i} s^2(X_{ik}) \hat{\xi}_{ik}^2 - \frac{\theta}{1+\theta n_i} \left(\sum_{k=1}^{n_i} s(X_{ik}) \hat{\xi}_{ik} \right)^2,$$

$$A_i^{(2)} := \hat{\xi}_i^\top W_i \hat{\varepsilon}_i = \sum_{k=1}^{n_i} s^2(X_{ik}) \hat{\xi}_{ik} \hat{\varepsilon}_{ik} - \frac{\theta}{1+\theta n_i} \left(\sum_{k=1}^{n_i} s(X_{ik}) \hat{\xi}_{ik} \right) \left(\sum_{k=1}^{n_i} s(X_{ik}) \hat{\varepsilon}_{ik} \right).$$

end

Calculate $A^{(3)}$ and $A^{(4)}$:

for $i \in \mathcal{I}$ **and** $j \in [n_i]$ **do**

$$A_{ij}^{(3)} := 2s(X_{ij}) \hat{\xi}_{ij}^2 - \frac{2\theta}{1+\theta n_i} B_i^{(1)} \hat{\xi}_{ij},$$

$$A_{ij}^{(4)} := 2s(X_{ij}) \hat{\xi}_{ij} \hat{\varepsilon}_{ij} - \frac{\theta}{1+\theta n_i} \left(B_i^{(2)} \hat{\xi}_{ij} + B_i^{(1)} \hat{\varepsilon}_{ij} \right).$$

end

Calculate $A_{\text{sum}}^{(1)} := \sum_{i=1}^I A_i^{(1)}$ and $A_{\text{sum.sq}}^{(2)} := \sum_{i=1}^I A_i^{(2)^2}$.

Calculate s -scores:

for $i \in \mathcal{I}$ **and** $j \in [n_i]$ **do**

$$U_{\text{SL}}^{(s)}(s, \theta)(X_{ij}) = -2 \left(A_{\text{sum}}^{(1)} \right)^{-3} \left[A_{\text{sum.sq}}^{(2)} A_{ij}^{(3)} - A_{\text{sum}}^{(1)} A_i^{(2)} A_{ij}^{(4)} \right].$$

end

Calculate the θ -score:

$$U_{\text{SL}}^{(\theta)}(s, \theta) = -2 \left(A_{\text{sum}}^{(1)} \right)^{-3} \left[A_{\text{sum.sq}}^{(2)} \left(\sum_{i=1}^I \frac{(B_i^{(1)})^2}{(1+\theta n_i)^2} \right) - A_{\text{sum}}^{(1)} \left(\sum_{i=1}^I \frac{A_i^{(2)} B_i^{(1)} B_i^{(2)}}{(1+\theta n_i)^2} \right) \right].$$

Output: The set of N s -scores $(X_{ij}, U_{\text{SL}}^{(s)}(s, \theta)(X_{ij}))_{i \in \mathcal{I}, j \in [n_i]}$ and θ -score $U_{\text{SL}}^{(\theta)}(s, \theta)$.

D Proof of results in Section 4

In this section we prove a stronger version of Theorem 4, extending it to non-identically grouped data of varying group size, and allowing group sizes to diverge at sufficiently slow rates with N . To allow for non-identical groups we work on a sequence of families of probability distributions $(\mathcal{P}_I)_{I \in \mathbb{N}}$ that the grouped observations follow (satisfying Assumption A1 with us suppressing dependence on i), and re-define the nuisance function error rates given in Assumptions A2 and A3 with

$$\mathcal{R}_l := \max_{k \in [K]} \max_{i \in \mathcal{I}_k} \max_{j \in [n_i]} \mathbb{E}_P \left[\left(\hat{l}^{(k)}(X_{ij}) - l_0(X_{ij}) \right)^2 \mid \hat{l}^{(k)} \right],$$

and similarly for $\mathcal{R}_m, \mathcal{R}_\sigma$ and \mathcal{R}_ρ . This notation then recovers that in Assumption A2 when the groups are identically distributed. We also introduce an additional parameter τ of the working correlation structure and re-define

$$V := \left(\frac{1}{N} \sum_{i=1}^I \mathbb{E}_P \left[\xi_i^\top W_i^* \xi_i \right] \right)^{-2} \left(\frac{1}{N} \sum_{i=1}^I \mathbb{E}_P \left[\left(\varepsilon_i^\top W_i^* \xi_i \right)^2 \right] \right),$$

with $W_i^* := W(\sigma^*, \rho^*)(X_i)$.

Assumption A3 Extended (Additional assumption on weight function). Additionally to the Assumptions (A3.1)-(A3.4) of Assumption A3 suppose we also make the assumption

(A3.5) The working correlation function ρ^* satisfies $\max_{k \in [K]} \max_{i \in \mathcal{I}_k} \max_{j \in [n_i]} \sum_{j' \neq j}^{n_i} \rho^*(X_{ij}, X_{ij'}) \lesssim n_i^\tau$ for some $\tau \in [0, 1]$ almost surely.

We can now specify the asymptotic regimes on the maximum group size $n_{\max} := \max_{i \in [I]} n_i$ for which Theorem 4 holds, in terms of the nuisance error rates $(\mathcal{R}_m, \mathcal{R}_l, \mathcal{R}_\sigma, \mathcal{R}_\rho)$, the working correlation metadata $(\alpha, \kappa, \gamma, \tau)$, and the true data generation mechanism metadata $(\delta, \mu_\Sigma, \mu_\Omega)$.

Assumption A4 (Maximum Group Size). The maximum group size $n_{\max} := \max_{i \in [I]} n_i$ of the grouped observations $\{S_i = (Y_i, D_i, X_i)\}_{i \in [I]}$ satisfies

$$n_{\max} = o\left(N^{\frac{\delta}{4(1+\kappa)+(2+\kappa)\delta} \wedge \frac{1}{3+2\kappa}} \wedge \mathcal{R}_m^{-\left(\frac{1}{\kappa+\mu_\Omega} \wedge \frac{1}{2(\kappa+\alpha)} \wedge \frac{1}{\gamma}\right)} \wedge (N\mathcal{R}_m^{-1})^{\frac{1}{2\gamma+\mu_\Omega}} \wedge \mathcal{R}_l^{-\frac{1}{\kappa+\mu_\Omega}} \wedge \mathcal{R}_\sigma^{-\left(\frac{1}{2\kappa+2\mu_\Sigma+\tau+1} \wedge \frac{1}{2\gamma+\tau+1}\right)} \wedge \mathcal{R}_\rho^{-\left(\frac{1}{2(\kappa+\mu_\Sigma+1)} \wedge \frac{1}{2(\gamma+1)}\right)} \wedge \left(N^{-1}\mathcal{R}_m^{-1}(\mathcal{R}_l \vee \mathcal{R}_m)^{-1}\right)^{\frac{1}{\kappa}} \right).$$

This assumption reduces to the rates given in Section 4 by setting $(\mu_\Sigma, \mu_\Omega, \gamma, \tau) = (1, 1, 1, 1)$ in the equicorrelated working correlation setting, and $(\mu_\Sigma, \mu_\Omega, \alpha, \kappa, \gamma, \tau) = (0, 1, 0, 0, 0, 0)$ in the autoregressive AR(1) working correlation setting.

We introduce the following additional notation relevant for the proofs in this section. For a distribution P governing the distribution of a random vector $A \in \mathbb{R}^d$, we write $\|A\|_{P,p} := (\mathbb{E}_P \|A\|_2^p)^{1/p}$, the subscript P in the expectation indicating the dependence on P ; we will similarly write \mathbb{P}_P . We extend this definition to random matrices A where then $\|A\|_{P,p} := (\mathbb{E}_P \|A\|_{\text{op}}^p)^{1/p}$, where $\|\cdot\|_{\text{op}}$ is the operator norm.

D.1 Proof of Theorem 4

The proof of Theorem 4 amounts to proving the two claims:

- (i) $\lim_{I \rightarrow \infty} \sup_{P \in \mathcal{P}_I} \sup_{t \in \mathbb{R}} \left| \mathbb{P}_P \left(\sqrt{N/V} (\hat{\beta} - \beta) \leq t \right) - \Phi(t) \right| = 0;$
- (ii) $\lim_{I \rightarrow \infty} \sup_{P \in \mathcal{P}_I} \sup_{t \in \mathbb{R}} \left| \mathbb{P}_P \left(\sqrt{N/\hat{V}} (\hat{\beta} - \beta) \leq t \right) - \Phi(t) \right| = 0.$

We will prove each in turn.

Proof of Theorem 4 (i). Consider separating the nuisance functions into ‘fixed effects’ functions $\eta := (l, m)$ and weight functions $\nu := (\sigma, \rho)$ (analogous to the working covariance in the GEE setup or random effects in the ML setup). Define the terms

$$\begin{aligned} \varphi(S_i; \tilde{\beta}, \eta, \nu) &:= (D_i - m(X_i))^\top W_i(\nu) (Y_i - l(X_i) - (D_i - m(X_i)) \tilde{\beta}), \\ \phi(S_i; \eta, \nu) &:= (D_i - m(X_i))^\top W_i(\nu) (D_i - m(X_i)), \end{aligned}$$

where $W_i(\nu) := W(\nu)(X_i)$. Further, recall the notation that we denote the k^{th} fold estimators for the ‘fixed effects’ $\eta_0 := (l_0, m_0)$ as $\hat{\eta}^{(k)} := (\hat{l}^{(k)}, \hat{m}^{(k)})$, and for the weights the estimators $\hat{\nu}^{(k)} := (\hat{\sigma}^{(k)}, \hat{\rho}^{(k)})$ that converge to a deterministic limit $\nu^* := (\sigma^*, \rho^*)$, each with corresponding weights $\hat{W}_i = W(\hat{\nu}^{(k)})(X_i)$ (with k -dependence implicit through i) and $W_i^* = W(\nu^*)(X_i)$. By scaling we may assume without loss of generality that in Assumption A3.1 we have $c^* = 1$.

Then $\hat{\beta}$ in Algorithm 1 can be expanded as

$$\begin{aligned}
N^{\frac{1}{2}}(\hat{\beta} - \beta) &= \left(N^{-1} \sum_{k=1}^K \sum_{i \in \mathcal{I}_k} \phi(S_i; \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) \right)^{-1} \left(N^{-\frac{1}{2}} \sum_{k=1}^K \sum_{i \in \mathcal{I}_k} \varphi(S_i; \beta, \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) \right) \\
&= \left(\frac{N^{-1} n_{\max}^{\kappa} \sum_{i=1}^I \mathbb{E}_P [\varphi^2(S_i; \beta, \eta_0, \nu^*)]}{\left(N^{-1} \sum_{i=1}^I \mathbb{E}_P [\phi(S_i; \eta_0, \nu^*)] \right)^2} \right)^{\frac{1}{2}} \left(\frac{N^{-1} n_{\max}^{\gamma} \sum_{i=1}^I \phi(S_i; \eta_0, \nu^*) + M_2}{N^{-1} n_{\max}^{\gamma} \sum_{i=1}^I \mathbb{E}_P [\phi(S_i; \eta_0, \nu^*)]} \right)^{-1} \\
&\quad \cdot \left(\frac{N^{-\frac{1}{2}} n_{\max}^{\frac{\kappa}{2}} \sum_{i=1}^I \varphi(S_i; \beta, \eta_0, \nu^*) + M_1}{\left(N^{-1} n_{\max}^{\kappa} \sum_{i=1}^I \mathbb{E}_P [\varphi^2(S_i; \beta, \eta_0, \nu^*)] \right)^{\frac{1}{2}}} \right), \tag{27}
\end{aligned}$$

where

$$\begin{aligned}
M_1 &:= N^{-\frac{1}{2}} n_{\max}^{\frac{\kappa}{2}} \sum_{k=1}^K \sum_{i \in \mathcal{I}_k} \left[\varphi(S_i; \beta, \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) - \varphi(S_i; \beta, \eta_0, \nu^*) \right], \\
M_2 &:= N^{-1} n_{\max}^{\gamma} \sum_{k=1}^K \sum_{i \in \mathcal{I}_k} \left[\phi(S_i; \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) - \phi(S_i; \eta_0, \nu^*) \right].
\end{aligned}$$

We claim that by Lemma 17 we have

$$\lim_{I \rightarrow \infty} \sup_{P \in \mathcal{P}_I} \sup_{t \in \mathbb{R}} \left| \mathbb{P}_P \left(\left(\sum_{i=1}^I \mathbb{E}_P [\varphi^2(S_i; \beta, \eta_0, \nu^*)] \right)^{-\frac{1}{2}} \left(\sum_{i=1}^I \varphi(S_i; \beta, \eta_0, \nu^*) \right) \leq t \right) - \Phi(t) \right| = 0. \tag{28}$$

We must verify the condition for Lemma 17 to hold; we will show that

$$\sup_{P \in \mathcal{P}_I} \frac{\sum_{i=1}^I \mathbb{E}_P \left[|\varphi(S_i; \beta, \eta_0, \nu^*)|^{2+2\Delta} \right]}{\left(\sum_{i=1}^I \mathbb{E}_P [\varphi^2(S_i; \beta, \eta_0, \nu^*)] \right)^{1+\Delta}} = o(1), \tag{29}$$

holds for $\Delta = \frac{\delta}{4}$. First, note that for an arbitrary $P \in \mathcal{P}_I$,

$$\begin{aligned}
\sum_{i=1}^I \mathbb{E}_P \left[|\varphi(S_i; \beta, \eta_0, \nu^*)|^{2+2\Delta} \right] &= \sum_{i=1}^I \mathbb{E}_P \left[\left| \varepsilon_i^\top W_i^* \xi_i \right|^{2+2\Delta} \right] \\
&\leq \sum_{i=1}^I \mathbb{E}_P \left[\|W_i^*\|_{\text{op}}^{2+2\Delta} \|\xi_i\|_2^{2+2\Delta} \|\varepsilon_i\|_2^{2+2\Delta} \right] \\
&\lesssim \sum_{i=1}^I \|\xi_i\|_{P,4+4\Delta}^{2+2\Delta} \|\varepsilon_i\|_{P,4+4\Delta}^{2+2\Delta} \lesssim \sum_{i=1}^I n_i^{2+2\Delta} \leq N n_{\max}^{1+2\Delta}, \\
\sum_{i=1}^I \mathbb{E}_P \left[\varphi^2(S_i; \beta, \eta_0, \nu^*) \right] &= \sum_{i=1}^I \mathbb{E}_P \left[\text{tr} \left(\Sigma_i W_i^* \xi_i \xi_i^\top W_i^* \right) \right] \\
&\geq \sum_{i=1}^I \mathbb{E}_P \left[\Lambda_{\min} \left(W_i^{*\frac{1}{2}} \Sigma_i W_i^{*\frac{1}{2}} \right) \|W_i^*\|_{\text{op}} \text{tr} \left(\xi_i \xi_i^\top \right) \right] \\
&\gtrsim \sum_{i=1}^I n_i^{-\kappa} \mathbb{E}_P \left[\text{tr} \left(\xi_i \xi_i^\top \right) \right] = \sum_{i=1}^I n_i^{-\kappa} \mathbb{E}_P \left[\text{tr} \left(\Omega_i \right) \right] \\
&\geq \sum_{i=1}^I n_i^{1-\kappa} \mathbb{E}_P \left[\Lambda_{\min} \left(\Omega_i \right) \right] \gtrsim \sum_{i=1}^I n_i^{1-\kappa} \geq N n_{\max}^{-\kappa},
\end{aligned}$$

where the bounding of the numerator follows by the Cauchy–Schwarz inequality and Assumptions A1.1 and A3.3, and the bounding of the denominator follows by Assumptions A1.2 and A3.3. Therefore,

$$\sup_{P \in \mathcal{P}_I} \frac{\sum_{i=1}^I \mathbb{E}_P \left[|\varphi(S_i; \beta, \eta_0, \nu^*)|^{2+2\Delta} \right]}{\left(\sum_{i=1}^I \mathbb{E}_P \left[\varphi^2(S_i; \beta, \eta_0, \nu^*) \right] \right)^{1+\Delta}} \leq N^{-\Delta} n_{\max}^{(1+\kappa)+(2+\kappa)\Delta} = o(1)$$

by Assumption A4, with $\Delta = \frac{\delta}{4}$. Therefore, the condition (29) holds, and so the claim (28) is proved.

