BAYESIAN METHODS FOR CUMULATIVE PROBABILITY MODELS AND POPULATION PHARMACOKINETIC MODELS

By

Nathan Thomas James

Dissertation

Submitted to the Faculty of the

Graduate School of Vanderbilt University

in partial fulfillment of the requirements

for the degree of

DOCTOR OF PHILOSOPHY

in

Biostatistics

May 31, 2022

Nashville, Tennessee

Approved:

Bryan Shepherd, Ph.D.

Leena Choi, Ph.D.

Matthew Shotwell, Ph.D.

Sara Van Driest, M.D., Ph.D.

For my parents, Thomas & Sue and Marie and my son, Eben

ACKNOWLEDGMENTS

I would like to thank my committee members, Matt Shotwell, Sara Van Driest, and especially my committee chair, Bryan Shepherd, and advisor, Leena Choi, for their time and insightful feedback. I would also like to thank Jeffrey Blume, Frank Harrell, Robert Greevy and Prince Kannenkeril for their advice and mentorship; Cole Beck, Elizabeth McNeer, Michael Williams, and Hannah Weeks for their work on the 'EHRtoPKPD' system used to process much of data used in this dissertation; members of my PhD cohort Sarah Lotspeich, Elizabeth Sigworth, and Valerie Welty; and the students, faculty, and staff in the Vanderbilt Department of Biostatistics.

Additional thanks go to John Koethe for providing the biomarker data used in chapter 2; Yuqi Tian and Chun Li who provided helpful comments and review of early versions of the work in that chapter; all the co-authors of chapter 4 which first appeared as an original article in the *British Journal of Clinical Pharmacology*; the Vanderbilt Clinical Laboratory staff for their assistance in obtaining the remnant specimens; Ahmed Elboraie and Rachel Tyndale for assistance with the *CYP2A6* genetic risk score; and the patients and families who participated in these studies.

Portions of this dissertation were supported by funding from the United States National Institutes of Health (R01AI093234, R01GM124109, R01HD084461, P30AI110527, K23100700, K23AT002508, P30AI54999, and UL1TR000445) and by an appointment to the Research Participation Program at the Office of Biostatistics, Center for Drug Evaluation and Research, U.S. Food and Drug Administration, administered by the Oak Ridge Institute for Science and Education through an interagency agreement between the U.S. Department of Energy and FDA.

Finally, I'd like to thank my family for their support. My greatest thanks goes to my wife, Amy, whose encouragement, patience, and love were a constant. I'm so grateful that you're in my life. I couldn't have done it without you!

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

Introduction

Bayesian methods offer many advantages in clinical trials, healthcare evaluation, and drug development, but their usage still lags that of classical 'frequentist' methods in these areas. While some of this lag may be related to disagreements concerning the foundations of inference and other philosophical issues, there is also a more practical component: despite advances in recent decades, much less methodological and computational work has been done to show how Bayesian methods can be applied for real world analyses. The goal of this dissertation is to demonstrate the use of Bayesian methodology in two application areas, semi-parametric modeling and population pharmacokinetic (PK) analysis.

We begin by reviewing Bayesian inference and estimation and introducing the two application areas. Chapter 2 details an extension to the classical cumulative probability model (CPM) to perform Bayesian semi-parametric regression modeling. The aim is to show how CPMs can be reparameterized to handle a large number of ordinal categories and to characterize the performance of these models with continuous or mixed continuous and discrete outcome data. Chapter 3 is focused on the application of variational inference (VI), specifically automatic differentiation variational inference (ADVI), to population PK models. While Markov chain Monte Carlo (MCMC) remains the gold standard for Bayesian estimation, it can suffer from slow convergence and require long computation time to obtain results. This creates a bottleneck during iterative model exploration and development. Therefore, we seek an approximate estimation method that balances accuracy and speed. ADVI offers one potential solution. Chapter 4 describes a frequentist population PK analysis of the sedative dexmedetomidine in a pediatric cohort using real-world data collected from electronic health records (EHRs) and remnant specimens. This analysis confirmed the importance of weight and age for dexmedetomidine PK but did not find evidence for hypothesized pharmacogenetic effects due to limitations of the observational EHR data. To address these limitations, Chapter 5 applies the ADVI methodology from Chapter 3 along with MCMC estimation to a Bayesian reanalysis. The goal is to stabilize estimation by incorporating external evidence from previous studies with the larger observational dataset. We conclude by summarizing the work and discussing future research directions.

1.1 Bayesian Inference

1.1.1 Background

Bayesian inference has a wide range of definitions. To give just a few recent examples, McElreath (2016) states that "Bayesian inference is no more than counting the numbers of ways things can happen, according to our assumptions" [1] and Kruschke's (2015) definition involves two main ideas, the assignment of possibilities to parameter values in descriptive mathematical models and reallocation of credibility across the possibilities [2]. Lunn et al. (2013) also state that it involves two key ingredients: using probability to express uncertainty about an unknown quantity of interest and a focus on learning about the unknown quantity given observed data [3]. Gelman et al. (2014) define it as "the process of fitting a probability model to a set of data and summarizing the result by a probability distribution on the parameters of the model and on unobserved quantities such as predictions for new observations" which involves setting up a full probability model (a joint probability distribution for all observable and unobservable quantities), conditioning on observed data to calculate the posterior distribution, and evaluating the fit of the model and the implications of the resulting posterior [4]. Wasserman (2004) enumerates three postulates of the Bayesian approach versus the 'classical' or 'frequentist' paradigm: (1) probability describes degree of belief, not limiting frequency; (2) we can make probability statements about parameters, even though they are fixed constants; and (3) we make inferences about a parameter θ by producing a probability distribution for θ, from which point and interval estimates are extracted [5]. While all these definitions differ slightly, they are rooted in the unifying idea of quantifying prior information about unknown parameters, θ, using probability, $p(\theta)$, and then combining the prior information with observed data, z, based on Bayes' Theorem:

$$p(\theta|z) = \frac{p(z|\theta)p(\theta)}{\int p(z|\theta)p(\theta)\,d\theta} \tag{1.1}$$

The term $p(z|\theta)$ in Equation 1.1 is the assumed model for the data (also called the likelihood when viewed as a function of θ) and $p(\theta)$ is a probability distribution which represents prior knowledge (or ignorance) about the parameters. The combination of these two components yields the posterior $p(\theta|z)$ which represents an updated probability distribution for the parameters conditional on the observed data. All Bayesian inference is based on the posterior distribution of the parameters. Moreover, we can make predictions about unknown (but potentially observable) data, $\tilde{z}$. Assuming past and future observations are conditionally independent:

$$p(\tilde{z}|z) = \int p(\tilde{z}|\theta)p(\theta|z)\,d\theta \tag{1.2}$$

where the first term in the integral on the right side of the equation is the assumed model for the data (likelihood) and the second term is the posterior distribution.

Performing inference in a Bayesian framework has several advantages. One strength is the ability to directly answer inferential questions using posterior probabilities. A Bayesian posterior interval can be interpreted intuitively as the probability that the parameter is within the bounds of the interval (an interpretation that is often incorrectly used for frequentist confidence intervals). Another benefit of Bayesian inference is the ability to combine external evidence, in the form of the prior, with evidence from the observed data likelihood. By developing appropriate priors Bayesian models can include information from different sources such as previous studies, pre-clinical models, elicited clinical expertise, and observational data. While this aspect of Bayesian inference may be seen as injecting a subjective element into an objective inferential procedure an alternate view is that the choice of prior is an assumption that should be justified in the same way as other model assumptions. Bayesian models also provide a natural framework for handling missing data (since missing values are treated in the same way as other unobserved quantities), hierarchical modeling (since correlation between units can be modeled using shared parameters or hyperparameters), and prediction (as demonstrated in Equation 1.2) and do not rely on asymptotic approximations.

The main challenge in implementing Bayesian methods is calculation of the posterior since the integral in Equation 1.1 is high-dimensional and analytically intractable for many problems of interest. However, many methods have been developed to provide numerical solutions. In this dissertation, we will utilize two such methods, MCMC and VI.

1.1.2 Markov Chain Monte Carlo Estimation

MCMC is a simulation-based approach used to sequentially draw samples from the posterior distribution. The goal is to construct a Markov chain, a stochastic process with the property that future states depend only on the current state, whose stationary distribution is equal to the target posterior distribution, $p(\theta|z)$. The chain is constructed using iterative Monte Carlo sampling from a proposal distribution. Defined rules for accepting or rejecting each proposed sample ensure that the Markov chain converges to the correct distribution[1].

For this dissertation, we use a type of MCMC called Hamiltonian Monte Carlo (HMC). HMC augments the posterior distribution, $p(\theta|z)$, with an independent auxiliary variable, $p(\xi)$. The joint distribution, $p(\theta,\xi|z) = p(\xi)p(\theta|z)$, can then be used to simulate dynamics in a 'physical' system in which the parameters, θ, represent 'position' and the auxiliary variables, ξ, represent 'momentum.'

[1] Some types of MCMC such as Gibbs sampling and slice sampling can be viewed as special cases where every draw from the proposal distribution is accepted.

In brief, the steps of the HMC algorithm are:

1. Update ξ with a random draw from its prior distribution (e.g., $\xi \sim N(0, M)$ where M is a prespecified 'mass' matrix)

2. Update the joint distribution (θ, ξ) with L 'leapfrog steps' as follows:

   a. Use the gradient of the log-posterior to make a half-step update of ξ

   $$\xi \leftarrow \xi + \frac{1}{2}\varepsilon \frac{dp(\theta|z)}{d\theta}$$

   b. Use the momentum vector, ξ, to update the position vector, θ

   $$\theta \leftarrow \theta + \varepsilon M^{-1}\xi$$

   c. Use the gradient of the log-posterior again to make another half-step update of ξ

   $$\xi \leftarrow \xi + \frac{1}{2}\varepsilon \frac{dp(\theta|z)}{d\theta}$$

   The leapfrog steps are a discrete approximation of a continuous Hamiltonian dynamic system and ε is a scaling factor for these steps.

3. Letting $(\theta^{t-1}, \xi^{t-1})$ be the values of the joint distribution before step 2 and (θ^*, ξ^*) be the values at the end of step 2 compute the acceptance ratio:

   $$r = \frac{p(\theta^*|z)p(\xi^*)}{p(\theta^{t-1}|z)p(\xi^{t-1})}$$

4. Update the position vector:

   $$\theta^t = \begin{cases} \theta^* & \text{with probability } \min(r, 1) \\ \theta^{t-1} & \text{otherwise.} \end{cases}$$

In practice, HMC converges to the target distribution in fewer iterations than some other types of MCMC such as the Metropolis-Hastings algorithm, especially for high-dimensional distributions. Once the chain has converged, all subsequent draws from the chain can be considered samples from the posterior. Finally, the posterior samples are used to produce summary statistics such as means, quantiles, intervals, etc. for the parameters or to make predictions. Additional details and references for MCMC and HMC can be found in the

texts by Brooks (2011) [6], Gelman et al. (2014) [4], and Turkman et al. (2019) [7].

1.1.3 Variational Inference

MCMC methods use samples from a stationary Markov chain to numerically approximate the exact posterior distribution $p(\theta|z)$. In contrast, the goal of VI is to find the joint distribution, $q(\theta;\phi)$, in a specified distributional family indexed by variational parameters, $\phi \in \Phi$, which has the smallest Kullback-Leibler (KL) divergence to the posterior distribution. KL divergence is a non-symmetric measure of the degree of difference between a probability distribution, $f(\theta)$, and a reference distribution, $g(\theta)$, $KL(f(\theta)||g(\theta)) = \int f(\theta)\log\frac{f(\theta)}{g(\theta)}\,d\theta = E_f\left[\log\left(\frac{f(\theta)}{g(\theta)}\right)\right]$. The best variational approximation, $q(\theta;\phi^*)$, is found by solving the optimization problem, $\phi^* = \underset{\phi\in\Phi}{\arg\min}\{KL(q(\theta;\phi)||p(\theta|z))\}$, which involves the unknown posterior distribution $p(\theta|z)$. Using KL as the divergence measure allows the optimization to be solved even though the posterior is unknown. Decomposing the KL divergence to be minimized we have:

$$\begin{aligned}
KL(q(\theta;\phi)||p(\theta|z)) &= E_q\left[\log\frac{q(\theta;\phi)}{p(\theta|z)}\right] = E_q[\log q(\theta;\phi) - \log p(\theta|z)] \\
&= E_q\left[\log q(\theta;\phi) - \log\left(\frac{p(\theta)p(z|\theta)}{p(z)}\right)\right] \\
&= E_q\left[\log\frac{q(\theta;\phi)}{p(\theta)}\right] - E_q[\log p(z|\theta)] + E_q[\log p(z)] \\
&= KL(q(\theta;\phi)||p(\theta)) - E_q[\log p(z|\theta)] + \log p(z).
\end{aligned}$$

The final equality shows the dependence of the optimization on the marginal likelihood (or evidence), $\log p(z)$. However, because all expectations are with respect to $q(\theta;\phi)$ (which $\log p(z)$ does not depend on), finding ϕ^* is equivalent to finding $\phi^* = \underset{\phi\in\Phi}{\arg\max}\{E_q[\log p(z|\theta)] - KL(q(\theta;\phi)||p(\theta))\}$. Further, since the KL divergence ≥ 0 for all distributions, we have:

$$\begin{gathered}
KL(q(\theta;\phi)||p(\theta|z)) \geq 0 \\
KL(q(\theta;\phi)||p(\theta)) - E_q[\log p(z|\theta)] + \log p(z) \geq 0 \\
\log p(z) \geq \underbrace{E_q[\log p(z|\theta)] - KL(q(\theta;\phi)||p(\theta))}_{ELBO}.
\end{gathered}$$

Because of this relationship, the term that is maximized to find ϕ^* is also called the evidence lower bound (ELBO).

The optimization problem can be solved using several approaches. One common method uses a 'mean-field' approximation which assumes mutual independence between each of the j marginal distributions in

the specified variational distributional family, $q(\theta;\phi) = \Pi_j q_j(\theta_j;\phi_j)$. Then, a coordinate ascent algorithm iteratively optimizes each marginal factor, $q_j(\theta_j;\phi_j)$, using complete conditional densities given all the other parameters, θ_{-j}, and the observed data. Early work in VI focused on deriving closed form solutions for the optimal updates in special cases such as conjugate and exponential family models [8–12]. However, deriving and coding these updates can be time-consuming in itself, requiring alternative methods for models outside of the special cases [13] and undercutting the aim of quickly iterating over multiple models without constraints on model and prior parameterization.

As a solution, Kucukelbir et al. (2017) [14] developed a VI method called automatic differentiation variational inference. The ADVI methodology combines advances in probabilistic programming and automatic differentiation to make VI more accessible – by avoiding derivation of complete conditionals for each new model – and more general – since it requires only a differentiable probability model (compared to more stringent conjugacy or exponential family requirements). Briefly, all unknown parameters are first transformed to a common, unconstrained real coordinate space. This transformation allows the use of a Gaussian variational family on the transformed space for all problems. Next a stochastic gradient ascent algorithm is developed to optimize the variational objective function in the transformed space. A second transformation facilitates calculation of expectations using Monte Carlo integration while gradients are calculated using automatic differentiation. Further details and references for VI can be found in the review article by Blei et al. (2017) [15]. We explain the ADVI algorithm in more depth in Chapter 3.

1.2 Semi-parametric Bayesian Models

In Chapter 2, we explore semi-parametric Bayesian models using the framework of cumulative probability models. Cumulative probability models for ordinal outcomes – traditionally denoted cumulative link models [16] – have been discussed extensively in the literature using both frequentist and Bayesian implementations. Since these models are characterized by adding probabilities, not link functions, we prefer the nomenclature cumulative probability model (CPM). Under the frequentist paradigm, Walker and Duncan (1967) [17] and McCullagh (1980) [18] described these models as an extension of dichotomous outcome regression models such as logistic and probit regression. A Bayesian CPM for ordinal regression was explored by Albert and Chib (1993; 1997) [19,20] and Johnson and Albert (1999) [21]. Additional Bayesian CPM extensions including partial proportional odds [22], mixture link models [23], location-scale ordinal regression, and multivariate ordinal outcomes are described by Congdon (2005) [24]. In all these settings, the number of ordered outcome categories is implicitly assumed to be much smaller than the sample size. However, continuous data where each distinct value is its own category are also ordinal and can therefore be fit using CPMs.

In the continuous outcome setting, Liu et al. (2017) [25] demonstrate the equivalence between CPMs and semiparametric linear transformation models of the form:

$$Y = H(\beta^T X + \varepsilon) \quad \text{with} \quad \varepsilon \sim F_\varepsilon \tag{1.3}$$

where $H(\cdot)$ is an increasing function, β is a vector of regression coefficients, X is a vector of covariates, and ε are errors distributed according to known F_ε. Harrell (2015) [26], Liu et al. (2017) [25], and Tian et al. (2019) [27] describe non-parametric maximum likelihood estimation [28] for the unspecified transformation $H(\cdot)$ and β parameters. The models have several favorable characteristics including invariance to monotonic outcome transformations for the regression coefficient estimates and the ability to handle mixed continuous and discrete outcomes such as those that arise from a lower or upper limit of detection. In addition CPMs directly model the full conditional cumulative distribution function (CDF); this allows estimates of conditional means, quantiles, and other statistics to be calculated from a single model fit. Further, because only the $H(\cdot)$ part of the model is nonparametric, CPMs are semiparametric regression models which balance the robustness of fully nonparametric models and the efficiency of fully parametric models.

There is an extensive literature on Bayesian semiparametric regression models. The aim stated by Gelfand (1999) [29] in his discussion of general approaches for these models is, "to enrich the class of standard parametric hierarchical models by wandering nonparametrically near (in some sense) the standard class but retaining the linear structure." For example, Brunner (1995) [30] describes Bayesian linear regression models with symmetric unimodal error densities and Kottas and Gelfand (2001) [31] describe Bayesian semiparametric median regression. DeYoreo and Kottas (2020) [32] explore Bayesian nonparametric density regression for ordinal responses by modeling the joint density between the outcome and covariates using latent continuous random variables. Song and Lu (2012) [33] develop a semiparametric transformation nonlinear mixed model which estimates the transformation, $H(\cdot)$, and incorporates possible nonlinear relationships between X and β and random effects using Bayesian P-splines. Tang et al. (2018) [34] describe semiparametric Bayesian analysis for transformation linear mixed models using a similar Bayesian P-spline approach to estimate the transformation with a focus on nonparametric estimation of random effects. For survival outcomes, Mallick and Walker (2003) [35] describe a linear transformation model for mean survival time where the transformation, $H(\cdot)$, and the error distribution, F_ε, are estimated nonparametrically using mixtures of incomplete beta functions and a Pólya tree distribution, respectively. Lin et al. (2012) [36] detail a semiparametric Bayesian transformation model for median survival. Hanson and colleagues [37,38] and Ibrahim et al. (2010) [39] describe other Bayesian nonparametric survival models. Additional details on general Bayesian nonparametric models can be found in the texts by Müller et al. (2015) [40] and Hjort et al. (2010) [41].

1.3 Pharmacokinetics

Pharmacokinetic (PK) models are often used to characterize the absorption, distribution, metabolism, and elimination (ADME) of a drug in the body over time. They are simply characterized as "how the body acts on the drug." There are several approaches to PK modeling, including and non-compartmental analyses [42] and physiologically based PK modeling. For this work we focus on population PK analysis with compartmental models.

1.3.1 Compartmental Models

Compartmental PK models approximate the complex ADME process by grouping body systems into a set of compartments. Each hypothetical compartment is comprised of tissues, organs, and fluids assumed to have similar kinetic behavior [43]. The movement of a drug through the compartments is modeled by a system of differential equations whose solutions provide the amount of drug in each compartment over time. For all but the simplest models, the solutions are nonlinear in the parameters. For example, assuming a one-compartment model with linear elimination, the concentration at time t for a single intravenous (IV) infusion dose, d, administered at time t_d for duration t_{inf} can be defined as:

$$f_1(D=\{d,t_d,t_{inf}\},\psi=\{Cl,V\},t) = \begin{cases} \frac{d}{t_{inf}}\frac{1}{Cl}\left(1-e^{-\frac{Cl}{V}(t-t_d)}\right) & \text{if } t-t_d \leq t_{inf} \\ \frac{d}{t_{inf}}\frac{1}{Cl}\left(1-e^{-\frac{Cl}{V}t_{inf}}\right)e^{-\frac{Cl}{V}(t-t_d-t_{inf})} & \text{otherwise,} \end{cases} \tag{1.4}$$

where the vector D contains the dose history and the vector ψ defines a set of PK parameters, clearance, Cl, and volume of distribution, V. This model is represented schematically in Figure 1.1 and the corresponding concentration over time (without error) in the single compartment is shown in Figure 1.2. The goal of compartmental PK modeling is to determine the parameter values of this system given known dose history and observed concentrations (which are subject to stochastic error).

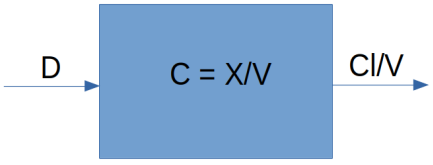

D: Given Dose
X: Amount of Drug
C: Drug Concentration
V: Volume of Distribution
Cl: Clearance

Figure 1.1: One-compartment pharmacokinetic model. Dose (D) and drug concentration (C) are observed quantities. Clearance (Cl) and volume of distribution (V) are parameters to be estimated.

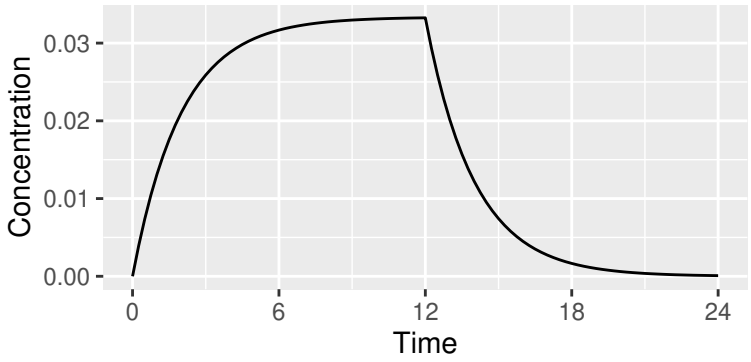

Figure 1.2: Concentration over time for a one-compartment model with single infusion (dose $D = 8$, infusion start time $t_d = 0$, infusion duration $t_{inf} = 12$, clearance $Cl = 20$, volume of distibution $V = 40$).

1.3.2 Population Pharmacokinetic Models

The structural model in Equation 1.4 describes the change in concentration over time for a single individual. Population PK models extend this idea to analyze data from multiple individuals. The PK profiles for all individuals are assumed to follow the same structural model, but PK parameter values can differ from subject to subject. Population PK models can estimate both population and individual-specific parameters and quantify variability within and between individuals. In addition they can incorporate covariates such as weight, age, and organ function, that may affect each individual's PK profile, which is of interest in many studies.

The foundational paper by Lewis Sheiner and Stuart Beal (1977) developed the general framework for population PK modeling [44]. In their paper population parameters were estimated directly using maximum likelihood and first-order Taylor series approximation (FO). This paper introduced important concepts in population PK estimation such as partially pooled interindividual variability that is a balance between complete homogeneity (identical PK parameters for every individual) and complete heterogeneity (distinct, independent PK parameters for each individual) and led to development of the nonlinear mixed effect model (NONMEM) software that has become a gold standard for population PK analyses. Additional papers by Sheiner and Beal [45–47] showed how to perform population PK analysis for different models and further evaluated the performance of the NONMEM approach. Further literature expanded the scope of these analyses to include covariates for the individual effects [48–50], compare strategies for covariate building [51], and develop more accurate frequentist approximations [52].

Bonate (2005) provides a summary of the history and development of population PK modeling [53] with references to many seminal papers. In addition, many texts, such as those by Gibaldi and Perrier (1982) [42], Davidian and Giltinan (1998) [54], Bonate (2011) [55], and Rosenbaum (2017) [43], provide more details on the theory, derivation, and estimation of population PK models.

1.3.3 Bayesian Population Pharmacokinetic Models

Fully Bayesian population PK models were first described in the mid-1990s by Wakefield (1994) and others [56–61]. They advocate the Bayesian approach for several reasons including calculation of the full posterior distribution to better account for uncertainty, principled incorporation of prior information, and inference that does not rely on asymptotic arguments for estimates and tests. The Bayesian paradigm is also well-suited for prediction and provides the ability to assess sensitivity to modeling assumptions.

The hierarchical Bayesian population PK model described by Wakefield (1996) [60] and Wakefield, Aarons, and Racine-Poon (1999) [62] has with three stages or levels: an individual level describing each subject's PK profile and intra-individual error, a population level describing how PK parameters vary between individuals based on fixed and random effects, and a third level with priors describing the distribution of parameters at the first two stages.

Let y_{ij} be the observed concentration for subject, $i = 1,\ldots,N$, at the jth time point, t_{ij}, $j = 1,\ldots,n_i$, where n_i is the number of observed concentrations for subject i. $y_i = (y_{i1},\ldots,y_{in_i})$ and $y = (y_1,\ldots,y_N)$ are observed concentrations for subject i and the entire study population, respectively. Stage one models data at the individual level:

$$y_{ij} = f_1(D_i, \psi_{ij}, t_{ij}) + \varepsilon_{ij} \tag{1.5}$$

$$\varepsilon_{ij} \sim N(0, \sigma_{ij}^2), \quad \sigma_{ij}^2 = f_1(D_i, \psi_{ij}, t_{ij})^{\gamma}\sigma_y^2, \quad \gamma \geq 0. \tag{1.6}$$

The function $f_1(\cdot)$ is the solution to the system of differential equations defining the compartmental PK model (such as Equation 1.4) where D_i is dose history such as dose amounts and time of administration for subject i, ψ_{ij} is the p-dimensional PK parameter vector for subject i at t_{ij}, and ε_{ij} are independent random variables with mean zero representing intra-individual variability (residual error). For the intra-individual term in Equation 1.6, setting $\gamma = 0$ gives an additive error model and $\gamma = 2$ yields a constant coefficient of variation (CV) or proportional error model. A combined additive and proportional error model can be constructed by adding two error terms to $f_1(\cdot)$, $\varepsilon_{ij,add}$ and $\varepsilon_{ij,prop}$ where the corresponding $\sigma_{ij,add}^2 = \sigma_{y,add}^2$ and $\sigma_{ij,prop}^2 = f_1(D_i, \psi_{ij}, t_{ij})^2\sigma_{y,prop}^2$. Alternatively, a lognormal or exponential error model for stage 1 is:

$$\log y_{ij} = \log f_1(D_i, \psi_{ij}, t_{ij}) + \varepsilon_{ij}$$

$$\varepsilon_{ij} \sim N(0, \sigma_y^2)$$

The first stage is summarized by the density $p_1(y|\psi, \sigma_y^2)$.

Stage two models differences between individuals at the population level. A multivariate regression model for the latent PK parameters of individual i is:

$$\psi_{ij} = f_2(\mu, X_{ij}) + \eta_i, \tag{1.7}$$

where μ is the population regression coefficients (fixed effects), X_{ij} is the design matrix comprised of individual-specific covariates at time t_{ij}, and η_i is individual-specific deviations (random effects) with $E[\eta_i] = 0$ and covariance matrix Ω. The η_i are usually assumed to follow a multivariate normal or multivariate Student-t distribution. If the function f_2 is linear, $f_2(\mu, X_{ij}) = \mu X_{ij}$. Further the PK parameters ψ_{ij} are frequently transformed since they must be strictly positive and are often right-skewed [55]. For example, if the PK parameter vector in a one-compartment IV model is redefined as $\psi_{ij} = \{\log Cl_{ij}, \log V_{ij}\}$ the stage two model would typically be multivariate lognormal or multivariate log-Student-t. The second stage is represented by the density $p_2(\psi|X, \mu, \Omega)$.

Finally, Stage three specifies priors for the parameters from Stages one and two, $p_3(\mu, \Omega, \sigma_y^2) = p(\mu)p(\Omega)p(\sigma_y^2)$ where $p(\mu)$ is the prior for the fixed effects at Stage two, $p(\Omega)$ is the prior for random effects at Stage two, and $p(\sigma_y^2)$ is the prior for residual error variance at Stage one. The hyperparameters are fixed to represent prior knowledge (or lack thereof) about the parameters. Combining all three stages the posterior is $p(\theta = \{\mu, \Omega, \sigma_y^2\}|z = \{y, X, D\}) \propto p_1(y|\psi, \sigma_y^2)p_2(\psi|X, \mu, \Omega)p_3(\mu, \Omega, \sigma_y^2)$ where θ contains all the parameters to be estimated and z contains all observed data.

Inference for Bayesian population PK models is usually performed using MCMC sampling methods such as those implemented through the `PKBugs` interface to `WinBUGS` or using the HMC sampler with the `Torsten` library and `Stan` probabilistic programming language [63–66]. However MCMC can be slow to converge and may scale poorly with larger sample sizes, complex models, or high-dimensional parameter spaces. Slow convergence time, in turn, presents a bottleneck in the iterative model development process of fitting, checking, adjusting, and refitting models [4,55,67,68]. Therefore, it is of interest to explore alternative strategies for Bayesian inference which can balance speed and accuracy. In Chapter 3, we investigate ADVI for Bayesian population PK models as a potential solution to the model development bottleneck.

CHAPTER 2

Bayesian Cumulative Probability Models for Continuous and Mixed Outcomes

2.1 Introduction

In this paper, we develop Bayesian CPMs for continuous and mixed outcomes. They are distinguished from other Bayesian semiparametric approaches by their use of a simpler parametric prior specification. Bayesian CPMs inherit many of the properties of CPMs estimated using non-parametric maximum likelihood estimation and have additional benefits: interpretation using posterior probabilities, inference for quantities of interest without using asymptotic approximations, and the ability to incorporate available prior information. A primary challenge when implementing Bayesian CPMs for continuous outcomes is the specification of priors for the intercept parameters used to estimate the unspecified transformation $H(\cdot)$ (Equation 1.3) and we describe several proposed strategies. Through simulations, we explore characteristics of Bayesian CPMs using several model specifications and prior combinations. A case study of HIV biomarker data with outcomes that are both right-skewed and censored at a lower limit of detection provides a real-world example. We conclude with a discussion, including advantages, current limitations, and potential extensions, and provide some recommendations for using Bayesian CPMs.

2.2 Methods

2.2.1 Cumulative Probability Model Formulation

Let Y_i be the outcome for unit $i = 1,\ldots,n$ with p covariates $X_i = (X_{i1},\ldots,X_{ip})$ such that each Y_i falls into one of $j = 1,\ldots,J$ ordered categories. The Y_i can be modeled using a $Categorical(\pi_i)$ – or $Multinomial(1,\pi_i)$ – distribution where $\pi_i = (\pi_{i1},\ldots,\pi_{iJ})$ are the probabilities of unit i being in category j and $\sum_{j=1}^{J}\pi_{ij} = 1$. The value of π_{ij} is dependent on x_i, but we suppress the conditional notation for clarity. The cumulative probability of falling into category j or a lower category is $Pr(Y_i \leq j) = \eta_{ij} = \sum_{k=1}^{j}\pi_{ik}$. The CPM relates the cumulative probabilities to the observed covariates through a monotonically increasing link function $G^{-1}(\eta_{ij}) = \gamma_j - x_i'\beta$. Common choices for the link function are logit, $G^{-1}(p) = \log\left(\frac{p}{1-p}\right)$; probit, $G^{-1}(p) = \Phi^{-1}(p)$ where $\Phi^{-1}(p)$ is the quantile function for a standard normal distribution; and loglog, $G^{-1}(p) = -\log(-\log(p))$. For observed data $\{y_i, x_i\}$ the model can be expressed as

$$Pr(y_i \leq j|x_i,\beta,\gamma) = \eta_{ij} = G(\gamma_j - x_i'\beta), \tag{2.1}$$

where the γ_j are ordered continuous intercept parameters $-\infty \equiv \gamma_0 < \gamma_1 < \cdots < \gamma_{J-1} < \gamma_J \equiv \infty$, β is a vector of p coefficients, and the function $G(\cdot)$ is a CDF defined as the inverse of the link function: standard logistic, standard normal, and standard Gumbel for the logit, probit, and loglog links, respectively. For identifiability, the linear predictor $x_i'\beta$ does not include an intercept. The conditional probabilities of category membership are

$$\pi_{ij} = \eta_{i,j} - \eta_{i,j-1} = G(\gamma_j - x_i'\beta) - G(\gamma_{j-1} - x_i'\beta) \tag{2.2}$$

The likelihood for an independent and identically distributed sample of outcomes $y = (y_1, \ldots, y_n)$ with corresponding covariates $x = (x_1, \ldots, x_n)$ is

$$p(y|x, \gamma, \beta) = \prod_{j=1}^{J} \prod_{i:y_i=j} [G(\gamma_j - x_i'\beta) - G(\gamma_{j-1} - x_i'\beta)] \tag{2.3}$$

For continuous data with no ties $J = n$; letting $r(y_i)$ be the rank of y_i, the likelihood reduces to

$$p(y|x, \gamma, \beta) = \prod_{i=1}^{n} [G(\gamma_{r(y_i)} - x_i'\beta) - G(\gamma_{r(y_i)-1} - x_i'\beta)] \tag{2.4}$$

To complete the model specification we define priors for the parameters $p(\beta, \gamma)$. We assume a priori independence between β and γ so $p(\beta, \gamma) = p(\beta)p(\gamma)$. To simplify the model formulation we also assume noninformative priors for the regression coefficients, $p(\beta) \propto 1$; however weakly informative or informative priors can also be used.

Specifying priors for γ is more challenging because of the ordering restriction and dimensionality. Several approaches have been suggested in the traditional CPM setting where $J \ll n$. McKinley et al. (2015) [69] and Congdon (2005) [24] describe a sequentially truncated prior distribution: $p(\gamma) = p(\gamma_1) \prod_{j=2}^{J-1} p(\gamma_j|\gamma_{j-1})$ where $\gamma_1 \in \mathbb{R}$ and the support of γ_j for $j = 2, \ldots, J-1$ is (γ_{j-1}, ∞). For example using normal and truncated normal priors, $p(\gamma_1) \sim N(0, \sigma_\gamma^2)$ and $p(\gamma_j|\gamma_{j-1}) \sim N(0, \sigma_\gamma^2) I(\gamma_{j-1}, \infty)$. A second approach described by Albert and Chib (1997) [20] defines the prior on a transformation of the intercepts to an unconstrained space; first normalizing γ_0 to 0 so $0 \equiv \gamma_0 < \gamma_1 < \cdots < \gamma_{J-1} < \gamma_J \equiv \infty$ and then letting $\delta_1 = \log(\gamma_1)$ and $\delta_j = \log(\gamma_j - \gamma_{j-1})$, $2 \leq j \leq J-1$ a multivariate prior can be assigned, e.g. $\delta \sim N_{J-1}(\mu_0, \Sigma_0)$. Both approaches provide priors that satisfy the ordering restriction, but may be cumbersome when the number of distinct categories is high. The first requires specification of the distribution and its hyperparameters, then sampling from the sequential series of $J-2$ truncated distributions; the second requires specification of the $J-1$ dimensional μ_0 vector and the $J-1 \times J-1$ dimensional covariance matrix Σ_0.

We instead adopt a third approach which defines a prior on π_i for a prespecified covariate vector and utilizes the transformation defined by $G(\cdot)$ to induce a prior on γ [70]. Let $\pi_{\cdot j} \equiv Pr(r(y) = j|x = 0)$ be the probability of being in category j when all covariates are 0 and $\pi_{\cdot} = (\pi_{\cdot 1}, \ldots, \pi_{\cdot J})$. It may be useful to center the covariates by using $x' = x - \bar{x}$ in place of x. Then $\pi_{\cdot j}$ is the probability of being in category j when all covariates are at their mean value. From Equation 2.2 it follows that

$$\pi_{\cdot j} = G(\gamma_j - 0) - G(\gamma_{j-1} - 0) = G(\gamma_j) - G(\gamma_{j-1}) \tag{2.5}$$

These equations define a transformation $h(\gamma) = \pi_{\cdot}$ between the intercept parameters and probabilities of category membership when $X = 0$. Conversely,

$$\sum_{k=1}^{j} \pi_{\cdot k} = \sum_{k=1}^{j} [G(\gamma_k) - G(\gamma_{k-1})] = G(\gamma_j) \tag{2.6}$$

so $G^{-1}\left(\sum_{k=1}^{j} \pi_{\cdot k}\right) = \gamma_j$ defines the inverse transformation $h^{-1}(\pi_{\cdot}) = \gamma$. Because y has a multinomial distribution a conjugate Dirichlet distribution with hyperparameters α is a natural choice of prior for $\pi_{\cdot}$. Setting $p(\pi_{\cdot}|\alpha) \propto \prod_{j=1}^{J} \pi_{\cdot j}^{\alpha_j - 1}$ the posterior distribution is

$$\begin{aligned} p(\gamma, \beta|x, y) &\propto p(\gamma)p(\beta)p(y|x, \gamma, \beta) && (2.7) \\ &\propto p(h(\gamma))|\mathscr{J}|p(\beta)p(y|x, \gamma, \beta) && (2.8) \\ &\propto p(\pi_{\cdot}|\alpha)|\mathscr{J}|p(\beta)p(y|x, \gamma, \beta) && (2.9) \end{aligned}$$

where $\mathscr{J}$ is the Jacobian of the transformation $h(\gamma) = \pi_{\cdot}$. Letting $\Omega = \sum_{j=1}^{J} \pi_{\cdot j} = 1$ be the constraint that all category probabilities sum to 1, the entries in $\mathscr{J}$ where $\mathscr{J}_{r,c}$ is the term in row r and column c are

$$\begin{aligned} \mathscr{J}_{j,1} &= \frac{\partial \pi_{\cdot j}}{\partial \Omega} = 1, \\ \mathscr{J}_{j+1,j+1} &= \frac{\partial \pi_{\cdot j+1}}{\partial \gamma_j} = \frac{\partial}{\partial \gamma_j}\left[G(\gamma_{j+1}) - G(\gamma_j)\right] = -g(\gamma_j), \\ \mathscr{J}_{j,j+1} &= \frac{\partial \pi_{\cdot j}}{\partial \gamma_j} = \frac{\partial}{\partial \gamma_j}\left[G(\gamma_j) - G(\gamma_{j-1})\right] = g(\gamma_j), \end{aligned}$$

where $j = 1, \ldots, J-1$, $g(\cdot)$ is the density function of the distribution $G(\cdot)$, and $\mathscr{J}_{r,c} = 0$ for all other entries; the form of the Jacobian is

$$\begin{vmatrix} 1 & g(\gamma_1) & 0 & 0 & \cdots & 0 \\ 1 & -g(\gamma_1) & g(\gamma_2) & 0 & \cdots & 0 \\ 1 & 0 & -g(\gamma_2) & g(\gamma_3) & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 0 & 0 & 0 & -g(\gamma_{j-1}) & g(\gamma_j) \\ 1 & 0 & 0 & 0 & 0 & -g(\gamma_j) \end{vmatrix} \tag{2.10}$$

While it is possible to define separate α_j parameters for each category, we restrict our attention to symmetric Dirichlet distributions which use a single α value for all categories (i.e., $\alpha_1 = \alpha_2 = \cdots = \alpha_J$) so $\alpha = \alpha 1$ where 1 is a $J-1$ dimensional vector of 1s. The symmetric Dirichlet prior on $\pi_\cdot$ along with the inverse transformation $h^{-1}(\cdot)$ defined in Equation 2.6 induces a prior for γ with α controlling the concentration of the induced prior. For example, Figure 2.1 shows induced γ priors assuming a probit link for several combinations of concentration parameter and number of categories. The priors are approximately distributed around the intercepts that result under an assumption of equal probability for all categories when $X = 0$; that is, for $J = n$ the values $G^{-1}(\sum_{k=1}^{j} 1/n) = G^{-1}(j/n) = \hat{\gamma}_{j|X=0}$. The prior choices correspond to several options for a multinomial-Dirichlet model [4]: a uniform Dirichlet ($\alpha = 1$), the multivariate Jeffreys prior ($\alpha = 1/2$), an overall objective prior recommended by Berger et al. (2015) [71] ($\alpha = 1/J$), and two additional 'reciprocal' priors ($\alpha = 1/(2 + (J/3))$ and $\alpha = 1/(0.8 + 0.35J)$). The last two priors were found using a trial-and-error procedure in a simulation study with the aim of minimizing the difference between the posterior mean and mode intercept estimates and the corresponding maximum likelihood intercept estimates. As the number of categories increases the uniform and Jeffreys priors more strongly favor intercepts assuming equal probability for all categories; in contrast, the three reciprocal priors are adjusted to maintain the same degree of concentration relative to the equal probability intercepts, $\hat{\gamma}_{j|X=0}$.

In multiparameter models, the choice of an objective reference prior depends on the parameter or statistic of interest (e.g. β, conditional CDF, conditional mean) [71]. Without prior information, we seek a value of α with minimal impact on inference for a variety of settings and quantities of interest while still producing posterior estimates that can be sampled well by the MCMC algorithm.

2.2.2 Estimation

The model in (2.9) is implemented using the `R` interface to Stan [72] which performs MCMC sampling using no-U-Turn Hamiltonian Monte Carlo [4,73]. The `R` package `bayesCPM` which implements the Bayesian CPM model described in this paper is available at https://github.com/ntjames/bayesCPM/tree/master/pkg.

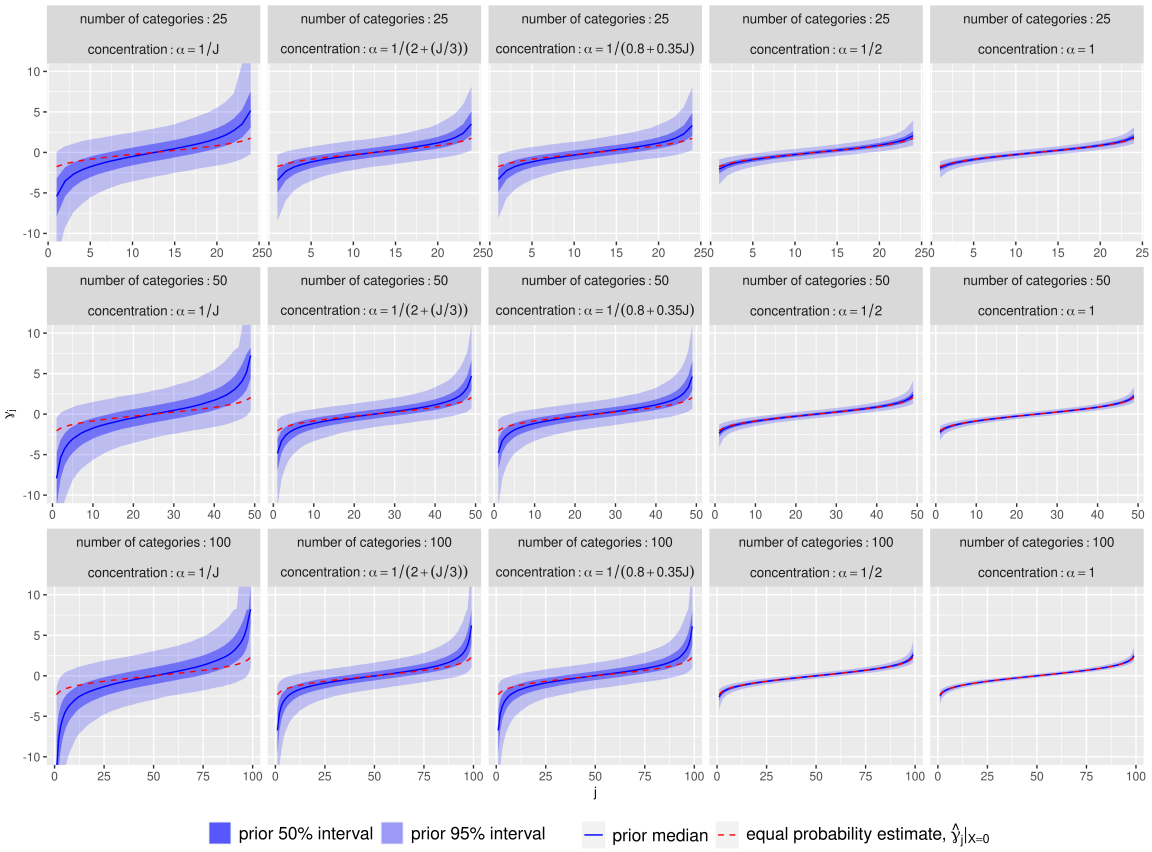

Figure 2.1: Induced γ priors under probit link, $G^{-1}(\cdot) = \Phi^{-1}(\cdot)$. Each subplot displays the median and credible intervals of γ_j for $j = 1, \ldots, J-1$

2.2.3 Posterior Conditional Quantities

Using the S draws from the posterior distribution, $(\tilde{\gamma}^{(s)}, \tilde{\beta}^{(s)})$ where $s = 1,\ldots,S$, it is straightforward to calculate the distribution of the posterior conditional CDF, mean, quantiles or other functions of the parameters. For example, the distribution of the posterior conditional CDF at y_j with covariates x can be approximated by the S values $\tilde{F}^{(s)}(y_j|x) = G(\tilde{\gamma}^{(s)}_{r(y_j)} - x^T\tilde{\beta}^{(s)})$ and the complete conditional CDF can be obtained by a step function connecting $\tilde{F}^{(s)}(y_j|x)$ for $j = 1,\ldots,J$. The posterior mean distribution conditional on covariate vector x is approximated by $\tilde{E}^{(s)}[Y|x] = \sum_{j=1}^{J} y_j \tilde{f}^{(s)}(y_j|x)$ where $\tilde{f}^{(s)}(y_j|x) = \tilde{F}^{(s)}(y_j|x) - \tilde{F}^{(s)}(y_{j-1}|x)$ and $\tilde{F}^{(s)}(y_0|x) \equiv 0$ so $\tilde{f}^{(s)}(y_1|x) = \tilde{F}^{(s)}(y_1|x)$. Note that for mixed continuous/discrete outcomes, such as those arising from a detection limit, data below or above the limit do not have a known y_j value; in this case a value must be assigned to calculate the conditional mean. To estimate the q^{th} posterior conditional quantile we first find $y_j^{(s)} = \inf\{y : \tilde{F}^{(s)}(y|x) \geq q\}$ and the next smallest value $y_{j-1}^{(s)}$, then use linear interpolation to find quantile $y_q^{(s)}$ where $y_{j-1}^{(s)} < y_q^{(s)} < y_j^{(s)}$. For each of these functionals, point and interval estimates can be obtained by summarizing the S values obtained from the posterior parameter draws without using asymptotic approximations. For example, the mean of the posterior conditional CDF distribution is $\frac{1}{S}\sum_{s=1}^{S} \tilde{F}^{(s)}(y_j|x)$ and the 2.5% and 97.5% percentiles of the $y_q^{(s)}$ values are the bounds of a 95% credible interval for the q^{th} posterior conditional quantile.

2.3 Simulations

2.3.1 Set-up

To evaluate the properties of the Bayesian CPM for continuous and mixed outcomes we generate data from several simulation scenarios:

$$1.\ Y = \exp(X_1\beta_1 + X_2\beta_2 + \varepsilon) \quad \varepsilon \sim N(0,1)$$

$$2.\ Y = \exp(X_1\beta_1 + X_2\beta_2 + \varepsilon) \quad \varepsilon \sim Logistic(0,1/3)$$

$$3.\ Y = X_1\beta_1 + X_2\beta_2 + \varepsilon \quad \varepsilon \sim Gumbel(0,1)$$

where $\beta_1 = 1$, $\beta_2 = -0.5$, $X_1 \sim Bernoulli(0.5)$ and $X_2 \sim N(0,1)$. For each scenario a second set of simulations was used to evaluate a mixed discrete/continuous outcome with a lower limit of detection; for scenario (1) and (2) values of $Y < 1$ to were set to 1, for scenario (3) values of $Y < 0$ were set to 0. The uncensored and censored outcome data based on (1) and (3) were evaluated using a Bayesian CPM with the properly specified probit and loglog links, respectively. For scenario (2) a logit link Bayesian CPM (which implies $\varepsilon \sim Logistic(0,1)$) was used. In each of the six outcome models, three α concentration hyperparameters $(1/J,\ 1/(2+(J/3))$, and $1/(0.8+0.35J))$ were considered for $p(\pi_\cdot|\alpha)$ for 18 model and prior combinations.

Sample sizes $n = 25, 50, 100, 200$ and 400 were used under each model/prior combination for a total of 90 simulation models. 1,000 datasets were generated under each simulation model.

We examine the average percent bias of the posterior median for parameters β_1 and β_2 and five γ_j parameters corresponding to y values spaced across the range of the data: For scenario (1) $y = \{y_1 = e^{-1}, y_2 = e^{-0.33}, y_3 = e^{0.5}, y_4 = e^{1.33}, y_5 = e^{2}\}$, for scenario (2) $y = \{y_1 = e^{-0.5}, y_2 = e^{0}, y_3 = e^{0.5}, y_4 = e^{1}, y_5 = e^{1.5}\}$ and for scenario (3) $y = \{y_1 = -0.3, y_2 = 0, y_3 = 0.5, y_4 = 1.5, y_5 = 2.5\}$. For the censored outcomes, estimates are only available for the values of y above the censoring threshold. We also calculate average percent bias of the conditional CDF for y when $X_1 = 1$ and $X_2 = 1$, and the conditional median, mean and 20th percentile at $(X_1 = 1, X_2 = 1)$ and $(X_1 = 1, X_2 = 0)$.

2.3.2 Results

A Bayesian CPM was fit to each of the 1000 simulation datasets for each scenario/prior/sample size combination. For each simulation dataset, the median of the posterior distribution of the parameter or conditional CDF, mean, or quantile was used as a point estimate. These point estimates were compared to the true value from the generating model and the results averaged over all simulation datasets. Each model was run with 2 MCMC chains using 2000 warmup and 2000 sampling iterations each. Retaining only the sampling iterations from each chain for inference resulted in a total of 4000 posterior parameter vector draws per model.

In general, the Bayesian CPM had reasonable performance in estimating parameters and conditional quantities for the simulation settings explored; especially for larger sample sizes. However, performance was poor for some quantities and may be sensitive to the conditioning covariate values and censoring threshold. The three α values produced similar results for most scenarios and no prior choice was best across all parameters and quantities of interest.

2.3.2.1 Parameters

For scenario (1) using a properly specified probit link CPM, the average percent bias in the posterior median for β_1, β_2, and γ_{y_k} is shown in Figure 2.2 for the uncensored and censored outcome data. Average percent bias was largest for the smallest sample sizes, but the direction and magnitude of the bias depended on the outcome, concentration prior and parameter. Across both outcomes, the estimates of β_1 and all γs were larger using the $\alpha = 1/J$ concentration prior than the $\alpha = 1/(0.8 + 0.35J)$ or $\alpha = 1/(2 + (J/3))$ concentration priors while the β_2 estimates were smaller with $\alpha = 1/J$. For the β parameters, the priors $\alpha = 1/(0.8 + 0.35J)$ and $\alpha = 1/(2 + (J/3))$ produced less biased estimates than $\alpha = 1/J$ for both outcomes. The situation was more complex for the γ parameters. With the uncensored outcome, the $\alpha = 1/J$ prior estimate was less biased for

γ_{y_1} and γ_{y_2}, but more biased for γ_{y_3}, γ_{y_4}, and γ_{y_5}; with the censored outcome, the $\alpha = 1/J$ prior estimate was less biased for γ_{y_3}, but more biased for γ_{y_4} and γ_{y_5}.

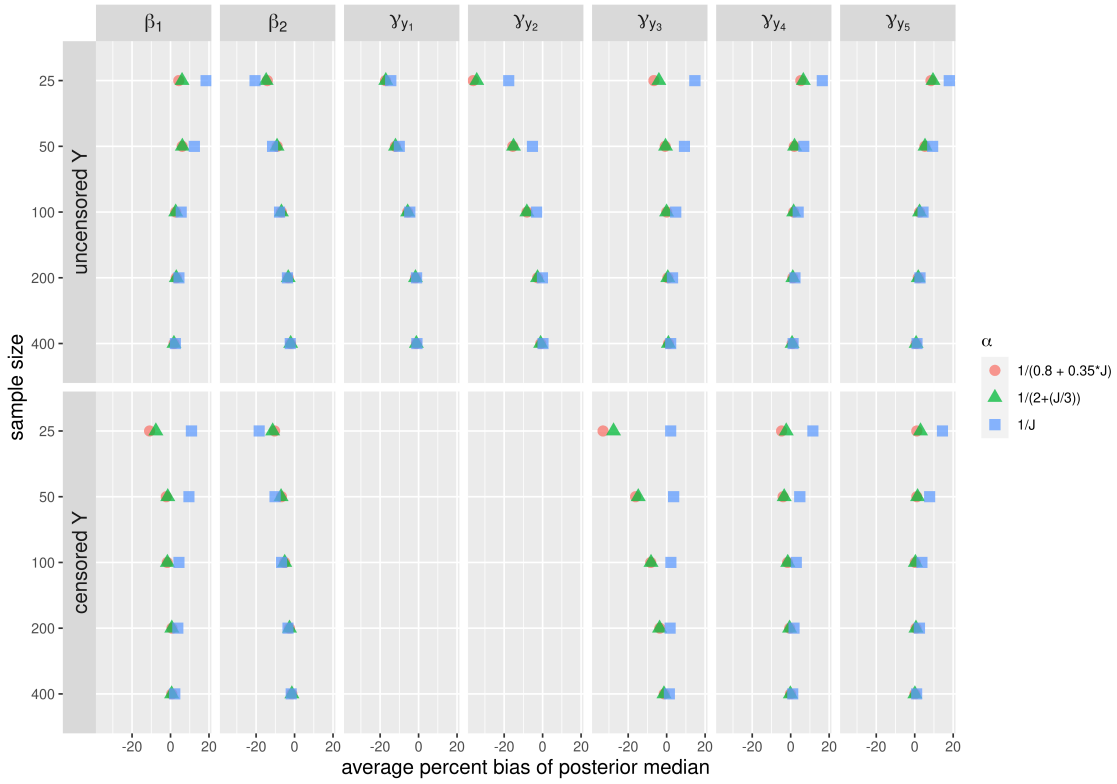

Figure 2.2: Percent bias in parameters for simulations using probit link

Figure 2.3 shows the average percent bias in the posterior median for β_1, β_2, and γ_{y_k} for scenario (2) using a logit link CPM. Unlike scenario (1), the assumed scale of the latent variable with logit link ($\varepsilon \sim Logistic(0,1)$) does not match the scale from the simulation model ($\varepsilon \sim Logistic(0,1/3)$). In this case it can be shown that the CPM parameter estimates are proportional to the parameters from the generating simulation model. Assume latent $Y^* = \beta' X + a\varepsilon$ with known $\varepsilon \sim F_\varepsilon$ and constant scaling factor $a > 0$, and observed $Y = H(Y^*)$ with increasing function $H(t)$. Then $Y = H(\beta' X + a\varepsilon) = H'(\xi' X + \varepsilon)$ where $\xi = a^{-1}\beta$ and $H'(t) = H(at)$ so $Pr(Y \leq y|X) = Pr(H'(\xi' X + \varepsilon) \leq y|X) = F_\varepsilon(H'^{-1}(y) - \xi' X)$. Using a CPM with link function F_ε^{-1} to analyze the observed outcome Y results in estimates of $\xi = a^{-1}\beta$ for the linear predictor coefficients and $H'^{-1} = a^{-1}H^{-1}$ for the intercept function. To compare the CPM model estimates (e.g. ξ, H'^{-1}) to the generating model parameters (β, H^{-1}) it is necessary to rescale by a. Conceptually this is equivalent to rescaling ε for the latent Y^* to match the assumed scale before fitting the CPM. Outside of simulations, the scale factor is not known but can be assumed to equal 1 without loss of generality because Y^* is latent; therefore rescaling is not necessary in practice. In general, simulation results were similar to those in Figure 2.2; bias was small with moderate sample sizes.

For scenario (3) using the correctly specified loglog link with an identity transformation, overall trends resembled those in scenario (1) (see Figure 2.18).

