

Distributed Bayesian Varying Coefficient Modeling Using a Gaussian Process Prior

Rajarshi Guhaniyogi

*Department of Statistics
Texas A & M University
College Station, TX 77843-3143, USA*

RAJGUHANIYOGI@TAMU.EDU

Cheng Li

*Department of Statistics and Data Science
National University of Singapore
Singapore 117546, Singapore*

STALIC@NUS.EDU.SG

Terrance D. Savitsky

*U.S. Bureau of Labor Statistics
Office of Survey Methods Research
Washington, DC 20212, USA*

SAVITSKY.TERRANCE@BLS.GOV

Sanvesh Srivastava

*Department of Statistics and Actuarial Science
University of Iowa
Iowa City, IA 52242, USA*

SANVESH-SRIVASTAVA@UIOWA.EDU

Editor: Robert McCulloch

Abstract

Varying coefficient models (VCMs) are widely used for estimating nonlinear regression functions for functional data. Their Bayesian variants using Gaussian process priors on the functional coefficients, however, have received limited attention in massive data applications, mainly due to the prohibitively slow posterior computations using Markov chain Monte Carlo (MCMC) algorithms. We address this problem using a divide-and-conquer Bayesian approach. We first create a large number of data subsamples with much smaller sizes. Then, we formulate the VCM as a linear mixed-effects model and develop a data augmentation algorithm for obtaining MCMC draws on all the subsets in parallel. Finally, we aggregate the MCMC-based estimates of subset posteriors into a single Aggregated Monte Carlo (AMC) posterior, which is used as a computationally efficient alternative to the true posterior distribution. Theoretically, we derive minimax optimal posterior convergence rates for the AMC posteriors of both the varying coefficients and the mean regression function. We provide quantification on the orders of subset sample sizes and the number of subsets. The empirical results show that the combination schemes that satisfy our theoretical assumptions, including the AMC posterior, have better estimation performance than their main competitors across diverse simulations and in a real data analysis.

Keywords: Varying coefficient models, distributed Bayesian computations, data augmentation, multivariate Gaussian process, posterior convergence rates.

1. Introduction

We first introduce the motivation of studying Bayesian varying coefficient models with a Gaussian process prior for massive data applications. Then, we outline our main contributions in this work and discuss the related literature.

1.1 Varying Coefficient Models Using a GP Prior

VCMs are a flexible and popular extension of the linear regression model (Hastie and Tibshirani, 1993), in which the regression coefficients can be smooth functions that capture nonlinear dependence of the response function on the covariates. VCMs are extensively used in practice, including time-series (Chen and Tsay, 1993; Cai et al., 2000), longitudinal (Wu et al., 1998; Ruppert et al., 2003), spatial (Gelfand et al., 2003), and spatiotemporal data analysis (Lu et al., 2009). Bayesian VCMs combine the flexibility of nonparametric models and the interpretability of parametric models and provide uncertainty estimates in inference and predictions via MCMC draws from the posterior distribution; therefore, they are well-suited for the Bayesian analysis of massive time-series, healthcare, and spatial/spatiotemporal databases.

We focus on Bayesian VCMs in which the varying coefficients are assigned a multivariate GP prior. Without loss of generality, we assume that all functional variables (responses or covariates) are defined on the d -dimensional indexing space $[0, 1]^d$ ($d \in \mathbb{N}$). The index is time and $d = 1$ in purely times series applications, whereas $d = 2$ and the index is a spatial location in purely spatial applications. The two indices are combined in spatiotemporal applications, where $d = 3$ and the index is a space-time tuple. More generally, for a sample of indexes $\{u_i : i = 1, \dots, n\}$ ($n \in \mathbb{N}$) from $[0, 1]^d$, we observe the s_i -dimensional i th response vector $y(u_i) \in \mathbb{R}^{s_i}$ ($s_i \in \mathbb{N}$) and the matrix of i th covariate functions $X(u_i) \in \mathbb{R}^{s_i \times p}$, where $p \in \mathbb{N}$ is the number of covariates. We consider a VCM with the form

$$y(u_i) = X(u_i)\beta(u_i) + \epsilon(u_i), \quad \epsilon(u_i) \stackrel{\text{i.i.d.}}{\sim} N(0, \tau^2 I_{s_i}), \quad i = 1, \dots, n, \quad (1)$$

where $\beta(u) = \{\beta_1(u), \dots, \beta_p(u)\}^T \in \mathbb{R}^p$ for $u \in [0, 1]^d$ is the vector of varying regression coefficients, 0 and I_s are a zero vector and an identity matrix of dimension s , and $\epsilon(u_i) \in \mathbb{R}^{s_i}$ ($i = 1, \dots, n$) are idiosyncratic normal errors. The responses in (1) are allowed to have different dimensions, but $s_i = s$ for every i ($s \in \mathbb{N}$) in a typical scientific application.

The VCM setup in (1) has advantages over its peers in the literature. For example, the varying coefficients $\beta(u)$ provide a more flexible and realistic modeling of responses and predictors with space or space-time indices, so they perform better in practice than fitting deterministic trends in covariates, such as polynomial regression (Gelfand et al., 2003). There are some methods for modeling the varying coefficients $\beta(u)$ for $u \in [0, 1]^d$ in (1), but we use a multivariate GP prior distribution on $\beta(u)$; see Section 2 for a detailed description. In fact, VCMs in the existing literature typically rely on basis expansions using local polynomials, P-splines, and trees for modeling the varying coefficients $\beta(u)$ (Li and Racine, 2007; Marx, 2010; Berger et al., 2019). The specification of the number of basis functions or the height of the tree and the choice of knots or split locations is usually difficult in practice. In comparison, Bayesian inference using GP priors only requires a tuning-free prior specification for the covariance parameters. Even low-dimensional structures in the

data are conveniently modeled using a GP prior projected on a moderately large number of inducing points (Quiñonero-Candela and Rasmussen, 2005).

While the VCM in (1) with a GP prior on $\beta(u)$ has such advantages in the modeling of structured data, there are several practical considerations, including inefficient posterior computations, that have severely restricted its application in massive data settings. Posterior sampling involving GPs are already prohibitively slow if the sample size is large. In fact, the cost per MCMC iteration for updating $\beta(u)$ scales as $O(p^2n^2)$ for storage and $O(p^3n^3)$ for computations. As a result, the simple posterior sampling scheme for inference in (1) using Gibbs and slice sampling as originally proposed in Gelfand et al. (2003) becomes infeasible in practice. Even with the low-rank GP approximation techniques using r inducing points (Quiñonero-Candela and Rasmussen, 2005; Álvarez et al., 2012), one can only reduce the cost per MCMC iteration from $O(n^3p^3)$ to $O(npr^2)$ (Álvarez and Lawrence, 2011). Furthermore, r is chosen to be sufficiently large, typically of the order $O\{(\log n)^d\}$, to achieve satisfactory approximation accuracy (Burt et al., 2019). The use of deep GP priors in (1) further worsens the computational burden and cannot be used in practice (Damianou and Lawrence, 2013; Duvenaud et al., 2014). Finally, variational inference has been widely used in machine learning for inference in applications based on VCMs with multivariate GPs involving big data, but MCMC-based inference has remained relatively unexplored in this context (Álvarez and Lawrence, 2011; Álvarez et al., 2019; Yousefi et al., 2019). MCMC based inference has the natural advantage of accurately characterizing the uncertainty of inference and prediction in VCMs with functional data, having strong local features. This is crucial in the spatio-temporal application of interest in Section 5.3, which aims at understanding local features in the space-time varying relationship between sea surface temperature and sea salinity for the Atlantic ocean based on large functional data. Our posterior inference algorithm fills this gap, providing a scalable MCMC-based alternative.

Addressing the computational bottlenecks for the VCM in (1) with a GP prior on $\beta(u)$, we develop a three stage distributed Bayesian inferential approach for efficient computation with functional response and covariates obtained at a large number of indices. The first stage of the algorithm constructs k subsets by randomly selecting m samples without replacement from the full data, where k is large and posterior computations with m is tractable. The second step obtains k MCMC-based approximations of the full data posterior distribution by fitting the VCM in (1) with a GP prior on $\beta(u)$ on all the subsets in parallel. This step has two main novelties. First, we compensate for the missing $(1 - m/n)$ -fraction of the full data in each subset by appropriately modifying the subset likelihood. Second, we reformulate the VCM in (1) with a GP prior on $\beta(u)$ as a linear mixed-effects model using parameter expansion. This leads to an MCMC algorithm that has closed-form full conditional distributions for all the parameters, except those used for defining the covariance function of the GP prior. We draw these parameters using elliptical slice sampling (ESS, Nishihara et al., 2014), which bypasses the proposal tuning problems of Metropolis-Hastings algorithm. The parameter expanded DA with the ESS step constitutes our DA-type algorithm for posterior inference and predictions on the subsets.

The subset posterior computations are tractable because $m \ll n$ and parameter updating is efficient due to the closed-form full conditionals; however, posterior computations a subset condition on m samples only. The third stage of the algorithm develops a combination scheme that aggregates MCMC-based approximations of the true posterior distribution

from the k subsets into the AMC posterior, which uses information from all the n samples. This step has several theoretical novelties. First, we identify regularity assumptions under which the AMC posterior distributions of both the varying coefficients and the mean regression function have minimax optimal posterior convergence rates in the L_2 norm toward their truth. Development of such guarantees in VCMs with multivariate latent GPs remains an open problem since their proposal in Gelfand et al. (2003). Second, our results provide quantification on the orders of the subset size m , the number of subsets k , and modification of the subset likelihood according to the underlying smoothness of varying coefficients. Finally, our theory only requires a weak condition on the combination scheme, so it encompasses a few existing combination methods, including the AMC posterior proposed in this paper as well as the double parallel Monte Carlo (DPMC, Xue and Liang, 2019), Wasserstein posterior (WASP, Srivastava et al., 2015), and posterior interval estimation (PIE, Li et al., 2017) algorithms. The minimax optimality of the AMC posterior distribution implies that it can be used for principled Bayesian inference in massive data settings with very large n if m and k are chosen appropriately.

1.2 Related Work

The theoretical and computational properties of frequentist estimation methods for VCMs have been studied extensively. The theoretical results focus mainly on VCMs that use local polynomial smoothing, regularized basis expansions, and boosted trees (Hastie and Tibshirani, 1993; Fan and Zhang, 1999; Huang et al., 2002; Zhou and Hooker, 2019); see Park et al. (2015) for a recent review. The software for fitting VCMs is also well-developed (Wood, 2017). On the other hand, Bayesian VCMs have been widely applied to different types of data (Gelfand et al., 2003; Bakar et al., 2015; Hamm et al., 2015; Datta et al., 2016), but the literature on their theoretical properties is sparsely populated. Recently, Bai et al. (2019) have studied the theoretical properties of Bayesian VCMs based on regularized basis expansions; however, their model is different from the VCM with multivariate GPs considered in this paper, and their main focus is on the high dimensional variable selection problem, which is essentially different from our focus on applications with massive n . Furthermore, frequentist properties of the posterior distribution of regression function obtained using a univariate GP prior are known (van der Vaart and van Zanten, 2011), but their extensions to a multivariate GP prior, similar to the one used for Bayesian inference in (1), are non-trivial and have not been studied.

Furthermore, Bayesian VCMs with multivariate response functions have not been studied extensively in the literature. There are some extensions of factor models based on independent GP priors that are used for modeling multivariate responses. One such example is a spatial factor model (Ren and Banerjee, 2013) that is defined as

$$y(u_i) = \beta(u_i) + \epsilon(u_i), \quad \beta(u_i) = L\nu(u_i), \quad y(u_i) \in \mathbb{R}^s, \quad s \in \mathbb{N}, \quad u_i \in [0, 1]^2, \quad (2)$$

for $i = 1, \dots, n$, where L is a p -by- q factor loading matrix and $\nu(\cdot) = \{\nu_1(\cdot), \dots, \nu_q(\cdot)\}^T$ is a vector of q spatial factors, all following mutually independent univariate GPs $\nu_1(\cdot), \dots, \nu_q(\cdot)$. Gu and Shen (2020) and Ren and Banerjee (2013) also specify identifiability constraints on L in (2) for valid frequentist estimation and Bayesian inference using MCMC, respectively. Compared to (1), every $y(u_i)$ in (2) has the same dimension and the covariate matrix

$X(u_i)$ is unobserved due to the unsupervised nature of the model. While we also propose a very similar formulation of the varying coefficients $\beta(u)$ using the linear model of co-regionalization (LMC) approach, our main focus is posterior inference on the regression coefficient $\beta(u)$ and prediction of $y(u)$ in (1), which does not require inference on L or $\nu(\cdot)$.

We now turn our focus to distributed Bayesian inference in (1). The strategy of modifying the subset likelihoods for obtaining better uncertainty characterization in distributed Bayesian inference for parametric models has been discussed in Minsker et al. (2017). Each subset contains only m/n -fraction of the full data, so the posterior distribution computed from the usual likelihood overestimates the uncertainty relative to the true posterior distribution. Thus, the modification of the subset likelihood is essential for accurate uncertainty quantification in parametric models (Minsker et al., 2017); however, the subset likelihood modification strategy for parametric models cannot be straightforwardly applied for Bayesian VCMs, due to the lack of any supporting theoretical result. One of our main contributions is to identify the subset likelihood modification and to justify it through rigorous theoretical results for VCMs; see Sections 3.2 and 4. If we use the asymptotic posterior L_2 -risk of a combined posterior distribution for quantifying its performance, then the likelihood modifications required for asymptotic optimality are different in the parametric models and VCMs based on an ‘‘appropriately tuned’’ multivariate GP prior.

The AMC algorithm belongs to the class of divide-and-conquer (or distributed) methods for Bayesian inference. These methods have been studied extensively for scalable Bayesian inference in parametric models (Scott et al., 2016; Entezari et al., 2017; Minsker et al., 2017; Li et al., 2017; Srivastava et al., 2018; Xue and Liang, 2019; Jordan et al., 2019) and nonparametric regression using univariate GP priors (Zhang et al., 2015; Shang and Cheng, 2017; Shang et al., 2019; Szabó and van Zanten, 2019, 2020; Zhang and Williamson, 2019; Guhaniyogi et al., 2017). Unfortunately, the literature fails to address distributed Bayesian inference in (1) using multivariate GP priors, which is our main focus. All these methods consist of three main steps: dividing the massive data set into smaller computationally manageable subsets, performing statistical estimation on the subsets in parallel, and combining the subset estimates into a global estimate, which is used as an alternative to the true posterior distribution. Existing distributed Bayesian inference methods differ mainly in the third step that computes the global estimate. Given some minimal requirements are met in the combination step, our theoretical results can be used to obtain posterior convergence rates for any of these combination schemes used for distributed Bayesian inference in (1). We demonstrate that such requirements are indeed met by the combination schemes used in the AMC, DPMC, PIE, and WASP algorithms.

2. Model Setup and Prior Specification

In this section, we describe the varying-coefficient model setup and its equivalent formulation as a linear mixed-effects model. We then provide the prior specification and outline the main data augmentation algorithm that is used to fit the model.

2.1 Model Reformulation

Consider a general VCM setup based on (1) with p predictors out of which q predictors have varying coefficients such that $p \geq q$. Without loss of generality, assume that the first q

predictors have varying coefficients, so that $\beta(u) = \{\beta_{\text{va}}(u), \beta_{\text{nv}}\}^T$ in (1), where $\beta_{\text{va}}(u) \in \mathbb{R}^q$ for every $u \in [0, 1]^d$ and $\beta_{\text{nv}} \in \mathbb{R}^{p-q}$ are the varying and non-varying coefficients blocks. For performing Bayesian inference on $\beta(\cdot)$, a typical strategy is to assign multivariate GP prior and Gaussian prior distributions on $\beta_{\text{va}}(\cdot)$ and β_{nv} , respectively, and obtain MCMC draws from the posterior distribution of $\beta(\cdot)$ using (1).

The most important part of the prior specification is to choose a cross-covariance function for the multivariate GP prior on $\beta_{\text{va}}(\cdot)$ that is flexible and leads to simple posterior computations. Versatile constructions exist for specifying the cross covariance of $\beta_{\text{va}}(u)$ (Gaspari and Cohn, 1999; Majumdar and Gelfand, 2007; Wackernagel, 2006; Zhang, 2007; Genton and Kleiber, 2015; Bourotte et al., 2016). We, however, adopt the LMC technique (Álvarez et al., 2012) for inducing correlation among the components of $\beta_{\text{va}}(u)$ due to its simplicity and relatively efficient computation. Under the LMC framework, we set $\beta_{\text{va}}(\cdot) = \alpha_{\text{va}} + \Gamma\nu(\cdot)$, where $\alpha_{\text{va}} \in \mathbb{R}^q$, $\Gamma \in \mathbb{R}^{q \times q}$, and $\nu(u) = \{\nu_1(u), \dots, \nu_q(u)\}^T$ is a vector of q independent GPs indexed by $[0, 1]^d$ with mean functions 0 and correlation functions $\rho_1(\cdot, \cdot), \dots, \rho_q(\cdot, \cdot)$ with parameters $\theta_1, \dots, \theta_q$, respectively. The independent GP priors on $\nu_1(\cdot), \dots, \nu_q(\cdot)$ induces a multivariate GP prior on $\beta_{\text{va}}(\cdot)$. Specifically, given α_{va} , Γ , and $\theta_1, \dots, \theta_q$, $\beta_{\text{va}}(\cdot) = \alpha_{\text{va}} + \Gamma\nu(\cdot)$ is a q -variate GP with mean function α_{va} and covariance function $C(u, u')$ defined as

$$C(u, u') = \text{Cov}\{\Gamma\nu(u), \Gamma\nu(u')\} = \sum_{a=1}^q \Gamma_a \rho_a(u, u') \Gamma_a^T = \sum_{a=1}^q \Gamma_a \{R(u, u')\}_{aa} \Gamma_a^T, \quad (3)$$

where $u, u' \in [0, 1]^d$, Γ_a is the a th column of Γ , and $R(u, u') = \text{diag}\{\rho_1(u, u'), \dots, \rho_q(u, u')\}$ is a q -by- q diagonal matrix of correlations determined by $\theta^T = (\theta_1^T, \dots, \theta_q^T)$.

We now reformulate the VCM in (1) with a GP prior imposed on $\beta_{\text{va}}(\cdot)$ using the LMC technique as linear mixed-effects model. Define $\alpha = (\alpha_{\text{va}}, \beta_{\text{nv}})^T \in \mathbb{R}^p$ and $Z(u_i) \in \mathbb{R}^{s_i \times q}$ to be the matrix that includes the first q columns of $X(u_i)$ ($i = 1, \dots, n$). Reformulate (1) as

$$y(u_i) = X(u_i)\alpha + Z(u_i)\Gamma\nu(u_i) + \epsilon(u_i), \quad \nu(\cdot) \sim \text{GP}\{0, R(\cdot, \cdot)\}, \quad i = 1, \dots, n, \quad (4)$$

where $R(\cdot, \cdot) = \text{diag}\{\rho_1(\cdot, \cdot), \dots, \rho_q(\cdot, \cdot)\}$ is the correlation ‘‘function’’ for $\nu(\cdot)$. The models in (1) and (4) are equivalent if we let $\beta_{\text{va}}(u) = \alpha_{\text{va}} + \Gamma\nu(u)$ for all $u \in [0, 1]^d$. The parameters α and Γ in (4) cannot be estimated uniquely from the data $\{y(u_i), X(u_i) : i = 1, \dots, n\}$ but the vector $\{\beta(u_1), \dots, \beta(u_n)\}$ is still estimable if the design matrix formed by $\{X(u_i) : i = 1, \dots, n\}$ as the row blocks is of full column rank. The prior distributions on the unknown parameters $\alpha, \Gamma, \tau^2, \theta$ are spelled out in Section 2.2.

Many widely used models are obtained as special cases of (4). If $\rho_a(u, u') = 1_{u=u'}$ for every a , where $1_{u=u'}$ equals 1 if $u = u'$ and 0 otherwise, then we recover the linear-mixed effects model using (4), where $\Gamma\Gamma^T$ equals the covariance matrix of the random effects. If $s_i = 1, p = q$, Γ is a diagonal matrix, and $X(u_i) = Z(u_i)$, then (1) reduces to

$$y(u_i) = X(u_i)\{\alpha + \Gamma\nu(u_i)\} + \epsilon(u_i) \equiv X(u_i)\beta(u_i) + \epsilon(u_i), \quad i = 1, \dots, n, \quad (5)$$

where α and $\Gamma\nu(\cdot)$ model the global and local effects, respectively, the diagonal entries of Γ determine the scale of local effects, and $\beta(\cdot)$ is the p -by-1 varying coefficients vector. The spatiotemporal varying coefficient model is a special case of (5) when $u \in [0, 1]^3$ (Gelfand et al., 2003; Gelfand and Banerjee, 2010). Finally, assuming u to be the time domain in (5) yields a regression model for longitudinal data analysis.

2.2 Prior Specification

The parameters (α, Γ, τ^2) are jointly assigned a noninformative prior with density $p(\alpha, \Gamma, \tau^2) \propto 1/\tau^2$. If γ represents the q^2 -dimensional vector formed by stacking the columns of Γ , then this prior is a limiting case of the normal-inverse-gamma prior distribution on $\{(\alpha, \gamma), \tau^2\}$, where $(p + q^2)$ -variate normal prior distribution is assigned on (α, γ) . We are also not concerned with the identifiability of Γ or α since they are intermediate latent variables enabling efficient estimation of $\beta_{\text{va}}(u)$ for every $u \in [0, 1]^d$.

As far as the choice of $\rho_a(\cdot, \cdot)$ ($a = 1, \dots, q$) is concerned, two types of correlation functions are used in this paper. The first one is the exponential correlation function defined as $\rho_a(u, u') = e^{-\phi_a \|u - u'\|_2}$ for any $u, u' \in [0, 1]^d$, where $\phi_a > 0$, $\|\cdot\|_2$ is the Euclidean norm, and $\theta_a = \{\phi_a\}$ ($a = 1, \dots, q$). We also use Gneiting's correlation function for varying coefficient modeling of spatiotemporal data presented in Section 5.3, which is defined as

$$\rho_a(u, u') = \frac{1}{(\psi_a |t - t'|^2 + 1)^{\kappa_a}} e^{-\frac{\phi_a \|h - h'\|_2}{(\psi_a |t - t'|^2 + 1)^{\kappa_a/2}}}, \quad u, u' \in [0, 1]^3, \quad (6)$$

where $u = (h, t), u' = (h', t')$ are space-time tuples, $h, h' \in [0, 1]^2, t, t' \in [0, 1], \phi_a > 0, \psi_a > 0, \kappa_a \in [0, 1]$, and $\theta_a = (\phi_a, \psi_a, \kappa_a)^T$ (Gneiting, 2002). For the exponential correlation function, we put a Uniform($\underline{c}_{0a}, \bar{c}_{0a}$) prior on ϕ_a . The parameters ϕ_a, ψ_a , and κ_a are assigned Uniform($\underline{c}_{1a}, \bar{c}_{1a}$), Uniform($\underline{c}_{2a}, \bar{c}_{2a}$), and Uniform($\underline{c}_{3a}, \bar{c}_{3a}$) priors respectively for the Gneiting's correlation function. The parameters for the uniform priors satisfy $0 < \underline{c}_{ia} < \bar{c}_{ia}$ for $i = 0, 1, 2$ and $0 < \underline{c}_{3a} < \bar{c}_{3a} \leq 1$.

2.3 The DA-type Algorithm

The DA-type algorithm for posterior inference on $\beta(\cdot), \tau^2$ and prediction of $y(\cdot)$ has six parts. Let \mathcal{U}^* be a given subset of $[0, 1]^d$ where the draws of $\beta(\cdot)$ and $y(\cdot)$ are required, \mathcal{D} be the training data, and $\nu_n = \{\nu(u_1), \dots, \nu(u_n)\}$, where $\nu(u_i)$ is defined in (4). The first part of the DA-type algorithm is the Imputation (I) step that draws ν_n given \mathcal{D} and $(\alpha, \Gamma, \tau^2, \theta)$. The second part of the DA-type algorithm is the Prediction (P) step that has five sub parts. It uses the ν_n to draw $(\alpha, \Gamma, \tau^2, \theta)$ given \mathcal{D} and $\{\beta(u^*), y(u^*) : u^* \in \mathcal{U}^*\}$ given $(\alpha, \Gamma, \tau^2, \theta)$. The I and P steps are repeated until convergence to the stationary distribution of the Markov chain for $(\tau^2, \{\beta(u^*), y(u^*) : u^* \in \mathcal{U}^*\})$; see Appendix A for derivation of the six parts and their analytic forms.

This DA-type algorithm is slow in moderately large data sets. The computational complexity of the I step is $O(n^3 p^3)$ if we update multivariate GPs. Low rank GP methods provide some computational relief, though the computational gain is not substantial if one needs to maintain the inferential accuracy. The sparse iterative methods for sampling from GPs lead to only marginal improvements because the number of iterations have to be relatively large for guaranteeing accurate approximation (Chow and Saad, 2014). Due to the slow I step, the DA-type algorithm using the full data is extremely inefficient in applications with a large n . The next section presents an extension of the DA-type algorithm using divide-and-conquer technique that overcomes these inefficiencies while retaining its simplicity and numerical stability.

3. Distributed Varying Coefficient Modeling Using a GP Prior

Our distributed model fitting of the Bayesian VCM consists of three steps described below.

3.1 First Step: Constructing Training Data Subsets

The first step of the distributed extension of the DA-type algorithm in Section 2.3 constructs k subsets from the training data. The default scheme for constructing subsets is to randomly sub-sample without replacement from the training data, ensuring that each subset provides a reliable representation of the full data and all observations specific to a sample are on the same subset. The size of a subset m is set to be moderately large so that $p, q \ll m$ and posterior computations are efficient on any subset. The ‘‘optimal choice’’ of k depends on the smoothness of the regression function, which we study in Section 4. Let \mathcal{D}_j be the training data on subset j ($j = 1, \dots, k$), u_{ji} be the i th index in subset j , and $y(u_{ji})$, $X(u_{ji})$, $Z(u_{ji})$ be the corresponding observations with dimensions s_{ji} , s_{ji} -by- p , s_{ji} -by- q , for $i = 1, \dots, m$. Similarly, the subset j versions of parameters $(\beta, \alpha, \Gamma, \tau^2, \theta)$ are denoted by $(\beta_j, \alpha_j, \Gamma_j, \tau_j^2, \theta_j)$. The correlation function ρ_{ja} equals ρ_a but replaces θ_a by θ_{ja} . The GP with correlation function ρ_{ja} is denoted as $\tilde{\nu}_{ja}(\cdot)$, and $\tilde{\nu}_j(\cdot) = \{\tilde{\nu}_{j1}(\cdot), \dots, \tilde{\nu}_{jq}(\cdot)\}^T$.

The VCM in (4) has a natural extension to subset j . For $i = 1, \dots, m$,

$$\begin{aligned} y(u_{ji}) &= X(u_{ji})\alpha_j + Z(u_{ji})\Gamma_j\tilde{\nu}_j(u_{ji}) + \epsilon(u_{ji}), \\ \tilde{\nu}_{ja}(\cdot) &\sim \text{GP}\{0, \rho_{ja}(\cdot, \cdot)\}, \quad a = 1, \dots, q, \quad \epsilon(u_{ji}) \sim N(0, \tau_j^2), \end{aligned} \quad (7)$$

which reduces to the subset j extension of (1) if $\{\beta_j(u)\}_{\text{va}} = (\alpha_j)_{\text{va}} + \Gamma_j\tilde{\nu}_j(u)$, $u \in [0, 1]^d$, where $\beta_j(u) = [\{\beta_j(u)\}_{\text{va}}, (\beta_j)_{\text{nv}}]$ and $\alpha_j = \{(\alpha_j)_{\text{va}}, (\beta_j)_{\text{nv}}\}$ are represented in terms of their varying and non-varying coefficients blocks. The prior distributions for $(\alpha_j, \Gamma_j, \tau_j^2)$ and θ_j in (7) are the same as defined in Section 2.2 for (α, Γ, τ^2) and θ , respectively. If we obtain MCMC draws of the parameters and predictions using the likelihood in (7) on each subset directly, then we condition on an (m/n) -fraction of the full data, resulting in wider credible intervals for parameters than those obtained using the full data posterior distribution. The next step fixes this problem by using a modified likelihood based on (7) that compensates for the missing $(1 - m/n)$ -fraction of the full data.

3.2 Second Step: Posterior Sampling on the Subsets

We now consider the inference on each subset using the DA-type algorithm based on Section 2.3 and (7). The GP realizations $\tilde{\nu}_j(u_{j1}), \dots, \tilde{\nu}_j(u_{jm})$ in (7) are looked upon as the ‘‘missing’’ data, and marginalizing over them recovers the subset j version of (1) with a GP prior on $[\{\beta_j(u_{j1})\}_{\text{va}}, \dots, \{\beta_j(u_{jm})\}_{\text{va}}]$. For $a = 1, \dots, q$, we let $\tilde{s}_j = \sum_{i=1}^m s_{ji}$ and define

$$\begin{aligned} \tilde{\nu}_{ja} &= \{\tilde{\nu}_a(u_{j1}), \dots, \tilde{\nu}_a(u_{jm})\}^T \in \mathbb{R}^m, \quad \tilde{\nu}_j = (\tilde{\nu}_{j1}^T, \dots, \tilde{\nu}_{jq}^T)^T \in \mathbb{R}^{mq}, \\ \Gamma_j &= (\Gamma_{j1}, \dots, \Gamma_{jq}), \quad \Gamma_{j1}, \dots, \Gamma_{jq} \in \mathbb{R}^q, \\ \tilde{Z}_{ja} &= \text{diag}\{Z(u_{j1})\Gamma_{ja}, \dots, Z(u_{jm})\Gamma_{ja}\}, \\ y_j &= \{y(u_{j1})^T, \dots, y(u_{jm})^T\}^T, \quad X_j = \{X(u_{j1})^T, \dots, X(u_{jm})^T\}^T. \end{aligned} \quad (8)$$

If $\tilde{\nu}_j$ is known, then the full conditional for drawing $(\alpha_j, \Gamma_j, \tau_j^2)$ given $\tilde{\nu}_j$, \mathcal{D}_j is available in closed-form; therefore, $\tilde{\nu}_j$ is an auxiliary variable that simplifies the forms of the full

conditionals if it is known and its marginalization preserves the Bayesian VCM with a GP prior on $\{\beta_j(\cdot)\}_{\text{nv}}$.

The I step of the DA-type algorithm on subset j uses this property of $\tilde{\nu}_j$ for simplifying the form of the modified likelihood. Assume that $(\alpha_j, \Gamma_j, \tau_j^2, \theta_j)$ are given. Let $p(\tilde{\nu}_j \mid \mathcal{D}_j, \alpha_j, \Gamma_j, \tau_j^2, \theta_j)$ denote the conditional density of $\tilde{\nu}_j$ given \mathcal{D}_j and $(\alpha_j, \Gamma_j, \tau_j^2, \theta_j)$ based on (7). Then, the I step

- (a) draws $\tilde{\nu}_j$ given \mathcal{D}_j and $(\alpha_j, \Gamma_j, \tau_j^2, \theta_j)$ from $N_{mq}(\mu_{\tilde{\nu}_j}, \Sigma_{\tilde{\nu}_j})$, where $\mu_{\tilde{\nu}_j}$ and $\Sigma_{\tilde{\nu}_j}$ are defined in terms of their blocks as

$$\begin{aligned} (\mu_{\tilde{\nu}_j})_a &= R_{ja}^T \tilde{Z}_{ja}^T \left(\sum_{c=1}^q \tilde{Z}_{jc} R_{jc} \tilde{Z}_{jc}^T + \tau_j^2 I_{\tilde{s}_j} \right)^{-1} (y_j - X_j \alpha_j), \quad a = 1, \dots, q, \\ (\Sigma_{\tilde{\nu}_j})_{aa} &= R_{ja} - R_{ja}^T \tilde{Z}_{ja}^T \left(\sum_{c=1}^q \tilde{Z}_{jc} R_{jc} \tilde{Z}_{jc}^T + \tau_j^2 I_{\tilde{s}_j} \right)^{-1} \tilde{Z}_{ja} R_{ja}, \quad a = 1, \dots, q, \\ (\Sigma_{\tilde{\nu}_j})_{ab} &= -R_{ja}^T \tilde{Z}_{ja}^T \left(\sum_{c=1}^q \tilde{Z}_{jc} R_{jc} \tilde{Z}_{jc}^T + \tau_j^2 I_{\tilde{s}_j} \right)^{-1} \tilde{Z}_{jb} R_{jb}, \quad a \neq b \in \{1, \dots, q\}, \\ (R_{ja})_{i'i'} &= \rho_{ja}(u_{ji}, u_{ji'}), \quad i, i' = 1, \dots, m, \quad a = 1, \dots, q. \end{aligned} \quad (9)$$

If we substitute the I step draw of $\tilde{\nu}_j$ in (7) and compute the likelihood of $(\alpha_j, \Gamma_j, \tau_j^2, \theta_j)$ given $(\mathcal{D}_j, \tilde{\nu}_j)$, then this is equivalent to computing the likelihood of $(\alpha_j, \Gamma_j, \tau_j^2, \theta_j)$ after marginalizing over $\tilde{\nu}_j$ in (7) using Monte Carlo. Denote this Monte Carlo based likelihood as L_j , and we use it as the likelihood of $(\alpha_j, \Gamma_j, \tau_j^2, \theta_j)$ given \mathcal{D}_j in the P step.