Further, $M_1 = o_{\mathcal{P}}(1)$ by Lemma 8 and (as shown already) for arbitrary $P \in \mathcal{P}_I$

$$N^{-1} n_{\max}^{\kappa} \sum_{i=1}^I \mathbb{E}_P \left[\varphi^2(S_i; \beta, \eta_0, \nu^*) \right] \gtrsim 1$$

therefore, by Lemma 18 in conjunction with (28) we have that

$$\lim_{I \rightarrow \infty} \sup_{P \in \mathcal{P}_I} \sup_{t \in \mathbb{R}} \left| \mathbb{P}_P \left(\frac{N^{-1} n_{\max}^{\frac{\kappa}{2}} \sum_{i=1}^I \varphi(S_i; \beta, \eta_0, \nu^*) + M_1}{\left(N^{-1} n_{\max}^{\kappa} \sum_{i=1}^I \mathbb{E}_P \left[\varphi^2(S_i; \beta, \eta_0, \nu^*) \right] \right)^{\frac{1}{2}}} \leq t \right) - \Phi(t) \right| = 0. \quad (30)$$

Next, we claim that

$$\frac{N^{-1} n_{\max}^{\gamma} \sum_{i=1}^I \phi(S_i; \eta_0, \nu^*) + M_2}{N^{-1} n_{\max}^{\gamma} \sum_{i=1}^I \mathbb{E}_P \left[\phi(S_i; \eta_0, \nu^*) \right]} = 1 + o_{\mathcal{P}}(1), \quad (31)$$

which we proceed to show. We claim that

$$N^{-1} n_{\max}^{\gamma} \sum_{i=1}^I \phi(S_i; \eta_0, \nu^*) = N^{-1} n_{\max}^{\gamma} \sum_{i=1}^I \mathbb{E}_P \left[\phi(S_i; \eta_0, \nu^*) \right] + o_{\mathcal{P}}(1).$$

This follows by Lemma 15 alongside

$$\begin{aligned}
& \mathbb{E}_P \left[\left(N^{-1} n_{\max}^\gamma \sum_{i=1}^I (\phi(S_i; \eta_0, \nu^*) - \mathbb{E}_P [\phi(S_i; \eta_0, \nu^*)]) \right)^2 \right] \\
& \leq N^{-2} n_{\max}^{2\gamma} \sum_{i=1}^I \mathbb{E}_P \left[(\phi(S_i; \eta_0, \nu^*) - \mathbb{E}_P [\phi(S_i; \eta_0, \nu^*)])^2 \right] \\
& \leq N^{-2} n_{\max}^{2\gamma} \sum_{i=1}^I \mathbb{E}_P \left[(\phi(S_i; \eta_0, \nu^*))^2 \right] \leq N^{-2} n_{\max}^{2\gamma} \sum_{i=1}^I \mathbb{E}_P \left[\|W_i^*\|_{\text{op}}^2 \|\xi_i\|_2^4 \right] \\
& \lesssim N^{-2} n_{\max}^{2\gamma} \sum_{i=1}^I n_i^2 \leq N^{-1} n_{\max}^{1+2\gamma} = o(1),
\end{aligned}$$

by the Cauchy–Schwarz inequality and Assumptions A1.1, A3.3 and A4. Because also $M_2 = o_P(1)$ by Lemma 9, as well as that

$$\begin{aligned}
N^{-1} n_{\max}^\gamma \sum_{i=1}^I \mathbb{E}_P [\phi(S_i; \eta_0, \nu^*)] &= N^{-1} n_{\max}^\gamma \sum_{i=1}^I \mathbb{E}_P \left[\text{tr} \left(W_i^* \xi_i \xi_i^\top \right) \right] \\
&\geq N^{-1} n_{\max}^\gamma \sum_{i=1}^I \mathbb{E}_P \left[\Lambda_{\min} (W_i^*) \text{tr} \left(\xi_i \xi_i^\top \right) \right] \\
&\gtrsim N^{-1} n_{\max}^\gamma \sum_{i=1}^I n_i^{-\gamma} \mathbb{E}_P [\text{tr} (\Omega_i)] \\
&\gtrsim N^{-1} n_{\max}^\gamma \sum_{i=1}^I n_i^{1-\gamma} \geq 1,
\end{aligned}$$

we have, by Lemma 12, proved the claim (31).

Finally, combining (30) with (31) with Lemma 18 completes the proof;

$$\lim_{I \rightarrow \infty} \sup_{P \in \mathcal{P}_I} \sup_{t \in \mathbb{R}} \left| \mathbb{P}_P \left(N^{\frac{1}{2}} V^{-\frac{1}{2}} (\hat{\beta} - \beta) \leq t \right) - \Phi(t) \right| = 0.$$

□

Proof of Theorem 4 (ii). Note that

$$\hat{V} = \left(N^{-1} \sum_{k=1}^K \sum_{i \in \mathcal{I}_k} \phi(S_i; \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) \right)^{-2} \left(N^{-1} \sum_{k=1}^K \sum_{i \in \mathcal{I}_k} \varphi^2(S_i; \hat{\beta}, \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) \right),$$

and so we can decompose

$$N^{\frac{1}{2}}\hat{V}^{-\frac{1}{2}}(\hat{\beta} - \beta) = \left(\frac{N^{-1} \sum_{k=1}^K \sum_{i \in \mathcal{I}_k} \varphi^2(S_i; \hat{\beta}, \hat{\eta}^{(k)}, \hat{\nu}^{(k)})}{N^{-1} \sum_{i=1}^I \mathbb{E}_P [\varphi^2(S_i; \beta, \eta_0, \nu^*)]} \right)^{-\frac{1}{2}} \cdot \left(\sum_{i=1}^I \mathbb{E}_P [\varphi^2(S_i; \beta, \eta_0, \nu^*)] \right)^{-\frac{1}{2}} \left(\sum_{k=1}^K \sum_{i \in \mathcal{I}_k} \varphi(S_i; \beta, \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) \right) \quad (32)$$

$$= \left(\frac{N^{-1} n_{\max}^{\kappa} \sum_{i=1}^I \mathbb{E}_P [\varphi^2(S_i; \beta, \eta_0, \nu^*)] + M_3 + M_4}{N^{-1} n_{\max}^{\kappa} \sum_{i=1}^I \mathbb{E}_P [\varphi^2(S_i; \beta, \eta_0, \nu^*)]} \right)^{-\frac{1}{2}}. \quad (33)$$

$$\cdot \left(\frac{N^{-\frac{1}{2}} n_{\max}^{\frac{\kappa}{2}} \sum_{i=1}^I \varphi(S_i; \beta, \eta_0, \nu^*) + M_1}{\left(N^{-1} n_{\max}^{\kappa} \sum_{i=1}^I \mathbb{E}_P [\varphi^2(S_i; \beta, \eta_0, \nu^*)] \right)^{\frac{1}{2}}} \right), \quad (34)$$

where

$$M_3 := N^{-1} n_{\max}^{\kappa} \sum_{k=1}^K \sum_{i \in \mathcal{I}_k} \left[\varphi^2(S_i; \hat{\beta}, \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) - \varphi^2(S_i; \beta, \eta_0, \nu^*) \right],$$

$$M_4 := N^{-1} n_{\max}^{\kappa} \sum_{i=1}^I \left[\varphi^2(S_i; \beta, \eta_0, \nu^*) - \mathbb{E}_P [\varphi^2(S_i; \beta, \eta_0, \nu^*)] \right].$$

Note that the latter term in this decomposition is considered in the proof of Theorem 4; specifically the result (30). Therefore it remains for us to show that

$$\frac{N^{-1} n_{\max}^{\kappa} \sum_{i=1}^I \mathbb{E}_P [\varphi^2(S_i; \beta, \eta_0, \nu^*)] + M_3 + M_4}{N^{-1} n_{\max}^{\kappa} \sum_{i=1}^I \mathbb{E}_P [\varphi^2(S_i; \beta, \eta_0, \nu^*)]} = 1 + o_{\mathcal{P}}(1). \quad (35)$$

From Lemmas 10 and 11 we have that M_3 and M_4 are both $o_{\mathcal{P}}(1)$. Further (as shown in the proof of Theorem 4),

$$N^{-1} n_{\max}^{\kappa} \sum_{i=1}^I \mathbb{E}_P [\varphi^2(S_i; \beta, \eta_0, \nu^*)] \gtrsim 1,$$

and so combining these results with Lemmas 12 and 14 yields the result (35). Finally, the results (30) and (35) in conjunction with Lemma 18 completes the proof. \square

D.2 Auxiliary results for the proof of Theorem 4: bounding terms

Lemma 7. *Under the setup and notation of Theorem 4, we have that*

$$\|\hat{W}_i - W_i^*\|_{P|_{S_{\mathcal{I}_k^c}, 2}} \lesssim n_i^{\frac{1+\tau}{2}} \mathcal{R}_{\sigma}^{\frac{1}{2}} + n_i \mathcal{R}_{\rho}^{\frac{1}{2}}.$$

Proof. We define

$$\Sigma_i(\sigma, \rho) := \left((\mathbb{1}_{\{j=k\}} + \mathbb{1}_{\{j \neq k\}} \rho(X_{ij}, X_{ik})) \sigma(X_{ij}) \sigma(X_{ik}) \right)_{(j,k) \in [n_i]^2},$$

and further $\Sigma_i^* := \Sigma_i(\sigma^*, \rho^*)$ and $\hat{\Sigma}_i := \Sigma_i(\hat{\sigma}^{(k)}, \hat{\rho}^{(k)})$ (suppressing dependence on $k \in [K]$).

Note that, as $\hat{W}_i - W_i^* \asymp \hat{\Sigma}_i^{-1} - \Sigma_i^{*-1}$,

$$\|\hat{W}_i - W_i^*\|_{\text{op}} \lesssim \|\hat{\Sigma}_i^{-1}(\Sigma_i^* - \hat{\Sigma}_i)\Sigma_i^{*-1}\|_{\text{op}} \leq \|\Sigma_i^{*-1}\|_{\text{op}}\|\hat{\Sigma}_i^{-1}\|_{\text{op}}\|\hat{\Sigma}_i - \Sigma_i^*\|_{\text{op}} \lesssim \|\hat{\Sigma}_i - \Sigma_i^*\|_{\text{op}},$$

and further,

$$\begin{aligned} \|\hat{\Sigma}_i - \Sigma_i^*\|_{\text{op}}^2 &\leq \|\hat{\Sigma}_i - \Sigma_i^*\|_F^2 \\ &= \sum_{j=1}^{n_i} \left(\hat{\sigma}^{(k)}(X_{ij})^2 - \sigma^*(X_{ij})^2 \right)^2 \\ &\quad + \sum_{j=1}^{n_i} \sum_{j'=1, j' \neq j}^{n_i} \left(\hat{\rho}^{(k)}(X_{ij}, X_{ij'}) \hat{\sigma}^{(k)}(X_{ij}) \hat{\sigma}^{(k)}(X_{ij'}) - \rho^*(X_{ij}, X_{ij'}) \sigma^*(X_{ij}) \sigma^*(X_{ij'}) \right)^2 \\ &\lesssim \sum_{j=1}^{n_i} \left(\hat{\sigma}^{(k)}(X_{ij}) - \sigma^*(X_{ij}) \right)^2 \left(1 + \sum_{j'=1, j' \neq j}^{n_i} \hat{\rho}^{(k)}(X_{ij}, X_{ij'})^2 \right) \\ &\quad + \sum_{j=1}^{n_i} \sum_{j'=1, j' \neq j}^{n_i} \left(\hat{\rho}^{(k)}(X_{ij}, X_{ij'}) - \rho^*(X_{ij}, X_{ij'}) \right)^2, \end{aligned}$$

where $\|\cdot\|_F$ denotes the Frobenius norm, and so

$$\mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\|\hat{\Sigma}_i - \Sigma_i^*\|_{\text{op}}^2 \right] \lesssim n_i^{1+\tau} \mathcal{R}_\sigma + n_i^2 \mathcal{R}_\rho,$$

and hence

$$\|\hat{W}_i - W_i^*\|_{P|S_{\mathcal{I}_k^c}, 2} = \left(\mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\|\hat{W}_i - W_i^*\|_{\text{op}}^2 \right] \right)^{\frac{1}{2}} \lesssim n_i^{\frac{1+\tau}{2}} \mathcal{R}_\sigma^{\frac{1}{2}} + n_i \mathcal{R}_\rho^{\frac{1}{2}}.$$

completing the proof. \square

Lemma 8. *Under the setup of Theorem 4,*

$$M_1 := N^{-\frac{1}{2}} n_{\max}^{\frac{\kappa}{2}} \sum_{k=1}^K \sum_{i \in \mathcal{I}_k} \left[\varphi(S_i; \beta, \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) - \varphi(S_i; \beta, \eta_0, \nu^*) \right] = o_{\mathcal{P}}(1).$$