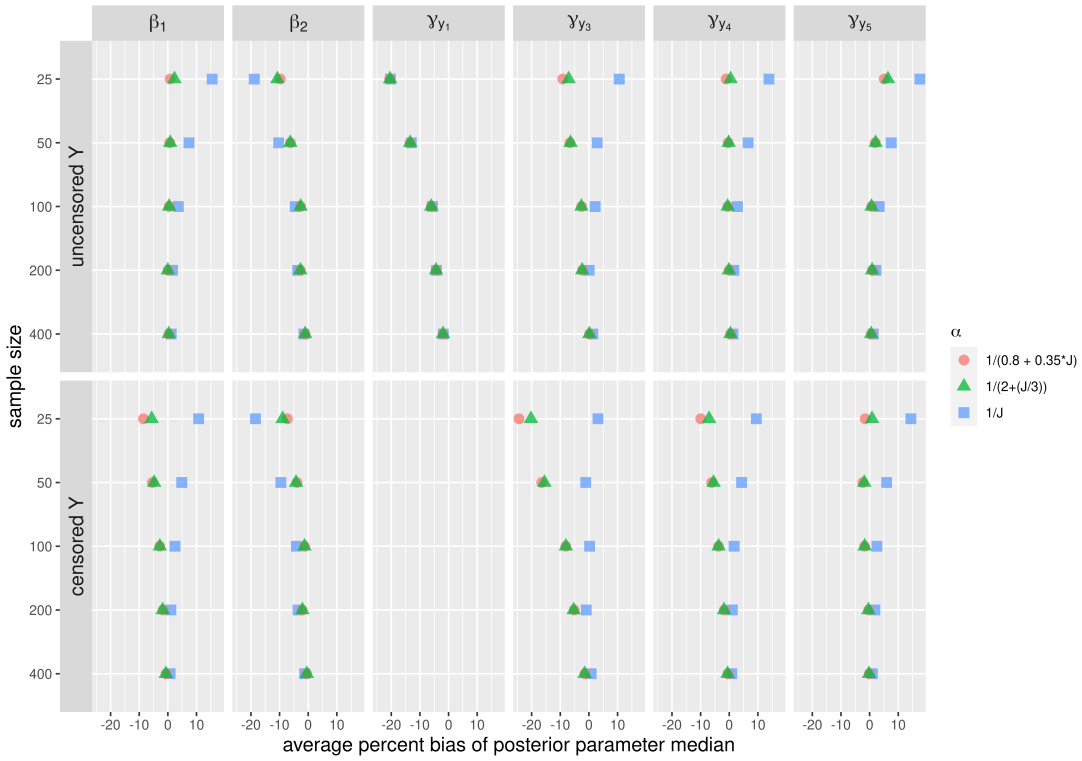

Figure 2.3: Percent bias in parameters for simulations using logit link

2.3.2.2 Conditional CDF

Figure 2.4 shows the average percent bias in the posterior conditional CDF, $F(y|X_1 = 1, X_2 = 1)$, for scenario (1). At the values $y_1 = e^{-1}, y_2 = e^{-0.33}, y_3 = e^{0.5}, y_4 = e^{1.33}, y_5 = e^{2}$ the true conditional CDF values were around 0.07, 0.20, 0.50, 0.8, and 0.93, respectively. For the uncensored outcome, the conditional CDF estimates had larger percent bias when $y < e^{0.5}$, especially for the sample sizes $n = 25$ and $n = 50$. This is not surprising, as it is difficult to estimate a conditional CDF at the tail of distribution with a small sample size. In addition, for conditional CDF estimates at $y < e^{0.5}$, the concentration prior $\alpha = 1/J$ produced estimates that were lower than the other reciprocal priors. The direction of the bias did not show a consistent trend across sample sizes. Similar patterns were seen for the censored outcome, less biased estimates for the CDF at higher y values and larger sample sizes. The results were much the same for scenarios (2) and (3) under both outcomes: larger average percent bias for the conditional estimates of $F(y|X_1 = 1, X_2 = 1)$ for lower values of y and smaller sample sizes (Figures 2.19 and 2.20).

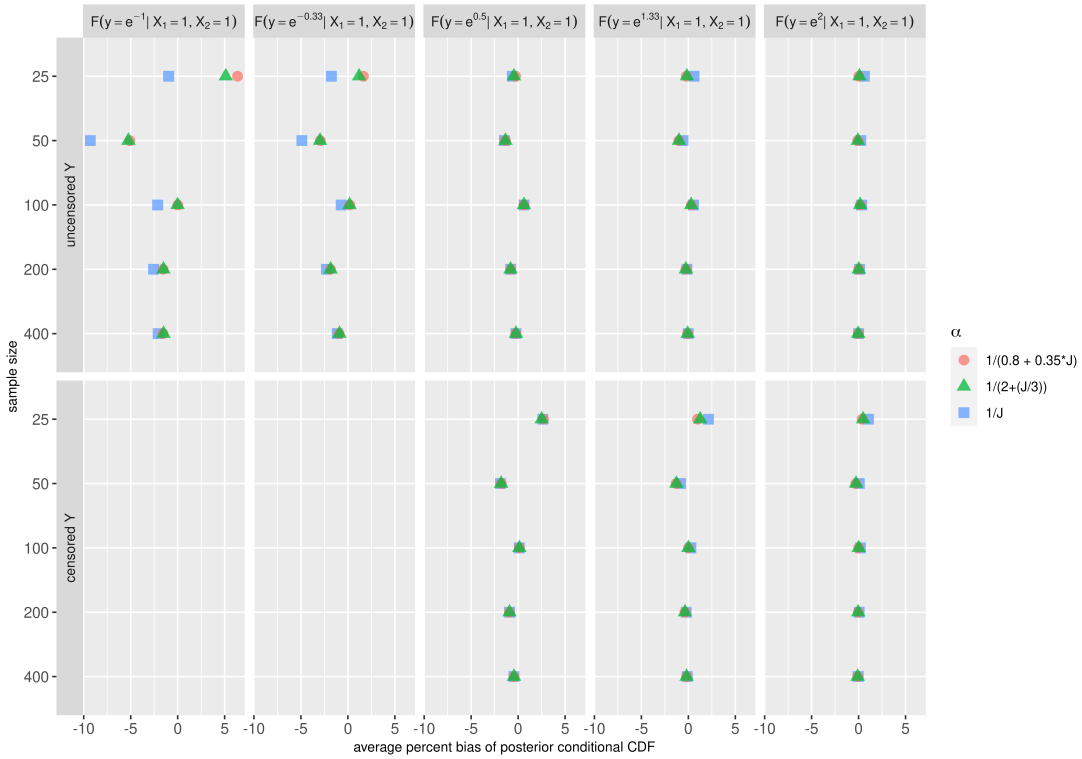

Figure 2.4: Percent bias in conditional CDF for simulations using probit link

2.3.2.3 Conditional Mean

The top row of Figure 2.5 presents the average percent bias in the posterior conditional mean for the uncensored simulation outcomes at $(X_1 = 1, X_2 = 0)$ and $(X_1 = 1, X_2 = 1)$ in scenario (1). For this scenario, the average percent bias was less than 5% for all sample sizes and priors. In contrast, the bottom row of Figure 2.5 shows the bias in posterior conditional mean estimates for the censored outcomes where a value of $y = 1$ was used in the conditional mean calculation for outcomes censored at $Y < 1$. Using the censoring threshold value for censored observations results in inflated average percent bias compared to the uncensored case depending on where the threshold falls in relation to the true conditional distribution. For example, the average percent bias of $E(Y|X_1 = 1, X_2 = 1)$ for the censored outcome in scenario (1) was around 40% even for the largest sample size. Results were similar for scenario (2) (see Figure 2.21).

For scenario (3) the average percent bias for the uncensored outcome ranges from -12.5% to -1.0% with larger bias for the $\alpha = 1/J$ prior and smaller n (Figure 2.6). As in the first two scenarios, the censored outcome (which replaced outcomes less than 0 with a value of $y = 0$) showed a positive shift in average percent bias at $(X_1 = 1, X_2 = 1)$ although to a much smaller degree than scenario (1).

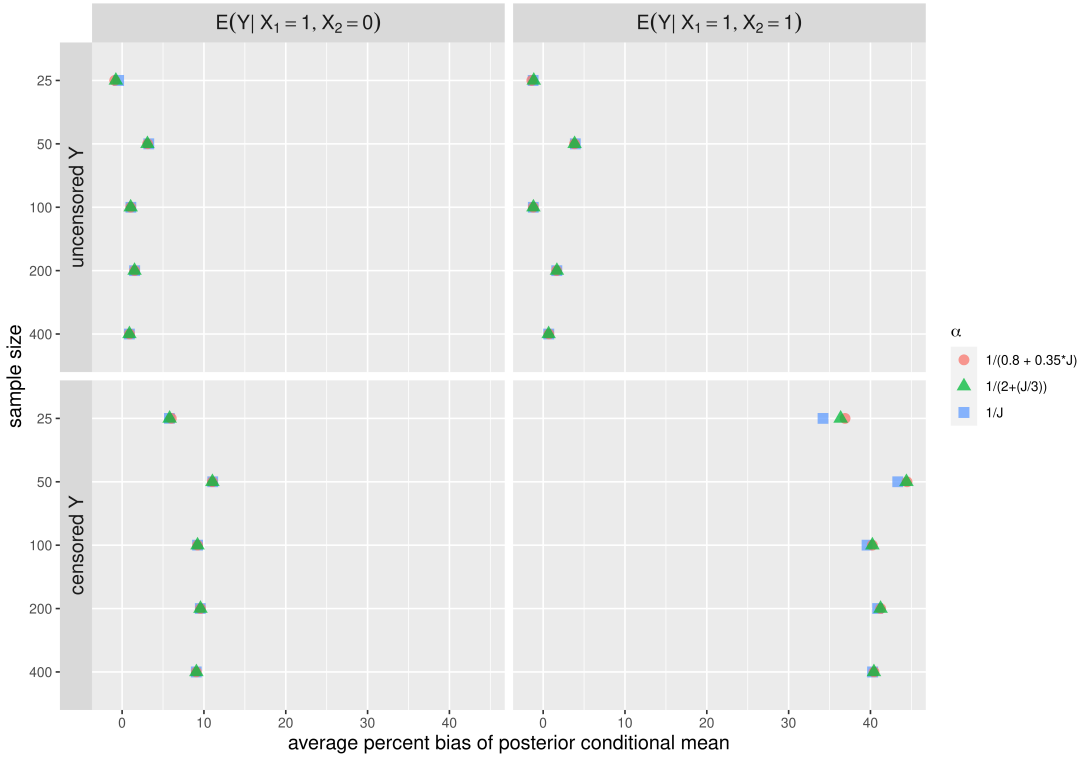

Figure 2.5: Percent bias in conditional mean for simulations using probit link

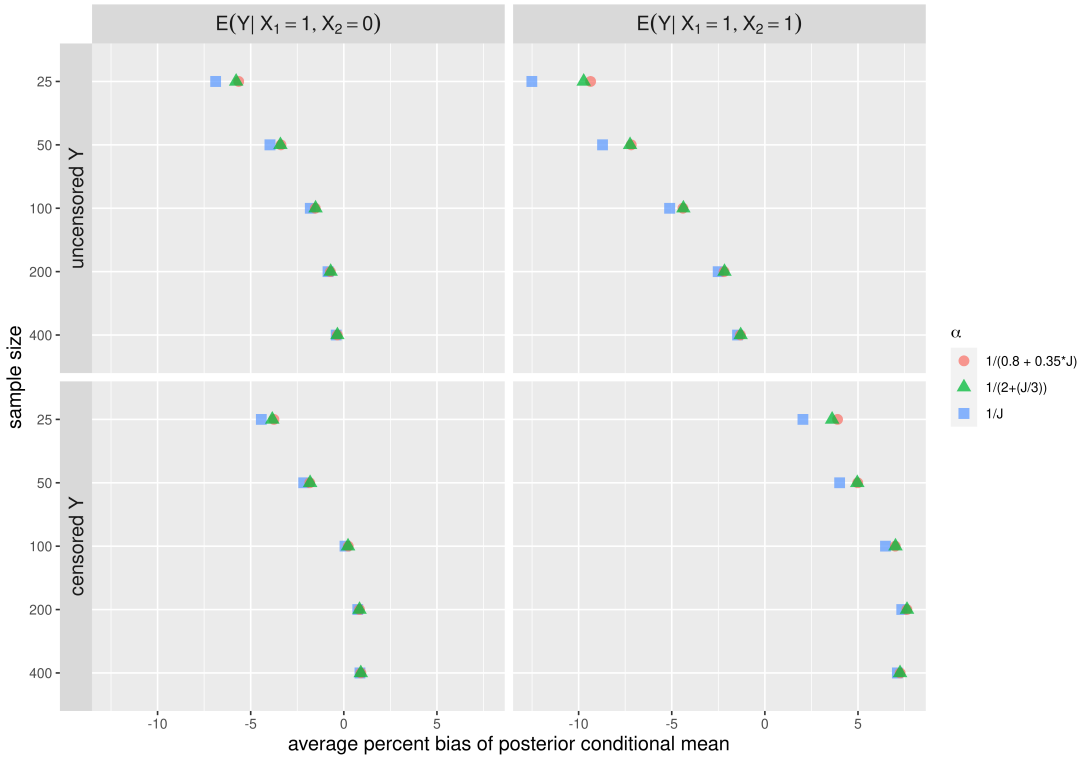

Figure 2.6: Percent bias in conditional mean for simulations using loglog link

2.3.2.4 Conditional Median and Quantiles

The simulation results for the conditional posterior median in scenario (1) are shown in Figure 2.7. Across outcomes, the conditional median estimates had a positive average percent bias for both $(X_1 = 1, X_2 = 0)$ and $(X_1 = 1, X_2 = 1)$ with smaller bias for larger sample sizes where there was more information to estimate the center of the distribution. There were negligible differences in average percent bias of the median estimates for the three α concentration parameter priors. The pattern looked similar for scenario (2) (Figure 2.22). For both outcomes under scenario (3), average percent bias in the conditional median estimate was smaller at $(X_1 = 1, X_2 = 0)$ than $(X_1 = 1, X_2 = 1)$. There were only small differences between the three α concentration parameters except with the smaller sample sizes (Figure 2.23).

Figure 2.8 presents the results for the posterior conditional 20th percentile in scenario (1). The uncensored outcome estimates were quite biased (between 25% and 90%) for the smaller sample sizes. The magnitude of the bias varied based on the values of the conditioning variables, X_1 and X_2, with larger bias when the conditional distribution was further from $\beta_1 = \beta_2 = 0$. The estimates of the conditional 20th percentile for the censored outcome in scenario (1) were similar to the uncensored outcome. When $(X_1 = 1, X_2 = 1)$ the true conditional $Q^{0.2}$ falls below the censoring threshold and does not have a specific numeric value. In this case percent bias could not be computed. For scenario (2) the estimates of the posterior conditional 20th percentile with the uncensored outcome were again positively biased for the smaller sample sizes with more bias for the $\alpha = 1/J$ concentration prior (Figure 2.24). Under scenario (3) the uncensored outcome estimates of the conditional 20th percentile had reasonably small average percent bias for all the priors and sample sizes except at $(X_1 = 1, X_2 = 1)$ when $n = 25$. Similar to scenario (1), the censored outcome estimates showed small average percent bias at $(X_1 = 1, X_2 = 0)$, but the true conditional 20th percentile fell below the censoring threshold for $(X_1 = 1, X_2 = 1)$ precluding calculation of percent bias (Figure 2.25).

2.3.2.5 Computation Time

Simulations were performed using `R` version 3.6.0 (2019-04-26) and `rstan` (Version 2.19.2) on a high-performance computing cluster running under CentOS Linux 7 (Core) with 1.90GHz or 2.40GHz Intel Xeon CPUs and up to 3 GB of memory per compute node. MCMC sampling time for the three scenarios is shown in Figure 2.9. Per chain sampling time increased approximately exponentially with sample size and was similar across scenarios and priors.

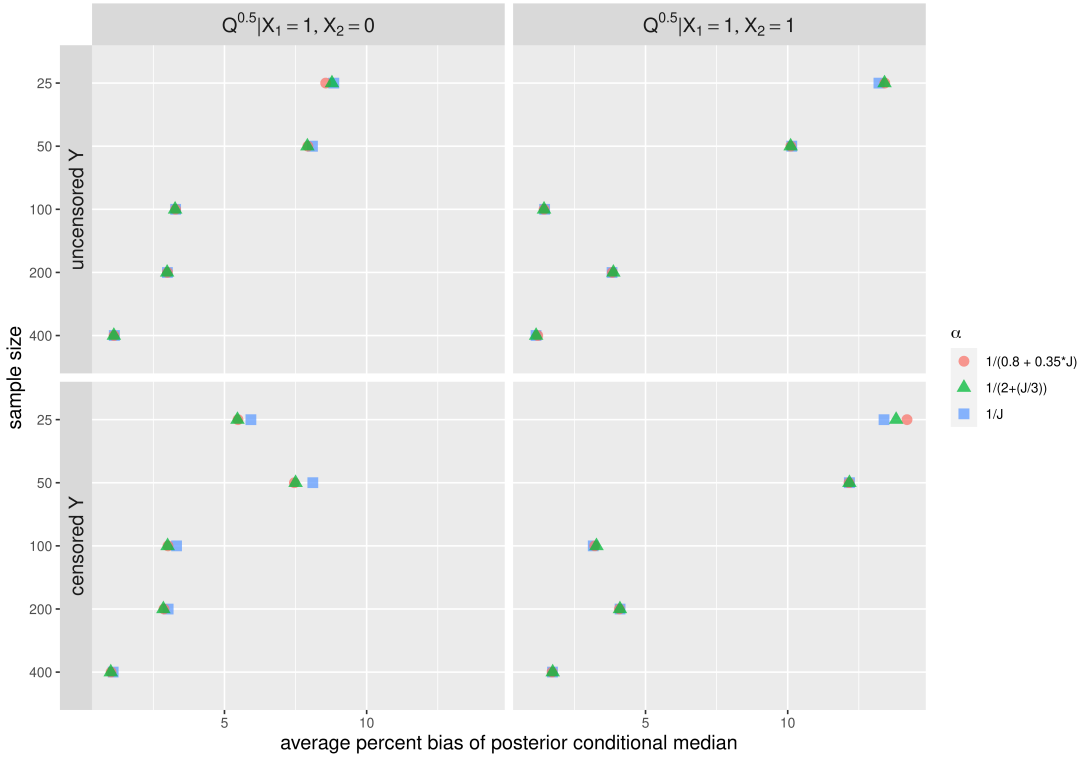

Figure 2.7: Percent bias in conditional median for simulations using probit link

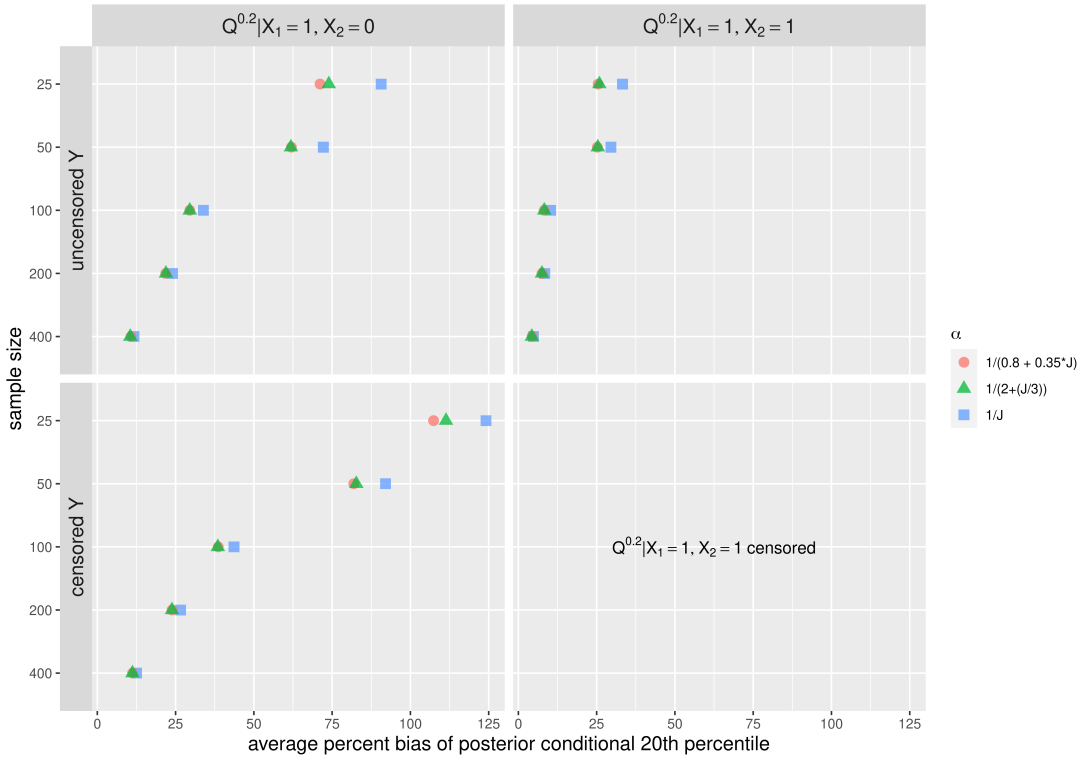

Figure 2.8: Percent bias in conditional 20th percentile for simulations using probit link

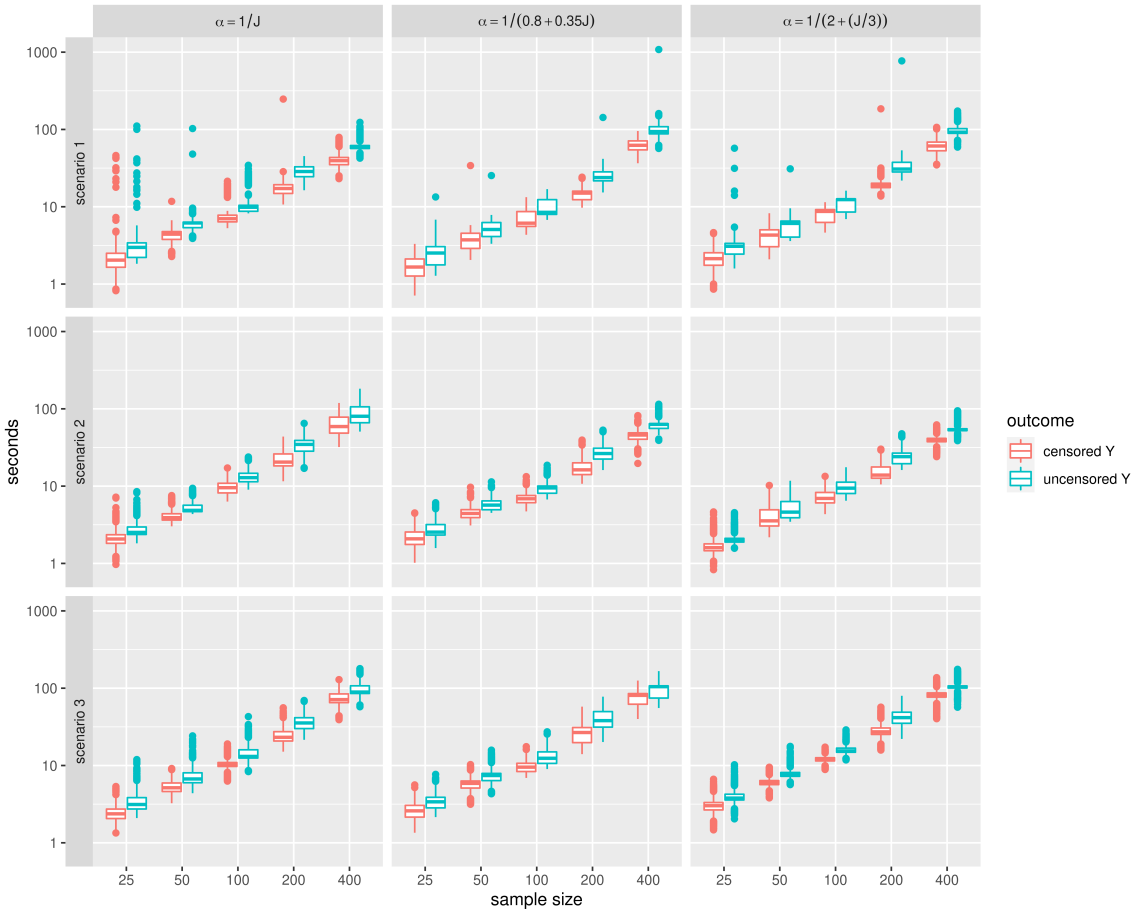

Figure 2.9: Per chain MCMC sampling time for three simulation scenarios. Each boxplot shows the sampling times required to produce 4000 posterior draws under the specified model/prior/sample size combination for 1000 simulation datasets

2.4 Case Study

2.4.1 Background and Methods

The data for the case study were collected from 216 HIV-positive adults on antiretroviral therapy in two cohort studies (Vanderbilt Lipoatrophy and Neuropathy Cohort (LiNC), n=147; Adiposity and Immune Activation Cohort (AIAC), n=69). Further details on the study design and cohorts are provided in Koethe et al. (2012; 2015) [74,75]. Because people living with HIV have increased risk of diabetes and cardiovascular disease, the aim of the analysis was to estimate the association between body mass index (BMI) and several inflammation biomarkers in this population, adjusting for additional covariates: age, sex, race, smoking status, study location and CD4 cell count.

We examine the biomarkers Interleukin 6 (IL-6) and Interleukin 1 beta (IL-1-β); both are right-skewed with 3% and 39% of values censored below the lower limit of detection, respectively. Censored values are set to 0. To account for skewness and censoring we fit Bayesian CPMs using logit, probit, and loglog link functions, noninformative β priors and a concentration parameter of either $\alpha = 1/J$ or $\alpha = 1/(0.8 + 0.35J)$ for the Dirichlet prior to estimate the association between BMI and the conditional mean, median, and 90th percentile of each biomarker.

We evaluate convergence using $\hat{R}$ scale reduction factor [4] and traceplots of MCMC draws. Model comparison is performed using the difference in expected log predictive density (ELPD) calculated using leave-one-out cross-validation [76]. Model fit is assessed with graphical checks of the posterior predictive distribution and posterior predictive p-values [4,77].

2.4.2 Results

For each of the two biomarker outcomes six model specifications were fit: probit, logit, or loglog link with $\alpha = 1/(0.8 + 0.35J)$ or $\alpha = 1/J$. Each model sampled from 2 chains with 2000 warmup and 4000 total iterations to produce 4000 posterior sample draws for each parameter. For all models, traceplots showed no issues with mixing or stationarity; further, all $\hat{R}$ potential scale reduction values were < 1.01 indicating likely convergence. Table 2.1 shows the difference in ELPD for the IL-6 and IL-1-β biomarker models. Based on the difference in ELPD, the CPM with loglog link and $\alpha = 1/(0.8 + 0.35J)$ was used for the both outcomes, however there is little difference in ELPD along the top several models.

2.4.2.1 IL-6 biomarker

A graphical check of 10 draws from the posterior predictive distribution compared to the observed IL-6 distribution (Figure 2.10) did not indicate any serious model misfit. In addition, there were no major

Table 2.1: Difference in Expected Log Pointwise Predictive Density for IL-6 Models and IL-1-β Models

Model	ELPD diff.	SE diff.
IL-6		
loglog link, $\alpha = 1/(0.8+0.35J)$	0.00	0.00
logit link, $\alpha = 1/J$	-1.94	6.65
probit link, $\alpha = 1/(0.8+0.35J)$	-1.99	6.64
logit link, $\alpha = 1/(0.8+0.35J)$	-4.96	6.68
loglog link, $\alpha = 1/J$	-7.16	5.38
probit link, $\alpha = 1/J$	-7.40	6.87
IL-1-β		
loglog link, $\alpha = 1/(0.8+0.35J)$	0.00	0.00
probit link, $\alpha = 1/(0.8+0.35J)$	-2.82	4.80
logit link, $\alpha = 1/(0.8+0.35J)$	-6.45	5.26
probit link, $\alpha = 1/J$	-7.25	4.85
loglog link, $\alpha = 1/J$	-10.81	4.03
logit link, $\alpha = 1/J$	-11.70	4.61

discrepancies between the model and data based on the posterior predictive p-values for the test quantities variance, skewness, and proportion of observations censored below the lower limit of detection (Table 2.2) so the CPM was able to reproduce these aspects of the observed data fairly well.

Table 2.2: Posterior Predictive p-values for IL-6 Model

Test quantity	Posterior predictive p-value
variance	0.43
skewness	0.34
proportion censored	0.53

The median posterior estimates of the covariate parameters along with 50% and 95% credible intervals for the IL-6 model are shown in Figure 2.11a. Age and BMI were positively associated with increased IL-6, while CD4 count, male gender, and the Lipoatrophy and Neuropathy cohort were negatively associated with IL-6. The relationship between IL-6 and smoking and nonwhite race was more equivocal. Figure 2.11b shows the posterior median γ estimates along with the 50% and 95% credible intervals. Plotting the γ estimates against the observed IL-6 values (Figure 2.12) gives the estimated transformation, $\hat{H}$.

The estimated relationship between BMI and the posterior conditional mean (using 0 for censored values), median, and 90th percentile of IL-6 (for a white, male, nonsmoker with average age and CD4 count in the LiNC study) is shown in Figure 2.13 along with 95% credible intervals. Higher BMI was associated with higher IL-6.

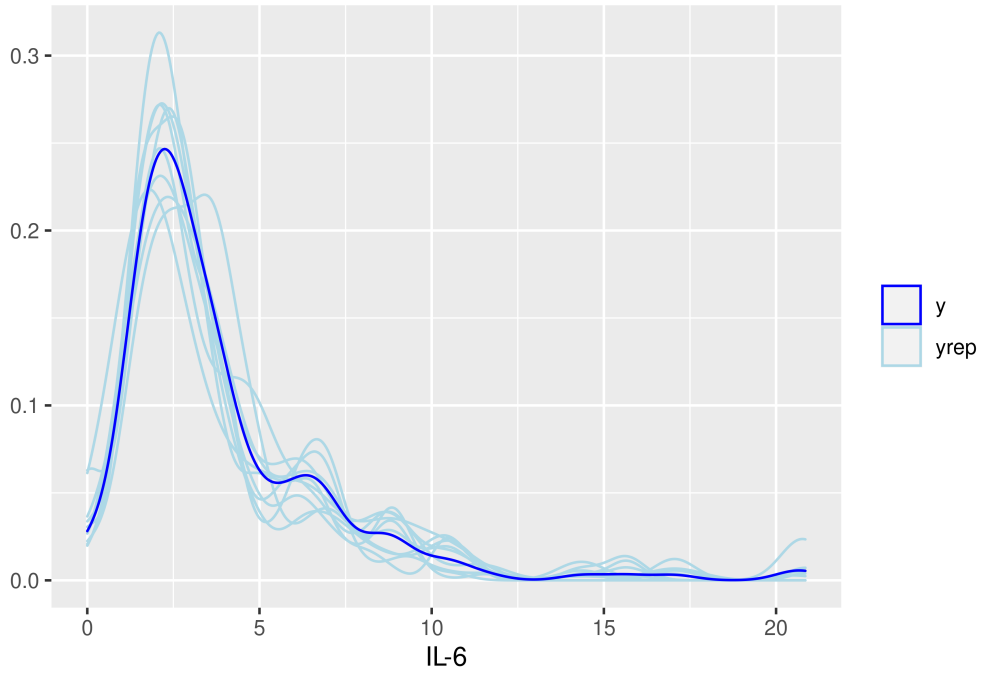

Figure 2.10: Observed outcome (y) and 10 posterior predictive distribution draws (y_{rep}) for IL-6 model

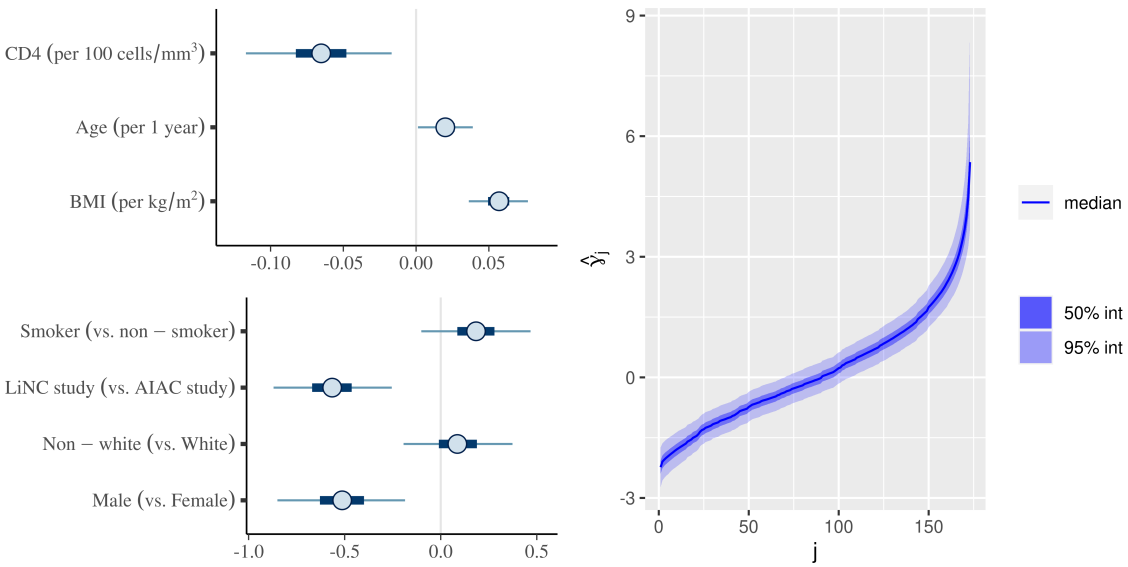

Figure 2.11: (a) Posterior median β estimates and (b) posterior median γ estimates with 50% and 95% credible intervals for IL-6 model

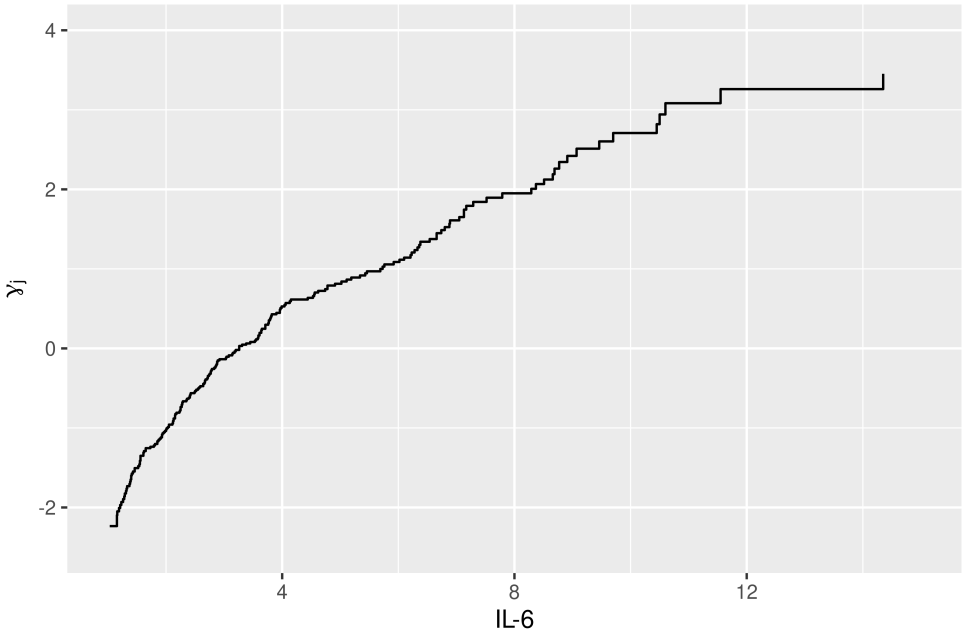

Figure 2.12: Estimated transformation for IL-6 model

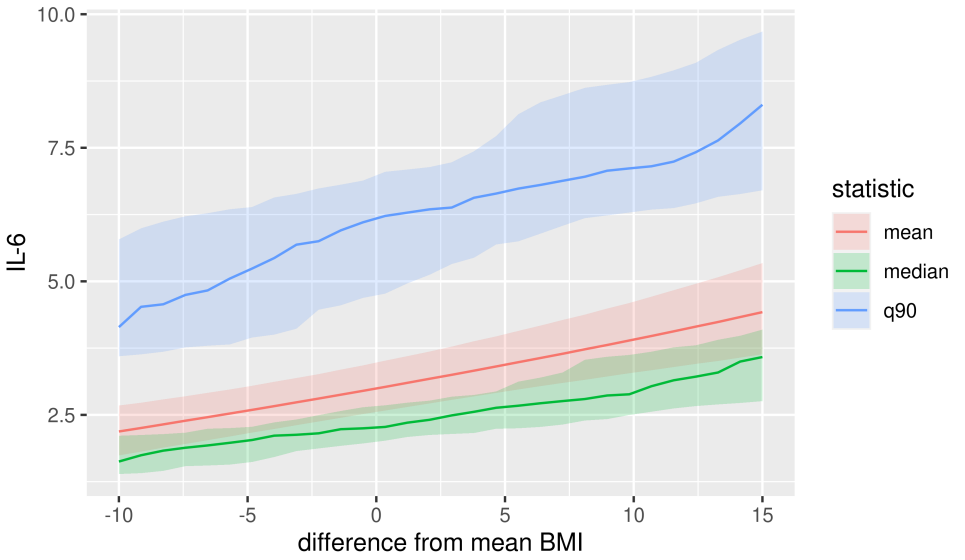

Figure 2.13: Difference from mean BMI vs. IL-6 mean, median, and 90th percentile for a white, male, nonsmoker with average age and CD4 count in the Lipoatrophy and Neuropathy cohort

2.4.2.2 IL-1-β biomarker

As with the IL-6 biomarker, comparing the observed IL-1-β distribution to draws from the posterior predictive distribution (Figure 2.14) did not reveal any serious model misfit. The posterior predictive p-values for variance, skewness, and proportion of observations below the lower limit of detection are shown in Table 2.3. There was no indication of serious discrepancy between the model and data for variance and proportion of censored observations although the posterior predictive p-value for skewness was more extreme indicating a moderate degree of misfit. This seems reasonable given the high level of right-skewness for IL-1-β.

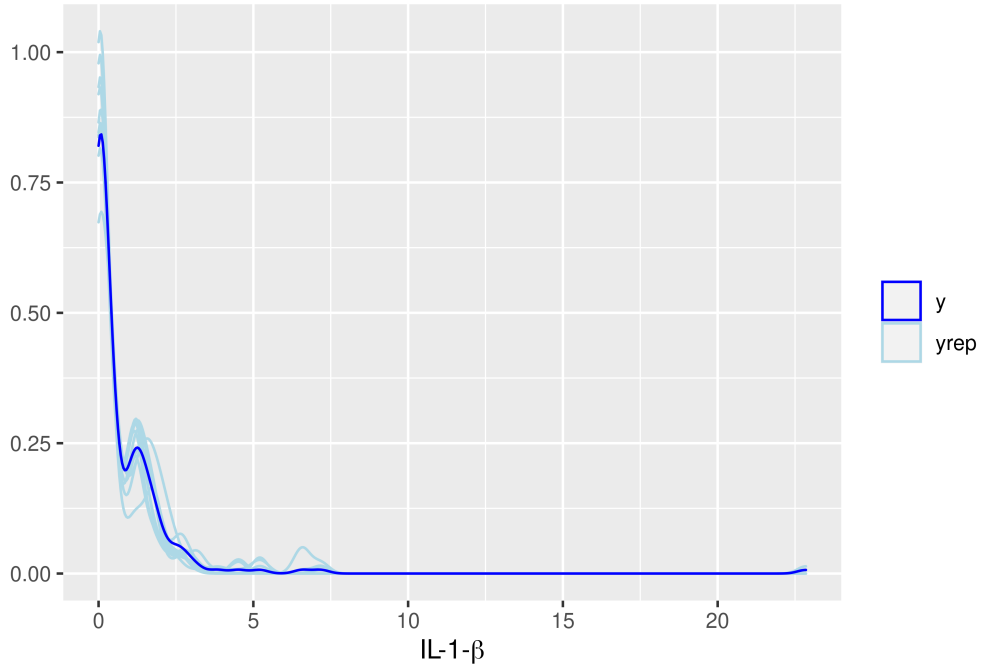

Figure 2.14: Observed outcome (y) and 10 posterior predictive distribution draws (y_{rep}) for IL-1-β model

Table 2.3: Posterior Predictive p-values for IL-1-β Model

Test quantity	Posterior predictive p-value
variance	0.38
skewness	0.15
proportion censored	0.63

The median posterior estimates of the covariate parameters along with 50% and 95% credible intervals for the IL-1-β model are shown in Figure 2.15a. In contrast to IL-6, there was weak association between all covariates (except study cohort) and IL-1-β level. Figure 2.15b shows the posterior median γ estimates along with the 50% and 95% credible intervals. Plotting the γ estimates against the observed IL-1-β values gives the estimated transformation, $\hat{H}$ (Figure 2.16).

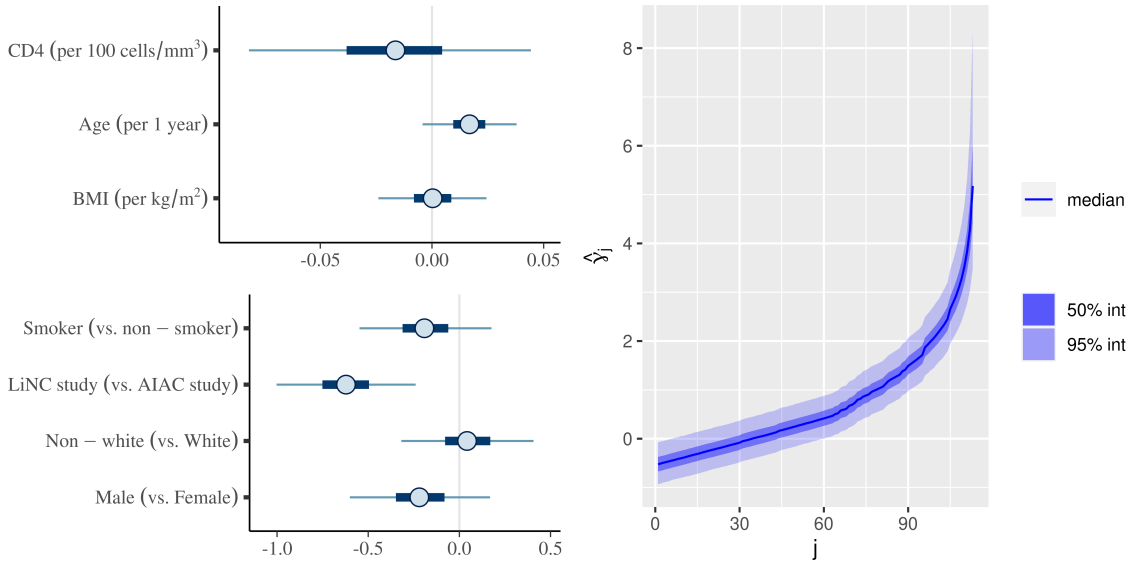

Figure 2.15: (a) Posterior median β estimates and (b) posterior median γ estimates with 50% and 95% credible intervals for IL-1-β model

Figure 2.17 displays the estimated relationship between BMI and the posterior conditional mean (plugging in 0 for censored values), median, and 90th percentile of IL-1-β (for a white, male, nonsmoker with average age and CD4 count in the Lipoatrophy and Neuropathy cohort) along with 95% credible intervals. The plot confirms little association between BMI and IL-1-β.

2.5 Discussion

Although Bayesian CPM models have been frequently applied to ordinal data when the number of outcome categories is much smaller than the sample size, the extension to continuous or mixed outcomes where the number of categories is close or equal to the sample size can be accomplished with only a few modifications to the prior specification. These modifications provide a versatile model with several advantages including the ability to handle both continuous and discrete ordered outcomes and estimation of the full conditional CDF, along with quantiles and other functionals using a single model fit. Inference is based on posterior probability statements and does not require asymptotic assumptions. In addition, the CPM does not require specification of a transformation to meet distributional assumptions since the transformation is estimated nonparametrically. As a result its parameter estimates are invariant to monotonic transformations of the data.

Our implementation of a Bayesian CPM performed reasonably well for the simple simulation scenarios considered. However, the model can produce biased estimates for quantiles far from the median and conditional quantities further from the model where $X = 0$ and this bias can be exacerbated by censoring. The model seems best suited for cases when the data are fairly dense and are sufficient to describe the posterior CDF well. In our

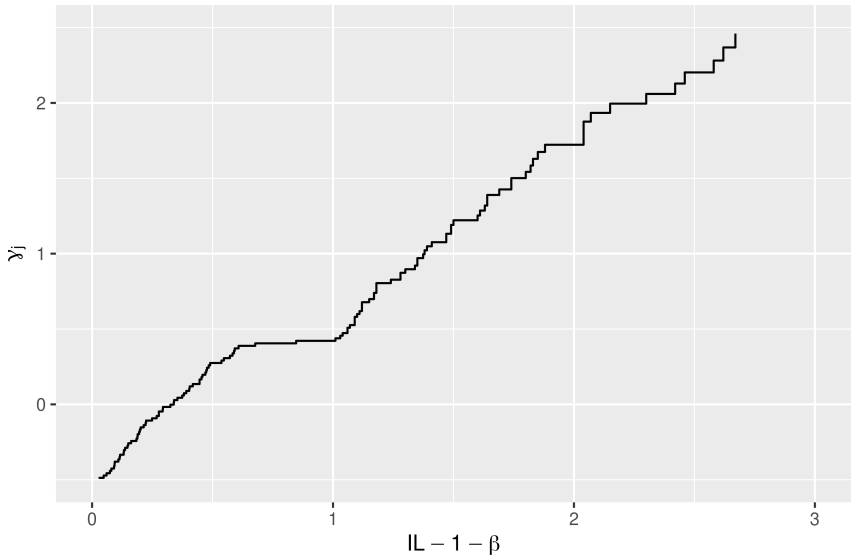

Figure 2.16: Estimated transformation for IL-1-β model

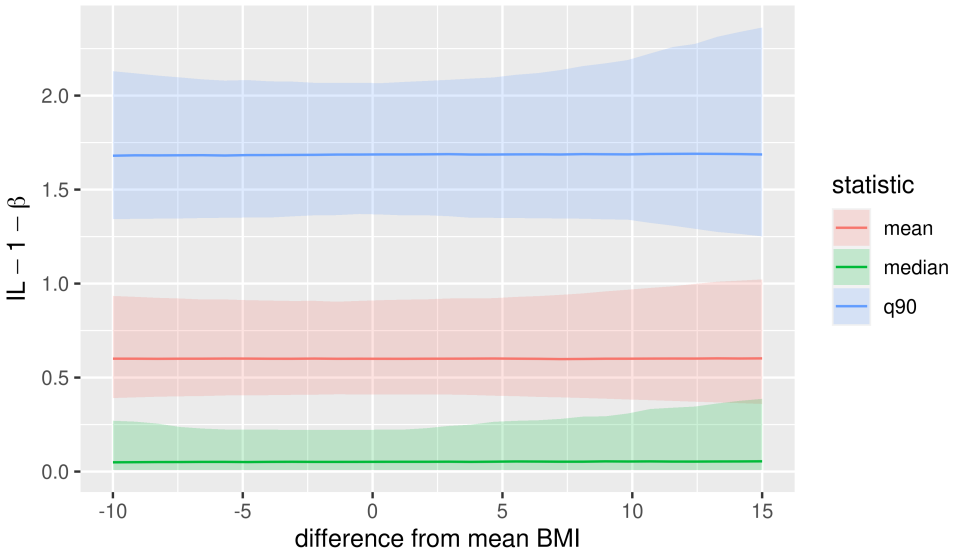

Figure 2.17: Difference from mean BMI vs. IL-1-β mean, median, 90th percentile for a white, male, nonsmoker with average age and CD4 count in the Lipoatrophy and Neuropathy cohort

simulations, a sample size of 50 or 100 was required for reasonably unbiased estimates of parameters and other posterior quantities. The choice of Dirichlet prior concentration with magnitude $\alpha \approx \frac{1}{J}$ has minimal impact on the bias of posterior estimates, except with small sample sizes. Much larger concentration parameters (e.g., $\alpha = 1/2$) may be too informative. As with all Bayesian models estimated with MCMC, checks of model convergence, model fit, and the posterior distribution are important. This is especially true when modeling a mixed continuous/discrete or when interest lies in quantities conditional on covariates far from the observed mean values.

Finally, there are several of limitations of the current model that present an opportunity for improvement. First, the number of distinct outcome values is assumed to be known *a priori*, that is we condition on J categories. In practice, the number of distinct continuous outcome values is unlikely to be available before data collection, so the prior cannot be specified without reference to the observed data. Relatedly, because the number of categories is fixed, the model cannot accommodate new observations for an unobserved category; once the initial prior is set, there is no way to add categories and all predictions are assumed to fall into one of the original categories. It may be possible to overcome this limitation by substituting the Dirichlet prior for a infinite-dimensional Bayesian nonparametric analog, such as a Dirichlet process prior, at the expense of additional complexity and computation time. Next, the choice of link function, and the implied error distribution on the scale of the latent untransformed data is also assumed to be known. If primary interest is not inference for the parameters, specification of the link could be avoided by either estimating the link nonparametrically, although other assumptions may be required for identifiability [33–35], or using a more flexible mixture link function [23].

2.6 Appendix

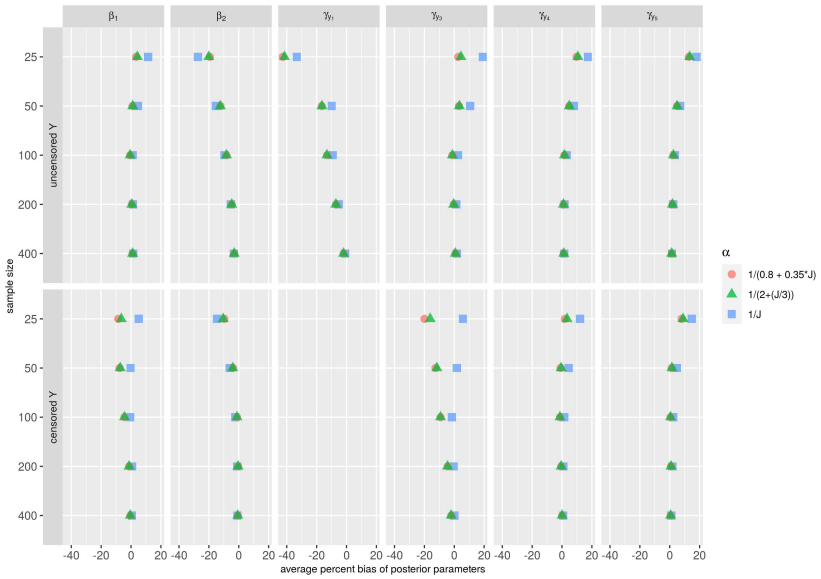

Figure 2.18: Bias in parameters for simulations using loglog link

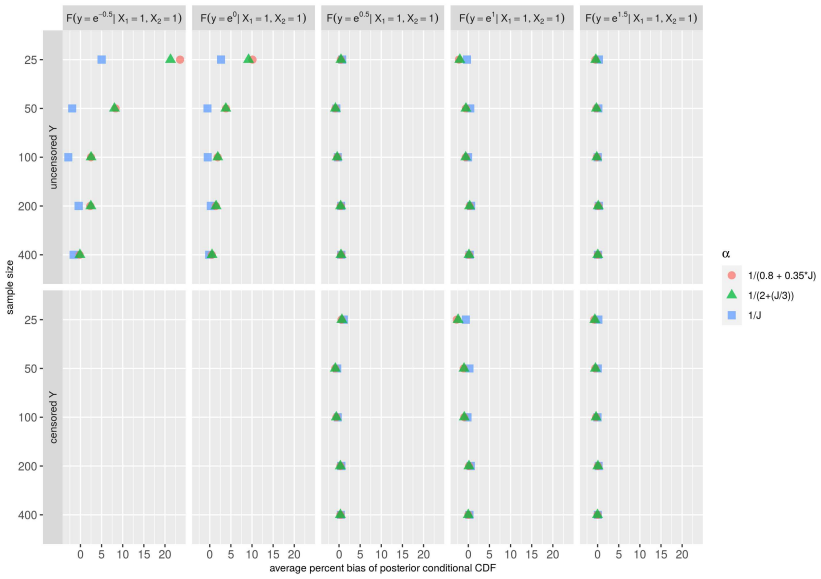

Figure 2.19: Percent bias in conditional CDF for simulations using logit link

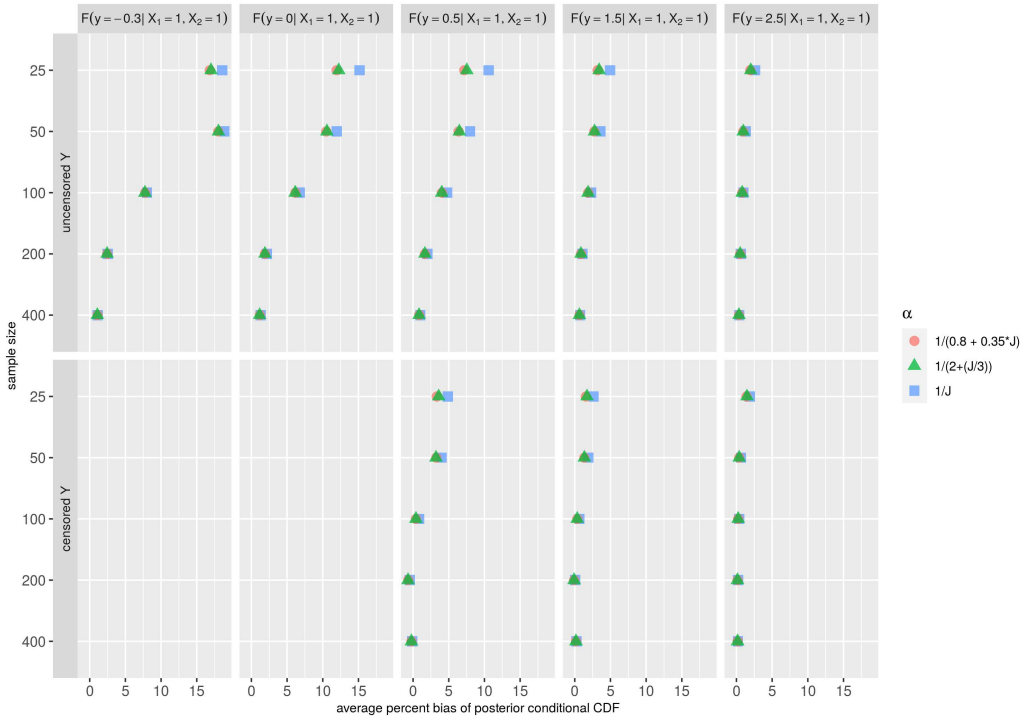

Figure 2.20: Percent bias in conditional CDF for simulations using loglog link

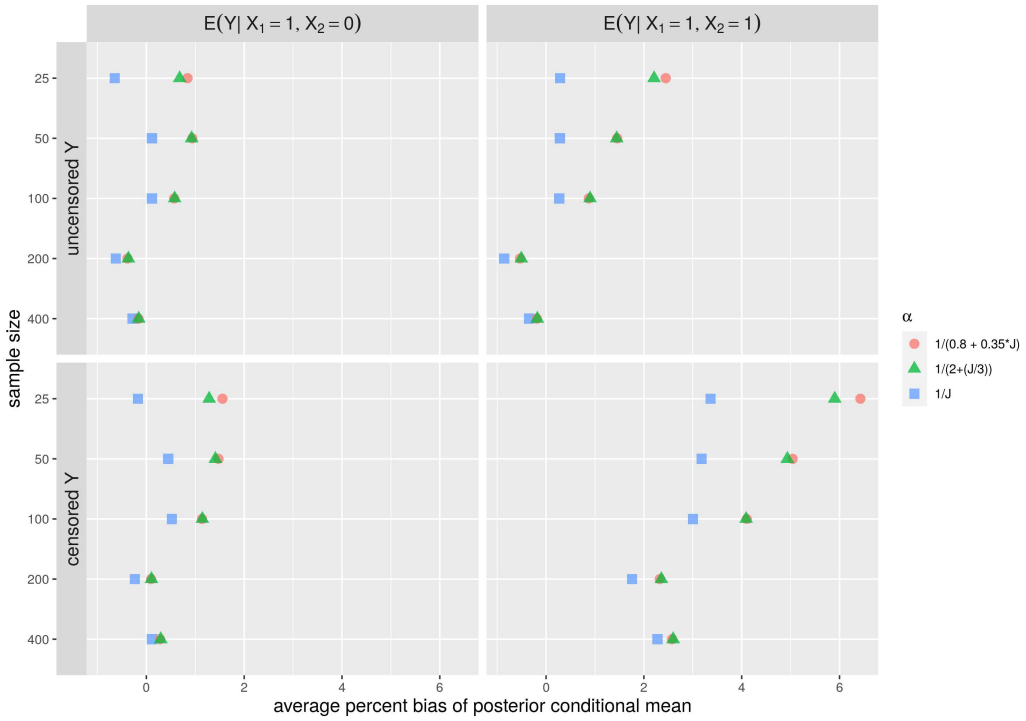

Figure 2.21: Percent bias in conditional mean for simulations using logit link

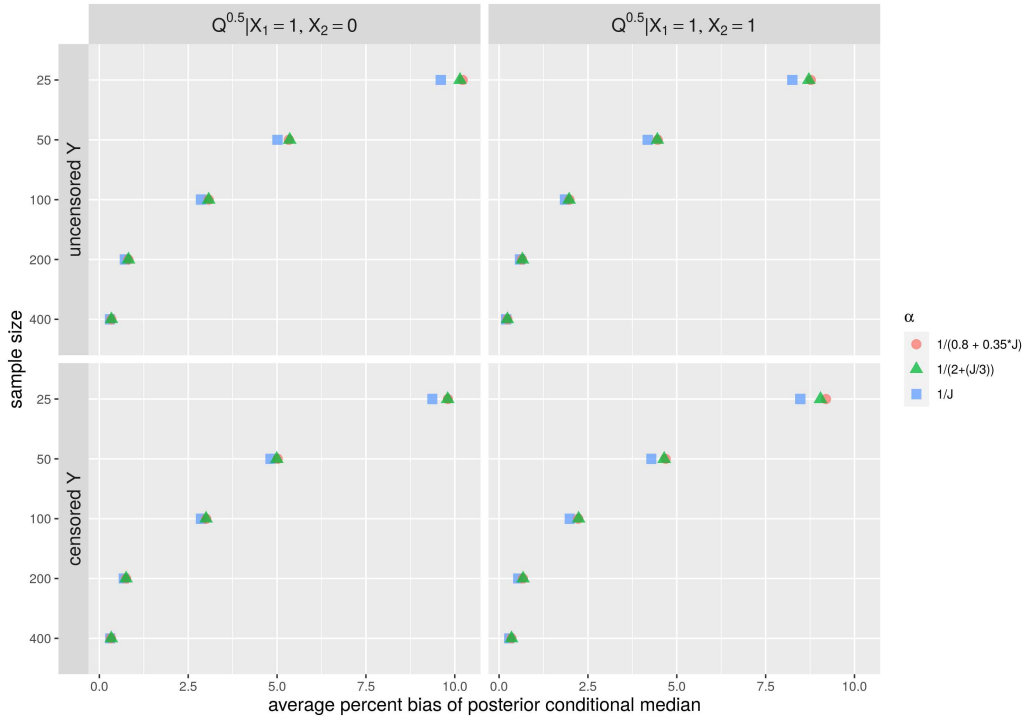

Figure 2.22: Percent bias in conditional median for simulations using logit link

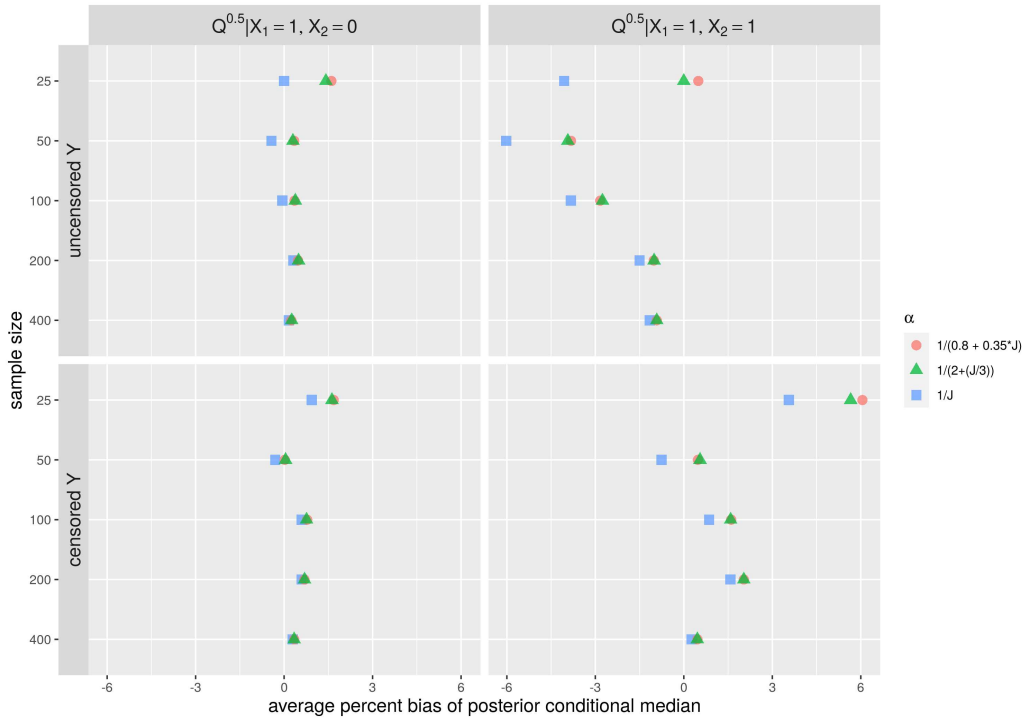

Figure 2.23: Percent bias in conditional median for simulations using loglog link

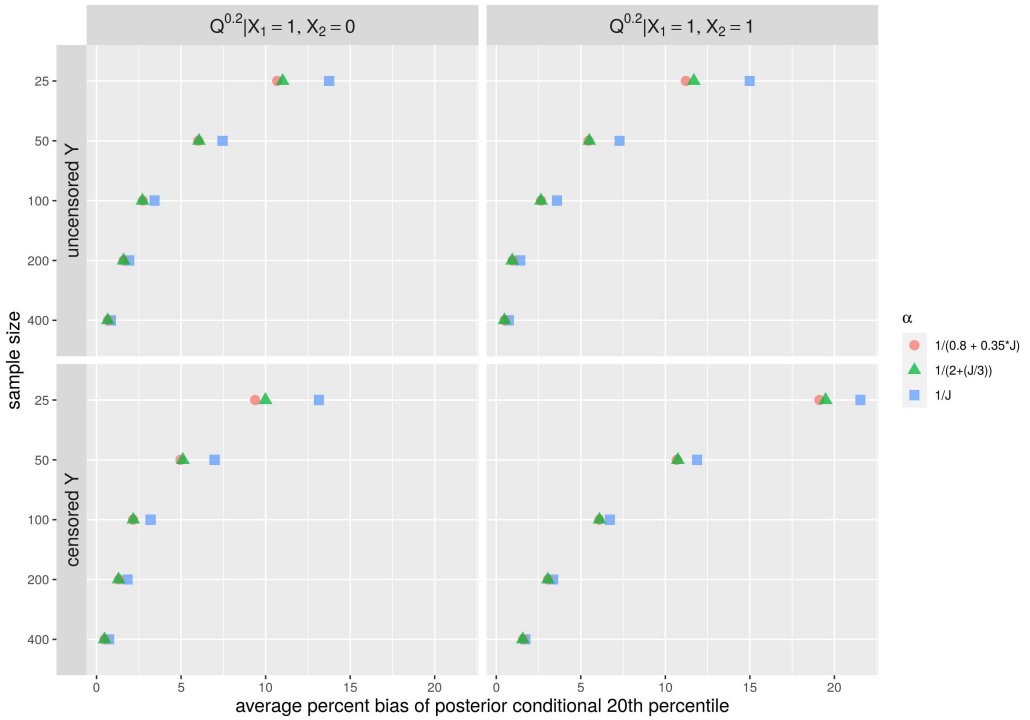

Figure 2.24: Bias in conditional 20th percentile for simulations using logit link

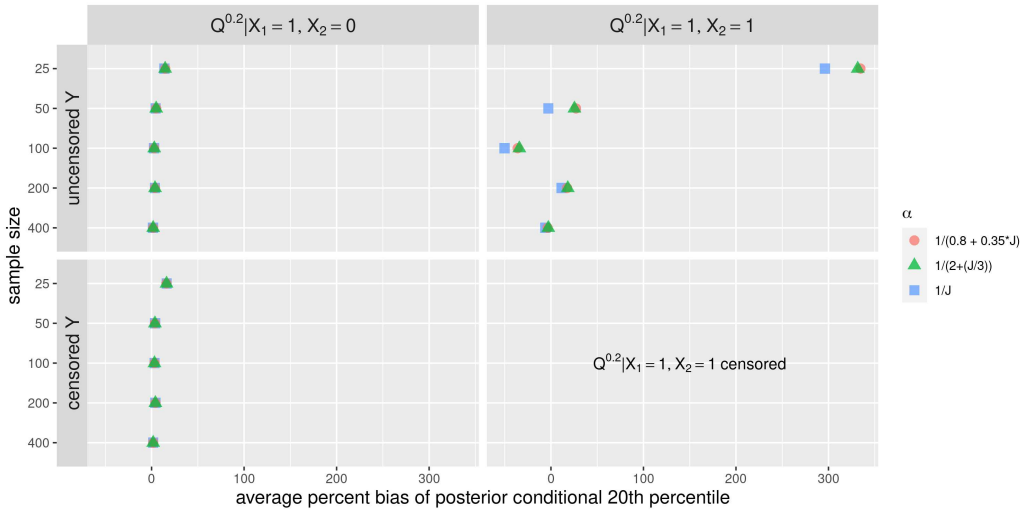

Figure 2.25: Bias in conditional 20th percentile for simulations using loglog link

CHAPTER 3

Bayesian Population Pharmacokinetic Modeling Using Automatic Differentiation Variational Inference

3.1 Introduction

Population pharmacokinetic (PK) models are used to study the variability in different individuals' PK profiles – the amount of drug in the body over time – and how the PK profile is affected by covariates such as weight, age, and organ function. Historically, classical 'frequentist' methods were used to perform estimation and inference for these models. However, fully Bayesian population PK models, first described in the mid-1990s by Wakefield (1994) and others [56–61], have several advantages. These include the ability to incorporate prior information when it is available; accounting for uncertainty in a principled manner using an explicitly defined, flexible hierarchical model structure; straightforward computation of predictive distributions and functionals; and interpretation in terms of posterior probabilities.