The subset j draws $(\alpha_j, \Gamma_j, \tau_j^2, \theta_j)$ given \mathcal{D}_j using a modified version of L_j . Since \mathcal{D}_j contains an (m/n) -fraction of the full data, we raise L_j to a power of δ_n , where δ_n is a deterministic sequence dependent on n . Let $L_j^{\delta_n}$ be this modified likelihood, and the modification is equivalent to replicating \mathcal{D}_j for δ_n -times. The power δ_n is chosen such that $L_j^{\delta_n}$ compensates for the missing $(1 - m/n)$ -fraction of the full data on subset j . The modified posterior density for drawing $(\alpha_j, \Gamma_j, \tau_j^2, \theta_j)$ given \mathcal{D}_j is defined as

$$\pi_m(\alpha_j, \Gamma_j, \tau_j^2, \theta_j \mid \mathcal{D}_j) = \frac{L_j^{\delta_n} p(\alpha_j, \Gamma_j, \tau_j^2) p(\theta_j)}{\int L_j^{\delta_n} p(\alpha_j, \Gamma_j, \tau_j^2) p(\theta_j) d\alpha_j d\Gamma_j d\tau_j^2 d\theta_j}, \quad j = 1, \dots, k, \quad (10)$$

where the denominator is finite due to the choice of prior distributions. The method of raising subset likelihoods to a power is known as the *stochastic approximation* (Minsker et al., 2014). We choose $\delta_n = n/m$ following the same choice in parametric models (Minsker et al., 2014; Entezari et al., 2017; Li et al., 2017; Srivastava et al., 2018), which is equivalent to replicating the subset data for n/m times, such that the subset posterior variances of parameters are comparable to the full data posterior variance. The theoretical impact from δ_n will be further discussed in Section 4.

The P step draws $\alpha_j, \Gamma_j, \tau_j^2$, and θ_j given $(\mathcal{D}_j, \tilde{\nu}_j)$ in a sequence of three steps using (10). Define $\gamma_j = (\Gamma_{j1}^T, \dots, \Gamma_{jq}^T)^T$, the column-wise vectorization of Γ_j , $b_j = (\alpha_j^T, \gamma_j^T)^T$, and

$$W_j = (W_{j1}^T, \dots, W_{jm}^T)^T, \quad W_{ji} = [X_{ji} \ \{\tilde{\nu}_1(u_{ji}), \dots, \tilde{\nu}_q(u_{ji})\} \otimes Z_{ji}] \in \mathbb{R}^{s_{ji} \times (p+q^2)}, \quad (11)$$

for $i = 1, \dots, m$, where $W_j \in \mathbb{R}^{\tilde{s}_j \times (p+q^2)}$ and \otimes is the Kronecker product. Then, the P step draws $\alpha_j, \Gamma_j, \tau_j^2$, and θ_j as follows:

(b) draw τ_j^2 given $\tilde{\nu}_j$ and \mathcal{D}_j as

$$\tau_j^2 \sim \frac{\delta_n \|y_j - \hat{y}_j\|_2^2}{\chi_{\delta_n \tilde{s}_j - p - q^2}^2}, \quad \hat{y}_j = W_j (W_j^T W_j)^{-1} W_j^T y_j, \quad (12)$$

where $\chi_{\delta_n \tilde{s}_j - p - q^2}^2$ is a chi-square random variable with $\delta_n \tilde{s}_j - p - q^2$ as its degrees of freedom.

(c) draw $b_j = (\alpha_j^T, \gamma_j^T)^T$ given $\tau_j^2, \tilde{\nu}_j$, and \mathcal{D}_j from $N\{(W_j^T W_j)^{-1} W_j^T y_j, \tau_j^2 (W_j^T W_j)^{-1}\}$; and

(d) draw $\theta_{j1}, \dots, \theta_{jq}$ given $\tilde{\nu}_j$ and \mathcal{D}_j using ESS (Algorithm 1 in Nishihara et al. 2014) with the modified log-likelihood for $\theta_1, \dots, \theta_q$ defined as

$$\log L(\theta_{j1}, \dots, \theta_{jq}) = -\frac{\delta_n m q}{2} \log 2\pi - \frac{\delta_n}{2} \sum_{a=1}^q \log \det(R_{ja}) - \frac{\delta_n}{2} \sum_{a=1}^q \tilde{\nu}_{ja}^T R_{ja}^{-1} \tilde{\nu}_{ja}, \quad (13)$$

where $\tilde{\nu}_{ja}$ is defined in (8), R_{ja} is a m -by- m matrix defined in (9) and depends on θ_{ja} . The form of the likelihood of $\theta_{j1}, \dots, \theta_{jq}$ depends on the correlation functions of the univariate GPs; see Appendix B for the exact details of the likelihood for the two correlation functions used in this paper.

In most applications, the goal is to perform inference on $\beta(u^*)$ and predict $y(u^*)$ for $u^* \in \mathcal{U}^*$, where $\mathcal{U}^* = \{u_1^*, \dots, u_l^*\}$ is a known subset of $[0, 1]^d$, also known as the testing set. This is done by using the parameter draws from parts (b)–(d) in the P step as follows:

(e) draw $\nu_{ja}^* = \{\nu_a(u_1^*), \dots, \nu_a(u_l^*)\}^T$ given $\tilde{\nu}_j, \theta_j$, and \mathcal{D}_j from $N(\mu_{ja}^*, \Sigma_{ja}^*)$, where

$$\begin{aligned} \mu_{ja}^* &= R_{ja^*}^T R_{ja}^{-1} \tilde{\nu}_{ja}, \quad \Sigma_{ja}^* = R_{ja^{**}} - R_{ja^*}^T R_{ja}^{-1} R_{ja^*}, \\ (R_{ja^{**}})_{i'i''} &= \rho_{ja}(u_{i'}^*, u_{i''}^*), \quad (R_{ja^*})_{ii'} = \rho_{ja}(u_i, u_{i'}^*), \end{aligned} \quad (14)$$

for $a = 1, \dots, q$, $i', i'' = 1, \dots, l$, and $i = 1, \dots, m$, and set $\{\beta_j(u^*)\}_{\text{nv}} = (\alpha_j)_{\text{nv}}$ and $\{\beta_j(u^*)\}_{\text{va}} = (\alpha_j)_{\text{va}} + \Gamma_j \nu_j(u^*)$, $u^* \in \mathcal{U}^*$; and

(f) draw $y_j(u^*)$ given $\alpha_j, \Gamma_j, \tau_j^2, X(u^*), \beta(u^*)$ independently from $N(\mu_{y_j}^*, \tau_j^2 I_{s^*})$ for every $u^* \in \mathcal{U}^*$, where $\mu_{y_j}^* = X(u^*) \beta_j(u^*)$ and s^* is the dimension of y at u^* .

The I and P steps, including the parts (a)–(f), are run in parallel on the k subsets until convergence of the Markov chain for $(\tau_j^2, \{\beta_j(u^*), y_j(u^*) : u^* \in \mathcal{U}^*\})$ to its stationary distribution; see Appendix B for derivation of the six parts and their analytic forms.

The AMC sampler cycles through steps (a)–(f) on subset j to obtain posterior draws of $\beta_j(u^*), \tau_j^2$, and $y_j(u^*)$, $u^* \in \mathcal{U}^*$ ($j = 1, \dots, k$). Let T be the number of post-burnin draws collected on every subset. Denote the parameter and prediction samples obtained from subset j at the t th iteration as $\{\beta_j^{(t)}(u^*), \tau_j^{2(t)}, y_j^{(t)}(u^*)\}$ ($t = 1, \dots, T; u \in \mathcal{U}^*$), which are called the j th subset posterior draws. We assume that the marginal j th subset posterior

draws for β , τ^2 , and y^* follow their invariant distributions denoted as $\Pi_\beta(\cdot | \mathcal{D}_j)$, $\Pi_{\tau^2}(\cdot | \mathcal{D}_j)$, and $\Pi_{y^*}(\cdot | \mathcal{D}_j)$, which are called j th subset posterior distributions and their densities are obtained using the joint density in (10). We develop next an algorithm for combining the collection of k subset posterior draws such that the combined draw follows the AMC posterior distribution that conditions on the full data.

3.3 Third Step: Aggregation of Subset Posterior Draws

We aggregate the subset posterior draws for $\beta(\cdot)$, $y(\cdot)$, and τ^2 using centering and scaling operations. Let $\beta_j^{*(t)} = \{\beta_j^{(t)}(u_1^*), \dots, \beta_j^{(t)}(u_l^*)\}$, $y_j^{*(t)} = \{y_j^{(t)}(u_1^*), \dots, y_j^{(t)}(u_l^*)\}$, and $\log \tau_j^{2(t)}$ be the t th draws for $\beta(\cdot)$, $y(\cdot)$, and $\log \tau^2$ on subset j ($j = 1, \dots, k$). Let $\xi \in \{\beta(\cdot), y(\cdot), \log \tau^2\}$ and $\xi_j^{(t)}$ be its t th draw on subset j . Define the empirical mean vector and covariance matrix of ξ draws on subset j as

$$\mu_{j\xi} = \frac{1}{T} \sum_{t=1}^T \xi_j^{(t)}, \quad \Sigma_{j\xi} = \frac{1}{T} \sum_{t=1}^T (\xi_j^{(t)} - \mu_{j\xi}) (\xi_j^{(t)} - \mu_{j\xi})^T, \quad j = 1, \dots, k. \quad (15)$$

We now summarize the algorithm for obtaining draws from the AMC posterior using the subset posterior draws. First, define the combined empirical mean and covariance matrix for ξ draws using the subset posterior empirical means and covariance matrices in (15) as

$$\mu_\xi = \frac{1}{k} \sum_{j=1}^k \mu_{j\xi}, \quad \Sigma_\xi = \frac{1}{k} \sum_{j=1}^k \Sigma_{j\xi}. \quad (16)$$

Second, center and scale the j th subset posterior draws of ξ as

$$q_{j\xi}^{(t)} = \Sigma_{j\xi}^{-1/2} (\xi_j^{(t)} - \mu_{j\xi}), \quad t = 1, \dots, T; j = 1, \dots, k. \quad (17)$$

Third, rescale and recenter the ξ draws from all the subsets in (17) as

$$\xi_{t'} = \mu_\xi + \Sigma_\xi^{1/2} q_{j\xi}^{(t)}, \quad t' = t + (j - 1)T; j = 1, \dots, k, \quad (18)$$

to obtain t' th draws from the AMC posterior distribution of ξ . The τ^2 draws are obtained by taking the exponential of draws from the AMC posterior of $\log \tau^2$.

The AMC aggregation algorithm for subset posteriors bears close connections to a few recently devised combination methods such as the Double Parallel Monte Carlo (DPMC, Xue and Liang, 2019) and Wasserstein posterior (WASP, Xu and Srivastava, 2021). All three algorithms agree on the combination of the subset posterior means but differ in their approach to combining subset posterior covariance matrices. The scaling and re-scaling steps are absent in DPMC because it relies on the asymptotic normality of subset posterior distributions. On the other hand, the combination algorithms of AMC and WASP have the same three steps, except the former and latter compute the combined covariance matrix as the arithmetic and geometric means of subset posterior covariance matrices. The computation of the geometric mean in WASP requires an iterative algorithm (Álvarez-Esteban et al., 2016), so the AMC algorithm is computationally simpler. Finally, AMC, WASP

and DPMC are applicable for aggregating posterior distributions of multivariate quantities, whereas the PIE algorithm is developed for posterior distributions of scalar quantities only.

An important contribution of this article is to theoretically establish that any aggregation step of subset posteriors constructed in Step 2 with the modified data likelihood that satisfies a simple assumption (see Assumption (A.6) in Section 4) will lead to an optimal estimation of the regression function. We establish that this assumption is satisfied not only by AMC, but also by DPMC, WASP and PIE aggregation methods for divide-and-conquer Bayesian inference mentioned above.

4. Theoretical Properties of AMC

This section derives the posterior convergence rates for the varying coefficients and the mean regression function for the three stage AMC framework under certain regularity assumptions on the smoothness of the latent GPs. This setup allows the study of VCMs with coefficients modeled using GPs with full-rank and low-rank covariance functions from a common framework.

We first make the following assumption on n, k, m and the sampling scheme of index u .

- (A.1) $c_1 n \leq km \leq c_2 n$ for some constants $0 < c_1 \leq 1 \leq c_2$. The sampled indices in the full data $\{u_1, \dots, u_n\}$ and a single testing index u^* are drawn independently from the Lebesgue measure on $[0, 1]^d$. The subset indexes $\{u_{j1}, \dots, u_{jm} : j = 1, \dots, k\}$ are drawn independently without replacement from $\{u_1, \dots, u_n\}$.

For the VCM in (4), we simplify the model setup by first assuming that $p = q$ and $X(\cdot) \equiv Z(\cdot)$; that is, every covariate function has a varying coefficient. The frequentist minimax rates for Gaussian process model with $s > 1$ are unknown, so we only consider the case where $s = 1$ because it enables direct comparison of our posterior convergence rates with existing frequentist results. A fully Bayesian asymptotic theory involves the full posterior distribution of $\alpha, \Gamma, \tau^2, \theta$ and the latent GP realizations $\nu_1(\cdot), \dots, \nu_q(\cdot)$; however, α and Γ are not identifiable, so their posteriors do not contract to any point mass as the sample size n increases to infinity. Additionally, given that our indexing space $[0, 1]^d$ is fixed and bounded, it is known that the length-scale parameters in θ are also not identifiable in fixed-domain asymptotics (Zhang, 2004). To alleviate the technical difficulties from such non-identifiable parameters, we make the following assumption:

- (A.2) $s = 1, p = q$ and $X(\cdot) \equiv Z(\cdot)$. α, Γ, τ^2 , and θ are all fixed at their true values $\alpha_0 = 0, \Gamma_0, \tau_0^2$, and θ_0 . Γ_0 is a full-rank $q \times q$ matrix. The observed response function satisfies $y(u) = Z(u)\Gamma_0\nu_0(u) + \epsilon(u)$, $\mathbb{E}[\epsilon(u)] = 0$, $\text{Var}[\epsilon(u)] = \tau_0^2$, for all $u \in [0, 1]^d$, where $\nu_0(\cdot) = \{\nu_{01}(\cdot), \dots, \nu_{0q}(\cdot)\}^T$ are the true latent functions.

Assuming $\alpha \equiv \alpha_0 = 0$ is equivalent to assuming that α is fixed at any true value α_0 , since we can always redefine $y'(u) = y(u) - X(u)\alpha_0$ and call $y'(u)$ the response function. While τ^2 is assumed fixed in Assumption (A.2), it is possible to generalize our technical proofs such that τ^2 has a prior in a bounded interval $[\underline{\tau}^2, \bar{\tau}^2]$ for some constants $0 < \underline{\tau}^2 < \bar{\tau}^2 < \infty$ (van der Vaart and van Zanten, 2008). For our theory on convergence rates, we only require the error to have mean zero and variance τ_0^2 , but do not require the true error distribution

to be normal. In other words, our convergence theory also works when our model of normal error is misspecified.

We now define some notations for the subset data. For $i = 1, \dots, m$ and $j = 1, \dots, k$, let $y_j = \{y(u_{j1}), \dots, y(u_{jm})\}^T$ and $\epsilon_j = \{\epsilon(u_{j1}), \dots, \epsilon(u_{jm})\}^T$. For $a = 1, \dots, q$, let $\tilde{\nu}_{ja} = \{\nu_a(u_{j1}), \dots, \nu_a(u_{jm})\}^T$. We have assumed that θ is known, so $\theta_{ja} = \theta_a = \theta_{0a}$ and $\rho_a = \rho_{ja}$ ($a = 1, \dots, q$). Since $s = 1$ and $p = q$, $Z(\cdot)$ is a q -dimensional row vector of functions. For $a = 1, \dots, q$, let Γ_{0a} be the a th column of Γ_0 and $\tilde{Z}_a(\cdot) = Z(\cdot)\Gamma_{0a}$. Let $\tilde{Z}(\cdot) = Z(\cdot)\Gamma_0 = \{\tilde{Z}_1(\cdot), \dots, \tilde{Z}_q(\cdot)\}$, which is still a q -dimensional row vector of functions, and $\tilde{Z}_{ja} = \{\tilde{Z}_a(u_{j1}), \dots, \tilde{Z}_a(u_{jm})\}$ for $a = 1, \dots, q$. With the stochastic approximation described in Section 3.2, our ‘‘working model’’ of VCM on the j th subset data under Assumption (A.2) can be written as

$$y_j = \sum_{a=1}^q \tilde{Z}_{ja} \tilde{\nu}_{ja} + \epsilon_j, \quad \epsilon_j \sim N\left(0, \frac{\tau_0^2}{\delta_n} I_m\right), \quad \tilde{\nu}_{ja} \sim N(0, \lambda_n^{-1} R_{ja}), \quad (19)$$

where R_{ja} is a m -by- m matrix with entries $(R_{ja})_{ii'} = \rho_a(u_{ji}, u_{ji'})$, $i, i' \in \{1, \dots, m\}$, $a = 1, \dots, q$, and ρ_a has its parameter θ_a fixed at θ_{0a} . $\delta_n = n/m$ is the same as used in the stochastic approximation definition in Section 3.2. The stochastic approximation with $\delta_n = n/m$ is crucial for ensuring that the AMC posterior of varying coefficients $\beta(\cdot)$ converges to the truth at a polynomial rate of n rather than m , such that the AMC posterior can be a valid approximation to the full data posterior, which converges at a polynomial rate of n . In (19), we have also added an additional tuning parameter $\lambda_n > 0$ that only depends on n and is only used for theory development. The value of λ_n helps offering minimax optimal rate and is specified later in Theorem 1. In practice, we simply set $\lambda_n = 1$ that provides a nearly optimal rate, and the model (19) becomes the same model as the VCM in (4). This can also be seen in Theorem 1 and in the ensuing discussion.

We focus on the posterior convergence behavior of the varying coefficients $\beta(\cdot) = \alpha_0 + \Gamma_0 \nu(\cdot)$ towards the truth $\beta_0(\cdot) = \alpha_0 + \Gamma_0 \nu_0(\cdot)$, as well as the mean function $w(\cdot) = Z(\cdot)\beta(\cdot)$ towards the truth $w_0(\cdot) = Z(\cdot)\beta_0(\cdot)$. We introduce some concepts for reproducing kernel Hilbert space (RKHS) that will be used for stating the assumptions on ν 's. Let $L_2(\mathrm{d}u)$ be the class of all square-integrable functions on $[0, 1]^d$ with respect to the Lebesgue measure, with the inner product given by $\langle f, g \rangle_{L_2(\mathrm{d}u)} = \int_{[0, 1]^d} f(u)g(u) \mathrm{d}u$ and the $L_2(\mathrm{d}u)$ -norm given by $\|f\|_2^2 = \langle f, f \rangle_{L_2(\mathrm{d}u)}$, for any generic $f, g \in L_2(\mathrm{d}u)$. For the correlation function $\rho_a(\cdot, \cdot)$ with parameters θ_{0a} , we assume that $\sup_{u, u' \in [0, 1]^d} \rho_a(u, u') < \infty$ ($a = 1, \dots, q$), which means that all correlation functions are trace class kernels. For each $a = 1, \dots, q$, by the Mercer's theorem, there exists an orthonormal sequence of eigenfunctions $\{\varphi_{ah}\}_{h=1}^\infty$ in $L_2(\mathrm{d}u)$ with eigenvalues $\mu_{a1} \geq \mu_{a2} \geq \dots \geq 0$, such that $\int_{[0, 1]^d} \rho_a(\cdot, u') \varphi_{ah}(u') \mathrm{d}u' = \mu_{ah} \varphi_{ah}(\cdot)$ for all $h = 1, 2, \dots$, and $\rho_a(u, u') = \sum_{h=1}^\infty \mu_{ah} \varphi_{ah}(u) \varphi_{ah}(u')$ for any $u, u' \in [0, 1]^d$. The RKHS \mathbb{H}_a attached to the correlation function ρ_a is the space of all functions $f \in L_2(\mathrm{d}u)$ such that the \mathbb{H}_a -norm $\|f\|_{\mathbb{H}_a}^2 = \sum_{h=1}^\infty \langle f, \varphi_{ah} \rangle_{L_2(\mathrm{d}u)}^2 / \mu_{ah} < \infty$, for $a = 1, \dots, q$.

For two positive sequences $\{a_n\}_{n \geq 1}$ and $\{b_n\}_{n \geq 1}$, the relation $\limsup_{n \rightarrow \infty} a_n/b_n \leq c$ for some constant $c > 0$ is denoted by $a_n \lesssim b_n$, or $b_n \gtrsim a_n$. If $a_n \lesssim b_n$ and $b_n \lesssim a_n$, then we say that $a_n \asymp b_n$.

We impose the following assumption on the eigenfunctions and eigenvalues of ρ_a ($a = 1, \dots, q$) as well as $Z(\cdot)$:

- (A.3) (i) There exists a constant $C_\varphi > 0$ such that $|\varphi_{ah}(u)| \leq C_\varphi$ for all $u \in [0, 1]^d$, $a = 1, \dots, q$, and $h = 1, 2, \dots$
 (ii) The largest h th eigenvalue of ρ_1, \dots, ρ_q , defined by $\mu_{h*} = \max_{a=1, \dots, q} \mu_{ah}$, satisfies $\mu_{h*} \lesssim h^{-2\nu/d}$ for every $h = 1, 2, \dots$ and some constant $\nu > d/2$.
- (A.4) For $a = 1, \dots, q$, the true latent functions satisfy $\nu_{0a} \in \mathbb{H}_a$.
- (A.5) $|Z_a(u)| \leq C_Z$ for some finite constant $C_Z > 0$ for all $u \in [0, 1]^d$ and $a = 1, \dots, q$. Let $\bar{h} = \lceil n^{3d/(2\nu-d)} \rceil$ with ν given in Assumption (A.3). For any $u \in [0, 1]^d$, define the $q\bar{h}$ -variate function

$$W(u) = \left\{ Z(u)\Gamma_{01}\varphi_{11}(u), \dots, Z(u)\Gamma_{01}\varphi_{1\bar{h}}(u), \dots, \right. \\ \left. Z(u)\Gamma_{0q}\varphi_{q1}(u), \dots, Z(u)\Gamma_{0q}\varphi_{q\bar{h}}(u) \right\}^\top \in \mathbb{R}^{q\bar{h}},$$

and the matrix $\Omega = \mathbb{E}_u \{W(u)W(u)^\top\} \in \mathbb{R}^{q\bar{h} \times q\bar{h}}$, where \mathbb{E}_u is the expectation with respect to the Lebesgue measure on $[0, 1]^d$. Then, the smallest and the largest eigenvalues of Ω are bounded away from zero and infinity by constants.

Consider the covariance function

$$C(u, u') = \sum_{h=0}^n a_h \cos(h\pi|u - u'|) = \sum_{h=0}^n a_h \{ \cos(h\pi u) \cos(h\pi u') + \sin(h\pi u) \sin(h\pi u') \},$$

for $u, u' \in [0, 1]$ with $a_h \geq 0$ for all $h = 1, 2, \dots$ and $\sum_{h=1}^\infty a_h < \infty$, then by the Mercer's theorem, $\{\varphi_0(u) = 1, \varphi_{2h-1}(u) = \cos(h\pi u), \varphi_{2h}(u) = \sin(h\pi u) : h = 1, 2, \dots\}$ are the eigenfunctions and $\{\mu_0 = a_0, \mu_{2h-1} = \mu_{2h} = a_h : h = 1, 2, \dots\}$ are the eigenvalues. In this example, Assumption (A.3) (i) is satisfied since the trigonometric functions are uniformly bounded by 1. The commonly used Matérn covariance function with smoothness parameter κ takes the form $C(u, u') = \frac{2^{1-\kappa}}{\Gamma(\kappa)} (\sqrt{2\kappa\theta}\|u - u'\|)^\kappa K_\kappa(\sqrt{2\kappa\theta}\|u - u'\|)$ for $u, u' \in [0, 1]^d$, where $\Gamma(\cdot)$ is the gamma function and K_κ is the modified Bessel function of the second kind. Then Assumption (A.3) (i) is satisfied for Matérn with $d = 1$ and $\kappa = 1/2$ as the eigenfunctions are again the trigonometric functions as shown in Section 3.4.1 of Van Trees (2001). It is also known in the literature that the decay rate of eigenvalues for the Matérn covariance function on $[0, 1]^d$ with smoothness parameter ν satisfies Assumption (A.3) (ii) with $\nu = \kappa + d/2$ (Ritter et al., 1995; Schaback and Wendland, 2002; Santin and Schaback, 2016).

Assumption (A.4) assumes the smoothness of the true underlying functions $\nu_{01}(\cdot), \dots, \nu_{0q}(\cdot)$. Given the eigenvalue condition in Assumption (A.3), for any given set of nonzero constants c_1, \dots, c_q , the RKHS attached to the covariance function $\sum_{a=1}^q c_a \rho_a(\cdot, \cdot)$ is norm equivalent to the ν -smooth Sobolev space on $[0, 1]^d$.

Assumption (A.5) is a technical condition that makes $\nu(\cdot)$ estimable from subset data. Similar conditions have been used in varying-coefficient modeling literature. For example, in the VCMs based on basis expansions where the dimension increases with n , the bounded eigenvalue condition in (A.5) is comparable to Condition (C1) in Wei et al. (2011) and Assumption (A5) in Bai et al. (2019), both of which have imposed bounded eigenvalue conditions on the covariance matrices involving the products of regressors and basis functions.

Given the fixed full-rank Γ_0 as in Assumption (A.2), combining subset draws of $\beta(\cdot)$ using the method described in Section 3.3 is equivalent to combining subset draws of $\nu(\cdot)$. For $a = 1, \dots, q$, let $\bar{\nu}_a(\cdot)$ be a random function drawn from the AMC posterior of $\nu_a(\cdot)$. We need the following assumption for the combination scheme in AMC.

(A.6) For any $u \in [0, 1]^d$, for each $a = 1, \dots, q$, the AMC posterior mean and variance of $\bar{\nu}_a(u)$ satisfy

$$\begin{aligned} \mathbb{E}_{\bar{\nu}_a|y,u}\{\bar{\nu}_a(u)\} &= \frac{1}{k} \sum_{j=1}^k \mathbb{E}_{\bar{\nu}_{ja}|y_j,u_j}\{\nu_{ja}(u)\} + O_p(n^{-1/2}), \\ \text{Var}_{\bar{\nu}_a|y,u}\{\bar{\nu}_a(u)\} &\leq \frac{\bar{c}}{k} \sum_{j=1}^k \text{Var}_{\bar{\nu}_{ja}|y_j,u_j}\{\nu_{ja}(u)\}, \end{aligned} \quad (20)$$

for some constant $\bar{c} > 0$, where $\mathbb{E}_{\bar{\nu}_{ja}|y_j,u_j}$ and $\text{Var}_{\bar{\nu}_{ja}|y_j,u_j}$ are the j th subset posterior mean and variance of $\nu_a(\cdot)$ given the j th subset data from (19), $\mathbb{E}_{\bar{\nu}_a|y,u}$ and $\text{Var}_{\bar{\nu}_a|y,u}$ denote the AMC posterior mean and variance of $\nu_a(\cdot)$ given the full data, and the term $O_p(n^{-1/2})$ holds uniformly over all $u \in [0, 1]^d$ and all $a = 1, \dots, q$ in the probability of the observed data.

Assumption (A.6) imposes very weak conditions on the combination method for the aggregated Bayesian posterior. It only requires that the AMC posterior mean is roughly unbiased compared to the average of subset posterior means, and the AMC posterior variance to be upper bounded by the average of subset posterior variances. These relations can be verified for many existing combination methods in the divide-and-conquer Bayes literature for parametric models. In particular, the $O_p(n^{-1/2})$ term in (20) is exactly zero in parametric models for the PIE algorithm (Li et al., 2017), the Wasserstein posterior (Xu and Srivastava, 2021), the DPMC posterior (Xue and Liang, 2019), and our proposed AMC method. Furthermore, if the model is parametric, then in all four methods, the subset posterior and the combined posterior variances satisfy $\text{Var}_{\bar{\nu}_a|y,u}\{\bar{\nu}_a(u)\} = n^{-1} \mathcal{I}_0^{-1} + o_p(n^{-1})$, where \mathcal{I}_0 is a fixed information matrix that does not depend on n ; see the theory in Li et al. (2017), Xu and Srivastava (2021) and Xue and Liang (2019). Therefore, for parametric models, the combined posterior using either of these methods can recover the exact asymptotic variance of the true posterior distribution after setting $\bar{c} = 1$ and changing the inequality to equality in (20). Based on these observations about the combined posterior means and variances, it is expected that in the VCM setup, the rates of convergence of the AMC, Wasserstein, and DPMC posterior distributions to the true posterior distribution should be similar to each other.

The following theorem is our main result on the convergence rate in L_2 norm of the AMC posterior distribution for the varying coefficients and the mean regression function; see Appendix C for the proof. Although the convergence results are presented for the AMC posterior distribution, they are not unique to the AMC posterior. Rather, they hold for any other combined posterior distribution that is built under the three step framework described in Sections 3.1, 3.2 and 3.3, with the combination step following Assumption (A.6), including the Wasserstein posterior and the DPMC posterior.

Theorem 1 *Suppose that Assumptions (A.1)–(A.6) hold for the VCM in (19). Let $\bar{\beta}(\cdot) = \alpha_0 + \Gamma_0 \bar{\nu}(\cdot)$ and $\bar{w}(\cdot) = Z(\cdot) \bar{\beta}(\cdot)$, where $\bar{\nu}(\cdot) = \{\bar{\nu}_1(\cdot), \dots, \bar{\nu}_q(\cdot)\}^T$ is a q -variate random function drawn from the AMC posterior of $\nu(\cdot)$. Let \mathbb{E}_{u^*} , $\mathbb{E}_{y,u}$, and $\mathbb{E}_{\bar{\beta}|y,u}$ be the expectations with respect to the distribution of testing point u^* , the true data generating distribution (with randomness from y_{ji} and u_{ji} , $j = 1, \dots, k$ and $i = 1, \dots, m$), and the AMC posterior distribution of the varying coefficients $\bar{\beta}(\cdot)$ given the full data.*

(i) *If $\lambda_n = 1$ and $m \gtrsim n^{(d/v)+\eta}$ for some constant $\eta \in (0, \frac{v-d}{v}]$, then the AMC posterior satisfies*

$$\begin{aligned} \mathbb{E}_{u^*} \mathbb{E}_{y,u} \mathbb{E}_{\bar{\beta}|y,u} \|\bar{\beta}(u^*) - \beta_0(u^*)\|_2^2 &\lesssim n^{-(2v-d)/(2v)}, \\ \text{and } \mathbb{E}_{u^*} \mathbb{E}_{y,u} \mathbb{E}_{\bar{\beta}|y,u} \{\bar{w}(u^*) - w_0(u^*)\}^2 &\lesssim n^{-(2v-d)/(2v)}, \end{aligned}$$

(ii) *If $\lambda_n \asymp n^{d/(2v+d)}$ and $m \gtrsim n^{2d/(2v+d)+\eta}$ for some constant $\eta \in (0, \frac{2v-d}{2v+d}]$, then the AMC posterior satisfies*

$$\begin{aligned} \mathbb{E}_{u^*} \mathbb{E}_{y,u} \mathbb{E}_{\bar{\beta}|y,u} \|\bar{\beta}(u^*) - \beta_0(u^*)\|_2^2 &\lesssim n^{-2v/(2v+d)}, \\ \text{and } \mathbb{E}_{u^*} \mathbb{E}_{y,u} \mathbb{E}_{\bar{\beta}|y,u} \{\bar{w}(u^*) - w_0(u^*)\}^2 &\lesssim n^{-2v/(2v+d)}. \end{aligned}$$

Theorem 1 gives the upper bounds for the posterior convergence rates of both the q -dimensional varying coefficients and the mean regression function in the L_2 norm. To the best of our knowledge, such convergence result for varying coefficients $\beta(\cdot)$ is new in the literature of Bayesian varying-coefficient models with multivariate latent GPs. Theorem 1 provides theoretical guarantees for the distributed extension of VCM proposed in Gelfand et al. (2003), which is developed in Section 3 and scales to massive data settings.