Proof. For $k \in [K]$ and $i \in \mathcal{I}_k$, define the $(\hat{l}^{(k)}, \hat{m}^{(k)})$ errors

$$\begin{aligned} \mathcal{E}_m^{(k)}(X_i) &:= \hat{m}^{(k)}(X_i) - m(X_i), & \mathcal{E}_l^{(k)}(X_i) &:= \hat{l}^{(k)}(X_i) - l(X_i), \\ \mathcal{E}_{l-\beta m}^{(k)}(X_i) &:= \mathcal{E}_l^{(k)}(X_i) - \beta \mathcal{E}_m^{(k)}(X_i). \end{aligned}$$

The term M_1 can then be decomposed as $M_1 = \sum_{k=1}^K M_{1,k}$, where

$$\begin{aligned} M_{1,k} &:= \underbrace{N^{-\frac{1}{2}} n_{\max}^{\frac{\kappa}{2}} \sum_{i \in \mathcal{I}_k} \mathcal{E}_{l-\beta m}^{(k)}(X_i)^\top \hat{W}_i \mathcal{E}_m^{(k)}(X_i)}_{\text{I}} + \underbrace{N^{-\frac{1}{2}} n_{\max}^{\frac{\kappa}{2}} \sum_{i \in \mathcal{I}_k} \mathcal{E}_{l-\beta m}^{(k)}(X_i)^\top \hat{W}_i \xi_i}_{\text{II}} \\ &\quad + \underbrace{N^{-\frac{1}{2}} n_{\max}^{\frac{\kappa}{2}} \sum_{i \in \mathcal{I}_k} \mathcal{E}_m^{(k)}(X_i)^\top W_i^* \varepsilon_i}_{\text{III}} + \underbrace{N^{-\frac{1}{2}} n_{\max}^{\frac{\kappa}{2}} \sum_{i \in \mathcal{I}_k} \left(\xi_i + \mathcal{E}_m^{(k)}(X_i) \right)^\top \left(\hat{W}_i - W_i^* \right) \varepsilon_i}_{\text{IV}}. \end{aligned}$$

We show that, for an arbitrary $k \in [K]$, each of these four terms are $o_{\mathcal{P}}(1)$ in turn:

Term I: Note that for an arbitrary $P \in \mathcal{P}_I$,

$$\begin{aligned}
& \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\left| N^{-\frac{1}{2}} n_{\max}^{\frac{\kappa}{2}} \sum_{i \in \mathcal{I}_k} \mathcal{E}_{l-\beta m}^{(k)}(X_i)^\top \hat{W}_i \mathcal{E}_m^{(k)}(X_i) \right| \right] \\
& \leq N^{-\frac{1}{2}} n_{\max}^{\frac{\kappa}{2}} \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\sum_{i \in \mathcal{I}_k} \left| \mathcal{E}_{l-\beta m}^{(k)}(X_i)^\top \hat{W}_i \mathcal{E}_m^{(k)}(X_i) \right| \right] \\
& \leq N^{-\frac{1}{2}} n_{\max}^{\frac{\kappa}{2}} \sum_{i \in \mathcal{I}_k} \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\left| \mathcal{E}_{l-\beta m}^{(k)}(X_i)^\top \hat{W}_i \mathcal{E}_{l-\beta m}^{(k)}(X_i) \right|^2 \right]^{\frac{1}{2}} \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\left| \mathcal{E}_m^{(k)}(X_i)^\top \hat{W}_i \mathcal{E}_m^{(k)}(X_i) \right|^2 \right]^{\frac{1}{2}} \\
& \lesssim N^{-\frac{1}{2}} n_{\max}^{\frac{\kappa}{2}} \sum_{i \in \mathcal{I}_k} \left\| \mathcal{E}_{l-\beta m}^{(k)}(X_i) \right\|_{P|S_{\mathcal{I}_k^c}, 2} \left\| \mathcal{E}_m^{(k)}(X_i) \right\|_{P|S_{\mathcal{I}_k^c}, 2} \\
& \leq N^{-\frac{1}{2}} n_{\max}^{\frac{\kappa}{2}} \sum_{i \in \mathcal{I}_k} n_i \mathcal{R}_m^2 (\mathcal{R}_l + \mathcal{R}_m)^{\frac{1}{2}} = N^{\frac{1}{2}} n_{\max}^{\frac{\kappa}{2}} \mathcal{R}_m^{\frac{1}{2}} (\mathcal{R}_l + \mathcal{R}_m)^{\frac{1}{2}} = o_{\mathcal{P}}(1),
\end{aligned}$$

by Assumption A4. The result that this term I is $o_{\mathcal{P}}(1)$ follows directly by Lemma 15.

Term II: First note that $\mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\mathcal{E}_{l-\beta m}^{(k)}(X_i)^\top \hat{W}_i \xi_i \right] = 0$ for any $P \in \mathcal{P}_I$ and $i \in \mathcal{I}_k$. Then we have

$$\begin{aligned}
& \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\left(N^{-\frac{1}{2}} n_{\max}^{\frac{\kappa}{2}} \sum_{i \in \mathcal{I}_k} \mathcal{E}_{l-\beta m}^{(k)}(X_i)^\top \hat{W}_i \xi_i \right)^2 \right] \\
& = N^{-1} n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\left(\mathcal{E}_{l-\beta m}^{(k)}(X_i)^\top \hat{W}_i \xi_i \right)^2 \right] \\
& = N^{-1} n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\mathcal{E}_{l-\beta m}^{(k)}(X_i)^\top \hat{W}_i \Omega_i \hat{W}_i \mathcal{E}_{l-\beta m}^{(k)}(X_i) \right] \\
& \leq N^{-1} n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\Lambda_{\max}(\Omega_i) \left\| \hat{W}_i \right\|_{\text{op}}^2 \left\| \mathcal{E}_{l-\beta m}^{(k)}(X_i) \right\|_{P|S_{\mathcal{I}_k^c}, 2}^2 \right] \\
& \lesssim N^{-1} n_{\max}^{\kappa + \mu \Omega} (\mathcal{R}_l + \mathcal{R}_m) = o_{\mathcal{P}}(1),
\end{aligned}$$

by Assumption A4. The result that this term II is $o_{\mathcal{P}}(1)$ follows directly by Lemma 15.

Term III: First note that $\mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\mathcal{E}_k^m(X_i)^\top \hat{W}_i \varepsilon_i \right] = 0$ for any $P \in \mathcal{P}_I$ and $i \in \mathcal{I}_k$. Then

$$\begin{aligned}
& \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\left(N^{-\frac{1}{2}} n_{\max}^{\frac{\kappa}{2}} \sum_{i \in \mathcal{I}_k} \mathcal{E}_k^m(X_i)^\top W_i^* \varepsilon_i \right)^2 \right] \\
& = N^{-1} n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\left(\mathcal{E}_k^m(X_i)^\top W_i^* \varepsilon_i \right)^2 \right] \\
& = N^{-1} n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\mathcal{E}_m^{(k)}(X_i)^\top W_i^* \Sigma_i W_i^* \mathcal{E}_m^{(k)}(X_i) \right] \\
& \leq N^{-1} n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\Lambda_{\max}(W_i^{*\frac{1}{2}} \Sigma_i W_i^{*\frac{1}{2}}) \left\| W_i^* \right\|_{\text{op}} \left\| \mathcal{E}_m^{(k)}(X_i) \right\|_2^2 \right] \\
& \lesssim N^{-1} n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} n_i^{1+\alpha} \mathcal{R}_m \leq n_{\max}^{\kappa+\alpha} \mathcal{R}_m = o_{\mathcal{P}}(1),
\end{aligned}$$

by Assumption A4. The result that this term II is $o_{\mathcal{P}}(1)$ follows directly by Lemma 15.

Term IV: Similarly, note that $\mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\left(\xi_i + \mathcal{E}_m^{(k)}(X_i) \right)^\top (\hat{W}_i - W_i^*) \varepsilon_i \right] = 0$ for any $P \in \mathcal{P}_I$ and $i \in \mathcal{I}_k$. Then

$$\begin{aligned}
& \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\left(\left(\xi_i + \mathcal{E}_m^{(k)}(X_i) \right)^\top (\hat{W}_i - W_i^*) \varepsilon_i \right)^2 \right] \\
&= \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\left(\xi_i + \mathcal{E}_m^{(k)}(X_i) \right)^\top (\hat{W}_i - W_i^*) \Sigma_i (\hat{W}_i - W_i^*) \left(\xi_i + \mathcal{E}_m^{(k)}(X_i) \right) \right] \\
&\leq \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\Lambda_{\max}(\Sigma_i) \|\hat{W}_i - W_i^*\|_{\text{op}}^2 \left(\|\xi_i\|_2^2 + \|\mathcal{E}_m^{(k)}(X_i)\|_2^2 \right) \right] \\
&\leq \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\Lambda_{\max}(\Sigma_i) \|\hat{W}_i - W_i^*\|_{\text{op}} \left(\|\xi_i\|_2^2 + \|\mathcal{E}_m^{(k)}(X_i)\|_2^2 \right) \right] \\
&\lesssim n_i^{\mu_\Sigma} \|\hat{W}_i - W_i^*\|_{P|S_{\mathcal{I}_k^c}, 2} \left(\|\xi_i\|_{P|S_{\mathcal{I}_k^c}, 4}^2 + \|\mathcal{E}_m^{(k)}(X_i)\|_{P|S_{\mathcal{I}_k^c}, 4}^2 \right) \\
&\lesssim n_i^{1+\mu_\Sigma} \left(n_i^{\frac{1+\tau}{2}} \mathcal{R}_\sigma^{\frac{1}{2}} + n_i \mathcal{R}_\rho^{\frac{1}{2}} \right)
\end{aligned}$$

with the second inequality holding because $\|\hat{W}_i - W_i^*\|_{\text{op}} \leq \|\hat{W}_i\|_{\text{op}} + \|W_i^*\|_{\text{op}} \lesssim 1$ almost surely, and the fourth inequality obtained by invoking Lemma 7. Therefore

$$\begin{aligned}
& \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\left(N^{-\frac{1}{2}} n_{\max}^{\frac{\kappa}{2}} \sum_{i \in \mathcal{I}_k} \left(\xi_i + \mathcal{E}_m^{(k)}(X_i) \right)^\top (\hat{W}_i - W_i^*) \varepsilon_i \right)^2 \right] \\
&\leq N^{-1} n_{\max}^\kappa \sum_{i \in \mathcal{I}_k} \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\left(\left(\xi_i + \mathcal{E}_m^{(k)}(X_i) \right)^\top (\hat{W}_i - W_i^*) \varepsilon_i \right)^2 \right] \\
&\lesssim n_{\max}^{\kappa+\mu_\Sigma} \left(n_{\max}^{\frac{1+\tau}{2}} \mathcal{R}_\sigma^{\frac{1}{2}} + n_i \mathcal{R}_\rho^{\frac{1}{2}} \right) = o_{\mathcal{P}}(1),
\end{aligned}$$

by Assumption A4. The result that this term IV is $o_{\mathcal{P}}(1)$ follows directly by Lemma 15.

All four terms are therefore $o_{\mathcal{P}}(1)$ and so applying Lemma 12 completes the proof. \square

Lemma 9. *Under the setup of Theorem 4,*

$$M_2 := N^{-1} n_{\max}^\gamma \sum_{k=1}^K \sum_{i \in \mathcal{I}_k} \left[\phi(S_i; \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) - \phi(S_i; \eta_0, \nu^*) \right] = o_{\mathcal{P}}(1).$$

Proof. We adopt the same notation as in Lemma 8. Then we can decompose M_2 into $M_2 = \sum_{k=1}^K M_{2,k}$, where

$$\begin{aligned}
M_{2,k} &:= \underbrace{N^{-1} n_{\max}^\gamma \sum_{i \in \mathcal{I}_k} \mathcal{E}_m^{(k)}(X_i) \hat{W}_i \mathcal{E}_m^{(k)}(X_i)}_{\text{V}} - 2 \underbrace{N^{-1} n_{\max}^\gamma \sum_{i \in \mathcal{I}_k} \xi_i^\top \hat{W}_i \mathcal{E}_m^{(k)}(X_i)}_{\text{VI}} \\
&\quad + \underbrace{N^{-1} n_{\max}^\gamma \sum_{i \in \mathcal{I}_k} \xi_i^\top (\hat{W}_i - W_i^*) \xi_i}_{\text{VII}}.
\end{aligned}$$

For an arbitrary $k \in [K]$ we will show that each of these three terms are each $o_{\mathcal{P}}(1)$:

Term V: We have that

$$\begin{aligned} \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[N^{-1} n_{\max}^\gamma \sum_{i \in \mathcal{I}_k} \mathcal{E}_m^{(k)}(X_i)^\top \hat{W}_i \mathcal{E}_m^{(k)}(X_i) \right] &\lesssim N^{-1} n_{\max}^\gamma \sum_{i \in \mathcal{I}_k^c} \|\mathcal{E}_m^{(k)}(X_i)\|_{P|S_{\mathcal{I}_k^c}, 2}^2 \\ &\lesssim n_{\max}^\gamma \mathcal{R}_m = o_{\mathcal{P}}(1), \end{aligned}$$

and so (by Lemma 15) this term is $o_{\mathcal{P}}(1)$.

Term VI:

$$\begin{aligned} &\mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\left(N^{-1} n_{\max}^\gamma \sum_{i \in \mathcal{I}_k} \mathcal{E}_m^{(k)}(X_i)^\top \hat{W}_i \xi \right)^2 \right] \\ &= N^{-2} n_{\max}^{2\gamma} \sum_{i \in \mathcal{I}_k} \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\mathcal{E}_m^{(k)}(X_i)^\top \hat{W}_i \Omega_i \hat{W}_i \mathcal{E}_m^{(k)}(X_i)^\top \right] \\ &\lesssim N^{-2} n_{\max}^{2\gamma} \sum_{i \in \mathcal{I}_k} n_i^{1+\mu_\Omega} \mathcal{R}_m \leq N^{-1} n_{\max}^{2\gamma+\mu_\Omega} \mathcal{R}_m = o_{\mathcal{P}}(1), \end{aligned}$$

and so (again by Lemma 15) this term is $o_{\mathcal{P}}(1)$.