Posterior distributions for Bayesian PK model parameters are usually obtained using Markov chain Monte Carlo (MCMC) sampling methods such as Gibbs sampling, the Metropolis-Hastings algorithm [63], and more recently Hamiltonian Monte Carlo (HMC) [65,66]. While MCMC remains the gold standard for Bayesian analyses – providing numerical approximations to the exact posterior distribution – it can be slow to converge and scale poorly with larger sample sizes, complex models, or high-dimensional parameter spaces.

Slow convergence equates to longer computing time, which creates a bottleneck in the iterative model development process of fitting, checking, adjusting, and refitting models [4,55,67,68]. Therefore, it is of interest to explore alternative strategies for Bayesian inference which balance speed and accuracy. In this paper, we explore Bayesian population PK model estimation using variational inference (VI), which uses optimization to find the distribution that is closest to the exact posterior within a specified variational family. As stated in Blei et al. (2017) [15], "Compared to MCMC, variational inference tends to be faster and easier to scale to large data," and is useful in "scenarios where we want to quickly explore many models." Thus, VI provides a potential solution to the computational bottleneck in Bayesian PK model development.

We first describe the specific type of VI used in this paper – mean-field automatic differentiation variational inference (ADVI). Using simulation, we compare HMC and ADVI for population and individual parameter estimation, fit time, posterior prediction, and model selection. We present a real-world PK case study of the analgesic ketolorac in a pediatric population using sparse data derived from electronic health records (EHRs)

and remnant specimens. We conclude by summarizing advantages and limitations of the ADVI approach and providing recommendations for using ADVI for Bayesian PK modeling.

3.2 Background

3.2.1 Automatic Differentiation Variational Inference

The basic principle behind ADVI is to transform all parameters to an unconstrained real coordinate space, then use a mean-field (or full-rank) Gaussian variational family in the real coordinate space. A second elliptical standardization transformation converts to standard Gaussian. Finally, stochastic gradient ascent optimization is performed on the transformed variational objective function using automatic differentiation to compute gradients and Monte Carlo integration to calculate expectations. Beginning with the evidence lower bound (ELBO) we have:

$$\begin{aligned} ELBO = \mathscr{Q}(\phi) &= E_q[\log p(z|\theta)] - KL(q(\theta;\phi)||p(\theta)) = E_q[\log p(z|\theta)] - E_q\left[\log \frac{q(\theta;\phi)}{p(\theta)}\right] \\ &= E_q\left[\log p(z|\theta) - \log \frac{q(\theta;\phi)}{p(\theta)}\right] = E_q\left[\log \frac{p(z|\theta)p(\theta)}{q(\theta;\phi)}\right] \\ &= E_q[\log p(z,\theta)] - E_q[\log q(\theta;\phi)] \end{aligned}$$

The transformation $T(\theta) = \zeta$ converts the parameters θ to an unconstrained real coordinate space and a mean-field Gaussian variational family is used for all transformed parameters $q(\zeta;\phi) = N(\zeta|\phi = \{\mu, diag(\sigma^2)\}) = \prod_j N(\zeta_j|\mu_j,\sigma_j^2) = \prod_j N(\zeta_j|\mu_j, \exp(\omega_j)^2)$. The ELBO becomes:

$$\mathscr{Q}(\phi) = E_{q(\zeta;\phi)}[\log p(z, T^{-1}(\zeta)) + \log|\det J_{T^{-1}}(\zeta)|] - E_{q(\zeta;\phi)}[q(\zeta;\phi)]$$

Stochastic optimization requires the gradients of $\mathscr{Q}(\phi)$ with respect to the variational parameters μ and ω. Elliptical standardization allows us to push the gradient inside the expectation. For mean-field ADVI, $S_\phi(\zeta) = \eta = diag(\exp(\omega))^{-1}(\zeta - \mu)$ so $q(\eta) = N(\eta|0,I) = \prod_j N(\eta_j|0,1)$ and the ELBO is:

$$\mathscr{Q}(\phi) = E_{q(\eta)}[\log p(z, T^{-1}(S_\phi^{-1}(\eta))) + \log|\det J_{T^{-1}}(S_\phi^{-1}(\eta))|] - E_{q(\zeta;\phi)}[q(\zeta;\phi)]$$

Applying the chain rule, the required gradients are:

$$\begin{aligned} \nabla_\mu \mathscr{Q}(\phi) &= E_{N(\eta)}[\nabla_\theta \log p(z,\theta)\nabla_\zeta T^{-1}(\zeta) + \nabla_\zeta \log|\det J_{T^{-1}}(\zeta)|] \\ \nabla_\omega \mathscr{Q}(\phi) &= E_{N(\eta)}[(\nabla_\theta \log p(z,\theta)\nabla_\zeta T^{-1}(\zeta) + \nabla_\zeta \log|\det J_{T^{-1}}(\zeta)|)\eta\ diag(\exp(\omega))] + 1 \end{aligned}$$

The gradients can be calculated using automatic differentiation and the expectations can be approximated using Monte Carlo integration ($E_{q(\eta)}[f(\eta)] = \int f(\eta)q(\eta)d\eta \approx \frac{1}{S}\sum_{s=1}^{S} f(\eta_s)$ where $\eta_s \sim q(\eta)$). Finally we iteratively update the variational parameters until the change in ELBO is below a given tolerance threshold

$$\mu^{(i+1)} \leftarrow \mu^{(i)} + diag(\rho^{(i)})\nabla_{\mu}\mathscr{Q}(\phi)$$
$$\omega^{(i+1)} \leftarrow \omega^{(i)} + diag(\rho^{(i)})\nabla_{\omega}\mathscr{Q}(\phi)$$

where ρ is a step-size sequence tuned to achieve the fastest convergence. Additional details on the ADVI algorithm and its implementation in `Stan` can be found in Kucukelbir et al. (2015; 2017) [14,78] and Turkman et al. (2019) [7].

3.2.2 Model Diagnostics

VI has been shown to provide reasonable results for many problems [14,79–84]. But unlike MCMC, it lacks theoretical guarantees related to the quality of the approximation. We use several strategies to assess the accuracy of VI, focusing on the ADVI approximation for population PK models. First, we design a simulation study to compare ADVI and MCMC estimation for population and individual parameters, computation time, and predicted concentration. We examine percent bias for posterior median, 90% posterior credible interval widths, the ratio of MCMC to ADVI fit time, and the ratio of posterior predicted C_{max} concentrations. We also explore two VI diagnostics developed by Yao et al. (2018), the Pareto Smoothed Importance Sampling (PSIS) $\hat{k}$ value and variational simulation based calibration (VSBC) [85]. The $\hat{k}$ diagnostic is a single value that measures how well the variational approximation matches the full joint posterior; VSBC constructs histograms for individual parameters which can be used to assess average bias. Vehtari et al. (2021) found that values of $\hat{k} > 0.7$ indicate a discrepancy between the VI approximation and the full posterior distribution [86]. However, Yao and colleagues note, "while the VI posterior may be a poor approximation to the full posterior, point estimates that are derived from it may still have good statistical properties." Unlike the global PSIS $\hat{k}$ statistic, the VSBC diagnostic can be used to assess the average calibration of individual parameter point estimates over all datasets that could be constructed from a given model. Asymmetry in the VSBC histogram for a parameter indicates bias in the VI approximation to the marginal posterior. The authors caution that the results should be interpreted conservatively; failure of the diagnostic indicates the VI approximation may perform poorly with observed data, but passing the diagnostic does not guarantee good performance. The Appendix provides more details on both diagnostics. Finally, in the ketolorac case study, we compare VI and MCMC using Bayesian posterior predictive checks for training and validation data and population PK diagnostic plots of observed vs. predicted values, random effects vs. covariates, and visual predictive check (VPC) plots.

3.3 Simulation Study

3.3.1 Methods

3.3.1.1 Set-up

We examine two simple PK model scenarios: one-compartment with IV infusion and two-compartment with IV infusion. For dosing we specify a 12-hour infusion starting at time $t = 0$ with a rate of 0.7 units/kg/hr for all subjects. Both scenarios include combined additive and proportional residual error, allometric weight scaling factors, random effects and correlations between random effects for the main PK parameters and a binary covariate, X_1, with effect $\beta_{Cl_{X_1}} = 0.4$ on the logarithm of total clearance. Under each scenario, we simulate dense and sparse sampling schemes. For the dense scheme, we sample 1 observation at 0.25 hours, 1 observation every 0.75 hours for 0.75 through 3 hours (4 observations), 1 observation every hour for hours 4 through 8 (5 observations), and 1 observation every 1.5 hours for hours 9 through 24 (11 observations), yielding a total of 21 observations per subject (Figure 3.1). For the sparse sampling scheme, each subject has a varying number of observations: 2, 3, 4, or 5 observations with probabilities 0.3, 0.4, 0.2, and 0.1, respectively. The chosen number of observations are then used to randomly select observations from those simulated using the dense sampling scheme. For each scenario, we simulate 250 datasets with 100 subjects per dataset using dense sampling and 250 datasets with 300 subjects each using sparse sampling. Figures 3.39, 3.40, 3.41, and 3.42 show the first nine simulated datasets for the four combinations of scenario (one or two compartments) and sampling scheme (sparse or dense). Additional details for the simulation scenarios are provided in the Appendix.

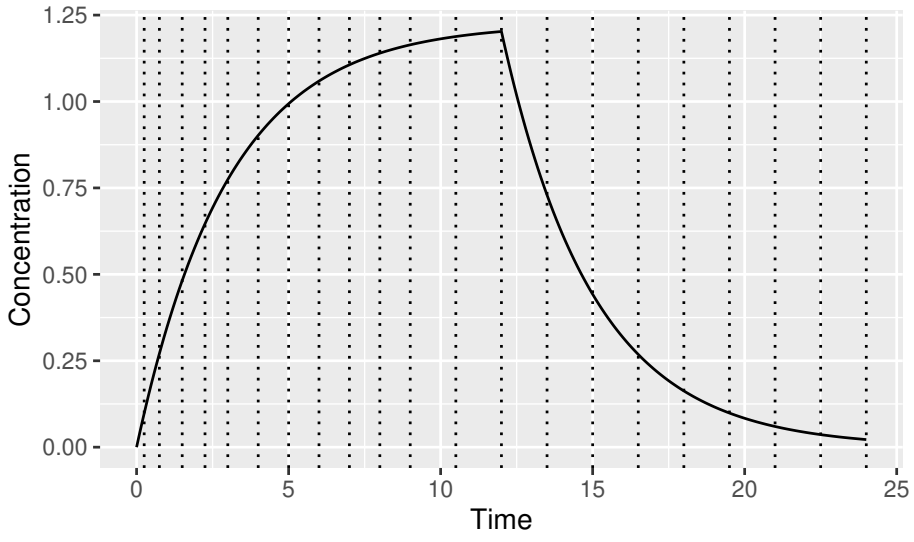

Figure 3.1: One-compartment model with dense sampling scheme

3.3.1.2 Modeling

For each simulated dataset, the correctly specified model (referred to as M1) is fit using both mean-field ADVI and the No-U-Turns (NUTS) HMC sampler. We consider four different priors for each model; the first three are informative with two concentrating their density around the true values for the main population PK parameters, but with different amounts of variability, and a third misspecified prior centered at incorrect values with moderate variability. Plots of the prior densities for the first three priors in scenarios 1 and 2 are shown in Figures 3.2 and 3.3, respectively. The final prior is non-informative. Simulations were performed using `R` version 3.6.0 (2019-04-26) and the `CmdStanR` interface (version 0.4.0) to the `CmdStan` probabilistic programming language (version 2.25.0) along with the PK library `Torsten` (version 0.88) [65,66,87]. Additional details on the priors and the settings used for the ADVI and NUTS HMC algorithms are provided in the Appendix.

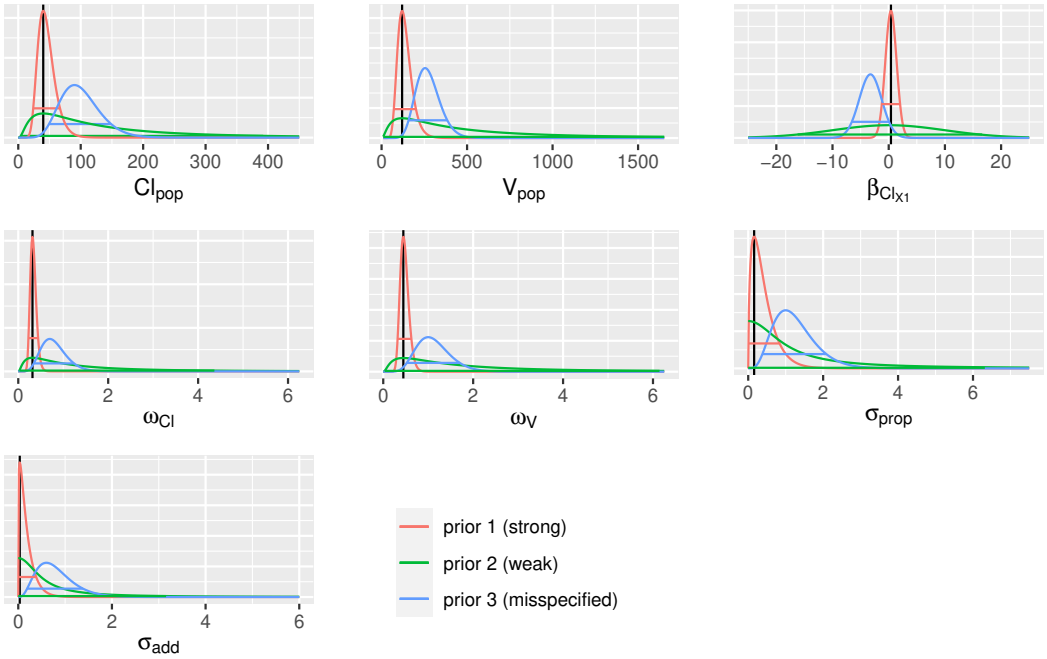

Figure 3.2: Informative priors for one-compartment with IV infusion models (Scenario 1). For Cl_{pop}, V_{pop}, $\beta_{Cl_{X1}}$, ω_{Cl}, and ω_V, the mode of the strong and weak priors are at the true parameter value (indicated with a vertical black line); for σ_{prop} and σ_{add} a Gamma prior with median at the true parameter or Half-Cauchy priors with mode 0 are used. Misspecified priors are either Gamma or Normal centered at incorrect values. Horizontal lines indicate the 90% highest density interval for each prior.

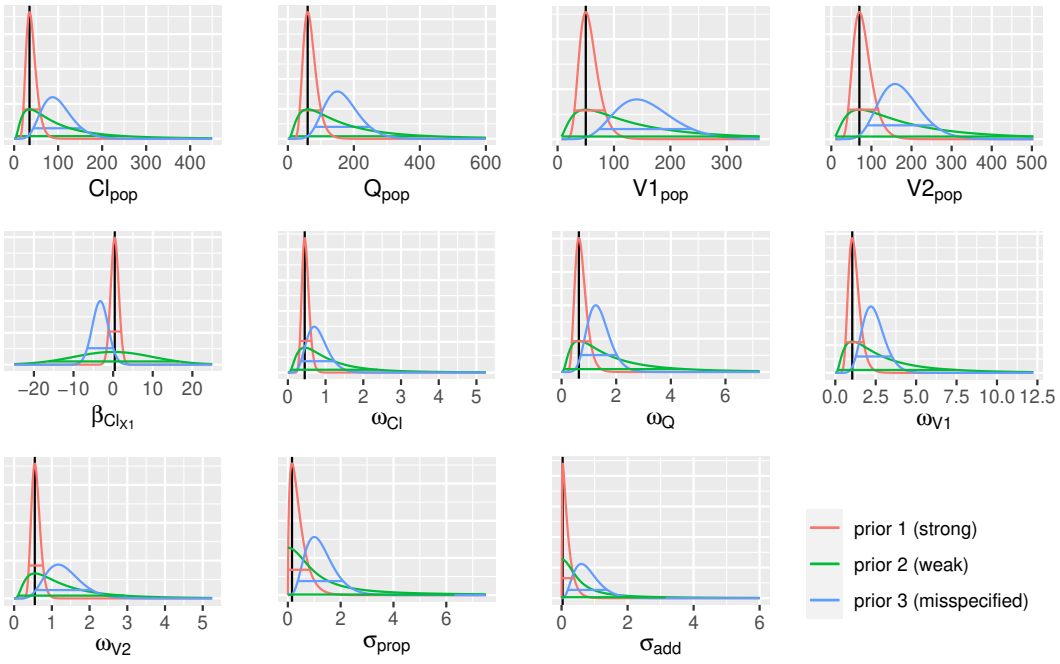

Figure 3.3: Informative priors for two-compartment with IV infusion models (Scenario 2). For Cl_{pop}, Q_{pop}, $V1_{pop}$, $V2_{pop}$, $\beta_{Cl_{X1}}$, ω_{Cl}, and ω_V the mode of the strong and weak priors are at the true parameter value (indicated with a vertical black line); for σ_{prop} and σ_{add} a Gamma prior with median at the true parameter or Half-Cauchy priors with mode 0 are used. Misspecified priors are either Gamma or Normal centered at incorrect values. Horizontal lines indicate the 90% highest density interval for each prior.

3.3.1.3 Model Selection

In addition to exploring the performance of ADVI for a single correctly specified model, we also examine how it performs for misspecified covariate models vis-a-vis model selection. As noted by Wakefield et al. (1999), choosing which components to include for each parameter in stage 2 is "a very difficult problem"; a multiple regression where the dependent variable, ψ, is both multivariate and latent [62].

The clearance component (excluding random effects) of the correctly specified model (M1) with $\beta_{Cl_{X_1}} = 0.4$ is given by: $\log(Cl) = \log(Cl_{pop}) + 0.75 \times \ln\left(\frac{weight}{70}\right) + \beta_{Cl_{X_1}} X_1$. To explore model selection, two misspecified models are also fit to each simulated dataset. For the first misspecified model (M2), the covariate X_1 in model M1 is replaced by the covariate $X_2 \sim Exponential(\lambda = 0.2)$ unrelated to the true data-generating model. The clearance component for M2 is given by: $\log(Cl) = \log(Cl_{pop}) + 0.75 \times \ln\left(\frac{weight}{70}\right) + \beta_{Cl_{X_2}} X_2$. The second misspecified model (M3) includes both X_1 and X_2 and its clearance component is summarized as: $\log(Cl) = \log(Cl_{pop}) + 0.75 \times \ln\left(\frac{weight}{70}\right) + \beta_{Cl_{X_1}} X_1 + \beta_{Cl_{X_2}} X_2$. The same sparse or dense sampling schemes and priors are used for all three models (M1, M2, and M3).

For model selection, each dataset is fit using the correctly specified model (M1) and the two misspecified

models (M2, M3). We use the strategies described below to rank the models. Under each strategy, we calculate how frequently each of the 6 possible rankings – from best (M1,M3,M2) to worst (M2,M3,M1) – are chosen. Three model selection strategies are considered.

- Strategy 1: For both ADVI and MCMC, calculate deviance information criteria (DIC) [88], widely applicable information criteria (WAIC) [76,89], and PSIS leave-one-out information criteria (PSIS-LOOIC) [86]. Models with lower information criteria have better out-of-sample predictive performance.
- Strategy 2: For ADVI, compare the three models using the maximized ELBO. As explained by Blei et al. (2017) [15], the premise for this strategy is that the ELBO is an approximation to the marginal likelihood and can provide a basis for model comparison (marginal likelihoods are also used to construct Bayes factors for model comparison). However, they also note that selection based on a bound is not justified in theory. Penny et al. (2007) [90] discuss this approach as a surrogate for model evidence in the context of neuroimaging. Beal and Ghahramini (2003) [8] show ELBO outperforms BIC in finding the correct model structure for conjugate-exponential models.
- Strategy 3: For ADVI, use k-fold leave-subject-out cross-validation. This strategy leverages the reduced time to find VI approximations to fit multiple models on the data partitioned into training and validation datasets [91]. Specifically, we partition the data by defining K groups each containing approximately $\frac{1}{K}$ of the subjects (not observations). Each of the groups then contains all observations for subjects assigned to that group. For $k = 1, \ldots, K$ folds an estimate of out-of-sample predictive fit is the cross-validation log pointwise predictive density $lppd_{cv} = \sum_{k=1}^{K} \sum_{i \in k} \log p_{post(-k)}(y_i)$ where $p_{post(-k)}$ is the posterior distribution for a model fit *without* the observations in fold k and the inner sum evaluates the log predictive density over the observations y_i in fold k. If $p_{post(-k)}$ is summarized by S simulation draws $\theta^{-k,s}$ then $lppd_{cv}$ can be computed as $\sum_{k=1}^{K} \sum_{i \in k} \log \left(\frac{1}{S} \sum_{s=1}^{S} p(y_i | \theta^{-k,s}) \right)$. Models with higher $lppd_{cv}$ have better out-of-sample predictive performance.

The criteria employed in the three strategies have somewhat different motivations and predictive aims. The information criteria in Strategy 1 are approximations of leave-observation-out cross-validation and ask the predictive question, "How well can the model predict unobserved concentrations for subjects in the current data?" The ELBO used in Strategy 2 is a bound on the marginal likelihood which can be used to estimate relative posterior probabilities for model comparison, but is not meant to directly estimate predictive accuracy [92]. The leave-subject-out $lppd_{cv}$ used in Strategy 3 partitions the data at the subject level and asks the question, "How well can the model predict unobserved concentrations for new unobserved subjects?"

3.3.2 Results

The Bayesian population PK models for each scenario (one or two compartments) were fit using both estimation methods (ADVI and MCMC) with the two sampling schemes (dense and sparse) and four priors (strong, weak, misspecified, and non-informative) described above. We compare posterior median and credible interval width for population parameters, the ratio of MCMC to ADVI total fit time, posterior median and credible interval width for individual parameters, VI diagnostics, and individual and population predicted concentrations.

3.3.2.1 Population Parameter Estimation

We first look at population parameters, which are global values for all individuals in the model. These include the fixed effects (μ) and covariance matrix (Ω) from stage two (Equation 1.7) and the variance (σ_y^2) from stage one (Equation 1.6). Specifically, for the one-compartment scenario we examine two parameters: Cl_{pop}, the population clearance for a subject with 70kg weight and covariate $X_1 = 0$ and V_{pop}, the population volume of distribution for a 70kg subject. For the two-compartment scenario we examine four parameters: Cl_{pop}, the population total clearance for a subject with 70kg weight and covariate $X_1 = 0$; Q_{pop}, the population inter-compartmental clearance for a 70kg subject; and $V1_{pop}$ and $V2_{pop}$, the first (central) and second (peripheral) compartment volume of distribution, respectively, for a 70kg subject. We also examine the corresponding variance and covariance terms (e.g., ω_{Cl}^2, ω_V^2, and $\omega_{Cl,V}$ for the one-compartment scenario) which estimate the inter-individual variability in PK parameters and the variance terms (σ_{add}^2 and σ_{prop}^2) which estimate the intra-individual variability in concentration.

For the one-compartment model scenario, the posterior medians for the parameters Cl_{pop} and V_{pop} were fairly accurate for both estimation methods, with larger percent bias for ADVI (Figure 3.4). The posterior medians for $\beta_{Cl_{X1}}$ also had reasonably small percent bias (half of estimates were within $\pm 12.5\%$) and were similar for ADVI and MCMC. The median point estimates for inter-individual parameter variability (Ω) were more accurate for MCMC than ADVI. Specifically, ADVI underestimates the variance terms (ω_{Cl}^2 and ω_V^2), but overestimates the covariance ($\omega_{Cl,V}$), especially under sparse sampling. For the intra-individual variance, both estimation methods underestimated the additive error variance (σ_{add}^2) and overestimated the proportional error variance (σ_{prop}^2).

In the more complex two-compartment model scenario, posterior medians for population parameters were less consistent (Figure 3.5). For Cl_{pop}, most estimates had bias of less than $\pm 5\%$ across priors, sampling schemes and estimation methods, except for ADVI estimation with dense sampling. Using ADVI, the population inter-compartmental clearance, Q_{pop}, was underestimated with median percent bias between -12.5% and -50% while the volume of distribution for the central compartment, $V1_{pop}$, was overestimated with median

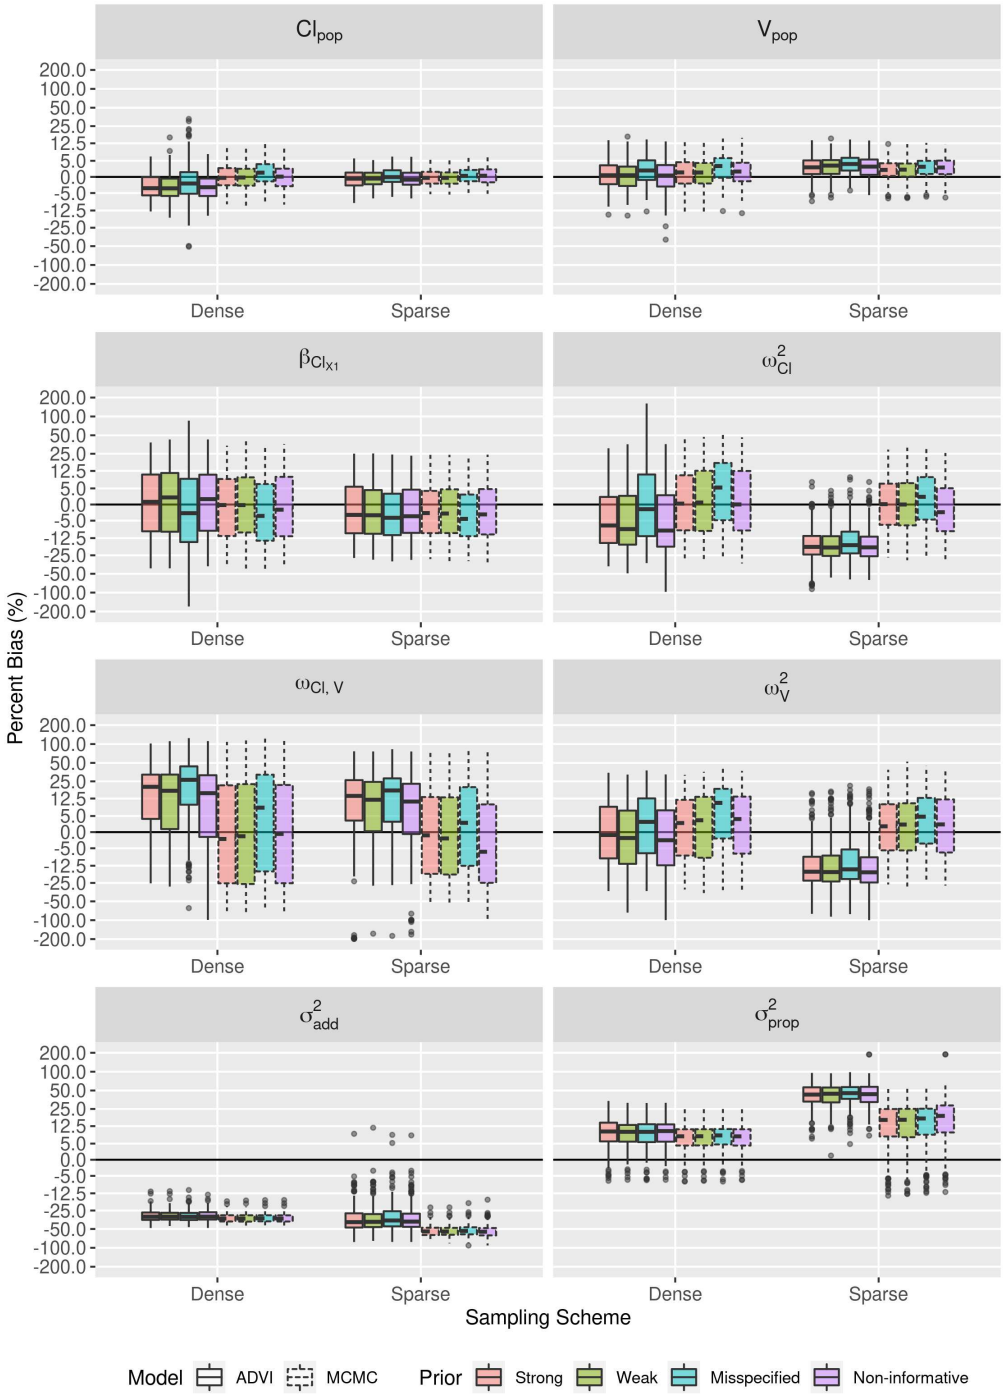

Figure 3.4: One-compartment PK model simulations - posterior medians for population and variance parameters for correctly specified model (M1). Extreme outliers greater than 10 times the interquartile range were removed.

percent bias between 2.5% and 25%. The volume of distribution for the peripheral compartment, $V2_{pop}$, was overestimated for all priors, sampling schemes, and estimation methods, except for ADVI estimation with

dense sampling. The posterior medians for $\beta_{Cl_{X1}}$ were fairly accurate and similar for all combinations of prior, sampling scheme, and estimation method. For most settings, the distribution of ADVI estimates was underestimated compared to MCMC estimates for Q_{pop} and $V2_{pop}$, but overestimated for $V1_{pop}$.

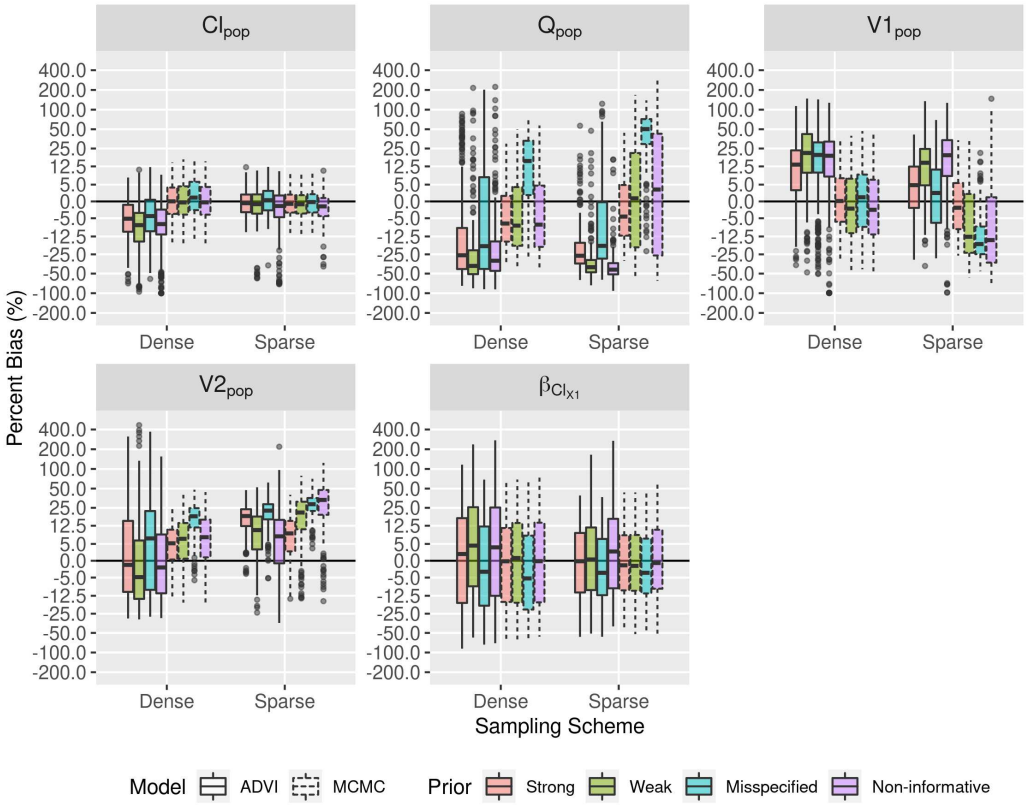

Figure 3.5: Two-compartment PK model simulations - posterior median population parameter estimates for correctly specified model (M1). Extreme outliers greater than 10 times the interquartile range were removed.

The posterior medians of the variance components in the two-compartment scenario have larger percent bias than the population parameters (Figure 3.6). The distributions of these estimates have thicker tails, with the middle 50% of estimates spanning a much wider range, and also show evidence of skewness or multimodality.

In contrast to the median estimates, the 90% equal-tailed credible interval widths were much smaller for ADVI than MCMC for all parameters except σ^2_{add} and σ^2_{prop} (Figures 3.7, 3.8, and 3.9).

Overall, posterior medians for population parameters were more accurate for the MCMC estimation method compared to ADVI, for the simpler one-compartment model compared to the two-compartment model, and for the fixed effect parameters (μ) compared to the variance parameters (Ω,σ^2_y). The uncertainty of these estimates, measured by posterior credible interval widths, was underestimated for ADVI compared to MCMC.

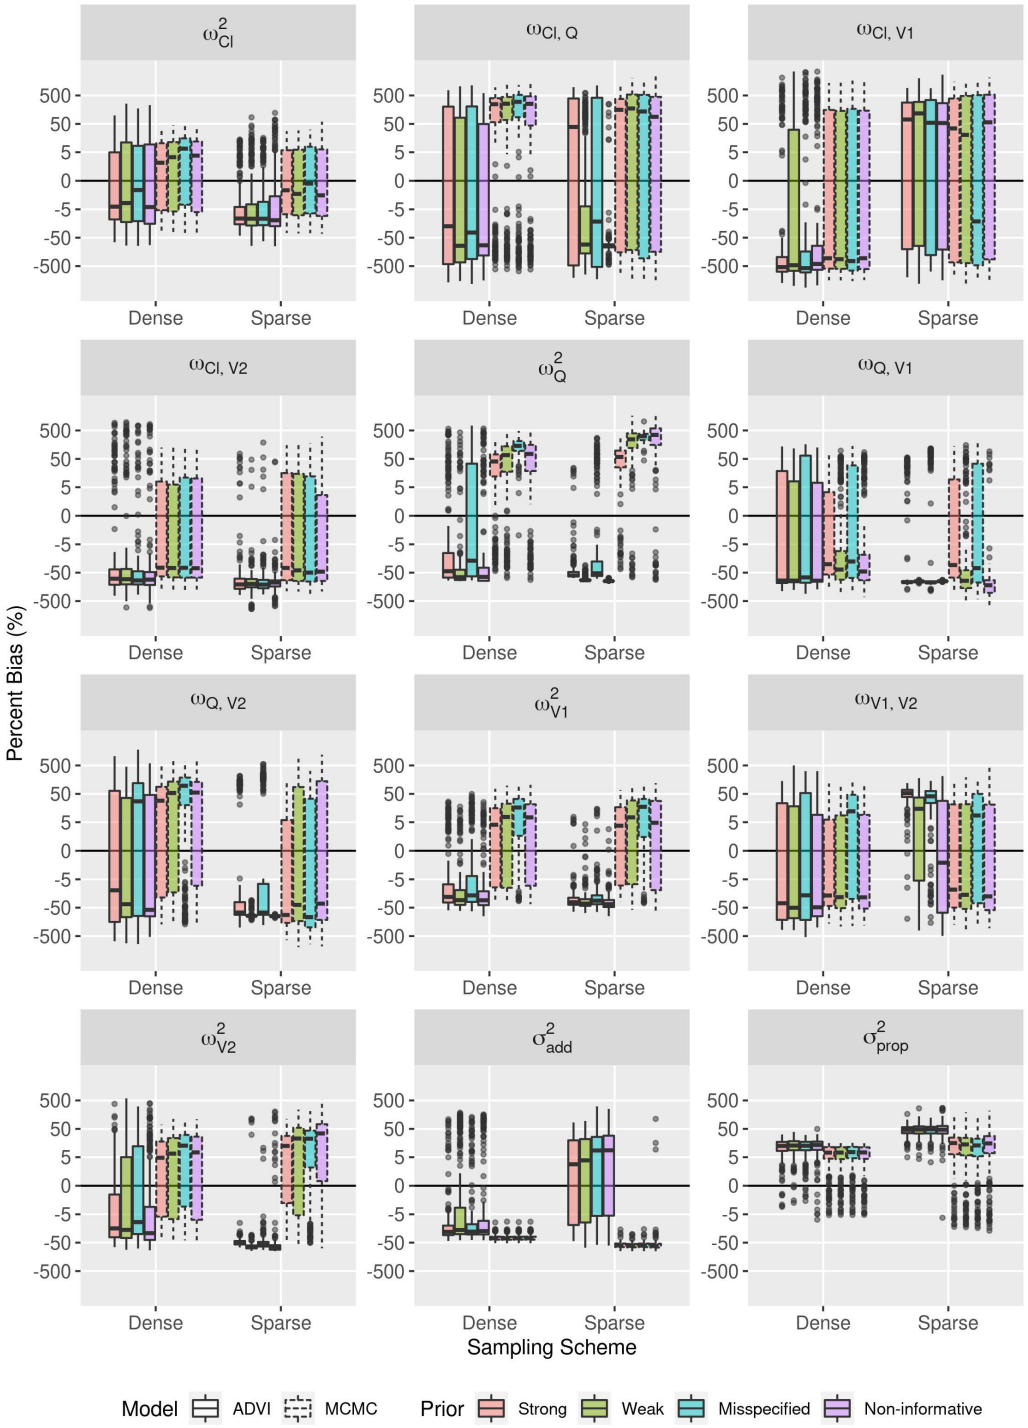

Figure 3.6: Two-compartment PK model simulations - posterior median variance parameter estimates for correctly specified model (M1). Extreme outliers greater than 10 times the interquartile range were removed.

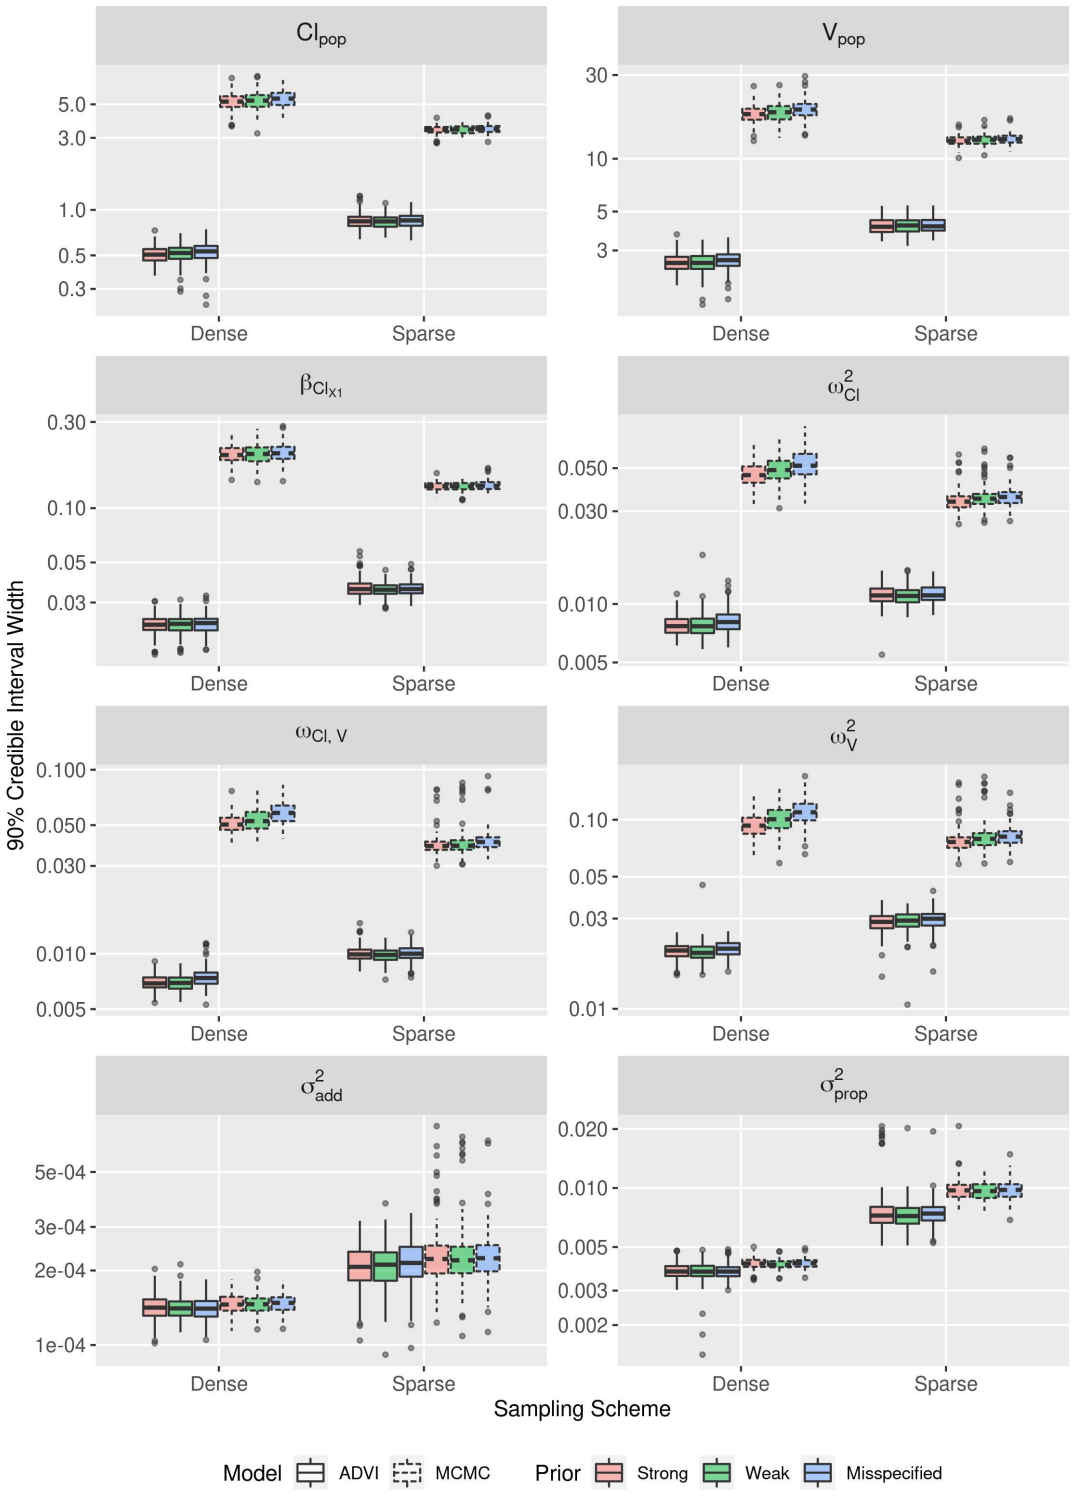

Figure 3.7: One-compartment PK model simulations - 90% equal-tailed credible interval widths for population parameters for correctly specified model (M1) with informative priors. Note: credible interval widths for non-informative prior are not shown due to a large proportion of extreme outliers.

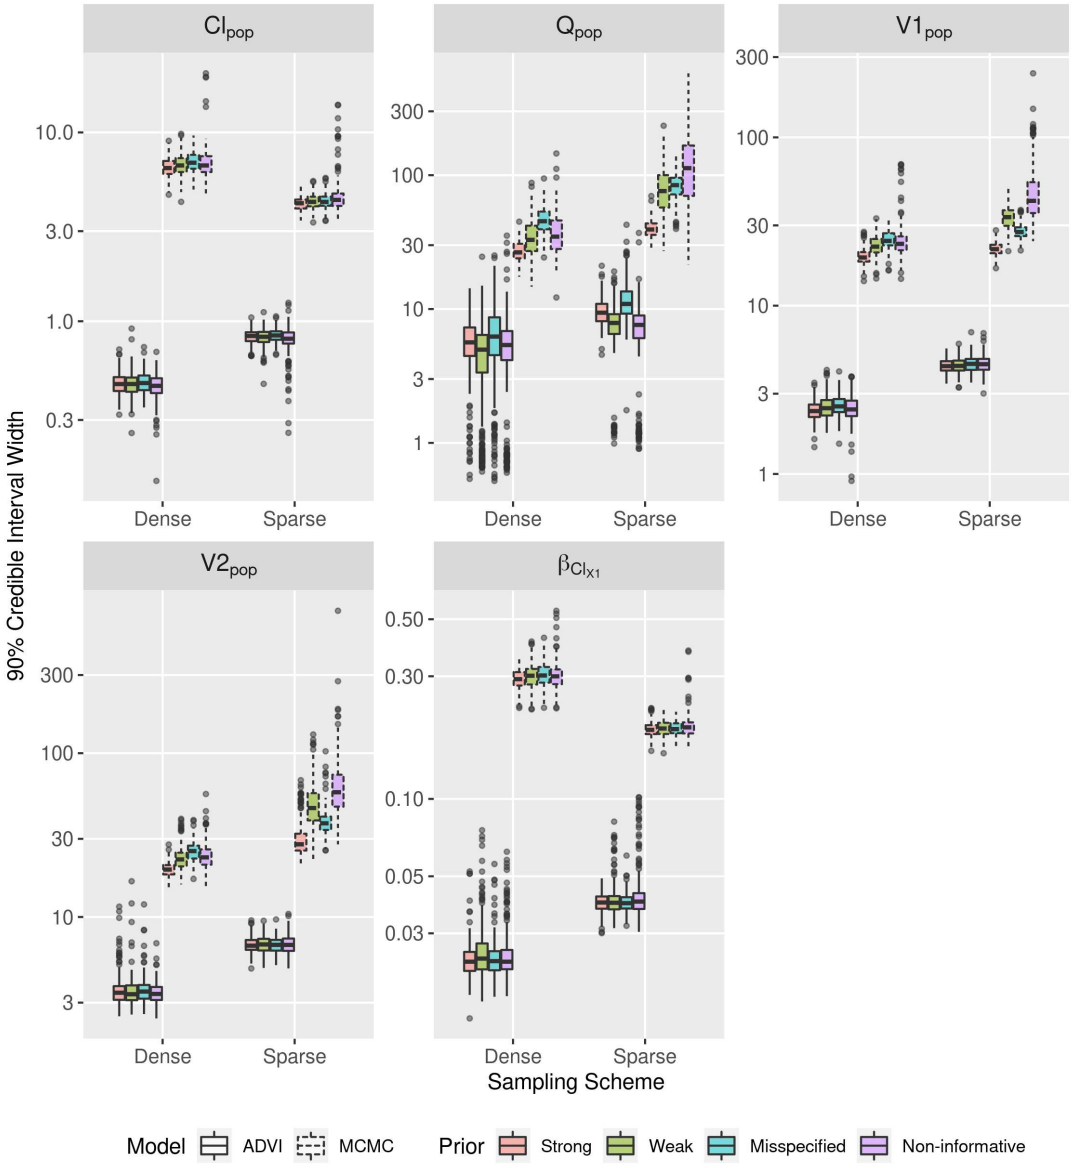

Figure 3.8: Two-compartment PK model simulations - 90% equal-tailed credible interval widths for population parameters for correctly specified model (M1) with informative priors.

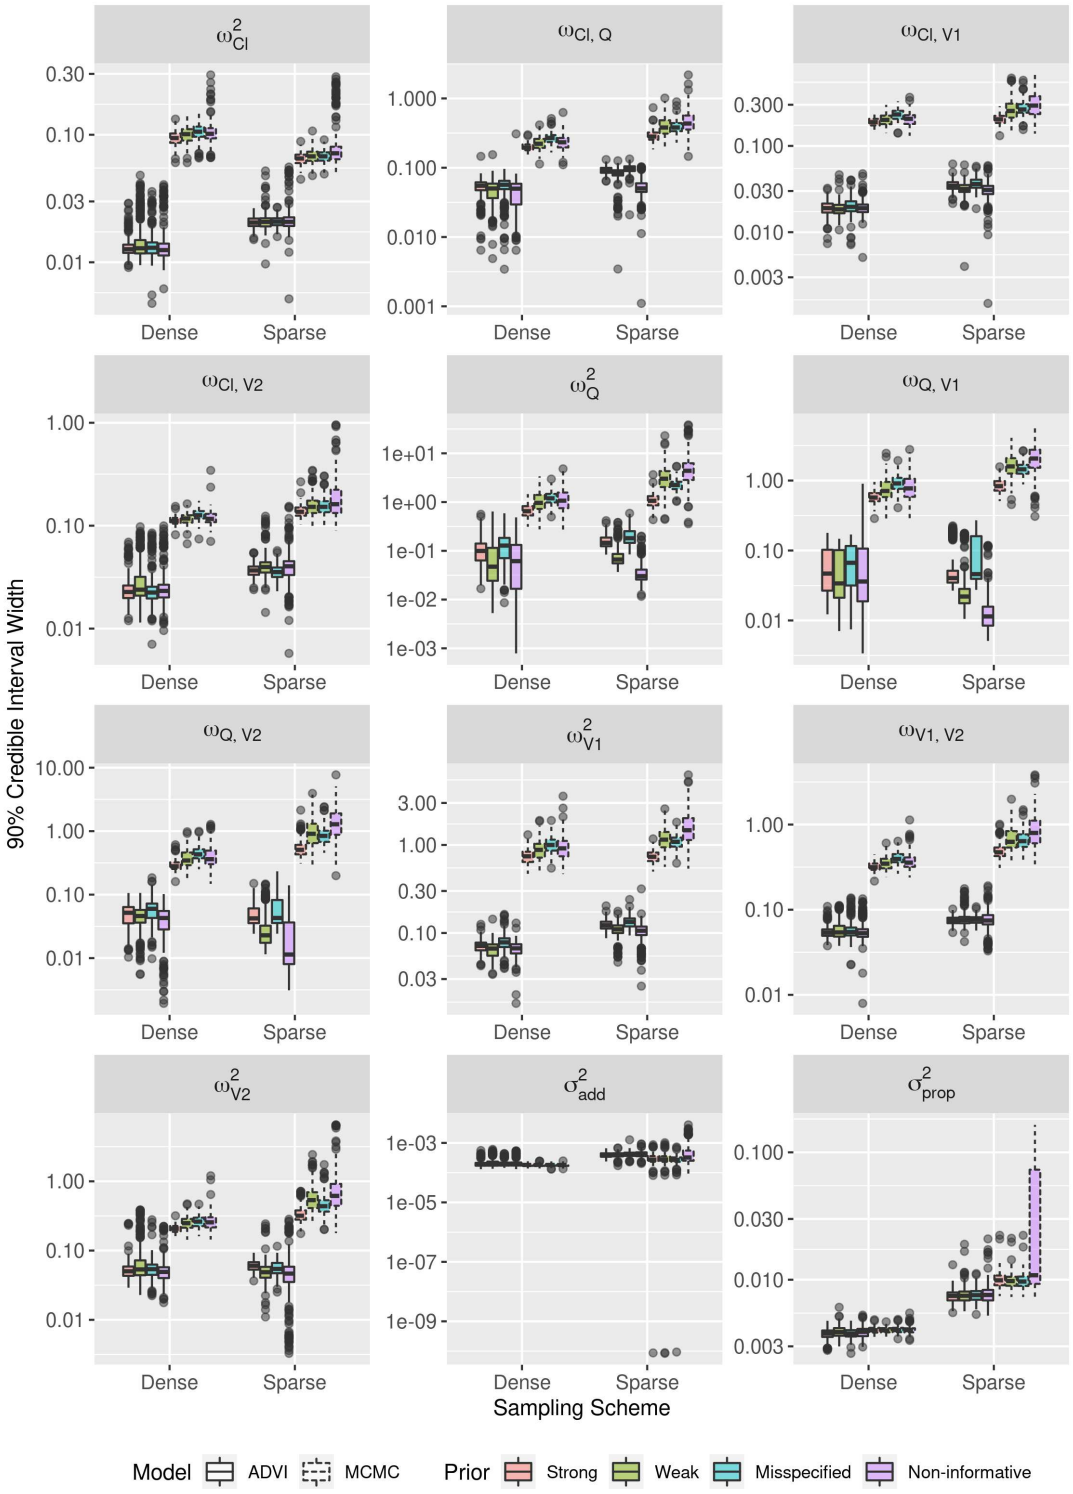

Figure 3.9: Two-compartment PK model simulations - posterior population parameter 90% equal-tailed credible interval widths for correctly specified model (M1) for informative priors.