When the tuning parameter λ_n is chosen appropriately as in Theorem 1 (ii), the AMC posterior of the varying coefficients $\beta(\cdot)$ converges to the underlying truth in the L_2 norm at the rate $n^{-v/(2v+d)}$. This rate is known as the minimax optimal posterior convergence rate in the L_2 norm for the simple Gaussian process regression (van der Vaart and van Zanten, 2011). In the extreme case of $m = n$ and $k = 1$, Theorem 1 also implies that the rate $n^{-v/(2v+d)}$ is the convergence rate of the full data posterior distribution of $\beta(\cdot)$; therefore, we have shown that the AMC posterior from our distributed Bayesian method can quantify the posterior uncertainty in the same order as the full data posterior. If tuning from λ_n is not available (that is, $\lambda_n = 1$), then Theorem 1 (i) shows that the AMC posterior converges in the L_2 norm at least at the rate $n^{-(2v-d)/(4v)}$, which is slightly slower than the optimal rate in part (ii). Furthermore, Theorem 1 also gives sufficient conditions for the subset size m in the two scenarios. For part (i), the order of m is $m \gtrsim n^{(d/v)+\eta}$, which is meaningful when $v > d$ since $m \leq n$. Since Assumption (A.1) says that mk and n have the same order, this implies that the number of subsets k can increase no faster than $n^{(v-d)/v}$. For part (ii), the order of m is $m \gtrsim n^{2d/(2v+d)+\eta}$ and this works for all $v > d/2$ as in Assumption (A.3). As a result, the number of subsets k can increase no faster than $n^{(2v-d)/(2v+d)}$.

Our convergence rates in Theorem 1 are also comparable to similar theoretical results on distributed Bayesian inference in non-parametric regression models using an univariate GP (without a VCM formulation). This includes the recent works of Guhaniyogi et al. (2017) and Szabó and van Zanten (2019). While Bai et al. (2019) have also shown the posterior

contraction rates for the varying coefficients, their Bayesian model is based on basis series expansion instead of multivariate latent GPs as in Gelfand et al. (2003) and our current paper. Furthermore, their main focus is on high dimensional variable selection, which is different from the big n problem considered here.

5. Experiments

This section evaluates the performance of methods based on the divide-and-conquer technique for inference and predictions in Bayesian VCMs using a simulation study and a real data analysis. The simulation settings, including the details of data generation, competing methods and the metrics for comparison are described in the first subsection. The second subsection presents simulation results for different methods for a comprehensive comparison. The third subsection presents an application of the Bayesian spatiotemporal VCM to a large dataset of sea surface temperature and salinity in the North Atlantic Ocean.

5.1 Setup

Data generation:

To assess performance of distributed methods, we design two simulation studies, referred to as *Simulation 1* and *Simulation 2*, with $n = 3000$ and $n = 9000$ samples, respectively. The sample size in *Simulation 1* is moderately large to ensure that posterior computation of VCMs using the full data, although exorbitantly slow, are tractable and their results serve as the benchmark. In both simulations, the cardinality of the set of indexes \mathcal{U}^* , where the function estimation and prediction are evaluated, is set at 300. Our simulation studies consider $d = 2$, with sample indices u_1, \dots, u_n , and the indices in the set \mathcal{U}^* , u_1^*, \dots, u_{300}^* are simulated independently from the uniform distribution on $[0, 1]^2$. Both simulations assume $p = 3$ predictors, with all p predictors have varying coefficients, i.e., $q = p = 3$. We simulate a bivariate response function at all indices (i.e., $s_i = 2$) using the varying coefficient model (4) as

$$y(u) = X(u)\beta_0(u) + \epsilon(u), \quad \beta_0(\cdot) = \alpha_0 + \Gamma_0\nu(\cdot), \quad u \in \{u_1, \dots, u_n, u_1^*, \dots, u_{300}^*\}, \quad (21)$$

where $X(u)$ s are 2×3 predictor matrices at each index, with each of their entries is independently simulated from $N(0, 1)$. To construct the varying coefficients, entries of the 3×3 matrix Γ_0 are independently simulated from $\text{uniform}(0, 3)$ and α_0^T is fixed at $(-2, 2, -2)$. The components $\nu_1(\cdot), \nu_2(\cdot), \nu_3(\cdot)$ of the LMC coefficient vector $\nu(\cdot) = (\nu_1(\cdot), \nu_2(\cdot), \nu_3(\cdot))^T$, are drawn from independent GPs with 0 mean and correlation functions $\rho_a(u, u') = e^{-\phi_a \|u - u'\|_2}$, where $\phi_a = a$ for $a = 1, 2, 3$. $\epsilon(u)$ s are idiosyncratic errors following i.i.d. $N(0, \tau^2)$. Both simulations set the error variance τ^2 at 0.1. Each simulation is replicated ten times.

Competing methods:

We compare the performance of AMC algorithm with two sets of competitors. The first set of competitors include distributed Bayesian methods which follow the same three step algorithm as AMC, with the main difference appearing in the third step involving the subset posterior combination. As part of our comparison endeavor with such distributed Bayesian methods, we include DPMC, PIE, WASP and CMC algorithms as competitors.

For each of these competitors, we first create k subsets of sizes $m = 500$ and $m = 1000$, using subsampling without replacement in both simulations and vary k as 10, 20 and 30, 60 in *Simulations 1* and *2*, respectively. Second, we use the DA-type algorithm developed in Section 3.2 in parallel to obtain posterior samples of τ^2 , $\beta(u^*)$, and $y(u^*)$ for $u^* \in \mathcal{U}^*$ from all the subsets. The sampling algorithm uses a sparse GP based on the FITC approximation with $r = 400$ inducing points (Quiñonero-Candela and Rasmussen, 2005; Álvarez et al., 2012). The imputation of \tilde{v}_j in part (a) of our DA-type algorithm is done using the Lanczos algorithm of Chow and Saad (2014) for computational tractability. The sampling algorithm in each subset runs for 10,000 iterations, and the Markov chain is thinned by collecting every fifth posterior sample after discarding the first 5,000 posterior samples as burn-in. Finally, we combine subset posterior samples for τ^2 , $\beta(u^*)$, and $y(u^*)$ for $u^* \in \mathcal{U}^*$ using CMC, DPMC, WASP, PIE and the AMC algorithms (described in Section 3.3). The combination steps of AMC, DPMC, WASP, and PIE satisfy the Assumption (A.6) in Section 4, so we expect similar empirical performance for these four methods. In contrast, no such theoretical guarantee exists for the performance of the CMC algorithm.

It is also instructive to compare performance of AMC posterior with the second set of competitors which include the Bayesian VCMs on the full data. To this end, we compare with the true posterior distribution computed using full data, which sets the performance benchmark for the distributed methods. Since there is no open source implementation available for the Bayesian VCMs with bivariate response vector, we implement it by ourselves following the DA-type algorithm discussed in Section 2.3. Additionally, Finley and Banerjee (2020) offer `spSVC` in the `spBayes` R package for fitting spatial VCMs, which are special cases of (1) with $d = 2$, and fits to our simulation settings. Unfortunately, the current software support is limited to univariate responses in VCMs; hence we implement `spSVC` function marginally on each component of the bivariate response vector, and refer to this competitor as `spSVC`. Ignoring correlation between the two components in the bivariate response will presumably lead to a loss in inferential accuracy of `spSVC` compared to the other competitors. Both the DA-type algorithm of Section 2.3 and implementation of `spSVC` are prohibitively slow when $n = 9000$, so we present their results only when $n = 3000$.

Comparison metrics:

The point estimation of the varying coefficients and predictions at \mathcal{U}^* from all methods are compared using mean square error (MSE), and mean square prediction error (MSPE), respectively. Further, the coverage and length of 95% credible and predictive intervals (CIs, PIs) from the competing methods help assessing uncertainty in function estimation and in prediction, respectively. Let $\beta_0(u^*) = \{\beta_{01}(u^*), \dots, \beta_{0p}(u^*)\}$, $y(u^*) = \{y_1(u^*), \dots, y_s(u^*)\}$ be the true values of $\beta(\cdot)$, $y(\cdot)$ at u^* , where p and s are their dimensions and $u^* \in \mathcal{U}^* \subset [0, 1]^d$. Let $\hat{\beta}(u^*)$, $\hat{y}(u^*)$ be the posterior means of $\beta(u^*)$, $y(u^*)$, respectively. Then, the MSE in estimating $\beta(\cdot)$ and MSPE in predicting $y(\cdot)$ are defined as

$$\begin{aligned} \text{MSE} &= \frac{1}{|\mathcal{U}^*|} \sum_{i=1}^{|\mathcal{U}^*|} \sum_{j=1}^p \{\hat{\beta}_j(u_i^*) - \beta_{0j}(u_i^*)\}^2, \\ \text{MSPE} &= \frac{1}{|\mathcal{U}^*|} \sum_{i=1}^{|\mathcal{U}^*|} \sum_{j=1}^s \{\hat{y}_j(u_i^*) - y_j(u_i^*)\}^2. \end{aligned} \tag{22}$$

We evaluate point-wise coverage and length of the 95% CIs and PIs obtained from the posterior distributions of $\beta(\cdot)$ and $y(\cdot)$, respectively, for every $u \in \mathcal{U}^*$, for all competitors. As discussed before, the simulation settings assume $s = 2$, $p = 2$, and $|\mathcal{U}^*| = 300$. Let ESS_{DA} be the effective sample size of any DA-type algorithm that runs for T_{DA} hours, where DA can signify any of the competitors discussed above. Then, following Johndrow et al. (2019), we define the computational efficiency of a DA-type algorithm, including spSVC, AMC, CMC, DPMC, WASP, or the full data posterior (referred to as the true posterior) as

$$\text{Computational Efficiency}_{\text{DA}} = \log_2 \text{ESS}_{\text{DA}}/T_{\text{DA}}, \quad (23)$$

where ESS_{DA} is computed using the coda R package (Plummer, 2003). We do not compute (23) for the PIE combination algorithm since it is not designed to provide MCMC samples for the parameters.

5.2 Simulated Data Analysis

Table 1 and 2 show the performance of all methods in terms of estimating the true varying coefficient $\beta_0(\cdot)$ and prediction of $y(\cdot)$, respectively. As expected, the empirical performance of combined posterior obtained using AMC, WASP, PIE, and DPMC are very similar in terms of all the comparison metrics. Specifically, for both *Simulation 1* and 2, they yield similar MSE, MSPE, close to nominal coverage and similar length of 95% CIs and PIs, which validates Theorem 1 empirically. The true posterior being the gold standard, achieves little lower MSE and a bit narrower 95% CIs and PIs than its distributed competitors; however, the computational efficiency of the true posterior is much smaller than that of AMC, WASP, PIE, and DPMC because it requires much longer to finish an iteration compared to its divide-and-conquer competitors. Increasing the size of each subset m leads to better inference, where as varying the number of subsets k with a fixed value of m does not seem to have much impact on the inference. Among the distributed methods, the CI and PI lengths are slightly larger for PIE, perhaps due to the marginal combination of subset posterior distributions. Although WASP shows marginally narrower 95% CIs and PIs compared to AMC and DPMC (all maintaining close to the nominal coverage), the subset posterior combination step of AMC and DPMC are much more computationally convenient. On the other hand, the CI and PI lengths of CMC are very small compared to that of the WASP, which results in poor coverage for every m and n and deteriorates as k increases. Except CMC, all other methods show similar performance for inference on the error variance τ^2 (Table 3). We also find the performance of spSVC to be excellent in predicting $y(u^*)$ s, but becomes extremely poor in inference on $\beta(u^*)$ s. The poor performance of spSVC in inference on $\beta(u^*)$ s is mainly because the marginal model ignores the dependence between $y_1(\cdot)$ and $y_2(\cdot)$. On the other hand, spSVC shows excellent performance in predicting $y(u^*)$ s because the marginal spSVC model still uses three GPs for predicting $y_1(\cdot)$ and $y_2(\cdot)$.

AMC, DPMC, PIE, and WASP satisfy Assumption (A.6) on the combination of subset posterior distributions, whereas CMC does not; therefore, we conclude that methods that satisfy our theoretical assumptions show superior empirical performance. Furthermore, the AMC and DPMC combination algorithms are the simplest among the distributed competitors that offer combination of subset posteriors of all parameters jointly. Hence, they are simple and computationally convenient alternatives to the full data posterior distribution in Bayesian VCMs for massive data.

Table 1: Summary of the results for inference on $\beta(\cdot)$. The CI coverage and their lengths are averaged across 10 simulation replications, dimensions, and $u \in \mathcal{U}^*$. A ‘-’ for the true posterior and spSVC corresponding to $n = 9000$ indicates that the results are missing due to intractable posterior computations. On the other hand, ‘-’ in reporting the computational efficiency for PIE is due to the lack of definition.

		$n = 3000$							
		Coverage at 95% Nominal Level				95% CI Length			
		$m = 500$		$m = 1000$		$m = 500$		$m = 1000$	
		$k = 10$	$k = 20$	$k = 10$	$k = 20$	$k = 10$	$k = 20$	$k = 10$	$k = 20$
True Posterior		0.96				2.53			
Marginal spSVC		0.27				3.22			
AMC		0.97	0.97	0.96	0.97	3.18	3.19	2.88	2.89
PIE		0.97	0.97	0.96	0.97	3.24	3.24	3.02	2.91
CMC		0.56	0.42	0.57	0.39	1.10	0.79	1.09	0.67
WASP		0.96	0.96	0.96	0.96	3.06	3.06	2.78	2.78
DPMC		0.97	0.97	0.96	0.97	3.19	3.20	2.89	2.89
		MSE				Computational Efficiency			
True Posterior		0.40				2.10			
Marginal spSVC		7.24				6.68			
AMC		0.56	0.56	0.48	0.48	9.02	8.41	8.79	8.37
PIE		0.57	0.57	0.52	0.48	-	-	-	-
CMC		0.56	0.56	0.51	0.48	4.70	4.56	4.27	2.60
WASP		0.57	0.57	0.49	0.48	9.02	8.41	8.79	8.37
DPMC		0.57	0.57	0.49	0.48	9.02	8.41	8.79	8.36
		$n = 9000$							
		Coverage at 95% Nominal Level				95% CI Length			
		$m = 500$		$m = 1000$		$m = 500$		$m = 1000$	
		$k = 30$	$k = 60$	$k = 30$	$k = 60$	$k = 30$	$k = 60$	$k = 30$	$k = 60$
True Posterior		-				-			
Marginal spSVC		-				-			
AMC		0.98	0.98	0.98	0.98	3.04	3.05	2.75	2.75
PIE		0.98	0.98	0.98	0.98	3.07	3.08	2.77	2.77
CMC		0.34	0.25	0.34	0.24	0.58	0.41	0.52	0.37
WASP		0.97	0.97	0.97	0.97	2.91	2.92	2.64	2.63
DPMC		0.98	0.98	0.98	0.98	3.05	3.06	2.75	2.75
		MSE				Computational Efficiency			
True Posterior		-				-			
AMC		0.45	0.45	0.38	0.38	10.40	9.67	9.94	9.70
PIE		0.46	0.46	0.39	0.39	-	-	-	-
CMC		0.45	0.45	0.39	0.39	5.76	4.03	5.47	4.01
WASP		0.45	0.45	0.38	0.38	10.40	9.67	9.94	9.70
DPMC		0.45	0.45	0.38	0.38	10.40	9.67	9.94	9.70

5.3 Real Data Analysis

We illustrate the performance of the combined posterior distributions obtained using AMC, DPMC, PIE, WASP or CMC combination technique for the space-time varying coefficient modeling, where the indices are $u = (h, t)$ with h and t denoting the spatial locations and time points of the response and covariates. VCMs are widely used in a variety of spatial applications, mostly without the temporal dimension; see, for example, Wheeler and Calder (2007); Finley et al. (2014); Banerjee and Johnson (2006). On the contrary, their applications in large data settings are limited, perhaps due to the demanding computations. This section specifically considers the problem of capturing the spatio-temporal association (with uncertainties) between the sea surface temperature (SST) and sea surface salinity (SSS) in the Atlantic Ocean between $0^\circ - 70^\circ$ north latitudes and $0^\circ - 80^\circ$ west longitudes using the spatiotemporal VCM. This implies that $s_i = 1$, $p = 2$, $d = 3$, and the space-time tuples lie in a fixed and bounded domain for the spatiotemporal model based on (1). The data on SST and SSS are obtained from the Hadley center observations under the met office in UK (www.metoffice.gov.uk/hadobs, more description available in Kennedy et al. (2011)). We specifically consider 72000 space-time observations on SST and SSS over the 12 months in 2018 and randomly set aside $|\mathcal{U}^*| = 600$ space-time tuples for prediction, which form

Table 2: Summary of the results for $y(\cdot)$ prediction. The PI coverage and their lengths are averaged across 10 simulation replications, dimensions $a = 1, 2$, and $u \in \mathcal{U}^*$. A ‘-’ for the true posterior and spSVC corresponding to $n = 9000$ indicates that the results are missing due to intractable posterior computations. On the other hand, ‘-’ in reporting the computational efficiency for PIE is due to the lack of definition.

$n = 3000$									
Coverage at 95% Nominal Level					95% PI Length				
$m = 500$		$m = 1000$			$m = 500$		$m = 1000$		
	$k = 10$	$k = 20$	$k = 10$	$k = 20$	$k = 10$	$k = 20$	$k = 10$	$k = 20$	
True Posterior	0.96				4.11				
Marginal spSVC	1.00				1.66				
AMC	0.97	0.97	0.96	0.96	5.05	5.05	4.61	4.62	
PIE	0.97	0.97	0.97	0.96	5.12	5.11	4.81	4.65	
CMC	0.55	0.42	0.57	0.39	1.75	1.25	1.73	1.08	
WASP	0.96	0.96	0.95	0.95	4.87	4.86	4.46	4.46	
DPMC	0.97	0.97	0.96	0.96	5.06	5.06	4.61	4.63	
MSPE					Computational Efficiency				
True Posterior	1.30				2.10				
Marginal spSVC	0.03				6.67				
AMC	1.79	1.76	1.53	1.54	9.02	8.41	8.79	8.37	
PIE	1.82	1.79	1.67	1.55	-	-	-	-	
CMC	1.78	1.75	1.65	1.53	5.72	4.10	5.55	4.11	
WASP	1.79	1.76	1.53	1.54	9.02	8.41	8.79	8.37	
DPMC	1.79	1.76	1.53	1.54	9.02	8.41	8.79	8.37	
$n = 9000$									
Coverage at 95% Nominal Level					95% PI Length				
$m = 500$		$m = 1000$			$m = 500$		$m = 1000$		
	$k = 30$	$k = 60$	$k = 30$	$k = 60$	$k = 30$	$k = 60$	$k = 30$	$k = 60$	
True Posterior	-				-				
Marginal spSVC	-				-				
AMC	0.98	0.98	0.97	0.97	4.80	4.82	4.39	4.38	
PIE	0.97	0.97	0.97	0.97	4.85	4.87	4.42	4.42	
CMC	0.33	0.24	0.32	0.23	0.92	0.65	0.83	0.58	
WASP	0.96	0.96	0.96	0.96	4.62	4.63	4.23	4.22	
DPMC	0.98	0.98	0.97	0.97	4.82	4.84	4.40	4.39	
MSPE					Computational Efficiency				
True Posterior	-				-				
Marginal spSVC	-				-				
AMC	1.41	1.40	1.22	1.21	10.40	9.67	9.94	9.70	
PIE	1.45	1.45	1.25	1.24	-	-	-	-	
CMC	1.40	1.40	1.23	1.23	5.76	4.03	5.47	4.02	
WASP	1.41	1.40	1.22	1.21	10.40	9.67	9.94	9.70	
DPMC	1.41	1.40	1.22	1.21	10.40	9.67	9.94	9.70	

Table 3: The 95% credible intervals for inference on τ^2 . The lower and upper ends of CIs are averaged across 10 simulation replications.

$n = 3000$				
$m = 500$		$m = 1000$		
	$k = 10$	$k = 20$	$k = 10$	$k = 20$
True Posterior	(0.0891, 0.1067)			
Marginal spSVC	(0.0975, 0.1141)			
AMC	(0.0707, 0.103)	(0.0695, 0.1016)	(0.0778, 0.1323)	(0.0798, 0.1041)
PIE	(0.0717, 0.1037)	(0.0703, 0.1021)	(0.0821, 0.1316)	(0.0801, 0.1042)
CMC	(0.0799, 0.0909)	(0.0802, 0.0881)	(0.0878, 0.0982)	(0.0884, 0.094)
WASP	(0.0708, 0.1029)	(0.0696, 0.1016)	(0.0799, 0.1268)	(0.0799, 0.1041)
DPMC	(0.0707, 0.1031)	(0.0695, 0.1017)	(0.0778, 0.1347)	(0.0798, 0.1041)
$n = 9000$				
$m = 500$		$m = 1000$		
	$k = 30$	$k = 60$	$k = 30$	$k = 60$
True Posterior	-			
Marginal spSVC	-			
AMC	(0.0709, 0.103)	(0.0704, 0.1022)	(0.0799, 0.1038)	(0.0795, 0.1031)
PIE	(0.0718, 0.1037)	(0.0713, 0.1029)	(0.0803, 0.1041)	(0.0799, 0.1034)
CMC	(0.0819, 0.0879)	(0.0821, 0.0864)	(0.0887, 0.0932)	(0.0888, 0.0919)
WASP	(0.071, 0.1029)	(0.0704, 0.1021)	(0.08, 0.1038)	(0.0795, 0.1031)
DPMC	(0.0709, 0.103)	(0.0704, 0.1022)	(0.0799, 0.1037)	(0.0795, 0.1031)

the set \mathcal{U}^* of size 600. Full scale Bayesian inference of spatio-temporal VCMs with data at this scale is extremely challenging and has been sparsely dealt with in the literature.

The global association between SST and SSS is well established in the fields of Oceanography and Geophysics (Millero et al., 1998; Key et al., 2004; Lee et al., 2006). In fact, salinity influences the depth to which water masses sink and how far they extend through the ocean. The location and depth of these water masses controls how heat are transported between the tropics and high latitudes. Both SST and SSS are also key in understanding how oceans interact with the atmosphere. Monsoons are driven by exchanges at the air-ocean boundary, affecting almost half of the world’s human population each year. Likewise, El Niño has profound effects on humankind and is, to an unknown extent, governed by ocean salinity and temperature. Earlier work with nonlinear regression models to ascertain relationships between SST and SSS (Xiong et al., 2013; Becker and Pauly, 1996) reveal significant positive association between these two climate indicators. Although some of these prior studies reveal such associations to be spatially varying (Weldeab et al., 2006), there is still a dearth of model based analysis of spatio-temporally varying associations between SST and SSS.

We compare performance of the combined posterior obtained using AMC, DPMC, WASP, PIE combination schemes (all following theoretically guaranteed optimal performance) along with the other distributed competitor CMC, popularly used in the machine learning literature for distributed inference with massive data. We have also attempted to fit spatially varying coefficient model on the full data using the `spSVC` function in the `spBayes` package in R; however, full data posterior computations using `spSVC` fails due to the large sample size. Gneiting’s correlation function (6) is employed in the spatiotemporal VCM due to its flexibility in modeling space-time correlations (Gneiting, 2002). The values of k and m are set to be 400 and 2500, respectively, and the results for the distributed methods follow from three step strategy described in Section 3. Because the true varying coefficients are unknown, we only make assessment of point prediction and predictive uncertainties for all the methods using MSPE and coverage of 95% PIs for the space-time tuples in \mathcal{U}^* , respectively. Computational efficiency of all methods are also reported.

Figure 1 presents the posterior mean of the spatially varying coefficient corresponding to SSS in January, May and September for AMC, PIE, WASP and DPMC. From the equator to the pole, the annual excess precipitation over evaporation increases, and thus salinity decreases along with SST, with latitude. However, in lower latitude, due to the pronounced salt accumulation as a result of excess heating and oceanic currents, SSS surges, which results in lower $\beta_1(s)$ values. This trend becomes more prominent during the months of summer or fall (columns 2 and 3). In general, SSS decreases in comparison with SST during winter, except for the Brazilian coast, which shows lower coefficient values even in winter due to the strong North Brazil Current (Weldeab et al., 2006). The increase in latitude shows a considerable drop of SSS compared to SST leading to higher $\beta_1(s)$ values. The estimates appear to be consistent over all the four combination approaches (AMC, WASP, PIE and DPMC) following the three step algorithm.

Turning our attention to the predictive inference, Table 4 demonstrates comparable point estimation and predictive uncertainties from AMC, PIE, WASP and DPMC combination schemes. In contrast, the machine learning competitor CMC shows high MSPE and considerably wider credible intervals at all space-time tuples. The computationally efficiency metric for other methods also supercede CMC by a large margin. As a result, the space-time varying coefficient figures corresponding to CMC also appear to be different from the other competitors, and hence it has not been included under Figure 1. As a whole,

the data analysis reinforces our findings on the three step divide and conquer approaches with the theoretically guaranteed combination schemes AMC, DPMC, WASP and PIE as simple, computationally efficient, flexible, and fully Bayesian inferential tools for inference in large spatiotemporal data with the VCM model. We emphasize that the scalability of all these approaches depend on the Bayesian VCM model fitted in each subset. Using more computationally efficient variants of GPs for $\nu(\cdot)$ in each subset, a much higher degree of scalability is achievable.

Table 4: Summary of the results for prediction of sea surface temperature at 600 space-time tuples using the spatio-temporal VCM fitted at each subset. The PI coverage and their lengths are based on point-wise 95% predictive intervals and are averaged across the 600 space-time tuples. Computational efficiency for all methods are also reported. Omission of computational efficiency for PIE is due to the lack of definition.

	Coverage	MSPE	95% PI Length	Computational Efficiency
AMC	0.99	2.92	6.60	9.99
PIE	0.99	2.93	5.61	-
CMC	0.90	74.95	24.71	1.06
WASP	0.98	2.92	5.26	9.99
DPMC	0.99	2.92	6.53	10.09

6. Discussion

Bayesian varying coefficient models with multivariate Gaussian process prior on varying coefficients are extremely popular in functional regression models with a wide variety of applications, since they combine the flexibility of a nonparametric model and interpretability of a linear regression model. Unfortunately, full scale Bayesian inference with VCMs are relatively less explored with big data, due to the computation and storage becoming prohibitively burdensome. This article proposes a three step distributed framework that divides the data into (possibly) overlapping subsets, fits posterior inference with VCM in each subset after appropriately modifying the subset likelihood, and finally aggregates inference from subsets with a combination algorithm to derive a pseudo posterior distribution that approximates the full posterior. This article is presumably the first approach to design a principled distributed Bayesian algorithm on VCMs with large data. Additionally, we claim threefold contribution in the literature of distributed Bayesian inference on VCMs. First, we identify the modification required in subset posterior likelihoods and justify such modification with rigorous theoretical guarantee. Second, a new subset posterior combination algorithm, referred to as the Aggregated Monte Carlo algorithm, is proposed. Unlike a few other subset posterior combination algorithms proposed in the context of univariate Gaussian process regression models, AMC jointly combines subset posterior of all parameters and yet offers straightforward implementation. Finally, the major contribution of the article becomes theoretically establishing minimax optimal convergence rate for the varying coefficients and regression mean function under the three stage distributed Bayesian framework. The theoretical results explicitly offer the choices of k and m as functions of n , smoothness of the fitted Gaussian process and the smoothness of the true varying co-

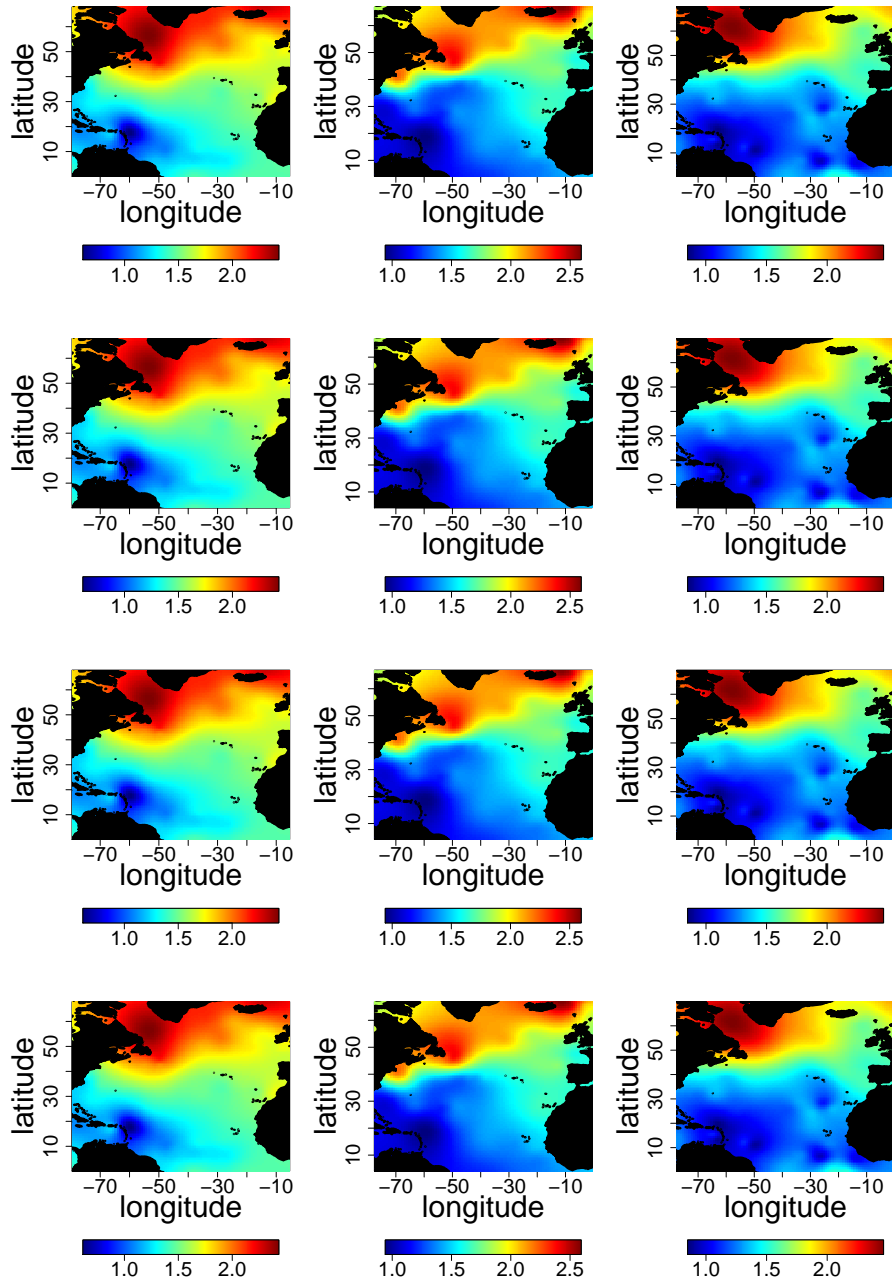


Figure 1: The four rows represent space-time varying coefficient corresponding to the STVC regression model with SST as the response and SSS as the predictor. Rows 1-4 correspond to the coefficient maps from AMC, WASP, PIE and DPMC respectively. In each row, columns 1, 2, 3 present coefficient maps in January, May and September respectively.

efficients to guarantee optimal inference. Furthermore, the theoretical results are proved under mild assumption on the subset posterior combination algorithm, which is found to hold for DPMC, WASP and PIE, along with the proposed AMC algorithm. Simulation studies demonstrate similar point estimation and prediction, along with close to nominal coverage for all the distributed methods that satisfy theoretical assumptions. We fit space-time varying coefficient models to delineate local variability, as well as seasonal variability in the relationship between SST and SSS in the northern Atlantic ocean.

As a first attempt to principled distributed inference with VCMs, we employ FITC approximation of the Gaussian process for each coefficient in every subset inference and are able to seamlessly scale Bayesian VCMs for $\sim 10^5$ observations, even with moderate dimensional multivariate varying coefficients. As an immediate future work, we plan to fit a more computationally convenient approximation to the GPs on varying coefficients (Gramacy and Apley, 2015) to ameliorate scalability. On the theoretical front, our proof techniques do not depend the normality of error terms in the VCM model. Our conjecture is that if we assume the normal error assumption, then by using the techniques in van der Vaart and van Zanten (2011), it is to possible to improve the rate in Theorem 1 to an adaptive minimax optimal rate $n^{-v/(2v+d)}$ without the tuning parameter λ_n , and to relax the function space from RKHS in Assumption (A.4) to the larger space of functions that are less smooth. We leave this direction for future research.

Acknowledgements

The research of Rajarshi Guhaniyogi and Sanvesh Srivastava is partially supported by grants from the Office of Naval Research (ONR-BAA N000141812741) and the National Science Foundation (DMS-1854667/1854662). The research of Cheng Li is supported by the Singapore Ministry of Education Academic Research Funds Tier 1 grants R-155-000-201-114 and R-155-000-223-114. The four authors have contributed equally to this work and their names appear in the alphabetical order.