Term VII:

$$\begin{aligned} &\mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\left| N^{-1} n_{\max}^\gamma \sum_{i \in \mathcal{I}_k} \xi_i^\top (\hat{W}_i - W_i^*) \xi \right| \right] \\ &\leq N^{-1} n_{\max}^\gamma \sum_{i \in \mathcal{I}_k} \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\|\hat{W}_i - W_i^*\|_{\text{op}} \|\xi_i\|_2^2 \right] \\ &\lesssim n_{\max} \left(n_{\max}^{\frac{1+\tau}{2}} \mathcal{R}_\sigma^{\frac{1}{2}} + n_{\max} \mathcal{R}_\rho^{\frac{1}{2}} \right) = o_{\mathcal{P}}(1), \end{aligned}$$

and so similarly this term is $o_{\mathcal{P}}(1)$.

Combining these results with Lemma 12 proves the claim. \square

Lemma 10. *Under the setup of Theorem 4,*

$$M_3 := N^{-1} n_{\max}^\kappa \sum_{k=1}^K \sum_{i \in \mathcal{I}_k} \left[\varphi^2(S_i; \hat{\beta}, \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) - \varphi^2(S_i; \beta, \eta_0, \nu^*) \right] = o_{\mathcal{P}}(1).$$

Proof. First, decompose the M_3 term as $M_3 = \sum_{k=1}^K M_{3,k}$, where

$$\begin{aligned}
M_{3,k} &:= N^{-1} n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \left[\varphi^2(S_i; \hat{\beta}, \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) - \varphi^2(S_i; \beta, \eta_0, \nu^*) \right] \\
&= N^{-1} n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \left(\varphi(S_i; \hat{\beta}, \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) - \varphi(S_i; \beta, \eta_0, \nu^*) \right)^2 \\
&\quad + N^{-1} n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \varphi(S_i; \beta, \eta_0, \nu^*) \left(\varphi(S_i; \hat{\beta}, \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) - \varphi(S_i; \beta, \eta_0, \nu^*) \right) \\
&\leq 2N^{-1} n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \left(\varphi(S_i; \beta, \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) - \varphi(S_i; \beta, \eta_0, \nu^*) \right)^2 \\
&\quad + 2N^{-1} n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \left(\varphi(S_i; \hat{\beta}, \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) - \varphi(S_i; \beta, \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) \right)^2 \\
&\quad + N^{-1} n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \varphi(S_i; \beta, \eta_0, \nu^*) \left(\varphi(S_i; \beta, \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) - \varphi(S_i; \beta, \eta_0, \nu^*) \right) \\
&\quad + N^{-1} n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \varphi(S_i; \beta, \eta_0, \nu^*) \left(\varphi(S_i; \hat{\beta}, \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) - \varphi(S_i; \beta, \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) \right) \\
&= 2N^{-1} n_{\max}^{\kappa} \underbrace{\sum_{i \in \mathcal{I}_k} \left(\varphi(S_i; \beta, \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) - \varphi(S_i; \beta, \eta_0, \nu^*) \right)^2}_{\text{VIII}} \\
&\quad + N^{-1} n_{\max}^{\kappa} \underbrace{\sum_{i \in \mathcal{I}_k} \varphi(S_i; \beta, \eta_0, \nu^*) \left(\varphi(S_i; \beta, \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) - \varphi(S_i; \beta, \eta_0, \nu^*) \right)}_{\text{IX}} \\
&\quad + 2N^{-1} n_{\max}^{\kappa} \underbrace{\sum_{i \in \mathcal{I}_k} \phi^2(S_i; \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) (\hat{\beta} - \beta)^2}_{\text{X}} \\
&\quad + N^{-1} n_{\max}^{\kappa} \underbrace{\sum_{i \in \mathcal{I}_k} \varphi(S_i; \beta, \eta_0, \nu^*) \phi(S_i; \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) (\hat{\beta} - \beta)}_{\text{XI}}
\end{aligned}$$

We claim that each term is $o_{\mathcal{P}}(1)$, and prove this for each term separately:

Term VIII: We decompose this term into the four subsequent terms

$$\begin{aligned}
&2N^{-1} n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \left(\varphi(S_i; \beta, \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) - \varphi(S_i; \beta, \eta_0, \nu^*) \right)^2 \\
&\lesssim \underbrace{N^{-1} n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \left| \mathcal{E}_{l-\beta m}^{(k)}(X_i)^\top \hat{W}_i \mathcal{E}_m^{(k)}(X_i) \right|^2}_{\text{VIIIa}} + \underbrace{N^{-1} n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \left(\mathcal{E}_{l-\beta m}^{(k)}(X_i)^\top \hat{W}_i \xi_i \right)^2}_{\text{VIIIb}} \\
&+ \underbrace{N^{-1} n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \left(\mathcal{E}_m^{(k)}(X_i)^\top W_i^* \varepsilon_i \right)^2}_{\text{VIIIc}} + \underbrace{N^{-1} n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \left(\left(\xi_i + \mathcal{E}_m^{(k)}(X_i) \right)^\top (\hat{W}_i - W_i^*) \varepsilon_i \right)^2}_{\text{VIIId}}.
\end{aligned}$$

Note that

$$\begin{aligned} & \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\left(N^{-1} n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \left| \mathcal{E}_{l-\beta m}^{(k)}(X_i)^\top \hat{W}_i \mathcal{E}_m^{(k)}(X_i) \right|^2 \right)^{\frac{1}{2}} \right] \\ & \leq N^{-\frac{1}{2}} n_{\max}^{\frac{\kappa}{2}} \sum_{i \in \mathcal{I}_k} \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\left| \mathcal{E}_{l-\beta m}^{(k)}(X_i)^\top \hat{W}_i \mathcal{E}_m^{(k)}(X_i) \right| \right] = o_{\mathcal{P}}(1) \end{aligned}$$

as shown when dealing with Term I in the decomposition for Lemma 8, and so the square root of VIIIa and hence VIIIa are both $o_{\mathcal{P}}(1)$. Further, the expectation under the law of $P|S_{\mathcal{I}_k^c}$ of the terms VIIIb, VIIIc and VIIId are each shown in Lemma 8 to be $o_{\mathcal{P}}(1)$, and so the terms VIIIb, VIIIc and VIIId are all $o_{\mathcal{P}}(1)$, thus (again by Lemma 15) the term V is $o_{\mathcal{P}}(1)$.

Term IX: Note that this term can be further decomposed into the four terms

$$\begin{aligned} & N^{-1} n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \varphi(S_i; \beta, \eta_0, \nu^*) \left(\varphi(S_i; \beta, \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) - \varphi(S_i; \beta, \eta_0, \nu^*) \right) \\ & = N^{-1} n_{\max}^{\kappa} \underbrace{\sum_{i \in \mathcal{I}_k} (\xi_i^\top W_i^* \varepsilon_i) \left(\mathcal{E}_{l-\beta m}^{(k)}(X_i)^\top \hat{W}_i \mathcal{E}_m^{(k)}(X_i) \right)}_{\text{IXa}} \\ & \quad + \underbrace{N^{-1} n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} (\xi_i^\top W_i^* \varepsilon_i) \left(\mathcal{E}_{l-\beta m}^{(k)}(X_i)^\top \hat{W}_i^* \xi_i \right)}_{\text{IXb}} + \underbrace{N^{-1} n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} (\xi_i^\top W_i^* \varepsilon_i) \left(\mathcal{E}_m^{(k)}(X_i)^\top W_i^* \varepsilon_i \right)}_{\text{IXc}} \\ & \quad + \underbrace{N^{-1} n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \left((\xi_i + \mathcal{E}_m^{(k)}(X_i))^\top W_i^* \varepsilon_i \right) \left(\xi_i^\top (\hat{W}_i - W_i^*) \varepsilon_i \right)}_{\text{IXd}}. \end{aligned}$$

We claim that each of these four terms are $o_{\mathcal{P}}(1)$; we will show each in turn. For the first term IXa note that $\mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[(\xi_i^\top W_i^* \varepsilon_i) \left(\mathcal{E}_{l-\beta m}^{(k)}(X_i)^\top \hat{W}_i \mathcal{E}_m^{(k)}(X_i) \right) \right] = 0$ for all $P \in \mathcal{P}_I$ and $i \in \mathcal{I}_k$. We consider the cases of $\delta \in [0, 4)$ and $\delta \geq 4$ separately. If $\delta \in [0, 4)$ then by the von Bahr–Esseen inequality

$$\begin{aligned} & \mathbb{E}_P \left[\left| N^{-1} n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} (\xi_i^\top W_i^* \varepsilon_i) \left(\mathcal{E}_{l-\beta m}^{(k)}(X_i)^\top \hat{W}_i \mathcal{E}_m^{(k)}(X_i) \right) \right|^{1+\frac{\delta}{4}} \right] \\ & \leq (2 - I^{-1}) N^{-1-\frac{\delta}{4}} n_{\max}^{\kappa(1+\frac{\delta}{4})} \sum_{i \in \mathcal{I}_k} \mathbb{E}_P \left[\left| (\xi_i^\top W_i^* \varepsilon_i) \left(\mathcal{E}_{l-\beta m}^{(k)}(X_i)^\top \hat{W}_i \mathcal{E}_m^{(k)}(X_i) \right) \right|^{1+\frac{\delta}{4}} \right] \\ & \lesssim N^{-1-\frac{\delta}{4}} n_{\max}^{\kappa(1+\frac{\delta}{4})} \sum_{i \in \mathcal{I}_k} \mathbb{E}_P \left[\|\xi_i\|_2^{1+\frac{\delta}{4}} \|\varepsilon_i\|_2^{1+\frac{\delta}{4}} \|\mathcal{E}_{l-\beta m}^{(k)}(X_i)\|_2^{1+\frac{\delta}{4}} \|\mathcal{E}_m^{(k)}(X_i)\|_2^{1+\frac{\delta}{4}} \right] \\ & \lesssim N^{-1-\frac{\delta}{4}} n_{\max}^{\kappa(1+\frac{\delta}{4})} \sum_{i \in \mathcal{I}_k} \|\xi_i\|_{P, 4+\delta}^{1+\frac{\delta}{4}} \|\varepsilon_i\|_{P, 4+\delta}^{1+\frac{\delta}{4}} \|\mathcal{E}_{l-\beta m}^{(k)}(X_i)\|_{P|S_{\mathcal{I}_k^c}, 4+\delta}^{1+\frac{\delta}{4}} \|\mathcal{E}_m^{(k)}(X_i)\|_{P|S_{\mathcal{I}_k^c}, 4+\delta}^{1+\frac{\delta}{4}} \\ & \lesssim N^{-1-\frac{\delta}{4}} n_{\max}^{\kappa(1+\frac{\delta}{4})} \sum_{i \in \mathcal{I}_k} n_i^{2+\frac{\delta}{2}} \leq N^{-\frac{\delta}{4}} n_{\max}^{(1+\kappa)+(2+\kappa)\frac{\delta}{4}}, \end{aligned}$$

for any $P \in \mathcal{P}_I$, and thus

$$\begin{aligned} \sup_{P \in \mathcal{P}_I} \mathbb{E}_P \left[\left| N^{-1} n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} (\xi_i^\top W_i^* \varepsilon_i) \left(\mathcal{E}_{l-\beta m}^{(k)}(X_i)^\top \hat{W}_i \mathcal{E}_m^{(k)}(X_i) \right) \right|^{1+\frac{\delta}{4}} \right] & \lesssim N^{-\frac{\delta}{4}} n_{\max}^{(1+\kappa)+(2+\kappa)\frac{\delta}{4}} \\ & = o(1), \end{aligned}$$

by Assumption A4. If on the other hand $\delta \geq 4$ then note that

$$\begin{aligned}
& \mathbb{E}_P \left[\left(N^{-1} n_{\max}^\kappa \sum_{i \in \mathcal{I}_k} (\xi_i^\top W_i^* \varepsilon_i) \left(\mathcal{E}_{l-\beta m}^{(k)}(X_i)^\top \hat{W}_i \mathcal{E}_m^{(k)}(X_i) \right) \right)^2 \right] \\
&= N^{-2} n_{\max}^{2\kappa} \sum_{i \in \mathcal{I}_k} \mathbb{E}_P \left[\left| (\xi_i^\top W_i^* \varepsilon_i) \left(\mathcal{E}_{l-\beta m}^{(k)}(X_i)^\top \hat{W}_i \mathcal{E}_m^{(k)}(X_i) \right) \right|^2 \right] \\
&\lesssim N^{-2} n_{\max}^{2\kappa} \sum_{i \in \mathcal{I}_k} \mathbb{E}_P \left[\|\xi_i\|_2^2 \|\varepsilon_i\|_2^2 \|\mathcal{E}_{l-\beta m}^{(k)}(X_i)\|_2^2 \|\mathcal{E}_m^{(k)}(X_i)\|_2^2 \right] \\
&\lesssim N^{-2} n_{\max}^{2\kappa} \sum_{i \in \mathcal{I}_k} \|\xi_i\|_{P,8}^2 \|\varepsilon_i\|_{P,8}^2 \|\mathcal{E}_{l-\beta m}^{(k)}(X_i)\|_{P|S_{\mathcal{I}_k^c},8}^2 \|\mathcal{E}_m^{(k)}(X_i)\|_{P|S_{\mathcal{I}_k^c},8}^2 \\
&\lesssim N^{-2} n_{\max}^{2\kappa} \sum_{i \in \mathcal{I}_k} n_i^4 \leq N^{-1} n_{\max}^{3+2\kappa},
\end{aligned}$$

for all $P \in \mathcal{P}_I$, and thus

$$\sup_{P \in \mathcal{P}_I} \mathbb{E}_P \left[\left(N^{-1} n_{\max}^\kappa \sum_{i \in \mathcal{I}_k} (\xi_i^\top W_i^* \varepsilon_i) \left(\mathcal{E}_{l-\beta m}^{(k)}(X_i)^\top \hat{W}_i \mathcal{E}_m^{(k)}(X_i) \right) \right)^2 \right] \lesssim N^{-1} n_{\max}^{3+2\kappa} = o(1),$$

by Assumption A4. Hence, by Lemma 15, the term IXa is $o_P(1)$.

For the term IXb, note that as $\mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[(\xi_i^\top W_i \varepsilon_i) \left(\mathcal{E}_{l-\beta m}^{(k)}(X_i)^\top \hat{W}_i \xi_i \right) \right] = 0$ for any $P \in \mathcal{P}_I$ and $i \in \mathcal{I}_k$, and so the same methodology as for the term IXa can be used for IXb, showing that this term is also $o_P(1)$.