3.3.2.2 Computation Time

Simulations were performed on a high-performance computing cluster running under CentOS Linux 7 (Core) with 1.90GHz or 2.40GHz Intel Xeon CPUs and up to 4 GB of memory per compute node. The computation time to fit each model is related to several factors. For ADVI, the step size, number of samples used for gradient estimation, and relative tolerance threshold all affect the time to convergence. Similarly, the NUTS HMC sampling time is controlled by the number of warmup and sampling iterations, the adaptation target acceptance rate, and several parameters related to the underlying Hamiltonian system. In addition, while the Markov chain is guaranteed to converge asymptotically, there is no criteria that can assure convergence with a finite number of iterations. Given these factors, it is difficult to make a definitive comparison of fit time between the two methods. However, examining the fit time ratios (Figures 3.10 and 3.11) reveals total fit times for MCMC estimation were several times longer than ADVI estimation under both simulation scenarios, with larger differences for the two-compartment model scenario.

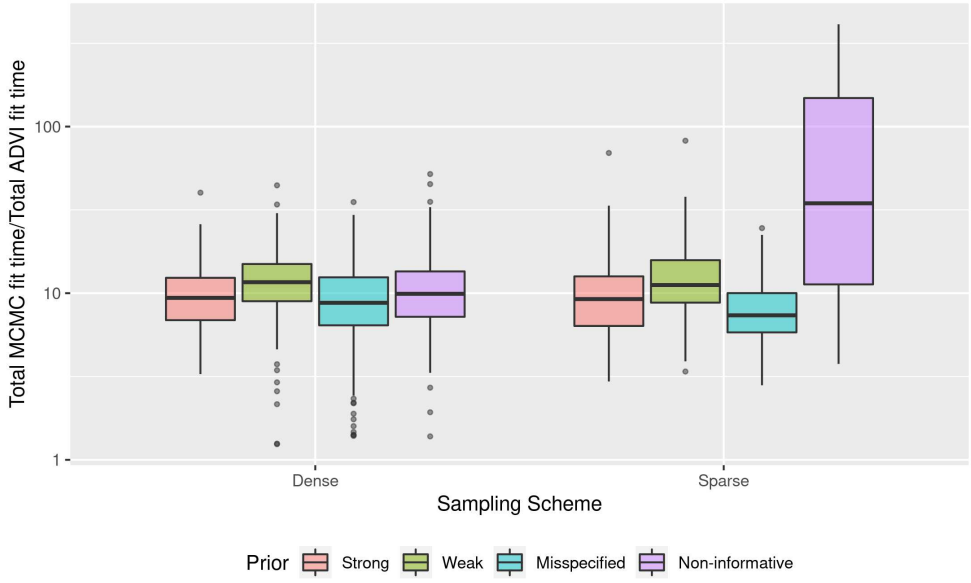

Figure 3.10: One-compartment PK model simulations - relative fit time (MCMC vs. ADVI) for correctly specified model (M1). The median (IQR) relative fit time for dense sampling was 9.37 (6.89 - 12.4) for the strong prior, 11.6 (8.95 - 15.0) for the weak prior, 8.76 (6.41 - 12.5) for the misspecified prior, and 9.92 (7.22 - 13.5) for the non-informative prior; for sparse sampling the median (IQR) relative fit time was 9.22 (6.36 - 12.6) for the strong prior, 11.2 (8.78 - 15.8) for the weak prior, 7.37 (5.82 - 10.0) for the misspecified prior, and 34.7 (11.3 - 148.0) for the non-informative prior.

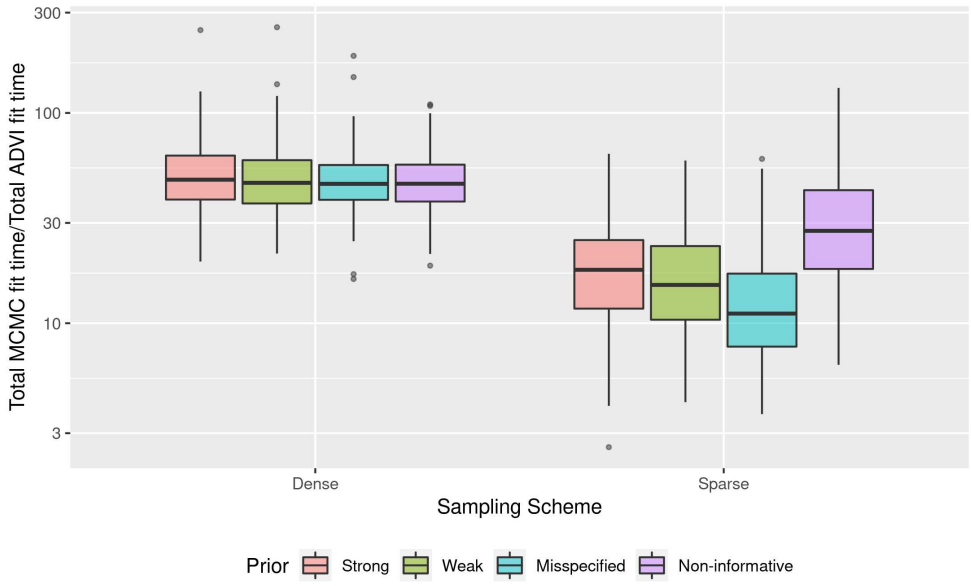

Figure 3.11: Two-compartment PK model simulations - relative fit time (MCMC vs. ADVI) for correctly specified model (M1). The median (IQR) relative fit time for dense sampling was 48.2 (38.7 - 62.7) for the strong prior, 46.5 (37.1 - 59.6) for the weak prior, 46.0 (38.6 - 56.6) for the misspecified prior, and 46.1 (37.9 - 56.8) for the non-informative prior; for sparse sampling the median (IQR) relative fit time was 18.0 (11.7 - 24.9) for the strong prior, 15.2 (10.4 - 23.3) for the weak prior, 11.1 (7.74 - 17.2) for the misspecified prior, and 27.5 (18.1 - 42.9) for the non-informative prior.

3.3.2.3 Individual Parameter Estimation

In addition to population parameters shared by all subjects, the model also estimates individual-level PK parameters (ψ_i) which combine population level fixed effects and subject specific random effects (Equation 1.7). For the one-compartment model, we estimate an individual clearance, Cl_i, and an individual central compartment volume of distribution, V_i, for each of the $i = 1, \ldots, N$ subjects. Estimates of percent bias in the posterior median individual parameters for the one-compartment model are shown in Figure 3.12. For both ADVI and MCMC, the posterior medians for Cl_i and V_i are centered around the true individual parameter values. For each sampling scheme, the distribution of percent bias for the posterior medians is similar for both estimation methods and all four priors. For the one-compartment simulation the 90% credible interval widths for individual parameter estimates (Figure 3.13) are also alike for both estimation methods.

Similar individual estimates are produced for the two-compartment scenario for total and inter-compartmental clearance and central and peripheral volume of distribution. Posterior medians for individual parameters for the two-compartment model are shown in Figure 3.14. Percent bias is small and similar for both estimation methods for most individual clearance estimates Cl_i. For the individual inter-compartmental clearance parameter, Q_i, bias from ADVI is more negative on average than MCMC, while for volume of distribution for the central

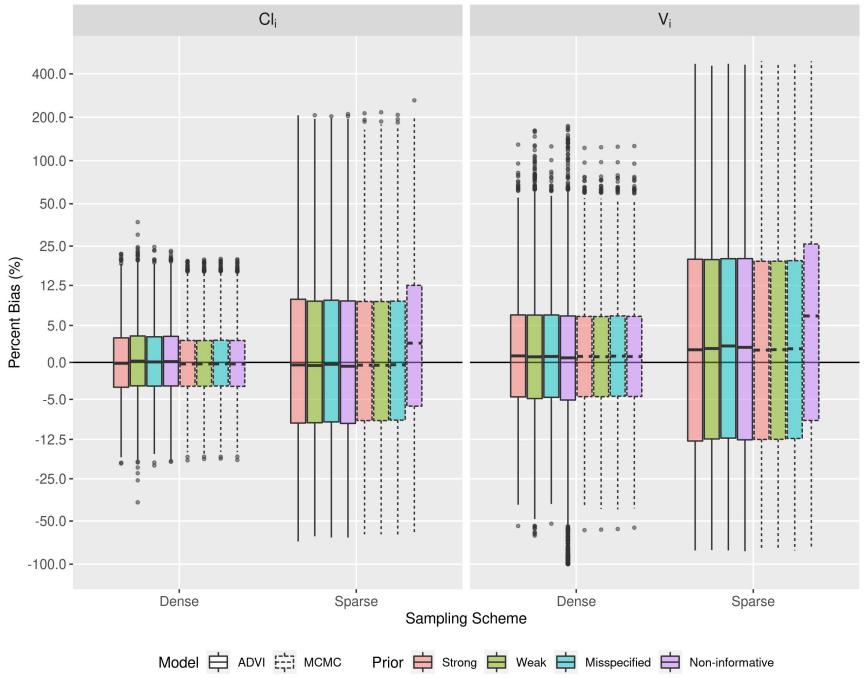

Figure 3.12: One-compartment PK model simulations - posterior median individual parameter estimates for correctly specified model (M1). Boxplots for the dense sampling scheme each have $250 \times 100 = 25{,}000$ individual parameter estimates while each boxplot for the sparse sampling scheme has $250 \times 300 = 75{,}000$ individual parameter estimates. Extreme outliers greater than 15 times interquartile range were removed.

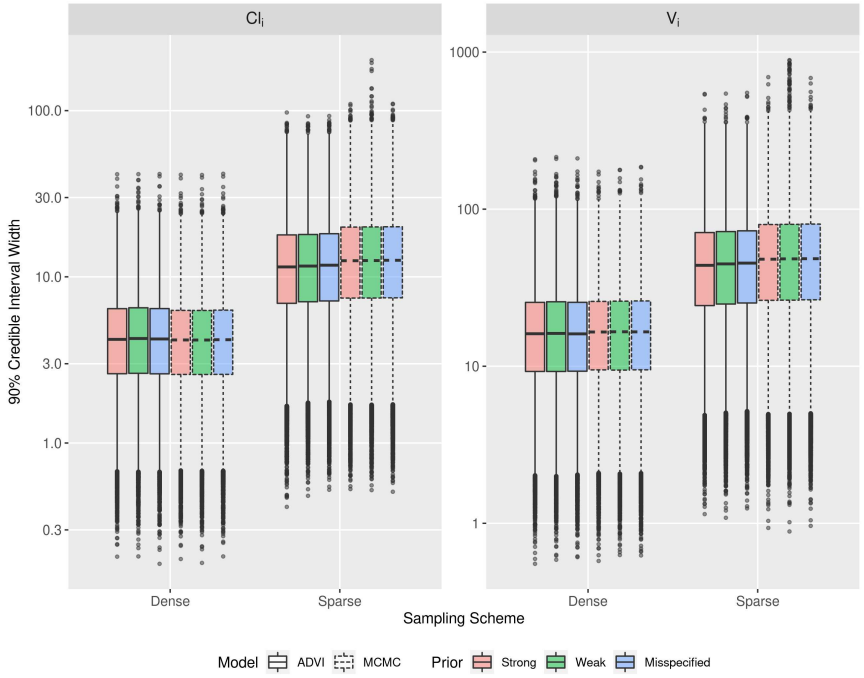

Figure 3.13: One-compartment PK model simulations - posterior individual parameter 90% equal-tailed credible interval widths for correctly specified model (M1) with informative priors. Note: credible interval widths for non-informative prior are not shown due to a large proportion of extreme outliers.

compartment, $V1_i$, estimates from ADVI are larger than those from MCMC. The 90% credible interval widths in Figure 3.15 are similar for Cl_i, but somewhat underestimated for the other parameters with the largest differences between estimation methods seen for Q_i.

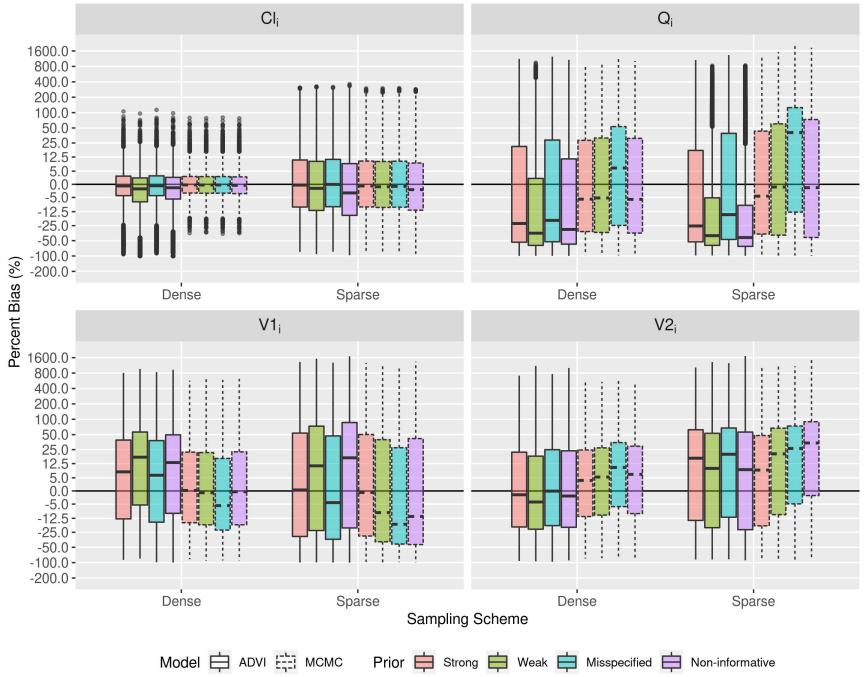

Figure 3.14: Two-compartment PK model simulations - posterior median individual parameter estimates for correctly specified model (M1). Boxplots for the dense sampling scheme each have $250 \times 100 = 25{,}000$ individual parameter estimates while each boxplot for the sparse sampling scheme has $250 \times 300 = 75{,}000$ individual parameter estimates. Extreme outliers greater than 15 times interquartile range were removed.

Overall the individual parameter estimates are similar between estimation methods, with larger differences for the two-compartment scenario. Also, unlike the population parameters, both the medians and credible intervals are similar for ADVI and MCMC.

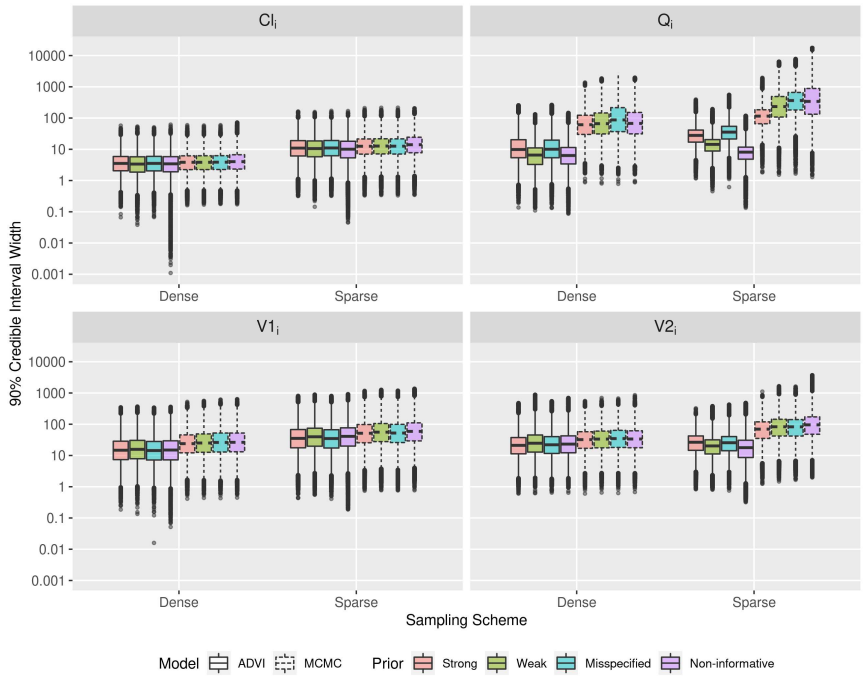

Figure 3.15: Two-compartment PK model simulations - posterior individual parameter 90% equal-tailed credible interval widths for correctly specified model (M1) for informative priors.

3.3.2.4 Diagnostics

We also examine the performance of the VI diagnostics. The failure of ADVI estimation to capture all aspects of the posterior well is evident when examining the global PSIS $\hat{k}$ diagnostic (Figures 3.43 and 3.61); all values are much greater than 0.7, with larger values for the sparse sampling scheme and the two-compartment simulation scenario.

3.3.2.5 Variational Simulation Based Calibration

Figure 3.16 shows the results of VSBC for the one-compartment model simulations using a dense sampling scheme with strong informative priors. Asymmetry in the histogram for a parameter indicates bias in the VI approximation to the marginal posterior. All parameters, with the possible exception of $\beta_{Cl_{X1}}$ show bias. Similarly, Figure 3.17 shows the VSBC histograms for the two-compartment model simulations using a dense sampling scheme with strong informative priors. VSBC plots for additional simulation settings are shown in Figures 3.44 - 3.48 for the one-compartment scenario and Figures 3.62 - 3.66 for the two-compartment scenario. Using the VSBC diagnostic shows some degree of bias in nearly all the population parameters confirming the results from the simulation study plots of median parameter estimates.

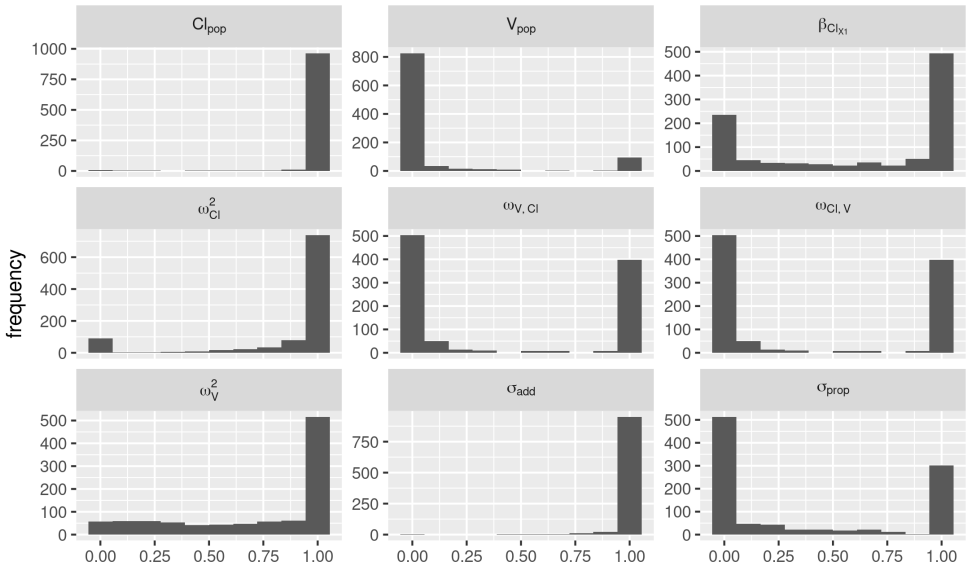

Figure 3.16: One-compartment PK model simulations - variational simulation based calibration for dense sampling and strong informative priors.

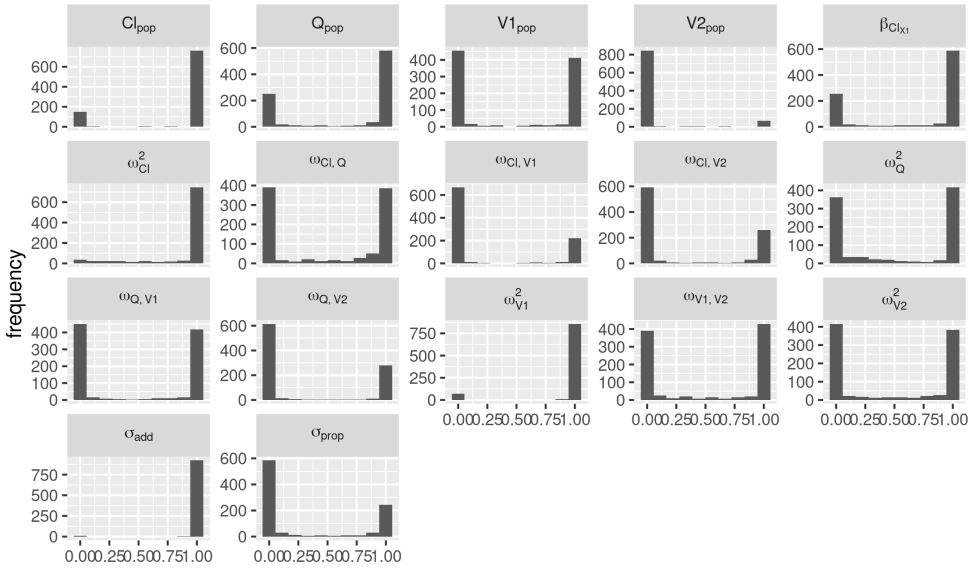

Figure 3.17: Two-compartment PK model simulations - variational simulation based calibration for dense sampling and strong informative priors.

3.3.2.6 Predicted Concentrations

We examine two types of posterior predicted concentrations. Individual predictions are made using population parameters and individual-specific random effects and provide predictions for existing subjects (e.g., "What is the C_{max} for subject 14 in this dataset?"). Population predictions also use the estimated population parameters, but draw new random effects to provide predictions for unobserved subjects (e.g., "What is the C_{max} for a new subject with the same covariates as subject 14?"). ADVI and MCMC produce similar individual predicted concentrations. For example, Figure 3.18 shows the observed values and individual predicted concentrations under dense sampling for 20 subjects from the first simulated dataset in the one-compartment scenario estimated using the true model (M1) with weak priors. The solid lines represent the medians and the shaded regions represent the 90% credible intervals. The ADVI predicted median and 90% credible intervals for the concentration profiles overlap almost completely with the MCMC predictions. Figure 3.19 shows the same example for sparse sampling; a similar pattern emerges with nearly identical predictions. Individual concentration predictions for 20 subjects from the first simulated dataset in the two-compartment scenario (estimated with weak priors) for dense and sparse sampling are shown in Figures 3.20 and 3.21, respectively. The posterior predicted concentration profiles are still similar across estimation methods; however, some discrepancies are seen especially for the credible intervals in the sparse sampling setting (e.g., subjects 24, 25, 34 and 35 in Figure 3.21). Figures 3.49-3.54 show the individual concentration predictions for the one-compartment scenario under the other prior settings; the trends between estimation methods are all similar except for sparse sampling with MCMC which has large credible intervals. The individual concentration predictions for the two-compartment scenario under the other prior settings are shown in Figures 3.67-3.72.

Population concentration predictions were also close for both estimation methods and simulation scenarios. Figures 3.22 and 3.23 show population predictions for subjects from one simulated dataset under dense and sparse sampling with weak priors for the one-compartment simulation and Figures 3.24 and 3.25 show population predictions under dense and sparse sampling with weak priors for the two-compartment simulation. As with the individual predictions, the population predictions are similar for both ADVI and MCMC. Figures 3.55-3.60 and 3.73-3.78 show population concentration predictions under the other priors for the one- and two-compartment scenarios, respectively.

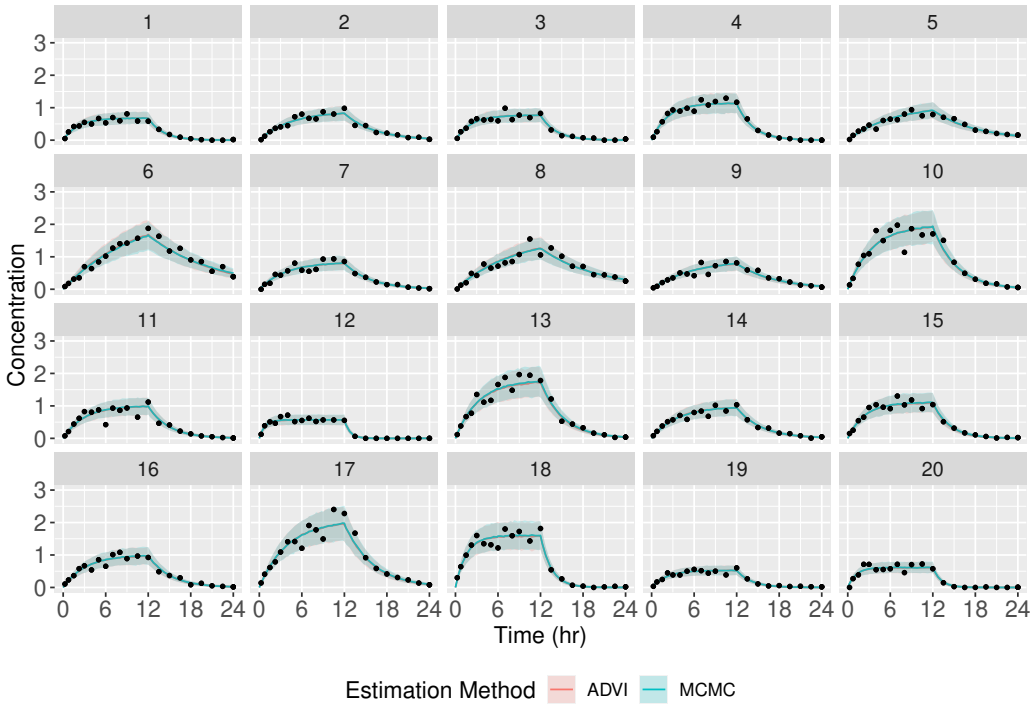

Figure 3.18: One-compartment PK model simulations - observed concentrations and individual predicted concentrations for the first simulated dataset with dense sampling and weak priors.

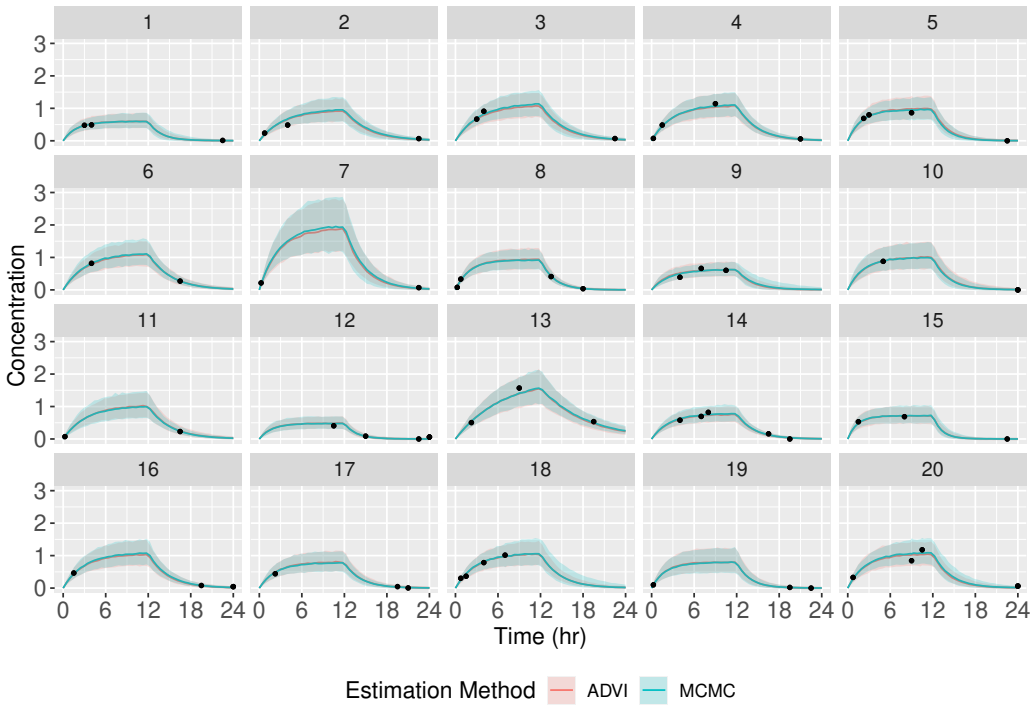

Figure 3.19: One-compartment PK model simulations - observed concentrations and individual predicted concentrations for the first simulated dataset with sparse sampling and weak priors.

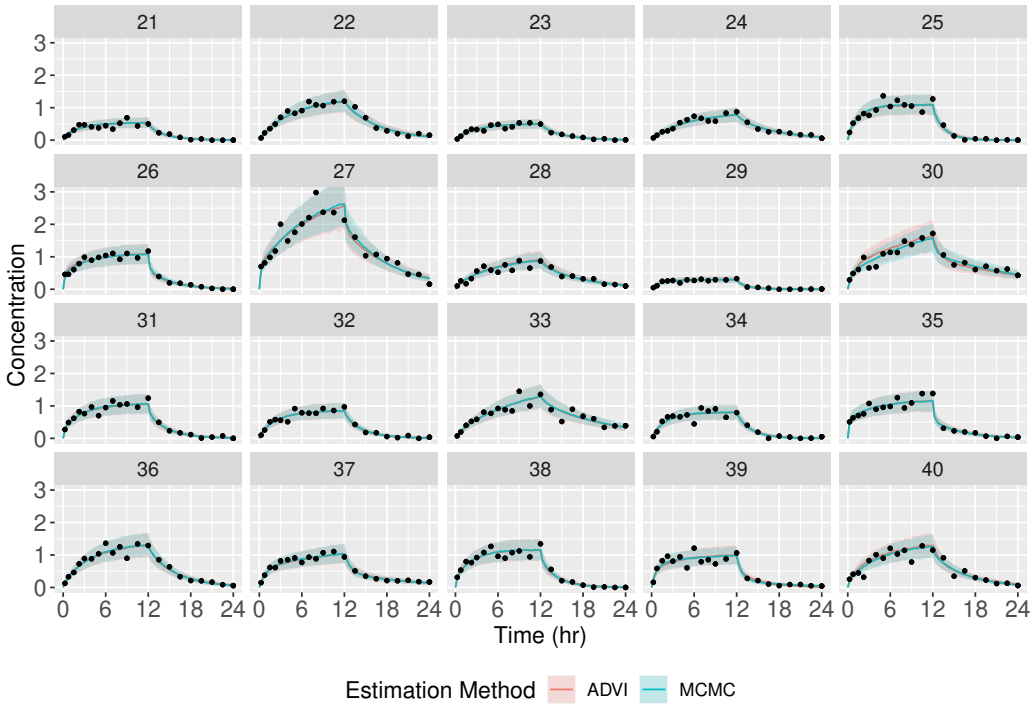

Figure 3.20: Two-compartment PK model simulations - observed concentrations and individual predicted concentrations for the first simulated dataset with dense sampling and weak priors.

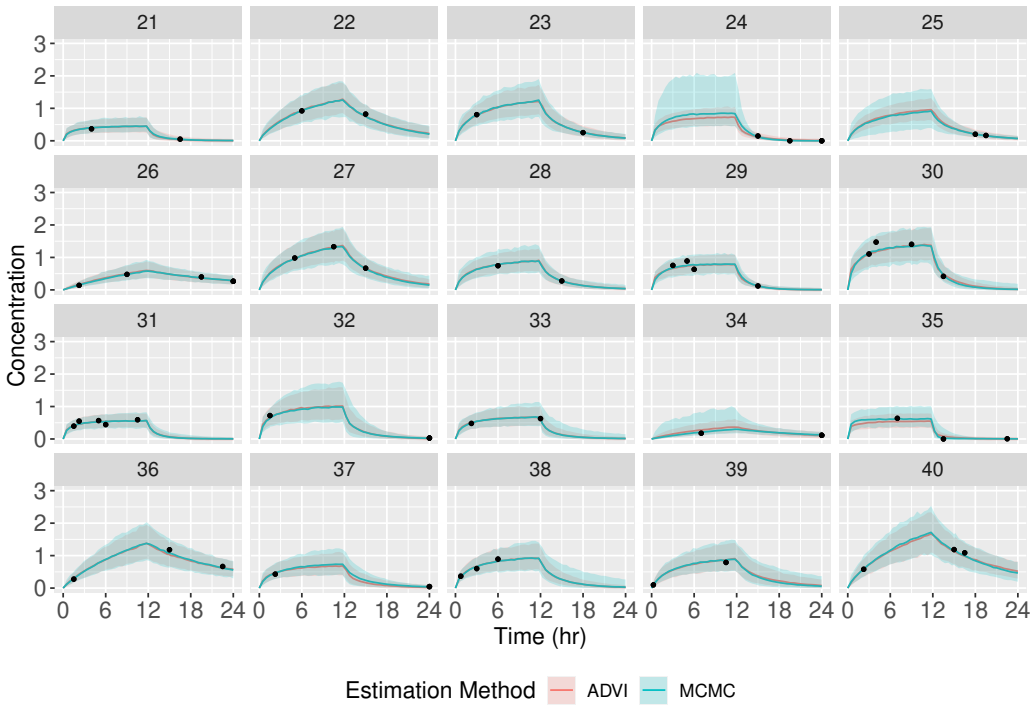

Figure 3.21: Two-compartment PK model simulations - observed concentrations and individual predicted concentrations for the first simulated dataset with sparse sampling and weak priors.

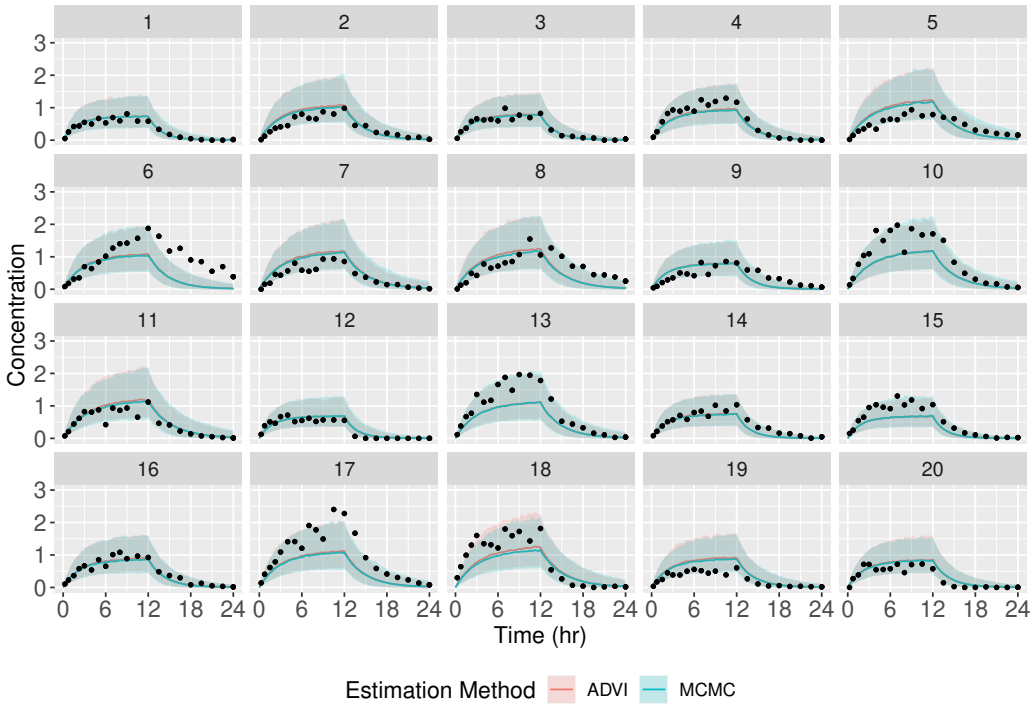

Figure 3.22: One-compartment PK model simulations - observed concentrations and population predicted concentrations for the first simulated dataset with dense sampling and weak priors.

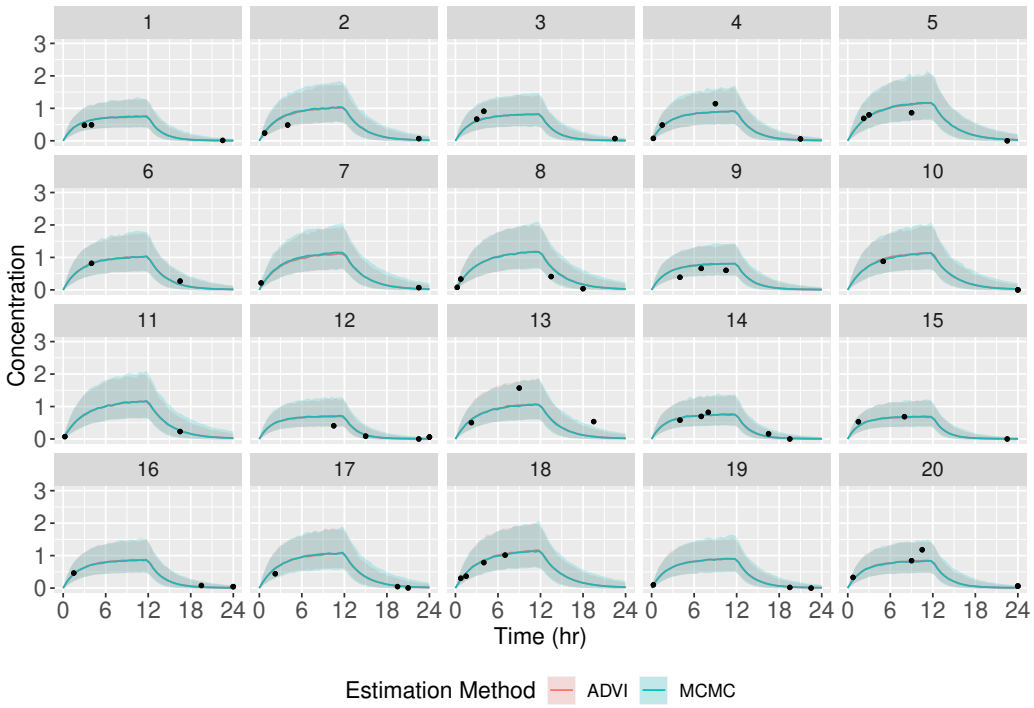

Figure 3.23: One-compartment PK model simulations - observed concentrations and population predicted concentrations for the first simulated dataset with sparse sampling and weak priors.

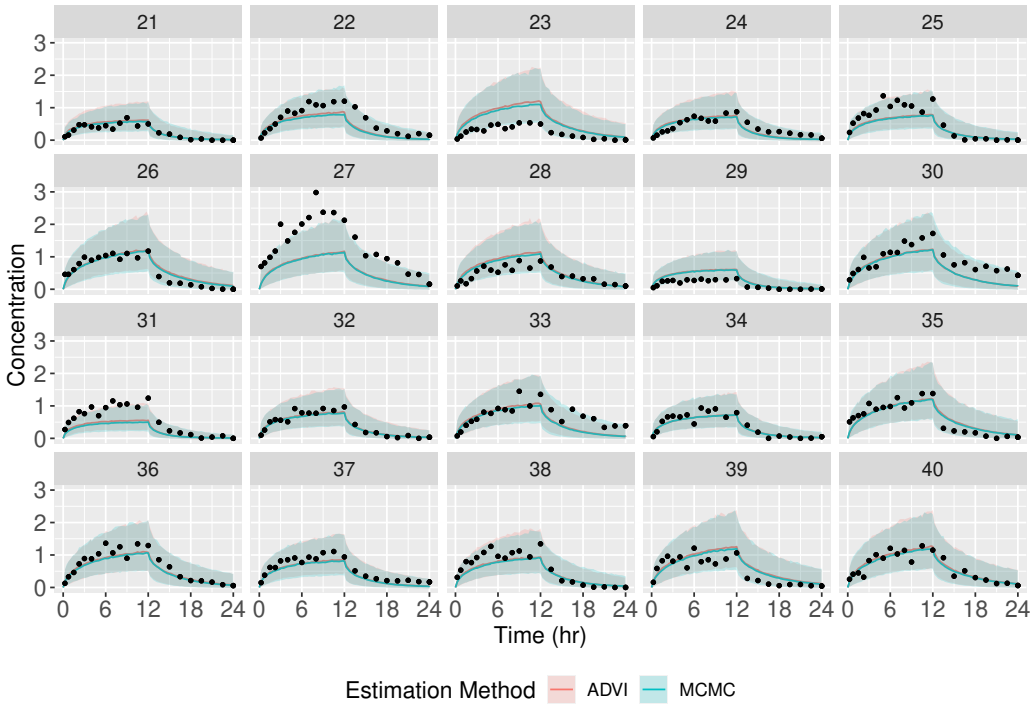

Figure 3.24: Two-compartment PK model simulations - observed concentrations and population predicted concentrations for the first simulated dataset with dense sampling and weak priors.

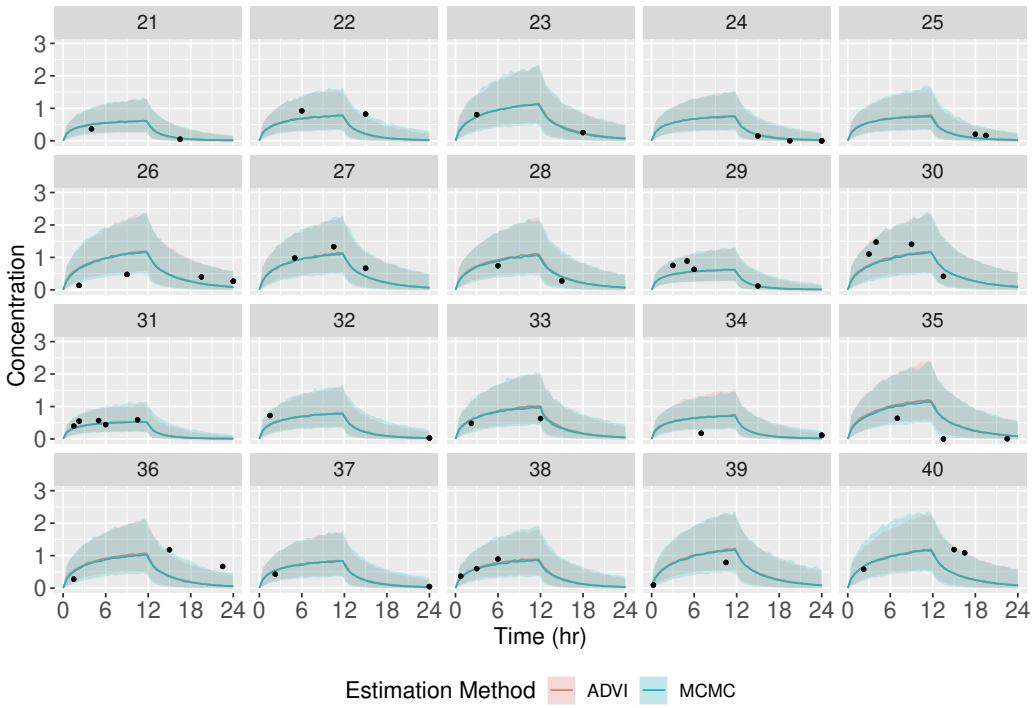

Figure 3.25: Two-compartment PK model simulations - observed concentrations and population predicted concentrations for the first simulated dataset with sparse sampling and weak priors.

The posterior concentration predictions at $t = 12$ estimate the maximum concentration (C_{max}) for each subject. We compare estimation methods across all simulations by calculating the ratio of posterior median C_{max} estimated by MCMC to posterior median C_{max} estimated by ADVI. In addition, we compare the ratio of the 90% credible interval widths for C_{max} estimated by MCMC vs. ADVI. Figures 3.26 and 3.27 show the ratios of median C_{max} and C_{max} credible intervals estimated by the two methods for the one-compartment scenario and Figures 3.28 and 3.29 show the same ratios for the two-compartment scenario. Medians and credible intervals for predicted individual and population concentrations are similar using both MCMC and ADVI.

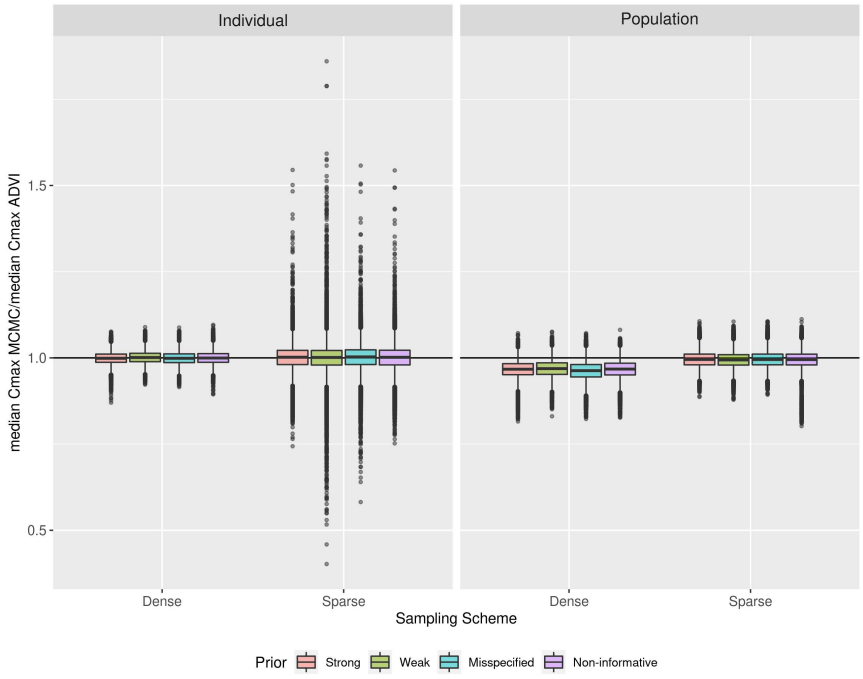

Figure 3.26: One-compartment PK model simulations - individual and population predicted median C_{max} ratios.

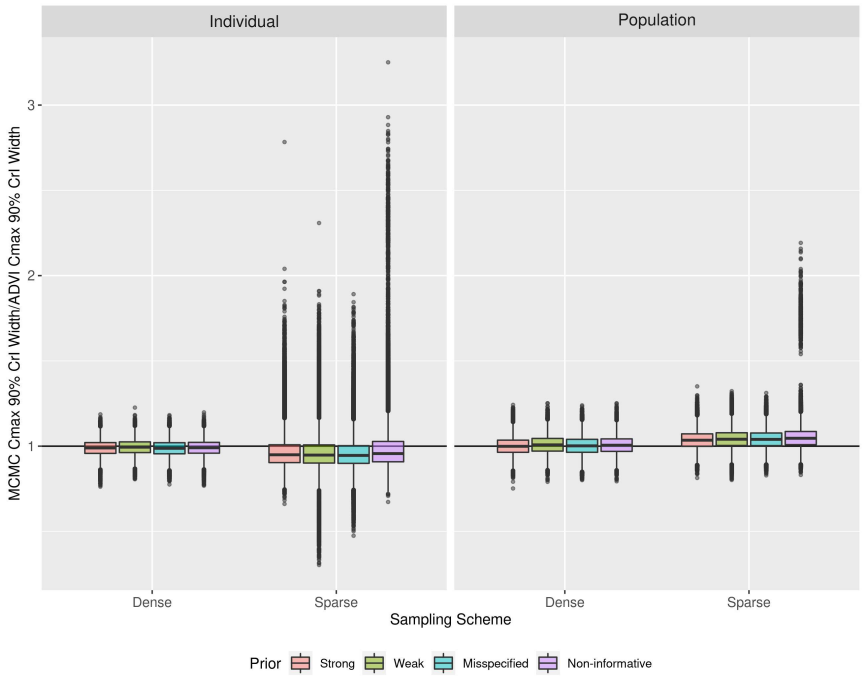

Figure 3.27: One-compartment PK model simulations - individual and population predicted C_{max} 90% equal-tailed credible interval width ratios.

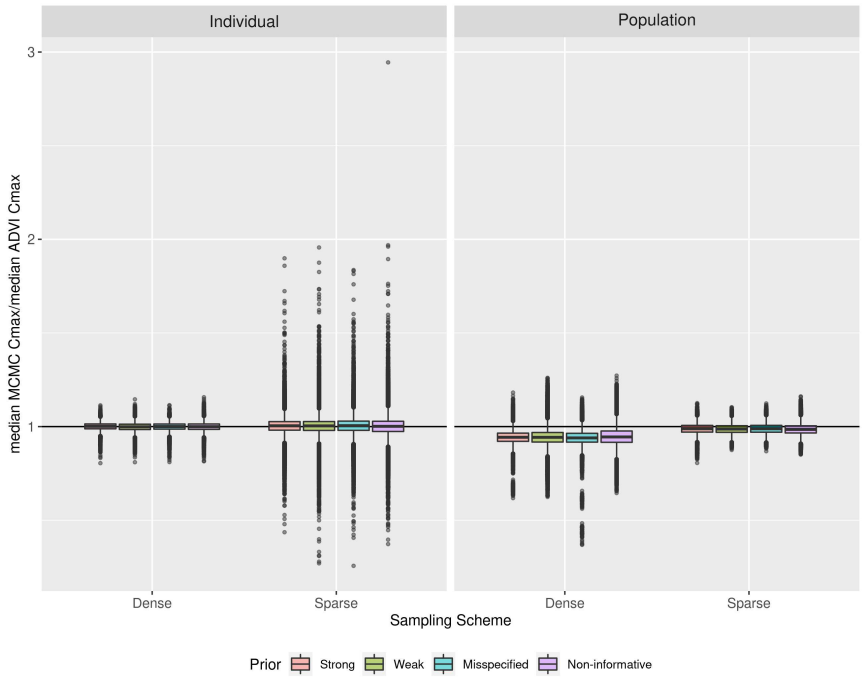

Figure 3.28: Two-compartment PK model simulations - individual and population predicted median C_{max} ratios.

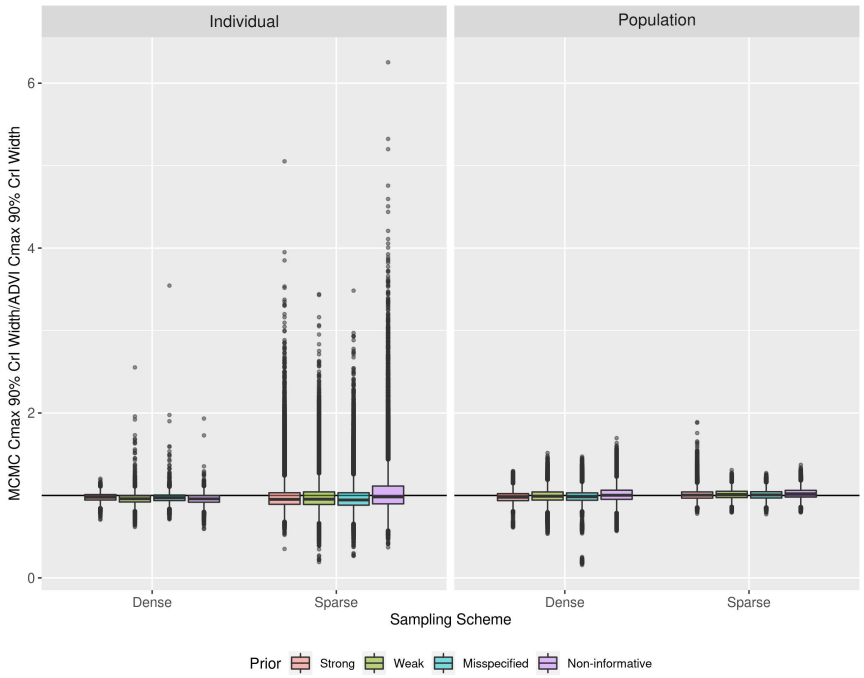

Figure 3.29: Two-compartment PK model simulations - individual and population predicted C_{max} 90% equal-tailed credible interval width ratios.

3.3.2.7 Model Selection

For model selection, we compare the three strategies – information criteria (Strategy 1), ELBO (Strategy 2), and 5-fold leave-subject-out cross-validation (Strategy 3) – by ranking each model. We assume the best ranking is the correctly specified model (M1), followed by the model with both the correct and noise variables (M3) and the model with only the incorrect noise covariate (M2).

For the one-compartment model scenario, using MCMC estimation with Strategy 1 and sparse sampling, the best order is selected for between 18 and 44% of simulated model fits depending on the information criteria and priors used, with the second best order 'M3,M1,M2' being selected for between 13 and 53% of fits (Table 3.3). Strategy 1 performs worse for dense sampling in terms of ranking the models. Using Strategy 1 with ADVI, the best ranking is selected for around 18 - 27% of model fits (Table 3.4). Using model selection Strategy 2, the best ranking is selected for between 40% and 75% of simulations with better performance under sparse sampling (Table 3.5). The percentage of fits where the worst model (M2) is ranked as best, is much lower for this strategy. Finally, Strategy 3, leave-subject-out cross-validation, outperforms both Strategies 1 and 2 in ability to correctly rank models (Table 3.6). For both sparse and dense sampling, the correct order is chosen for around 60% of simulations with order 'M3,M1,M2' selected for nearly all the others.

Under the two-compartment scenario, Strategy 1 using information criteria to rank the models performs poorly

using both MCMC and ADVI estimation with the best order selected for less than 25% of fits using MCMC and less than 23% of fits using ADVI (Tables 3.7 and 3.8). As with the one-compartment scenario, performance is slightly better for sparse sampling. Using information criteria with ADVI and dense sampling incorrectly ranks the worst model (M2) as best for more than 30% of fits for most combinations of sampling scheme and prior. Strategy 2 selects the correct model order for between 45% and 85% of fits for sparse sampling, but performs poorly for dense sampling (Table 3.9). Using Strategy 3, the correct order or the second best order is selected for more than 96% of fits under sparse sampling and between 42% and 95% of fits with dense sampling (Table 3.10).

Because the model selection simulations produced some non-intuitive results we also plot the values for WAIC calculated using MCMC and ADVI from Strategy 1, $-2 \times ELBO$ from Strategy 2, and $-2 \times lppd_{cv}$ from Strategy 3 for the two-compartment scenario to provide more insight on the model ranking[1]. While dense sampling may be expected to perform better than sparse sampling, examining the plot of WAIC calculated using MCMC (Figure 3.79) provides one explanation for the apparent discrepancy. When there is a large amount of information, as with dense sampling, all three models can predict well and have similar predictive accuracy. However, if all the models perform similarly, ranking the models to select the correct order is more difficult. The number of covariates is the same for M1 and M2, and M3 has only one more covariate so the penalty for effective number of parameters is not very substantial. The plot of WAIC calculated using ADVI (Figure 3.80) shows worse estimated predictive accuracy overall using the approximation; WAIC values are larger and more variable than when using MCMC. The mediocre predictive performance for Strategy 1 may also be related to fact that the criteria do not account for the hierarchical structure of the data [76].

Figure 3.81 shows the results for Strategy 2 and Figure 3.82 shows the results for Strategy 3. For these strategies, sparse sampling also performs as well or better than dense sampling, even when the predictive accuracy under dense sampling is not similar for all models. A potential explanation is related to the design of the simulation scenarios, specifically the number of subjects in the dense and sparse scenarios and the number of samples per subject. For dense sampling each dataset contained 100 subjects with 21 samples each and the sparse sampling datasets contained 300 subjects with between 2 and 5 samples each. While the dense data contains more observations overall, the sparsely sampled data contains more subjects and may be better able to estimate covariate effects. Several papers provide evidence to support this explanation. In the context of optimal design for population PK studies, Duffull et al. (2005) compared scenarios with a fixed total number of samples, but varying numbers of subjects and samples per subject using both balanced (e.g., 20 subjects with 12 samples each; 30 subjects with 8 samples each; 48 subjects with 5 samples each, etc.) and unbalanced (e.g.,

[1]The criteria from Strategy 2 and 3 are multiplied by -2 to be on approximately the same scale and have similar interpretation to the information criteria in Strategy 1; lower values of information criteria, $-2 \times ELBO$, and $-2 \times lppd_{cv}$ indicate better fit.

66 subjects with 3 samples each and 21 subjects with 2 samples each) designs. They found higher efficiency (measured using D-optimality) for designs with more subjects, but fewer samples per subject [93]. Using a similar approach, Retout et al. (2007) found that nonlinear mixed effects model designs with fewer samples per subject but a larger number of subjects had higher efficiency and greater power [94].

In order to investigate the model selection results in more detail, two additional sensitivity analyses were performed by increasing or decreasing size of the true $\beta_{Cl_{X_1}}$ effect used simulate the datasets. For the first sensitivity analysis, we set $\beta_{Cl_{X_1}} = 0.1$, reducing the separation between the true and misspecified models and making model selection more difficult. The second sensitivity analysis increased $\beta_{Cl_{X_1}}$ to 1 to make model selection easier.

The results of these sensitivity analyses using ADVI estimation are shown in Tables 3.11, 3.12, and 3.13 and Figures 3.83, 3.84, and 3.85 for $\beta_{Cl_{X_1}} = 0.1$; and Tables 3.14, 3.15, and 3.16 and Figures 3.86, 3.87, and 3.88 for $\beta_{Cl_{X_1}} = 1$. The three strategies all did poorly when the true effect size was reduced to 0.1 in the first sensitivity analysis with no more than 33% of datasets choosing the correct order using any of the strategies. For the second sensitivity analysis, the results using the information criteria based on ADVI were still poor with the best model ranking selected for less than 25% of datasets. Using ELBO, between 60% and 80% of datasets ranked the models correctly using sparse sampling and between 31% and 40% ranked the models correctly using dense sampling, exhibiting the same trend seen for the main model selection simulations with $\beta_{Cl_{X_1}} = 0.4$. With 5-fold cross-validation, the correct ranking was selected between for between 62% and 88% of datasets, with nearly all the rest of the datasets choosing the second best order.