Appendix A. Draws from the True Posterior Distribution

Consider the linear mixed effects reformulation of the Bayesian VCM in (4) and the prior distributions on $(\alpha, \Gamma, \tau^2, \theta_1, \dots, \theta_q)$ again. Define total number of observations as $s = \sum_{i=1}^n s_i$, \tilde{Z}_a , ν_a , R_a ($a = 1, \dots, q$), and y , X , \tilde{Z} , R , ν , ϵ as follows:

$$\begin{aligned} \tilde{Z}_a &= \text{diag}\{Z(u_1)\Gamma_a, \dots, Z(u_n)\Gamma_a\}, \quad \nu_a^\top = \{\nu_a(u_1), \dots, \nu_a(u_n)\}, \quad \nu^\top = (\nu_1^\top, \dots, \nu_q^\top), \\ (R_a)_{ii'} &= \rho_a(u_i, u_{i'}), \quad i, i' = 1, \dots, n, \quad R = \text{diag}(R_1, \dots, R_q), \quad \epsilon^\top = (\epsilon_1^\top, \dots, \epsilon_n^\top), \\ y^\top &= \{y(u_1), \dots, y(u_n)\}, \quad X^\top = [X(u_1)^\top, \dots, X(u_n)^\top], \quad \tilde{Z} = [\tilde{Z}_1, \dots, \tilde{Z}_q], \end{aligned} \quad (24)$$

where \tilde{Z}_a , R_a , X , and Z are s -by- n , n -by- n , s -by- p , and s -by- nq matrices, respectively, R is a nq -by- nq block diagonal matrix, and ν , ν_a , and y are nq -by-1, n -by-1, and s -by-1 vectors, respectively. Using these definitions, the Bayesian VCM in (4) and the prior distributions are re-written as

$$\begin{aligned} y &= X\alpha + \sum_{a=1}^q \tilde{Z}_a \nu_a + \epsilon = \tilde{Z}\nu + \epsilon, \quad \epsilon \sim N(0, \tau^2 I), \quad \nu_a \sim N(0, R_a), \\ \nu_a &\sim N(0, R_a), \quad \nu \sim N(0, R), \quad p(\alpha, \Gamma, \tau^2) \propto \tau^{-2}, \quad p(\theta_a) = \text{Uniform}(\underline{c}_a, \bar{c}_a), \end{aligned} \quad (25)$$

where the uniform distribution of θ_a is assumed to be component-wise if θ_a is a vector. In this case, \underline{c}_a and \bar{c}_a are also vectors of the same dimension as θ_a . The prior distribution on the latent variables ν and parameters $\alpha, \Gamma, \tau^2, \theta_1, \dots, \theta_q$ are assumed to have the form

$$p(\nu, \alpha, \Gamma, \tau^2, \theta_1, \dots, \theta_q) = \prod_{a=1}^q \{p(\nu_a | \theta_a) p(\theta_a)\} p(\alpha, \Gamma, \tau^2) \equiv p(\nu, \theta_1, \dots, \theta_q) p(\alpha, \Gamma, \tau^2). \quad (26)$$

Our sampling algorithm for drawing $\beta(\cdot)$, $y(\cdot)$, and τ^2 from their respective full data posterior distributions is based on (25) and (26).

First, we derive the full conditional of ν . Assume that $y, \alpha, \Gamma, \tau^2, \theta_1, \dots, \theta_q$ are given. Using (25), the joint distribution of (y, ν) is an $(s + nq)$ -variate Gaussian distribution with mean $(X\alpha, 0)$, where 0 is an nq -by-1 vector, and covariance matrix \bar{C} , where the blocks corresponding to the marginal covariance matrices of y , ν , respectively, and their cross covariance matrix are

$$(\bar{C})_{yy} = \tilde{Z}R\tilde{Z}^\top + \tau^2 I = \sum_{a=1}^q \tilde{Z}_a R_a \tilde{Z}_a^\top + \tau^2 I, \quad (\bar{C})_{\nu\nu} = R, \quad (\bar{C})_{y\nu} = \tilde{Z}R. \quad (27)$$

This implies that ν given $y, \alpha, \Gamma, \tau^2, \theta_1, \dots, \theta_q$ follows $N(\mu_\nu, \Sigma_\nu)$, where

$$\mu_\nu = R^\top \tilde{Z}^\top (\bar{C})_{yy}^{-1} (y - X\alpha), \quad \Sigma_\nu = R - R^\top \tilde{Z}^\top (\bar{C})_{yy}^{-1} \tilde{Z}R. \quad (28)$$

Second, we derive the full conditional of (α, Γ, τ^2) . Assuming y, ν are given, define

$$W_i = [X(u_i), \nu^\top(u_i) \otimes Z(u_i)], \quad i = 1, \dots, n, \quad W^\top = [W_1^\top, \dots, W_n^\top], \quad b^\top = (\alpha^\top, \gamma^\top), \quad (29)$$

where γ is the column-wise vectorization of Γ . Rewrite the Bayesian VCM in (25) as

$$y = Wb + \epsilon, \quad \epsilon \sim N(0, \tau^2 I), \quad p(b, \tau^2) \propto \tau^{-2}. \quad (30)$$

If $\hat{b} = (W^T W)^{-1} W^T y$ is the least squares estimate of b and $\hat{y} = W\hat{b}$ is the mean estimate of y based on \hat{b} , then

$$\begin{aligned} p(\tau^2, b | y) &\propto \frac{1}{(\tau^2)^{s/2}} \exp \left\{ -\frac{(y - Wb)^T (y - Wb)}{2\tau^2} \right\} \frac{1}{\tau^2} \\ &= \frac{1}{(\tau^2)^{s/2+1}} \exp \left\{ -\frac{(y - \hat{y})^T (y - \hat{y}) + (\hat{y} - Wb)^T (\hat{y} - Wb)}{2\tau^2} \right\} \\ &= \frac{1}{(\tau^2)^{s/2+1}} \exp \left(-\frac{\|y - \hat{y}\|_2^2}{2\tau^2} \right) \exp \left\{ -\frac{(\hat{b} - b)^T (W^T W)(\hat{b} - b)}{2\tau^2} \right\}. \end{aligned} \quad (31)$$

Marginalizing over b in (31) implies that

$$p(\tau^2 | y) \propto \frac{1}{(\tau^2)^{(s-p-q^2)/2+1}} \exp \left(-\frac{\|y - \hat{y}\|_2^2}{2\tau^2} \right), \quad \tau^2 | y, \nu \sim \frac{\|y - \hat{y}\|_2^2}{\chi_{s-p-q^2}^2}, \quad (32)$$

and (31) assuming τ^2 is given implies that

$$p(b | y, \tau^2) \propto \exp \left\{ -\frac{(\hat{b} - b)^T (W^T W)(\hat{b} - b)}{2\tau^2} \right\}, \quad b | y, \nu, \tau^2 \sim N \left\{ \hat{b}, \tau^2 (W^T W)^{-1} \right\}; \quad (33)$$

therefore, (32) and (33) imply that the distribution of (b, τ^2) given y lies in the Normal-Inverse-Gamma family. The α draw corresponds to the first p elements of b and the remaining elements of b are “un-vectorized” into the q -by- q matrix Γ .

Finally, we derive the ESS algorithm for drawing $\theta_1, \dots, \theta_q$ given ν . Using (26), we have that

$$\log p(\theta_1, \dots, \theta_q | \nu) \propto -\frac{1}{2} \sum_{a=1}^q \log |R_a(\theta_a)| - \frac{1}{2} \sum_{a=1}^q \nu_a^T R_a^{-1}(\theta_a) \nu_a + \sum_{a=1}^q \log 1_{\underline{c}_a \leq \theta_a \leq \bar{c}_a}, \quad (34)$$

but ESS cannot be applied directly for sampling $\theta_1, \dots, \theta_q$ given ν due the range restrictions on θ_a s imposed by the uniform prior distributions. We address this problem by first transforming $\theta_1, \dots, \theta_q$ to $\bar{\theta}_1, \dots, \bar{\theta}_q$ as

$$\bar{\theta}_a = \log \frac{\theta_a - \underline{c}_a}{\bar{c}_a - \theta_a}, \quad \theta_a = \underline{c}_a + \frac{\bar{c}_a - \underline{c}_a}{1 + e^{-\bar{\theta}_a}}, \quad (35)$$

where each $\bar{\theta}_a \in (-\infty, \infty)$ and the mapping of θ_a to $\bar{\theta}_a$ is done component-wise if θ_a is a vector. The form is $\log p(\theta_1, \dots, \theta_q | \nu)$ in (34) is modified using the Jacobian of the transform as

$$\begin{aligned} \log p(\bar{\theta}_1, \dots, \bar{\theta}_q | \nu) &\propto -\frac{1}{2} \sum_{a=1}^q \log |R_a(\bar{\theta}_a)| - \frac{1}{2} \sum_{a=1}^q \nu_a^T R_a^{-1}(\bar{\theta}_a) \nu_a + \sum_{a=1}^q \log(\bar{c}_a - \underline{c}_a) \\ &\quad + \sum_{a=1}^q \log \bar{\theta}_a - 2 \sum_{a=1}^q \log(1 + e^{\bar{\theta}_a}). \end{aligned} \quad (36)$$

The $\bar{\theta}_a$ s in (36) are supported on $(-\infty, \infty)$, so we apply the ESS algorithm using the proposal $N(0, 4I)$ for $(\bar{\theta}_1, \dots, \bar{\theta}_q)$, where I is an identity whose dimension is determined by that of $(\theta_1, \dots, \theta_q)$; see Algorithm 1 in Nishihara et al. (2014).

In summary, the sampling algorithm for drawing from the posterior distribution of $(\alpha, \Gamma, \tau^2, \theta_1, \dots, \theta_q)$ starts from an initial value of parameters $(\alpha^{(0)}, \Gamma^{(0)}, \tau^{2(0)}, \theta_1^{(0)}, \dots, \theta_q^{(0)})$ and cycles through the following four steps for $t = 0, 1, \dots, \infty$:

- (a) draw $\nu^{(t+1)}$ given $y, \alpha^{(t)}, \Gamma^{(t)}, \tau^{2(t)}, \theta_1^{(t)}, \dots, \theta_q^{(t)}$ from $N(\mu_\nu, \Sigma_\nu)$, where μ_ν, Σ_ν are defined in (28);
- (b) draw $\tau^{2(t+1)}$ given $y, \nu^{(t+1)}$ using (32);
- (c) draw $(\alpha^{(t+1)}, \gamma^{(t+1)})$ given $y, \nu^{(t+1)}, \tau^{2(t+1)}$ using (33) and the vectorization of $\gamma^{(t+1)}$ is reversed to obtain $\Gamma^{(t+1)}$; and
- (d) draw $\theta_1^{(t+1)}, \dots, \theta_q^{(t+1)}$ given $\nu^{(t+1)}$ using ESS, the likelihood in (36), and the relation between $\bar{\theta}_a$ and θ_a in (35).

In practice, the interest also lies in drawing $\beta(u^*)$ and $y(u^*)$ for $u^* \in \mathcal{U}^*$, where \mathcal{U}^* is a known subset of $[0, 1]^d$ with l elements. This is accomplished by the addition of two extra steps after steps (a)–(d). Let $\nu_a^* = \{\nu_a(u_1^*), \dots, \nu_a(u_l^*)\}$ ($a = 1, \dots, q$) and $\nu(u^*) = \{\nu_1(u^*), \dots, \nu_q(u^*)\}$, $u^* \in \mathcal{U}^*$. Then, the GP prior on $\nu_a(\cdot)$ ($a = 1, \dots, q$) implies that the ν_a^* given ν_a and θ_a is drawn from $N(\mu_a^*, \Sigma_a^*)$, where

$$\begin{aligned} \mu_a^* &= R_{a**}^\top R_a^{-1} \nu_a, \quad \Sigma_a^* = R_{a**} - R_{a**}^\top R_a^{-1} R_{a**}, \quad a = 1, \dots, q, \\ (R_{a**})_{ii'} &= \rho_a(u_i, u_{i'}), \quad (R_{a**})_{i'i''} = \rho_a(u_i^*, u_{i''}^*), \quad i = 1, \dots, n, \quad i', i'' = 1, \dots, l. \end{aligned} \quad (37)$$

Given $(\alpha, \Gamma, \tau^2, \theta_1, \dots, \theta_q)$ and $u^* \in \mathcal{U}^*$, the draws of $\beta(u^*)$ is obtained as

$$\beta(u^*) = [\{\beta(u^*)\}_{\text{va}}^\top, \{\beta(u^*)\}_{\text{nv}}^\top]^\top, \quad \{\beta(u^*)\}_{\text{va}} = \alpha_{\text{va}} + \Gamma \nu(u^*), \quad \{\beta(u^*)\}_{\text{nv}} = \alpha_{\text{nv}}, \quad (38)$$

where nv, v correspond to the non-varying and varying coefficients indices, and the draw of $y(u^*)$ is obtained as

$$y(u^*) \sim N(\mu_{y^*}, \tau^2 I), \quad \mu_{y^*} = X(u^*) \beta(u^*); \quad (39)$$

therefore, at the $(t+1)$ iteration of the full data sampling algorithm, the following two steps are added if \mathcal{U}^* and $X(u^*)$ for every $u^* \in \mathcal{U}^*$ are known:

- (e) draw $\nu_a^{*(t+1)}$ given $\nu_a^{(t+1)}, \theta_a^{(t+1)}$ using (37) ($a = 1, \dots, q$) and obtain $\beta(u^*)^{(t+1)}$ given $\alpha^{(t+1)}, \Gamma^{(t+1)}, \nu_a^{*(t+1)}$ using (38) for every $u^* \in \mathcal{U}^*$; and
- (f) draw $y(u^*)^{(t+1)}$ given $X(u^*), \beta(u^*)^{(t+1)}, \tau^{2(t+1)}$ using (39) for every $u^* \in \mathcal{U}^*$.

We also take advantage of the low-rank structure in the covariance matrices if inducing points are used defining the covariance functions of GP (Quiñonero-Candela and Rasmussen, 2005; Álvarez and Lawrence, 2011).

Appendix B. Draws from the j th Subset Posterior Distribution

Consider the reformulation of (25) on subset j . Define total number of observations on subset j as $\tilde{s}_j = \sum_{i=1}^m s_{ji}$, \tilde{Z}_{ja} , $\tilde{\nu}_{ja}$, R_{ja} , y_j , X_j , \tilde{Z}_j , R_j , $\tilde{\nu}_j$, ϵ_j are the subset j counterparts of \tilde{Z}_a , ν_a , R_a , y , X , \tilde{Z} , R , ν , ϵ , where the dimensions of subset j variables are obtained by replacing s by \tilde{s}_j and n by m in their full data counterparts. Using these definitions, the Bayesian VCM in on subset j and the prior distributions are re-written as

$$y_j = X_j \alpha_j + \tilde{Z}_{j1} \tilde{\nu}_{j1} + \cdots + \tilde{Z}_{jq} \tilde{\nu}_{jq} + \epsilon_j = \tilde{Z}_j \tilde{\nu}_j + \epsilon_j, \quad \epsilon_j \sim N(0, \tau_j^2 I), \quad \tilde{\nu}_{ja} \sim N(0, R_{ja}), \\ \tilde{\nu}_{ja} \sim N(0, R_{ja}), \quad \tilde{\nu} \sim N(0, R_j), \quad p(\alpha_j, \Gamma_j, \tau_j^2) \propto \tau_j^{-2}, \quad p(\theta_{aj}) = \text{Uniform}(\underline{c}_a, \bar{c}_a). \quad (40)$$

where the prior distributions on θ_{ja} s remain unchanged. The prior distribution of the subset j latent variables $\tilde{\nu}_j$ and parameters $\alpha_j, \Gamma_j, \tau_j^2, \theta_{j1}, \dots, \theta_{jq}$ also have the same form as in (26)

$$p(\tilde{\nu}_j, \alpha_j, \Gamma_j, \tau_j^2, \theta_{j1}, \dots, \theta_{jq}) = \prod_{a=1}^q \{p(\tilde{\nu}_{ja} | \theta_{ja}) p(\theta_{ja})\} p(\alpha_j, \Gamma_j, \tau_j^2) \\ \equiv p(\tilde{\nu}_j, \theta_{j1}, \dots, \theta_{jq}) p(\alpha_j, \Gamma_j, \tau_j^2). \quad (41)$$

The sampling algorithm for drawing $\beta(\cdot)$, $y(\cdot)$, and τ^2 from their respective subset j posterior distributions is based on a modified form of (40) and (41). The conditional density of $(\alpha_j, \Gamma_j, \tau_j^2, \theta_{j1}, \dots, \theta_{jq})$ given y_j after stochastic approximation is defined following (10) as

$$\pi(\alpha_j, \Gamma_j, \tau_j^2, \theta_{j1}, \dots, \theta_{jq} | y_j) \propto \{p(y_j | \tilde{\nu}_j, \alpha_j, \Gamma_j, \tau_j^2) p(\tilde{\nu}_j, \theta_{j1}, \dots, \theta_{jq})\}^{\delta_n} p(\alpha_j, \Gamma_j, \tau_j^2), \quad (42)$$

where $p(\tilde{\nu}_j, \theta_{j1}, \dots, \theta_{jq})$ is the subset j version of $p(\nu, \theta_1, \dots, \theta_q)$ in (26). The model in (40) implies that the joint distribution of $(y_j, \tilde{\nu}_j)$ given $\alpha_j, \Gamma_j, \tau_j^2, \theta_{j1}, \dots, \theta_{jq}$ is an $(\tilde{s}_j + mq)$ -variate Gaussian distribution with mean $(X_j \alpha_j, 0)$, where 0 is an mq -by-1 vector, and covariance matrix \bar{C}_j , where the blocks corresponding to the marginal covariance matrices of $y_j, \tilde{\nu}_j$, respectively, and their cross covariance matrix are

$$(\bar{C}_j)_{y_j y_j} = \tilde{Z}_j R_j \tilde{Z}_j^T + \tau_j^2 I = \sum_{a=1}^q \tilde{Z}_{ja} R_{ja} \tilde{Z}_{ja}^T + \tau_j^2 I, \quad (\bar{C}_j)_{\tilde{\nu}_j \tilde{\nu}_j} = R_j, \quad (\bar{C}_j)_{y_j \tilde{\nu}_j} = \tilde{Z}_j R_j. \quad (43)$$

This implies that $\tilde{\nu}_j$ given $y_j, \alpha_j, \Gamma_j, \tau_j^2, \theta_{j1}, \dots, \theta_{jq}$ follows $N(\mu_{\tilde{\nu}_j}, \Sigma_{\tilde{\nu}_j})$, where

$$\mu_{\tilde{\nu}_j} = R_j^T \tilde{Z}_j^T (\bar{C}_j)_{y_j y_j}^{-1} (y_j - X_j \alpha_j), \quad \Sigma_{\tilde{\nu}_j} = R_j - R_j^T \tilde{Z}_j^T (\bar{C}_j)_{y_j y_j}^{-1} \tilde{Z}_j R_j. \quad (44)$$

The conditional density of $(\alpha_j, \Gamma_j, \tau_j^2)$ given y_j is obtained using (42). The subset j counterpart of W and b in (29) are W_j and $b_j^T = (\alpha_j^T, \gamma_j^T)$, where γ_j is the column-wise vectorization of Γ_j . The conditional density in (42) together with (30) implies that

$$\pi(\tau_j^2, b_j | y_j) \propto \frac{1}{(\tau_j^2)^{\tilde{s}_j \delta_n / 2}} \exp \left\{ -\frac{(y_j - W_j b_j)^T (y_j - W_j b_j)}{2\tau_j^2 / \delta_n} \right\} \frac{1}{\tau_j^2}$$

$$\begin{aligned}
 &= \frac{1}{(\tau_j^2)^{\bar{s}_j \delta_n / 2 + 1}} \exp \left\{ -\frac{(y_j - \hat{y}_j)^\top (y_j - \hat{y}_j) + (\hat{y}_j - W_j b_j)^\top (\hat{y}_j - W_j b_j)}{2\tau_j^2 / \delta_n} \right\} \\
 &= \frac{1}{(\tau_j^2)^{\bar{s}_j \delta_n / 2 + 1}} \exp \left(-\frac{\delta_n \|y_j - \hat{y}_j\|_2^2}{2\tau_j^2} \right) \exp \left\{ -\frac{(\hat{b}_j - b_j)^\top \delta_n (W_j^\top W_j) (\hat{b}_j - b_j)}{2\tau_j^2} \right\}, \tag{45}
 \end{aligned}$$

where $\hat{b}_j = (W_j^\top W_j)^{-1} W_j^\top y_j$ is the least squares estimate of b_j and $\hat{y}_j = W_j \hat{b}_j$ is the mean estimate of y_j based on \hat{b}_j . Marginalizing over b_j in (45) implies that

$$\pi(\tau_j^2 | y_j) \propto \frac{1}{(\tau_j^2)^{(\bar{s}_j \delta_n - p - q^2) / 2 + 1}} \exp \left(-\frac{\delta_n \|y_j - \hat{y}_j\|_2^2}{2\tau_j^2} \right), \quad \tau_j^2 | y_j \sim \frac{\delta_n \|y_j - \hat{y}_j\|_2^2}{\chi_{\bar{s}_j \delta_n - p - q^2}^2}, \tag{46}$$

and (45) assuming τ_j^2 is given implies that

$$\begin{aligned}
 \pi(b_j | y_j, \tau_j^2) &\propto \exp \left\{ -\frac{(\hat{b}_j - b_j)^\top \delta_n (W_j^\top W_j) (\hat{b}_j - b_j)}{2\tau_j^2} \right\}, \\
 b_j | y_j, \tau_j^2 &\sim N \left\{ \hat{b}_j, \frac{\tau_j^2}{\delta_n} (W_j^\top W_j)^{-1} \right\}; \tag{47}
 \end{aligned}$$

therefore, (46) and (47) imply that the distribution of (b_j, τ_j^2) given y_j lies in the Normal-Inverse-Gamma family. The α_j draw corresponds to the first p elements of b_j and the remaining elements of b are ‘‘un-vectorized’’ into the q -by- q matrix Γ_j .

Finally, the form of the full conditionals of $\theta_{j1}, \dots, \theta_{jq}$ and $\beta(u^*), y(u^*)$ for $u^* \in \mathcal{U}^*$ remain the same as in steps (d)–(f) of the sampling algorithm for drawing from the full data posterior distribution, except the log likelihood for $\theta_{j1}, \dots, \theta_{jq}$ is multiplied by a factor δ_n in step (d). Let $\beta_j(u^*), y_j(u^*), \nu_j(u^*) = \{\nu_{j1}(u^*), \dots, \nu_{jq}(u^*)\}$, and $R_{ja^*}, R_{ja^{**}}, \nu_{ja}^*$ ($a = 1, \dots, q$) be the subset j counterparts of $\beta(u^*), y(u^*), \nu(u^*)$ and $R_{a^*}, R_{a^{**}}, \nu_a^*$. The full conditional of ν_j in (42) implies that the conditional likelihood of $\theta_{j1}, \dots, \theta_{jq}$ given $\tilde{\nu}_j$ remains unchanged after stochastic approximation. Similarly, the GP prior on $\nu_a(\cdot)$ remains unchanged except that θ_{ja} replaces θ_a ($a = 1, \dots, q$). This implies that the steps for drawing from the full conditionals of $\theta_{j1}, \dots, \theta_{jq}$ and $\beta(u^*), y(u^*)$ for $u^* \in \mathcal{U}^*$ have the same form as in the steps (d)–(f).

In summary, the sampling algorithm for drawing from the j th subset posterior distribution of $(\alpha, \Gamma, \tau^2, \theta_1, \dots, \theta_q)$ and $\beta(u^*), y(u^*)$ for $u^* \in \mathcal{U}^*$ starts from an initial value of parameters $(\alpha_j^{(0)}, \Gamma_j^{(0)}, \tau_j^{2(0)}, \theta_{j1}^{(0)}, \dots, \theta_{jq}^{(0)})$ and cycles through the following six steps for $t = 0, 1, \dots, \infty$:

- (a*) draw $\tilde{\nu}_j^{(t+1)}$ given $y_j, \alpha_j^{(t)}, \Gamma_j^{(t)}, \tau_j^{2(t)}, \theta_{jq}^{(t)}, \dots, \theta_{jq}^{(t)}$ from $N(\mu_{\tilde{\nu}_j}, \Sigma_{\tilde{\nu}_j})$, where $\mu_{\tilde{\nu}_j}, \Sigma_{\tilde{\nu}_j}$ are defined in (44);
- (b*) draw $\tau_j^{2(t+1)}$ given $y_j, \tilde{\nu}_j^{(t+1)}$ using (46);
- (c*) draw $(\alpha_j^{(t+1)}, \gamma_j^{(t+1)})$ given $y_j, \tilde{\nu}_j^{(t+1)}, \tau_j^{2(t+1)}$ using (47) and the vectorization of $\gamma_j^{(t+1)}$ is reversed to obtain $\Gamma_j^{(t+1)}$;

- (d*) draw $\theta_{j1}^{(t+1)}, \dots, \theta_{jq}^{(t+1)}$ given $\tilde{\nu}_j^{(t+1)}$ using ESS as in step (d) of the full data posterior sampling algorithm after multiplying the log likelihood by δ_n ;
- (e*) draw $\nu_{ja}^{*(t+1)}$ and obtain $\beta_j(u^*)^{(t+1)}$ given $\tilde{\nu}_{ja}^{(t+1)}, \alpha^{(t+1)}, \Gamma^{(t+1)}, \theta_{ja}^{(t+1)}$ as in step (e) of the full data posterior sampling algorithm; and
- (f*) draw $y_j(u^*)^{(t+1)}$ given $X(u^*), \beta_j(u^*)^{(t+1)}, \tau_j^{2(t+1)}$ as in step (f) of the full data posterior sampling algorithm.

Appendix C. Proof of Theorem 1

We use $\|\cdot\|$ to denote the matrix operator norm, and $\text{tr}(\cdot)$ to denote the trace of a matrix or a covariance function. For two semi-positive definite matrices A_1 and A_2 , $A_1 \preceq A_2$ means $A_2 - A_1$ is semi-positive definite. For abbreviation, we write the Bayes L_2 -risk in estimating $\beta_0(\cdot)$ using the combined posterior from Theorem 1 as

$$\mathbb{L}(\bar{\Pi}) \equiv \mathbb{E}_{u^*} \mathbb{E}_{y,u} \mathbb{E}_{\bar{\beta}|y,u} \|\bar{\beta}(u^*) - \beta_0(u^*)\|_2^2, \quad (48)$$

where $\bar{\beta}(\cdot)$ be a q -variate random function drawn from the AMC posterior of $\beta(\cdot)$. Let \mathbb{E}_{u^*} , \mathbb{E}_{y_j, u_j} , $\mathbb{E}_{y,u}$, $\mathbb{E}_{\beta_j|y_j, u_j}$ (similarly $\mathbb{E}_{\nu_j|y_j, u_j}$ and $\mathbb{E}_{\nu_{ja}|y_j, u_j}$ for $a = 1, \dots, q$), $\mathbb{E}_{\bar{\beta}|y,u}$ (similarly $\mathbb{E}_{\bar{\nu}|y,u}$ and $\mathbb{E}_{\bar{\nu}_a|y,u}$ for $a = 1, \dots, q$) respectively be the expectations with respect to the distribution of u^* , the true data generating distribution of the subset data (y_j, u_j) , the true data generating distribution of the full data (y, u) , the subset posterior distribution of β (similarly ν and ν_a) given y_j, u_j after stochastic approximation, and the AMC posterior distribution of $\bar{\beta}$ (similarly $\bar{\nu}$ and $\bar{\nu}_a$) given the full data. Notations for the variances are similarly defined.

We begin by decomposing this Bayes L_2 -risk. Given the relation $\beta(\cdot) = \Gamma_0 \nu(\cdot)$, it follows that

$$\mathbb{L}(\bar{\Pi}) \leq \|\Gamma_0\|^2 \mathbb{E}_{u^*} \mathbb{E}_{y,u} \mathbb{E}_{\bar{\nu}|y,u} \|\bar{\nu}(u^*) - \nu_0(u^*)\|_2^2, \quad (49)$$

where $\|\Gamma_0\| < \infty$ given Assumption (A.2). Therefore, it suffices to study the Bayes L_2 risk of $\nu(\cdot)$.

Based on Assumption (A.6), we have the following decomposition: For any $u^* \in [0, 1]^d$,

$$\begin{aligned} & \mathbb{E}_{\bar{\nu}|y,u} \|\bar{\nu}(u^*) - \nu_0(u^*)\|_2^2 \\ &= \mathbb{E}_{\bar{\nu}|y,u} \|\bar{\nu}(u^*) - \mathbb{E}_{\bar{\nu}|y,u} \{\bar{\nu}(u^*)\} + \mathbb{E}_{\bar{\nu}|y,u} \{\bar{\nu}(u^*)\} - \nu_0(u^*)\|_2^2 \\ &= \mathbb{E}_{\bar{\nu}|y,u} \|\bar{\nu}(u^*) - \mathbb{E}_{\bar{\nu}|y,u} \{\bar{\nu}(u^*)\}\|_2^2 + \|\mathbb{E}_{\bar{\nu}|y,u} \{\bar{\nu}(u^*)\} - \nu_0(u^*)\|_2^2 \\ &= \sum_{a=1}^q \text{Var}_{\bar{\nu}_a|y,u} \{\bar{\nu}_a(u^*)\} + \|\mathbb{E}_{\bar{\nu}|y,u} \{\bar{\nu}(u^*)\} - \nu_0(u^*)\|_2^2 \\ &\leq \frac{\bar{c}}{k} \sum_{a=1}^q \sum_{j=1}^k \text{Var}_{\nu_{ja}|y_j, u_j} \{\nu_{ja}(u^*)\} \\ &\quad + \sum_{a=1}^q \left\{ \frac{1}{k} \sum_{j=1}^k \left[\mathbb{E}_{\nu_{ja}|y_j, u_j} \{\nu_{ja}(u^*)\} - \nu_{0a}(u^*) \right] + O_p(n^{-1/2}) \right\}^2 \end{aligned}$$

$$\begin{aligned}
 &\leq \frac{\bar{c}}{k} \sum_{a=1}^q \sum_{j=1}^k \text{Var}_{\nu_{ja}|y_j, u_j} \{\nu_{ja}(u^*)\} \\
 &\quad + 2 \sum_{a=1}^q \left\{ \frac{1}{k} \sum_{j=1}^k \left[\mathbb{E}_{\nu_{ja}|y_j, u_j} \{\nu_{ja}(u^*)\} - \nu_{0a}(u^*) \right] \right\}^2 + O_p(n^{-1}), \tag{50}
 \end{aligned}$$

where the last inequality follows from $(a+b)^2 \leq 2(a^2+b^2)$. Using this inequality again, we have that

$$\begin{aligned}
 &\mathbb{E}_{u^*} \mathbb{E}_{y, u} \mathbb{E}_{\bar{\nu}|y, u} \|\bar{\nu}(u^*) - \nu_0(u^*)\|_2^2 \\
 &\leq \frac{\bar{c}}{k} \sum_{a=1}^q \sum_{j=1}^k \mathbb{E}_{u^*} \mathbb{E}_{y_j, u_j} \text{Var}_{\nu_{ja}|y_j, u_j} \{\nu_{ja}(u^*)\} \\
 &\quad + 2 \sum_{a=1}^q \mathbb{E}_{u^*} \mathbb{E}_{y, u} \sum_{a=1}^q \left\{ \frac{1}{k} \sum_{j=1}^k \left[\mathbb{E}_{\nu_{ja}|y_j, u_j} \{\nu_{ja}(u^*)\} - \nu_{0a}(u^*) \right] \right\}^2 + O_p(n^{-1}) \\
 &= \frac{\bar{c}}{k} \sum_{a=1}^q \sum_{j=1}^k \mathbb{E}_{u^*} \mathbb{E}_{y_j, u_j} \text{Var}_{\nu_{ja}|y_j, u_j} \{\nu_{ja}(u^*)\} \\
 &\quad + 2 \sum_{a=1}^q \mathbb{E}_{u^*} \mathbb{E}_{y, u} \left\{ \frac{1}{k} \sum_{j=1}^k \left[\mathbb{E}_{\nu_{ja}|y_j, u_j} \{\nu_{ja}(u^*)\} - \mathbb{E}_{\nu_{ja}, y_j|u_j} \{\nu_{ja}(u^*)\} \right] \right\} \\
 &\quad + \frac{1}{k} \sum_{j=1}^k \left\{ \mathbb{E}_{\nu_{ja}, y_j|u_j} \{\nu_{ja}(u^*)\} - \nu_{0a}(u^*) \right\}^2 + O_p(n^{-1}) \\
 &\leq \frac{\bar{c}}{k} \sum_{a=1}^q \sum_{j=1}^k \mathbb{E}_{u^*} \mathbb{E}_{y_j, u_j} \text{Var}_{\nu_{ja}|y_j, u_j} \{\nu_{ja}(u^*)\} \\
 &\quad + 4 \sum_{a=1}^q \mathbb{E}_{u^*} \mathbb{E}_{y, u} \left(\frac{1}{k} \sum_{j=1}^k \left[\mathbb{E}_{\nu_{ja}|y_j, u_j} \{\nu_{ja}(u^*)\} - \mathbb{E}_{\nu_{ja}, y_j|u_j} \{\nu_{ja}(u^*)\} \right] \right)^2 \\
 &\quad + 4 \sum_{a=1}^q \mathbb{E}_{u^*} \mathbb{E}_u \left(\frac{1}{k} \sum_{j=1}^k \left[\mathbb{E}_{\nu_{ja}, y_j|u_j} \{\nu_{ja}(u^*)\} - \nu_{0a}(u^*) \right] \right)^2 + O_p(n^{-1}). \tag{51}
 \end{aligned}$$