For the term IXc, note that

$$\begin{aligned}
& \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\left| N^{-1} n_{\max}^\kappa \sum_{i \in \mathcal{I}_k} (\xi_i^\top W_i^* \varepsilon_i) \left(\mathcal{E}_m^{(k)}(X_i)^\top W_i^* \varepsilon_i \right) \right| \right] \\
&\leq N^{-1} n_{\max}^\kappa \sum_{i \in \mathcal{I}_k} \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\left| (\xi_i^\top W_i^* \varepsilon_i) \left(\mathcal{E}_m^{(k)}(X_i)^\top W_i^* \varepsilon_i \right) \right| \right] \\
&\leq N^{-1} n_{\max}^\kappa \sum_{i \in \mathcal{I}_k} \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\left((\xi_i^\top W_i^* \varepsilon_i)^2 \right)^{\frac{1}{2}} E_{P|S_{\mathcal{I}_k^c}} \left[\left(\mathcal{E}_m^{(k)}(X_i)^\top W_i^* \varepsilon_i \right)^2 \right]^{\frac{1}{2}} \right] \\
&= N^{-1} n_{\max}^\kappa \sum_{i \in \mathcal{I}_k} \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\xi_i^\top W_i^* \Sigma_i W_i^* \xi_i \right]^{\frac{1}{2}} E_{P|S_{\mathcal{I}_k^c}} \left[\mathcal{E}_m^{(k)}(X_i)^\top W_i^* \Sigma_i W_i^* \mathcal{E}_m^{(k)}(X_i) \right]^{\frac{1}{2}} \\
&\leq N^{-1} n_{\max}^\kappa \sum_{i \in \mathcal{I}_k} n_i^{1+\alpha} \mathcal{R}_m^{\frac{1}{2}} \leq n_{\max}^{\kappa+\alpha} \mathcal{R}_m^{\frac{1}{2}} = o_P(1),
\end{aligned}$$

Similarly, for the term IXd,

$$\begin{aligned}
& \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\left| N^{-1} n_{\max}^\kappa \sum_{i \in \mathcal{I}_k} (\xi_i^\top W_i^* \varepsilon_i) \left((\xi_i + \mathcal{E}_m^{(k)}(X_i))^\top (\hat{W}_i - W_i^*) \varepsilon_i \right) \right| \right] \\
&\leq N^{-1} n_{\max}^\kappa \sum_{i \in \mathcal{I}_k} \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\left| (\xi_i^\top W_i^* \varepsilon_i) \left((\xi_i + \mathcal{E}_m^{(k)}(X_i))^\top (\hat{W}_i - W_i^*) \varepsilon_i \right) \right| \right] \\
&\leq N^{-1} n_{\max}^\kappa \sum_{i \in \mathcal{I}_k} \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\left((\xi_i^\top W_i^* \varepsilon_i)^2 \right)^{\frac{1}{2}} E_{P|S_{\mathcal{I}_k^c}} \left[\left((\xi_i + \mathcal{E}_m^{(k)}(X_i))^\top (\hat{W}_i - W_i^*) \varepsilon_i \right)^2 \right]^{\frac{1}{2}} \right] \\
&\lesssim n_{\max}^{\kappa+\frac{\alpha}{2}+\frac{\mu\Sigma}{2}} \left(n_{\max}^{\frac{1+\tau}{2}} \mathcal{R}_\sigma^{\frac{1}{2}} + n_{\max} \mathcal{R}_\rho^{\frac{1}{2}} \right) = o_P(1),
\end{aligned}$$

by Assumption A4. Hence, by Lemma 15, the term IXd is $o_P(1)$.

Term X: Note that

$$\begin{aligned}
& \mathbb{E}_P \left[\left| 2N^{-1}n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \phi^2(S_i; \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) (\hat{\beta} - \beta)^2 \right| \right] \\
& \leq 2N^{-1}n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \mathbb{E}_P \left[\phi^2(S_i; \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) (\hat{\beta} - \beta)^2 \right] \\
& \lesssim N^{-1}n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \left(\mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\|\xi_i\|_2^4 (\hat{\beta} - \beta)^2 \right] + \mathbb{E}_{P|S_{\mathcal{I}_k^c}} \left[\|\mathcal{E}_m^{(k)}(X_i)\|_2^4 (\hat{\beta} - \beta)^2 \right] \right) \\
& \leq N^{-1}n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \left(\|\xi_i\|_{P,4+\delta}^4 \|\hat{\beta} - \beta\|_{P, \frac{2(4+\delta)}{\delta}}^2 + \|\mathcal{E}_m^{(k)}(X_i)\|_{P,4+\delta}^4 \|\hat{\beta} - \beta\|_{P, \frac{2(4+\delta)}{\delta}}^2 \right) \\
& \leq N^{-2}n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} n_i^2 \left(\left\| n_i^{-\frac{1}{2}} \xi_i \right\|_{P,4+\delta}^4 \left\| N^{\frac{1}{2}} (\hat{\beta} - \beta) \right\|_{P, \frac{2(4+\delta)}{\delta}}^2 \right. \\
& \quad \left. + \left\| n_i^{-\frac{1}{2}} \mathcal{E}_m^{(k)}(X_i) \right\|_{P,4+\delta}^4 \left\| N^{\frac{1}{2}} (\hat{\beta} - \beta) \right\|_{P, \frac{2(4+\delta)}{\delta}}^2 \right) \\
& \lesssim N^{-2}n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} n_i^2 \leq N^{-1}n_{\max}^{1+\kappa},
\end{aligned}$$

by using the fact that any finite moment of a normal distribution is finite. Thus

$$\sup_{P \in \mathcal{P}_I} \mathbb{E}_P \left[\left| 2N^{-1}n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \phi^2(S_i; \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) (\hat{\beta} - \beta)^2 \right| \right] \lesssim N^{-1}n_{\max}^{1+\kappa} = o(1),$$

by Assumption A4. By Lemmas 12 and 15 we have that term X is $o_P(1)$.

Term XI: Note that

$$\begin{aligned}
& \mathbb{E}_P \left[\left| N^{-1}n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \varphi(S_i; \beta, \eta_0, \nu^*) \phi(S_i; \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) (\hat{\beta} - \beta) \right| \right] \\
& \leq N^{-1}n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \mathbb{E}_P \left[\|\xi_i\|_2^2 \left(\|\xi_i\|_2 + \|\mathcal{E}_m^{(k)}(X_i)\|_2 \right) \|\varepsilon_i\|_2 |\hat{\beta} - \beta| \right] \\
& \leq N^{-1}n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \left(\|\xi_i\|_{P,4+\delta}^3 \|\varepsilon_i\|_{P,4+\delta} \|\hat{\beta} - \beta\|_{P, \frac{4+\delta}{\delta}} \right. \\
& \quad \left. + \|\xi_i\|_{P,4+\delta}^2 \|\varepsilon_i\|_{P,4+\delta} \|\mathcal{E}_m^{(k)}(X_i)\|_{P,4+\delta} \|\hat{\beta} - \beta\|_{P, \frac{4+\delta}{\delta}} \right) \\
& \lesssim N^{-\frac{3}{2}}n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} n_i^2 \left\| N^{\frac{1}{2}} (\hat{\beta} - \beta) \right\|_{P, \frac{4+\delta}{\delta}} \\
& \lesssim N^{-\frac{3}{2}}n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} n_i^2 \leq N^{-\frac{1}{2}}n_{\max}^{1+\kappa},
\end{aligned}$$

by Hölder's inequality (twice), and thus

$$\sup_{P \in \mathcal{P}_I} \mathbb{E}_P \left[\left| N^{-1}n_{\max}^{\kappa} \sum_{i \in \mathcal{I}_k} \varphi(S_i; \beta, \eta_0, \nu^*) \phi(S_i; \hat{\eta}^{(k)}, \hat{\nu}^{(k)}) (\hat{\beta} - \beta) \right| \right] \lesssim N^{-\frac{1}{2}}n_{\max}^{1+\kappa} = o(1)$$

and Assumption A4. It then follows by Lemmas 12 and 15 that the term XI is $o_P(1)$. As all terms are $o_P(1)$, applying Lemma 12 completes the proof. \square

Lemma 11. *Under the setup of Theorem 4,*

$$M_4 := N^{-1} n_{\max}^{\kappa} \sum_{i=1}^I [\varphi^2(S_i; \beta, \eta_0, \nu^*) - \mathbb{E}_P [\varphi^2(S_i; \beta, \eta_0, \nu^*)]] = o_{\mathcal{P}}(1).$$

Proof. We will consider the two cases $\delta \in (0, 4)$ and $\delta \geq 4$ separately. If $\delta \in (0, 4)$ then, by the von Bahr–Esseen inequality, for all $P \in \mathcal{P}_I$

$$\begin{aligned} \mathbb{E}_P \left[|M_4|^{1+\frac{\delta}{4}} \right] &\leq N^{-(1+\frac{\delta}{4})} n_{\max}^{\kappa(1+\frac{\delta}{4})} \mathbb{E}_P \left[\left| \sum_{i=1}^I \varphi^2(S_i; \beta, \eta_0, \nu^*) - \mathbb{E}_P [\varphi^2(S_i; \beta, \eta_0, \nu^*)] \right|^{1+\frac{\delta}{4}} \right] \\ &\leq 2N^{-(1+\frac{\delta}{4})} n_{\max}^{\kappa(1+\frac{\delta}{4})} \sum_{i=1}^I \mathbb{E}_P \left[\left| \varphi^2(S_i; \beta, \eta_0, \nu^*) - \mathbb{E}_P [\varphi^2(S_i; \beta, \eta_0, \nu^*)] \right|^{1+\frac{\delta}{4}} \right] \\ &\leq 2N^{-(1+\frac{\delta}{4})} n_{\max}^{\kappa(1+\frac{\delta}{4})} \sum_{i=1}^I \mathbb{E}_P \left[|\varphi(S_i; \beta, \eta_0, \nu^*)|^{2+\frac{\delta}{2}} \right] \\ &\leq 2N^{-(1+\frac{\delta}{4})} n_{\max}^{\kappa(1+\frac{\delta}{4})} \sum_{i=1}^I \mathbb{E}_P \left[\|W_i^*\|_{\text{op}}^{2+\frac{\delta}{2}} \|\xi_i\|_2^{2+\frac{\delta}{2}} \|\varepsilon_i\|_2^{2+\frac{\delta}{2}} \right] \\ &\lesssim N^{-(1+\frac{\delta}{4})} n_{\max}^{\kappa(1+\frac{\delta}{4})} \sum_{i=1}^I \|\xi_i\|_{P, 4+\delta}^{2+\frac{\delta}{2}} \|\varepsilon_i\|_{P, 4+\delta}^{2+\frac{\delta}{2}} \\ &\lesssim N^{-(1+\frac{\delta}{4})} n_{\max}^{\kappa(1+\frac{\delta}{4})} \sum_{i=1}^I n_i^{2+\frac{\delta}{2}} \leq N^{-\frac{\delta}{4}} n_{\max}^{(1+\kappa)+(2+\kappa)\frac{\delta}{4}}, \end{aligned}$$

and so

$$\sup_{P \in \mathcal{P}_I} \mathbb{E}_P \left[|M_4|^{1+\frac{\delta}{4}} \right] \lesssim N^{-\frac{\delta}{4}} n_{\max}^{(1+\kappa)+(2+\kappa)\frac{\delta}{4}} = o(1),$$

by Assumption A4.

If $\delta \geq 4$ then, again for all $P \in \mathcal{P}_I$,

$$\begin{aligned} \mathbb{E}_P [M_4^2] &\leq N^{-2} n_{\max}^{2\kappa} \sum_{i=1}^I \mathbb{E}_P [\varphi^4(S_i; \beta, \eta_0, \nu^*)] \\ &\leq N^{-2} n_{\max}^{2\kappa} \sum_{i=1}^I \mathbb{E}_P \left[\left| \xi_i^\top W_i^* \varepsilon_i \right|^4 \right] \\ &\leq N^{-2} n_{\max}^{2\kappa} \sum_{i=1}^I \mathbb{E}_P \left[\|W_i^*\|_{\text{op}}^4 \|\xi_i\|_2^4 \|\varepsilon_i\|_2^4 \right] \\ &\lesssim N^{-2} n_{\max}^{2\kappa} \sum_{i=1}^I \|\xi_i\|_{P, 8}^4 \|\varepsilon_i\|_{P, 8}^4 \\ &\lesssim N^{-2} n_{\max}^{2\kappa} \sum_{i=1}^I n_i^4 \leq N^{-1} n_{\max}^{3+2\kappa}, \end{aligned}$$

and so

$$\sup_{P \in \mathcal{P}_I} \mathbb{E}_P [M_4^2] \lesssim N^{-1} n_{\max}^{3+2\kappa} = o(1),$$

by Assumptions A1.1 and A3.3.

Thus, in each case we have by applying Markov's inequality that $M_4 = o_{\mathcal{P}}(1)$. \square

D.3 Auxiliary uniform convergence results

For the uniform convergence results of Theorem 4 we require the following lemmata. Given a sequence of families of probability distributions $(\mathcal{P}_I)_{I \in \mathbb{N}}$ we use the shorthand $(A_I)_{I \in \mathbb{N}}$ to denote a sequence of random variables $(A_{P,I})_{P \in \mathcal{P}_I, I \in \mathbb{N}}$.