3.4 Case Study

3.4.1 Background and Methods

Ketorolac is a nonsteroidal anti-inflammatory drug used to treat pain after surgery. Data for ketorolac IV dosing and patient characteristics were gathered from EHRs and combined with remnant specimens collected using opportunistic sampling to determine blood plasma concentration. Using real-world data allows us to collect information for several hundred subjects (compared to less than 20 for many reported pediatric studies). However, the remnant concentration samples are sparse and not collected with the goal of estimating the PK profile. Therefore, we leverage the Bayesian framework to combine prior information from several small designed clinical studies with our larger observational dataset. The primary goals are to find a model that adequately describes individual PK time-concentration profiles and parameters and to determine whether any covariates have an effect on PK profile.

Based on previous literature, ketorolac PK can be described well using a two-compartment model. We compare

a model with only allometric scaling to a model that also includes a maturation factor based on postmenstrual age (pma). The compartmental PK model for multiple IV bolus doses is given by:

$$D = \{d_k, t_{d_k}\} \quad \psi = \{Cl, Q, V_1, V_2\}$$

$$\alpha = \frac{\frac{QCl}{V_2 V_1}}{\beta} \quad \beta = \frac{1}{2}\left[\frac{Q}{V_1} + \frac{Q}{V_2} + \frac{Cl}{V_1} - \sqrt{\left(\frac{Q}{V_1} + \frac{Q}{V_2} + \frac{Cl}{V_1}\right)^2 - 4\frac{QCl}{V_2 V_1}}\right] \quad A = \frac{1}{V_1}\frac{\alpha - \frac{Q}{V_2}}{\alpha - \beta} \quad B = \frac{1}{V_1}\frac{\beta - \frac{Q}{V_2}}{\beta - \alpha}$$

$$f(D, \psi, t) = \sum_{k=1}^{K} d_k (A e^{-\alpha(t - t_{d_k})} + B e^{-\beta(t - t_{d_k})}),$$

where d_k and t_{d_k} are the amounts and administration times, respectively, for the $k = 1, \ldots, K$ doses. Using this compartmental model, the stage 1 model for observed concentration y_{ij} for individual i at time j including combined proportional and additive error can be defined as:

$$y_{ij} = f(D_i, \psi_i, t_{ij}) \times (1 + \varepsilon_{ij}^{prop}) + \varepsilon_{ij}^{add}$$

$$\varepsilon_{ij}^{prop} \sim N(0, \sigma_{prop}^2) \quad \varepsilon_{ij}^{add} \sim N(0, \sigma_{add}^2)$$

The stage 2 model with only allometric scaling is:

$$\log(Cl_i) = \log(Cl_{pop}) + 0.75 \times \log\left(\frac{weight_i}{15}\right) + \eta_{Cl_i},\ \eta_{Cl} \sim N(0, \omega_{Cl}^2)$$

$$\log(Q_i) = \log(Q_{pop}) + 0.75 \times \log\left(\frac{weight_i}{15}\right) + \eta_{Q_i},\ \eta_{Q} \sim N(0, \omega_{Q}^2)$$

$$\log(V_{1i}) = \log(V_{1\,pop}) + 1 \times \log\left(\frac{weight_i}{15}\right) + \eta_{V_{1i}},\ \eta_{V_1} \sim N(0, \omega_{V_1}^2)$$

$$\log(V_{2i}) = \log(V_{2\,pop}) + 1 \times \log\left(\frac{weight_i}{15}\right) + \eta_{V_{2i}},\ \eta_{V_2} \sim N(0, \omega_{V_2}^2)$$

$$\{\eta_{Cl}, \eta_Q, \eta_{V_1}, \eta_{V_2}\} = \eta \sim N_4(0, \Omega)$$

The stage 2 model with allometric scaling and maturation factor is identical except for the clearance model:

$$\log(Cl_i) = \log(Cl_{pop}) + 0.75 \times \log\left(\frac{weight_i}{15}\right) + \log\left(\frac{1}{1 + (TM_{50}/pma_i)^{Hill}}\right) + \eta_{Cl_i}$$

In the simulation study, models using non-informative priors had worse performance for both estimation methods, especially with sparse sampling, so we use an informative prior based on previously published pediatric ketorolac studies and literature on PK scaling for children (Figure 3.30 and Table 3.17). The ketorolac dataset was split into a training dataset with 320 subjects used to fit the models and a validation dataset with 86 subjects.

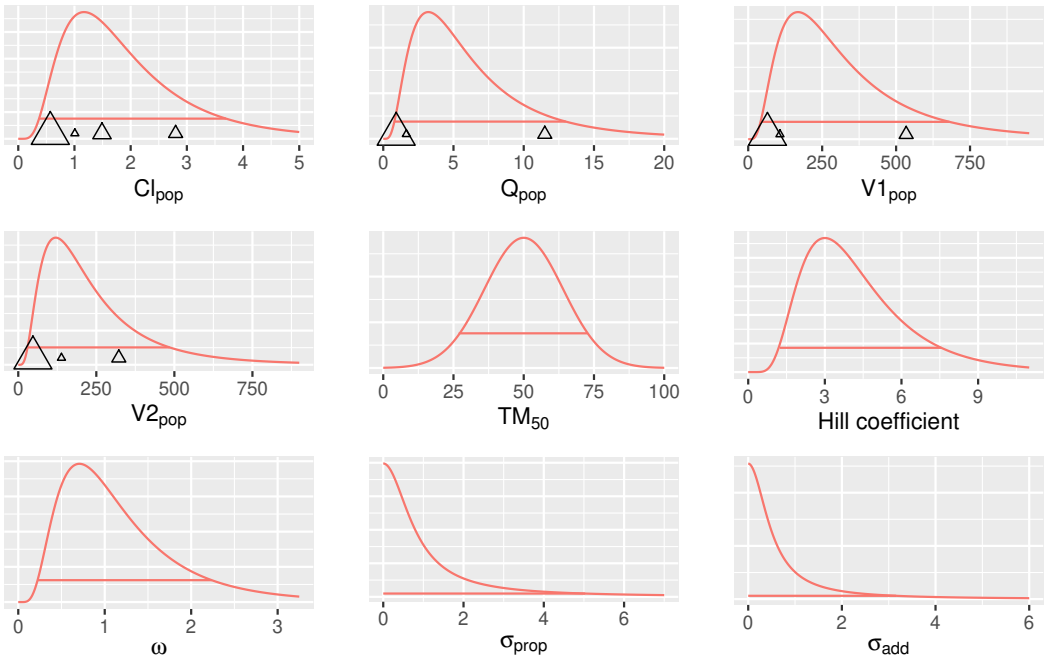

Figure 3.30: Priors for ketorolac PK models. For Cl_{pop}, Q_{pop}, $V1_{pop}$, and $V2_{pop}$ open triangles indicate point estimates from previous studies; the size of the triangle corresponds to the sample size of the previous study.

3.4.2 Results

For the model with only allometric scaling, the mean-field ADVI algorithm took around 6 minutes and 26 seconds to converge while NUTS HMC took 358 minutes (5 hours, 58 minutes) for 4 parallel chains running concurrently with 4000 warmup iterations and 250 sampling iterations. Adding a sigmoid Hill maturation factor on clearance, the mean-field ADVI algorithm took 6 minutes and 33 seconds to converge. The NUTS HMC sampler took 341 minutes (5 hours, 41 minutes) for 4 parallel chains with 4000 warmup iterations and 250 sampling iterations. We use 5-fold leave-subject-out cross-validation with the training data to compare the two models. The log pointwise predictive density ($lppd_{cv}$) computed using 5-fold cross-validation was -359.35 for the model with only allometric scaling and -353.2 for the model with allometric scaling and maturation factor. We present results from the second model.

Posterior population parameter density estimates are shown in Figure 3.31. The ADVI and MCMC parameter estimates mirror trends seen in the two-compartment simulation study with some discrepancies in the central tendency of the distributions, especially for the variance components, and smaller posterior uncertainty for all population parameters, except σ^2_{add} and σ^2_{prop}, using ADVI compared to MCMC. Posterior individual clearance (Cl_i) estimates for 20 individuals are shown in Figure 3.32. In contrast to the population parameters, individual clearance density estimates for many subjects are similar for both estimation methods.

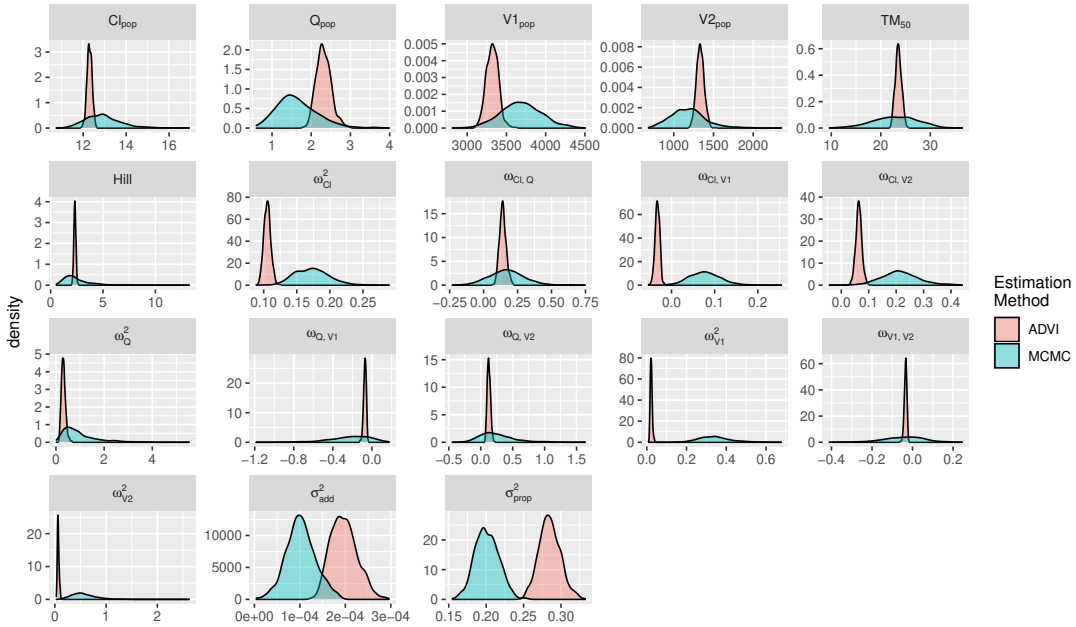

Figure 3.31: Posterior parameter estimates from ketorolac PK model with fixed allometric scaling and age maturation parameters.

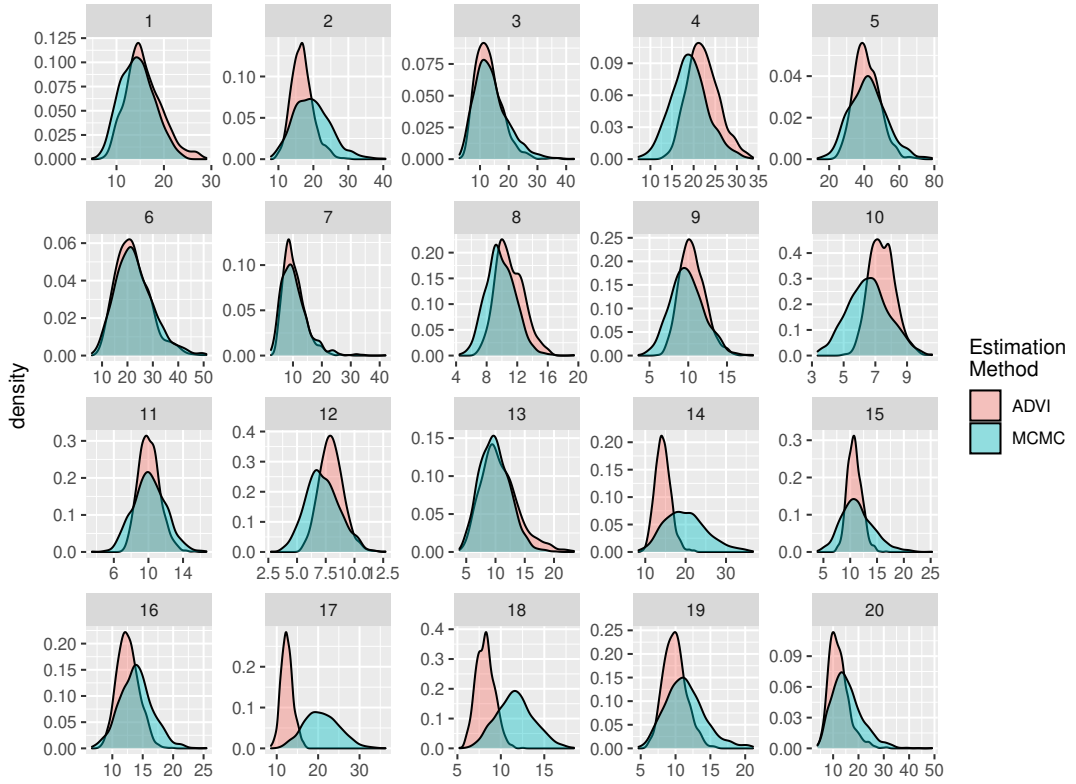

Figure 3.32: Individual posterior clearance (Cl_i) for 20 individuals from ketorolac PK model with fixed allometric scaling and age maturation parameters.

3.4.3 Posterior Checks and Diagnostics

Individual and population posterior predictive plots for 20 subjects from the training data using the model with allometric scaling and age maturation are shown in Figures 3.33 and 3.34. Individual PK profile estimates are similar for the majority of subjects using both estimation methods, however misestimation of individual parameters with ADVI, especially an incorrect mode, translates to misestimation in the individual PK profile. For example, compare the individual clearance estimates for subjects 14, 17 and 18 in Figure 3.32 to the individual predictions for these subjects in Figure 3.33. In contrast, population predictions are much more similar between estimation methods.

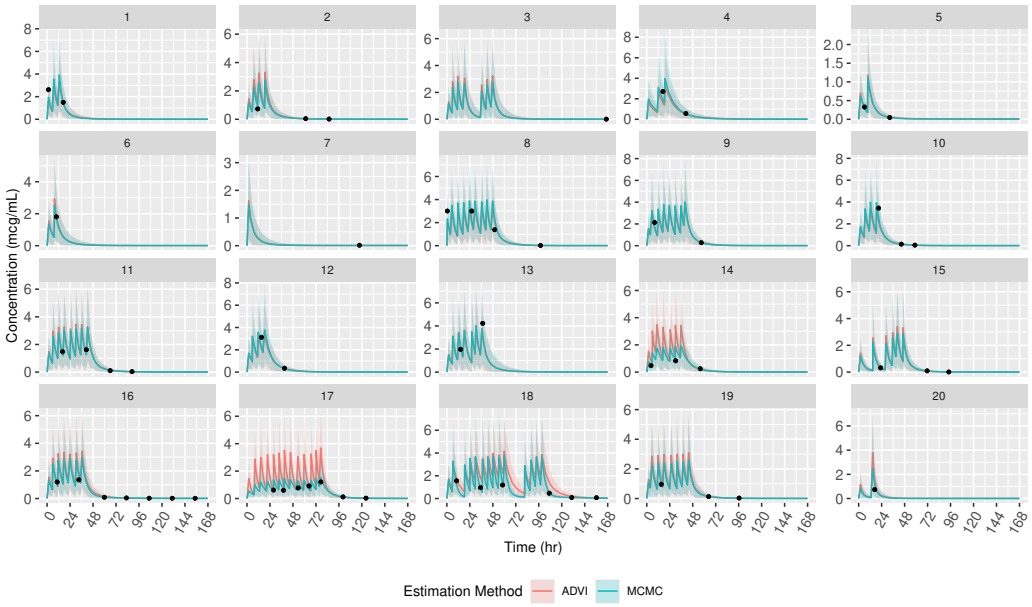

Figure 3.33: Individual posterior predictions for 20 subjects from training dataset using ketorolac PK model with fixed allometric scaling and age maturation parameters.

In addition to predictions for the individuals in the training data, population predictions for subjects in the validation data are shown in Figure 3.35. As with the training data population predictions, the ADVI and MCMC estimates are nearly identical. There are several examples where the ADVI and MCMC estimates are similar, but both methods are far from the observed concentration values (e.g., the first observed concentrations for subjects 1 and 8 in the training set and concentrations for subjects 332, 335, and 337 in the validation dataset). This may be related to errors in the EHR data. In a study of vancomycin in neonates van der Meer et al. (2012) [95] found that errors in patient records such as incorrect "time of administration, amount of dose given, registration or recording of dosing information and sampling times, and drug assay, or measurement, errors" had limited influence on outcomes of population PK modeling but can have detrimental effects on individual maximum a posteriori Bayesian estimation. The persistent and consistent deviation from predicted

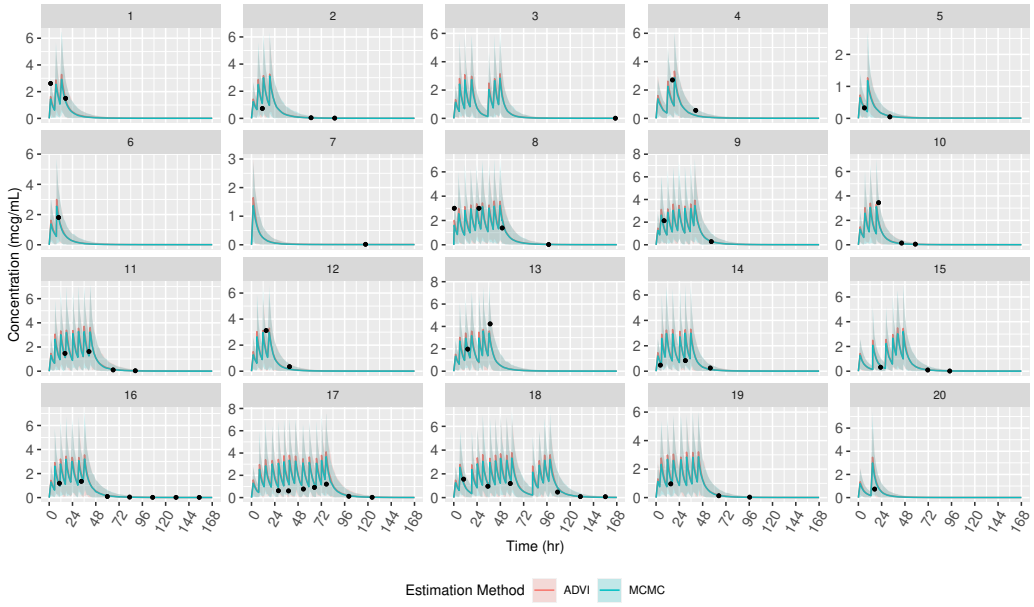

Figure 3.34: Population posterior predictions for 20 subjects from training dataset using ketorolac PK model with fixed allometric scaling and age maturation parameters.

values for some subjects (e.g., 335) may also be due to model misspecification, with an unmeasured covariate influencing the PK profile for these subjects.

We also examine several commonly used graphical PK model checks including observed vs. predicted concentration plots, random effects vs. covariates, and visual predictive checks (VPCs). Observed vs. predicted plots are shown in Figure 3.36 with results from MCMC estimates in the top row and ADVI estimates in the bottom row. The overall trend for population predicted vs. observed concentrations is similar for both estimation methods. For individual predicted vs. observed concentrations, the relationship is similar for predicted concentrations between 0.01 and around 1 mcg/mL, however ADVI estimation shows a trend toward overestimation for larger concentrations.

Plots of the random effects for clearance (η_{Cl_i}) vs. weight and postmenstrual age are displayed in Figure 3.37. Plots using ADVI and MCMC estimates are again much the same; both indicate some unexplained weight and age effects among the heaviest and oldest children in the cohort.

Finally, we compare the VPCs produced using MCMC and ADVI estimates. Figure 3.38 shows VPCs with observed values for the 10th, 50th and 90th percentiles (solid lines) and theoretical model-based values (dashed lines) for the same percentiles along with 90% prediction intervals for the model-based estimates (shaded regions). For a well-fit model, the observed values should be similar to the theoretical values. In this example, there is a discrepancy between the observed and theoretical values due to the lower limit of quantification for

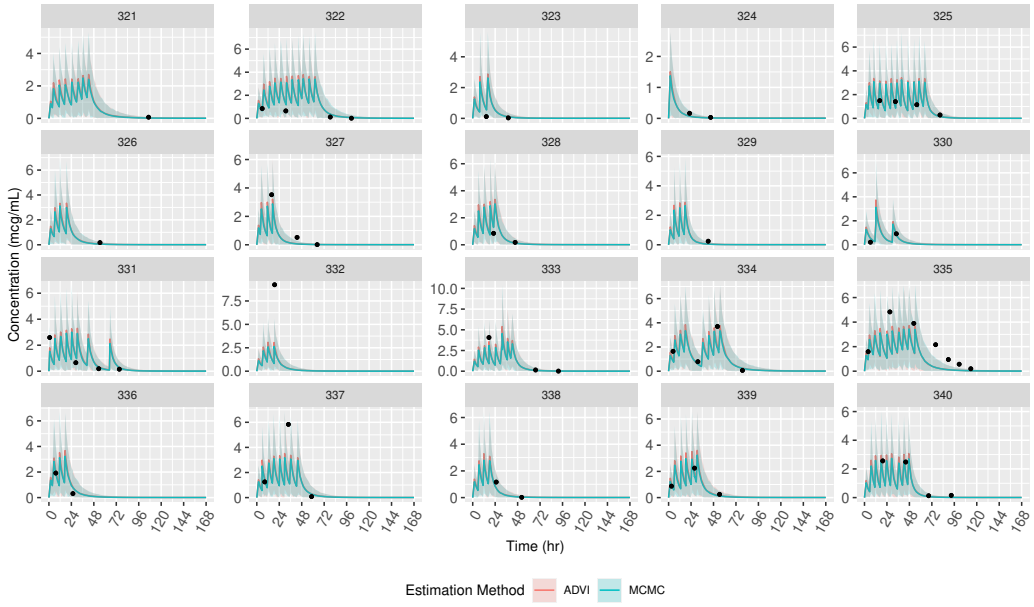

Figure 3.35: Population posterior predictions for 20 subjects from validation dataset using ketorolac PK model with fixed allometric scaling and age maturation parameters.

observed concentrations (0.01 mcg/mL); the model accounts for values censored below this threshold and is able to make predictions for values < 0.01 mcg/mL. As with the other graphical PK checks both estimation methods produce nearly identical plots.

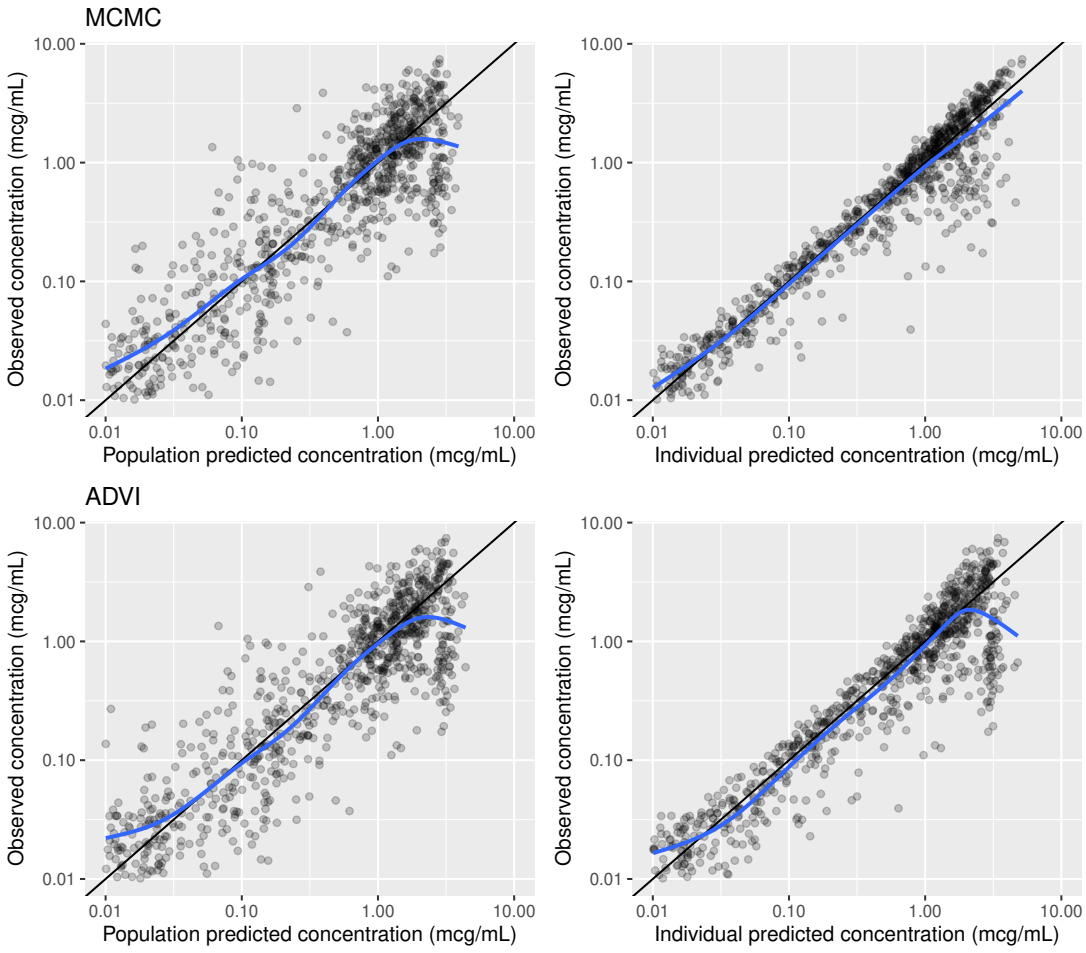

Figure 3.36: Observed concentration vs. population and individual posterior predictions from ketorolac PK model with fixed allometric scaling and age maturation parameters.

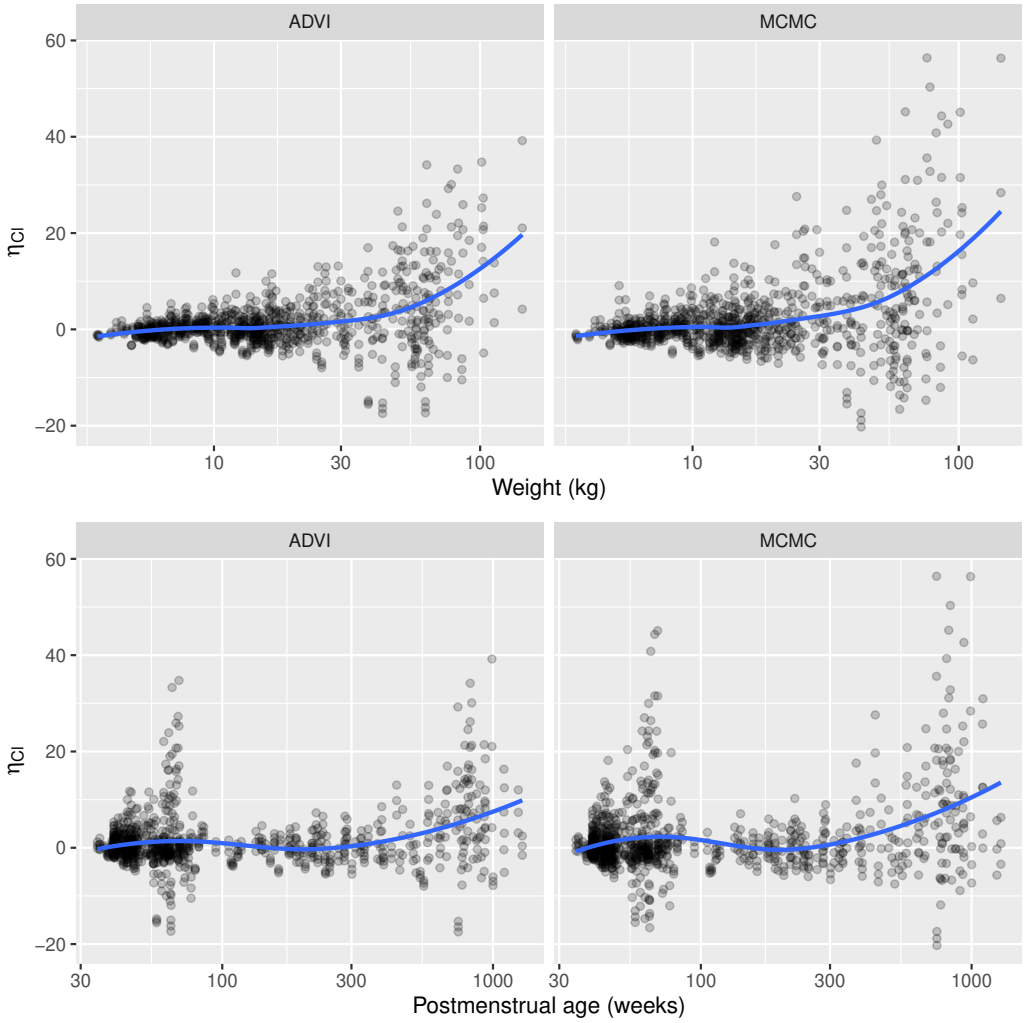

Figure 3.37: Random effects for clearance vs. covariates for ketorolac PK model with fixed allometric scaling and age maturation parameters.

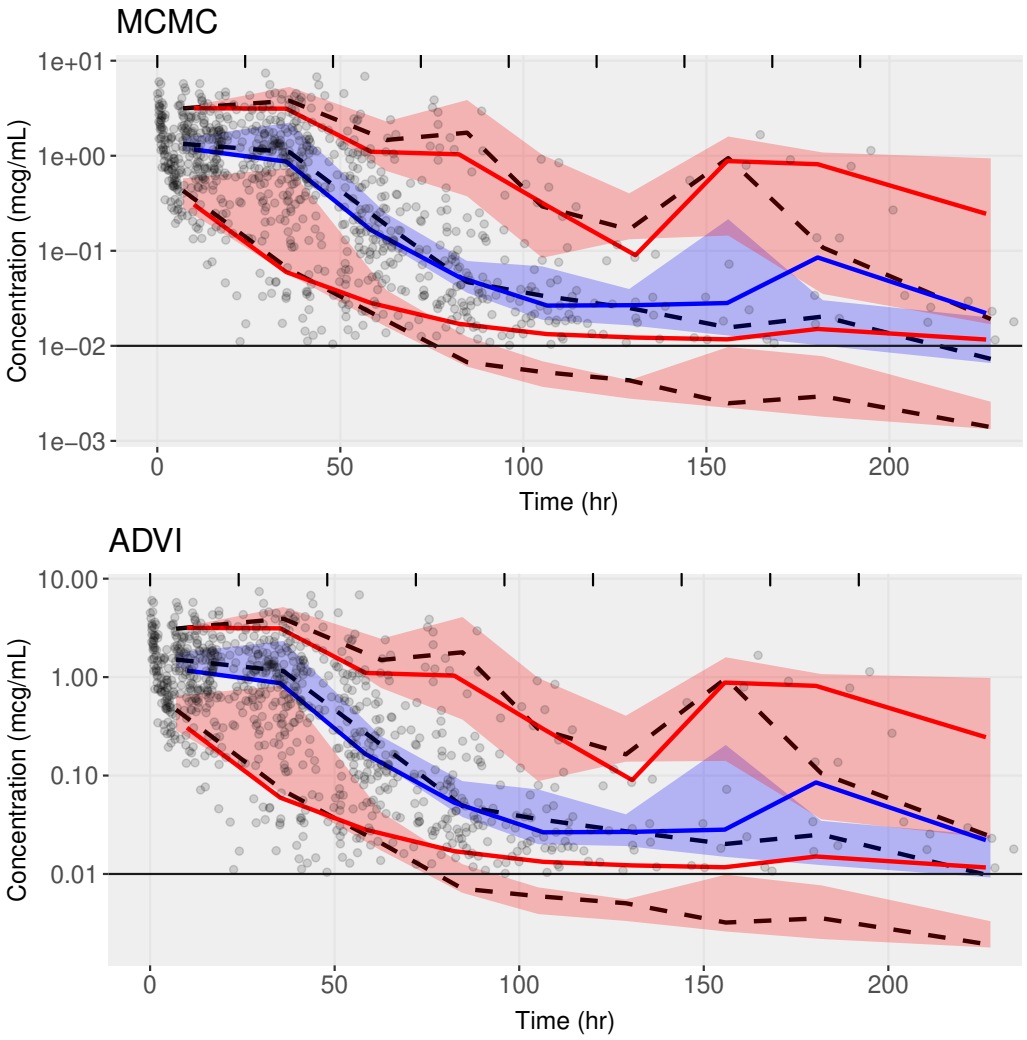

Figure 3.38: Visual predictive check for ketorolac PK model with fixed allometric scaling and age maturation parameters. Observed values for 10th, 50th and 90th percentile (solid lines) and theoretical values (dashed lines) along with 90% prediction interval for theoretical percentiles (shaded regions) are shown. Filled circles indicate observed values. The lower limit of quantification of 0.01 mcg/mL is represented by a horizontal gray line. Black tickmarks along the top of the plot indicate bin boundaries.

3.5 Discussion

ADVI can provide a good balance between speed and accuracy when fitting Bayesian population PK models, but involves trade-offs. The greatest advantage of ADVI is reduced computation time. In the simulation study, total MCMC fit times were around 10 times longer than ADVI fit times for one-compartment models and more than 45 times longer for two-compartment models. For the ketorolac case study, the ADVI models took around 6.5 minutes to fit compared to nearly 6 hours using MCMC. The reduced time to fit models with ADVI can be leveraged for model selection using cross-validation. Specifically, the model selection strategy using 5-fold leave-subject-out cross-validation with ADVI was more accurate than using either deviance information criteria, the information criteria provided by the `Stan` package `loo` (WAIC and PSIS-LOOIC) or the maximized variational objective function (ELBO). Importantly, we found that $lppd_{cv}$ calculated using 5-fold cross-validation with ADVI estimation was a competitive strategy across scenarios, sampling schemes, and sensitivity analyses. In addition, while both ADVI and MCMC estimation can fail to converge, the decreased ADVI fit time means models with poor parameterization or initialization values can be interrogated and discarded faster.

For population parameters, posterior median estimates can be fairly accurate, but credible intervals are underestimated compared to MCMC. This result is similar to those in a recent analysis by Park et al. (2021) comparing ADVI and MCMC for two simple Bayesian PK models[2] [96]. Underestimation of posterior parameter variability is a well-known issue for mean-field VI methods [15,97–99]. One possible solution is to expand the variational family by relaxing the mean-field assumption, however more flexible variational families (e.g., full-rank VI which allows correlation between variational parameters) can result in more difficult optimization and longer fit times. In addition, some previous work has shown that a more complex variational family may not necessarily outperform a simpler one [98,99]. Giordano et al. (2018) present a method to correct the variance estimates of mean-field VI using "linear response methods" from statistical physics [99]. The linear response variational Bayes (LR-ADVI) covariance estimates show good performance in matching MCMC. While not as computational intensive as MCMC methods, LR-ADVI requires more calculation than ADVI estimates. LR-ADVI also assumes that the mean-field VI correctly approximates the true posterior means; as seen in the simulation study this assumption may be incorrect for nonlinear population PK models.

Compared to population parameters, individual-specific posterior median and credible interval estimates are generally more similar between estimation methods. Again, this result mirrors those in the Park et al. (2021) population PK model, with individual-specific posterior distributions for total clearance closely matching using both ADVI and MCMC [96]. However there are more discrepancies between ADVI and MCMC

[2]Park et al. consider a one-compartment oral administration model for a single subject and a hierarchical one-compartment oral administration model using 12 subjects with dense sampling and no covariates

individual-level estimates for two-compartment models than one-compartment models. Beyond parameter estimates, posterior predicted concentrations produced using ADVI appear similar to those produced with MCMC, especially for new subjects. The trend of somewhat better estimation for population predictions was also seen in model selection. This is likely because the criteria used in Strategy 1 are focused on predictions for subjects in the current dataset (individual predictions) while the criteria for Strategy 3 is focused on predictions for new, unobserved subjects (population predictions). Concentration predictions are useful for personalized dosing (e.g., maintaining concentration in a known therapeutic range) or as the input to a pharmacodynamic model (e.g., modeling the relationship between predicted concentration and a physiological outcome such as pain response or sedation). Individual and population predictions and individual random effects are also used in common PK graphical model checks. In the case study, these graphical checks were similar for both estimation methods.

The conclusions regarding predicted concentrations are based on visualization of a small number of simulated subjects' PK profiles (Figures 3.20 - 3.25) along with comparisons of the maximum concentration (C_{max}) for all subjects (Figures 3.26 - 3.29). The individual plots show where the two estimated PK profiles for a subject are similar or different, especially in relationship to dosing and observed concentration data, but there are too many simulated subjects and datasets (e.g., $250 \times 100 = 25{,}000$ plots for each prior and sampling scheme) to examine all the individual plots. The plots showing ratios of C_{max} medians and credible interval widths allow comparisons that include all the simulation data, but are focused on a single feature of the PK profile which may not be relevant for all drugs. More research is needed to define and assess criteria which quantify the differences between methods in relation to predicted concentrations.

Results were generally similar across the four priors, with more biased parameter estimates using non-informative and misspecified priors for some settings. The non-informative priors also had more extreme outliers, indicating computational problems, especially when used with sparse sampling. Relative fit time for MCMC vs. ADVI was also larger for sparse sampling with non-informative priors. Improper non-informative priors can result in improper posteriors [62] or more difficult estimation. In our simulations, this may be related to the use of (improper) non-informative priors for all parameters simultaneously. An alternate approach would define non-informative priors for only some of the parameters while using informative priors for other parameters. While there may be a desire to "let the data speak" without external evidence, for population PK models some level of prior information is available in nearly all cases. For example, a non-informative prior for CL_{pop} implies that all positive values are equally plausible (in a given parameterization), but a value of 100,000 L/h or 100,000,000 L/h is not reasonable. Our recommendation is to begin with weakly informative, but proper, priors which give extremely low probability to implausible values. If the resulting posterior does

not produce reasonable results (scientific or computational), then the model or prior should be adjusted. A potential concern is that the analysis may be overly influenced by a subjective prior choice, however this disregards the important role of model checking and iterative model development in Bayesian analysis. As stated by Box and Tiao (1992), "inferences that are unacceptable *must* come from inappropriate assumption and not from inadequacies of the inferential system. Thus all parts of the model, including the prior distribution, are exposed to appropriate criticism" [67]. An analogous situation involving non-robustness or subjective influence can occur under any statistical paradigm if two investigators with different assumptions arrive at different results. The prior need not be considered any more subjective than other modeling assumptions.

There are several caveats for this study. First, the `Torsten` library (version 0.88) used for the PK models and the ADVI algorithm implemented using the `CmdStanR` interface (version 0.4.0) to `CmdStan` (version 2.25.0) are both in relatively early stages of development. The authors of `Torsten` include the warning, "The current version of Torsten is a prototype. It is being released for review and comment, and to support limited research applications. It has not been rigorously tested and should not be used for critical applications without further testing or cross-checking by comparison with other methods." Similarly, the ADVI algorithm includes the message, "This procedure has not been thoroughly tested and may be unstable or buggy." In addition, the results of the simulation study are dependent on the specific parameterization used for the PK models. For example, the structural two-compartment model can be parameterize in terms of clearances and apparent volumes of distribution (Cl, Q, V_1, V_2; used for this study), micro-rate constants ($V = V_1$, $k = Cl/V_1$, $k_{12} = Q/V_1$, $k_{21} = k_{12}V_1/V2$) and macro-rate constants (α,β,A,B in section 3.6.2.2) [100,101]; the random effects variance-covariance matrix was decomposed using standard deviations and a correlation matrix and a non-centered parameterization [102]; and there are several ways to include proportional and additive residual error components [103]. Reparameterization changes the geometry of the posterior parameter space and could affect estimation as a result. We recommend using a more thoroughly understood estimation technique, such as MCMC, to confirm final model results and obtain parameter estimates.

In addition to VI, many other approximations could be considered to reduce computation time for Bayesian population PK analysis. Although formal evaluation of other approaches is beyond the scope of this work, we briefly comment on a few. Maximum a posteriori (MAP) estimation [104] provides posterior estimates quickly, but has several issues compared to ADVI. First, MAP provides only a point estimate of the posterior distribution mode without quantifying or propagating estimate uncertainty. In addition, Maier et al. (2020) demonstrate that in a nonlinear model, MAP parameter estimates do not translate to the most probable value on the scale of observed data [105]. Giordano, Broderick, and Jordan (2018) [99] also show that MAP estimates can perform poorly for certain components in hierarchical models. Another potential solution is the Laplace

approximation. Both VI and the Laplace approximation require optimization and use the inverse Hessian of the optimization objective function, however VI uses a more expressive approximating distributional family [99]. Opper and Archambeau (2009) show that Gaussian VI is not equivalent to the Laplace approximation [106].

For this chapter, only two relatively simple population PK scenarios were explored using mean-field ADVI. Further work is required to assess extensions such as full-rank ADVI or LR-ADVI and more complex situations which may benefit from ADVI, such as combined pharmacokinetic-pharmacodynamic models or PK models without analytic solutions, such as physiologically based PK (PBPK) models [107]. For example, PBPK analysis attempts to model drug kinetics mechanistically by using a large number of compartments with underlying physiological meanings (e.g., fat, liver, kidneys, stomach, muscle, etc.). While PBPK models provide much more detailed information about the PK processes in the body, they are also more complicated and time-consuming to estimate especially under the Bayesian paradigm [108]. Thus a faster estimation method such as ADVI could be of benefit, but poor performance of the approximate estimates could result due to the increased complexity of the PBPK model.

ADVI is a valuable estimation technique for Bayesian population PK modeling. It is best used when fitting many models, such as during model development and refinement or for model selection using cross-validation, or when the goal is to evaluate models quickly and approximate results are acceptable.

3.6 Appendix

3.6.1 Diagnostics for Variational Inference

3.6.1.1 Pareto Smoothed Importance Sampling $\hat{k}$

The PSIS $\hat{k}$ is based on the idea that the VI approximation (or a Bayesian MCMC estimate) can be improved by using importance sampling [86]. Using $s = 1, \ldots, S$ samples from the optimal VI density $q^*(\theta^s)$ as proposal samples along with the joint density $p(\theta^s, y)$, the importance sampling estimate of an integrable function $E_p[h(\theta)] \approx \frac{\sum_{s=1}^{S} h(\theta^s) r^s}{\sum_{s=1}^{S} r^s}$ where the importance ratios are defined as $r^s = \frac{p(\theta^s, y)}{q^*(\theta^s)}$. PSIS stabilizes the importance sampling by fitting a generalized Pareto distribution to the largest r_i, which are then replaced by their expectation under the fitted distribution. The estimated shape parameter $\hat{k}$ provides a diagnostic measure between the posterior $p(\theta|y)$ and $q^*(\theta)$.

3.6.1.2 Variational Simulation Based Calibration

VSBC is a method to assess the calibration of a VI point estimate. It works by simulating data from the prior predictive distribution, fitting a VI approximation to the simulated data and comparing the posterior parameter estimates to the priors. First, samples drawn from the prior distribution $\theta_m^{(0)} \sim p(\theta)$ are used to generate $m = 1, \ldots, M$ datasets under the specified model likelihood $y_m \sim p(y|\theta_m^{(0)})$. Recall that θ (and $\theta_m^{(0)}$) are j-dimensional vectors; $p(\theta_j)$ is the j^{th} marginal parameter prior and $\theta_{jm}^{(0)}$ is the m^{th} draw from $p(\theta_j)$. For each simulated dataset, the VI approximation to the posterior, $p(\theta|y_m)$, is found and marginal calibration probabilities are calculated for each parameter: $p_{jm} = Pr_{\theta|y_m}(\theta_j \leq \theta_{jm}^{(0)})$. Using $s = 1, \ldots, S$ samples from the posterior VI approximation $p_{jm} \approx \frac{1}{S} \sum_{s=1}^{S} I[\theta_j^s \leq \theta_{jm}^{(0)}]$. Asymmetry in the distribution of $p_{j:}$ indicates bias in the VI approximation to the marginal posterior $\theta_j|y.$ and can be assessed visually or using a Kolmogorov-Smirnov test for equality between the distribution $p_{j:}$ and $1 - p_{j:}$.

3.6.2 Simulation Study

3.6.2.1 Scenario 1: One-compartment Intravenous Infusion Model

Stage 1:

$$f_1(D=\{d,t_d,t_{inf}\},\psi=\{Cl,V\},t) = \begin{cases} \frac{d}{t_{inf}}\frac{1}{Cl}\left(1-e^{-\frac{Cl}{V}(t-t_d)}\right) & \text{if } t-t_d \leq t_{inf} \\ \frac{d}{t_{inf}}\frac{1}{Cl}\left(1-e^{-\frac{Cl}{V}t_{inf}}\right)e^{-\frac{Cl}{V}(t-t_d-t_{inf})} & \text{otherwise} \end{cases}$$

$$y_{ij} = f_1(D_i,\psi_i,t_{ij}) \times (1+\varepsilon_{ij}^{prop}) + \varepsilon_{ij}^{add}$$
$$\varepsilon_{ij}^{prop} \sim N(0,\sigma^2_{prop}) \quad \varepsilon_{ij}^{add} \sim N(0,\sigma^2_{add})$$

Stage 2:

$$\log(Cl_i) = \log(Cl_{pop}) + 0.75 \times \log\left(\frac{weight_i}{70}\right) + \beta_{Cl_{X_1}} X_{1i} + \eta_{Cl_i},\ \eta_{Cl} \sim N(0,\omega^2_{Cl})$$
$$\log(V_i) = \log(V_{pop}) + 1 \times \log\left(\frac{weight_i}{70}\right) + \eta_{V_i},\ \eta_V \sim N(0,\omega^2_V)$$

Stage 3:

$$Cl_{pop} = 40\ \ V_{pop} = 120\ \ \beta_{Cl_{X_1}} = 0.4\ \ \sigma^2_{add} = 0.001\ \ \sigma^2_{prop} = 0.025$$
$$\Omega = \begin{bmatrix} \omega^2_{Cl} = 0.1 & \omega_{Cl,V} = 0.05 \\ & \omega^2_V = 0.2 \end{bmatrix}$$

Covariates:

$$weight \sim Unif(3,80)$$
$$X_1 \sim Bernoulli(\theta = 0.4)$$

Table 3.1: Scenario 1 Priors

Parameter	Strong	Weak	Misspecified	Non-Informative
Cl_{pop}	$logNormal(\mu = \log(40) + 0.3^2, \sigma = 0.3)$	$logNormal(\mu = \log(40) + 1^2, \sigma = 1)$	$Gamma(\alpha = 10, \beta = 0.1)$	$Uniform(0, \infty)$
V_{pop}	$logNormal(\mu = \log(120) + 0.3^2, \sigma = 0.3)$	$logNormal(\mu = \log(120) + 1.1^2, \sigma = 1.1)$	$Gamma(\alpha = 15, \beta = 0.055)$	$Uniform(0, \infty)$
$\beta_{Cl_{X_1}}$	$Normal(\mu = 0.4, \sigma = 1)$	$Normal(\mu = 0, \sigma = 10)$	$Normal(\mu = -3.25, \sigma = 2)$	$Uniform(-\infty, \infty)$
ω_{Cl}	$logNormal(\mu = \log(\sqrt{0.1}) + 0.2^2, \sigma = 0.2)$	$logNormal(\mu = \log(\sqrt{0.1}) + 1.1^2, \sigma = 1.1)$	$Gamma(\alpha = 8, \beta = 10)$	$Uniform(0, \infty)$
ω_V	$logNormal(\mu = \log(\sqrt{0.2}) + 0.2^2, \sigma = 0.2)$	$logNormal(\mu = \log(\sqrt{0.2}) + 1.1^2, \sigma = 1.1)$	$Gamma(\alpha = 9, \beta = 8)$	$Uniform(0, \infty)$
ρ	$LKJ(\eta = 0.8)$	$LKJ(\eta = 1)$	$LKJ(\eta = 1)$	$LKJ(\eta = 1)$
σ_{prop}	$Gamma(\alpha = 1 + 4\sqrt{0.025}, \beta = 4)$	$Half - Cauchy(\mu = 0, \sigma = 1)$	$Gamma(\alpha = 5, \beta = 4)$	$Uniform(0, \infty)$
σ_{add}	$Gamma(\alpha = 1 + 7\sqrt{0.001}, \beta = 7)$	$Half - Cauchy(\mu = 0, \sigma = 0.5)$	$Gamma(\alpha = 4, \beta = 5)$	$Uniform(0, \infty)$

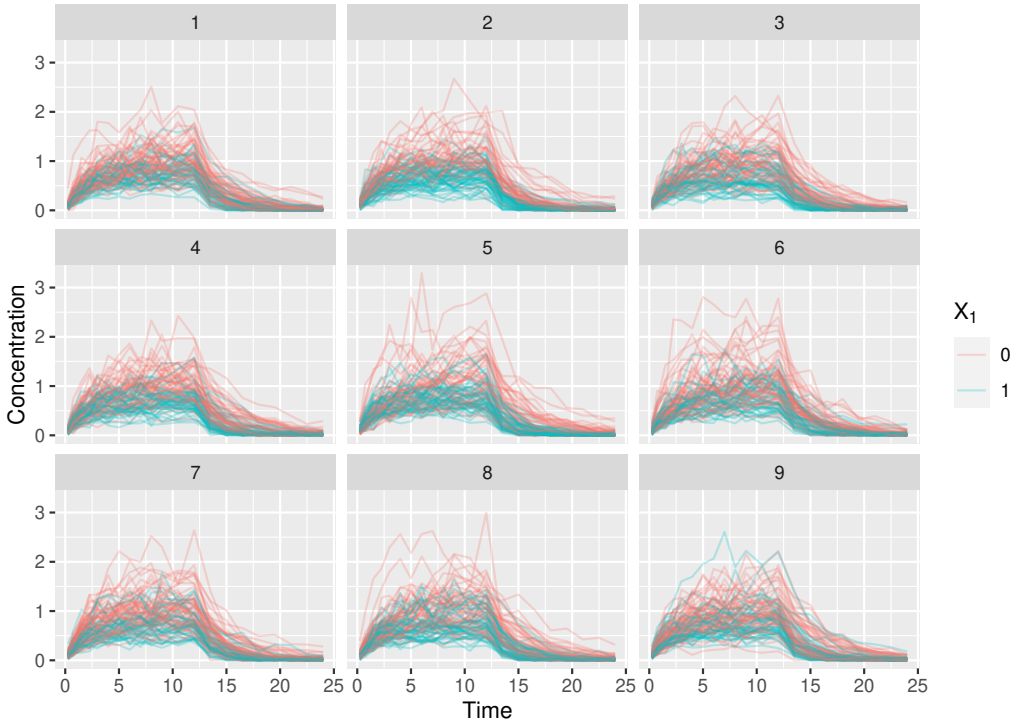

Figure 3.39: Nine simulation datasets with dense sampling for one-compartment model

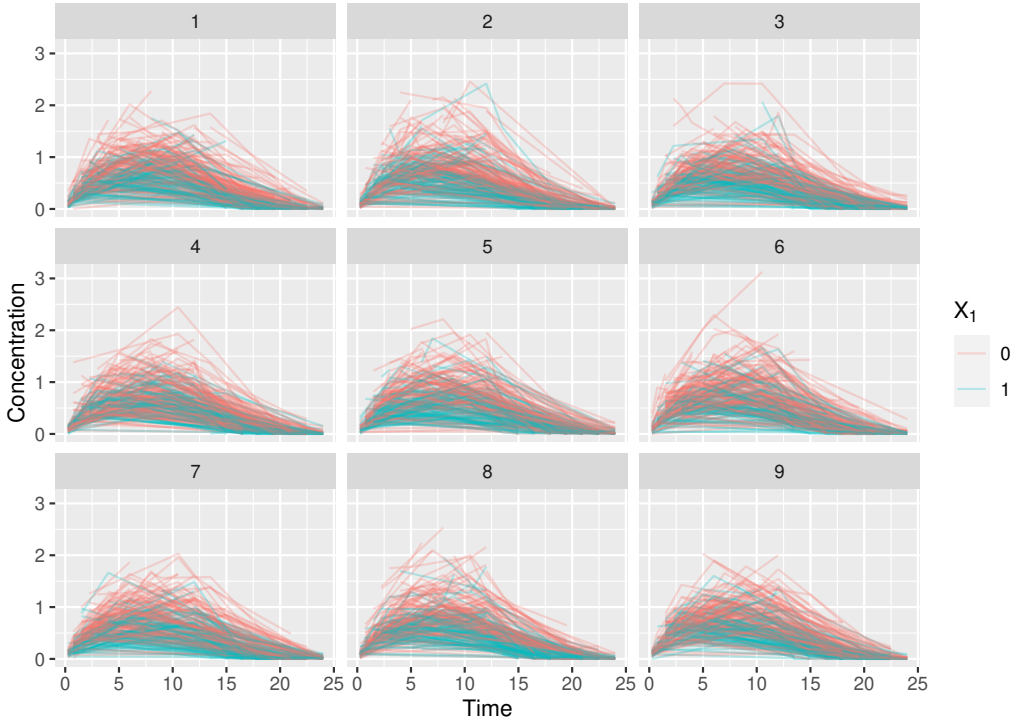

Figure 3.40: Nine simulation datasets with sparse sampling for one-compartment model

3.6.2.2 Scenario 2: Two-compartment Intravenous Infusion Model

Stage 1:

$$D = \{d, t_d, t_{inf}\} \quad \psi = \{Cl, Q, V_1, V_2\}$$

$$\alpha = \frac{\frac{QCl}{V_2 V_1}}{\beta} \quad \beta = \frac{1}{2}\left[\frac{Q}{V_1} + \frac{Q}{V_2} + \frac{Cl}{V_1} - \sqrt{\left(\frac{Q}{V_1} + \frac{Q}{V_2} + \frac{Cl}{V_1}\right)^2 - 4\frac{QCl}{V_2 V_1}}\right] \quad A = \frac{1}{V_1}\frac{\alpha - \frac{Q}{V_2}}{\alpha - \beta} \quad B = \frac{1}{V_1}\frac{\beta - \frac{Q}{V_2}}{\beta - \alpha}$$

$$f_2(D, \psi, t) = \begin{cases} \frac{d}{t_{inf}}\left[\frac{A}{\alpha}(1 - e^{-\alpha(t - t_d)}) + \frac{B}{\beta}(1 - e^{-\beta(t - t_d)})\right] & \text{if } t - t_d \leq t_{inf} \\ \frac{d}{t_{inf}}\left[\frac{A}{\alpha}(1 - e^{-\alpha t_{inf}})e^{-\alpha(t - t_d - t_{inf})} + \frac{B}{\beta}(1 - e^{-\beta t_{inf}})e^{-\beta(t - t_d - t_{inf})}\right] & \text{otherwise} \end{cases}$$

$$y_{ij} = f_2(D_i, \psi_i, t_{ij}) \times (1 + \varepsilon_{ij}^{prop}) + \varepsilon_{ij}^{add}$$

$$\varepsilon_{ij}^{prop} \sim N(0, \sigma_{prop}^2) \quad \varepsilon_{ij}^{add} \sim N(0, \sigma_{add}^2)$$

Stage 2:

$$\log(Cl_i) = \log(Cl_{pop}) + 0.75 \times \log\left(\frac{weight_i}{70}\right) + \beta_{Cl_{X_1}} X_{1i} + \eta_{Cl_i},\ \eta_{Cl} \sim N(0, \omega_{Cl}^2)$$

$$\log(Q_i) = \log(Q_{pop}) + 0.75 \times \log\left(\frac{weight_i}{70}\right) + \eta_{Q_i},\ \eta_Q \sim N(0, \omega_Q^2)$$

$$\log(V_{1i}) = \log(V_{1\,pop}) + 1 \times \log\left(\frac{weight_i}{70}\right) + \eta_{V_{1i}},\ \eta_{V_1} \sim N(0, \omega_{V_1}^2)$$

$$\log(V_{2i}) = \log(V_{2\,pop}) + 1 \times \log\left(\frac{weight_i}{70}\right) + \eta_{V_{2i}},\ \eta_{V_2} \sim N(0, \omega_{V_2}^2)$$

Stage 3:

$$Cl_{pop} = 35 \ \ Q_{pop} = 60 \ \ V_{1\,pop} = 50 \ \ V_{2\,pop} = 70 \ \ \beta_{Cl_{X_1}} = 0.4 \ \ \sigma_{add}^2 = 0.001 \ \ \sigma_{prop}^2 = 0.025$$

$$\Omega = \begin{bmatrix} \omega_{Cl}^2 = 0.2 & \omega_{Cl,Q} = -0.02 & \omega_{Cl,V_1} = -0.01 & \omega_{Cl,V_2} = -0.05 \\ & \omega_Q^2 = 0.4 & \omega_{Q,V_1} = 0.3 & \omega_{Q,V_2} = 0.1 \\ & & \omega_{V_1}^2 = 1.1 & \omega_{V_1,V_2} = 0.2 \\ & & & \omega_{V_2}^2 = 0.3 \end{bmatrix}$$