Furthermore, using the sampling independence between subsets of u_j as in Assumption (A.1), the second term in (51) can be simplified as

$$\begin{aligned}
 &4 \sum_{a=1}^q \mathbb{E}_{u^*} \mathbb{E}_{y, u} \left(\frac{1}{k} \sum_{j=1}^k \left[\mathbb{E}_{\nu_{ja}|y_j, u_j} \{\nu_{ja}(u^*)\} - \mathbb{E}_{\nu_{ja}, y_j|u_j} \{\nu_{ja}(u^*)\} \right] \right)^2 \\
 &= \frac{4}{k^2} \sum_{a=1}^q \sum_{j=1}^k \mathbb{E}_{u^*} \mathbb{E}_{y_j, u_j} \left[\mathbb{E}_{\nu_{ja}|y_j, u_j} \{\nu_{ja}(u^*)\} - \mathbb{E}_{\nu_{ja}, y_j|u_j} \{\nu_{ja}(u^*)\} \right]^2 \\
 &\quad + \frac{4}{k^2} \sum_{a=1}^q \sum_{j \neq j'} \mathbb{E}_{u^*} \mathbb{E}_u \mathbb{E}_{u_j|u} \mathbb{E}_{u_{j'}|u} \left(\left[\mathbb{E}_{y_j|u_j} \mathbb{E}_{\nu_{ja}|y_j, u_j} \{\nu_{ja}(u^*)\} - \mathbb{E}_{\nu_{ja}, y_j|u_j} \{\nu_{ja}(u^*)\} \right] \right. \\
 &\quad \left. \left[\mathbb{E}_{y_{j'}|u_{j'}} \mathbb{E}_{\nu_{ja}|y_{j'}, u_{j'}} \{\nu_{ja}(u^*)\} - \mathbb{E}_{\nu_{ja}, y_{j'}|u_{j'}} \{\nu_{ja}(u^*)\} \right] \right)
 \end{aligned}$$

$$\begin{aligned}
 & \times \left[\mathbb{E}_{y_{j'}|u_{j'}} \mathbb{E}_{\nu_{j'a}|y_{j'},u_{j'}} \{ \nu_{j'a}(u^*) \} - \mathbb{E}_{\nu_{j'a},y_{j'}|u_{j'}} \{ \nu_{j'a}(u^*) \} \right] \Big) \\
 &= \frac{4}{k^2} \sum_{a=1}^q \sum_{j=1}^k \mathbb{E}_{u^*} \mathbb{E}_{y_j, u_j} \left[\mathbb{E}_{\nu_{ja}|y_j, u_j} \{ \nu_{ja}(u^*) \} - \mathbb{E}_{\nu_{ja}, y_j | u_j} \{ \nu_{ja}(u^*) \} \right]^2 \\
 &= \frac{4}{k^2} \sum_{a=1}^q \sum_{j=1}^k \mathbb{E}_{u^*} \mathbb{E}_{u_j} \left(\text{Var}_{y_j|u_j} \left[\mathbb{E}_{\nu_{ja}|y_j, u_j} \{ \nu_{ja}(u^*) \} \right] \right), \tag{52}
 \end{aligned}$$

where the second term after the first equal sign is zero due to the independence between u_j and $u_{j'}$ given the full data u according to Assumption (A.1). The third term in (51) can be bounded from above by the Cauchy-Schwarz inequality:

$$\begin{aligned}
 & 4 \sum_{a=1}^q \mathbb{E}_{u^*} \mathbb{E}_u \left(\frac{1}{k} \sum_{j=1}^k \left[\mathbb{E}_{\nu_{ja}, y_j | u_j} \{ \nu_{ja}(u^*) \} - \nu_{0a}(u^*) \right] \right)^2 \\
 & \leq \frac{4}{k} \sum_{a=1}^q \sum_{j=1}^k \mathbb{E}_{u^*} \mathbb{E}_{u_j} \left[\mathbb{E}_{\nu_{ja}, y_j | u_j} \{ \nu_{ja}(u^*) \} - \nu_{0a}(u^*) \right]^2. \tag{53}
 \end{aligned}$$

We can plug in (52), and (53) to (51) to further obtain that

$$\begin{aligned}
 & \mathbb{E}_{u^*} \mathbb{E}_{y, u} \mathbb{E}_{\bar{\nu}|y, u} \|\bar{\nu}(u^*) - \nu_0(u^*)\|_2^2 \\
 & \leq \frac{4}{k} \sum_{a=1}^q \sum_{j=1}^k \mathbb{E}_{u^*} \mathbb{E}_{u_j} \left[\mathbb{E}_{\nu_{ja}, y_j | u_j} \{ \nu_{ja}(u^*) \} - \nu_{0a}(u^*) \right]^2 \\
 & \quad + \frac{4}{k^2} \sum_{a=1}^q \sum_{j=1}^k \mathbb{E}_{u^*} \mathbb{E}_{u_j} \left(\text{Var}_{y_j|u_j} \left[\mathbb{E}_{\nu_{ja}|y_j, u_j} \{ \nu_{ja}(u^*) \} \right] \right) \\
 & \quad + \frac{\bar{c}}{k} \sum_{a=1}^q \sum_{j=1}^k \mathbb{E}_{u^*} \mathbb{E}_{y_j, u_j} \text{Var}_{\nu_{ja}|y_j, u_j} \{ \nu_{ja}(u^*) \} + O_p(n^{-1}) \\
 & \leq \frac{4}{k} \sum_{j=1}^k \mathbb{E}_{u^*} \mathbb{E}_{u_j} \left\| \mathbb{E}_{\nu_j, y_j | u_j} \{ \nu_j(u^*) \} - \nu_0(u^*) \right\|_2^2 \\
 & \quad + \frac{4}{k^2} \sum_{j=1}^k \mathbb{E}_{u^*} \mathbb{E}_{u_j} \text{tr} \left(\text{Var}_{y_j|u_j} \left[\mathbb{E}_{\nu_j|y_j, u_j} \{ \nu_j(u^*) \} \right] \right) \\
 & \quad + \frac{\bar{c}}{k} \sum_{j=1}^k \mathbb{E}_{u^*} \mathbb{E}_{y_j, u_j} \text{tr} \left(\text{Var}_{\nu_j|y_j, u_j} \{ \nu_j(u^*) \} \right) + O_p(n^{-1}). \tag{54}
 \end{aligned}$$

Now we provide the detailed expressions for three terms on the right-hand side of (54). We define the following notations. Let Γ_{0a} be the a th column of Γ_0 , $a = 1, \dots, q$. In the derivation below, $\rho_a(\cdot, \cdot)$ is the correlation function with its parameter fixed at θ_{0a} , $a = 1, \dots, q$. For $j = 1, \dots, k$ and $a = 1, \dots, q$,

$$R_{ja}(u^*) = \{ \rho_{ja}(u_{j1}, u^*), \dots, \rho_a(u_{jm}, u^*) \}^T \in \mathbb{R}^m,$$

$$\begin{aligned}
 R_j(u^*) &= \text{diag}\{R_{j1}^T(u^*), \dots, R_{jq}^T(u^*)\} \in \mathbb{R}^{q \times qm}, \\
 R(u, u') &= \text{diag}\{\rho_1(u, u'), \dots, \rho_q(u, u')\} \in \mathbb{R}^{q \times q}, \quad u, u' \in [0, 1]^d, \\
 \tilde{Z}(\cdot) &= Z(\cdot)\Gamma_0 \in \mathbb{R}^{1 \times q}, \quad \tilde{Z}_a(\cdot) = Z(\cdot)\Gamma_{0a} \in \mathbb{R}, \\
 \tilde{Z}_{ja} &= \text{diag}\{\tilde{Z}_a(u_{j1}), \dots, \tilde{Z}_a(u_{jm})\} \in \mathbb{R}^{m \times m}, \quad \tilde{Z}_j = \{\tilde{Z}_{j1}, \dots, \tilde{Z}_{jq}\} \in \mathbb{R}^{m \times qm}, \\
 (R_{ja})_{ii'} &= \rho_a(u_{ji}, u_{ji'}), \quad i, i' \in \{1, \dots, m\}, \quad \tilde{R}_{jj} = \text{diag}(R_{j1}, \dots, R_{jq}) \in \mathbb{R}^{qm \times qm}, \\
 \tilde{\nu}_{ja} &= \{\nu_a(u_{j1}), \dots, \nu_a(u_{jm})\}^T \in \mathbb{R}^m, \quad \tilde{\nu}_j^T = (\tilde{\nu}_{1j}^T, \dots, \tilde{\nu}_{qj}^T) \in \mathbb{R}^{qm}, \\
 \nu_{0ja} &= \{\nu_{0a}(u_{j1}), \dots, \nu_{0a}(u_{jm})\}^T \in \mathbb{R}^m.
 \end{aligned} \tag{55}$$

Based on the subset model (19) and the standard GP derivation, we have that for $a = 1, \dots, q$,

$$\mathbb{E}_{\nu_j, y_j | u_j} \{\nu_j(u^*)\} - \nu_0(u^*) = R_j(u^*) \tilde{Z}_j^T \left(\tilde{Z}_j \tilde{R}_{jj} \tilde{Z}_j^T + \frac{\tau_0^2 \lambda_n}{k} I_m \right)^{-1} \tilde{Z}_j \tilde{\nu}_j - \nu_0(u^*), \tag{56}$$

$$\begin{aligned}
 \text{Var}_{y_j | u_j} \left[\mathbb{E}_{\nu_j | y_j, u_j} \{\nu_j(u^*)\} \right] &= \text{Var}_{y_j | u_j} \left[R_j(u^*) \tilde{Z}_j^T \left(\tilde{Z}_j \tilde{R}_{jj} \tilde{Z}_j^T + \frac{\tau_0^2 \lambda_n}{k} I_m \right)^{-1} y_j \right] \\
 &= \tau_0^2 R_j(u^*) \tilde{Z}_j^T \left(\tilde{Z}_j \tilde{R}_{jj} \tilde{Z}_j^T + \frac{\tau_0^2 \lambda_n}{k} I_m \right)^{-2} \tilde{Z}_j R_j(u^*)^T,
 \end{aligned} \tag{57}$$

$$\text{Var}_{\nu_j | y_j, u_j} \{\nu_j(u^*)\} = \lambda_n^{-1} \left\{ R(u^*, u^*) - R_j(u^*) \tilde{Z}_j^T \left(\tilde{Z}_j \tilde{R}_{jj} \tilde{Z}_j^T + \frac{\tau_0^2 \lambda_n}{k} I_m \right)^{-1} \tilde{Z}_j R_j(u^*)^T \right\}. \tag{58}$$

In the following lemmas, we give upper bounds for (56), (57), and (58), respectively. We first introduce some additional notations. In Assumption (A.5), we assume that the largest eigenvalue of Ω is bounded above by a constant \bar{c}_Ω and its smallest eigenvalue is bounded below by a constant \underline{c}_Ω , for $0 < \underline{c}_\Omega < \bar{c}_\Omega < \infty$ and for all $a = 1, \dots, q$. By the definition of $\tilde{Z}(\cdot) = Z(\cdot)\Gamma_0$ and $\tilde{Z}_a(\cdot) = Z(\cdot)\Gamma_{0a}$ ($a = 1, \dots, q$) and Assumption (A.5), we have that $|\tilde{Z}_a(u)| \leq C_Z \|\Gamma_{0a}\|_2 \leq C_Z \|\Gamma_0\|$ for all $a = 1, \dots, q$ and all $u \in [0, 1]^d$. We define the positive constant $\tilde{C}_Z \equiv C_Z \|\Gamma_0\|$ for the derivations below. We define the \mathbb{R}^q -valued functional space $L_2^q(\text{d}u)$ for $\mathbf{f} = (f_1, \dots, f_q)^T$ with $f_1, \dots, f_q \in L_2(\text{d}u)$. The q -variate RKHS \mathbb{H} is defined to the space of functions with the finite \mathbb{H} -norm $\|\mathbf{f}\|_{\mathbb{H}}^2 = \sum_{a=1}^q \|f_a\|_{\mathbb{H}_a}^2$ if $\mathbf{f} = (f_1, \dots, f_q)^T \in L_2^q(\text{d}u)$ and $f_a \in \mathbb{H}_a$ for $a = 1, \dots, q$. This q -variate RKHS satisfies the reproducing property that if $\mathbf{f} \in \mathbb{H}$, then for any $u \in [0, 1]^d$ and any $\mathbf{c} \in \mathbb{R}^q$, $\langle \mathbf{f}, R(u, \cdot) \mathbf{c} \rangle_{\mathbb{H}} = \mathbf{f}(u)^T \mathbf{c}$; see Section 3.2 of Álvarez et al. (2012). For abbreviation, we let $\Lambda(c, h) = \sum_{i=1}^h \mu_{i*} / (\mu_{i*} + c)$ for any $c > 0$, any positive integer h , and μ_{i*} as defined in Assumption (A.3). For $a = 1, \dots, q$ and any positive integer h , let $\text{tr}(\rho_a) = \sum_{i=1}^{\infty} \mu_{ai}$, $\text{tr}(\rho_{a,h}) = \sum_{i=h+1}^{\infty} \mu_{ai}$, $\text{Tr}(\rho) = \sum_{a=1}^q \text{tr}(\rho_a)$, and $\text{Tr}(\rho, h) = \sum_{a=1}^q \text{tr}(\rho_{a,h})$.

Lemma 2 *Suppose that Assumptions (A.1)–(A.6) hold. Then for every $j = 1, \dots, k$,*

$$\begin{aligned}
 &\mathbb{E}_{u^*} \mathbb{E}_{u_j} \left\| \mathbb{E}_{\nu_j, y_j | u_j} \{\nu_j(u^*)\} - \nu_0(u^*) \right\|_2^2 \\
 &\leq 8 \frac{\tau_0^2 \lambda_n}{\underline{c}_\Omega k m} \|\nu_0\|_{\mathbb{H}}^2 + 8 \frac{k m}{\tau_0^2 \lambda_n} C_\varphi^4 \tilde{C}_Z^4 \bar{c}_\Omega \underline{c}_\Omega^{-2} \|\nu_0\|_{\mathbb{H}}^2 \text{Tr}(\rho) \text{Tr}(\rho, \bar{h}) + \mu_{(\bar{h}+1)*} \|\nu_0\|_{\mathbb{H}}^2
 \end{aligned}$$

$$+ 2q\bar{h}\|\nu_0\|_{\mathbb{H}}^2 \text{Tr}(\rho) \exp \left\{ -\frac{m}{8(B^2 + B)} \right\},$$

where $B = C_{\varphi}^2 \tilde{C}_{\tilde{Z}}^2 \underline{c}_{\Omega}^{-1} q \Lambda (\tau_0^2 \lambda_n / (\underline{c}_{\Omega} k m), \bar{h}) + 1$.

Proof [of Lemma 2] Consider the subset posterior distribution of $\nu_j(\cdot)$ on the subset j . The j th subset posterior distribution of $\nu_j(u^*)$ has the mean

$$\hat{\nu}_j(u^*) = R_j(u^*) \tilde{Z}_j^T \left(\tilde{Z}_j \tilde{R}_{jj} \tilde{Z}_j^T + \frac{\tau_0^2 \lambda_n}{k} I_m \right)^{-1} y_j. \quad (59)$$

We view $\hat{\nu}_j(u^*) = \{\hat{\nu}_{j1}(u^*), \dots, \hat{\nu}_{jq}(u^*)\}^T$ as a q -variate function of u^* . Define the $[0, 1]^d \mapsto \mathbb{R}^q$ operators Δ_j ($j = 1, \dots, k$) and Δ as

$$\begin{aligned} \Delta_j(\cdot) &= R_j(\cdot) \tilde{Z}_j^T \left(\tilde{Z}_j \tilde{R}_{jj} \tilde{Z}_j^T + \frac{\tau_0^2 \lambda_n}{k} I_m \right)^{-1} y_j - \nu_0(\cdot) \equiv \hat{\nu}_j(\cdot) - \nu_0(\cdot), \\ \Delta(\cdot) &= \frac{1}{k} \sum_{j=1}^k \{\hat{\nu}_j(\cdot) - \nu_0(\cdot)\} = \frac{1}{k} \sum_{j=1}^k \Delta_j(\cdot). \end{aligned} \quad (60)$$

For $a = 1, \dots, q$, Δ_{ja} and Δ_a denote the a th components of Δ_j and Δ , respectively. We can recognize that for $j = 1, \dots, k$, $\hat{\nu}_j(\cdot)$ in (60) is the solution of q -variate function to the optimization problem (see Section 3.2 in Álvarez et al. 2012)

$$\operatorname{argmin}_{\nu \in \mathbb{H}} \left[\sum_{i=1}^m \left\{ y(u_{ji}) - \tilde{Z}(u_{ji}) \nu(u_{ji}) \right\}^2 + \frac{\tau_0^2 \lambda_n}{k} \|\nu\|_{\mathbb{H}}^2 \right]. \quad (61)$$

Now fix $a \in \{1, \dots, q\}$. Taking the Frechét derivative of (61) with respect to ν_a and plugging in $\hat{\nu}_j$ gives

$$\sum_{i=1}^m \left\{ \tilde{Z}(u_{ji}) \hat{\nu}_j(u_{ji}) - y(u_{ji}) \right\} \tilde{Z}_a(u_{ji}) \rho_a(u_{ji}, \cdot) + \frac{\tau_0^2 \lambda_n}{k} \hat{\nu}_{ja}(\cdot) = 0. \quad (62)$$

Stacking this up across $a = 1, \dots, q$ gives

$$\sum_{i=1}^m \left\{ \tilde{Z}(u_{ji}) \hat{\nu}_j(u_{ji}) - y(u_{ji}) \right\} R(u_{ji}, \cdot) \tilde{Z}(u_{ji})^T + \frac{\tau_0^2 \lambda_n}{k} \hat{\nu}_j(\cdot) = 0. \quad (63)$$

Since the true model assumes that

$$y(u_{ji}) = \tilde{Z}(u_{ji}) \nu_0(u_{ji}) + \epsilon(u_{ji}) = \langle \nu_0, R(u_{ji}, \cdot) \tilde{Z}(u_{ji})^T \rangle_{\mathbb{H}} + \epsilon(u_{ji}),$$

so

$$\begin{aligned} \tilde{Z}(u_{ji}) \hat{\nu}_j(u_{ji}) - y(u_{ji}) &= \tilde{Z}(u_{ji}) \{\hat{\nu}_j(u_{ji}) - \nu_0(u_{ji})\} - \epsilon(u_{ji}) \\ &= \langle \Delta_j(\cdot), R(u_{ji}, \cdot) \tilde{Z}(u_{ji})^T \rangle_{\mathbb{H}} - \epsilon(u_{ji}), \end{aligned}$$

Hence, we take expectations $\mathbb{E}_{\nu_j, y_j | u_j}$ on both sides of (62) and obtain that

$$\begin{aligned}
 0 &= \sum_{i=1}^m \mathbb{E}_{\nu_j, y_j | u_j} \left\{ \tilde{Z}(u_{ji}) \hat{\nu}_j(u_{ji}) - y(u_{ji}) \right\} R(u_{ji}, \cdot) \tilde{Z}(u_{ji})^\top + \frac{\tau_0^2 \lambda_n}{k} \mathbb{E}_{\nu_j, y_j | u_j} \{ \hat{\nu}_j(\cdot) \} \\
 &= \sum_{i=1}^m \left\langle \mathbb{E}_{\nu_j, y_j | u_j} \{ \Delta_j(\cdot) \}, R(u_{ji}, \cdot) \tilde{Z}(u_{ji})^\top \right\rangle_{\mathbb{H}} R(u_{ji}, \cdot) \tilde{Z}(u_{ji})^\top \\
 &\quad - \mathbb{E}_{\nu_j, y_j | u_j} \{ \epsilon(u_{ji}) \} R(u_{ji}, \cdot) \tilde{Z}(u_{ji})^\top + \frac{\tau_0^2 \lambda_n}{k} \mathbb{E}_{\nu_j, y_j | u_j} \{ \hat{\nu}_j(\cdot) \} \\
 &= \sum_{i=1}^m \left\langle \mathbb{E}_{\nu_j, y_j | u_j} \{ \Delta_j(\cdot) \}, R(u_{ji}, \cdot) \tilde{Z}(u_{ji})^\top \right\rangle_{\mathbb{H}} R(u_{ji}, \cdot) \tilde{Z}(u_{ji})^\top + \frac{\tau_0^2 \lambda_n}{k} \mathbb{E}_{\nu_j, y_j | u_j} \{ \hat{\nu}_j(\cdot) \}.
 \end{aligned} \tag{64}$$

Using (60), $\mathbb{E}_{\nu_j, y_j | u_j} \{ \hat{\nu}_j(\cdot) \} = \mathbb{E}_{\nu_j, y_j | u_j} \{ \Delta_j(\cdot) \} + \nu_0$, and dividing by m in (64), we obtain that

$$\begin{aligned}
 &\frac{1}{m} \sum_{i=1}^m \left\langle \mathbb{E}_{\nu_j, y_j | u_j} \{ \Delta_j(\cdot) \}, R(u_{ji}, \cdot) \tilde{Z}(u_{ji})^\top \right\rangle_{\mathbb{H}} R(u_{ji}, \cdot) \tilde{Z}(u_{ji})^\top + \frac{\tau_0^2 \lambda_n}{km} \mathbb{E}_{\nu_j, y_j | u_j} \{ \Delta_j(\cdot) \} \\
 &= - \frac{\tau_0^2 \lambda_n}{km} \nu_0(\cdot).
 \end{aligned} \tag{65}$$

If we define the j th subset covariance operator $\hat{\Sigma}_j = m^{-1} \sum_{j=1}^m R(u_{ji}, \cdot) \tilde{Z}(u_{ji})^\top \otimes R(u_{ji}, \cdot) \tilde{Z}(u_{ji})^\top$, then (65) reduces to

$$\begin{aligned}
 &\left(\hat{\Sigma}_j + \frac{\tau_0^2 \lambda_n}{km} I_m \right) \mathbb{E}_{\nu_j, y_j | u_j} \{ \Delta_j(\cdot) \} = - \frac{\tau_0^2 \lambda_n}{km} \nu_0(\cdot) \\
 &\implies \left\| \mathbb{E}_{\nu_j, y_j | u_j} (\Delta_j) \right\|_{\mathbb{H}} \leq \left\| \nu_0 \right\|_{\mathbb{H}}, \quad j = 1, \dots, k,
 \end{aligned} \tag{66}$$

where the last inequality follows because $\hat{\Sigma}_j$ is a positive semi-definite matrix.

The rest of the proof finds an upper bound for

$$\left\| \mathbb{E}_{\nu_j, y_j | u_j} (\Delta_j) \right\|_2^2 = \sum_{a=1}^q \left\| \mathbb{E}_{\nu_{ja}, y_j | u_j} (\Delta_{ja}) \right\|_2^2.$$

The main idea is to reduce this problem to a finite dimensional one indexed by a chosen $\bar{h} = \lceil \frac{2v}{2v-d} \rceil \in \mathbb{N}$ as specified in Assumption (A.5). For each $j = 1, \dots, k$ and $a = 1, \dots, q$, let $\delta_{ja} = \left(\delta_{ja1}, \dots, \delta_{ja\bar{h}}, \delta_{ja(\bar{h}+1)}, \dots, \delta_{ja\infty} \right)^\top$, such that

$$\begin{aligned}
 \mathbb{E}_{\nu_{ja}, y_j | u_j} \{ \Delta_{ja}(\cdot) \} &= \sum_{h=1}^{\infty} \delta_{jah} \phi_{ah}(\cdot), \quad \delta_{jah} = \left\langle \mathbb{E}_{\nu_{ja}, y_j | u_j} \{ \Delta_{ja}(\cdot) \}, \phi_h(\cdot) \right\rangle_{L_2(\mathrm{d}u)}, \\
 \left\| \mathbb{E}_{\nu_{ja}, y_j | u_j} (\Delta_{ja}) \right\|_2^2 &= \sum_{h=1}^{\infty} \delta_{jah}^2.
 \end{aligned} \tag{67}$$

Define the vectors

$$\delta_{ja}^\downarrow = (\delta_{ja1}, \dots, \delta_{ja\bar{h}})^\top, \quad \delta_{ja}^\uparrow = (\delta_{ja(\bar{h}+1)}, \dots, \delta_{ja\infty})^\top,$$

so $\|\mathbb{E}_{\nu_{ja}, y_j | u_j}(\Delta_j)\|_2^2 = \|\delta_{ja}^\downarrow\|_2^2 + \|\delta_{ja}^\uparrow\|_2^2$ and we upper bound $\|\mathbb{E}_{\nu_{ja}, y_j | u_j}(\Delta_{ja})\|_2^2$ by separately upper bounding $\|\delta_{ja}^\downarrow\|_2^2$ and $\|\delta_{ja}^\uparrow\|_2^2$. Using the expansion $\rho_a(u, u') = \sum_{h=1}^{\infty} \mu_h \varphi_{ah}(u) \varphi_{ah}(u')$ for any $u, u' \in [0, 1]^d$, we have the following upper bound for $\|\delta_{ja}^\downarrow\|_2^2$ and $\|\delta_{ja}^\uparrow\|_2^2$:

$$\begin{aligned} \|\delta_{ja}^\downarrow\|_2^2 &= \frac{\mu_{a1}}{\mu_{a1}} \sum_{h=1}^{\bar{h}} \delta_{jah}^2 \stackrel{(i)}{\leq} \mu_{a1} \sum_{h=1}^{\bar{h}} \frac{\delta_{jah}^2}{\mu_{ah}} \\ &\leq \mu_{a1} \sum_{h=1}^{\infty} \frac{\delta_{jah}^2}{\mu_{ah}} \stackrel{(ii)}{=} \mu_{a1} \left\| \mathbb{E}_{\nu_{ja}, y_j | u_j}(\Delta_{ja}) \right\|_{\mathbb{H}_a}^2 \stackrel{(iii)}{\leq} \text{tr}(\rho_a) \|\nu_{0a}\|_{\mathbb{H}_a}^2, \end{aligned} \quad (68)$$

$$\begin{aligned} \|\delta_{ja}^\uparrow\|_2^2 &= \frac{\mu_{a(\bar{h}+1)}}{\mu_{a(\bar{h}+1)}} \sum_{h=\bar{h}+1}^{\infty} \delta_{jah}^2 \leq \mu_{a(\bar{h}+1)} \sum_{h=\bar{h}+1}^{\infty} \frac{\delta_{jah}^2}{\mu_{ah}} \\ &\stackrel{(iv)}{\leq} \mu_{a(\bar{h}+1)} \left\| \mathbb{E}_{\nu_{ja}, y_j | u_j}(\Delta_{ja}) \right\|_{\mathbb{H}_a}^2 \end{aligned} \quad (69)$$

where (i) follows from the decreasing eigenvalues $\mu_{a1} \geq \mu_{a2} \geq \dots$, (ii) and (iv) follow because $\|\mathbb{E}_{\nu_{ja}, y_j | u_j}(\Delta_{ja})\|_{\mathbb{H}_a}^2 = \sum_{h=1}^{\infty} \delta_{jah}^2 / \mu_{ah}$, (iii) follow from $\mu_{a1} \leq \text{tr}(\rho_a)$.

Let

$$\delta_j = (\delta_{j1}^\top, \dots, \delta_{jq}^\top)^\top, \quad \delta_j^\downarrow = (\delta_{j1}^{\downarrow\top}, \dots, \delta_{jq}^{\downarrow\top})^\top, \quad \delta_j^\uparrow = (\delta_{j1}^{\uparrow\top}, \dots, \delta_{jq}^{\uparrow\top})^\top.$$

We now derive a more refined upper bound than (68) for $\|\delta_j^\downarrow\|_2^2 = \sum_{a=1}^q \|\delta_{ja}^\downarrow\|_2^2$, such that the upper bound converges to zero as $m \rightarrow \infty$. For a given positive integer \bar{h} , each $a = 1, \dots, q$, and each $j = 1, \dots, k$, let $M_a = \text{diag}(\mu_{a1}, \dots, \mu_{a\bar{h}}) \in \mathbb{R}^{\bar{h} \times \bar{h}}$, $M = \text{diag}\{M_1, \dots, M_q\} \in \mathbb{R}^{q\bar{h} \times q\bar{h}}$. Let $\Phi_j \in \mathbb{R}^{qm \times q\bar{h}}$ and $\Phi_{ja} \in \mathbb{R}^{m \times \bar{h}}$ be matrices such that

$$\begin{aligned} (\Phi_{ja})_{ih} &= \varphi_{ah}(u_{ji}), \quad i = 1, \dots, m, \quad h = 1, \dots, \bar{h}, \quad j = 1, \dots, k, \\ \Phi_j &= \text{diag}\{\Phi_{j1}, \dots, \Phi_{jq}\} \in \mathbb{R}^{qm \times q\bar{h}}. \end{aligned} \quad (70)$$

Let $\nu_{0a}(\cdot) = \sum_{h=1}^{\infty} \zeta_{0ah} \varphi_{ah}(\cdot)$ with $\zeta_{0ah} = \langle \nu_{0a}, \varphi_{ah} \rangle_{L_2(\text{d}u)}$ for $h = 1, 2, \dots$ and $a = 1, \dots, q$. Also define the tail error vector $v_{ja} = (v_{ja1}, \dots, v_{jam})^\top \in \mathbb{R}^m$ and $v_j = (v_{j1}^\top, \dots, v_{jq}^\top)^\top \in \mathbb{R}^{qm}$ ($j = 1, \dots, k$) such that

$$v_{jai} = \sum_{h=\bar{h}+1}^{\infty} \delta_{jah} \varphi_{ah}(u_{ji}), \quad i = 1, \dots, m.$$

For any $g \in \{1, \dots, \bar{h}\}$ and any $a \in \{1, \dots, q\}$, we take the \mathbb{H}_a -inner product between φ_{ag} and the a th component of (65) and obtain that for $j = 1, \dots, k$,

$$\frac{1}{m} \sum_{i=1}^m \left\langle \mathbb{E}_{\nu_j, y_j | u_j} \{ \Delta_j(\cdot) \}, R(u_{ji}, \cdot) \tilde{Z}(u_{ji})^\top \right\rangle_{\mathbb{H}} \left\langle \rho_a(u_{ji}, \cdot) \tilde{Z}_a(u_{ji}), \varphi_{ag}(\cdot) \right\rangle_{\mathbb{H}_a}$$

$$+ \frac{\tau_0^2 \lambda_n}{km} \left\langle \mathbb{E}_{\nu_j, y_j | u_j} \{ \Delta_{ja}(\cdot) \}, \varphi_{ag}(\cdot) \right\rangle_{\mathbb{H}_a} = - \frac{\tau_0^2 \lambda_n}{km} \langle \nu_{0a}, \varphi_{ag} \rangle_{\mathbb{H}_a}, \quad (71)$$

which implies that

$$\frac{1}{m} \sum_{i=1}^m \left[\sum_{b=1}^q \tilde{Z}_b(u_{ji}) \mathbb{E}_{\nu_j, y_j | u_j} \{ \Delta_{jb}(u_{ji}) \} \right] \tilde{Z}_a(u_{ji}) \varphi_{ag}(u_{ji}) + \frac{\tau_0^2 \lambda_n}{km} \frac{\delta_{jag}}{\mu_{ag}} = - \frac{\tau_0^2 \lambda_n}{km} \frac{\zeta_{0ag}}{\mu_{ag}}. \quad (72)$$