Lemma 12. *Let $(A_I)_{I \in \mathbb{N}}$ and $(B_I)_{I \in \mathbb{N}}$ be sequences of real-valued random variables. If $A_I = o_{\mathcal{P}}(1)$ and $B_I = o_{\mathcal{P}}(1)$ then $A_I + B_I = o_{\mathcal{P}}(1)$.*

Lemma 13. *Let $(A_I)_{I \in \mathbb{N}}$ and $(B_I)_{I \in \mathbb{N}}$ be sequences of real-valued random variables. Suppose that $A_I = o_{\mathcal{P}}(1)$ and $B_I \leq A_I$ for all $I \in \mathbb{N}$ uniformly over all $P \in \mathcal{P}_I$. Then $B_I = o_{\mathcal{P}}(1)$.*

Lemma 14. *Let $(A_I)_{I \in \mathbb{N}}$ and $(B_I)_{I \in \mathbb{N}}$ be sequences of real-valued random variables. If $A_I = o_{\mathcal{P}}(1)$ and $B_I = O_{\mathcal{P}}(1)$ then $A_I B_I = o_{\mathcal{P}}(1)$.*

Lemma 15. *Let $(A_I)_{I \in \mathbb{N}}$ and $(B_I)_{I \in \mathbb{N}}$ be a sequences of real-valued random variables. If there exists some $\alpha \geq 1$ such that $\mathbb{E}_P(|A_I|^\alpha | B_I) = o_{\mathcal{P}}(1)$, then $A_I = o_{\mathcal{P}}(1)$.*

For proofs of these four lemmas see, for example, [Lundborg et al. \[2022\]](#)

Lemma 16. *Let $(A_I^{(1)})_{I \in \mathbb{N}}, \dots, (A_I^{(K)})_{I \in \mathbb{N}}$ be sequences of real-valued random variables satisfying $A_I^{(k)} = o_{\mathcal{P}}(1)$ for all $k \in [K]$, where $K \in \mathbb{N}$ is finite. Then $\max_{k \in [K]} A_I^{(k)} = o_{\mathcal{P}}(1)$.*

Proof. For any $\epsilon > 0$

$$\sup_{P \in \mathcal{P}} \mathbb{P}_P \left(\left| \max_{k \in [K]} A_I^{(k)} \right| > \epsilon \right) \leq \sup_{P \in \mathcal{P}} \mathbb{P}_P \left(\max_{k \in [K]} |A_I^{(k)}| > \epsilon \right) \leq \sum_{k=1}^K \sup_{P \in \mathcal{P}} \mathbb{P}_P \left(|A_I^{(k)}| > \epsilon \right) = o(1),$$

completing the proof. □

Lemma 17. *Let $(A_{I,i})_{I \in \mathbb{N}, i \in [I]}$ be a triangular array of real-valued random variables satisfying:*

- $A_{I,1}, \dots, A_{I,I}$ are independent;
- $\mathbb{E}_P[A_{I,i}] = 0 \quad (\forall I \in \mathbb{N}, i \in [I]);$
- *There exists some $\Delta > 0$ such that*

$$\lim_{I \rightarrow \infty} \sup_{P \in \mathcal{P}_I} \frac{\sum_{i=1}^I \mathbb{E}_P \left[|A_{I,i}|^{2+\Delta} \right]}{\left(\sum_{i=1}^I \mathbb{E}_P \left[A_{I,i}^2 \right] \right)^{1+\frac{\Delta}{2}}} = 0.$$

Then $S_I := \left(\sum_{i=1}^I \mathbb{E}_P \left[A_{I,i}^2 \right] \right)^{-\frac{1}{2}} \left(\sum_{i=1}^I A_{I,i} \right)$ converges uniformly to $N(0,1)$, i.e.

$$\lim_{I \rightarrow \infty} \sup_{P \in \mathcal{P}_I} \sup_{t \in \mathbb{R}} |\mathbb{P}_P(S_I \leq t) - \Phi(t)| = 0.$$

Proof. For each $I \in \mathbb{N}$, let $P \in \mathcal{P}_I$ satisfy

$$\sup_{P \in \mathcal{P}_I} \sup_{t \in \mathbb{R}} |\mathbb{P}_P(S_I \leq t) - \Phi(t)| \leq \sup_{t \in \mathbb{R}} |\mathbb{P}_P(S_I \leq t) - \Phi(t)| + I^{-1}. \quad (36)$$

Let $T_{I,i} := \left(\sum_{i=1}^I \mathbb{E}_P \left[A_{I,i}^2 \right] \right)^{-\frac{1}{2}} (A_{I,i})$ for each $i \in [I]$. Note that by construction $T_{I,1}, \dots, T_{I,I}$ are independent and mean zero and $\sum_{i=1}^I \mathbb{E}_P(T_{I,i}^2) = 1$. By the Lindeberg–Feller central limit theorem for triangular arrays (Proposition 2.27, [van der Vaart \[1998\]](#)) we have

$$\limsup_{I \rightarrow \infty} \sup_{t \in \mathbb{R}} |\mathbb{P}_P(S_I \leq t) - \Phi(t)| = 0,$$

provided the condition for the Lindeberg–Feller theorem holds. This condition holds as, for all $\epsilon > 0$,

$$\begin{aligned} \sum_{i=1}^I \mathbb{E}_P \left[|T_{I,i}|^2 \mathbb{1}_{\{|T_{I,i}| > \epsilon\}} \right] &\leq \sum_{i=1}^I \left(\mathbb{E}_P \left[|T_{I,i}|^{2+\Delta} \right] \right)^{\frac{2}{2+\Delta}} \left(\mathbb{P}_P(|T_{I,i}| > \epsilon) \right)^{\frac{\Delta}{2+\Delta}} \\ &\leq \epsilon^{-\Delta} \sum_{i=1}^I \mathbb{E}_P \left[|T_{I,i}|^{2+\Delta} \right] \\ &= \epsilon^{-\Delta} \cdot \frac{\sum_{i=1}^I \mathbb{E}_P \left[|A_{I,i}|^{2+\Delta} \right]}{\left| \sum_{i=1}^I \mathbb{E}_P[A_{I,i}^2] \right|^{1+\frac{\Delta}{2}}} \end{aligned}$$

by Hölder’s inequality followed by Markov’s inequality. Taking limits of (36) therefore completes the proof. \square

Lemma 18 ([Shah and Peters \[2020\]](#), Lemma 20). *Let $(A_I)_{I \in \mathbb{N}}$ and $(B_I)_{I \in \mathbb{N}}$ be sequences of real-valued random variables. Suppose*

$$\limsup_{I \rightarrow \infty} \sup_{P \in \mathcal{P}_I} \sup_{t \in \mathbb{R}} |\mathbb{P}_P(A_I \leq t) - \Phi(t)| = 0.$$

Then we have:

(a)

$$\text{If } B_I = o_P(1) \text{ then } \limsup_{I \rightarrow \infty} \sup_{P \in \mathcal{P}_I} \sup_{t \in \mathbb{R}} |\mathbb{P}_P(A_I + B_I \leq t) - \Phi(t)| = 0;$$

(b)

$$\text{If } B_I = 1 + o_P(1) \text{ then } \limsup_{I \rightarrow \infty} \sup_{P \in \mathcal{P}_I} \sup_{t \in \mathbb{R}} |\mathbb{P}_P(A_I/B_I \leq t) - \Phi(t)| = 0.$$

D.4 Proof of Corollary 6

Proof of Theorem 6. We aim to show that Assumption [A3.1’](#) implies [A3.1](#) and [A3.2’](#) implies [A3.2](#). Note that for all $k \in [K]$, $i \in \mathcal{I}_k$, $j \in [n_i]$,

$$\mathbb{E}_P \left[\left(\hat{\sigma}^{(k)}(X_{ij}) - \sigma^*(X_{ij}) \right)^2 \right] = \mathbb{E}_P \left[\frac{\left(\hat{s}^{(k)}(X_{ij}) - s^*(X_{ij}) \right)^2}{\hat{s}^{(k)}(X_{ij})^2 s^*(X_{ij})^2} \right] \lesssim \mathbb{E}_P \left[\left(\hat{s}^{(k)}(X_{ij}) - s^*(X_{ij}) \right)^2 \right],$$

and so $\mathcal{R}_\sigma \lesssim \mathcal{R}_{s^*}$, hence Assumption [A3.1’](#) implies [A3.1](#).

To show that Assumption [A3.2’](#) implies [A3.2](#), we consider each case separately:

Equicorrelation Case: Note that for all $k \in [K]$,

$$\left| \hat{\rho}^{(k)} - \rho^* \right| = \frac{\left| \hat{\theta}^{(k)} - \theta^* \right|}{|1 + \theta^*| \cdot |1 + \hat{\theta}^{(k)}|} \lesssim \left| \hat{\theta}^{(k)} - \theta^* \right|,$$

and therefore $\mathcal{R}_\rho \lesssim \mathcal{R}_{\theta^*}$.

Hierarchical / Nested Case: Note that for all $k \in [K]$,

$$\begin{aligned} \left| \hat{\rho}_1^{(k)} - \rho_1^* \right| &= \frac{\left| \hat{\theta}_1^{(k)} + \hat{\theta}_2^{(k)} - \theta_1^* - \theta_2^* \right|}{\left| 1 + \hat{\theta}_1^{(k)} + \hat{\theta}_2^{(k)} \right| \cdot \left| 1 + \theta_1^* + \theta_2^* \right|} \lesssim \left| \hat{\theta}_1^{(k)} - \theta_1^* \right| + \left| \hat{\theta}_2^{(k)} - \theta_2^* \right| \leq \sqrt{2} \left\| \hat{\theta}^{(k)} - \theta^* \right\|_2, \\ \left| \hat{\rho}_2^{(k)} - \rho_2^* \right| &= \frac{\left| (1 + \theta_1^*) (\hat{\theta}_2 - \theta_2^*) - \theta_2^* (\hat{\theta}_1^{(k)} - \theta_1^*) \right|}{\left| 1 + \hat{\theta}_1^{(k)} + \hat{\theta}_2^{(k)} \right| \cdot \left| 1 + \theta_1^* + \theta_2^* \right|} \lesssim \sqrt{2} \left\| \hat{\theta}^{(k)} - \theta^* \right\|_2, \end{aligned}$$

and therefore $\mathcal{R}_\rho \lesssim \mathcal{R}_{\theta^*}$.

Longitudinal AR(1) Case: Note that for all $k \in [K]$, $i \in \mathcal{I}_k$, $j \in [n_i]$, $j' \in [n_i] \setminus \{j\}$,

$$\begin{aligned} &\left| \hat{\rho}^{(k)}(X_{ij}, X_{ij'}) - \rho^*(X_{ij}, X_{ij'}) \right| = \left| \hat{\theta}^{(k)|j-j'|} - \theta^{*|j-j'|} \right| \\ &\leq \left| \hat{\theta}^{(k)} - \theta^* \right| \cdot \left| \sum_{l=1}^{|j-j'|-1} \hat{\theta}^{(k)l} \theta^{*|j-j'|-1-l} \right| \leq \left| \hat{\theta}^{(k)} - \theta^* \right| \cdot |j - j'| \cdot \left(\hat{\theta}^{(k)} \vee \theta^* \right)^{|j-j'|-1} \\ &\leq \left| \hat{\theta}^{(k)} - \theta^* \right| \cdot |j - j'| (1 - c)^{|j-j'|-1} \leq \left(\log \left(\frac{1}{1-c} \right) \right)^{-1} (1 - c)^{\frac{1}{\log \left(\frac{1}{1-c} \right)} - 1} \left| \hat{\theta}^{(k)} - \theta^* \right| \\ &\lesssim \left| \hat{\theta}^{(k)} - \theta^* \right|, \end{aligned}$$

for some $c > 0$, with the last inequality following by elementary calculus (optimising the function $N \mapsto N(1-c)^{N-1}$). Therefore $\mathcal{R}_\rho \lesssim \mathcal{R}_{\theta^*}$.

Hence, for each working correlation, Assumption [A3.2'](#) implies [A3.2](#).

Finally, we must show that for each working correlation $\Lambda_{\max}(W_i^*) \lesssim 1$, $\Lambda_{\min}(W_i^*) \gtrsim n^{-\gamma}$ for all $i \in [I]$. We consider each separately.

Equicorrelation Case: Consider the decomposition of the matrix W_i^* as the sum of two matrices $W_i = W_i^{*(A)} + W_i^{*(B)}$

$$\begin{aligned} (W_i^{*(A)})_{jj'} &= \delta_{jj'} s^*(X_{ij}) s^*(X_{ij'}) \\ (W_i^{*(B)})_{jj'} &= -\frac{\theta}{1 + \theta n_i} s^*(X_{ij}) s^*(X_{ij'}), \end{aligned}$$

for $j, j' \in [n_i]$. Then, we have that

$$\Lambda_{\max}(W_i^*) \leq \max_{j \in [n_i]} s^*(X_{ij})^2 + 0 \lesssim 1,$$

and further that, by Weyl's inequality,

$$\begin{aligned} \Lambda_{\min}(W_i^*) &\geq \min_{j \in [n_i]} s^*(X_{ij})^2 - \frac{\theta^*}{1 + \theta^* n_i} \|s^*(X_i)\|_2^2 \\ &\geq \left(1 - \frac{\theta^* n_i}{1 + \theta^* n_i} \right) \min_{j \in [n_i]} s^*(X_{ij})^2 \gtrsim \frac{1}{1 + \theta^* n_i} \geq \frac{1}{1 + \theta^*} \cdot \frac{1}{n_i} \gtrsim n_i^{-1}. \end{aligned}$$

Hierarchical / Nested Case: First consider the decomposition of $W_i^* = W_i^{*(A)} + W_i^{*(B)}$ where

$$\begin{aligned}(W_i^{*(A)})_{jj'} &= (1 + \theta_1^* + \theta_2^*)\delta_{jj'}s^*(X_{ij})s^*(X_{ij'}), \\ (W_i^{*(B)})_{jj'} &= -\mathcal{B}_{i,jj'}s^*(X_{ij})s^*(X_{ij'}),\end{aligned}$$

for $j, j' \in [n_i]$ and \mathcal{B} as defined in Section 3.2. Then

$$\Lambda_{\max}(W_i^*) \leq \max_{j \in [n_i]} s^*(X_{ij})^2 + 0 \lesssim 1$$

Further, by considering W_i^{*-1} parametrised by $(\rho_1^*, \rho_2^*) = \left(\frac{\theta_1^* + \theta_2^*}{1 + \theta_1^* + \theta_2^*}, \frac{\theta_2^*}{1 + \theta_1^* + \theta_2^*}\right)$, we can similarly show that $\Lambda_{\max}(W_i^{*-1}) \lesssim n_i$, and therefore $\Lambda_{\min}(W_i^*) \gtrsim n_i^{-1}$, completing the proof.