Covariates:

$$weight \sim Unif(3, 80)$$

$$X_1 \sim Bernoulli(\theta = 0.4)$$

Table 3.2: Scenario 2 Priors

Parameter	Strong	Weak	Misspecified	Non-Informative
Cl_{pop}	$logNormal(\mu = \log(35) + 0.3^2, \sigma = 0.3)$	$logNormal(\mu = \log(35) + 0.9^2, \sigma = 0.9)$	$Gamma(\alpha = 8, \beta = 0.08)$	$Uniform(0, \infty)$
Q_{pop}	$logNormal(\mu = \log(60) + 0.3^2, \sigma = 0.3)$	$logNormal(\mu = \log(60) + 0.9^2, \sigma = 0.9)$	$Gamma(\alpha = 10, \beta = 0.06)$	$Uniform(0, \infty)$
$V1_{pop}$	$logNormal(\mu = \log(50) + 0.3^2, \sigma = 0.3)$	$logNormal(\mu = \log(50) + 0.9^2, \sigma = 0.9)$	$Gamma(\alpha = 9, \beta = 0.057)$	$Uniform(0, \infty)$
$V2_{pop}$	$logNormal(\mu = \log(70) + 0.3^2, \sigma = 0.3)$	$logNormal(\mu = \log(70) + 0.9^2, \sigma = 0.9)$	$Gamma(\alpha = 11, \beta = 0.063)$	$Uniform(0, \infty)$
$\beta_{Cl_{X_1}}$	$Normal(\mu = 0.4, \sigma = 1)$	$Normal(\mu = 0, \sigma = 10)$	$Normal(\mu = -3.25, \sigma = 2)$	$Uniform(-\infty, \infty)$
ω_{Cl}	$logNormal(\mu = \log(\sqrt{0.2}) + 0.2^2, \sigma = 0.2)$	$logNormal(\mu = \log(\sqrt{0.2}) + 0.8^2, \sigma = 0.8)$	$Gamma(\alpha = 8, \beta = 10)$	$Uniform(0, \infty)$
ω_{Q}	$logNormal(\mu = \log(\sqrt{0.4}) + 0.3^2, \sigma = 0.3)$	$logNormal(\mu = \log(\sqrt{0.4}) + 0.9^2, \sigma = 0.9)$	$Gamma(\alpha = 11, \beta = 8)$	$Uniform(0, \infty)$
ω_{V1}	$logNormal(\mu = \log(\sqrt{1.1}) + 0.3^2, \sigma = 0.3)$	$logNormal(\mu = \log(\sqrt{1.1}) + 0.9^2, \sigma = 0.9)$	$Gamma(\alpha = 12, \beta = 5)$	$Uniform(0, \infty)$
ω_{V2}	$logNormal(\mu = \log(\sqrt{0.3}) + 0.2^2, \sigma = 0.2)$	$logNormal(\mu = \log(\sqrt{0.3}) + 0.8^2, \sigma = 0.8)$	$Gamma(\alpha = 8, \beta = 6)$	$Uniform(0, \infty)$
ρ	$LKJ(\eta = 0.8)$	$LKJ(\eta = 1)$	$LKJ(\eta = 1)$	$LKJ(\eta = 1)$
σ_{prop}	$Gamma(\alpha = 1 + 4\sqrt{0.025}, \beta = 4)$	$Half - Cauchy(\mu = 0, \sigma = 1)$	$Gamma(\alpha = 5, \beta = 4)$	$Uniform(0, \infty)$
σ_{add}	$Gamma(\alpha = 1 + 7\sqrt{0.001}, \beta = 7)$	$Half - Cauchy(\mu = 0, \sigma = 0.5)$	$Gamma(\alpha = 4, \beta = 5)$	$Uniform(0, \infty)$

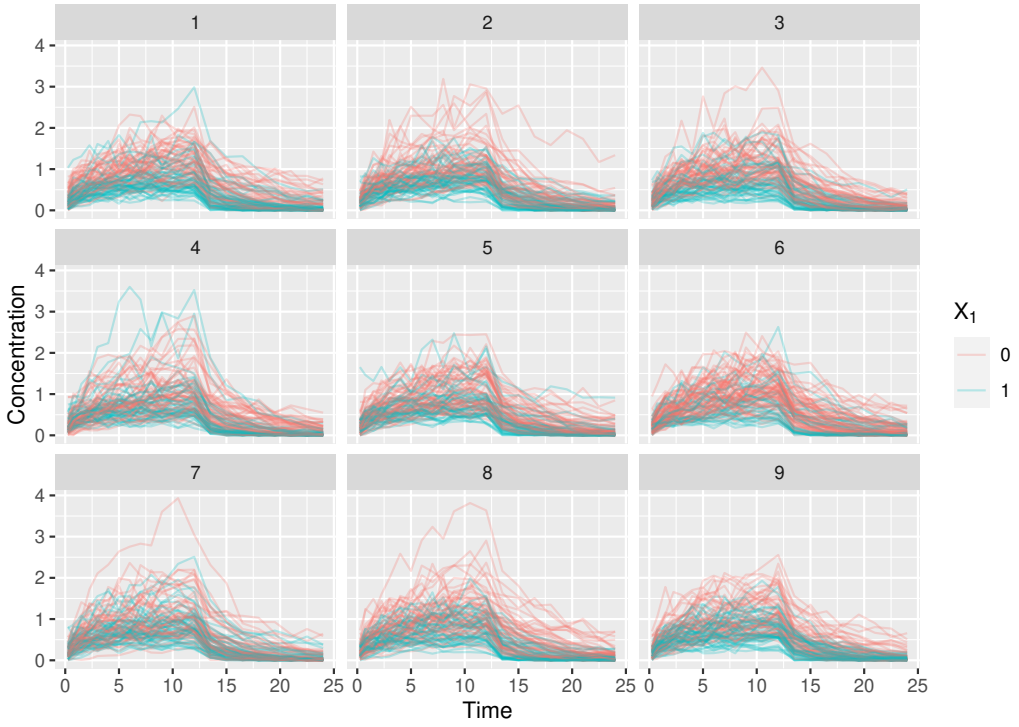

Figure 3.41: Nine simulation datasets with dense sampling for two-compartment model

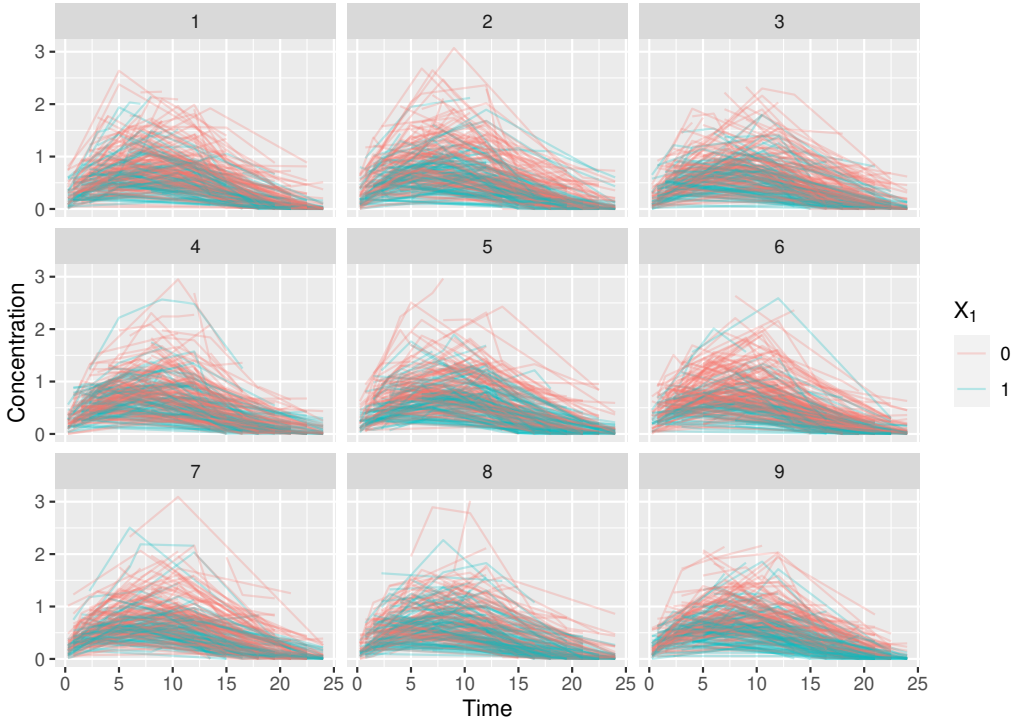

Figure 3.42: Nine simulation datasets with sparse sampling for two-compartment model

3.6.2.3 Estimation Procedure Settings

3.6.2.4 Model Parameterization

The variance-covariance matrix for the parameter random effects, Ω, is parameterized using a decomposition into standard deviations and a correlation matrix [109]; $\Omega = \begin{bmatrix} \omega_{Cl} & 0 \\ 0 & \omega_V \end{bmatrix} \times \rho \times \begin{bmatrix} \omega_{Cl} & 0 \\ 0 & \omega_V \end{bmatrix}$. This decomposition allows prior information about the variability of population PK parameters to be specified on the same scale as the parameter. In addition, a non-centered parameterization is used for efficiency and numerical stability [102]. In the `Stan` model implementation, a prior is specified on the Cholesky factor of the correlation matrix L_ρ where $\rho = L_\rho \times L'_\rho$ instead of directly on ρ. Note that for the LKJ prior $\eta < 1$ favors matrices with more correlation, while $\eta = 1$ is a uniform distribution over correlation matrices.

3.6.2.5 Automatic Differentiation Variational Inference Settings

The `cmdstanr $variational()` function was used for ADVI estimation. The main function was wrapped in a helper function to automatically retry the fit if the following occurred:

1. a non-zero return code indicating an error in the run

2. a message indicating "The maximum number of iterations is reached!"

3. a message indicating "The ELBO at a previous iteration is larger than the ELBO upon convergence!"

For each retry, the seed for the initialization function was changed. If a successful run did not result after more than 1/2 of the tries, the seed for the `$variational()` function was also changed for each remaining retry. ELBO values were extracted by diverting output from `cmdstanr $variational()` method to an external text file and then parsing the text file to find the ELBO at convergence (or after the maximum number of iterations and maximum number of tries to refit).

For the one-compartment simulations (scenario 1) the following settings were used:

- algorithm: "meanfield"
- init: an initialization function based on the true population parameters

```
init0 <- function(){
list(CL_pop = exp(rnorm(1, log(40), 0.1)),
     V_pop = exp(rnorm(1, log(120), 0.1)))
}
```

- maximum warmup adaptation iterations (adapt_iter): 50 (Default)
- convergence tolerance on the relative norm of the objective (tol_rel_obj): 0.01 (Default)
- maximum number of iterations (iter): 17500

- number of samples for Monte Carlo estimate of gradients (grad_samples): 3
- number of samples for Monte Carlo estimate of ELBO (elbo_samples): 100 (Default)
- evaluate ELBO every Nth iteration (eval_elbo): 100 (Default)
- the number of iterations between printed screen updates (refresh): 500
- maximum tries to refit: 5
- all other parameters used `$variational()` defaults

For the two-compartment simulations (scenario 2) the following settings were used:

- algorithm: “meanfield”
- init: an initialization function based on the true population parameters

```
init0 <- function(){
list(CL_pop = exp(rnorm(1, log(35), 0.1)),
     Q_pop = exp(rnorm(1, log(60), 0.1)),
     V1_pop = exp(rnorm(1, log(50), 0.1)),
     V2_pop = exp(rnorm(1, log(70), 0.1)))
}
```

- maximum warmup adaptation iterations (adapt_iter): 50 (Default)
- convergence tolerance on the relative norm of the objective (tol_rel_obj): 0.01 (Default)
- maximum number of iterations (iter): 20000
- number of samples for Monte Carlo estimate of gradients (grad_samples): 3
- number of samples for Monte Carlo estimate of ELBO (elbo_samples): 100 (Default)
- evaluate ELBO every Nth iteration (eval_elbo): 100 (Default)
- the number of iterations between printed screen updates (refresh): 500
- maximum tries to refit: 5 (10 for cross-validation models)
- all other parameters used `$variational()` defaults

For the case study the following settings were used:

- algorithm: “meanfield”
- init: an initialization function

```
# model 1
init0 <- function(){
list(CL_pop = exp(rnorm(1, log(11), 0.1)),
     Q_pop = exp(rnorm(1, log(1.5), 0.1)),
```

```
      V1_pop = exp(rnorm(1, log(3500), 0.1)),
      V2_pop = exp(rnorm(1, log(1500), 0.1)))
}
# model 2
init0 <- function(){
list(CL_pop = exp(rnorm(1, log(11), 0.4)),
      Q_pop = exp(rnorm(1, log(1.5), 0.4)),
      V1_pop = exp(rnorm(1, log(3500), 0.4)),
      V2_pop = exp(rnorm(1, log(1500), 0.4)),
      tm50 = rnorm(1, 40, 2),
      hill = exp(rnorm(1, log(3), 0.2)))
}
```

- maximum warmup adaptation iterations (adapt_iter): 50 (Default)
- convergence tolerance on the relative norm of the objective (tol_rel_obj): 0.005 (0.01 for cross-validation models)
- maximum number of iterations (iter): 10000
- number of samples for Monte Carlo estimate of gradients (grad_samples): 5
- number of samples for Monte Carlo estimate of ELBO (elbo_samples): 150
- evaluate ELBO every Nth iteration (eval_elbo): 100 (Default)
- the number of iterations between printed screen updates (refresh): 500
- maximum tries to refit: 5
- all other parameters used `$variational()` defaults

3.6.2.6 No U Turns Hamiltonian Monte Carlo Settings

The `cmdstanr` `$sample()` function was used for HMC estimation. The main function was wrapped in a helper function to automatically retry the fit if a non-zero return code resulted indicating an error in the run. For each retry, the seed for the initialization function was changed. If a successful run did not result after more than 1/2 of the tries, the seed for the `$sample()` function was also changed for each remaining retry.

For the one-compartment simulations (scenario 1) the following settings were used:

- algorithm: “hmc” (Default)
- engine: “nuts” (Default)
- init: an initialization function based on the true population parameters

```
init0 <- function(){
list(CL_pop = exp(rnorm(1, log(40), 0.1)),
     V_pop = exp(rnorm(1, log(120), 0.1)))
}
```

- number of warmup iterations to run per chain (iter_warmup): 1500
- number of post-warmup iterations to run per chain (iter_sampling): 250
- number of Markov chains to run (chains): 4
- maximum number of MCMC chains to run in parallel (parallel_chains): 4
- The number of iterations between printed screen updates (refresh): 500
- maximum tries to refit: 3
- all other parameters used `$sample()` defaults

For the two-compartment simulations (scenario 2) the following settings were used:

- algorithm: “hmc” (Default)
- engine: “nuts” (Default)
- init: an initialization function based on the true population parameters

```
init0 <- function(){
list(CL_pop = exp(rnorm(1, log(35), 0.1)),
     Q_pop = exp(rnorm(1, log(60), 0.1)),
     V1_pop = exp(rnorm(1, log(50), 0.1)),
     V2_pop = exp(rnorm(1, log(70), 0.1)))
}
```

- number of warmup iterations to run per chain (iter_warmup): 1750
- number of post-warmup iterations to run per chain (iter_sampling): 250
- number of Markov chains to run (chains): 4
- maximum number of MCMC chains to run in parallel (parallel_chains): 4
- The number of iterations between printed screen updates (refresh): 500
- maximum tries to refit: 3
- all other parameters used `$sample()` defaults

For the case study the following settings were used:

- algorithm: “meanfield”

- init: an initialization function

```
# model 1
init0 <- function(){
list(CL_pop = exp(rnorm(1, log(11), 0.1)),
     Q_pop = exp(rnorm(1, log(1.5), 0.1)),
     V1_pop = exp(rnorm(1, log(3500), 0.1)),
     V2_pop = exp(rnorm(1, log(1500), 0.1)))
}
# model 2
init0 <- function(){
list(CL_pop = exp(rnorm(1, log(11), 0.4)),
     Q_pop = exp(rnorm(1, log(1.5), 0.4)),
     V1_pop = exp(rnorm(1, log(3500), 0.4)),
     V2_pop = exp(rnorm(1, log(1500), 0.4)),
     tm50 = rnorm(1, 40, 2),
     hill = exp(rnorm(1, log(3), 0.2)))
}
```

- number of warmup iterations to run per chain (iter_warmup): 4000
- number of post-warmup iterations to run per chain (iter_sampling): 250
- number of Markov chains to run (chains): 4
- maximum number of MCMC chains to run in parallel (parallel_chains): 4
- The adaptation target acceptance statistic (adapt_delta): 0.9
- The number of iterations between printed screen updates (refresh): 500
- all other parameters used `$sample()` defaults

3.6.2.7 Variational Simulation Based Calibration

For VSBC $M = 1000$ datasets were simulated using the sparse or dense sampling scheme with strong, weak, or misspecified informative priors for the correctly specified model (M1) using the following procedure:

1. Make one draw from the defined prior distribution for each parameter and create a simulated data using `$sample()` and `fixed_param=TRUE`. Covariates and dosing history (but not concentration values) from a randomly selected simulation dataset are used for each VSBC dataset.

2. Find parameter estimates using the simulated data with ADVI estimation and draw approximate posterior samples.

3. Transform prior draw and approximate posterior samples for constrained parameters (e.g. CL_{pop}, V_{pop}, σ_{add}, etc.) using a log transformation.

4. Calculate the calibration probability as the proportion of the posterior samples that are less than or equal to the prior draw used to simulate the data.

5. Plot the histogram of $M = 1000$ calibration probabilities

3.6.2.8 Simulation Results

3.6.2.8.1 Scenario 1 Pareto Smoothed Importance Sampling $\hat{k}$

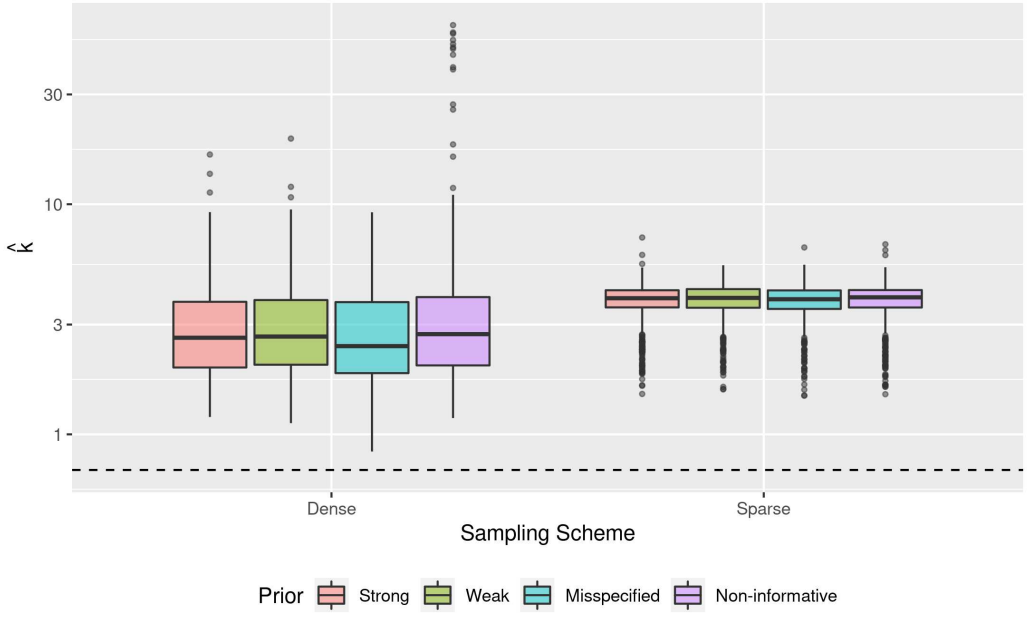

Figure 3.43: One-compartment PK model simulations - variational inference PSIS $\hat{k}$ values. Values above dashed line ($\hat{k} > 0.7$) indicate discrepancy between VI approximation and posterior distribution

3.6.2.8.2 Scenario 1 Variational Simulation Based Calibration

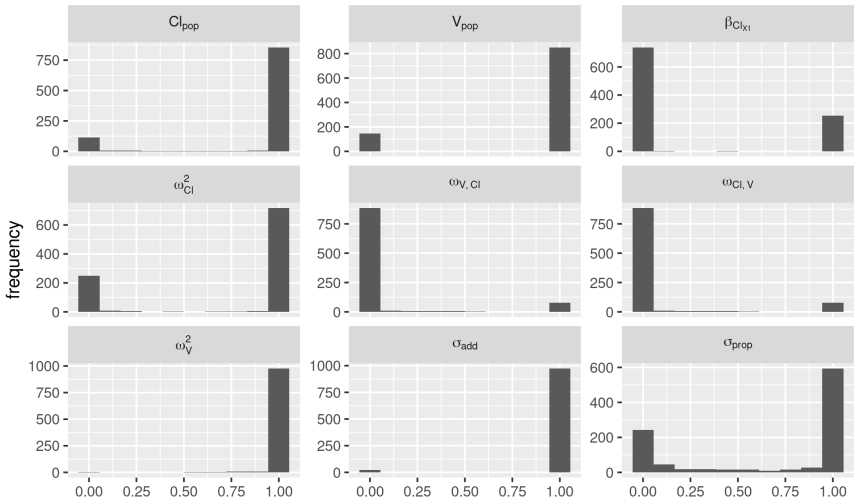

Figure 3.44: One-compartment PK model simulations - variational simulation based calibration for dense sampling and weak informative priors

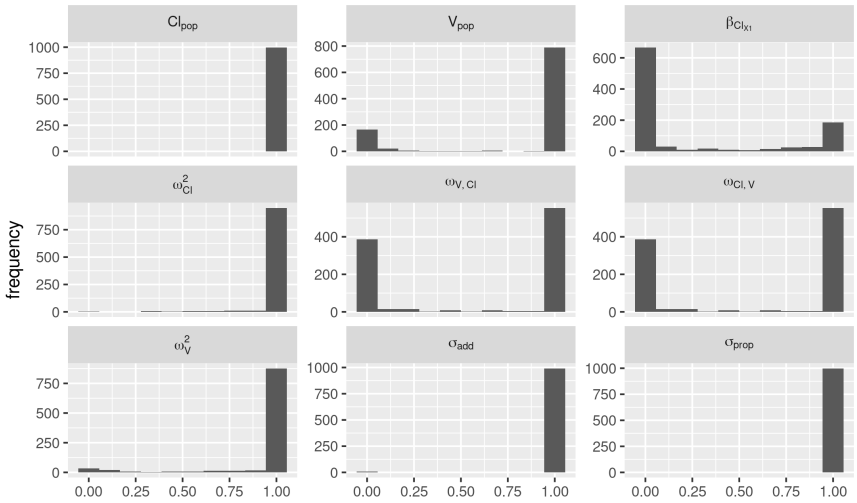

Figure 3.45: One-compartment PK model simulations - variational simulation based calibration for dense sampling and misspecified informative priors

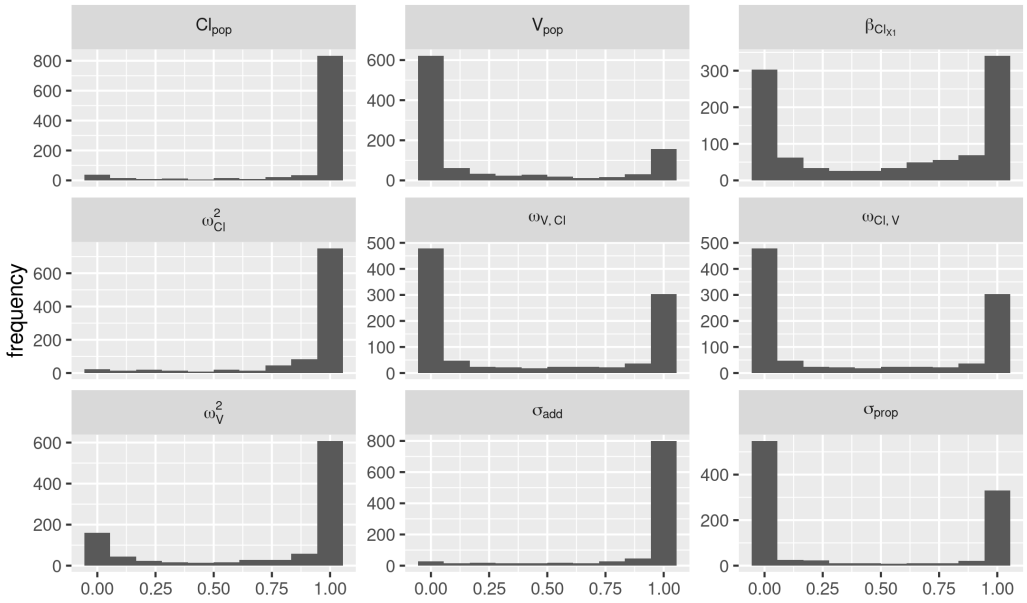

Figure 3.46: One-compartment PK model simulations - variational simulation based calibration for sparse sampling and strong informative priors

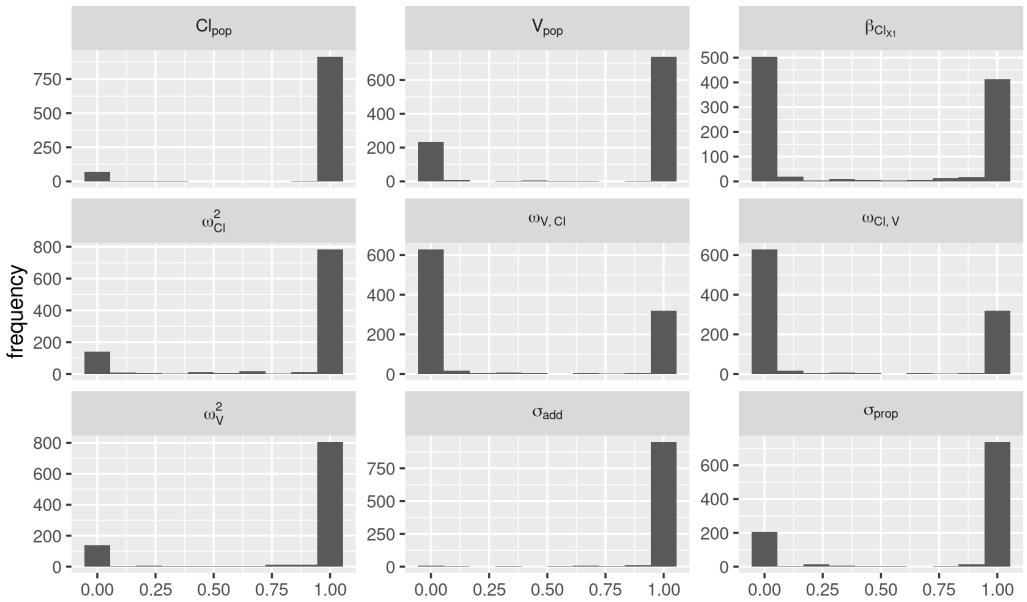

Figure 3.47: One-compartment PK model simulations - variational simulation based calibration for sparse sampling and weak informative priors

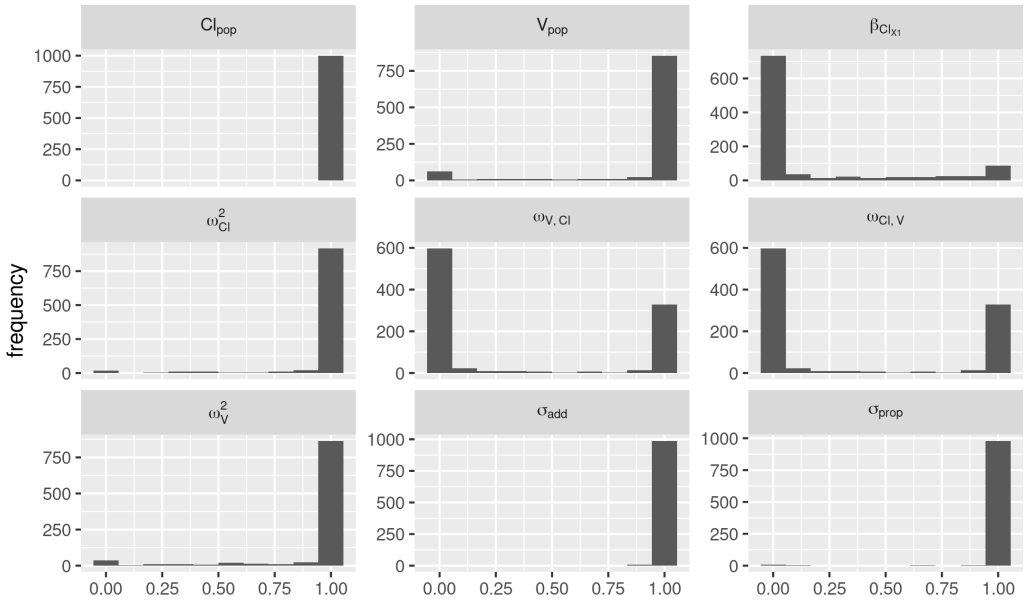

Figure 3.48: One-compartment PK model simulations - variational simulation based calibration for sparse sampling and misspecified informative priors

3.6.2.8.3 Scenario 1 Individual Concentration Predictions

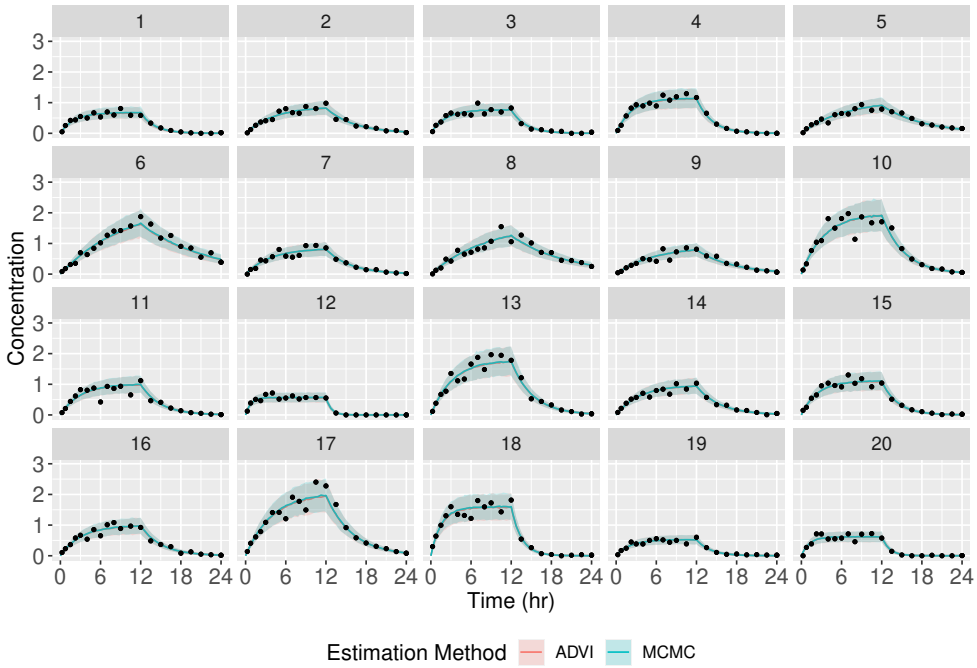

Figure 3.49: One-compartment PK model simulations - observed concentrations and individual concentration predictions for first simulated dataset with dense sampling and strong priors

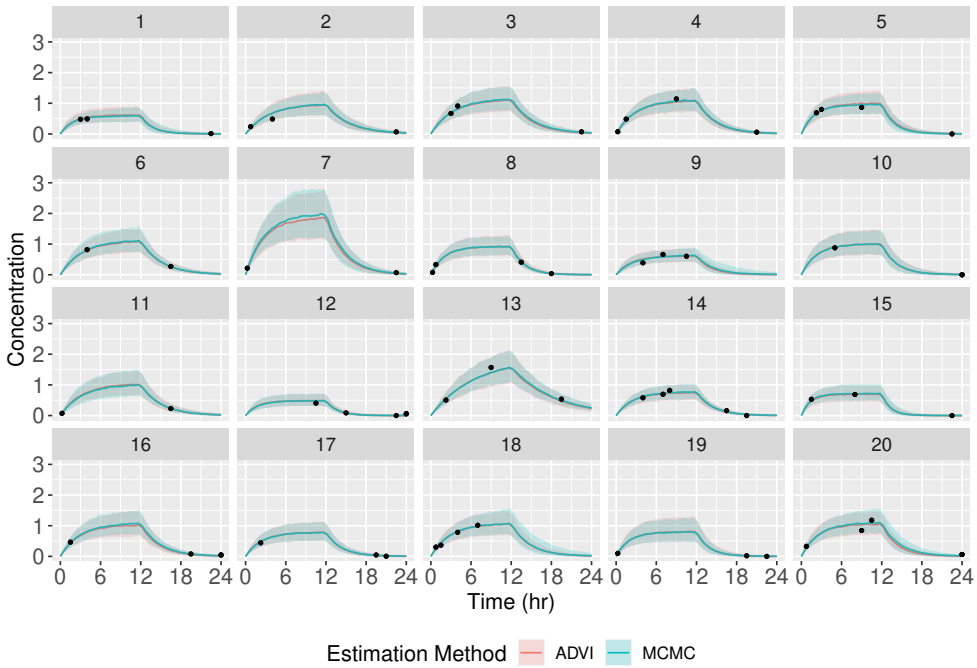

Figure 3.50: One-compartment PK model simulations - observed concentrations and individual concentration predictions for first simulated dataset with sparse sampling and strong priors

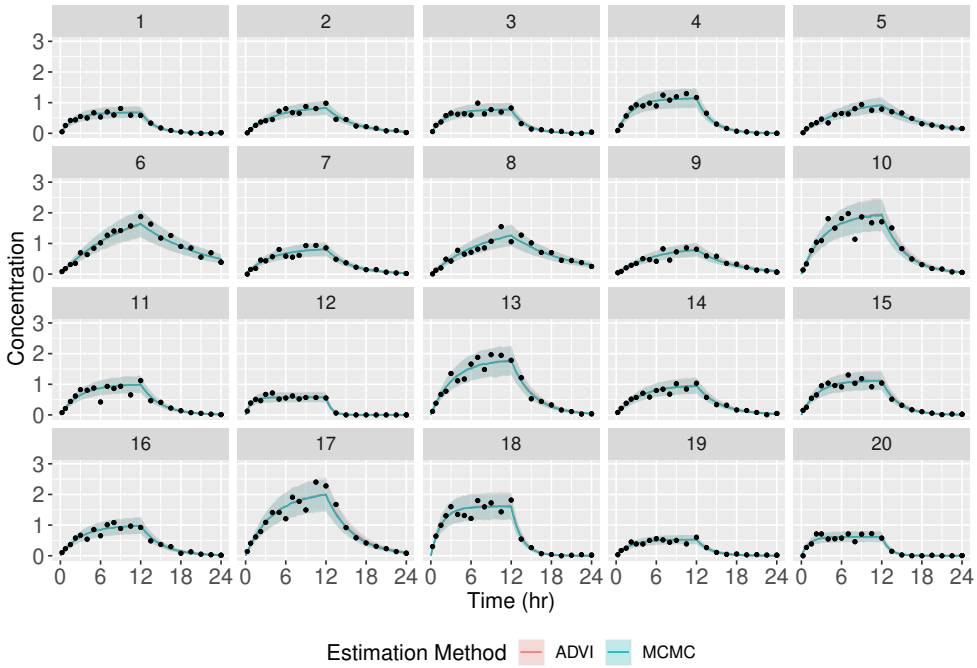

Figure 3.51: One-compartment PK model simulations - observed concentrations and individual concentration predictions for first simulated dataset with dense sampling and misspecified priors

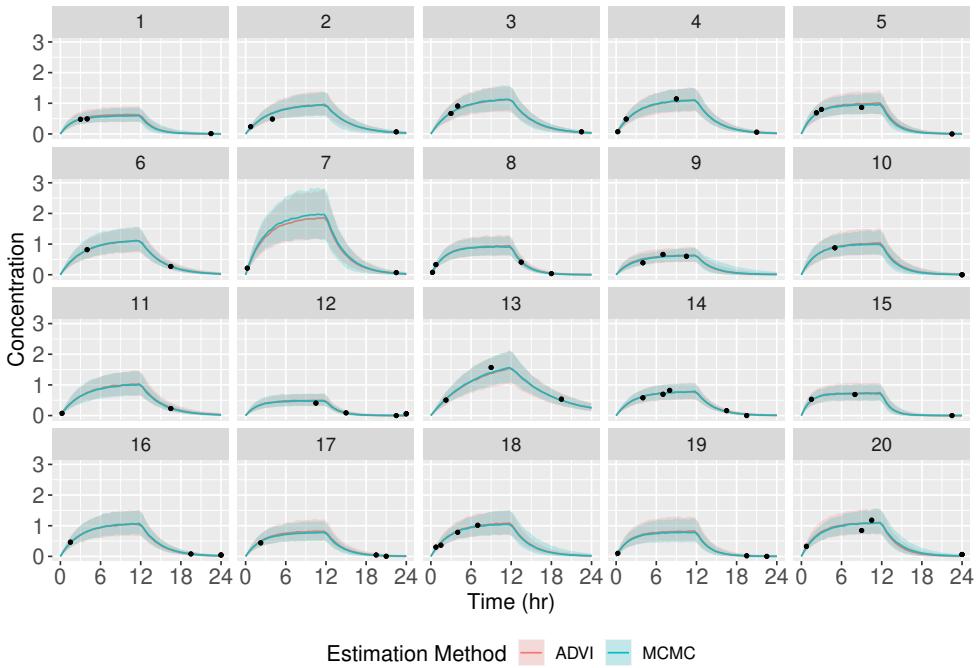

Figure 3.52: One-compartment PK model simulations - observed concentrations and individual concentration predictions for first simulated dataset with sparse sampling and misspecified priors

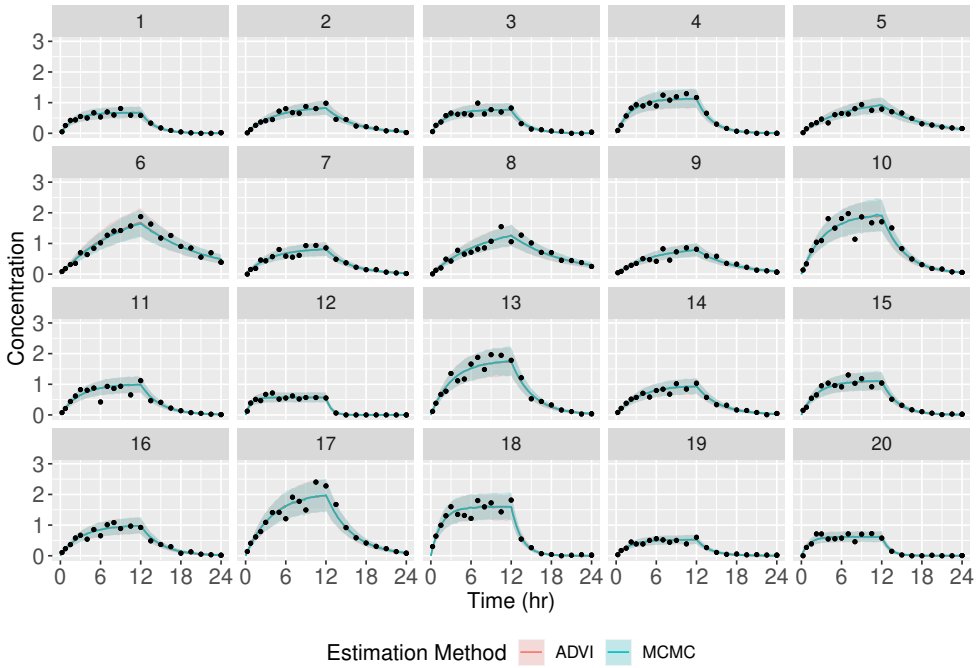

Figure 3.53: One-compartment PK model simulations - observed concentrations and individual concentration predictions for first simulated dataset with dense sampling and non-informative priors

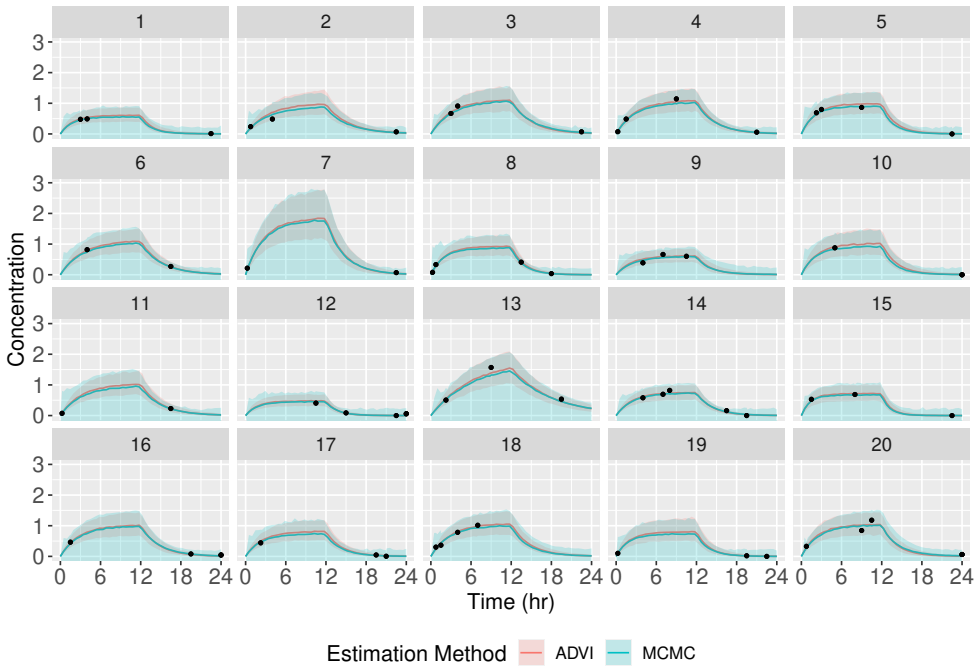

Figure 3.54: One-compartment PK model simulations - observed concentrations and individual concentration predictions for first simulated dataset with sparse sampling and non-informative priors

3.6.2.8.4 Scenario 1 Population Concentration Predictions

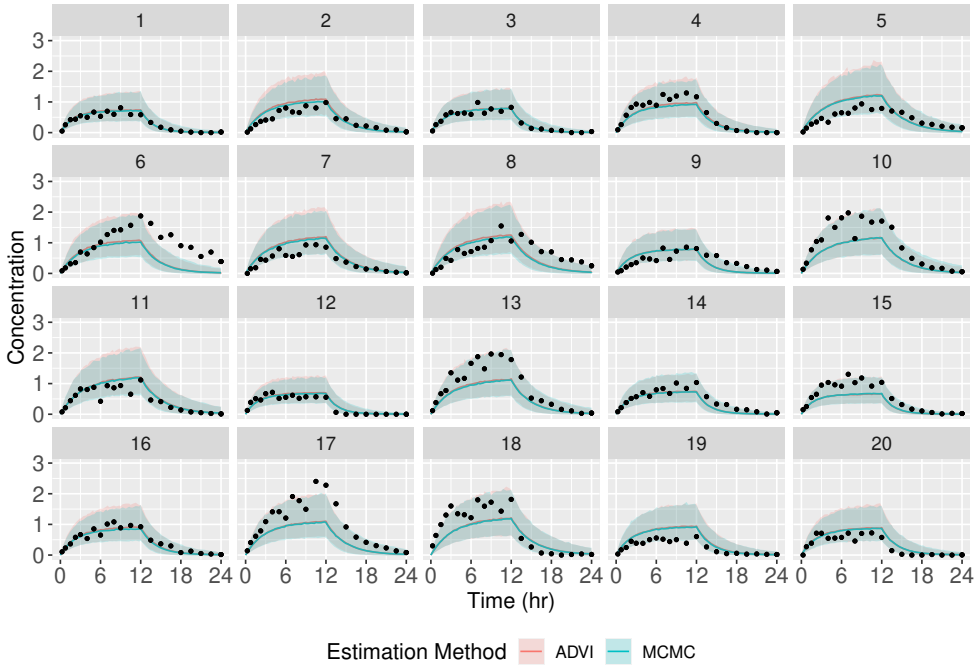

Figure 3.55: One-compartment PK model simulations - observed concentrations and population concentration predictions for first simulated dataset with dense sampling and strong priors

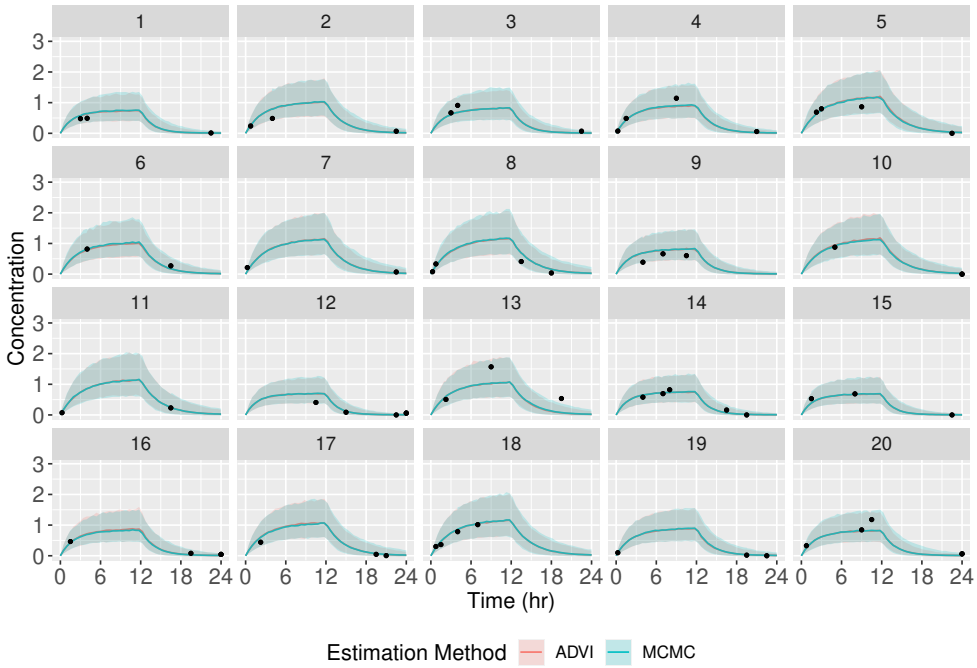

Figure 3.56: One-compartment PK model simulations - observed concentrations and population concentration predictions for first simulated dataset with sparse sampling and strong priors

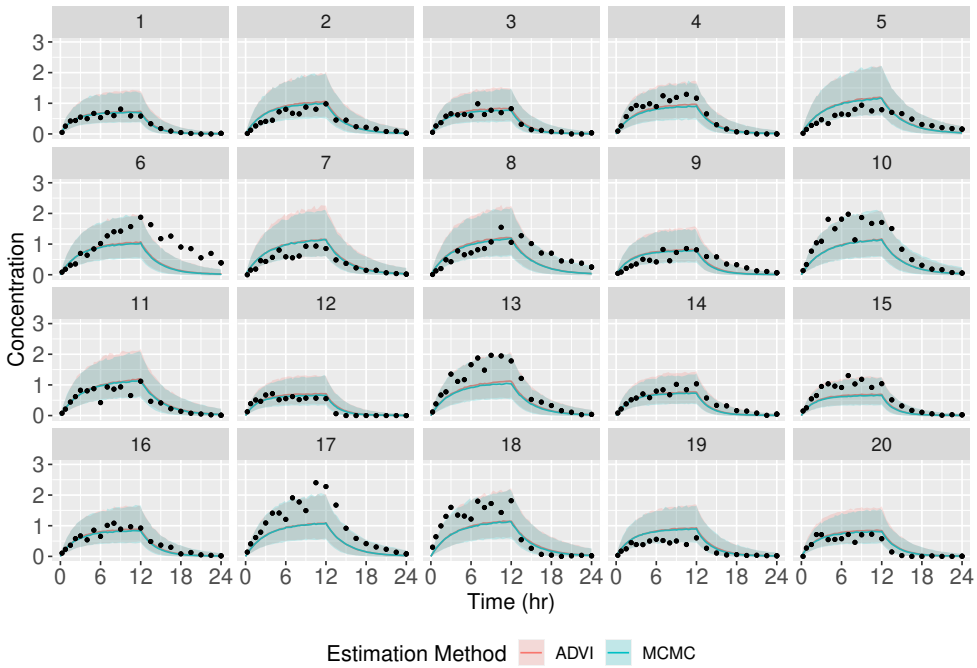

Figure 3.57: One-compartment PK model simulations - observed concentrations and population concentration predictions for first simulated dataset with dense sampling and misspecified priors

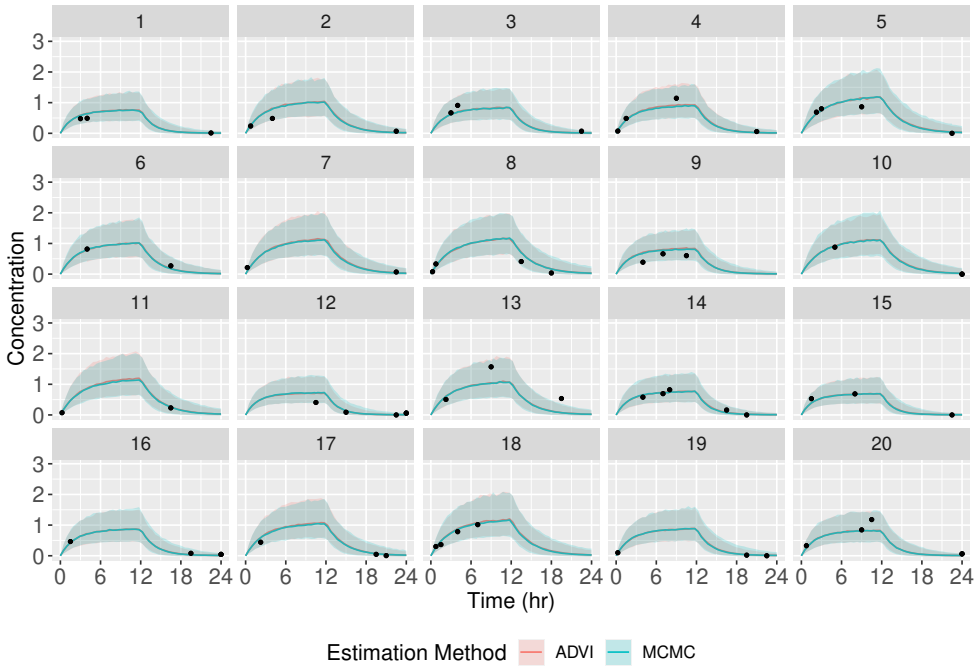

Figure 3.58: One-compartment PK model simulations - observed concentrations and population concentration predictions for first simulated dataset with sparse sampling and misspecified priors

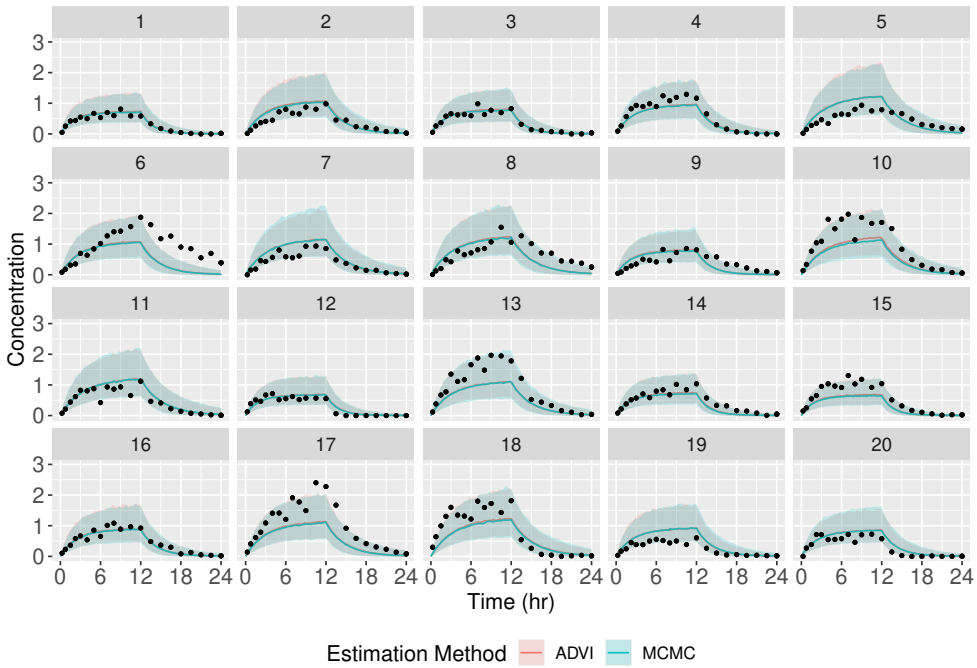

Figure 3.59: One-compartment PK model simulations - observed concentrations and population concentration predictions for first simulated dataset with dense sampling and non-informative priors

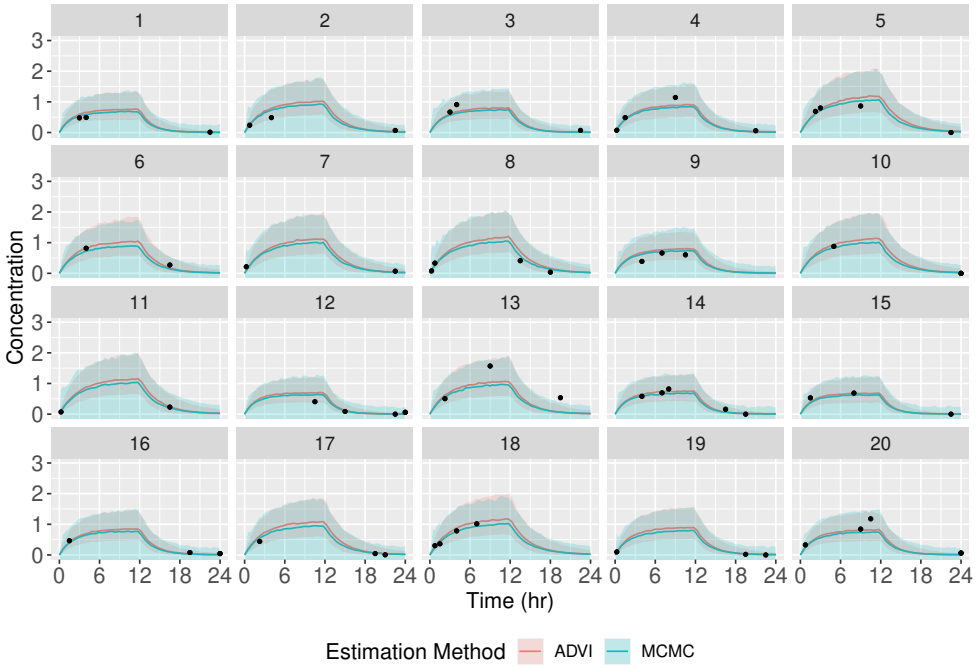

Figure 3.60: One-compartment PK model simulations - observed concentrations and population concentration predictions for first simulated dataset with sparse sampling and non-informative priors

3.6.2.8.5 Scenario 1 Model Selection Tables

Table 3.3: One-Compartment Scenario, Model Selection Strategy 1 (information criteria) - MCMC

	Strong Priors		Weak Priors		Misspecified Priors		Non-informative Priors	
Model Order	n	%	n	%	n	%	n	%
Sparse sampling								
DIC								
M1,M3,M2	91	36.4	85	34.0	95	38.0	34	24.3
M1,M2,M3	12	4.8	16	6.4	9	3.6	19	13.6
M3,M1,M2	110	44.0	83	33.2	106	42.4	19	13.6
M3,M2,M1	16	6.4	14	5.6	17	6.8	2	1.4
M2,M1,M3	9	3.6	33	13.2	10	4.0	40	28.6
M2,M3,M1	12	4.8	19	7.6	13	5.2	26	18.6
PSIS-LOOIC								
M1,M3,M2	89	35.6	109	43.6	98	39.2	51	20.4
M1,M2,M3	23	9.2	15	6.0	21	8.4	3	1.2
M3,M1,M2	108	43.2	102	40.8	98	39.2	133	53.2
M3,M2,M1	18	7.2	13	5.2	22	8.8	51	20.4
M2,M1,M3	4	1.6	7	2.8	2	0.8	1	0.4
M2,M3,M1	8	3.2	4	1.6	9	3.6	11	4.4
WAIC								
M1,M3,M2	81	32.4	88	35.2	85	34.0	44	17.6
M1,M2,M3	15	6.0	24	9.6	12	4.8	2	0.8
M3,M1,M2	103	41.2	97	38.8	96	38.4	133	53.2
M3,M2,M1	20	8.0	18	7.2	25	10.0	54	21.6
M2,M1,M3	16	6.4	6	2.4	14	5.6	4	1.6
M2,M3,M1	15	6.0	17	6.8	18	7.2	13	5.2
Dense sampling								
DIC								
M1,M3,M2	35	14.0	47	18.8	45	18.0	54	21.6
M1,M2,M3	33	13.2	33	13.2	35	14.0	31	12.4
M3,M1,M2	62	24.8	55	22.0	73	29.2	46	18.4
M3,M2,M1	59	23.6	43	17.2	38	15.2	43	17.2
M2,M1,M3	31	12.4	31	12.4	26	10.4	37	14.8
M2,M3,M1	30	12.0	41	16.4	33	13.2	39	15.6
PSIS-LOOIC								
M1,M3,M2	41	16.4	49	19.6	48	19.2	44	17.6
M1,M2,M3	35	14.0	30	12.0	32	12.8	42	16.8
M3,M1,M2	55	22.0	51	20.4	65	26.0	42	16.8
M3,M2,M1	58	23.2	39	15.6	45	18.0	42	16.8
M2,M1,M3	26	10.4	39	15.6	35	14.0	40	16.0
M2,M3,M1	35	14.0	42	16.8	25	10.0	40	16.0
WAIC								
M1,M3,M2	40	16.0	49	19.6	49	19.6	46	18.4
M1,M2,M3	32	12.8	35	14.0	33	13.2	35	14.0
M3,M1,M2	59	23.6	52	20.8	66	26.4	46	18.4
M3,M2,M1	59	23.6	37	14.8	43	17.2	44	17.6
M2,M1,M3	27	10.8	33	13.2	32	12.8	40	16.0
M2,M3,M1	33	13.2	44	17.6	27	10.8	39	15.6

Table 3.4: One-Compartment Scenario, Model Selection Strategy 1 (information criteria) - ADVI