The first term on the left-hand side of (72) can be rewritten as

$$\begin{aligned} & \frac{1}{m} \sum_{i=1}^m \left[\sum_{b=1}^q \tilde{Z}_b(u_{ji}) \mathbb{E}_{\nu_j, y_j | u_j} \{ \Delta_{jb}(u_{ji}) \} \right] \tilde{Z}_a(u_{ji}) \varphi_{ag}(u_{ji}) \\ &= \frac{1}{m} \sum_{i=1}^m \left[\sum_{b=1}^q \tilde{Z}_b(u_{ji}) \sum_{h=1}^{\bar{h}} \delta_{jbh} \varphi_{bh}(u_{ji}) \right] \tilde{Z}_a(u_{ji}) \varphi_{ag}(u_{ji}) \\ & \quad + \frac{1}{m} \sum_{i=1}^m \left[\sum_{b=1}^q \tilde{Z}_b(u_{ji}) \sum_{h=\bar{h}+1}^{\infty} \delta_{jbh} \varphi_{bh}(u_{ji}) \right] \tilde{Z}_a(u_{ji}) \varphi_{ag}(u_{ji}) \\ &= \frac{1}{m} \sum_{i=1}^m \sum_{b=1}^q \left(\tilde{Z}_{ja} \Phi_{ja} \right)_{ig} \left(\tilde{Z}_{jb} \Phi_{jb} \delta_{jb}^{\downarrow} \right)_i + \frac{1}{m} \sum_{i=1}^m \left(\tilde{Z}_{ja} \Phi_{ja} \right)_{ig} \left(\tilde{Z}_j v_j \right)_i \\ &= \frac{1}{m} \left(\Phi_{ja}^T \tilde{Z}_{ja}^T \sum_{b=1}^q \tilde{Z}_{jb} \Phi_{jb} \delta_{jb}^{\downarrow} \right)_g + \frac{1}{m} \left(\Phi_{ja}^T \tilde{Z}_{ja}^T \tilde{Z}_j v_j \right)_g \\ &= \frac{1}{m} \left(\Phi_{ja}^T \tilde{Z}_{ja}^T \tilde{Z}_j \Phi_j \delta_j^{\downarrow} \right)_g + \frac{1}{m} \left(\Phi_{ja}^T \tilde{Z}_{ja}^T \tilde{Z}_j v_j \right)_g, \end{aligned} \quad (73)$$

which implies that in (72), for $g = 1, \dots, \bar{h}$ and $a = 1, \dots, q$,

$$\frac{1}{m} \left(\Phi_{ja}^T \tilde{Z}_{ja}^T \tilde{Z}_j \Phi_j \delta_j^{\downarrow} \right)_g + \frac{1}{m} \left(\Phi_{ja}^T \tilde{Z}_{ja}^T \tilde{Z}_j v_j \right)_g + \frac{\tau_0^2 \lambda_n}{km} \frac{\delta_{jag}}{\mu_{ag}} = - \frac{\tau_0^2 \lambda_n}{km} \frac{\zeta_{0ag}}{\mu_{ag}}.$$

Let $\zeta_{0a}^{\downarrow} = (\zeta_{0a1}, \dots, \zeta_{0a\bar{h}})^T$ and $\zeta_0^{\downarrow} = (\zeta_{01}^{\downarrow}, \dots, \zeta_{0q}^{\downarrow})^T \in \mathbb{R}^{q\bar{h}}$. Stacking the last display over $g = 1, \dots, \bar{h}$ gives

$$\frac{1}{m} \Phi_{ja}^T \tilde{Z}_{ja}^T \tilde{Z}_j \Phi_j \delta_j^{\downarrow} + \frac{1}{m} \Phi_{ja}^T \tilde{Z}_{ja}^T \tilde{Z}_j v_j + \frac{\tau_0^2 \lambda_n}{km} M_a^{-1} \delta_{ja}^{\downarrow} = - \frac{\tau_0^2 \lambda_n}{km} M_a^{-1} \zeta_{0a}^{\downarrow}.$$

Then stacking this over $a = 1, \dots, q$ gives

$$\frac{1}{m} \Phi_j^T \tilde{Z}_j^T \tilde{Z}_j \Phi_j \delta_j^{\downarrow} + \frac{1}{m} \Phi_j^T \tilde{Z}_j^T \tilde{Z}_j v_j + \frac{\tau_0^2 \lambda_n}{km} M^{-1} \delta_j^{\downarrow} = - \frac{\tau_0^2 \lambda_n}{km} M^{-1} \zeta_0^{\downarrow},$$

which implies that

$$\left(\frac{1}{m} \Phi_j^T \tilde{Z}_j^T \tilde{Z}_j \Phi_j + \frac{\tau_0^2 \lambda_n}{km} M^{-1} \right) \delta_j^{\downarrow} = - \frac{\tau_0^2 \lambda_n}{km} M^{-1} \zeta_0^{\downarrow} - \frac{1}{m} \Phi_j^T \tilde{Z}_j^T \tilde{Z}_j v_j. \quad (74)$$

The proof is completed by showing that the right hand side expression in (74) gives an upper bound for $\|\delta_j^\downarrow\|_2^2$. Using the Ω matrix defined in Assumption (A.5), we define the matrix $Q = \left(I_{q\bar{h}} + \frac{\tau_0^2 \lambda_n}{km} \Omega^{-1} M^{-1}\right)^{1/2} \in \mathbb{R}^{q\bar{h} \times q\bar{h}}$. Then

$$\begin{aligned} \frac{1}{m} \Phi_j^\top \tilde{Z}_j^\top \tilde{Z}_j \Phi_j + \frac{\tau_0^2 \lambda_n}{km} M^{-1} &= \Omega + \frac{\tau_0^2 \lambda_n}{km} M^{-1} + \frac{1}{m} \Phi_j^\top \tilde{Z}_j^\top \tilde{Z}_j \Phi_j - \Omega \\ &= \Omega Q \left\{ I_{q\bar{h}} + Q^{-1} \Omega^{-1} \left(\frac{1}{m} \Phi_j^\top \tilde{Z}_j^\top \tilde{Z}_j \Phi_j - \Omega \right) Q^{-1} \right\} Q. \end{aligned}$$

and using this in (74) gives

$$\begin{aligned} &\left\{ I_{q\bar{h}} + Q^{-1} \Omega^{-1} \left(\frac{1}{m} \Phi_j^\top \tilde{Z}_j^\top \tilde{Z}_j \Phi_j - \Omega \right) Q^{-1} \right\} Q \delta_j^\downarrow \\ &= -\frac{\tau_0^2 \lambda_n}{km} Q^{-1} \Omega^{-1} M^{-1} \zeta_0^\downarrow - \frac{1}{m} Q^{-1} \Omega^{-1} \Phi_j^\top \tilde{Z}_j^\top \tilde{Z}_j v_j. \end{aligned} \quad (75)$$

Now we define the event

$$\mathcal{E}_{j1} = \left\{ \left\| Q^{-1} \Omega^{-1} \left(\frac{1}{m} \Phi_j^\top \tilde{Z}_j^\top \tilde{Z}_j \Phi_j - \Omega \right) Q^{-1} \right\| \leq 1/2 \right\}, \quad (76)$$

with the randomness in $\{u_{j1}, \dots, u_{jm}\}$. We have that

$$I_{q\bar{h}} + Q^{-1} \Omega^{-1} \left(\frac{1}{m} \Phi_j^\top \tilde{Z}_j^\top \tilde{Z}_j \Phi_j - \Omega \right) Q^{-1} \succeq \frac{1}{2} I_{q\bar{h}} \quad (77)$$

whenever \mathcal{E}_{j1} occurs. It is also clear that $Q \succeq I_{q\bar{h}}$. Therefore, when \mathcal{E}_{j1} occurs, (75) implies that

$$\begin{aligned} \|\delta_j^\downarrow\|_2^2 &\leq 4 \left\| \frac{\tau_0^2 \lambda_n}{km} Q^{-1} \Omega^{-1} M^{-1} \zeta_0^\downarrow + \frac{1}{m} Q^{-1} \Omega^{-1} \Phi_j^\top \tilde{Z}_j^\top \tilde{Z}_j v_j \right\|_2^2 \\ &\leq 8 \left\| \frac{\tau_0^2 \lambda_n}{km} Q^{-1} \Omega^{-1} M^{-1} \zeta_0^\downarrow \right\|_2^2 + 8 \left\| \frac{1}{m} Q^{-1} \Omega^{-1} \Phi_j^\top \tilde{Z}_j^\top \tilde{Z}_j v_j \right\|_2^2, \end{aligned} \quad (78)$$

where the last inequality follows because $(a+b)^2 \leq 2a^2 + 2b^2$ for any $a, b \in \mathbb{R}$.

For the first term on the right hand side of (78), we have that

$$\begin{aligned} &\left\| \frac{\tau_0^2 \lambda_n}{km} Q^{-1} \Omega^{-1} M^{-1} \zeta_0^\downarrow \right\|_2^2 \leq \left(\frac{\tau_0^2 \lambda_n}{km} \right)^2 \zeta_0^{\downarrow\top} \left\{ M \Omega \left(I_{q\bar{h}} + \frac{\tau_0^2 \lambda_n}{km} \Omega^{-1} M^{-1} \right) \Omega M \right\}^{-1} \zeta_0^\downarrow \\ &= \left(\frac{\tau_0^2 \lambda_n}{km} \right)^2 \zeta_0^{\downarrow\top} \left(M \Omega^2 M + \frac{\tau_0^2 \lambda_n}{km} \Omega M \right)^{-1} \zeta_0^\downarrow \\ &\leq \left(\frac{\tau_0^2 \lambda_n}{km} \right)^2 \zeta_0^{\downarrow\top} \left(\frac{\tau_0^2 \lambda_n}{km} \Omega M \right)^{-1} \zeta_0^\downarrow = \frac{\tau_0^2 \lambda_n}{\underline{c}_\Omega km} \sum_{a=1}^q \sum_{h=1}^{\bar{h}} \frac{\zeta_{0h}^2}{\mu_{ah}} \leq \frac{\tau_0^2 \lambda_n}{\underline{c}_\Omega km} \|\nu_0\|_{\mathbb{H}}^2. \end{aligned} \quad (79)$$

For the second term on right hand side of (78), it is equal to

$$\frac{1}{m} Q^{-1} \Omega^{-1} \Phi_j^\top \tilde{Z}_j^\top \tilde{Z}_j v_j = \left(M + \frac{\tau_0^2 \lambda_n}{km} \Omega^{-1} \right)^{-1/2} \cdot \frac{1}{m} M^{1/2} \Omega^{-1} \Phi_j^\top \tilde{Z}_j^\top \tilde{Z}_j v_j. \quad (80)$$

The first term in (80) has bounded matrix operator norm by Assumption (A.5):

$$\left\| \left(M + \frac{\tau_0^2 \lambda_n}{km} \Omega^{-1} \right)^{-1/2} \right\| = \max_{1 \leq a \leq q, 1 \leq h \leq \bar{h}} \frac{1}{\sqrt{\mu_{ah} + \frac{\tau_0^2 \lambda_n}{\bar{c}_\Omega km}}} \leq \sqrt{\frac{\bar{c}_\Omega km}{\tau_0^2 \lambda_n}}. \quad (81)$$

For the second term in (80), we repeatedly apply $|\varphi_{ai}| \leq C_\varphi$ and $|\tilde{Z}_a(u_{ji})| \leq \tilde{C}_Z$ in Assumptions (A.3) and (A.5) to obtain that

$$\begin{aligned} & \left\| \frac{1}{m} M^{1/2} \Omega^{-1} \Phi_j^T \tilde{Z}_j^T \tilde{Z}_j v_j \right\|_2^2 \leq \frac{\bar{c}_\Omega^{-2}}{m^2} \left(\tilde{Z}_j v_j \right)^T \tilde{Z}_j \Phi_j M \Phi_j^T \tilde{Z}_j^T \left(\tilde{Z}_j v_j \right) \\ &= \frac{\bar{c}_\Omega^{-2}}{m^2} \sum_{a=1}^q \left(\tilde{Z}_j v_j \right)^T \left(\tilde{Z}_{ja} \Phi_{ja} M_a \Phi_{ja}^T \tilde{Z}_{ja} \right) \left(\tilde{Z}_j v_j \right) \\ &\leq \frac{\bar{c}_\Omega^{-2}}{m^2} \sum_{a=1}^q \sum_{g=1}^{\bar{h}} \mu_{ag} \left\{ \sum_{i=1}^m \tilde{Z}_a(u_{ji}) \varphi_{ag}(u_{ji}) \left(\tilde{Z}_j v_j \right)_i \right\}^2 \\ &\leq \frac{C_\varphi^2 \tilde{C}_Z^2 \bar{c}_\Omega^{-2}}{m^2} \sum_{a=1}^q \sum_{g=1}^{\bar{h}} \mu_{ag} \left\{ \sum_{i=1}^m \left| \left(\tilde{Z}_j v_j \right)_i \right| \right\}^2 \\ &= \frac{C_\varphi^2 \tilde{C}_Z^2 \bar{c}_\Omega^{-2}}{m^2} \sum_{a=1}^q \sum_{g=1}^{\bar{h}} \mu_{ag} \left\{ \sum_{i=1}^m \sum_{b=1}^q \left| \tilde{Z}_b(u_{ji}) v_{jbi} \right| \right\}^2 \\ &\leq \frac{C_\varphi^2 \tilde{C}_Z^4 \bar{c}_\Omega^{-2}}{m^2} \sum_{a=1}^q \sum_{g=1}^{\bar{h}} \mu_{ag} \left\{ \sum_{i=1}^m \sum_{b=1}^q \sum_{h=\bar{h}+1}^{\infty} |\delta_{jbh} \varphi_{bh}(u_{ji})| \right\}^2 \\ &\leq C_\varphi^4 \tilde{C}_Z^4 \bar{c}_\Omega^{-2} \left(\sum_{a=1}^q \sum_{g=1}^{\bar{h}} \mu_{ag} \right) \left(\sum_{b=1}^q \sum_{h=\bar{h}+1}^{\infty} |\delta_{jbh}| \right)^2 \\ &\stackrel{(i)}{\leq} C_\varphi^4 \tilde{C}_Z^4 \bar{c}_\Omega^{-2} \left(\sum_{a=1}^q \text{tr}(\rho_a) \right) \cdot \left(\sum_{b=1}^q \sum_{h=\bar{h}+1}^{\infty} \mu_{bh} \right) \left(\sum_{b=1}^q \sum_{h=\bar{h}+1}^{\infty} \frac{\delta_{jbh}^2}{\mu_{bh}} \right) \\ &= C_\varphi^4 \tilde{C}_Z^4 \bar{c}_\Omega^{-2} \left(\sum_{a=1}^q \text{tr}(\rho_a) \right) \left(\sum_{a=1}^q \text{tr}(\rho_{a, \bar{h}}) \right) \left(\sum_{b=1}^q \left\| \mathbb{E}_{\nu_{jb, y_j | u_j}}(\Delta_{jb}) \right\|_{\mathbb{H}_b}^2 \right) \\ &= C_\varphi^4 \tilde{C}_Z^4 \bar{c}_\Omega^{-2} \text{Tr}(\rho) \text{Tr}(\rho, \bar{h}) \left\| \mathbb{E}_{\nu_{j, y_j | u_j}}(\Delta_j) \right\|_{\mathbb{H}}^2 \\ &\stackrel{(ii)}{\leq} C_\varphi^4 \tilde{C}_Z^4 \bar{c}_\Omega^{-2} \|\nu_0\|_{\mathbb{H}}^2 \text{Tr}(\rho) \text{Tr}(\rho, \bar{h}), \end{aligned} \quad (82)$$

where (i) is from the Cauchy-Schwarz inequality and the definition of $\text{tr}(\rho_a)$, and (ii) is from the relation (66). Combining (80), (81), and (82) leads to

$$\left\| \frac{1}{m} Q^{-1} \Omega^{-1} \Phi_j^T \tilde{Z}_j^T \tilde{Z}_j v_j \right\|_2^2 \leq \frac{\bar{c}_\Omega km}{\tau_0^2 \lambda_n} C_\varphi^4 \tilde{C}_Z^4 \bar{c}_\Omega^{-2} \|\nu_0\|_{\mathbb{H}}^2 \text{Tr}(\rho) \text{Tr}(\rho, \bar{h}). \quad (83)$$

Finally we combine these results. Note that (68) implies that

$$\|\delta_j^\downarrow\|_2^2 = \sum_{a=1}^q \|\delta_{ja}^\downarrow\|_2^2 \leq \sum_{a=1}^q \text{tr}(\rho_a) \|\nu_{0a}\|_{\mathbb{H}_a}^2 \leq \text{Tr}(\rho) \|\nu_0\|_{\mathbb{H}}^2. \quad (84)$$

(69) and (66) together imply that

$$\|\delta_j^\uparrow\|_2^2 = \sum_{a=1}^q \|\delta_{ja}^\uparrow\|_2^2 \leq \sum_{a=1}^q \mu_{a(\bar{h}+1)} \left\| \mathbb{E}_{\nu_{ja}, y_j | u_j}(\Delta_{ja}) \right\|_{\mathbb{H}_a}^2 \leq \mu_{(\bar{h}+1)*} \|\nu_0\|_{\mathbb{H}}^2. \quad (85)$$

Based on (56) and the definition of Δ_j , we combine (84), (69), (78), (79), (83), and Lemma 5 to obtain that

$$\begin{aligned} & \mathbb{E}_{u^*} \mathbb{E}_{u_j} \left\| \mathbb{E}_{\nu_j, y_j | u_j} \{ \nu_j(u^*) \} - \nu_0(u^*) \right\|_2^2 = \mathbb{E}_{u^*} \mathbb{E}_{u_j} \left\| \mathbb{E}_{\nu_j, y_j | u_j}(\Delta_j) \right\|_2^2 \\ & \leq \mathbb{E}_{u_j} \left(\|\delta_j^\downarrow\|_2^2 + \|\delta_j^\uparrow\|_2^2 \right) \\ & = \mathbb{E}_{u_j} \left\{ \|\delta_j^\downarrow\|_2^2 \mathbf{1}(\mathcal{E}_{j1}) + \|\delta_j^\downarrow\|_2^2 \mathbf{1}(\mathcal{E}_{j1}^c) + \|\delta_j^\uparrow\|_2^2 \right\} \\ & \leq \mathbb{E}_{u_j} \left\{ \|\delta_j^\downarrow\|_2^2 \mathbf{1}(\mathcal{E}_{j1}) \right\} + \text{Tr}(\rho) \|\nu_0\|_{\mathbb{H}}^2 \mathbb{P}_{u_j}(\mathcal{E}_{j1}^c) + \mathbb{E}_{u_j} \left(\|\delta_j^\uparrow\|_2^2 \right) \\ & \leq 8 \frac{\tau_0^2 \lambda_n}{\underline{c}_\Omega km} \|\nu_0\|_{\mathbb{H}}^2 + 8 \frac{km}{\tau_0^2 \lambda_n} C_\varphi^4 \tilde{C}_Z^4 \bar{c}_\Omega \underline{c}_\Omega^{-2} \|\nu_0\|_{\mathbb{H}}^2 \text{Tr}(\rho) \text{Tr}(\rho, \bar{h}) \\ & \quad + 2q\bar{h} \|\nu_0\|_{\mathbb{H}}^2 \text{Tr}(\rho) \exp \left\{ -\frac{m}{8(B^2 + B)} \right\} + \mu_{(\bar{h}+1)*} \|\nu_0\|_{\mathbb{H}}^2, \end{aligned}$$

where $B = C_\varphi^2 \tilde{C}_Z^2 \underline{c}_\Omega^{-1} q \Lambda (\tau_0^2 \lambda_n / (\underline{c}_\Omega km), \bar{h}) + 1$. This proves the conclusion. \blacksquare

Lemma 3 *Suppose that Assumptions (A.1)–(A.6) hold. Then for every $j = 1, \dots, k$,*

$$\begin{aligned} & \mathbb{E}_{u^*} \mathbb{E}_{u_j} \text{tr} \left(\text{Var}_{y_j | u_j} \left[\mathbb{E}_{\nu_j | y_j, u_j} \{ \nu_j(u^*) \} \right] \right) \\ & \leq 12 \frac{\tau_0^2 \lambda_n}{\underline{c}_\Omega km} \|\nu_0\|_{\mathbb{H}}^2 + \frac{24km}{\tau_0^2 \lambda_n} C_\varphi^4 \tilde{C}_Z^4 \bar{c}_\Omega \underline{c}_\Omega^{-2} \text{Tr}(\rho) \text{Tr}(\rho, \bar{h}) \left(\frac{km}{\lambda_n} + 2\|\nu_0\|_{\mathbb{H}}^2 \right) \\ & \quad + 12 \frac{C_\varphi^2 \tilde{C}_Z^2 \underline{c}_\Omega^{-2} \tau_0^2 q}{m} \Lambda(\tau_0^2 \lambda_n / (\bar{c}_\Omega km), \bar{h}) + 2\mu_{(\bar{h}+1)*} \left(\frac{km}{\lambda_n} + 2\|\nu_0\|_{\mathbb{H}}^2 \right) \\ & \quad + 4q\bar{h} \text{Tr}(\rho) \left(\frac{km}{\lambda_n} + 2\|\nu_0\|_{\mathbb{H}}^2 \right) \exp \left\{ -\frac{m}{8(B^2 + B)} \right\}, \end{aligned}$$

where $B = C_\varphi^2 \tilde{C}_Z^2 \underline{c}_\Omega^{-1} q \Lambda (\tau_0^2 \lambda_n / (\underline{c}_\Omega km), \bar{h}) + 1$.

Proof [of Lemma 3] We use the same notations as in the proof of Lemma 2. We further expand the functions $\Delta_{ja}(\cdot)$ defined in (60) for $a = 1, \dots, q$ and $j = 1, \dots, k$ as

$$\Delta_{ja}(\cdot) = \sum_{i=1}^{\infty} \tilde{\delta}_{jai} \varphi_{ai}(\cdot), \quad \Delta_{ja}^\downarrow(\cdot) = \sum_{i=1}^{\bar{h}} \tilde{\delta}_{jai} \varphi_{ai}(\cdot), \quad \Delta_{ja}^\uparrow(\cdot) = \sum_{i=\bar{h}+1}^{\infty} \tilde{\delta}_{jai} \varphi_{ai}(\cdot),$$

$$\begin{aligned}\tilde{\delta}_{ja}^\downarrow &= (\tilde{\delta}_{ja1}, \dots, \tilde{\delta}_{ja\bar{h}})^\top, & \tilde{\delta}_{ja}^\uparrow &= (\tilde{\delta}_{ja(\bar{h}+1)}, \dots, \tilde{\delta}_{ja\infty})^\top, \\ \tilde{\delta}_j^\downarrow &= (\tilde{\delta}_{j1}^\downarrow, \dots, \tilde{\delta}_{jq}^\downarrow)^\top, & \tilde{\delta}_j^\uparrow &= (\tilde{\delta}_{j1}^\uparrow, \dots, \tilde{\delta}_{jq}^\uparrow)^\top.\end{aligned}$$

From (57), we can see that

$$\begin{aligned}& \mathbb{E}_{u^*} \operatorname{tr} \left(\operatorname{Var}_{y_j|u_j} \left[\mathbb{E}_{\nu_j|y_j, u_j} \{ \nu_j(u^*) \} \right] \right) \\ &= \mathbb{E}_{u^*} \operatorname{tr} \left[\operatorname{Var}_{y_j|u_j} \left\{ R_j(u^*) \tilde{Z}_j^\top \left(\tilde{Z}_j \tilde{R}_{jj} \tilde{Z}_j^\top + \frac{\tau_0^2 \lambda_n}{k} I_m \right)^{-1} y_j \right\} \right] \\ &\leq \mathbb{E}_{u^*} \mathbb{E}_{y_j|u_j} \left\| R_j(u^*) \tilde{Z}_j^\top \left(\tilde{Z}_j \tilde{R}_{jj} \tilde{Z}_j^\top + \frac{\tau_0^2 \lambda_n}{k} I_m \right)^{-1} y_j - \nu_0(u^*) \right\|_2^2 \\ &= \mathbb{E}_{u^*} \mathbb{E}_{y_j|u_j} \|\Delta_j(u^*)\|_2^2 = \mathbb{E}_{u^*} \mathbb{E}_{y_j|u_j} \|\Delta_j^\downarrow(u^*)\|_2^2 + \mathbb{E}_{u^*} \mathbb{E}_{y_j|u_j} \|\Delta_j^\uparrow(u^*)\|_2^2 \\ &= \mathbb{E}_{y_j|u_j} \|\tilde{\delta}_j^\downarrow\|_2^2 + \mathbb{E}_{y_j|u_j} \|\tilde{\delta}_j^\uparrow\|_2^2.\end{aligned}\tag{86}$$

Therefore, we will find an upper bound for $\mathbb{E}_{u^*} \mathbb{E}_{y_j|u_j} \|\Delta_j(u^*)\|_2^2$ in the following. We start with finding a rough upper bound for $\mathbb{E}_{y_j|u_j} \|\Delta_j\|_{\mathbb{H}}^2$. Using the definition of $\hat{\nu}_j$ in (60) and the optimizer property in (61), we have that

$$\begin{aligned}\|\hat{\nu}_j\|_{\mathbb{H}}^2 &\stackrel{(i)}{\leq} \sum_{i=1}^m \frac{\left\{ y(u_{ji}) - \tilde{Z}(u_{ji}) \hat{\nu}_j(u_{ji}) \right\}^2}{\tau_0^2 \lambda_n / k} + \|\hat{\nu}_j\|_{\mathbb{H}}^2 \\ &\stackrel{(ii)}{\leq} \sum_{i=1}^m \frac{\left\{ y(u_{ji}) - \tilde{Z}(u_{ji}) \nu_0(u_{ji}) \right\}^2}{\tau_0^2 \lambda_n / k} + \|\nu_0\|_{\mathbb{H}}^2 \\ &\stackrel{(iii)}{\leq} \sum_{i=1}^m \frac{\{\epsilon(u_{ji})\}^2}{\tau_0^2 \lambda_n / k} + \|\nu_0\|_{\mathbb{H}}^2,\end{aligned}\tag{87}$$

where (i) follows because the term inside the summation is non-negative, (ii) follows because $\hat{\nu}_j$ minimizes the objective, and (iii) follows from our model assumption. Since the error variance is τ_0^2 , (87) implies that

$$\begin{aligned}\mathbb{E}_{y_j|u_j} \|\Delta_j\|_{\mathbb{H}}^2 &\leq 2 \mathbb{E}_{y_j|u_j} \|\hat{\nu}_j\|_{\mathbb{H}}^2 + 2 \mathbb{E}_{y_j|u_j} \|\nu_0\|_{\mathbb{H}}^2 \\ &\leq 2 \mathbb{E}_{y_j|u_j} \left[\sum_{i=1}^m \frac{\{\epsilon(u_{ji})\}^2}{\tau_0^2 \lambda_n / k} \right] + 4 \|\nu_0\|_{\mathbb{H}}^2 \\ &\leq \frac{2km}{\lambda_n} + 4 \|\nu_0\|_{\mathbb{H}}^2.\end{aligned}\tag{88}$$

Using this bound for $\mathbb{E}_{y_j|u_j} \|\Delta_j\|_{\mathbb{H}}^2$, we can find an upper bound for $\mathbb{E}_{y_j|u_j} \|\delta_j^\uparrow\|_2^2$:

$$\mathbb{E}_{y_j|u_j} \|\tilde{\delta}_j^\uparrow\|_2^2 = \sum_{a=1}^q \sum_{i=\bar{h}+1}^{\infty} \mathbb{E}_{y_j|u_j} (\tilde{\delta}_{jai}^2)$$

$$\begin{aligned}
 &= \sum_{a=1}^q \mu_{a(\bar{h}+1)} \sum_{i=\bar{h}+1}^{\infty} \frac{\mathbb{E}_{y_j|u_j}(\tilde{\delta}_{jai}^2)}{\mu_{a(\bar{h}+1)}} \leq \sum_{a=1}^q \mu_{a(\bar{h}+1)} \sum_{i=\bar{h}+1}^{\infty} \frac{\mathbb{E}_{y_j|u_j}(\tilde{\delta}_{jai}^2)}{\mu_{ai}} \\
 &= \sum_{a=1}^q \mu_{a(\bar{h}+1)} \mathbb{E}_{y_j|u_j} \|\Delta_{ja}^\uparrow\|_{\mathbb{H}_a}^2 \leq \sum_{a=1}^q \mu_{a(\bar{h}+1)} \mathbb{E}_{y_j|u_j} \|\Delta_{ja}\|_{\mathbb{H}_a}^2 \\
 &\leq \mu_{(\bar{h}+1)*} \left(\frac{2km}{\lambda_n} + 4\|\nu_0\|_{\mathbb{H}}^2 \right), \tag{89}
 \end{aligned}$$

and also an upper bound for $\mathbb{E}_{y_j|u_j} \|\tilde{\delta}_j\|_2^2$:

$$\begin{aligned}
 \mathbb{E}_{y_j|u_j} \|\tilde{\delta}_j\|_2^2 &= \sum_{a=1}^q \sum_{i=1}^{\infty} \mathbb{E}_{y_j|u_j}(\tilde{\delta}_{jai}^2) \\
 &\leq \sum_{a=1}^q \mu_{a1} \sum_{i=1}^{\infty} \frac{\mathbb{E}_{y_j|u_j}(\tilde{\delta}_{ji}^2)}{\mu_{ai}} = \sum_{a=1}^q \mu_{a1} \mathbb{E}_{y_j|u_j} \|\Delta_j^\uparrow\|_{\mathbb{H}_a}^2 \leq \mu_{1*} \mathbb{E}_{y_j|u_j} \|\Delta_j\|_{\mathbb{H}}^2 \\
 &\leq \text{Tr}(\rho) \left(\frac{2km}{\lambda_n} + 4\|\nu_0\|_{\mathbb{H}}^2 \right), \tag{90}
 \end{aligned}$$

Now we find an upper bound for $\mathbb{E}_{y_j|u_j} \|\tilde{\delta}_j^\downarrow\|_2^2$. Define the error vectors

$$\begin{aligned}
 \tilde{v}_{jai} &= \sum_{h=\bar{h}+1}^{\infty} \tilde{\delta}_{jai} \varphi_{ah}(u_{ji}), \\
 \tilde{v}_{ja} &= (\tilde{v}_{ja1}, \dots, \tilde{v}_{jam})^\top \in \mathbb{R}^m, \quad i = 1, \dots, m, \quad a = 1, \dots, q, \\
 \tilde{v}_j &= (\tilde{v}_{j1}^\top, \dots, \tilde{v}_{jq}^\top)^\top \in \mathbb{R}^{qm}.
 \end{aligned}$$

Now we use an argument similar to the derivation of (63), (64), (65), (71), (72), (73), and (74). Instead of taking the $\mathbb{E}_{\nu_j, y_j|u_j}$ as in (65), we do not take this expectation and keep the error term $\epsilon(u_{ji})$ all the way along the derivation. We can obtain the following relation similar to (74):

$$\left(\frac{1}{m} \Phi_j^\top \tilde{Z}_j^\top \tilde{Z}_j \Phi_j + \frac{\tau_0^2 \lambda_n}{km} M^{-1} \right) \tilde{\delta}_j^\downarrow = -\frac{\tau_0^2 \lambda_n}{km} M^{-1} \zeta_0^\downarrow - \frac{1}{m} \Phi_j^\top \tilde{Z}_j^\top \tilde{Z}_j \tilde{v}_j + \frac{1}{m} \Phi_j^\top \tilde{Z}_j^\top \epsilon_j. \tag{91}$$

We use the same Q matrix as defined in the proof of Lemma 2. Then (91) can be rewritten as

$$\begin{aligned}
 &\left\{ I_{q\bar{h}} + Q^{-1} \Omega^{-1} \left(\frac{1}{m} \Phi_j^\top \tilde{Z}_j^\top \tilde{Z}_j \Phi_j - \Omega \right) Q^{-1} \right\} Q \tilde{\delta}_j^\downarrow \\
 &= -\frac{\tau_0^2 \lambda_n}{km} Q^{-1} \Omega^{-1} M^{-1} \zeta_0^\downarrow - \frac{1}{m} Q^{-1} \Omega^{-1} \Phi_j^\top \tilde{Z}_j^\top \tilde{Z}_j \tilde{v}_j + \frac{1}{m} Q^{-1} \Omega^{-1} \Phi_j^\top \tilde{Z}_j^\top \epsilon_j. \tag{92}
 \end{aligned}$$

On the event \mathcal{E}_{j1} defined as in (76), using (77) and the fact $Q \succeq I_{q\bar{h}}$ (92) imply that

$$\mathbb{E}_{y_j|u_j} \|\tilde{\delta}_j^\downarrow\|_2^2 \leq \mathbb{E}_{y_j|u_j} \|Q \tilde{\delta}_j^\downarrow\|_2^2$$

$$\begin{aligned}
 &\leq 4 \mathbb{E}_{y_j|u_j} \left\| -\frac{\tau_0^2 \lambda_n}{km} Q^{-1} \Omega^{-1} M^{-1} \zeta_0^\downarrow - \frac{1}{m} Q^{-1} \Omega^{-1} \Phi_j^T \tilde{Z}_j^T \tilde{Z}_j \tilde{v}_j + \frac{1}{m} Q^{-1} \Omega^{-1} \Phi_j^T \tilde{Z}_j^T \epsilon_j \right\|_2^2 \\
 &\leq 12 \left\| \frac{\tau_0^2 \lambda_n}{km} Q^{-1} \Omega^{-1} M^{-1} \zeta_0^\downarrow \right\|_2^2 + 12 \mathbb{E}_{y_j|u_j} \left\| \frac{1}{m} Q^{-1} \Omega^{-1} \Phi_j^T \tilde{Z}_j^T \tilde{Z}_j \tilde{v}_j \right\|_2^2 \\
 &\quad + 12 \mathbb{E}_{y_j|u_j} \left\| \frac{1}{m} Q^{-1} \Omega^{-1} \Phi_j^T \tilde{Z}_j^T \epsilon_j \right\|_2^2, \tag{93}
 \end{aligned}$$

where the last inequality follows because $(a + b + c)^2 \leq 3a^2 + 3b^2 + 3c^2$ for any $a, b, c \in \mathbb{R}$. We bound the three terms on the right hand side of (93). The first term can be bounded as in (79). The second term can be bounded similar to the proof of Lemma 2: By Assumption (A.3), we have that

$$\frac{1}{m} Q^{-1} \Omega^{-1} \Phi_j^T \tilde{Z}_j^T \tilde{Z}_j \tilde{v}_j = \left(M + \frac{\tau_0^2 \lambda_n}{km} \Omega^{-1} \right)^{-1/2} \cdot \frac{1}{m} M^{1/2} \Omega^{-1} \Phi_j^T \tilde{Z}_j^T \tilde{Z}_j \tilde{v}_j. \tag{94}$$