Longitudinal AR(1) Case: By the Gershgorin circle theorem alongside the triangle inequality,

$$\begin{aligned}\Lambda_{\max}(W_i^*) &\leq (1 + 2\theta^*) \max_{j \in [n_i]} s^*(X_{ij})^2 \lesssim 1, \\ \Lambda_{\max}(W_i^{*-1}) &\leq \left(\min_{j \in [n_i]} s^*(X_{ij})^2\right)^{-1} + 2 \left(\min_{j \in [n_i]} s^*(X_{ij})^2\right)^{-1} \sum_{j=1}^{n_i-1} \theta^j \\ &= \left(1 + \frac{2\theta^*(1 - \theta^{*n_i-1})}{1 - \theta^*}\right) \left(\min_{j \in [n_i]} s^*(X_{ij})^2\right)^{-1} \\ &\leq \left(1 + \frac{2}{1 - \theta^*}\right) \left(\min_{j \in [n_i]} s^*(X_{ij})^2\right)^{-1} \lesssim 1,\end{aligned}$$

and so $\Lambda_{\min}(W_i^*) = (\Lambda_{\max}(W_i^{*-1}))^{-1} \gtrsim 1$ as required. \square

E Practical details of numerical results in Section 5

E.1 Details of weight estimators studied in 5.1

The ‘heteroscedastic ML’ estimator used in Section 5 follows Algorithm 1 but with weights constructed as follows. Adopting the notation in Algorithm 1 with weights as in (13), we construct

$$\begin{aligned}(\hat{\phi}_{\text{ML}}, \hat{\rho}_{\text{ML}}, \hat{\beta}_{\text{ML}}) &= \underset{(\phi, \rho, \beta) \in \mathbb{R}^5 \times (-1, 1) \times \mathbb{R}}{\operatorname{argmin}} l_{\text{ML}}(\phi, \rho, \beta), \\ l_{\text{ML}}(\phi, \rho, \beta) &:= \sum_{i \in \mathcal{I}_k^c} \left(-\log \det W(\phi, \rho)(X_i) + (\tilde{R}_i^Y - \beta \tilde{R}_i^D)^\top W(\phi, \rho)(X_i) (\tilde{R}_i^Y - \beta \tilde{R}_i^D) \right), \\ \{W(\phi, \rho)(X_i)\}^{-1} &= D_{\sigma_\phi}(X_i) C_{\theta_\rho} D_{\sigma_\phi}(X_i), \quad \sigma_\phi(x) = \exp(\phi^\top(1, x, x^2, x^3, x^4)),\end{aligned}$$

with working correlations $(C_{\theta_\rho})_{jk} = \mathbb{1}_{\{j=k\}} + \rho \cdot \mathbb{1}_{\{j \neq k\}}$ for examples 5.1.1 and 5.1.4 (equicorrelated) and $(C_{\theta_\rho})_{jk} = \rho^{|j-k|}$ for examples 5.1.2 and 5.1.3 (AR(1)). These are estimated using the nlme package [Pinheiro et al., 2022]. Weights are then constructed as

$$\{\hat{W}_{\text{ML}}^{(k)}(\cdot)\}^{-1} = D_{\hat{\sigma}_{\text{ML}}}(\cdot) C_{\hat{\theta}_{\rho_{\text{ML}}}} D_{\hat{\sigma}_{\text{ML}}}(\cdot), \quad \hat{\sigma}_{\text{ML}}(x) := \exp(\hat{\phi}_{\text{ML}}^\top(1, x, x^2, x^3, x^4))$$

and used in Algorithm 1 to construct the final $\hat{\beta}$. The ‘homoscedastic ML’ estimator is constructed as above fixing $\hat{\phi}_{\text{ML}} = \mathbf{0}$.

The ‘heteroscedastic GEE’ estimator also follow the setup of Algorithm 1 with weights constructed as follows. First we generate a generalised additive cubic spline estimator for the variance and estimator for the correlation parameter as

$$\hat{\sigma}_{\text{GEE}} := \operatorname{argmin}_{\sigma \in \Pi} \sum_{i \in \mathcal{I}_k^c} \sum_{j=1}^{n_i} (\tilde{\varepsilon}_{ij}^2 - \sigma^2(X_{ij}))^2 + \lambda J(\sigma^2), \quad J(f) = \int (f''(x))^2 dx,$$

$$\hat{\rho}_{\text{GEE}} := \operatorname{argmin}_{\rho \in (-1,1)} \sum_{i \in \mathcal{I}_k^c} \sum_{\substack{j,k=1 \\ j \neq k}}^{n_i} \left(\tilde{\varepsilon}_{ij} \tilde{\varepsilon}_{ik} - \hat{\sigma}_{\text{GEE}}(X_{ij}) \hat{\sigma}_{\text{GEE}}(X_{ik}) (C_{\theta_\rho})_{jk} \right)^2,$$

where Π is the set of $\mathbb{R}_{>0}$ valued functions that have absolutely continuous first derivative, λ is a tuning parameter that minimises the generalised cross-validation (GCV) score, and C_{θ_ρ} takes the equicorrelated or AR(1) structure for each example as outlined above. In all the simulations in Section 5.1 we fit $\hat{\sigma}_{\text{GEE}}$ using the `gam` package [Wood, 2017], and $\hat{\rho}_{\text{GEE}}$ using the `geepack` package [Halekoh et al., 2006]. Weights are then constructed as

$$\{\hat{W}_{\text{GEE}}^{(k)}(\cdot)\}^{-1} = D_{\hat{\sigma}_{\text{GEE}}}(\cdot) C_{\theta_{\hat{\rho}_{\text{GEE}}}} D_{\hat{\sigma}_{\text{GEE}}}(\cdot),$$

and used in Algorithm 1 to construct $\hat{\beta}$. The ‘homoscedastic GEE’ estimator is estimated in the same way except we set $\hat{\sigma}_{\text{GEE}} \equiv 1$.

E.2 Sandwich boosting with variable step size

In Section 3.2 we outlined the sandwich boosting algorithm 2 in terms of a fixed step size $\lambda^{(s)}$, with an optimal number of boosting iterations m_{stop} selected by cross-validation. In this setup $\lambda^{(s)}$ acts as a hyperparameter; if it is set to be too large one can be susceptible to under-fitting (being unable to attain a local minimum), and if set too small it can take an unreasonably large m_{stop} to achieve convergence. An adaptation of Algorithm 2 that may be helpful in practice would be to replace the $\lambda^{(s)}$ constant step size with a variable step size

$$\lambda_m^{(s)} \approx \operatorname{argmin}_{\lambda \in \Lambda} \hat{L}_{\text{SL}}(\hat{s}_m - \lambda \hat{u}_m, \hat{\theta}_m),$$

over some closed interval Λ of positive step sizes, attained by minimising the second order expansion of $\lambda \mapsto \hat{L}_{\text{SL}}(\hat{s}_m - \lambda \hat{u}_m, \hat{\theta}_m)$. The m th iterative update of the s -function is then taken as

$$\hat{s}_{m+1} = (\hat{s}_m - \mu^{(s)} \lambda_m^{(s)} \hat{u}_m) \vee \epsilon,$$

for some constant shrinkage parameter $\mu^{(s)} \in (0, 1)$, whose purpose is to avoid overfitting, and small constant $\epsilon > 0$ which ensures that \hat{s}_{m+1} is always strictly positive everywhere; in all our experiments, we set $\epsilon = 0.1$. With this setup boosting, can then be performed until a convergence criterion is achieved. This can eliminate the need to determine an approximately optimal $(\lambda^{(s)}, m_{\text{stop}})$, as well as achieve faster convergence of sandwich boosting. We implement this variable step size scheme in some of the numerical examples studied in Section 5.

E.3 Details of numerical results in Section 5

Sandwich boosting for the numerical results is setup as follows.

Increasing Model Complexity (Section 5.1.1): The variable step size adaptation of sandwich boosting is used, with $\Lambda = [0.001, 10]$ and $\lambda^{(\theta)} = 0.1$.

Increasing Covariance Misspecification (Section 5.1.2): A constant step size $\lambda^{(s)} = 10$ is used, with $\lambda^{(\theta)} = 0.1$ and $m_{\text{stop}} = 200$ selected by cross-validation.

Correlation Misspecification (Section 5.1.3): As this simulated example involves optimisation of the sandwich loss over a parametric weight class in terms of a single AR(1) parameter, there is no need to select a $(\lambda^{(s)}, m_{\text{stop}})$.

Conditional Variance Misspecification (Section 5.1.4): A constant step size $\lambda^{(s)} = 100$ is used, with $\lambda^{(\theta)} = 0.1$ and $m_{\text{stop}} = 500$ selected by cross-validation.

Real world data scenarios (Section 5.2): We carry out sandwich boosting for both with the variable step size adaptation within the interval $\Lambda = [0.01, 1]$, with s -shrinkage parameter $\mu^{(s)} = 0.1$ and θ -step size $\lambda^{(\theta)} = 0.1$.

When analysing the real world data discussed in Section 5.2, we alleviate the dependence of the cross-fitting method on the randomness of the sample splits by repeating the procedure of Algorithm 1 $\mathcal{S} = 50$ times for different random sample splits. Suppose for the repeat indexed by $s \in [\mathcal{S}]$ we have β -estimator $\hat{\beta}_s$ and asymptotic variance estimator \hat{V}_s . Then we aggregate the results of these repeats to produce a final estimate of β and associated estimate of its variance via

$$\hat{\beta} := \text{median}_{s \in [\mathcal{S}]} \hat{\beta}_s, \quad \hat{V} := \text{median}_{s \in [\mathcal{S}]} (\hat{V}_s + (\hat{\beta} - \hat{\beta}_s)^2);$$

see for example Chernozhukov et al. [2018], Emmenegger and Bühlmann [2023].

E.4 Choice of parameters for cross-fitting

As outlined in Section 5, for the simulations in Section 5.1 we perform cross-fitting with $K = 2$ folds and for $\mathcal{S} = 1$ random sample splits. The choice of (K, \mathcal{S}) does not affect the asymptotic mean squared error of the resulting estimator (see Theorem 4), and thus for sufficiently large sample size the choice of cross-fitting parameters does not affect the resulting mean squared error. Figure 6 (left panel) shows a simulation in Section 5.1 (specifically Example 5.1.4 with group size $n_i \equiv 8$) for cross-fitting parameters $K \in \{2, 5, 10\}$. We see the choice of K does not have a significant effect on the mean squared error of $\hat{\beta}$ at the sample sizes considered in Section 5.1, therefore justifying our choice of $(K, \mathcal{S}) = (2, 1)$ (although any other choices would be valid, at a further computational cost). We do however note that, while such a simple, computationally lean choice is justified for the simulations in Section 5.1, this does not necessarily hold for smaller sample sizes; see for example Figure 6 (right). In such settings a larger choice of either (or indeed both) K and \mathcal{S} can improve the mean squared error, by reducing the variance in the nuisance function estimators, in addition to the variance induced by the sample split(s) used for cross-fitting. Therefore, when analysing the real data scenarios of Section 5.2 we choose a larger choice of $(K, \mathcal{S}) = (5, 50)$ (see Section E.3 for details).

F Additional numerical results

Tables 5–8 present the coverage probabilities of nominal 95% confidence intervals based on the construction in Algorithm 1 for the simulation settings in Section 5.1.

G Further details on the generalised sandwich loss of Section 6

Consider the setup of Section 6 where we wish to estimate function $\beta(X)$ in model (18), a task for which by our assumption (19), it suffices to estimate coefficient vector ϕ . Below we

λ	Sandwich Boosting	Het ML	Hom ML	Het GEE	Hom GEE	Unweighted	Oracle
0.1	0.952	0.954	0.950	0.950	0.950	0.946	0.950
0.3	0.956	0.950	0.956	0.950	0.956	0.946	0.950
0.5	0.944	0.952	0.946	0.952	0.946	0.952	0.950
0.7	0.950	0.950	0.948	0.948	0.948	0.956	0.948
0.9	0.950	0.958	0.948	0.950	0.948	0.954	0.948
1.1	0.956	0.952	0.950	0.958	0.950	0.956	0.958
1.3	0.952	0.946	0.952	0.962	0.952	0.950	0.954
1.5	0.958	0.942	0.948	0.958	0.948	0.946	0.962
1.7	0.962	0.946	0.946	0.970	0.946	0.948	0.962
1.9	0.958	0.942	0.944	0.970	0.944	0.944	0.966
2.1	0.962	0.938	0.942	0.954	0.942	0.948	0.954
2.3	0.954	0.936	0.946	0.964	0.946	0.950	0.956
2.5	0.946	0.942	0.952	0.948	0.952	0.946	0.952

Table 5: Coverage of nominal 95% confidence intervals $\hat{C}(0.05)$ for β in Example 5.1.1.

η	Sandwich Boosting	Het ML	Hom ML	Het GEE	Hom GEE	Unweighted
10	0.982	0.974	0.974	0.974	0.974	0.958
20	0.976	0.978	0.984	0.952	0.984	0.952
30	0.982	0.962	0.970	0.948	0.972	0.946
40	0.984	0.960	0.976	0.930	0.974	0.944
50	0.972	0.946	0.968	0.910	0.966	0.940
60	0.972	0.942	0.958	0.910	0.960	0.932
70	0.972	0.932	0.958	0.892	0.958	0.934
80	0.970	0.934	0.962	0.906	0.960	0.952
90	0.970	0.908	0.964	0.884	0.964	0.956
100	0.968	0.914	0.960	0.868	0.958	0.932

Table 6: Coverage of nominal 95% confidence intervals $\hat{C}(0.05)$ for β in Example 5.1.2.

Group Size	Hom SL	Hom ML	Hom GEE	Unweighted
2^1	0.948	0.948	0.948	0.946
2^2	0.946	0.946	0.948	0.944
2^3	0.950	0.946	0.936	0.942
2^4	0.942	0.960	0.934	0.958
2^5	0.952	0.956	0.936	0.948
2^6	0.950	0.946	0.936	0.950
2^7	0.944	0.942	0.932	0.946
2^8	0.942	0.944	0.938	0.938

Table 7: Coverage of nominal 95% confidence intervals $\hat{C}(0.05)$ for β in Example 5.1.3.




Figure 6: Mean squared error of the sandwich boosted β -estimators for different cross-fitting parameter choices (the number of folds used for cross-fitting, $K \in \{2, 5, 10\}$, and the number of repeats over different sample splits, $\mathcal{S} \in \{1, 5\}$). We present these for two simulations: (Left) Example 5.1.4, with $I = 2^{12}$ groups each of size $n_i \equiv 8$; and (Right) The same data generating mechanism as in Example 5.1.4 but with $I = 500$ groups each of size $n_i \equiv 4$. For all settings nominal 95% confidence intervals for the mean squared errors (over 500 simulations) are plotted.