	Strong Priors		Weak Priors		Misspecified Priors		Non-informative Priors	
Model Order	n	%	n	%	n	%	n	%
Sparse sampling								
DIC								
M1,M3,M2	61	24.4	59	23.6	54	21.6	30	21.4
M1,M2,M3	56	22.4	60	24.0	53	21.2	33	23.6
M3,M1,M2	58	23.2	54	21.6	82	32.8	27	19.3
M3,M2,M1	32	12.8	31	12.4	35	14.0	21	15.0
M2,M1,M3	24	9.6	31	12.4	15	6.0	21	15.0
M2,M3,M1	19	7.6	15	6.0	11	4.4	8	5.7
PSIS-LOOIC								
M1,M3,M2	56	22.4	59	23.6	64	25.6	68	27.2
M1,M2,M3	41	16.4	48	19.2	40	16.0	49	19.6
M3,M1,M2	72	28.8	52	20.8	72	28.8	58	23.2
M3,M2,M1	34	13.6	43	17.2	35	14.0	34	13.6
M2,M1,M3	26	10.4	29	11.6	25	10.0	22	8.8
M2,M3,M1	21	8.4	19	7.6	14	5.6	19	7.6
WAIC								
M1,M3,M2	45	18.0	48	19.2	55	22.0	52	20.8
M1,M2,M3	58	23.2	58	23.2	57	22.8	59	23.6
M3,M1,M2	59	23.6	49	19.6	62	24.8	54	21.6
M3,M2,M1	40	16.0	41	16.4	40	16.0	33	13.2
M2,M1,M3	30	12.0	39	15.6	24	9.6	30	12.0
M2,M3,M1	18	7.2	15	6.0	12	4.8	22	8.8
Dense sampling								
DIC								
M1,M3,M2	56	22.4	51	20.4	60	24.0	54	21.6
M1,M2,M3	61	24.4	42	16.8	58	23.2	43	17.2
M3,M1,M2	19	7.6	47	18.8	23	9.2	38	15.2
M3,M2,M1	28	11.2	38	15.2	33	13.2	39	15.6
M2,M1,M3	51	20.4	40	16.0	46	18.4	35	14.0
M2,M3,M1	35	14.0	32	12.8	30	12.0	41	16.4
PSIS-LOOIC								
M1,M3,M2	53	21.2	52	20.8	44	17.6	52	20.8
M1,M2,M3	59	23.6	42	16.8	66	26.4	41	16.4
M3,M1,M2	30	12.0	53	21.2	35	14.0	53	21.2
M3,M2,M1	35	14.0	41	16.4	26	10.4	38	15.2
M2,M1,M3	44	17.6	26	10.4	45	18.0	29	11.6
M2,M3,M1	29	11.6	36	14.4	34	13.6	37	14.8
WAIC								
M1,M3,M2	55	22.0	53	21.2	46	18.4	54	21.6
M1,M2,M3	58	23.2	38	15.2	64	25.6	43	17.2
M3,M1,M2	29	11.6	50	20.0	33	13.2	49	19.6
M3,M2,M1	36	14.4	40	16.0	29	11.6	40	16.0
M2,M1,M3	43	17.2	30	12.0	45	18.0	28	11.2
M2,M3,M1	29	11.6	39	15.6	33	13.2	36	14.4

Table 3.5: One-Compartment Scenario, Model Selection Strategy 2 (ELBO) - ADVI only

	Strong Priors		Weak Priors		Misspecified Priors		Non-informative Priors	
Model Order	n	%	n	%	n	%	n	%
Sparse sampling								
M1,M3,M2	158	63.2	188	75.2	179	71.6	155	62.0
M1,M2,M3	39	15.6	45	18.0	39	15.6	33	13.2
M3,M1,M2	43	17.2	16	6.4	31	12.4	59	23.6
M3,M2,M1	5	2.0	1	0.4	1	0.4	3	1.2
M2,M1,M3	5	2.0	0	0.0	0	0.0	0	0.0
M2,M3,M1	0	0.0	0	0.0	0	0.0	0	0.0
Dense sampling								
M1,M3,M2	107	42.8	99	39.6	107	42.8	101	40.4
M1,M2,M3	34	13.6	45	18.0	35	14.0	32	12.8
M3,M1,M2	61	24.4	51	20.4	55	22.0	64	25.6
M3,M2,M1	39	15.6	39	15.6	36	14.4	40	16.0
M2,M1,M3	4	1.6	7	2.8	5	2.0	5	2.0
M2,M3,M1	5	2.0	9	3.6	12	4.8	8	3.2

Table 3.6: One-Compartment Scenario, Model Selection Strategy 3 (5-fold leave-subject-out cross-validation) - ADVI only

	Strong Priors		Weak Priors		Misspecified Priors		Non-informative Priors	
Model Order	n	%	n	%	n	%	n	%
Sparse sampling								
M1,M3,M2	160	64.0	172	68.8	155	62.0	213	85.2
M1,M2,M3	0	0.0	1	0.4	0	0.0	3	1.2
M3,M1,M2	90	36.0	76	30.4	95	38.0	34	13.6
M3,M2,M1	0	0.0	1	0.4	0	0.0	0	0.0
M2,M1,M3	0	0.0	0	0.0	0	0.0	0	0.0
M2,M3,M1	0	0.0	0	0.0	0	0.0	0	0.0
Dense sampling								
M1,M3,M2	154	61.6	147	58.8	167	66.8	142	56.8
M1,M2,M3	1	0.4	1	0.4	1	0.4	1	0.4
M3,M1,M2	95	38.0	102	40.8	82	32.8	106	42.4
M3,M2,M1	0	0.0	0	0.0	0	0.0	1	0.4
M2,M1,M3	0	0.0	0	0.0	0	0.0	0	0.0
M2,M3,M1	0	0.0	0	0.0	0	0.0	0	0.0

3.6.2.8.6 Scenario 2 Pareto Smoothed Importance Sampling $\hat{k}$

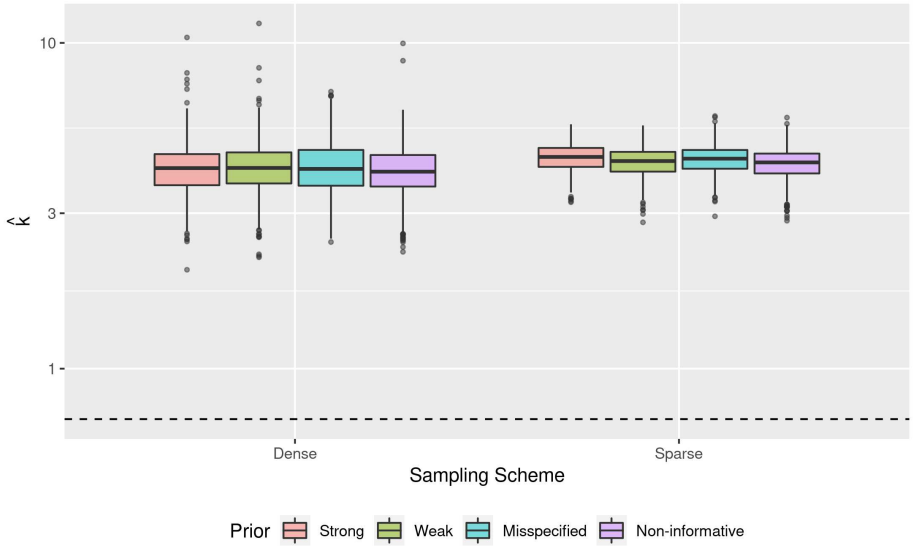

Figure 3.61: Two-compartment PK model simulations - variational inference PSIS $\hat{k}$ values. Values above dashed line ($\hat{k} > 0.7$) indicate discrepancy between VI approximation and posterior distribution

3.6.2.8.7 Scenario 2 Variational Simulation Based Calibration

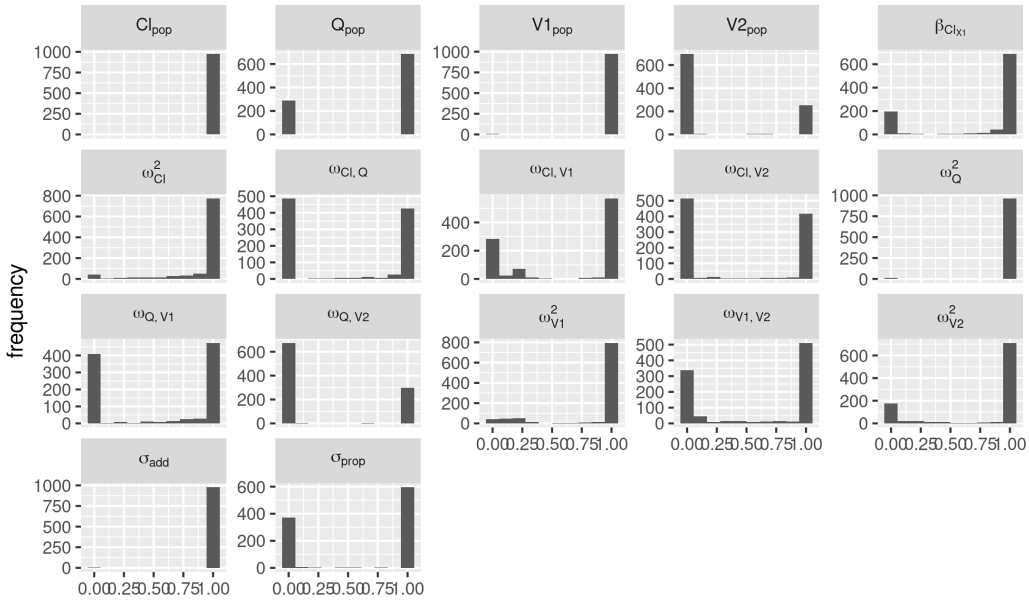

Figure 3.62: Two-compartment PK model simulations - variational simulation based calibration for dense sampling and weak informative priors

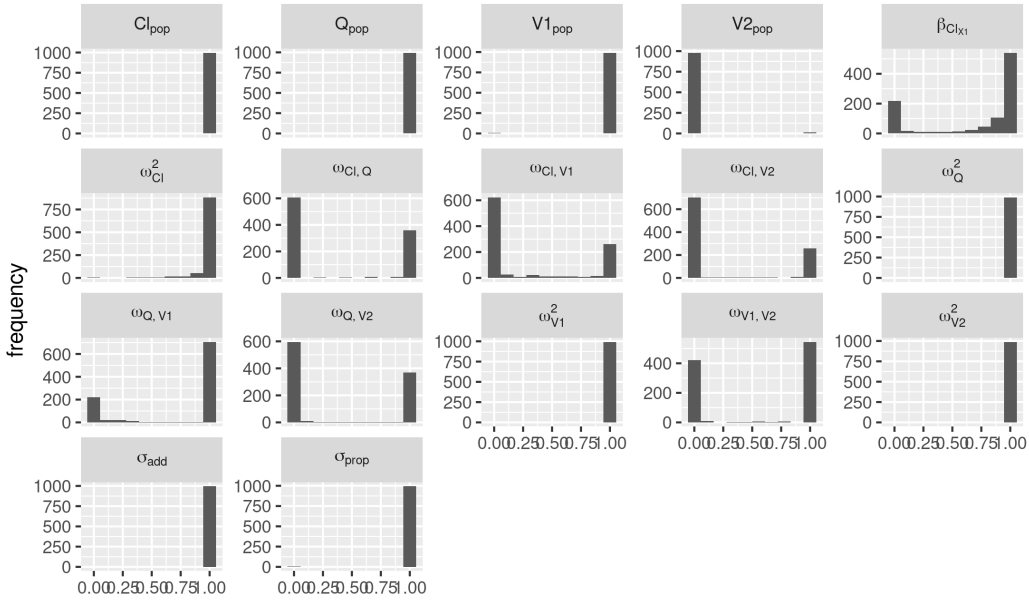

Figure 3.63: Two-compartment PK model simulations - variational simulation based calibration for dense sampling and misspecified informative priors

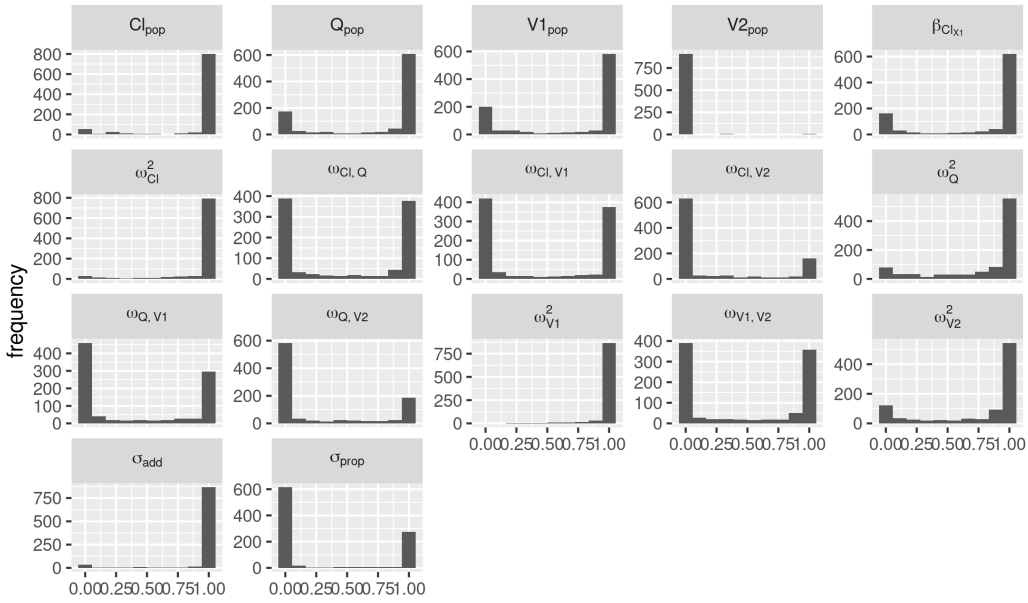

Figure 3.64: Two-compartment PK model simulations - variational simulation based calibration for sparse sampling and strong informative priors

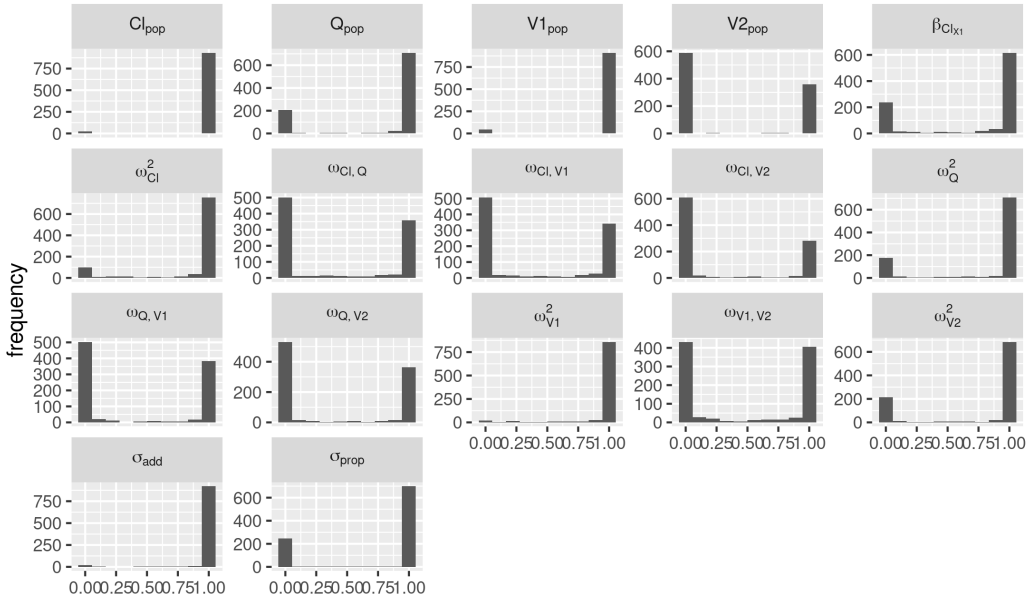

Figure 3.65: Two-compartment PK model simulations - variational simulation based calibration for sparse sampling and weak informative priors

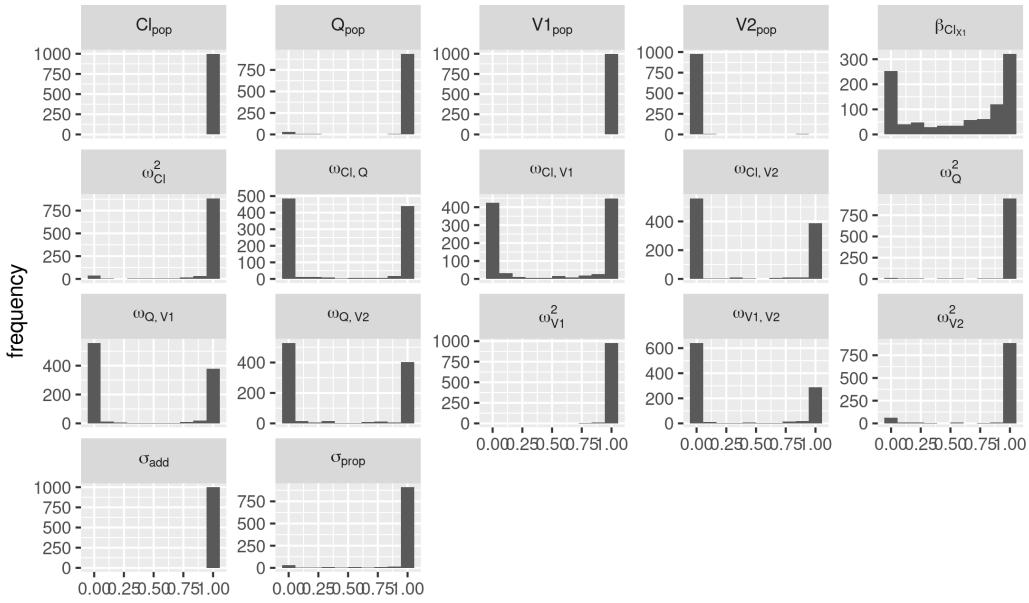

Figure 3.66: Two-compartment PK model simulations - variational simulation based calibration for sparse sampling and misspecified informative priors

3.6.2.8.8 Scenario 2 Individual Concentration Predictions

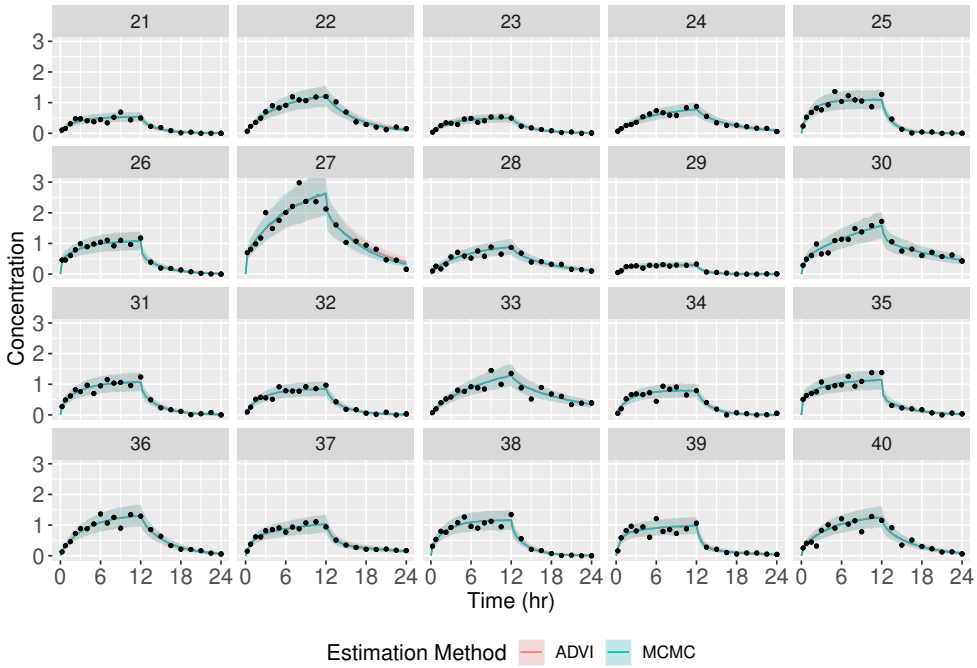

Figure 3.67: Two-compartment PK model simulations - observed concentrations and individual concentration predictions for first simulated dataset with dense sampling and strong priors

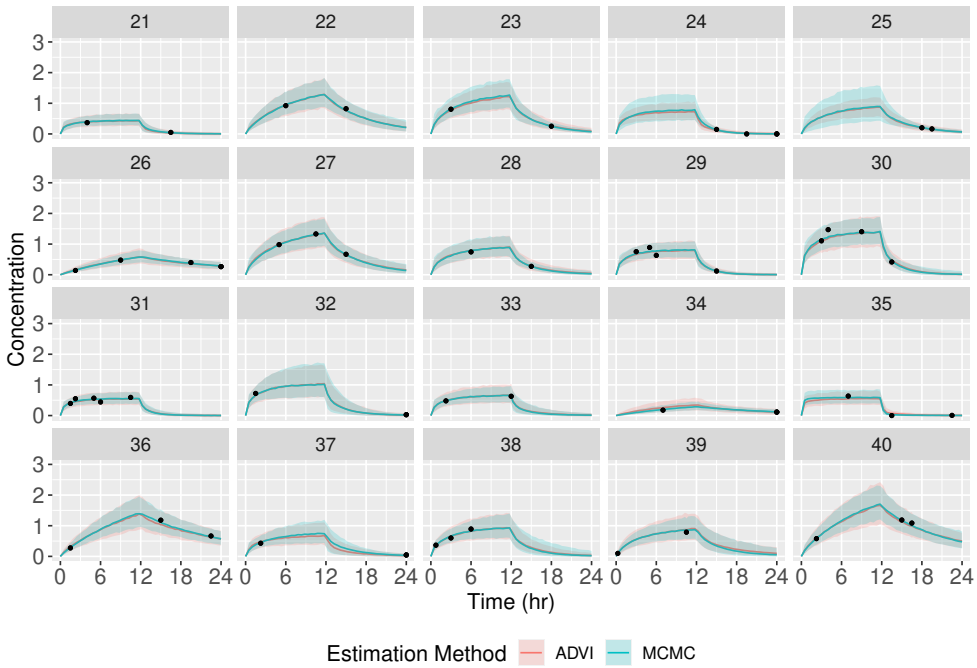

Figure 3.68: Two-compartment PK model simulations - observed concentrations and individual concentration predictions for first simulated dataset with sparse sampling and strong priors

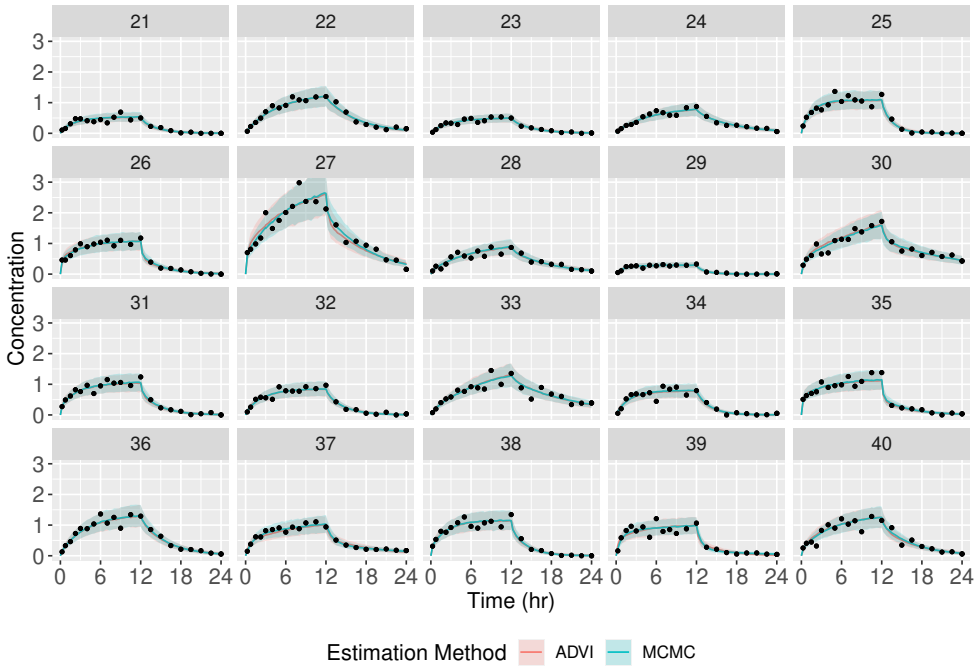

Figure 3.69: Two-compartment PK model simulations - observed concentrations and individual concentration predictions for first simulated dataset with dense sampling and misspecified priors

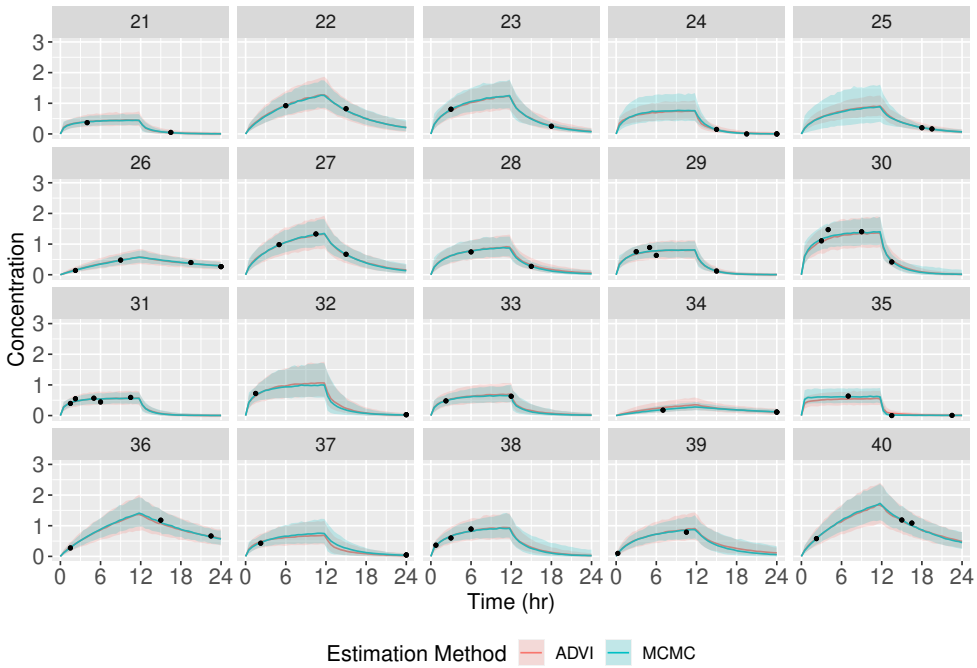

Figure 3.70: Two-compartment PK model simulations - observed concentrations and individual concentration predictions for first simulated dataset with sparse sampling and misspecified priors

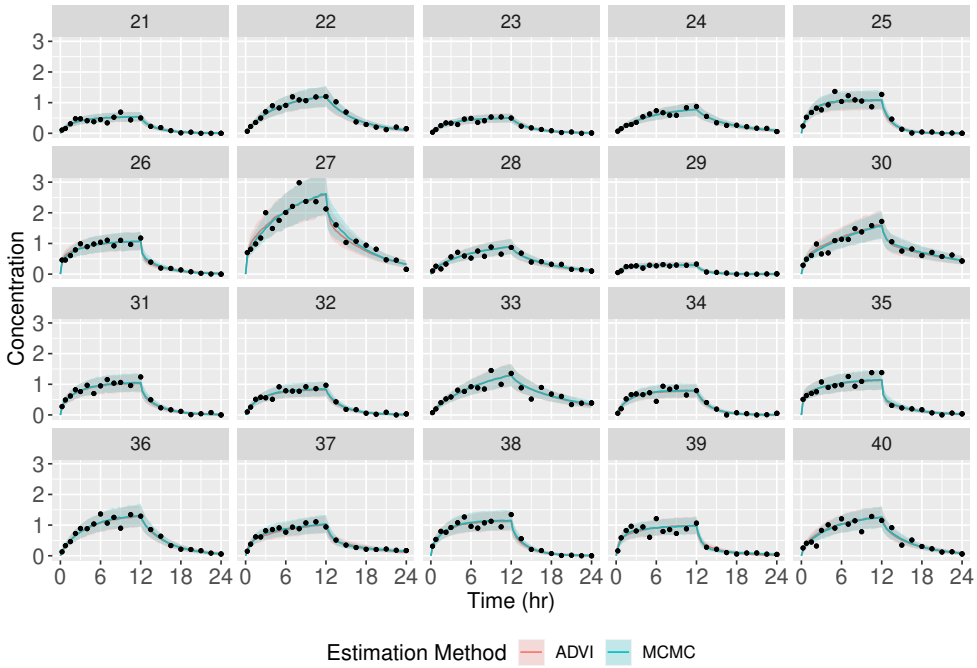

Figure 3.71: Two-compartment PK model simulations - observed concentrations and individual concentration predictions for first simulated dataset with dense sampling and non-informative priors

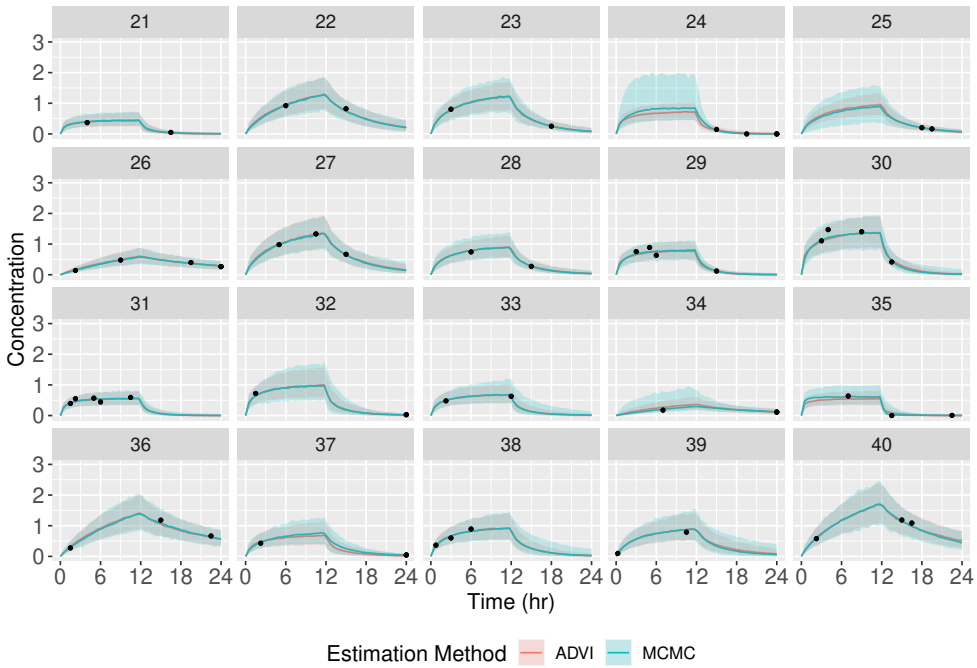

Figure 3.72: Two-compartment PK model simulations - observed concentrations and individual concentration predictions for first simulated dataset with sparse sampling and non-informative priors

3.6.2.8.9 Scenario 2 Population Concentration Predictions

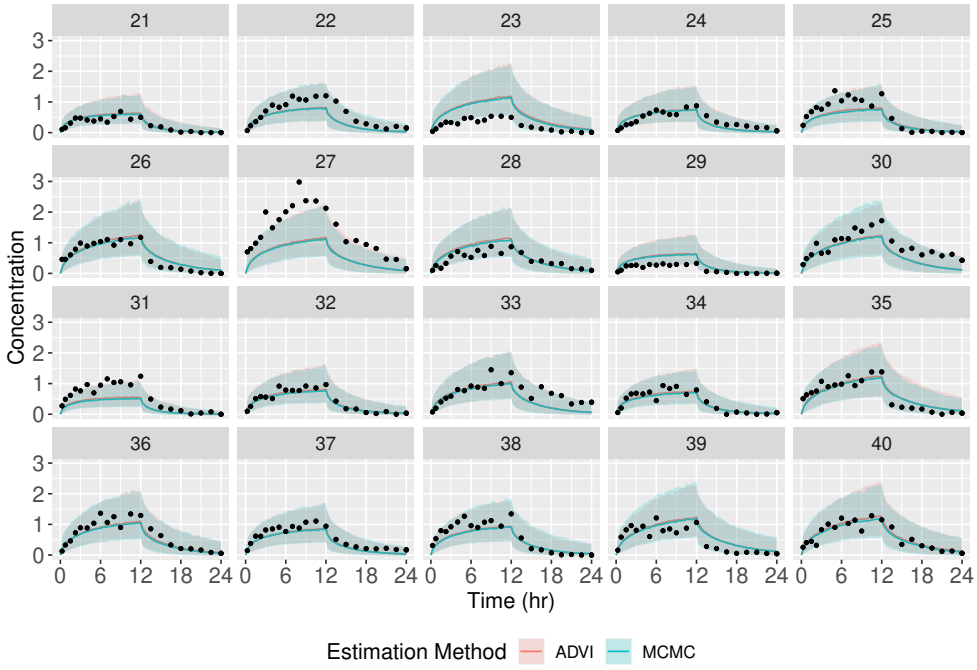

Figure 3.73: Two-compartment PK model simulations - observed concentrations and population concentration predictions for first simulated dataset with dense sampling and strong priors

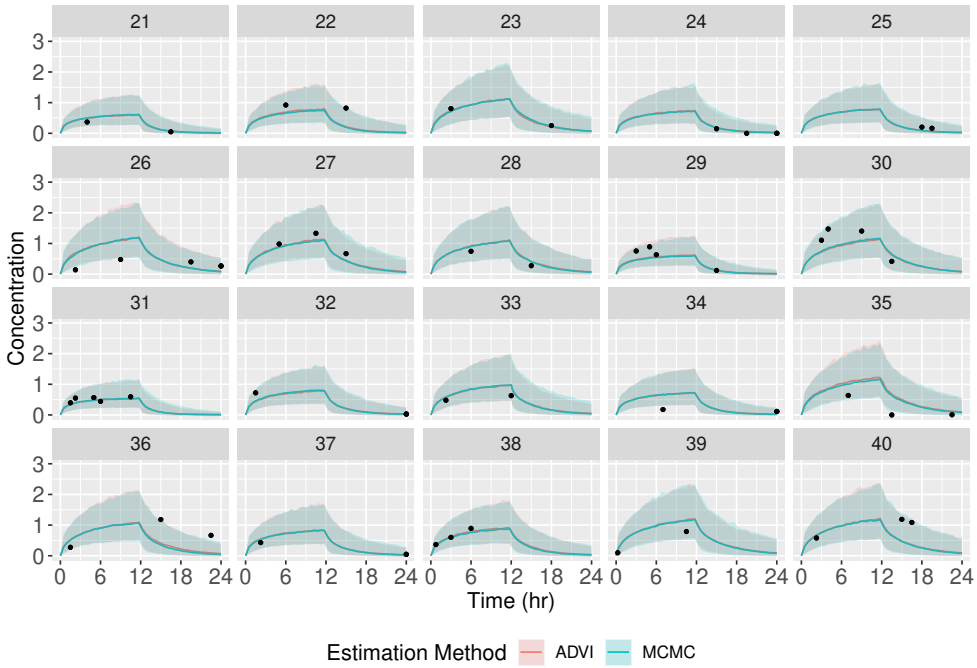

Figure 3.74: Two-compartment PK model simulations - observed concentrations and population concentration predictions for first simulated dataset with sparse sampling and strong priors

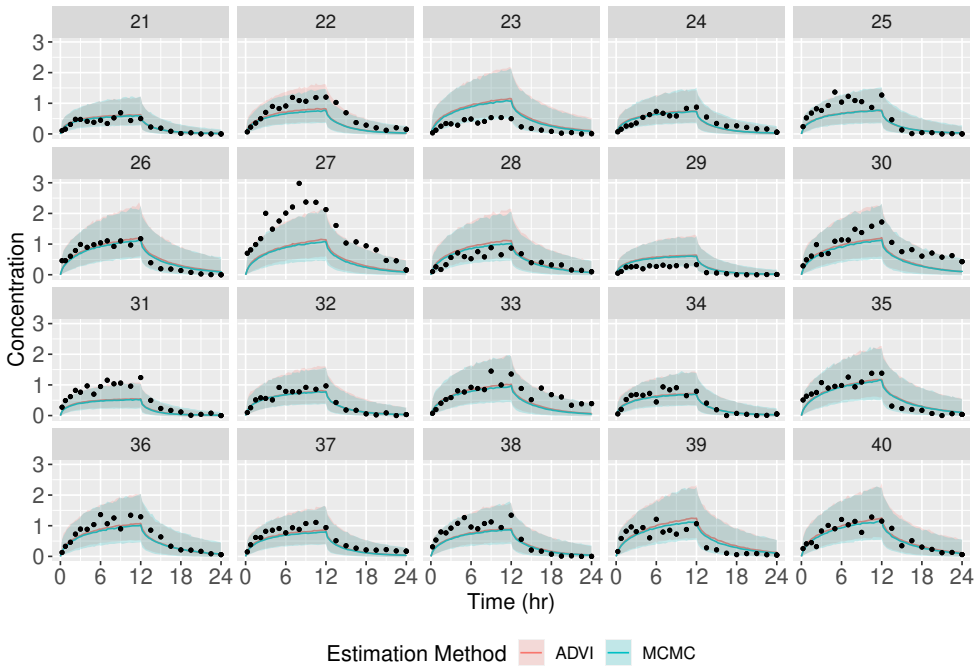

Figure 3.75: Two-compartment PK model simulations - observed concentrations and population concentration predictions for first simulated dataset with dense sampling and misspecified priors

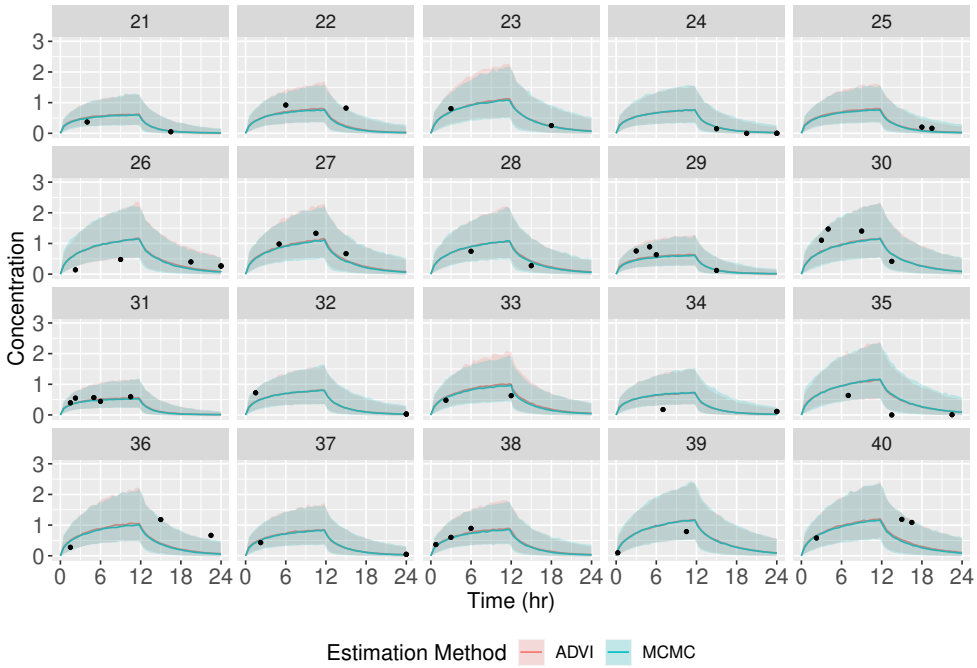

Figure 3.76: Two-compartment PK model simulations - observed concentrations and population concentration predictions for first simulated dataset with sparse sampling and misspecified priors

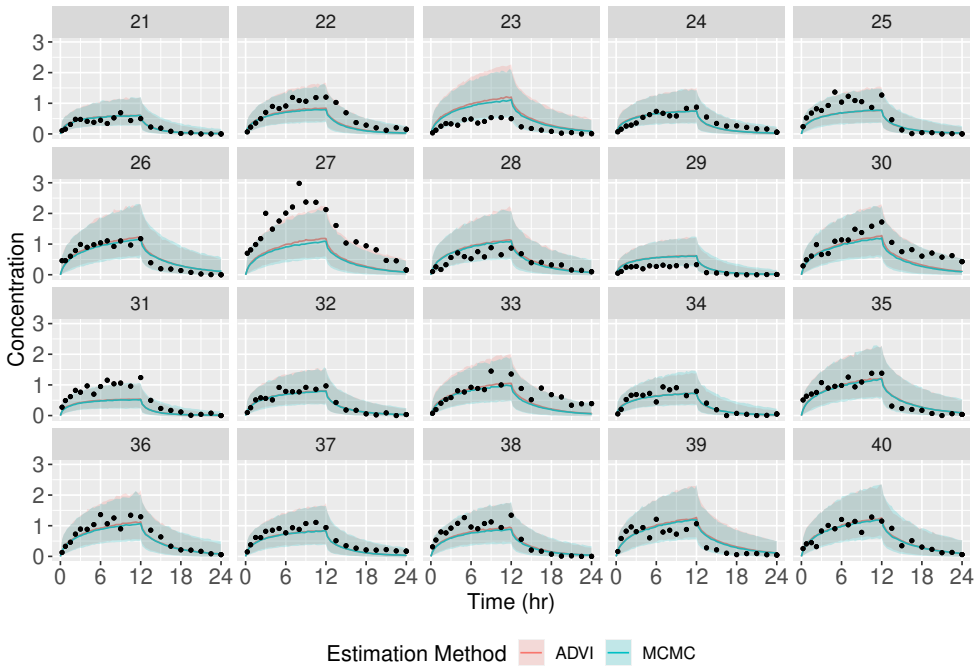

Figure 3.77: Two-compartment PK model simulations - observed concentrations and population concentration predictions for first simulated dataset with dense sampling and non-informative priors

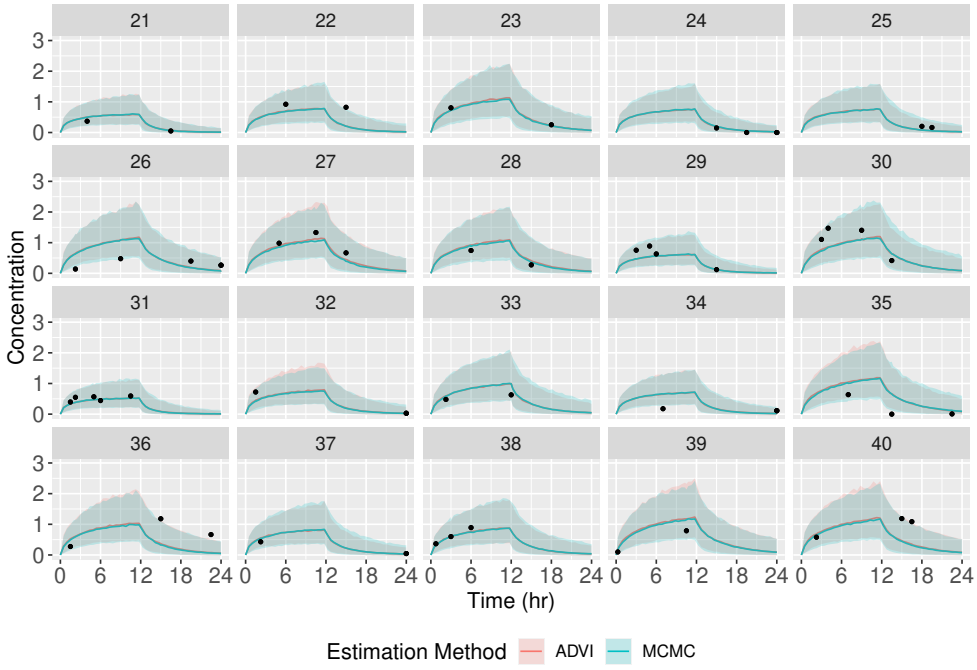

Figure 3.78: Two-compartment PK model simulations - observed concentrations and population concentration predictions for first simulated dataset with sparse sampling and non-informative priors

3.6.2.8.10 Scenario 2 Model Selection Tables and Plots

Table 3.7: Two-Compartment Scenario, Model Selection Strategy 1 (information criteria) - MCMC

	Strong Priors		Weak Priors		Misspecified Priors		Non-informative Priors	
Model Order	n	%	n	%	n	%	n	%
Sparse sampling								
DIC								
M1,M3,M2	50	20.0	44	17.6	49	19.7	44	17.7
M1,M2,M3	43	17.2	45	18.0	30	12.0	48	19.4
M3,M1,M2	48	19.2	40	16.0	42	16.9	28	11.3
M3,M2,M1	33	13.2	30	12.0	41	16.5	23	9.3
M2,M1,M3	39	15.6	43	17.2	40	16.1	41	16.5
M2,M3,M1	37	14.8	48	19.2	47	18.9	64	25.8
PSIS-LOOIC								
M1,M3,M2	61	24.4	60	24.0	58	23.3	43	17.3
M1,M2,M3	31	12.4	27	10.8	23	9.2	29	11.7
M3,M1,M2	61	24.4	58	23.2	78	31.3	71	28.6
M3,M2,M1	48	19.2	50	20.0	31	12.4	47	19.0
M2,M1,M3	24	9.6	16	6.4	25	10.0	17	6.9
M2,M3,M1	25	10.0	39	15.6	34	13.7	41	16.5
WAIC								
M1,M3,M2	53	21.2	54	21.6	43	17.3	46	18.5
M1,M2,M3	30	12.0	30	12.0	24	9.6	26	10.5
M3,M1,M2	53	21.2	51	20.4	69	27.7	68	27.4
M3,M2,M1	44	17.6	44	17.6	38	15.3	42	16.9
M2,M1,M3	29	11.6	31	12.4	40	16.1	25	10.1
M2,M3,M1	41	16.4	40	16.0	35	14.1	41	16.5
Dense sampling								
DIC								
M1,M3,M2	53	21.2	33	13.2	40	16.0	36	14.4
M1,M2,M3	32	12.8	34	13.6	37	14.8	42	16.8
M3,M1,M2	44	17.6	49	19.6	47	18.8	38	15.2
M3,M2,M1	29	11.6	47	18.8	33	13.2	46	18.4
M2,M1,M3	36	14.4	43	17.2	47	18.8	40	16.0
M2,M3,M1	56	22.4	44	17.6	46	18.4	48	19.2
PSIS-LOOIC								
M1,M3,M2	43	17.2	42	16.8	45	18.0	53	21.2
M1,M2,M3	37	14.8	37	14.8	38	15.2	28	11.2
M3,M1,M2	57	22.8	45	18.0	51	20.4	50	20.0
M3,M2,M1	37	14.8	47	18.8	45	18.0	43	17.2
M2,M1,M3	43	17.2	42	16.8	37	14.8	30	12.0
M2,M3,M1	33	13.2	37	14.8	34	13.6	46	18.4
WAIC								
M1,M3,M2	51	20.4	45	18.0	42	16.8	47	18.8
M1,M2,M3	29	11.6	35	14.0	39	15.6	27	10.8
M3,M1,M2	52	20.8	48	19.2	45	18.0	58	23.2
M3,M2,M1	31	12.4	44	17.6	56	22.4	42	16.8
M2,M1,M3	44	17.6	43	17.2	37	14.8	31	12.4
M2,M3,M1	43	17.2	35	14.0	31	12.4	45	18.0

Table 3.8: Two-Compartment Scenario, Model Selection Strategy 1 (information criteria) - ADVI

	Strong Priors		Weak Priors		Misspecified Priors		Non-informative Priors	
Model Order	n	%	n	%	n	%	n	%
Sparse sampling								
DIC								
M1,M3,M2	50	20.0	50	20.0	49	19.7	54	21.8
M1,M2,M3	48	19.2	48	19.2	49	19.7	46	18.5
M3,M1,M2	36	14.4	40	16.0	49	19.7	49	19.8
M3,M2,M1	46	18.4	27	10.8	32	12.9	29	11.7
M2,M1,M3	37	14.8	48	19.2	44	17.7	43	17.3
M2,M3,M1	33	13.2	37	14.8	26	10.4	27	10.9
PSIS-LOOIC								
M1,M3,M2	46	18.4	50	20.0	45	18.1	57	23.0
M1,M2,M3	42	16.8	43	17.2	34	13.7	42	16.9
M3,M1,M2	44	17.6	46	18.4	59	23.7	58	23.4
M3,M2,M1	42	16.8	38	15.2	35	14.1	24	9.7
M2,M1,M3	35	14.0	28	11.2	45	18.1	37	14.9
M2,M3,M1	41	16.4	45	18.0	31	12.4	30	12.1
WAIC								
M1,M3,M2	35	14.0	51	20.4	42	16.9	51	20.6
M1,M2,M3	47	18.8	43	17.2	43	17.3	45	18.1
M3,M1,M2	51	20.4	41	16.4	54	21.7	53	21.4
M3,M2,M1	40	16.0	35	14.0	38	15.3	32	12.9
M2,M1,M3	45	18.0	38	15.2	38	15.3	42	16.9
M2,M3,M1	32	12.8	42	16.8	34	13.7	25	10.1
Dense sampling								
DIC								
M1,M3,M2	26	10.4	48	19.2	30	12.0	38	15.2
M1,M2,M3	88	35.2	87	34.8	69	27.6	89	35.6
M3,M1,M2	16	6.4	17	6.8	21	8.4	10	4.0
M3,M2,M1	15	6.0	17	6.8	24	9.6	12	4.8
M2,M1,M3	84	33.6	60	24.0	83	33.2	83	33.2
M2,M3,M1	21	8.4	21	8.4	23	9.2	18	7.2
PSIS-LOOIC								
M1,M3,M2	28	11.2	48	19.2	30	12.0	38	15.2
M1,M2,M3	90	36.0	90	36.0	73	29.2	86	34.4
M3,M1,M2	18	7.2	19	7.6	20	8.0	8	3.2
M3,M2,M1	15	6.0	13	5.2	24	9.6	13	5.2
M2,M1,M3	80	32.0	60	24.0	81	32.4	86	34.4
M2,M3,M1	19	7.6	20	8.0	22	8.8	19	7.6
WAIC								
M1,M3,M2	28	11.2	48	19.2	30	12.0	36	14.4
M1,M2,M3	95	38.0	89	35.6	69	27.6	89	35.6
M3,M1,M2	15	6.0	17	6.8	19	7.6	8	3.2
M3,M2,M1	12	4.8	16	6.4	25	10.0	13	5.2
M2,M1,M3	79	31.6	61	24.4	86	34.4	84	33.6
M2,M3,M1	21	8.4	19	7.6	21	8.4	20	8.0

Table 3.9: Two-Compartment Scenario, Model Selection Strategy 2 (ELBO) - ADVI only

	Strong Priors		Weak Priors		Misspecified Priors		Non-informative Priors	
Model Order	n	%	n	%	n	%	n	%
Sparse sampling								
M1,M3,M2	197	78.8	183	73.2	212	85.1	114	45.6
M1,M2,M3	4	1.6	38	15.2	10	4.0	44	17.6
M3,M1,M2	47	18.8	16	6.4	24	9.6	55	22.0
M3,M2,M1	2	0.8	9	3.6	3	1.2	22	8.8
M2,M1,M3	0	0.0	3	1.2	0	0.0	8	3.2
M2,M3,M1	0	0.0	1	0.4	0	0.0	7	2.8
Dense sampling								
M1,M3,M2	42	16.8	52	20.8	35	14.0	43	17.2
M1,M2,M3	93	37.2	102	40.8	87	34.8	106	42.4
M3,M1,M2	21	8.4	18	7.2	34	13.6	15	6.0
M3,M2,M1	16	6.4	19	7.6	18	7.2	13	5.2
M2,M1,M3	61	24.4	41	16.4	49	19.6	58	23.2
M2,M3,M1	17	6.8	18	7.2	27	10.8	15	6.0

Table 3.10: Two-Compartment Scenario, Model Selection Strategy 3 (5-fold leave-subject-out cross-validation) - ADVI only

	Strong Priors		Weak Priors		Misspecified Priors		Non-informative Priors	
Model Order	n	%	n	%	n	%	n	%
Sparse sampling								
M1,M3,M2	163	65.2	137	57.6	144	66.7	139	57.9
M1,M2,M3	2	0.8	1	0.4	4	1.9	4	1.7
M3,M1,M2	85	34.0	97	40.8	67	31.0	93	38.8
M3,M2,M1	0	0.0	3	1.3	1	0.5	3	1.2
M2,M1,M3	0	0.0	0	0.0	0	0.0	1	0.4
M2,M3,M1	0	0.0	0	0.0	0	0.0	0	0.0
Dense sampling								
M1,M3,M2	129	51.6	26	10.4	201	80.4	39	15.6
M1,M2,M3	69	27.6	6	2.4	9	3.6	2	0.8
M3,M1,M2	37	14.8	81	32.4	38	15.2	193	77.2
M3,M2,M1	4	1.6	87	34.8	1	0.4	13	5.2
M2,M1,M3	4	1.6	13	5.2	1	0.4	0	0.0
M2,M3,M1	7	2.8	37	14.8	0	0.0	3	1.2

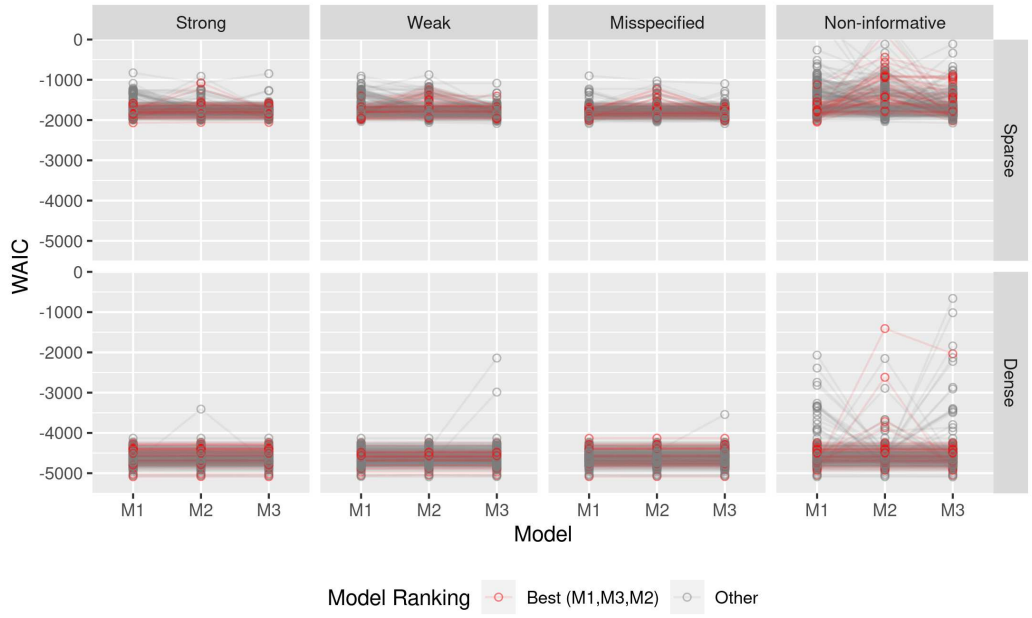

Figure 3.79: Two-Compartment Scenario, Model Selection Strategy 1 (information criteria) - MCMC WAIC values

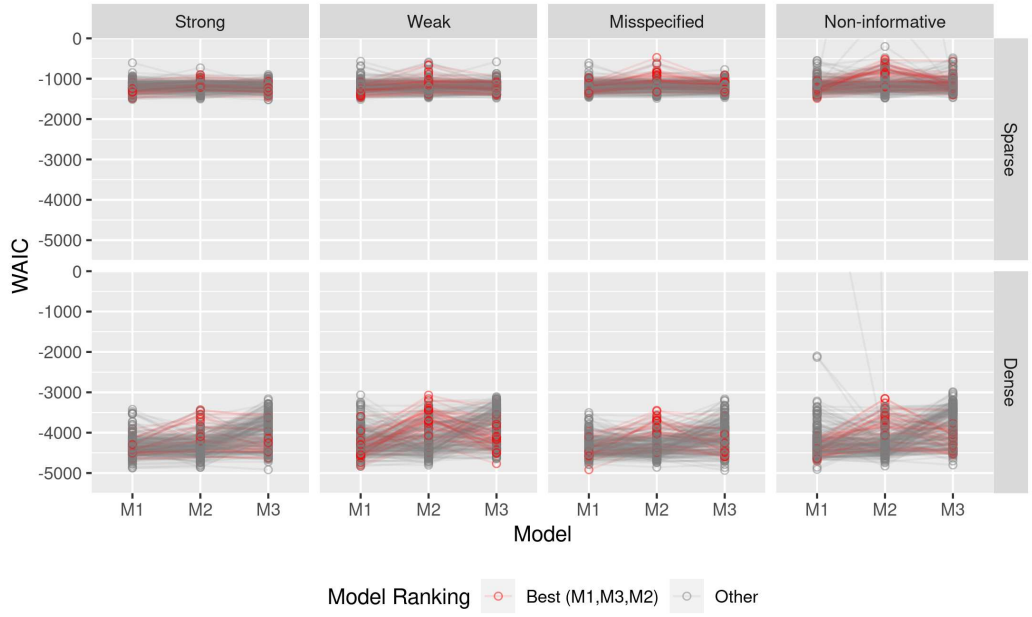

Figure 3.80: Two-Compartment Scenario, Model Selection Strategy 1 (information criteria) - ADVI WAIC values