The first term in (94) has bounded matrix operator norm in (81). For the second term in (80), we repeatedly apply Assumptions (A.3) and (A.5) to obtain that

$$\begin{aligned}
 &\mathbb{E}_{y_j|u_j} \left\| \frac{1}{m} M^{1/2} \Omega^{-1} \Phi_j^T \tilde{Z}_j^T \tilde{Z}_j \tilde{v}_j \right\|_2^2 \\
 &= \frac{C_\Omega^{-2}}{m^2} \mathbb{E}_{y_j|u_j} \sum_{a=1}^q \left(\tilde{Z}_j v_j \right)^T \left(\tilde{Z}_{ja} \Phi_{ja} M_a \Phi_{ja}^T \tilde{Z}_{ja} \right) \left(\tilde{Z}_j \tilde{v}_j \right) \\
 &\leq \frac{C_\Omega^{-2}}{m^2} \mathbb{E}_{y_j|u_j} \sum_{a=1}^q \sum_{g=1}^{\bar{h}} \mu_{ag} \left\{ \sum_{i=1}^m \tilde{Z}_a(u_{ji}) \varphi_{ag}(u_{ji}) \left(\tilde{Z}_j \tilde{v}_j \right)_i \right\}^2 \\
 &\leq \frac{C_\varphi^2 \tilde{C}_Z^2 C_\Omega^{-2}}{m^2} \mathbb{E}_{y_j|u_j} \sum_{a=1}^q \sum_{g=1}^{\bar{h}} \mu_{ag} \left\{ \sum_{i=1}^m \left| \left(\tilde{Z}_j \tilde{v}_j \right)_i \right| \right\}^2 \\
 &= \frac{C_\varphi^2 \tilde{C}_Z^2 C_\Omega^{-2}}{m^2} \mathbb{E}_{y_j|u_j} \sum_{a=1}^q \sum_{g=1}^{\bar{h}} \mu_{ag} \left\{ \sum_{i=1}^m \sum_{b=1}^q \left| \tilde{Z}_b(u_{ji}) \tilde{v}_{jbi} \right| \right\}^2 \\
 &\leq \frac{C_\varphi^2 \tilde{C}_Z^4 C_\Omega^{-2}}{m^2} \mathbb{E}_{y_j|u_j} \sum_{a=1}^q \sum_{g=1}^{\bar{h}} \mu_{ag} \left\{ \sum_{i=1}^m \sum_{b=1}^q \sum_{h=\bar{h}+1}^\infty \left| \tilde{\delta}_{jbb} \varphi_{bh}(u_{ji}) \right| \right\}^2 \\
 &\leq C_\varphi^4 \tilde{C}_Z^4 C_\Omega^{-2} \left(\sum_{a=1}^q \sum_{g=1}^{\bar{h}} \mu_{ag} \right) \mathbb{E}_{y_j|u_j} \left(\sum_{b=1}^q \sum_{h=\bar{h}+1}^\infty \left| \tilde{\delta}_{jbb} \right| \right)^2 \\
 &\stackrel{(i)}{\leq} C_\varphi^4 \tilde{C}_Z^4 C_\Omega^{-2} \left(\sum_{a=1}^q \text{tr}(\rho_a) \right) \cdot \left(\sum_{b=1}^q \sum_{h=\bar{h}+1}^\infty \mu_{bh} \right) \mathbb{E}_{y_j|u_j} \left(\sum_{b=1}^q \sum_{h=\bar{h}+1}^\infty \frac{\tilde{\delta}_{jbb}^2}{\mu_{bh}} \right) \\
 &= C_\varphi^4 \tilde{C}_Z^4 C_\Omega^{-2} \text{Tr}(\rho) \left(\sum_{a=1}^q \text{tr}(\rho_{a,\bar{h}}) \right) \mathbb{E}_{y_j|u_j} \left(\sum_{b=1}^q \|\Delta_{jb}\|_{\mathbb{H}_b}^2 \right) \\
 &= C_\varphi^4 \tilde{C}_Z^4 C_\Omega^{-2} \text{Tr}(\rho) \text{Tr}(\rho, \bar{h}) \mathbb{E}_{y_j|u_j} \|\Delta_j\|_{\mathbb{H}}^2
 \end{aligned}$$

$$\stackrel{(ii)}{\leq} C_\varphi^4 \tilde{C}_Z^4 \underline{c}_\Omega^{-2} \text{Tr}(\rho) \text{Tr}(\rho, \bar{h}) \left(\frac{2km}{\lambda_n} + 4\|\nu_0\|_{\mathbb{H}}^2 \right), \quad (95)$$

where (i) is from the Cauchy-Schwarz inequality and the definition of $\text{tr}(C)$, and (ii) is from the relation (88). Combining (94), (81), and (95) leads to

$$\mathbb{E}_{y_j|u_j} \left\| \frac{1}{m} Q^{-1} \Omega^{-1} \Phi_j^T \tilde{Z}_j^T \tilde{Z}_j \tilde{\nu}_j \right\|_2^2 \leq \frac{2km}{\tau_0^2 \lambda_n} C_\varphi^4 \tilde{C}_Z^4 \bar{c}_\Omega \underline{c}_\Omega^{-2} \text{Tr}(\rho) \text{Tr}(\rho, \bar{h}) \left(\frac{km}{\lambda_n} + 2\|\nu_0\|_{\mathbb{H}}^2 \right). \quad (96)$$

For the third term in (93), by Assumptions (A.3) and (A.5), we have that

$$\begin{aligned} & \mathbb{E}_{y_j|u_j} \left\| \frac{1}{m} Q^{-1} \Omega^{-1} \Phi_j^T \tilde{Z}_j^T \epsilon_j \right\|_2^2 \\ & \leq \frac{\underline{c}_\Omega^{-2}}{m^2} \mathbb{E}_{y_j|u_j} \left\{ \epsilon_j^T \tilde{Z}_j \Phi_j \left(I_{q\bar{h}} + \frac{\tau_0^2 \lambda_n}{km} \Omega^{-1} M^{-1} \right)^{-1} \Phi_j^T \tilde{Z}_j^T \epsilon_j \right\} \\ & \leq \frac{\underline{c}_\Omega^{-2}}{m^2} \mathbb{E}_{y_j|u_j} \sum_{a=1}^q \sum_{h=1}^{\bar{h}} \frac{1}{1 + \frac{\tau_0^2 \lambda_n}{\bar{c}_\Omega km \mu_{ah}}} \left\{ \sum_{i=1}^m \tilde{Z}_a(u_{ji}) \varphi_{ah}(u_{ji}) \epsilon(u_{ji}) \right\}^2 \\ & \stackrel{(i)}{=} \frac{\underline{c}_\Omega^{-2}}{m^2} \sum_{a=1}^q \sum_{h=1}^{\bar{h}} \frac{1}{1 + \frac{\tau_0^2 \lambda_n}{\bar{c}_\Omega km \mu_{ah}}} \mathbb{E}_{y_j|u_j} \left\{ \sum_{i=1}^m \tilde{Z}_a(u_{ji})^2 \varphi_{ah}(u_{ji})^2 \epsilon(u_{ji})^2 \right\} \\ & \leq \frac{C_\varphi^2 \tilde{C}_Z^2 \underline{c}_\Omega^{-2}}{m^2} \sum_{a=1}^q \sum_{h=1}^{\bar{h}} \frac{1}{1 + \frac{\tau_0^2 \lambda_n}{\bar{c}_\Omega km \mu_{ah}}} \mathbb{E}_{y_j|u_j} \left\{ \sum_{i=1}^m \epsilon(u_{ji})^2 \right\} \\ & \leq \frac{C_\varphi^2 \tilde{C}_Z^2 \underline{c}_\Omega^{-2} \tau_0^2 q}{m} \Lambda(\tau_0^2 \lambda_n / (\bar{c}_\Omega km), \bar{h}). \end{aligned} \quad (97)$$

where (i) follows from the independence between $\{\epsilon_{j1}, \dots, \epsilon_{jm}\}$. Therefore, we can obtain that

$$\begin{aligned} & \mathbb{E}_{u^*} \mathbb{E}_{u_j} \text{tr} \left(\text{Var}_{y_j|u_j} \left[\mathbb{E}_{\nu_j|y_j, u_j} \{ \nu_j(u^*) \} \right] \right) \\ & \stackrel{(i)}{\leq} \mathbb{E}_{y_j|u_j} \|\tilde{\delta}_j^\downarrow\|_2^2 + \mathbb{E}_{y_j|u_j} \|\tilde{\delta}_j^\uparrow\|_2^2 \\ & = \mathbb{E}_{y_j|u_j} \left\{ \|\tilde{\delta}_j^\downarrow\|_2^2 \mathbf{1}(\mathcal{E}_{j1}) \right\} + \mathbb{E}_{y_j|u_j} \left\{ \|\tilde{\delta}_j^\downarrow\|_2^2 \mathbf{1}(\mathcal{E}_{j1}^c) \right\} + \mathbb{E}_{y_j|u_j} \|\tilde{\delta}_j^\uparrow\|_2^2 \\ & \stackrel{(ii)}{\leq} 12 \left\| \frac{\tau_0^2 \lambda_n}{km} Q^{-1} \Omega^{-1} M^{-1} \zeta_0^\downarrow \right\|_2^2 + 12 \mathbb{E}_{y_j|u_j} \left\| \frac{1}{m} Q^{-1} \Omega^{-1} \Phi_j^T \tilde{Z}_j^T \tilde{Z}_j \tilde{\nu}_j \right\|_2^2 \\ & \quad + 12 \mathbb{E}_{y_j|u_j} \left\| \frac{1}{m} Q^{-1} \Omega^{-1} \Phi_j^T \tilde{Z}_j^T \epsilon_j \right\|_2^2 + \mu_{(\bar{h}+1)^*} \left(\frac{2km}{\lambda_n} + 4\|\nu_0\|_{\mathbb{H}}^2 \right) \\ & \quad + \text{Tr}(\rho) \left(\frac{2km}{\lambda_n} + 4\|\nu_0\|_{\mathbb{H}}^2 \right) \mathbb{P}_{u_j}(\mathcal{E}_{j1}^c) \\ & \stackrel{(iii)}{\leq} 12 \frac{\tau_0^2 \lambda_n}{\underline{c}_\Omega km} \|\nu_0\|_{\mathbb{H}}^2 + \frac{24km}{\tau_0^2 \lambda_n} C_\varphi^4 \tilde{C}_Z^4 \bar{c}_\Omega \underline{c}_\Omega^{-2} \text{Tr}(\rho) \text{Tr}(\rho, \bar{h}) \left(\frac{km}{\lambda_n} + 2\|\nu_0\|_{\mathbb{H}}^2 \right) \\ & \quad + 12 \frac{C_\varphi^2 \tilde{C}_Z^2 \underline{c}_\Omega^{-2} \tau_0^2 q}{m} \Lambda(\tau_0^2 \lambda_n / (\bar{c}_\Omega km), \bar{h}) + 2\mu_{(\bar{h}+1)^*} \left(\frac{km}{\lambda_n} + 2\|\nu_0\|_{\mathbb{H}}^2 \right) \end{aligned}$$

$$+ 4q\bar{h} \operatorname{Tr}(\rho) \left(\frac{km}{\lambda_n} + 2\|\nu_0\|_{\mathbb{H}}^2 \right) \exp \left\{ -\frac{m}{8(B^2 + B)} \right\},$$

where $B = C_\varphi^2 \tilde{C}_Z^2 \underline{c}_\Omega^{-1} q \Lambda (\tau_0^2 \lambda_n / (\underline{c}_\Omega km), \bar{h}) + 1$, (i) is from (86), (ii) is from (89), (90), and (93), and (iii) is from (79), (96), (97), and Lemma 5. This completes the proof of Lemma 3. \blacksquare

Lemma 4 *Suppose that Assumptions (A.1)–(A.6) hold. Then for every $j = 1, \dots, k$,*

$$\begin{aligned} & \mathbb{E}_{u^*} \mathbb{E}_{y_j, u_j} \operatorname{tr} \left[\operatorname{Var}_{\nu_j | y_j, u_j} \{ \nu_j(u^*) \} \right] \\ & \leq \frac{5\tau_0^2 \lambda_n}{2\underline{c}_\Omega km} \Lambda (\tau_0^2 \lambda_n / (\underline{c}_\Omega km), \bar{h}) + \frac{4\tilde{C}_Z^2 km}{\tau_0^2 \lambda_n^2} \operatorname{Tr}(\rho, \bar{h}) \operatorname{Tr}(\rho) + \lambda_n^{-1} \operatorname{Tr}(\rho, \bar{h}) \\ & \quad + 2\lambda_n^{-1} q\bar{h} \operatorname{Tr}(\rho) \exp \left\{ -\frac{m}{8(B^2 + B)} \right\}, \end{aligned}$$

where $B = C_\varphi^2 \tilde{C}_Z^2 \underline{c}_\Omega^{-1} q \Lambda (\tau_0^2 \lambda_n / (\underline{c}_\Omega km), \bar{h}) + 1$.

Proof [of Lemma 4] For each $a = 1, \dots, q$, we have the eigen-decomposition $\rho_a(u, u') = \sum_{i=1}^{\infty} \mu_{ai} \varphi_{ai}(u) \varphi_{ai}(u')$ for $u, u' \in [0, 1]^d$. This together with the expression of $\operatorname{Var}_{\nu_j | y_j, u_j} \{ \nu_j(u^*) \}$ in (58) and the orthonormal property of $\{\varphi_{ai}\}_{i=1}^{\infty}$ imply that

$$\begin{aligned} & \mathbb{E}_{u^*} \mathbb{E}_{y_j, u_j} \operatorname{tr} \left[\operatorname{Var}_{\nu_j | y_j, u_j} \{ \nu_j(u^*) \} \right] \\ & = \lambda_n^{-1} \mathbb{E}_{u^*} \mathbb{E}_{y_j, u_j} \sum_{a=1}^q \left\{ \rho_a(u^*, u^*) - R_{ja}(u^*)^T \tilde{Z}_{ja}^T \left(\tilde{Z}_j \tilde{R}_{jj} \tilde{Z}_j^T + \frac{\tau_0^2 \lambda_n}{k} I_m \right)^{-1} \tilde{Z}_{ja} R_{ja}(u^*) \right\} \\ & = \lambda_n^{-1} \sum_{a=1}^q \sum_{h=1}^{\infty} \mu_{ah} \mathbb{E}_{u^*} \{ \varphi_{ah}(u^*)^2 \} \\ & \quad - \lambda_n^{-1} \mathbb{E}_{u_j} \sum_{a=1}^q \sum_{i=1}^m \sum_{i'=1}^m \sum_{h=1}^{\infty} \sum_{h'=1}^{\infty} \mu_{ah} \mu_{ah'} \left\{ \tilde{Z}_{ja}^T \left(\tilde{Z}_j \tilde{R}_{jj} \tilde{Z}_j^T + \frac{\tau_0^2 \lambda_n}{k} I_m \right)^{-1} \tilde{Z}_{ja} \right\}_{i'i''} \\ & \quad \times \left[\varphi_{ah}(u_{ji}) \varphi_{ah'}(u_{ji'}) \mathbb{E}_{u^*} \{ \varphi_{ah}(u^*) \varphi_{ah'}(u^*) \} \right] \\ & = \lambda_n^{-1} \sum_{a=1}^q \sum_{h=1}^{\infty} \mu_{ah} - \lambda_n^{-1} \mathbb{E}_{u_j} \sum_{a=1}^q \sum_{i=1}^m \sum_{i'=1}^m \sum_{h=1}^{\infty} \mu_{ah}^2 \varphi_{ah}(u_{ji}) \varphi_{ah}(u_{ji'}) \\ & \quad \times \left\{ \tilde{Z}_{ja}^T \left(\tilde{Z}_j \tilde{R}_{jj} \tilde{Z}_j^T + \frac{\tau_0^2 \lambda_n}{k} I_m \right)^{-1} \tilde{Z}_{ja} \right\}_{ii'} \\ & = \lambda_n^{-1} \sum_{a=1}^q \sum_{h=1}^{\bar{h}} \mu_{ah} - \lambda_n^{-1} \mathbb{E}_{u_j} \sum_{a=1}^q \sum_{i=1}^m \sum_{i'=1}^m \sum_{h=1}^{\bar{h}} \mu_{ah}^2 \varphi_{ah}(u_{ji}) \varphi_{ah}(u_{ji'}) \\ & \quad \times \left\{ \tilde{Z}_{ja}^T \left(\tilde{Z}_j \tilde{R}_{jj} \tilde{Z}_j^T + \frac{\tau_0^2 \lambda_n}{k} I_m \right)^{-1} \tilde{Z}_{ja} \right\}_{ii'} \end{aligned}$$

$$\begin{aligned}
 & + \lambda_n^{-1} \sum_{a=1}^q \sum_{h=\bar{h}+1}^{\infty} \mu_{ah} - \lambda_n^{-1} \mathbb{E}_{u_j} \sum_{a=1}^q \sum_{i=1}^m \sum_{i'=1}^m \sum_{h=\bar{h}+1}^{\infty} \mu_{ah}^2 \varphi_{ah}(u_{ji}) \varphi_{ah}(u_{ji'}) \\
 & \times \left\{ \tilde{Z}_{ja}^T \left(\tilde{Z}_j \tilde{R}_{jj} \tilde{Z}_j^T + \frac{\tau_0^2 \lambda_n}{k} I_m \right)^{-1} \tilde{Z}_{ja} \right\} \\
 & \stackrel{(i)}{\leq} \lambda_n^{-1} \mathbb{E}_{u_j} \sum_{a=1}^q \sum_{h=1}^{\bar{h}} \left\{ \mu_{ah} - \mu_{ah}^2 \Phi_{jah}^T \tilde{Z}_{ja}^T \left(\tilde{Z}_j \tilde{R}_{jj} \tilde{Z}_j^T + \frac{\tau_0^2 \lambda_n}{k} I_m \right)^{-1} \tilde{Z}_{ja} \Phi_{jah} \right\} \\
 & + \lambda_n^{-1} \text{Tr}(\rho, \bar{h}) \\
 & = \lambda_n^{-1} \mathbb{E}_{u_j} \text{tr} \left\{ M - M \Phi_j^T \tilde{Z}_j^T \left(\tilde{Z}_j \tilde{R}_{jj} \tilde{Z}_j^T + \frac{\tau_0^2 \lambda_n}{k} I_m \right)^{-1} \tilde{Z}_j \Phi_j M \right\} + \lambda_n^{-1} \text{Tr}(\rho, \bar{h}), \quad (98)
 \end{aligned}$$

where Φ_{jah} denotes the h th column of the matrix Φ_{ja} defined in (70), and (i) follows because we dropped the last negative term to make it larger.

If we let

$$\tilde{M} = M - M \Phi_j^T \tilde{Z}_j^T \left(\tilde{Z}_j \tilde{R}_{jj} \tilde{Z}_j^T + \frac{\tau_0^2 \lambda_n}{k} I_m \right)^{-1} \tilde{Z}_j \Phi_j M,$$

then (98) has shown that

$$\mathbb{E}_{u^*} \mathbb{E}_{y_j, u_j} \left[\text{Var}_{\nu_j | y_j, u_j} \{ \nu_j(u^*) \} \right] \leq \lambda_n^{-1} \mathbb{E}_{u_j} \text{tr}(\tilde{M}) + \lambda_n^{-1} \text{Tr}(\rho, \bar{h}). \quad (99)$$

For $j = 1, \dots, k$, $a = 1, \dots, q$, and $h = 1, 2, \dots$, we define the following matrices

$$\begin{aligned}
 M_a^\dagger &= \text{diag} \left\{ \mu_{a(\bar{h}+1)}, \dots, \mu_{a\infty} \right\}, \quad M^\dagger = \text{diag} \left\{ M_1^\dagger, \dots, M_q^\dagger \right\}, \\
 \Phi_{jah} &= \{ \varphi_{ah}(u_{j1}), \dots, \varphi_{ah}(u_{jm}) \}^T, \\
 \Phi_{ja}^\dagger &= \left\{ \Phi_{ja(\bar{h}+1)}, \dots, \Phi_{ja\infty} \right\}, \quad \Phi_j^\dagger = \text{diag} \left\{ \Phi_{j1}^\dagger, \dots, \Phi_{jq}^\dagger \right\}, \\
 \tilde{R}_{jj}^\dagger &= \Phi_j^\dagger M^\dagger \Phi_j^\dagger, \quad \tilde{R}_{jj} = \Phi_j M \Phi_j + \tilde{R}_{jj}^\dagger,
 \end{aligned}$$

where $\Phi_j \in \mathbb{R}^{qm \times q\bar{h}}$ is defined in (70). Then the Woodbury formula (Harville, 1997) and the definition of Q imply that

$$\begin{aligned}
 \tilde{M} &= \left\{ M^{-1} + \Phi_j^T \tilde{Z}_j^T \left(\tilde{Z}_j \tilde{R}_{jj}^\dagger \tilde{Z}_j^T + \frac{\tau_0^2 \lambda_n}{k} I_m \right)^{-1} \tilde{Z}_j \Phi_j \right\}^{-1} \\
 &= \frac{\tau_0^2 \lambda_n}{km} \left\{ \Omega + \frac{\tau_0^2 \lambda_n}{km} M^{-1} + \frac{1}{m} \Phi_j^T \tilde{Z}_j^T \left(\frac{k}{\tau_0^2 \lambda_n} \tilde{Z}_j \tilde{R}_{jj}^\dagger \tilde{Z}_j^T + I_m \right)^{-1} \tilde{Z}_j \Phi_j - \Omega \right\}^{-1} \\
 &= \frac{\tau_0^2 \lambda_n}{km} Q^{-1} \left[I_{q\bar{h}} + Q^{-1} \Omega^{-1} \left\{ \frac{1}{m} \Phi_j^T \tilde{Z}_j^T \left(\frac{k}{\tau_0^2 \lambda_n} \tilde{Z}_j \tilde{R}_{jj}^\dagger \tilde{Z}_j^T + I_m \right)^{-1} \tilde{Z}_j \Phi_j - \Omega \right\} Q^{-1} \right]^{-1} Q^{-1} \Omega^{-1}. \quad (100)
 \end{aligned}$$

For $j = 1, \dots, k$, define the event $\mathcal{E}_{j2} = \left\{ \frac{k}{\tau_0^2 \lambda_n} \tilde{Z}_j \tilde{R}_{jj}^\dagger \tilde{Z}_j^T \preceq \frac{1}{4} I_m \right\}$. Since the matrix \tilde{R}_{jj}^\dagger is semi-positive definite, we have the relation that

$$\left\{ \text{tr} \left(\frac{k}{\tau_0^2 \lambda_n} \tilde{Z}_j \tilde{R}_{jj}^\dagger \tilde{Z}_j^T \right) \leq \frac{1}{4} \right\} \subseteq \left\{ \text{smax} \left(\frac{k}{\tau_0^2 \lambda_n} \tilde{Z}_j \tilde{R}_{jj}^\dagger \tilde{Z}_j^T \right) \leq \frac{1}{4} \right\} \subseteq \mathcal{E}_{j2},$$

$s_{\max}(A)$ is the maximum eigenvalue of the square matrix A . Therefore, by Markov's inequality and Assumption (A.5), we have that

$$\begin{aligned}
 \mathbb{P}_{u_j}(\mathcal{E}_{j2}^c) &\leq \mathbb{P}_{u_j} \left\{ \text{tr} \left(\frac{k}{\tau_0^2 \lambda_n} \tilde{Z}_j \tilde{R}_{jj}^\dagger \tilde{Z}_j^\top \right) > \frac{1}{4} \right\} \leq 4 \mathbb{E}_{u_j} \text{tr} \left(\frac{k}{\tau_0^2 \lambda_n} \tilde{Z}_j \tilde{R}_{jj}^\dagger \tilde{Z}_j^\top \right) \\
 &= \frac{4k}{\tau_0^2 \lambda_n} \mathbb{E}_{u_j} \text{tr} \left(\tilde{Z}_j \Phi_j^\dagger M^\dagger \Phi_j^\dagger \tilde{Z}_j^\top \right) \\
 &= \frac{4k}{\tau_0^2 \lambda_n} \sum_{a=1}^q \sum_{h=\bar{h}+1}^{\infty} \mu_{ah} \mathbb{E}_{u_j} \left\{ \sum_{i=1}^m Z_a(u_{ji})^2 \varphi_{ah}(u_{ji})^2 \right\} \\
 &\leq \frac{4\tilde{C}_Z^2 k}{\tau_0^2 \lambda_n} \sum_{a=1}^q \sum_{h=\bar{h}+1}^{\infty} \sum_{i=1}^m \mu_{ah} \mathbb{E}_{u_j} \{ \varphi_{ah}(u_{ji})^2 \} \\
 &= \frac{4\tilde{C}_Z^2 km}{\tau_0^2 \lambda_n} \text{Tr}(\rho, \bar{h}). \tag{101}
 \end{aligned}$$

On the event $\mathcal{E}_{j1} \cap \mathcal{E}_{j2}$ (with \mathcal{E}_{j1} defined in (76)), we have that

$$\begin{aligned}
 &I_{q\bar{h}} + Q^{-1} \Omega^{-1} \left\{ \frac{1}{m} \Phi_j^\top \tilde{Z}_j^\top \left(\frac{k}{\tau_0^2 \lambda_n} \tilde{Z}_j \tilde{R}_{jj}^\dagger \tilde{Z}_j^\top + I_m \right)^{-1} \tilde{Z}_j \Phi_j - \Omega \right\} Q^{-1} \\
 &\stackrel{(i)}{\succeq} I_{q\bar{h}} + Q^{-1} \Omega^{-1} \left\{ \frac{1}{m} \Phi_j^\top \tilde{Z}_j^\top \left(\frac{1}{4} I_m + I_m \right)^{-1} \tilde{Z}_j \Phi_j - \Omega \right\} Q^{-1} \\
 &= I_{\bar{h}} - \frac{1}{5} Q^{-2} + \frac{4}{5} Q^{-1} \Omega^{-1} \left\{ \frac{1}{m} \Phi_j^\top \tilde{Z}_j^\top \left(\frac{1}{4} I_m + I_m \right)^{-1} \tilde{Z}_j \Phi_j - \Omega \right\} Q^{-1} \\
 &\stackrel{(ii)}{\succeq} I_{q\bar{h}} - \frac{1}{5} I_{q\bar{h}} - \frac{4}{5} \cdot \frac{1}{2} I_{q\bar{h}} = \frac{2}{5} I_{q\bar{h}}, \tag{102}
 \end{aligned}$$

where (i) follows on the event \mathcal{E}_{j2} , and (ii) holds on the event \mathcal{E}_{j1} and from the fact $Q^{-2} \preceq I_{q\bar{h}}$. Therefore, from (100) and (102), we can obtain that

$$\begin{aligned}
 &\mathbb{E}_{u_j} \left\{ \text{tr}(\tilde{M}) \mathbf{1}(\mathcal{E}_{j1} \cap \mathcal{E}_{j2}) \right\} \leq \mathbb{E}_{u_j} \left\{ \frac{\tau_0^2 \lambda_n}{km} Q^{-1} \cdot \frac{5}{2} I_{q\bar{h}} \cdot Q^{-1} \Omega^{-1} \right\} \\
 &= \frac{5\tau_0^2 \lambda_n}{2km} \text{tr}(Q^{-2} \Omega^{-1}) \leq \frac{5\tau_0^2 \lambda_n}{2km} \text{tr} \left\{ \left(\Omega + \frac{\tau_0^2 \lambda_n}{km} M^{-1} \right)^{-1} \right\} \\
 &\leq \frac{5\tau_0^2 \lambda_n}{2km} \sum_{a=1}^q \sum_{h=1}^{\bar{h}} \frac{1}{\underline{c}_\Omega + \frac{\tau_0^2 \lambda_n}{km \mu_{ah}}} = \frac{5\tau_0^2 \lambda_n}{2\underline{c}_\Omega km} \Lambda(\tau_0^2 \lambda_n / (\underline{c}_\Omega km), \bar{h}). \tag{103}
 \end{aligned}$$

Therefore, by combining (99), (101), (103), and Lemma 5, we obtain that

$$\begin{aligned}
 &\mathbb{E}_{u^*} \mathbb{E}_{y_j, u_j} \text{tr} \left[\text{Var}_{\nu_j | y_j, u_j} \{ \nu_j(u^*) \} \right] \\
 &\leq \lambda_n^{-1} \mathbb{E}_{u_j} \text{tr}(\tilde{M}) + \lambda_n^{-1} \text{Tr}(\rho, \bar{h}) \\
 &\leq \lambda_n^{-1} \mathbb{E}_{u_j} \left\{ \text{tr}(\tilde{M}) \mathbf{1}(\mathcal{E}_{j1} \cap \mathcal{E}_{j2}) \right\} + \lambda_n^{-1} \mathbb{E}_{u_j} \left\{ \text{tr}(\tilde{M}) \mathbf{1}(\mathcal{E}_{j1}^c) \right\}
 \end{aligned}$$

$$\begin{aligned}
 & + \lambda_n^{-1} \mathbb{E}_{u_j} \left\{ \text{tr}(\tilde{M}) 1(\mathcal{E}_{j2}^c) \right\} + \lambda_n^{-1} \text{Tr}(\rho, \bar{h}) \\
 & \leq \frac{5\tau_0^2 \lambda_n}{2\underline{c}_\Omega km} \Lambda(\tau_0^2 \lambda_n / (\underline{c}_\Omega km), \bar{h}) + \lambda_n^{-1} \text{Tr}(\rho) \mathbb{P}_{u_j}(\mathcal{E}_{j1}^c) \\
 & \quad \frac{4\tilde{C}_Z^2 km}{\tau_0^2 \lambda_n^2} \text{Tr}(\rho, \bar{h}) \text{Tr}(\rho) + \lambda_n^{-1} \text{Tr}(\rho, \bar{h}) \\
 & \leq \frac{5\tau_0^2 \lambda_n}{2\underline{c}_\Omega km} \Lambda(\tau_0^2 \lambda_n / (\underline{c}_\Omega km), \bar{h}) + \frac{4\tilde{C}_Z^2 km}{\tau_0^2 \lambda_n^2} \text{Tr}(\rho, \bar{h}) \text{Tr}(\rho) + \lambda_n^{-1} \text{Tr}(\rho, \bar{h}) \\
 & \quad + 2\lambda_n^{-1} q \bar{h} \text{Tr}(\rho) \exp \left\{ -\frac{m}{8(B^2 + B)} \right\},
 \end{aligned}$$

where $B = C_\varphi^2 \tilde{C}_Z^2 \underline{c}_\Omega^{-1} q \Lambda(\tau_0^2 \lambda_n / (\underline{c}_\Omega km), \bar{h}) + 1$. This completes the proof of Lemma 4. \blacksquare

Lemma 5 *For the event \mathcal{E}_{j1} defined in (76), the probability of the event \mathcal{E}_{j1}^c is upper bounded by*

$$\mathbb{P}_{u_j}(\mathcal{E}_{j1}^c) \leq 2q\bar{h} \exp \left\{ -\frac{m}{8(B^2 + B)} \right\}, \quad (104)$$

where $B = C_\varphi^2 \tilde{C}_Z^2 \underline{c}_\Omega^{-1} q \Lambda(\tau_0^2 \lambda_n / (\underline{c}_\Omega km), \bar{h}) + 1$.