Group Size	Sandwich Boosting	Het ML	Hom ML	Het GEE	Hom GEE	Unweighted
2^1	0.938	0.936	0.932	0.948	0.932	0.926
2^2	0.938	0.938	0.940	0.954	0.940	0.922
2^3	0.938	0.936	0.938	0.956	0.938	0.926
2^4	0.938	0.940	0.944	0.954	0.944	0.924
2^5	0.938	0.940	0.936	0.966	0.936	0.934
2^6	0.942	0.942	0.948	0.968	0.948	0.928
2^7	0.946	0.948	0.950	0.966	0.950	0.916
2^8	0.952	0.952	0.952	0.968	0.952	0.926

Table 8: Coverage of nominal 95% confidence intervals $\hat{C}(0.05)$ for β in Example 5.1.4.

briefly outline how this may be performed. For intuition, suppose the regression functions (l_0, m_0) (defined as in the rest of the paper) are known, and define $R_i^Y := Y_i - l_0(X_i)$ and $R_i^D := D_i - m_0(X_i)$. Then

$$R_i^Y = M_i \phi + \varepsilon_i, \quad M_i := (\varphi_l(X_{ij}) R_{ij}^D)_{j \in [n], l \in [L]}.$$

An estimate of ϕ can then be obtained via a potentially weighted multivariate linear regression of the R_i^Y onto the M_i . In reality, estimates (\hat{l}, \hat{m}) would be used in place of the unknown (l_0, m_0) in a cross-fitting scheme similar to that in Section 3.1. We may obtain an estimate $\hat{V}_\phi(W)$ as a function of the weight function, and then seek to minimise the generalised sandwich loss objective (21) via sandwich boosting.

Algorithm 5 outlines how to calculate the s -scores for sandwich boosting in this setup for an arbitrary inverse working correlation $\mathcal{C}_{ijk}(\theta)$ as in (26). The order of computation of this score calculation can be reduced to $O(L^2 N)$ when using one of the working correlation parameterisations in Section C; see Algorithm 6 for the example of the equicorrelated working correlation.

Algorithm 5: Generalised sandwich boosting: Score calculation for arbitrary working correlation

Input: Index set \mathcal{I} for boosting training; Working correlation structure \mathcal{C}_i (26); Estimates of (grouped) errors $(\hat{\xi}_i, \hat{\varepsilon}_i, X_i)_{i \in \mathcal{I}}$; Estimates of (s, θ) at which scores to be calculated; $\Phi \in \mathbb{R}^{L \times L}$ matrix (see (20)).

Calculate $A^{(1)} \in \mathbb{R}^{|\mathcal{I}| \times L \times L}$ and $A^{(2)} \in \mathbb{R}^{|\mathcal{I}| \times L \times L}$:

for $i \in \mathcal{I}$ **do**

for $(l, l') \in [L]^2$ **do**

$$\begin{aligned} & \left| \begin{aligned} A_{ill'}^{(1)} &:= \sum_{k=1}^{n_i} \sum_{k'=1}^{n_i} \mathcal{C}_{ikk'}(\theta) s(X_{ik}) s(X_{ik'}) \hat{\xi}_{ik} \hat{\xi}_{ik'} \varphi_l(X_{ik}) \varphi_{l'}(X_{ik'}), \\ A_{ill'}^{(2)} &:= \sum_{k=1}^{n_i} \sum_{k'=1}^{n_i} \mathcal{C}_{ikk'}(\theta) s(X_{ik}) s(X_{ik'}) \hat{\xi}_{ik} \hat{\varepsilon}_{ik'} \varphi_l(X_{ik}) \varphi_{l'}(X_{ik'}). \end{aligned} \right. \end{aligned}$$

end

end

Calculate $A_{\text{sum}}^{(1)} := \sum_{i \in \mathcal{I}} A_i^{(1)}$ and $A_{\text{sum.sq}}^{(2)} := \sum_{i \in \mathcal{I}} A_i^{(2)} A_i^{(2)T}$.

Calculate $A^{(3)}$ and $A^{(4)}$:

for $i \in \mathcal{I}$ **and** $j \in [n_i]$ **do**

for $(l, l') \in [L]^2$ **do**

$$\begin{aligned} & \left| \begin{aligned} A_{ijll'}^{(3)} &:= \sum_{k=1}^{n_i} \mathcal{C}_{ijk}(\theta) s(X_{ik}) \hat{\xi}_{ij} \hat{\xi}_{ik} \left(\varphi_l(X_{ij}) \varphi_{l'}(X_{ik}) + \varphi_{l'}(X_{ij}) \varphi_l(X_{ik}) \right), \\ A_{ijll'}^{(4)} &:= \sum_{k=1}^{n_i} \mathcal{C}_{ijk}(\theta) s(X_{ik}) \left(\hat{\xi}_{ij} \hat{\varepsilon}_{ik} \varphi_l(X_{ij}) \varphi_{l'}(X_{ik}) + \hat{\xi}_{ik} \hat{\varepsilon}_{ij} \varphi_{l'}(X_{ij}) \varphi_l(X_{ik}) \right). \end{aligned} \right. \end{aligned}$$

end

end

Calculate s -scores:

for $i \in \mathcal{I}$ **and** $j \in [n_i]$ **do**

$$\begin{aligned} & \left| \begin{aligned} U_{\text{SL}}^{(s)}(s, \theta)(X_{ij}) &:= -\text{tr} \left(\Phi \cdot \left(A_{\text{sum}}^{(1)} \right)^{-1} \left[A_{ij}^{(3)} \left(A_{\text{sum}}^{(1)} \right)^{-1} A_{\text{sum.sq}}^{(2)} \right. \right. \\ & \quad \left. \left. + A_{\text{sum.sq}}^{(2)} \left(A_{\text{sum}}^{(1)} \right)^{-1} A_{ij}^{(3)} - A_{ij}^{(4)} A_i^{(2)} - A_i^{(2)} A_{ij}^{(4)} \right] \left(A_{\text{sum}}^{(1)} \right)^{-1} \right). \end{aligned} \right. \end{aligned}$$

end

Calculate $A_{\text{sum}}^{(5)}$ and $A^{(6)}$:

for $(l, l') \in [L]^2$ **do**

$$\begin{aligned} & \left| \begin{aligned} \left(A_{\text{sum}}^{(5)} \right)_{ll'} &= \sum_{i \in \mathcal{I}} \sum_{k=1}^{n_i} \sum_{k'=1}^{n_i} \frac{\partial \mathcal{C}_{ikk'}}{\partial \theta} s(X_{ik}) s(X_{ik'}) \hat{\xi}_{ik} \hat{\xi}_{ik'} \varphi_l(X_{ik}) \varphi_{l'}(X_{ik'}), \\ \text{for } i \in \mathcal{I} \text{ do} \\ & \left| \begin{aligned} A_{ill'}^{(6)} &:= \sum_{k=1}^{n_i} \sum_{k'=1}^{n_i} \frac{\partial \mathcal{C}_{ikk'}}{\partial \theta} s(X_{ik}) s(X_{ik'}) \hat{\xi}_{ik} \hat{\varepsilon}_{ik'} \varphi_l(X_{ik}) \varphi_{l'}(X_{ik'}). \end{aligned} \right. \\ \text{end} \end{aligned} \right. \end{aligned}$$

end

Calculate $A_{\text{sum.sq.diff}}^{(6)} := \sum_{i \in \mathcal{I}} A_i^{(2)} A_i^{(6)} + A_i^{(6)} A_i^{(2)}$.

Calculate the θ -score:

$$\begin{aligned} U_{\text{SL}}^{(\theta)}(s, \theta) &= -\text{tr} \left(\Phi \cdot \left(A_{\text{sum}}^{(1)} \right)^{-1} \left[A_{\text{sum}}^{(5)} \left(A_{\text{sum}}^{(1)} \right)^{-1} A_{\text{sum.sq}}^{(2)} + A_{\text{sum.sq}}^{(2)} \left(A_{\text{sum}}^{(1)} \right)^{-1} A_{\text{sum}}^{(5)} \right. \right. \\ & \quad \left. \left. - A_{\text{sum.sq.diff}}^{(6)} \right] \left(A_{\text{sum}}^{(1)} \right)^{-1} \right). \end{aligned}$$

Output: The set of N s -scores $\left(X_{ij}, U_{\text{SL}}^{(s)}(s, \theta)(X_{ij}) \right)_{i \in \mathcal{I}, j \in [n_i]}$ and θ -score $U_{\text{SL}}^{(\theta)}(s, \theta)$.

Algorithm 6: Generalised sandwich boosting: Score calculation for equicorrelated working correlation

Input: Index set \mathcal{I} for boosting training; Estimates of (grouped) errors $(\hat{\xi}_i, \hat{\varepsilon}_i, X_i)_{i \in \mathcal{I}}$; Estimates of (s, θ) at which scores to be calculated; $\Phi \in \mathbb{R}^{L \times L}$ matrix (see (20)).

Calculate $A^{(1)} \in \mathbb{R}^{|\mathcal{I}| \times L \times L}$, $A^{(2)} \in \mathbb{R}^{|\mathcal{I}| \times L \times L}$, $B^{(1)} \in \mathbb{R}^{|\mathcal{I}| \times L}$ and $B^{(2)} \in \mathbb{R}^{|\mathcal{I}| \times L}$:

for $i \in \mathcal{I}$ **do**

for $l \in [L]$ **do**

$$B_{il}^{(1)} := \sum_{k=1}^{n_i} s(X_{ik}) \hat{\xi}_{ik} \varphi_l(X_{ik}),$$

$$B_{il}^{(2)} := \sum_{k=1}^{n_i} s(X_{ik}) \hat{\varepsilon}_{ik} \varphi_l(X_{ik}),$$

end

for $(l, l') \in [L]^2$ **do**

$$A_{ill'}^{(1)} := \sum_{k=1}^{n_i} s^2(X_{ik}) \hat{\xi}_{ik}^2 \varphi_l(X_{ik}) \varphi_{l'}(X_{ik}) - \frac{\theta}{1+\theta n_i} B_{il}^{(1)} B_{il'}^{(1)},$$

$$A_{ill'}^{(2)} := \sum_{k=1}^{n_i} s^2(X_{ik}) \hat{\xi}_{ik} \hat{\varepsilon}_{ik} \varphi_l(X_{ik}) \varphi_{l'}(X_{ik}) - \frac{\theta}{1+\theta n_i} B_{il}^{(1)} B_{il'}^{(2)}.$$

end

end

Calculate $A_{\text{sum}}^{(1)} := \sum_{i \in \mathcal{I}} A_i^{(1)}$ and $A_{\text{sum.sq}}^{(2)} := \sum_{i \in \mathcal{I}} A_i^{(2)} A_i^{(2)T}$.

Calculate $A^{(3)}$ and $A^{(4)}$:

for $i \in \mathcal{I}$ **and** $j \in [n_i]$ **do**

for $(l, l') \in [L]^2$ **do**

$$A_{ijll'}^{(3)} := 2s(X_{ij}) \hat{\xi}_{ij}^2 \varphi_l(X_{ij}) \varphi_{l'}(X_{ij}) - \frac{\theta}{1+\theta n_i} \hat{\xi}_{ij} \left(B_{il}^{(1)} \varphi_{l'}(X_{ij}) + B_{il'}^{(1)} \varphi_l(X_{ij}) \right),$$

$$A_{ijll'}^{(4)} := 2s(X_{ij}) \hat{\xi}_{ij} \hat{\varepsilon}_{ij} \varphi_l(X_{ij}) \varphi_{l'}(X_{ij}) - \frac{\theta}{1+\theta n_i} \left(B_{il}^{(1)} \hat{\varepsilon}_{ij} \varphi_{l'}(X_{ij}) + B_{il'}^{(2)} \hat{\xi}_{ij} \varphi_l(X_{ij}) \right).$$

end

end

Calculate s -scores:

for $i \in \mathcal{I}$ **and** $j \in [n_i]$ **do**

$$U_{\text{SL}}^{(s)}(s, \theta)(X_{ij}) := -\text{tr} \left(\Phi \cdot \left(A_{\text{sum}}^{(1)} \right)^{-1} \left[A_{ij}^{(3)} \left(A_{\text{sum}}^{(1)} \right)^{-1} A_{\text{sum.sq}}^{(2)} + A_{\text{sum.sq}}^{(2)} \left(A_{\text{sum}}^{(1)} \right)^{-1} A_{ij}^{(3)} - A_{ij}^{(4)} A_i^{(2)} - A_i^{(2)} A_{ij}^{(4)} \right] \left(A_{\text{sum}}^{(1)} \right)^{-1} \right).$$

end

Calculate the θ -score:

$$U_{\text{SL}}^{(\theta)}(s, \theta) = \text{tr} \left(\Phi \cdot \left(A_{\text{sum}}^{(1)} \right)^{-1} \left[\left(\sum_{i \in \mathcal{I}} \frac{B_i^{(1)} B_i^{(1)T}}{(1+\theta n_i)^2} \right) \left(A_{\text{sum}}^{(1)} \right)^{-1} A_{\text{sum.sq}}^{(2)} + A_{\text{sum.sq}}^{(2)} \left(A_{\text{sum}}^{(1)} \right)^{-1} \left(\sum_{i \in \mathcal{I}} \frac{B_i^{(1)} B_i^{(1)T}}{(1+\theta n_i)^2} \right) - \sum_{i \in \mathcal{I}} \frac{A_i^{(2)} B_i^{(1)} B_i^{(2)T} + B_i^{(2)} B_i^{(1)T} A_i^{(2)}}{(1+\theta n_i)^2} \right] \left(A_{\text{sum}}^{(1)} \right)^{-1} \right).$$

Output: The set of N s -scores $(X_{ij}, U_{\text{SL}}^{(s)}(s, \theta)(X_{ij}))_{i \in \mathcal{I}, j \in [n_i]}$ and θ -score $U_{\text{SL}}^{(\theta)}(s, \theta)$.

References

- V. Chernozhukov, D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, and J. Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 2018. doi: 10.1111/ectj.12097.
- C. Emmenegger and P. Bühlmann. Plug-in machine learning for partially linear mixed-effects

- models with repeated measurements. *Scandinavian Journal of Statistics*, 2023. doi: 10.1111/sjos.12639.
- U. Halekoh, S. Højsgaard, and J. Yan. The r package geePack for generalized estimating equations. *Journal of Statistical Software*, 15/2:1–11, 2006.
- R. I. Jennrich. Asymptotic Properties of Non-Linear Least Squares Estimators. *The Annals of Mathematical Statistics*, 40(2):633 – 643, 1969. doi: 10.1214/aoms/1177697731.
- A. R. Lundborg, I. Kim, R. D. Shah, and R. J. Samworth. The projected covariance measure for assumption-lean variable significance testing. *arXiv preprint arXiv:2211.02039*, 2022.
- J. Pinheiro, D. Bates, and R Core Team. *nlme: Linear and Nonlinear Mixed Effects Models*, 2022. URL <https://CRAN.R-project.org/package=nlme>. R package version 3.1-161.
- R. D. Shah and J. Peters. The hardness of conditional independence testing and the generalised covariance measure. *The Annals of Statistics*, 48(3), June 2020. doi: 10.1214/19-aos1857.
- A. W. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 1998.
- S. Wood. *Generalized Additive Models: An Introduction with R*. Chapman and Hall/CRC, 2 edition, 2017.