Proof [of Lemma 5] For $j = 1, \dots, k$ and $i = 1, \dots, m$, let

$$\begin{aligned}
 W_{ji} = \{ & \tilde{Z}_1(u_{ji})\varphi_{11}(u_{ji}), \dots, \tilde{Z}_1(u_{ji})\varphi_{1\bar{h}}(u_{ji}), \dots, \\
 & \tilde{Z}_q(u_{ji})\varphi_{q1}(u_{ji}), \dots, \tilde{Z}_q(u_{ji})\varphi_{q\bar{h}}(u_{ji}) \}^T \in \mathbb{R}^{q\bar{h}},
 \end{aligned}$$

similar to $W(u)$ in Assumption (A.5). With some linear algebra, the matrix in the definition (76) can be rewritten as

$$\begin{aligned}
 & Q^{-1} \Omega^{-1} \left(\frac{1}{m} \Phi_j^T \tilde{Z}_j^T \tilde{Z}_j \Phi_j - \Omega \right) Q^{-1} = Q^{-1} \Omega^{-1} \left(\frac{1}{m} \sum_{i=1}^m W_{ji} W_{ji}^T - \Omega \right) Q^{-1} \\
 & = \frac{1}{m} \sum_{i=1}^m Q^{-1} \Omega^{-1} (W_{ji} W_{ji}^T - \Omega) Q^{-1}. \quad (105)
 \end{aligned}$$

Using Assumptions (A.3), (A.5), and the fact that $Q \succeq I_{q\bar{h}}$, we can obtain that for every $j = 1, \dots, k$ and $i = 1, \dots, m$,

$$\begin{aligned}
 & \left\| Q^{-1} \Omega^{-1} (W_{ji} W_{ji}^T - \Omega) Q^{-1} \right\| \\
 & \leq \left\| Q^{-1} \Omega^{-1} W_{ji} W_{ji}^T Q^{-1} \right\| + \left\| Q^{-2} \right\| \\
 & \leq W_{ji}^T Q^{-2} \Omega^{-1} W_{ji} + 1 \\
 & \leq W_{ji}^T \left(\Omega + \frac{\tau_0^2 \lambda_n}{km} M^{-1} \right)^{-1} W_{ji} + 1
 \end{aligned}$$

$$\begin{aligned}
 &\leq \underline{c}_\Omega^{-1} \sum_{a=1}^q \sum_{h=1}^{\bar{h}} \frac{\mu_{ah}}{\mu_{ah} + \frac{\tau_0^2 \lambda_n}{\underline{c}_\Omega km}} \tilde{Z}_a(u_{ji})^2 \varphi_{ah}(u_{ji})^2 + 1 \\
 &\leq C_\varphi^2 \tilde{C}_Z^2 \underline{c}_\Omega^{-1} q \Lambda(\tau_0^2 \lambda_n / (\underline{c}_\Omega km), \bar{h}) + 1 \equiv B.
 \end{aligned} \tag{106}$$

Furthermore,

$$\begin{aligned}
 &\left\| \left\{ Q^{-1} \Omega^{-1} (W_{ji} W_{ji}^\top - \Omega) Q^{-1} \right\}^2 \right\| \\
 &\leq \left\| Q^{-1} \Omega^{-1} (W_{ji} W_{ji}^\top - \Omega) Q^{-1} \right\|^2 \leq B^2.
 \end{aligned} \tag{107}$$

Now from (105), (106) and (107), we apply the matrix Bernstein inequality (Theorem 6.1.1 of Tropp 2015) to the sequence of $\left\{ Q^{-1} \Omega^{-1} (W_{ji} W_{ji}^\top - \Omega) Q^{-1} \right\}_{i=1}^m$ to obtain that

$$\begin{aligned}
 \mathbb{P}_{u_j}(\mathcal{E}_{j1}^c) &= \mathbb{P}_{u_j} \left(\left\| Q^{-1} \Omega^{-1} \left(\frac{1}{m} \Phi_j^\top \tilde{Z}_j^\top \tilde{Z}_j \Phi_j - \Omega \right) Q^{-1} \right\| > 1/2 \right) \\
 &\leq \mathbb{P}_{u_j} \left(\left\| \frac{1}{m} \sum_{i=1}^m Q^{-1} \Omega^{-1} (W_{ji} W_{ji}^\top - \Omega) Q^{-1} \right\| > 1/2 \right) \\
 &\leq 2q\bar{h} \exp \left\{ -\frac{(m/2)^2/2}{mB^2 + mB/6} \right\} \\
 &\leq 2q\bar{h} \exp \left\{ -\frac{m}{8(B^2 + B)} \right\},
 \end{aligned}$$

which completes the proof. \blacksquare

Proof [of Theorem 1] We prove part (i) and part (ii), respectively. We simplify the upper bounds in Lemmas 2, 3, and 4 with the choice $\lambda_n = 1$ and $\lambda_n \asymp n^{-(2\nu)/(2\nu+d)}$.

(i) With $\lambda_n = 1$, we first derive a bound for the quantity $\Lambda(\tau_0^2 \lambda_n / (ckm), \bar{h})$ for a generic constant $c > 0$ (which will be replaced by \bar{c}_Ω and \underline{c}_Ω in the upper bounds in Lemmas 2, 3, and 4). Using Assumption (A.3), there exists some constant $c_\mu > 0$ such that $\mu_{i*} \leq c_\mu i^{-2\nu/d}$ for $i = 1, 2, \dots$. Therefore, given the fact that $2\nu > d$, we have that if $\lambda_n = 1$,

$$\begin{aligned}
 \Lambda(\tau_0^2 \lambda_n / (ckm), \bar{h}) &\leq \sum_{h=1}^{\infty} \left(1 + \frac{\tau_0^2}{ckm\mu_{h*}} \right)^{-1} \leq \sum_{h=1}^{\infty} \left(1 + \frac{c_\mu c_2 \tau_0^2}{cn} h^{2\nu/d} \right)^{-1} \\
 &\leq \sum_{h \leq n^{d/(2\nu)}} \left(1 + \frac{c_\mu c_2 \tau_0^2}{cn} h^{2\nu/d} \right)^{-1} + \sum_{h > n^{d/(2\nu)}} \left(1 + \frac{c_\mu c_2 \tau_0^2}{cn} h^{2\nu/d} \right)^{-1} \\
 &\leq n^{d/(2\nu)} + \sum_{h > n^{d/(2\nu)}} \left(\frac{c_\mu c_2 \tau_0^2}{cn} h^{2\nu/d} \right)^{-1} \\
 &\leq n^{d/(2\nu)} + \left(\frac{c_\mu c_2 \tau_0^2}{cn} \right)^{-1} \sum_{h > n^{d/(2\nu)}} \int_h^{h+1} x^{-2\nu/d} dx
 \end{aligned}$$

$$\begin{aligned}
 &\leq n^{d/(2v)} + \left(\frac{c_\mu c_2 T_0^2}{cn} \right)^{-1} n^{-\frac{d}{2v} \cdot \left(\frac{2v}{d} - 1 \right)} \\
 &\asymp n^{d/(2v)} + n^{1 - \frac{2v-d}{2v}} \\
 &\asymp n^{d/(2v)}.
 \end{aligned} \tag{108}$$

Given the condition $m \gtrsim n^{(d/v)+\eta}$, in the exponent of (104), we have

$$\frac{m}{8(B^2 + B)} \asymp \frac{m}{\Lambda(\tau_0^2 \lambda_n / (c_\Omega km), \bar{h})^2} \gtrsim \frac{n^{(d/v)+\eta}}{n^{d/v}} = n^\eta.$$

This implies that for some positive constant $c' > 0$,

$$\exp \left\{ -\frac{m}{8(B^2 + B)} \right\} \leq \exp(-c' n^\eta). \tag{109}$$

With the choice $\bar{h} = \lceil n^{3d/(2v-d)} \rceil$, we have that

$$\begin{aligned}
 \text{Tr}(\rho, \bar{h}) &= \sum_{a=1}^q \sum_{h \geq \bar{h}} \mu_{ah} \leq q \sum_{h \geq \bar{h}} \mu_{h*} \leq q \sum_{h \geq \bar{h}} c_\mu h^{-2v/d} \\
 &\leq qc_\mu \sum_{h \geq \bar{h}} \int_h^{h+1} x^{-2v/d} dx \leq c_\mu \int_{n^{3d/(2v-d)}}^\infty x^{-2v/d} dx \\
 &= qc_\mu n^{-\frac{3d}{2v-d} \cdot \frac{2v-d}{d}} \asymp n^{-3}.
 \end{aligned} \tag{110}$$

Note that Assumption (A.3) (ii) and Assumption (A.4) imply that $\text{Tr}(\rho) = O(1)$ and $\|\nu_0\|_{\mathbb{H}} = O(1)$. Using the orders in (108), (109), and (110), the order of the upper bound in Lemma 2 can be quantified as

$$\begin{aligned}
 &\frac{4}{k} \sum_{j=1}^k \mathbb{E}_{u^*} \mathbb{E}_{u_j} \left\| \mathbb{E}_{\nu_j, y_j | u_j} \{ \nu_j(u^*) \} - \nu_0(u^*) \right\|_2^2 \\
 &\lesssim n^{-1} + n \cdot n^{-3} + n^{-3} + n^{\lceil 3d/(2v-d) \rceil} \cdot \exp(-c' n^\eta) \\
 &\lesssim n^{-1}.
 \end{aligned} \tag{111}$$

Similarly, the order of the upper bound in Lemma 3 can be quantified as

$$\begin{aligned}
 &\frac{4}{k^2} \sum_{j=1}^k \mathbb{E}_{u^*} \mathbb{E}_{u_j} \text{tr} \left(\text{Var}_{y_j | u_j} \left[\mathbb{E}_{\nu_j | y_j, u_j} \{ \nu_j(u^*) \} \right] \right) \\
 &\lesssim \frac{1}{kn} + \frac{n}{k} \cdot n^{-3} \cdot (n+1) \\
 &\quad + \frac{1}{km} \cdot n^{d/(2v)} + \frac{1}{k} \cdot n^{-3} (n+1) \\
 &\quad + \frac{1}{k} \cdot n^{\lceil 3d/(2v-d) \rceil} (n+1) \cdot \exp(-c' n^\eta) \\
 &\lesssim k^{-1} n^{-1} + n^{-(2v-d)/(2v)} + k^{-1} n^{-2} + k^{-1} n^{\lceil 3d/(2v-d) \rceil + 1} \cdot \exp(-c' n^\eta)
 \end{aligned}$$

$$\lesssim n^{-(2\nu-d)/(2\nu)}. \quad (112)$$

The order of the upper bound in Lemma 4 can be quantified as

$$\begin{aligned} & \frac{\bar{c}}{k} \sum_{j=1}^k \mathbb{E}_{u^*} \mathbb{E}_{y_j, u_j} \operatorname{tr} \left(\operatorname{Var}_{\nu_j | y_j, u_j} \{ \nu_j(u^*) \} \right) \\ & \lesssim n^{-1} \cdot n^{d/(2\nu)} + (n+1) \cdot n^{-3} + n^{\lceil 3d/(2\nu-d) \rceil} \cdot \exp(-c'n^\eta) \\ & \lesssim n^{-(2\nu-d)/(2\nu)}. \end{aligned} \quad (113)$$

Finally, we combine (111), (112), (113), (49), and (54) to obtain that

$$\mathbb{L}(\bar{\Pi}) \lesssim n^{-1} + n^{-(2\nu-d)/(2\nu)} + n^{-(2\nu-d)/(2\nu)} + n^{-1} \lesssim n^{-(2\nu-d)/(2\nu)}.$$

The rate for $\bar{w}(\cdot)$ follows trivially from the inequality

$$|\bar{w}(u^*) - w_0(u^*)| = |Z(u^*)\bar{\beta}(u^*) - Z(u^*)\beta_0(u^*)| \leq C_Z \|\bar{\beta}(u^*) - \beta_0(u^*)\|_2.$$

This proves the conclusion of Theorem 1 (i).

(ii) When $\lambda_n \asymp n^{d/(2\nu+d)}$, similar to part (i), we first derive a bound for the quantity $\Lambda_n(\tau_0^2 \lambda_n / (ckm), \bar{h})$ with $c > 0$ being a generic constant. Using Assumption (A.3), there exists some constant $c_\mu > 0$ such that $\mu_j \leq c_\mu j^{-2\nu/d}$. Therefore, we have that

$$\begin{aligned} \Lambda(\tau_0^2 \lambda_n / (km), \bar{h}) & \leq \sum_{h=1}^{\infty} \left(1 + \frac{\tau_0^2 \lambda_n}{ckm \mu_{h^*}} \right)^{-1} \leq \sum_{h=1}^{\infty} \left(1 + \frac{c_\mu c_2 \tau_0^2 \lambda_n h^{2\nu/d}}{cn} \right)^{-1} \\ & \leq \sum_{h \leq n^{d/(2\nu+d)}} \left(1 + \frac{c_\mu c_2 \tau_0^2 \lambda_n h^{2\nu/d}}{n} \right)^{-1} + \sum_{h > n^{d/(2\nu+d)}} \left(1 + \frac{c_\mu c_2 \tau_0^2 \lambda_n h^{2\nu/d}}{cn} \right)^{-1} \\ & \leq n^{d/(2\nu+d)} + \sum_{h > n^{d/(2\nu+d)}} \left(\frac{c_\mu c_2 \tau_0^2 \lambda_n h^{2\nu/d}}{cn} \right)^{-1} \\ & \leq n^{d/(2\nu+d)} + \left(\frac{c_\mu c_2 \tau_0^2 \lambda_n}{cn} \right)^{-1} \sum_{h > n^{d/(2\nu+d)}} \int_h^{h+1} x^{-2\nu/d} dx \\ & \leq n^{d/(2\nu+d)} + \left(\frac{c_\mu c_2 \tau_0^2 \lambda_n}{cn} \right)^{-1} n^{-\frac{d}{2\nu+d} \cdot \left(\frac{2\nu}{d} - 1 \right)} \\ & \asymp n^{d/(2\nu+d)} + \lambda_n^{-1} n^{1 - \frac{2\nu-d}{2\nu+d}} \\ & \asymp n^{d/(2\nu+d)}, \end{aligned} \quad (114)$$

where the last step follows because $\lambda_n \asymp n^{d/(2\nu+d)}$.

Given our condition $m \gtrsim n^{2d/(2\nu+d)+\eta}$, in the exponent of (104), we have

$$\frac{m}{8(B^2 + B)} \asymp \frac{m}{\Lambda(\tau_0^2 \lambda_n / (ckm), \bar{h})^2} \gtrsim \frac{n^{2d/(2\nu+d)+\eta}}{n^{2d/(2\nu+d)}} = n^\eta.$$

This implies that for some positive constant $c' > 0$,

$$\exp\left\{\frac{m}{8(B^2 + B)}\right\} \leq \exp(-c'n^\eta). \quad (115)$$

With the choice $\bar{h} = \lceil n^{3d/(2v-d)} \rceil$, $\text{Tr}(\rho, \bar{h})$ is upper bounded by n^{-3} as in (110).

Using the orders in (114), (115), and (110), the order of the upper bound in Lemma 2 can be quantified as

$$\begin{aligned} & \frac{4}{k} \sum_{j=1}^k \mathbb{E}_{u^*} \mathbb{E}_{u_j} \left\| \mathbb{E}_{\nu_j, y_j | u_j} \{\nu_j(u^*)\} - \nu_0(u^*) \right\|_2^2 \\ & \lesssim \frac{n^{d/(2v+d)}}{n} + \frac{n}{n^{d/(2v+d)}} \cdot n^{-3} + n^{-3} + n^{\lceil 3d/(2v-d) \rceil} \cdot \exp(-c'n^\eta) \\ & \lesssim n^{-2v/(2v+d)}. \end{aligned} \quad (116)$$

Similarly, the order of the upper bound in Lemma 3 can be quantified as

$$\begin{aligned} & \frac{4}{k^2} \sum_{j=1}^k \mathbb{E}_{u^*} \mathbb{E}_{u_j} \text{tr} \left(\text{Var}_{y_j | u_j} \left[\mathbb{E}_{\nu_j | y_j, u_j} \{\nu_j(u^*)\} \right] \right) \\ & \lesssim \frac{n^{d/(2v+d)}}{kn} + \frac{n}{kn^{d/(2v+d)}} \cdot n^{-3} \cdot \left(\frac{n}{n^{d/(2v+d)}} + 1 \right) \\ & \quad + \frac{1}{km} \cdot n^{d/(2v+d)} + \frac{1}{k} \cdot n^{-3} \left(\frac{2n}{n^{d/(2v+d)}} + 1 \right) \\ & \quad + \frac{1}{k} \cdot n^{\lceil 3d/(2v-d) \rceil} \left(\frac{n}{n^{d/(2v+d)}} + 1 \right) \cdot \exp(-c'n^\eta) \\ & \lesssim k^{-1} n^{-2v/(2v+d)} + k^{-1} n^{4v/(2v+d)-3} + n^{-2v/(2v+d)} + k^{-1} n^{2v/(2v+d)-3} \\ & \quad + k^{-1} n^{\lceil 3v/(2v-d) \rceil + 2v/(2v+d)} \cdot \exp(-c'n^\eta) \\ & \lesssim n^{-2v/(2v+d)}. \end{aligned} \quad (117)$$

The order of the upper bound in Lemma 4 can be quantified as

$$\begin{aligned} & \frac{\bar{c}}{k} \sum_{j=1}^k \mathbb{E}_{u^*} \mathbb{E}_{y_j, u_j} \text{tr} \left(\text{Var}_{\nu_j | y_j, u_j} \{\nu_j(u^*)\} \right) \\ & \lesssim \frac{1}{n} \cdot n^{d/(2v+d)} + \left(n \cdot n^{-2d/(2v+d)} + n^{-d/(2v+d)} \right) \cdot n^{-3} \\ & \quad + n^{-d/(2v+d)} \cdot n^{\lceil 3d/(2v-d) \rceil} \cdot \exp(-c'n^\eta) \\ & \lesssim n^{-2v/(2v+d)}. \end{aligned} \quad (118)$$

Finally, we combine (116), (117), and (118), (49), and (54) to obtain that

$$\mathbb{L}(\bar{\Pi}) \lesssim n^{-2v/(2v+d)} + n^{-2v/(2v+d)} + n^{-2v/(2v+d)} + n^{-1} \lesssim n^{-2v/(2v+d)}.$$

The rate for $\bar{w}(\cdot)$ follows similarly. This proves the conclusion of Theorem 1 part (ii). \blacksquare

References

- Mauricio A Álvarez and Neil D Lawrence. Computationally efficient convolved multiple output Gaussian processes. *Journal of Machine Learning Research*, 12(May):1459–1500, 2011.
- Mauricio A Álvarez, Lorenzo Rosasco, and Neil D Lawrence. Kernels for vector-valued functions: A review. *Foundations and Trends® in Machine Learning*, 4(3):195–266, 2012.
- Mauricio A Álvarez, Wil Ward, and Cristian Guarnizo. Non-linear process convolutions for multi-output Gaussian processes. In *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics, PMLR*, volume 89, pages 1969–1977, 2019.
- Pedro C Álvarez-Esteban, E del Barrio, JA Cuesta-Albertos, and C Matrán. A fixed-point approach to barycenters in Wasserstein space. *Journal of Mathematical Analysis and Applications*, 441(2):744–762, 2016.
- Ray Bai, Mary R Boland, and Yong Chen. Fast algorithms and theory for high-dimensional Bayesian varying coefficient models. *arXiv preprint arXiv:1907.06477*, 2019.
- K Shuvo Bakar, Philip Kokic, and Huidong Jin. A spatiodynamic model for assessing frost risk in south-eastern Australia. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 64(5):755–778, 2015.
- Sudipto Banerjee and Gregg A Johnson. Coregionalized single-and multiresolution spatially varying growth curve modeling with application to weed growth. *Biometrics*, 62(3):864–876, 2006.
- Gerd A Becker and Manfred Pauly. Sea surface temperature changes in the north sea and their causes. *ICES Journal of Marine Science*, 53(6):887–898, 1996.
- Moritz Berger, Gerhard Tutz, and Matthias Schmid. Tree-structured modelling of varying coefficients. *Statistics and Computing*, 29(2):217–229, 2019.
- Marc Bourotte, Denis Allard, and Emilio Porcu. A flexible class of non-separable cross-covariance functions for multivariate space–time data. *Spatial Statistics*, 18:125–146, 2016.
- David R Burt, Carl E Rasmussen, and Mark van der Wilk. Rates of convergence for sparse variational Gaussian process regression. In *Proceedings of the 36th International Conference on Machine Learning, PMLR*, volume 97, pages 862–871, 2019.
- Zongwu Cai, Jianqing Fan, and Qiwei Yao. Functional-coefficient regression models for nonlinear time series. *Journal of the American Statistical Association*, 95(451):941–956, 2000.
- Rong Chen and Ruey S Tsay. Functional-coefficient autoregressive models. *Journal of the American Statistical Association*, 88(421):298–308, 1993.

- Edmond Chow and Yousef Saad. Preconditioned Krylov subspace methods for sampling multivariate Gaussian distributions. *SIAM Journal on Scientific Computing*, 36(2):A588–A608, 2014.
- Andreas Damianou and Neil Lawrence. Deep Gaussian processes. In *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics, PMLR*, volume 31, pages 207–215, 2013.
- Abhirup Datta, Sudipto Banerjee, Andrew O Finley, Nicholas AS Hamm, Martijn Schaap, et al. Nonseparable dynamic nearest neighbor Gaussian process models for large spatio-temporal data with an application to particulate matter analysis. *The Annals of Applied Statistics*, 10(3):1286–1316, 2016.
- David Duvenaud, Oren Rippel, Ryan Adams, and Zoubin Ghahramani. Avoiding pathologies in very deep networks. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics, PMLR*, volume 33, pages 202–210, 2014.
- Reihaneh Entezari, Radu V. Craiu, and Jeffrey S. Rosenthal. Likelihood inflating sampling algorithm. *The Canadian Journal of Statistics*, 46(1):147–175, 2017.
- Jianqing Fan and Wenyang Zhang. Statistical estimation in varying coefficient models. *The Annals of Statistics*, 27(5):1491–1518, 1999.
- Andrew O Finley and Sudipto Banerjee. Bayesian spatially varying coefficient models in the spBayes R package. *Environmental Modelling & Software*, page 104608, 2020.
- Andrew O Finley, Sudipto Banerjee, Aaron R Weiskittel, Chad Babcock, and Bruce D Cook. Dynamic spatial regression models for space-varying forest stand tables. *Environmetrics*, 25(8):596–609, 2014.
- Gregory Gaspari and Stephen E Cohn. Construction of correlation functions in two and three dimensions. *Quarterly Journal of the Royal Meteorological Society*, 125(554):723–757, 1999.
- Alan E Gelfand and Sudipto Banerjee. Multivariate spatial process models. *Handbook of Spatial Statistics*, pages 495–515, 2010.
- Alan E Gelfand, Hyon-Jung Kim, CF Sirmans, and Sudipto Banerjee. Spatial modeling with spatially varying coefficient processes. *Journal of the American Statistical Association*, 98(462):387–396, 2003.
- Marc G Genton and William Kleiber. Cross-covariance functions for multivariate geostatistics. *Statistical Science*, 30(2):147–163, 2015.
- Tilmann Gneiting. Nonseparable, stationary covariance functions for space–time data. *Journal of the American Statistical Association*, 97(458):590–600, 2002.
- Robert B Gramacy and Daniel W Apley. Local Gaussian process approximation for large computer experiments. *Journal of Computational and Graphical Statistics*, 24(2):561–578, 2015.

- Mengyang Gu and Weining Shen. Generalized probabilistic principal component analysis of correlated data. *Journal of Machine Learning Research*, 21(13):1–41, 2020.
- Rajarshi Guhaniyogi, Cheng Li, Terrance D Savitsky, and Sanvesh Srivastava. Distributed Bayesian Kriging. *arXiv preprint arXiv:1712.09767*, 2017.
- NAS Hamm, AO Finley, M Schaap, and A Stein. A spatially varying coefficient model for mapping PM10 air quality at the European scale. *Atmospheric Environment*, 102:393–405, 2015.
- David A Harville. *Matrix Algebra from a Statistician’s Perspective*, volume 1. Springer, 1997.
- Trevor Hastie and Robert Tibshirani. Varying-coefficient models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 55(4):757–779, 1993.
- Jianhua Z Huang, Colin O Wu, and Lan Zhou. Varying-coefficient models and basis function approximations for the analysis of repeated measurements. *Biometrika*, 89(1):111–128, 2002.
- James E Johndrow, Aaron Smith, Natesh Pillai, and David B Dunson. MCMC for imbalanced categorical data. *Journal of the American Statistical Association*, 114(527):1394–1403, 2019.
- Michael I Jordan, Jason D Lee, and Yun Yang. Communication-efficient distributed statistical inference. *Journal of the American Statistical Association*, 114(526):668–681, 2019.
- JJ Kennedy, NA Rayner, RO Smith, DE Parker, and M Saunby. Reassessing biases and other uncertainties in sea surface temperature observations measured in situ since 1850: 2. Biases and homogenization. *Journal of Geophysical Research: Atmospheres*, 116(D14), 2011.
- Robert M Key, Alex Kozyr, Chris L Sabine, Kitack Lee, Rik Wanninkhof, John L Bullister, Richard A Feely, Frank J Millero, Calvin Mordy, and T-H Peng. A global ocean carbon climatology: Results from Global Data Analysis Project (GLODAP). *Global biogeochemical cycles*, 18(4), 2004.
- Kitack Lee, Lan T Tong, Frank J Millero, Christopher L Sabine, Andrew G Dickson, Catherine Goyet, Geun-Ha Park, Rik Wanninkhof, Richard A Feely, and Robert M Key. Global relationships of total alkalinity with salinity and temperature in surface waters of the world’s oceans. *Geophysical research letters*, 33(19), 2006.
- Cheng Li, Sanvesh Srivastava, and David B Dunson. Simple, scalable and accurate posterior interval estimation. *Biometrika*, 104(3):665–680, 2017.
- Qi Li and Jeffrey Scott Racine. *Nonparametric Econometrics: Theory and Practice*. Princeton University Press, 2007.
- Zudi Lu, Dag Johan Steinskog, Dag Tjøstheim, and Qiwei Yao. Adaptively varying-coefficient spatiotemporal models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(4):859–880, 2009.

- Anandamayee Majumdar and Alan E Gelfand. Multivariate spatial modeling for geostatistical data using convolved covariance functions. *Mathematical Geology*, 39(2):225–245, 2007.
- Brian D Marx. P-spline varying coefficient models for complex data. In *Statistical Modelling and Regression Structures*, pages 19–43. Springer, 2010.
- Frank J Millero, Kitack Lee, and Mary Roche. Distribution of alkalinity in the surface waters of the major oceans. *Marine Chemistry*, 60(1-2):111–130, 1998.
- Stanislav Minsker, Sanvesh Srivastava, Lizhen Lin, and David Dunson. Scalable and robust Bayesian inference via the median posterior. In *Proceedings of the 31st International Conference on Machine Learning, PMLR*, volume 32, pages 1656–1664, 2014.
- Stanislav Minsker, Sanvesh Srivastava, Lizhen Lin, and David B Dunson. Robust and scalable Bayes via a median of subset posterior measures. *Journal of Machine Learning Research*, 17:1–40, 2017.
- Robert Nishihara, Iain Murray, and Ryan P Adams. Parallel MCMC with generalized elliptical slice sampling. *Journal of Machine Learning Research*, 15(1):2087–2112, 2014.
- Byeong U Park, Enno Mammen, Young K Lee, and Eun Ryung Lee. Varying coefficient regression models: A review and new developments. *International Statistical Review*, 83(1):36–64, 2015.
- M. Plummer. JAGS: A program for analysis of Bayesian Graphical models using Gibbs sampling. In *Proceedings of the 3rd International Workshop on Distributed Statistical Computing, March*, pages 20–22. Citeseer, 2003.
- Joaquin Quiñonero-Candela and Carl Edward Rasmussen. A unifying view of sparse approximate Gaussian process regression. *Journal of Machine Learning Research*, 6(Dec):1939–1959, 2005.
- Qian Ren and Sudipto Banerjee. Hierarchical factor models for large spatially misaligned data: A low-rank predictive process approach. *Biometrics*, 69(1):19–30, 2013.
- Klaus Ritter, Grzegorz W. Wasilkowski, and Henryk Woźniakowski. Multivariate integration and approximation of random fields satisfying Sacks-Ylvisaker conditions. *The Annals of Applied Probability*, 5:518–540, 1995.
- David Ruppert, Matt P Wand, and Raymond J Carroll. *Semiparametric regression*. Number 12. Cambridge University Press, 2003.
- Gabriele Santin and Robert Schaback. Approximation of eigenfunctions in kernel-based spaces. *Advances in Computational Mathematics*, 42(4):973–993, 2016.
- Robert Schaback and Hodger Wendland. Approximation by positive definite kernels. *Advanced Problems in Constructive Approximation*, pages 203–222, 2002.

- Steven L. Scott, Alexander W. Blocker, Fernando V. Bonassi, Hugh A. Chipman, Edward I. George, and Robert E. McCulloch. Bayes and big data: the consensus Monte Carlo algorithm. *International Journal of Management Science and Engineering Management*, 11(2):78–88, 2016.
- Zuofeng Shang and Guang Cheng. Computational limits of a distributed algorithm for smoothing spline. *Journal of Machine Learning Research*, 16(108):1–37, 2017.
- Zuofeng Shang, Botao Hao, and Guang Cheng. Nonparametric Bayesian aggregation for massive data. *Journal of Machine Learning Research*, 20(140):1–81, 2019.
- Sanvesh Srivastava, Volkan Cevher, Quoc Dinh, and David Dunson. WASP: Scalable Bayes via barycenters of subset posteriors. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics, PMLR*, volume 38, pages 912–920, 2015.
- Sanvesh Srivastava, Cheng Li, and David B. Dunson. Scalable Bayes via barycenter in Wasserstein space. *Journal of Machine Learning Research*, 19(8):312–346, 2018.
- Botond Szabó and Harry van Zanten. An asymptotic analysis of distributed nonparametric methods. *Journal of Machine Learning Research*, 20(87):1–30, 2019.
- Botond Szabó and Harry van Zanten. Adaptive distributed methods under communication constraints. *The Annals of Statistics*, 48(4):2347–2380, 2020.
- Joel A Tropp. An introduction to matrix concentration inequalities. *Foundations and Trends in Machine Learning*, 9(1–2):1–230, 2015.
- Aad van der Vaart and Harry van Zanten. Rates of contraction of posterior distributions based on Gaussian process priors. *The Annals of Statistics*, 36(3):1435–1463, 2008.
- Aad van der Vaart and Harry van Zanten. Information rates of nonparametric Gaussian process methods. *Journal of Machine Learning Research*, 12(Jun):2095–2119, 2011.
- Harry L. Van Trees. *Detection, Estimation, and Modulation Theory*. John Wiley & Sons, 2001.
- Hans Wackernagel. *Multivariate Geostatistics: An Introduction with Applications*. Springer Science & Business Media, 2006.
- Fengrong Wei, Jian Huang, and Hongzhe Li. Variable selection and estimation in high-dimensional varying-coefficient models. *Statistica Sinica*, 21:1515–1540, 2011.
- S Weldeab, RR Schneider, and M Kölling. Deglacial sea surface temperature and salinity increase in the western tropical atlantic in synchrony with high latitude climate instabilities. *Earth and Planetary Science Letters*, 241(3–4):699–706, 2006.
- David C Wheeler and Catherine A Calder. An assessment of coefficient accuracy in linear regression models with spatially varying coefficients. *Journal of Geographical Systems*, 9(2):145–166, 2007.

- Simon N Wood. *Generalized Additive Models: An Introduction with R*. Chapman and Hall/CRC, 2017.
- Colin O Wu, Chin-Tsang Chiang, and Donald R Hoover. Asymptotic confidence regions for kernel smoothing of a varying-coefficient model with longitudinal data. *Journal of the American Statistical Association*, 93(444):1388–1402, 1998.
- Yu Jiu Xiong, Zhi He Chen, and Sheng Lin Tan. Relationship between salinity and sea surface temperature in Pearl River Estuary, China. In *2013 IEEE International Geoscience and Remote Sensing Symposium-IGARSS*, pages 1575–1578. IEEE, 2013.
- Yixiang Xu and Sanvesh Srivastava. Multivariate distributed Bayesian inference for linear mixed-effects models. *Journal of the Computational and Graphical Statistics*, 30(3):594–611, 2021.
- Jingnan Xue and Faming Liang. Double-parallel Monte Carlo for Bayesian analysis of big data. *Statistics and Computing*, 29(1):23–32, 2019.
- Fariba Yousefi, Michael T Smith, and Mauricio Álvarez. Multi-task learning for aggregated data using Gaussian processes. In *Advances in Neural Information Processing Systems*, pages 15050–15060, 2019.
- Hao Zhang. Inconsistent estimation and asymptotically equal interpolations in model-based geostatistics. *Journal of the American Statistical Association*, 99(465):250–261, 2004.
- Hao Zhang. Maximum-likelihood estimation for multivariate spatial linear coregionalization models. *Environmetrics*, 18(2):125–139, 2007.
- Michael Minyi Zhang and Sinead A Williamson. Embarrassingly parallel inference for Gaussian processes. *Journal of Machine Learning Research*, 20(169):1–26, 2019.
- Yuchen Zhang, John C Duchi, and Martin J Wainwright. Divide and conquer kernel ridge regression: A distributed algorithm with minimax optimal rates. *Journal of Machine Learning Research*, 16(102):3299–3340, 2015.
- Yichen Zhou and Giles Hooker. Tree boosted varying coefficient models. *arXiv preprint arXiv:1904.01058*, 2019.