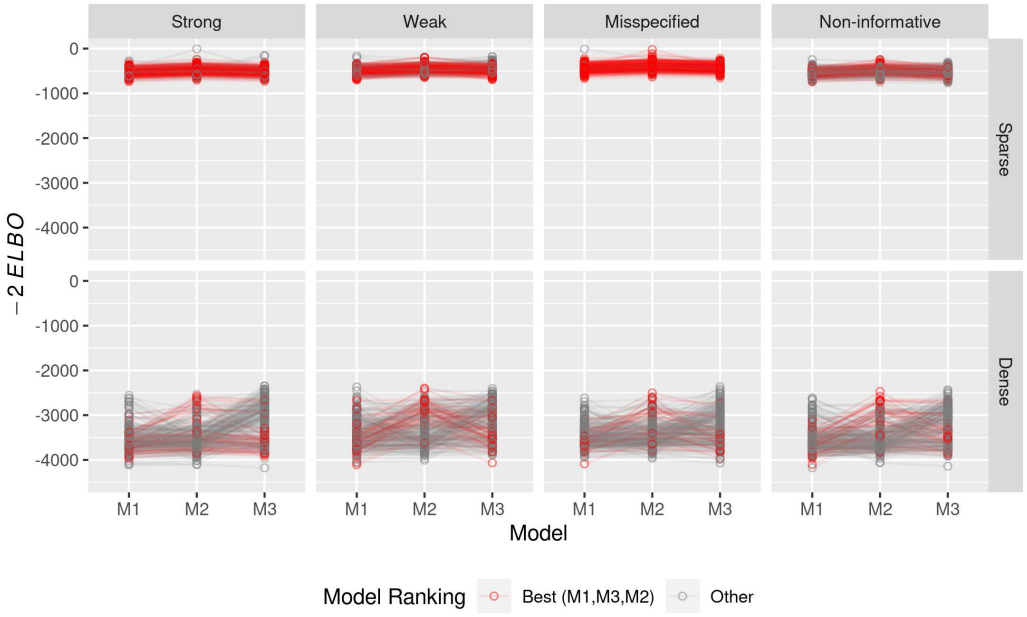

Figure 3.81: Two-Compartment Scenario, Model Selection Strategy 2 (ELBO) - ADVI $-2ELBO$ values

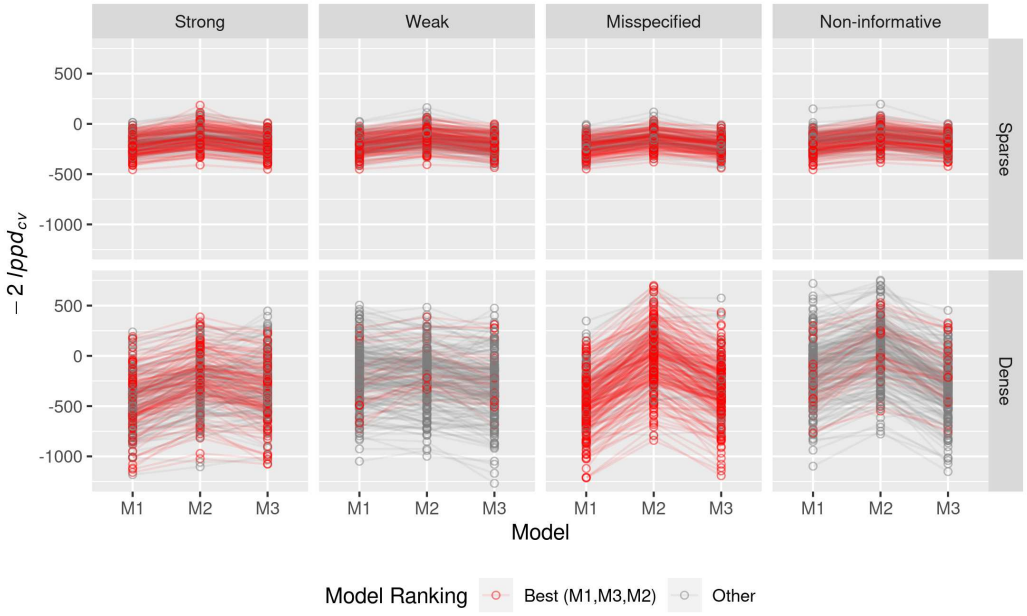

Figure 3.82: Two-Compartment Scenario, Model Selection Strategy 3 (5-fold leave-subject-out cross-validation) - ADVI $-2lppd_{cv}$ values

3.6.2.9 Model Selection Sensitivity Analysis Tables and Plots

Table 3.11: Two-Compartment Scenario, Model Selection Strategy 1 (information criteria) - ADVI [sensitivity analysis with $\beta_{Cl_{X1}} = 0.1$]

	Strong Priors		Weak Priors		Misspecified Priors		Non-informative Priors	
Model Order	n	%	n	%	n	%	n	%
Sparse sampling								
DIC								
M1,M3,M2	42	16.8	31	12.4	36	14.4	39	16.2
M1,M2,M3	47	18.8	40	16.0	39	15.6	53	22.1
M3,M1,M2	25	10.0	55	22.0	40	16.0	42	17.5
M3,M2,M1	38	15.2	30	12.0	49	19.6	20	8.3
M2,M1,M3	54	21.6	50	20.0	44	17.6	46	19.2
M2,M3,M1	44	17.6	44	17.6	42	16.8	40	16.7
PSIS-LOOIC								
M1,M3,M2	41	16.4	45	18.0	39	15.6	40	16.7
M1,M2,M3	32	12.8	36	14.4	31	12.4	39	16.2
M3,M1,M2	38	15.2	44	17.6	54	21.6	41	17.1
M3,M2,M1	54	21.6	48	19.2	61	24.4	39	16.2
M2,M1,M3	44	17.6	39	15.6	34	13.6	40	16.7
M2,M3,M1	41	16.4	38	15.2	31	12.4	41	17.1
WAIC								
M1,M3,M2	42	16.8	37	14.8	39	15.6	44	18.3
M1,M2,M3	37	14.8	42	16.8	36	14.4	43	17.9
M3,M1,M2	36	14.4	44	17.6	43	17.2	35	14.6
M3,M2,M1	44	17.6	45	18.0	60	24.0	37	15.4
M2,M1,M3	49	19.6	44	17.6	41	16.4	43	17.9
M2,M3,M1	42	16.8	38	15.2	31	12.4	38	15.8
Dense sampling								
DIC								
M1,M3,M2	63	25.2	47	18.8	48	19.2	31	12.4
M1,M2,M3	21	8.4	27	10.8	39	15.6	43	17.2
M3,M1,M2	53	21.2	56	22.4	42	16.8	38	15.2
M3,M2,M1	47	18.8	56	22.4	51	20.4	72	28.8
M2,M1,M3	30	12.0	30	12.0	42	16.8	26	10.4
M2,M3,M1	36	14.4	34	13.6	28	11.2	40	16.0
PSIS-LOOIC								
M1,M3,M2	53	21.2	46	18.4	43	17.2	33	13.2
M1,M2,M3	28	11.2	32	12.8	42	16.8	39	15.6
M3,M1,M2	55	22.0	56	22.4	44	17.6	38	15.2
M3,M2,M1	41	16.4	56	22.4	44	17.6	64	25.6
M2,M1,M3	31	12.4	26	10.4	45	18.0	30	12.0
M2,M3,M1	42	16.8	34	13.6	32	12.8	46	18.4
WAIC								
M1,M3,M2	54	21.6	45	18.0	48	19.2	32	12.8
M1,M2,M3	25	10.0	27	10.8	40	16.0	40	16.0
M3,M1,M2	56	22.4	59	23.6	41	16.4	37	14.8
M3,M2,M1	44	17.6	53	21.2	44	17.6	67	26.8
M2,M1,M3	30	12.0	29	11.6	47	18.8	29	11.6
M2,M3,M1	41	16.4	37	14.8	30	12.0	45	18.0

Table 3.12: Two-Compartment Scenario, Model Selection Strategy 2 (ELBO) - ADVI only [sensitivity analysis with $\beta_{Cl_{X1}} = 0.1$]

	Strong Priors		Weak Priors		Misspecified Priors		Non-informative Priors	
Model Order	n	%	n	%	n	%	n	%
Sparse sampling								
M1,M3,M2	56	22.4	47	18.8	46	18.4	53	22.1
M1,M2,M3	84	33.6	99	39.6	88	35.2	56	23.3
M3,M1,M2	16	6.4	9	3.6	16	6.4	23	9.6
M3,M2,M1	7	2.8	3	1.2	5	2.0	29	12.1
M2,M1,M3	59	23.6	62	24.8	60	24.0	40	16.7
M2,M3,M1	28	11.2	30	12.0	35	14.0	39	16.2
Dense sampling								
M1,M3,M2	67	26.8	55	22.0	51	20.4	26	10.4
M1,M2,M3	25	10.0	26	10.4	40	16.0	44	17.6
M3,M1,M2	48	19.2	50	20.0	35	14.0	40	16.0
M3,M2,M1	37	14.8	50	20.0	50	20.0	70	28.0
M2,M1,M3	29	11.6	31	12.4	37	14.8	25	10.0
M2,M3,M1	44	17.6	38	15.2	37	14.8	45	18.0

Table 3.13: Two-Compartment Scenario, Model Selection Strategy 3 (5-fold leave-subject-out cross-validation) - ADVI only [sensitivity analysis with $\beta_{Cl_{X1}} = 0.1$]

	Strong Priors		Weak Priors		Misspecified Priors		Non-informative Priors	
Model Order	n	%	n	%	n	%	n	%
Sparse sampling								
M1,M3,M2	46	21.8	47	33.1	26	28.3	48	22.5
M1,M2,M3	44	20.9	22	15.5	15	16.3	52	24.4
M3,M1,M2	41	19.4	21	14.8	16	17.4	26	12.2
M3,M2,M1	23	10.9	16	11.3	7	7.6	36	16.9
M2,M1,M3	36	17.1	22	15.5	17	18.5	31	14.6
M2,M3,M1	21	10.0	14	9.9	11	12.0	20	9.4
Dense sampling								
M1,M3,M2	38	15.2	49	19.6	18	7.2	51	20.4
M1,M2,M3	109	43.6	90	36.0	114	45.6	72	28.8
M3,M1,M2	7	2.8	8	3.2	8	3.2	22	8.8
M3,M2,M1	8	3.2	12	4.8	9	3.6	17	6.8
M2,M1,M3	67	26.8	62	24.8	78	31.2	63	25.2
M2,M3,M1	21	8.4	29	11.6	23	9.2	25	10.0

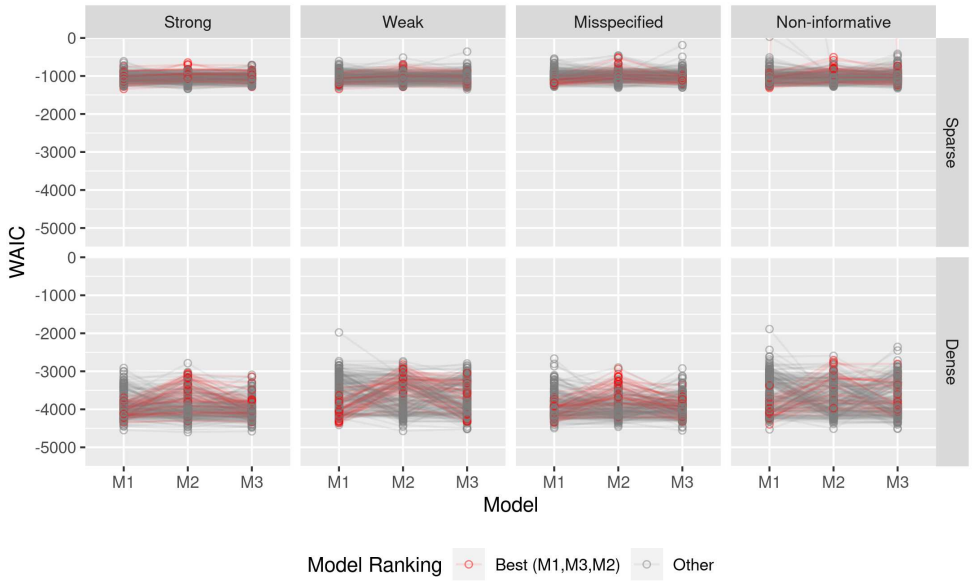

Figure 3.83: Two-Compartment Scenario, Model Selection Strategy 1 (information criteria) - ADVI WAIC values [sensitivity analysis with $\beta_{Cl_{X1}} = 0.1$]

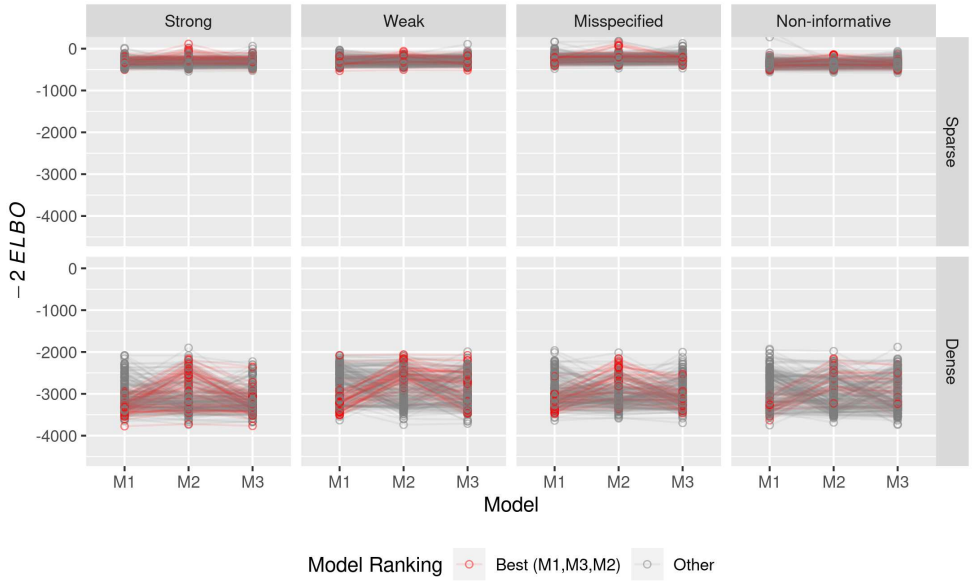

Figure 3.84: Two-Compartment Scenario, Model Selection Strategy 2 (ELBO) - ADVI $-2ELBO$ values [sensitivity analysis with $\beta_{Cl_{X1}} = 0.1$]

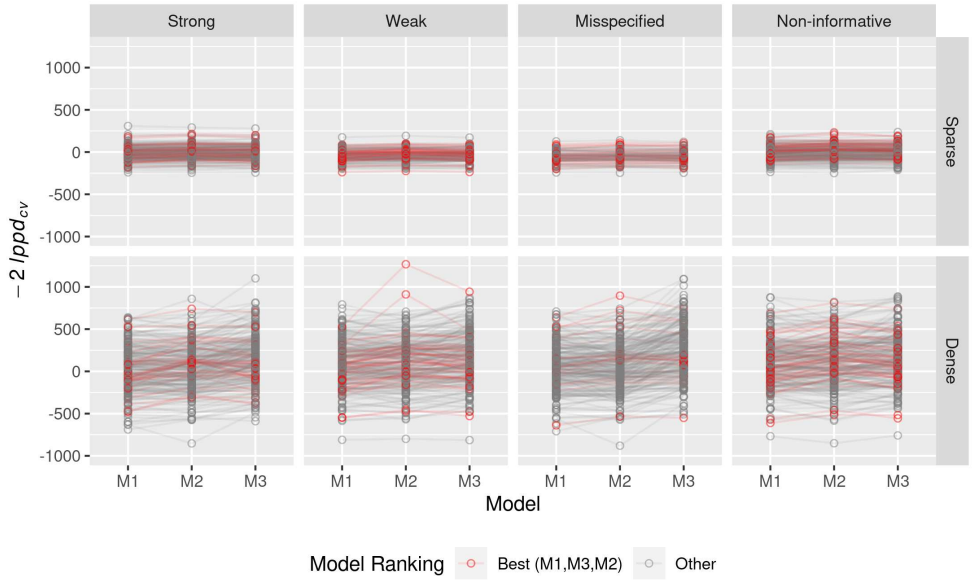

Figure 3.85: Two-Compartment Scenario, Model Selection Strategy 3 (5-fold leave-subject-out cross-validation) - ADVI $-2lppd_{cv}$ values [sensitivity analysis with $\beta_{Cl_{X1}} = 0.1$]

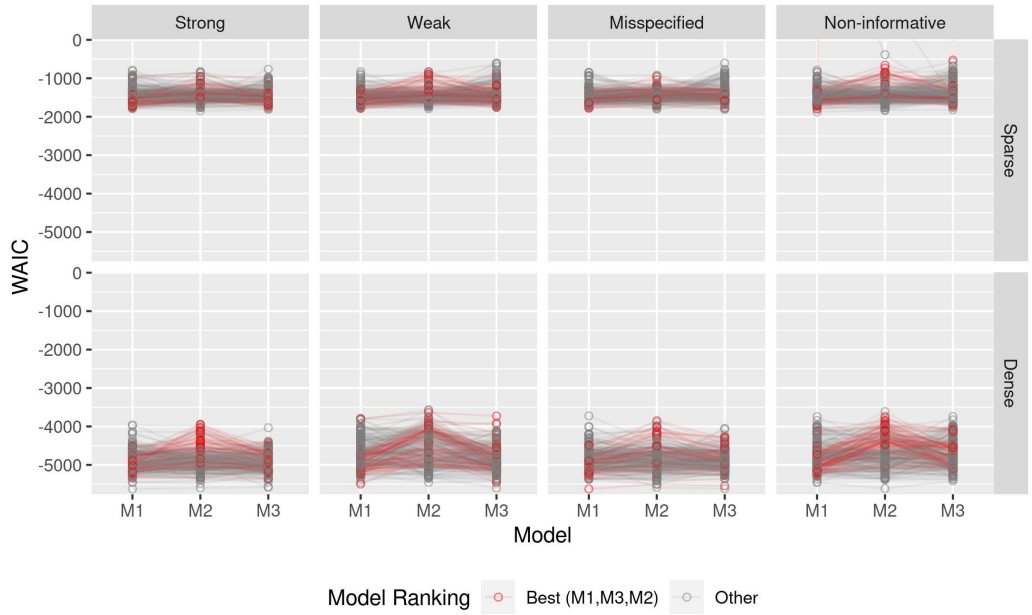

Figure 3.86: Two-Compartment Scenario, Model Selection Strategy 1 (information criteria) - ADVI WAIC values [sensitivity analysis with $\beta_{Cl_{X1}} = 1$]

Table 3.14: Two-Compartment Scenario, Model Selection Strategy 1 (information criteria) - ADVI [sensitivity analysis with $\beta_{Cl_{X1}} = 1$]

	Strong Priors		Weak Priors		Misspecified Priors		Non-informative Priors	
Model Order	n	%	n	%	n	%	n	%
Sparse sampling								
DIC								
M1,M3,M2	52	20.8	55	22.0	54	21.6	47	18.8
M1,M2,M3	49	19.6	66	26.4	58	23.2	62	24.8
M3,M1,M2	64	25.6	34	13.6	52	20.8	36	14.4
M3,M2,M1	32	12.8	29	11.6	21	8.4	30	12.0
M2,M1,M3	27	10.8	45	18.0	42	16.8	56	22.4
M2,M3,M1	26	10.4	21	8.4	23	9.2	19	7.6
PSIS-LOOIC								
M1,M3,M2	52	20.8	54	21.6	36	14.4	40	16.0
M1,M2,M3	45	18.0	52	20.8	62	24.8	54	21.6
M3,M1,M2	52	20.8	43	17.2	39	15.6	48	19.2
M3,M2,M1	38	15.2	30	12.0	38	15.2	34	13.6
M2,M1,M3	33	13.2	41	16.4	40	16.0	44	17.6
M2,M3,M1	30	12.0	30	12.0	35	14.0	30	12.0
WAIC								
M1,M3,M2	45	18.0	51	20.4	37	14.8	37	14.8
M1,M2,M3	51	20.4	53	21.2	58	23.2	61	24.4
M3,M1,M2	44	17.6	28	11.2	40	16.0	32	12.8
M3,M2,M1	49	19.6	33	13.2	35	14.0	29	11.6
M2,M1,M3	31	12.4	54	21.6	49	19.6	53	21.2
M2,M3,M1	30	12.0	31	12.4	31	12.4	38	15.2
Dense sampling								
DIC								
M1,M3,M2	51	20.4	60	24.0	45	18.0	59	23.6
M1,M2,M3	51	20.4	28	11.2	39	15.6	47	18.8
M3,M1,M2	39	15.6	80	32.0	48	19.2	52	20.8
M3,M2,M1	48	19.2	45	18.0	41	16.4	36	14.4
M2,M1,M3	26	10.4	14	5.6	36	14.4	36	14.4
M2,M3,M1	35	14.0	23	9.2	41	16.4	20	8.0
PSIS-LOOIC								
M1,M3,M2	44	17.6	56	22.4	46	18.4	60	24.0
M1,M2,M3	45	18.0	26	10.4	40	16.0	55	22.0
M3,M1,M2	51	20.4	84	33.6	52	20.8	49	19.6
M3,M2,M1	49	19.6	45	18.0	37	14.8	37	14.8
M2,M1,M3	25	10.0	18	7.2	38	15.2	33	13.2
M2,M3,M1	36	14.4	21	8.4	37	14.8	16	6.4
WAIC								
M1,M3,M2	47	18.8	57	22.8	44	17.6	63	25.2
M1,M2,M3	47	18.8	25	10.0	37	14.8	49	19.6
M3,M1,M2	45	18.0	82	32.8	53	21.2	50	20.0
M3,M2,M1	47	18.8	45	18.0	43	17.2	35	14.0
M2,M1,M3	27	10.8	19	7.6	35	14.0	36	14.4
M2,M3,M1	37	14.8	22	8.8	38	15.2	17	6.8

Table 3.15: Two-Compartment Scenario, Model Selection Strategy 2 (ELBO) - ADVI only [sensitivity analysis with $\beta_{Cl_{X1}} = 1$]

	Strong Priors		Weak Priors		Misspecified Priors		Non-informative Priors	
Model Order	n	%	n	%	n	%	n	%
Sparse sampling								
M1,M3,M2	150	60.0	201	80.4	202	80.8	154	61.6
M1,M2,M3	3	1.2	7	2.8	7	2.8	6	2.4
M3,M1,M2	96	38.4	42	16.8	38	15.2	88	35.2
M3,M2,M1	1	0.4	0	0.0	3	1.2	2	0.8
M2,M1,M3	0	0.0	0	0.0	0	0.0	0	0.0
M2,M3,M1	0	0.0	0	0.0	0	0.0	0	0.0
Dense sampling								
M1,M3,M2	102	40.8	79	31.6	79	31.6	81	32.4
M1,M2,M3	28	11.2	27	10.8	36	14.4	50	20.0
M3,M1,M2	74	29.6	100	40.0	72	28.8	73	29.2
M3,M2,M1	41	16.4	38	15.2	41	16.4	27	10.8
M2,M1,M3	1	0.4	2	0.8	9	3.6	10	4.0
M2,M3,M1	4	1.6	4	1.6	13	5.2	9	3.6

Table 3.16: Two-Compartment Scenario, Model Selection Strategy 3 (5-fold leave-subject-out cross-validation) - ADVI only [sensitivity analysis with $\beta_{Cl_{X1}} = 1$]

	Strong Priors		Weak Priors		Misspecified Priors		Non-informative Priors	
Model Order	n	%	n	%	n	%	n	%
Sparse sampling								
M1,M3,M2	155	62.0	170	68.0	159	64.1	195	78.0
M1,M2,M3	0	0.0	0	0.0	0	0.0	0	0.0
M3,M1,M2	95	38.0	80	32.0	89	35.9	55	22.0
M3,M2,M1	0	0.0	0	0.0	0	0.0	0	0.0
M2,M1,M3	0	0.0	0	0.0	0	0.0	0	0.0
M2,M3,M1	0	0.0	0	0.0	0	0.0	0	0.0
Dense sampling								
M1,M3,M2	219	87.6	202	80.8	216	86.4	164	65.6
M1,M2,M3	1	0.4	0	0.0	2	0.8	0	0.0
M3,M1,M2	30	12.0	48	19.2	32	12.8	86	34.4
M3,M2,M1	0	0.0	0	0.0	0	0.0	0	0.0
M2,M1,M3	0	0.0	0	0.0	0	0.0	0	0.0
M2,M3,M1	0	0.0	0	0.0	0	0.0	0	0.0

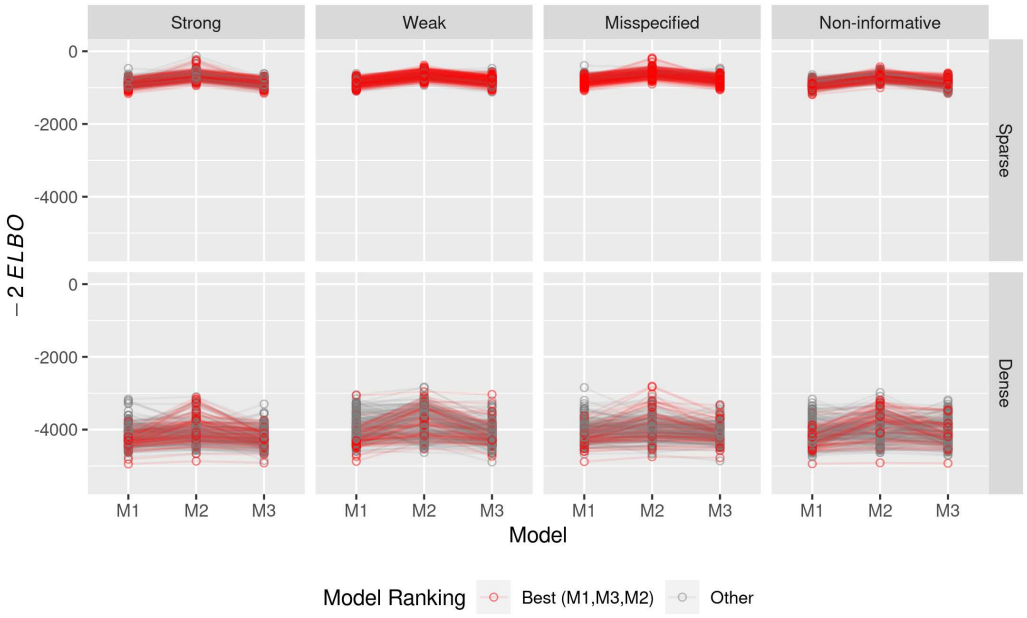

Figure 3.87: Two-Compartment Scenario, Model Selection Strategy 2 (ELBO) - ADVI $-2ELBO$ values [sensitivity analysis with $\beta_{Cl_{X1}} = 1$]

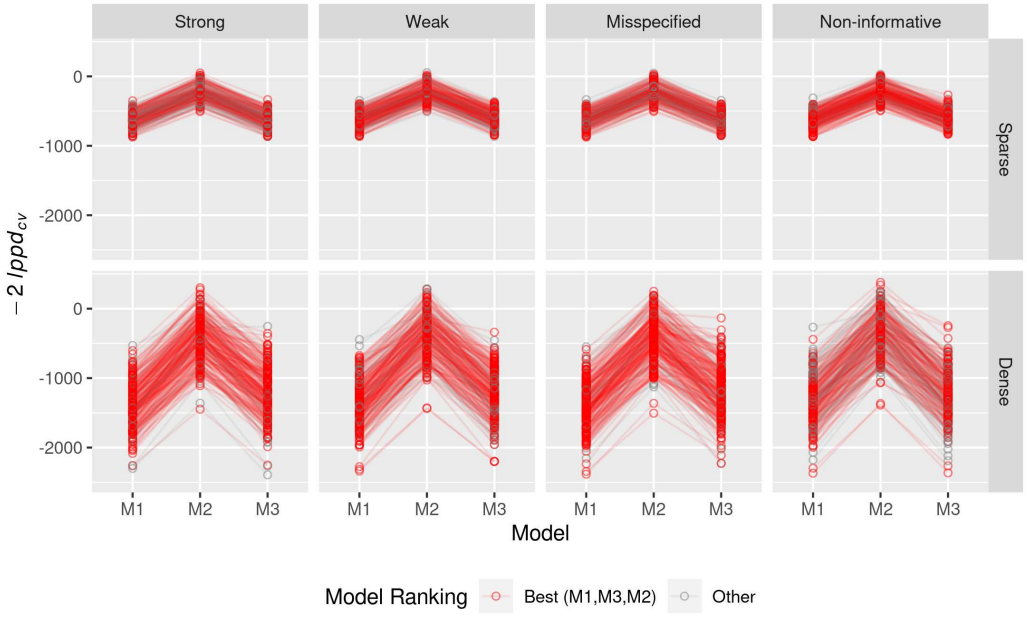

Figure 3.88: Two-Compartment Scenario, Model Selection Strategy 3 (5-fold leave-subject-out cross-validation) - ADVI $-2lppd_{cv}$ values [sensitivity analysis with $\beta_{Cl_{X1}} = 1$]

3.6.3 Case Study

Table 3.17: Ketorolac Case Study Priors

Parameter	Prior Distribution
Cl_{pop}	$logNormal(\mu = \log(1.17) + 0.6^2, \sigma = 0.6)$
Q_{pop}	$logNormal(\mu = \log(3.21) + 0.7^2, \sigma = 0.7)$
$V1_{pop}$	$logNormal(\mu = \log(168.27) + 0.7^2, \sigma = 0.7)$
$V2_{pop}$	$logNormal(\mu = \log(119.95) + 0.7^2, \sigma = 0.7)$
TM_{50}	$trunc - Normal(min = 0, \mu = 50, \sigma = 14)$ (for the second model)
$Hill$	$logNormal(\mu = \log(3) + 0.5^2, \sigma = 0.5)$ (for the second model)
ω_{Cl}	$logNormal(\mu = \log(\sqrt{0.5}) + 0.6^2, \sigma = 0.6)$
ω_{Q}	$logNormal(\mu = \log(\sqrt{0.5}) + 0.6^2, \sigma = 0.6)$
ω_{V1}	$logNormal(\mu = \log(\sqrt{0.5}) + 0.6^2, \sigma = 0.6)$
ω_{V2}	$logNormal(\mu = \log(\sqrt{0.5}) + 0.6^2, \sigma = 0.6)$
ρ	$LKJ(\eta = 1)$
σ_{prop}	$Half - Cauchy(\mu = 0, \sigma = 0.2)$
σ_{add}	$Half - Cauchy(\mu = 0, \sigma = 0.2)$

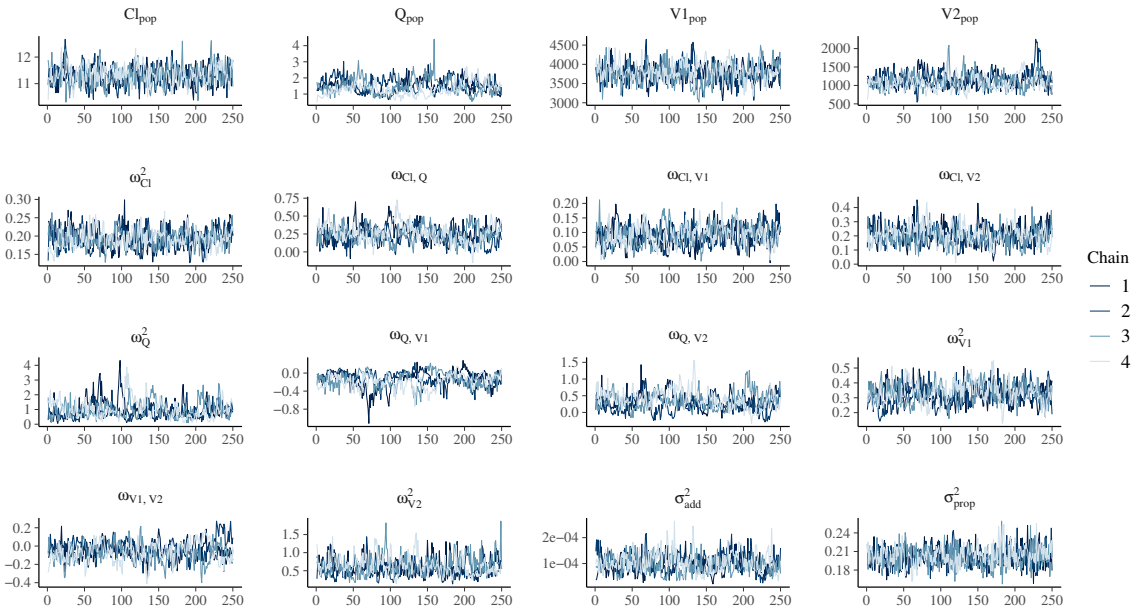

Figure 3.89: MCMC trace plots for ketorolac model with fixed allometric scaling only

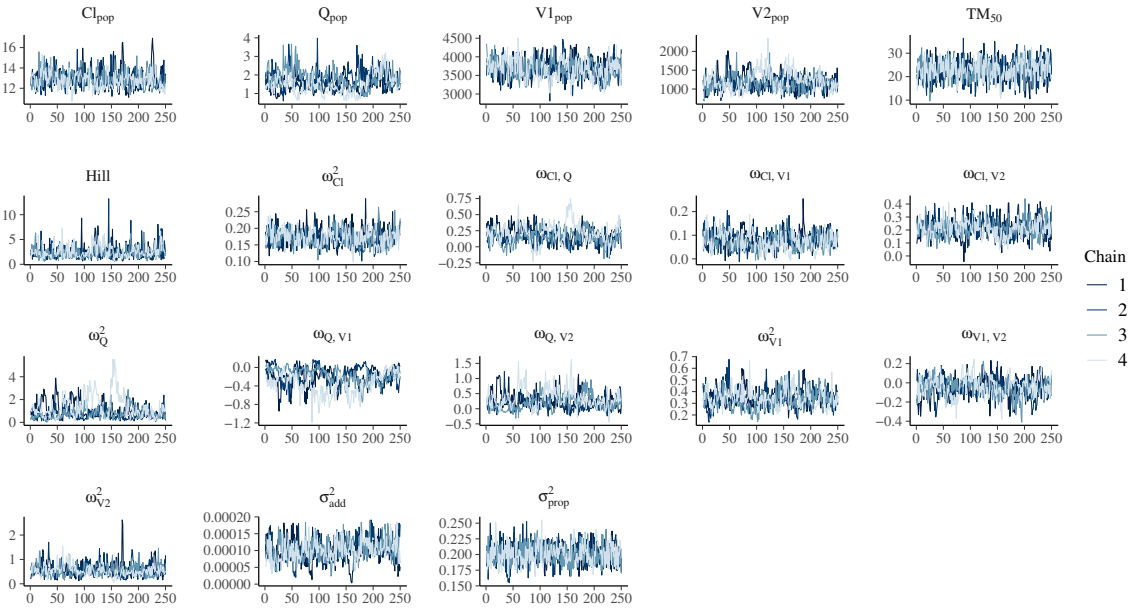

Figure 3.90: MCMC trace plots for ketorolac model with fixed allometric scaling and maturation factor

CHAPTER 4

Population Pharmacokinetic Analysis of Dexmedetomidine in Children using Real World Data[1]

4.1 Introduction

Dexmedetomidine is an $\alpha 2$-agonist with anxiolytic, sedative, and analgesic properties with minimal effects on respiratory depression [111,112]. It is routinely used as part of intraoperative anesthetic management during surgical repairs of congenital heart disease (CHD) and in the postoperative period in the intensive care unit (ICU) [113,114] and is commonly dosed as a continuous intravenous (IV) infusion using a fixed weight-based rate (e.g., starting at 0.3 mcg/kg/h). This dosing regimen will be adequate for some, but necessarily results in inappropriately low or high dosing for others. The proper dose for these latter individuals is not achieved until the initial sedation effect is observed, recognized as inadequate or excessive by the clinical team, and the dose adjusted accordingly. These patients are at risk for dose-related dexmedetomidine side effects, including bradycardia and hypotension, or use of additional sedative agents, including opioid analgesics. Accurate prediction of an individual's dexmedetomidine requirement (precision dosing) could help minimize titration time to achieve sedation and analgesia goals without toxicity.

Many population pharmacokinetic (PK) studies of dexmedetomidine in pediatric populations have been reported [115–127]. For example, Potts et al. (2009) [120] report on dexmedetomidine use in 95 pediatric ICU patients using data pooled from several previous studies, Su et al. (2016) [123] studied 59 children on mechanical ventilation after open heart surgery, Pérez-Guillé et al. (2018) [119] assessed 30 children undergoing ambulatory surgery, and Zuppa et al. (2019) [127] examined dexmedetomidine PK among 119 children undergoing cardiac surgery. Most have a small number of individuals and frequent specimen collection. For pediatric ICU populations, the median sample size is 29.5 (range 18-119), and the median number of total drug levels collected is 236.5 (range 89 – 1967) with a median of 9 per subject (range 2 – 16) [116,117,119,122–124,126,127]. Some studies have addressed small sample size with methods that combine information from multiple populations including pooled pediatric analyses [120], creating "universal" models for children and adults [115,118], and Bayesian analyses with informative priors [125]; however, even these models only include information from at most around 130 children.

Previously, studies have identified weight [115–120,122–127] and age [115,117,118,120,123–127], along

[1]This chapter is adapted from "Population Pharmacokinetic Analysis of Dexmedetomidine in Children using Real World Data from Electronic Health Records and Remnant Specimens" published in the *British Journal of Clinical Pharmacology* [110] and has been reproduced with the permission of the publisher and my co-authors Joseph H. Breeyear, Richard Caprioli, Todd Edwards, Brian Hachey, Prince J. Kannankeril, Jacob M. Keaton, Matthew D. Marshall, Sara L. Van Driest, and Leena Choi

with cardiac bypass [123,126,127], as important factors to explain inter-individual variability. However, lower sample size may limit identification of additional covariates impacting inter-individual variability. For example, although it is known that dexmedetomidine is rapidly distributed and metabolized in the liver by two pathways – direct glucuronidation by uridine 5'-diphosphate- glucuronosyltransferase (UGT) 1A4 and 2B10 and cytochrome P450 (CYP) 2A6 mediated aliphatic hydroxylation [112,128] – small studies of the impact of genetic variation or expression levels of these enzymes have failed to demonstrate pharmacogenetic associations [129,130]. A study including 260 children demonstrated that carriers of the T allele of *CYP2A6* rs835309 had significantly lower concentrations of dexmedetomidine (TT+TG vs. GG, p-value = .025) [131]. A newly developed weighted genetic risk score to predict *CYP2A6* activity raises the possibility of better capturing the impact of variants across this gene for pharmacogenetic analysis [132]. Study of a larger cohort may allow the identification of genetic biomarkers affecting dexmedetomidine PK, facilitating precision dosing based on genotype.

We combined data from electronic health records (EHRs) and remnant specimens collected during usual clinical care to perform a population PK analysis, similar to two previous pediatric fentanyl studies, and employing a system for constructing PK analysis datasets in R [133–136]. The major goals of this study were to develop a dexmedetomidine population PK model for children with data obtained from EHRs and remnant specimens and quantify genetic effects that were selected a priori based on previous studies and known metabolic pathways.

4.2 Methods

4.2.1 Study Design

This study was approved by the Vanderbilt University Medical Center (VUMC) Institutional Review Board and has been previously described [136]. In brief, pediatric patients undergoing surgery for CHD are offered enrollment in this observational study. Parents provide written consent for their child's participation, and informed assent is obtained when appropriate. Drug selection and dosing are determined by the primary clinical team; over the course of study enrollment, clinical leadership provided recommended protocols to guide clinicians in drug and dose selection for analgesia and sedation (included in the Appendix); however, final regimens were always at the discretion of the treating clinicians. Remnant specimens from clinical testing are obtained for drug concentration measurements, which are not disclosed to the clinical teams. Specimens were not collected in connection with dose administration or to monitor PK characteristics such as trough concentration or C_{max}. Enrollment with remnant specimen collection began in July 2012 and is ongoing. Data analyzed for this study were collected prior to October 2017. All study participants were

admitted to the pediatric cardiac ICU after surgery. Enrolled participants were excluded from the analysis if their surgery was cancelled, if there was missing genotype data, if extracorporeal membrane oxygenation (ECMO) treatment was required, or if they did not survive to hospital discharge. For those with multiple surgeries, data from the procedure with the highest number of measured serum drug concentrations were used, excluding all others. Drug concentrations were excluded if inadequate internal standard concentrations were detected and insufficient volume remained to repeat analysis, or if they were obtained before any documented dexmedetomidine dosing.

4.2.2 Data Collection

Demographic data (including parent-reported race) and medical history were documented at the time of study enrollment. Surgical and clinical data were extracted from the EHR prospectively. Dexmedetomidine dosing, including scheduled boluses, as-needed intermittent boluses, and continuous infusions after the postoperative admission to the ICU were determined from the EHR and the Vanderbilt Enterprise Data Warehouse. The Enterprise Data Warehouse contains nurse administration, nurse flowsheets, and pharmacy dispense data, enabling the computation of administered drug amounts over specific time periods. Study data were collected and managed using REDCap electronic data capture tools, a secure, web-based application hosted at Vanderbilt University [137].

4.2.3 Drug Concentration Measurement

For the purposes of drug concentration analysis, all remnant plasma specimens ≥ 100 μL from blood obtained for clinical testing of electrolyte or basic metabolic panels in study subjects were obtained from the Vanderbilt Clinical Chemistry Laboratory and stored at -20ºC until processing for drug concentration analysis in the Vanderbilt Mass Spectrometry Research Core. Specimen processing and mass spectrometry analysis have been previously described in detail [136]. Briefly, acetonitrile precipitation was followed by tandem mass spectrometry using a 16-drug assay. Dexmedetomidine assay accuracy is 89 – 112%, and the lower and upper limits of quantification (LLOQ and ULOQ) are 0.005 and 5 ng/mL.

4.2.4 Genotyping and *CYP2A6* Activity Score Prediction

Study participants provided a peripheral blood sample for genetic analysis. Genomic DNA was extracted through the Vanderbilt Technologies for Advanced Genomics (VANTAGE) Core laboratory and study participants were genotyped using either the Axiom™ Precision Medicine Research Array or the Precision Medicine Diversity Array according to manufacturer protocols at the Children's Hospital of Philadelphia DNA core. As part of genotype data quality control, variants were removed if genotype call rate was <98%, if minor allele

frequency was >20% different from 1000 Genomes phase 3 European reference populations, or for deviation from Hardy-Weinberg Equilibrium (p–value $< 1 \times 10^{-10}$, results shown in Table 4.7). Individuals were removed if their genotype call rate was <98%, the genetically estimated sex differed from parental-reported sex, or for relatedness (2nd degree or closer). Genotype data were imputed to the 1000 Genomes phase 3 reference panel. All non-synonymous, stop gain, and splice site variants in $UGT2B10$ and $UGT1A4$ with minor allele frequency >0.005 were included in the analysis ($UGT2B10$ variants: rs2942857, rs112561475, rs61750900; $UGT1A4$ variants rs2011425, rs3892221, rs6755571). For $CYP2A6$, enzyme activity was predicted using a polygenic score based on rs56113850, rs2316204, rs113288603, rs28399442, rs1801272, and rs28399433, using a previously published method [132].

4.2.5 Data Processing

Data was processed using the modularized *EHRtoPKPD* system for postmarketing population PK studies with real-world data from EHRs [133]. This system has been implemented in the `R` software [138] package `EHR` [135] which includes interactive checks for data quality to reconcile missing, duplicate, and other erroneous concentration or dosing information. Output from the `EHR` package was further cleaned by removing: (i) concentration measurements more than 150 hrs (approximately 50 times dexmedetomidine half-life) after the end of the final bolus or infusion dose, (ii) concentration measurements below the LLOQ if they are after the final bolus or infusion dose, except for the first such measurement, (iii) concentration measurements above the ULOQ, (iv) subjects whose only concentration measurements are below the LLOQ after applying criteria (i)-(iii), and (v) subjects with missing dose information indicated by increases in concentration without an accompanying dose and confirmed by manual chart review. Serum creatinine concentration was a time-varying covariate typically measured concurrently with dexmedetomidine concentration. If serum creatinine was not available when dexmedetomidine concentration was measured, we selected the serum creatinine concentration measured closest to the dexmedetomidine concentration data within 7 days. For each subject, weight varied little within the timeframe of available concentration data, so most weight data were the same as the baseline demographic measurements. When additional weight measurements were available, usually during infusion, weight measurements obtained at the same time as the dosing event were used. Measures of albumin concentration were available within a 7-day window for only 48 subjects, precluding use of albumin concentration as a covariate.

4.2.6 Population Pharmacokinetic Analysis

We performed population PK analysis of dexmedetomidine using nonlinear mixed-effects models implemented by `Monolix` (version 2020R1) [139] and estimated the parameters with the stochastic approximation

expectation-maximization (SAEM) algorithm. Observed concentrations below the LLOQ were considered to be censored between 0 and 0.005 ng/mL and were handled in the modeling using the likelihood (M3) method for interval censoring [140,141]. After the model parameters were estimated with SAEM, the objective function value (OFV) was calculated using Monte Carlo importance sampling with 10,000 samples from a Student-t proposal distribution and degrees of freedom chosen by testing a sequence of values (ν=1, 2, 5, 10, 15). Because the SAEM estimation method includes stochastic variability and can sometimes fail to converge in a setting with sparse sampling [142], we performed 5 runs with different random seeds for each model and selected the run with median OFV for model comparison. Additional details on the SAEM algorithm settings are included in the Appendix.

For model selection we used a likelihood ratio test to compare differences in estimated OFV for nested models and corrected Bayesian Information Criteria (BICc) to compare non-nested models [143]; relative standard errors, parameter estimate values, magnitude of random effects and change in CV% were also considered. In addition, we used several graphical methods for model evaluation including observed vs. population and individual predictions, individual weighted residuals vs. predicted concentration and time, correlations between samples from the conditional random effects distributions, samples from the conditional random effects distributions vs. covariates, and prediction corrected visual predictive checks [144,145].

All covariates were chosen a priori based on previous research and biological plausibility, including $UGT1A4$, $UGT2B10$, and $CYP2A6$ variants, age, sex, Society of Thoracic Surgery–European Association for Cardio-Thoracic Surgery (STAT) Congenital Heart Surgery Mortality score [146], cardiac bypass time, length of ICU stay, and serum creatinine. We explored the effects of covariates on total clearance and volume of distribution parameters, and graphical checks were examined for possible relationships between covariates and all main population PK parameters.

Model building proceeded in two stages; we first considered all covariates except $UGT1A4$, $UGT2B10$, and $CYP2A6$ to build an adequate model for dexmedetomidine PK and then examined the hypothesized association between the genotype variables and total clearance by adding these effects individually to the stage one model. For stage one we explored models with various structural, residual error, and inter-individual variance components and adjusted for non-genotype covariates. Following a strategy outlined by Bonate (2011), we began with richly parameterized inter-individual variability and residual error models including all random effects, all correlations between random effects, and combined additive and proportional residual error, and then simplified this structure [55]. We examined the structural model by comparing one- and two-compartment models without covariates. Following this we considered size and age maturation; these two covariates have been shown to be important factors in pediatric PK models with a large age range and in previous

dexmedetomidine studies [115,118,120,125,147–149]. For size, we employed an allometric weight model with fixed or estimated scaling parameters. For maturation, we considered an exponential age model, a sigmoid Hill maturation model, a body-weight dependent exponent model, and an age-dependent exponent model [150]. Next, we investigated whether other non-genotype covariates improved the model with size and maturation factors and refined the residual and inter-individual variance structure. Each covariate was considered as an exponentially linear or categorical term.

In the second stage we tested for the association between genotype and total clearance by including these effects in the model found in stage one. For $UGT1A4$ and $UGT2B10$, dichotomous models (coding individuals as having a loss-of-function variant or not) and additive models (counting the number of variants) were considered. For $CYP2A6$, the predicted enzyme activity was included as an exponential term [132]. Details of all models explored along with specific mathematical relationships, estimated OFV, and BICc are shown in Tables 4.10 – 4.17. Graphical checks for the model selection process are shown in Figures 4.11 – 4.35.

4.3 Results

4.3.1 Study Population and Specimens

We collected 4,369 residual plasma specimens from 620 subjects. After removing 89 subjects with unknown sample collection time, 108 subjects with no dosing information within 7 days of the first concentration measurement, and 12 subjects due to in hospital mortality or ECMO, the output of the EHR package pipeline contained 411 subjects with 2,172 dexmedetomidine concentration measurements. The further cleaning steps described above removed 14 subjects and 43 more subjects without genotype information were also removed. The study cohort flow diagram of data processing is shown in Figure 4.1, and the final study population of 354 subjects with 1,400 specimens is described in Tables 4.1 and 4.2. The median postnatal age was 16 months (interquartile range [IQR] 5 – 62), median postmenstrual age was 105 weeks (IQR 62 - 304) and median weight was 9.4 kg (IQR 6.0 – 18.2). The age and weight distributions are shown in Figures 4.9 and 4.10 and Table 4.8 shows postnatal age categories. There were 262 subjects (74%) with no variants of $UGT1A4$, 87 (25%) with 1 variant and 5 (1%) with 2 variants. For $UGT2B10$, 186 subjects (53%) had no variants, 117 (33%) had 1 variant and 51 (14%) had 2 or 3 variants. The $CYP2A6$ predicted activity score was available for 350 of the 354 subjects (median 2.04, IQR 2.00 – 2.21). There were 2,386 dexmedetomidine dosing events including 2,351 IV infusions and 35 bolus administrations. The median infusion rate was 0.6 mcg/kg/hr (IQR 0.5 – 1.0) and the median infusion duration was 2 hours (IQR 1 - 5); the median bolus dose was 1.0 mcg/kg (IQR 0.96 – 1.01). The top ten concomitant medications were acetaminophen (92.5%), cefazolin (92%), famotidine (89.1%), morphine (88.5%), furosemide (80.5%), fentanyl (77.9%), rocuronium (69.5%),

oxycodone (59.8%), heparin (58.3%), and lorazepam (51.4%). Table 4.9 includes all concomitant medications administered to at least 5% of subjects. The number of dexmedetomidine concentration measurements per subject varied from a minimum of 1 to a maximum of 18 with a median of 3 specimens (IQR 2 – 5). The median time of first dexmedetomidine measurement after dose start was 5 hours (IQR 4 – 11) and the median time of final dexmedetomidine measurement after dose start was 68 hours (IQR 39 – 131).

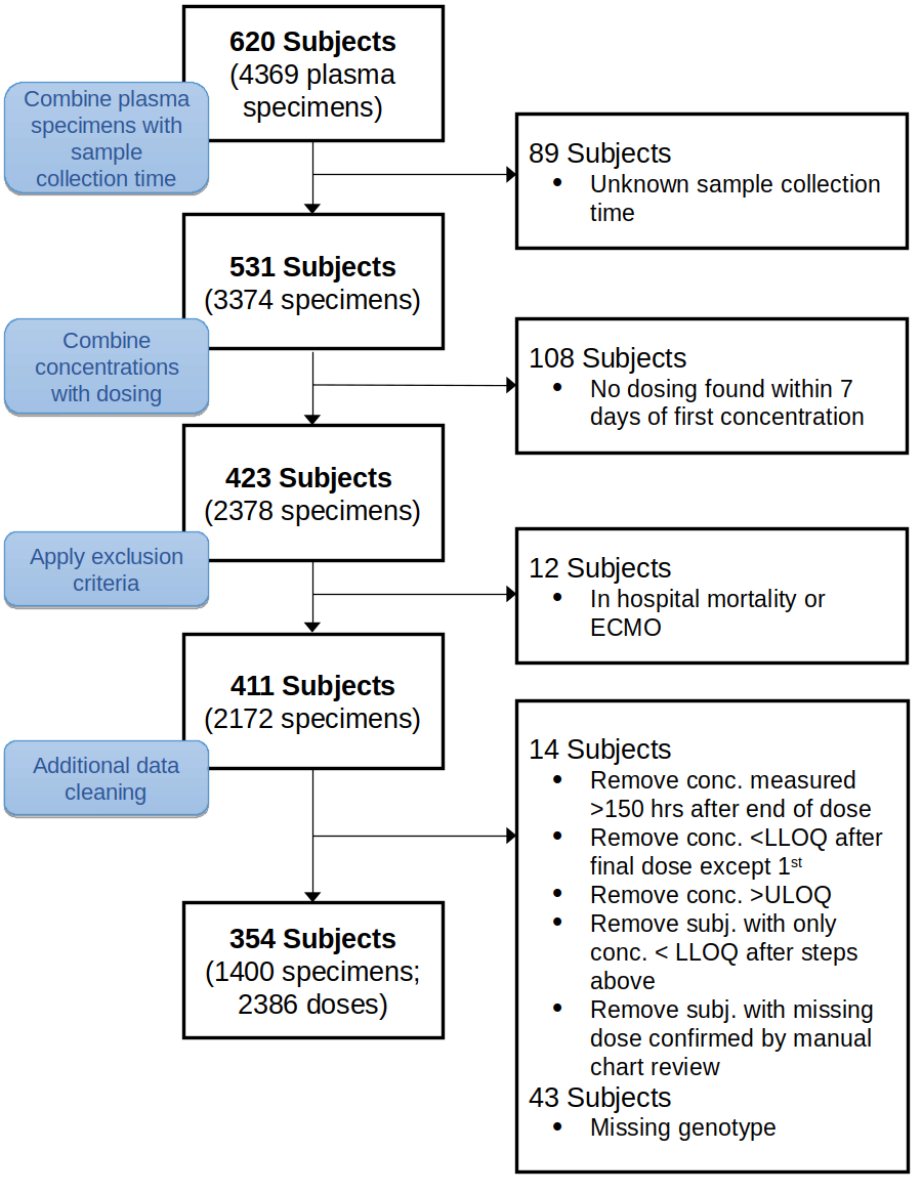

Figure 4.1: Study Cohort Flow Diagram in Data Processing with Exclusion Criteria. ECMO, extracorporeal membrane oxygenation; LLOQ, lower limit of quantification; ULOQ, upper limit of quantification.

Table 4.1: Dexmedetomidine Study Cohort Characteristics

	Cohort
n	354
Postnatal Age (months)	15.7 (5.34-61.77) [0.03-270.9]
Postmenstrual Age (weeks)	104.64 (61.75-303.75) [39.14-1200]
Median weight (kg)	9.41 (6-18.15) [2-138]
Gender	
Male	183 (52%)
Female	171 (48%)
Race	
White	292 (83%)
Black	40 (11%)
American Indian or Alaska Native	2 (1%)
Asian	6 (2%)
Other	5 (1%)
Unknown	8 (2%)
Median serum creatinine (mg/dL)	0.49 (0.44-0.56) [0.24-1.12]
STAT score	
1	154 (44%)
2	108 (31%)
3	41 (12%)
4	46 (13%)
5	5 (1%)
***UGT1A4* variants**	
0	262 (74%)
1	87 (25%)
2	5 (1%)
***UGT2B10* variants**	
0	186 (53%)
1	117 (33%)
2 or 3	51 (14%)
CYP2A6* score	2.04 (2-2.21) [1.58-2.43]
Cardiac bypass time (mins)	100 (71-143.75) [0-426]
Length of ICU hospitalization (days)	4 (2-6) [1-120]

Continuous variable summary statistics: median (interquartile range) [minimum, maximum]; Categorical variable summary statistics: n (%); ICU, Intensive Care Unit; STAT, Society of Thoracic Surgery–European Association for Cardio-Thoracic Surgery Congenital Heart Surgery Mortality score

* Among n = 350 subjects with available score.

Table 4.2: Dexmedetomidine Dosing and Specimen Sampling

	Cohort
Total dosing events*	2386
IV infusion doses	2351
Infusion duration (mins)	120 (60-306) [1-7536]
Infusion rate (mcg/kg/hr)	0.6 (0.5-1) [0.03-2]
Bolus doses	35
Bolus dose amount (mcg/kg)	1 (0.96-1.01) [0.06-4.21]
Dosing events per subject	5 (3-8) [1-37]
Total dexmedetomidine concentration measurements	1400
Concentration measurements below lower limit of quantification	120
Concentration measurements per subject	3 (2-5) [1-18]
First concentration measurement time after dose start (hr)	5.29 (3.73-11.27) [0-178.5]
Final concentration measurement time after dose start (hr)	67.62 (39.31-130.54) [2-659.87]

Continuous variable summary statistics: median (interquartile range) [minimum, maximum]

* A dosing event is defined as a bolus administration or an infusion interval with constant administration rate for a specific duration.

4.3.2 Population Pharmacokinetic Model

In the first stage of modeling, a two-compartment model with additive and proportional residual error was chosen as the base model based on BICc and graphical checks. The main PK parameters are total clearance (CL, L/h), volume of distribution for the central compartment ($V1$, L), inter-compartmental clearance (Q, L/h) and volume of distribution for the peripheral compartment ($V2$, L). The results for the base and covariate models without genotype are presented in Table 4.18. The coefficients of variation (CV) for CL, $V1$, Q, and $V2$ in the base model were 201%, 161%, 146%, and 672%, respectively. Including weight as covariate for all PK parameters with fixed allometric scaling parameters substantially improved the model fit (both OFV and BICc decreased by 406 from the base model, Table 4.18) and plots of individual predicted vs. observed concentration, individual weighted residuals vs. predicted concentration and random effects vs. covariates also improved (Figures 4.16 – 4.20). The CV for CL, $V1$, Q, and $V2$ were 123%, 168%, 91%, and 857%, respectively. Using estimated allometric parameters did not improve the model fit. Among the four models adjusting for both weight and age maturation, the model with sigmoid postmenstrual age maturation had the largest improvement in BICc compared to the model with only weight (difference of 8.2, Table 4.12). This model was further simplified by estimating models with fixed effects for $V2$ or $V2$ and Q, no additive residual error component, and several correlation structures for the covariance between random effects (Table 4.13). While the estimated CV% was large for some parameters, especially for $V2$, the models with fixed effects for $V2$ or $V2$ and Q had significantly worse model fit (both OFV and BICc increased by greater than 400). The final model without genetic effects includes proportional residual error, fixed allometric scaling for all parameters and sigmoid (Hill) maturation for total clearance, random effects for all PK parameters and correlation only between the random effects of CL and V_1 (Table 4.3). The condition number for this simplified model (calculated as the ratio of the maximum to the minimum eigenvalues) was 26.03, suggesting our final model was not ill-conditioned [55]. No further improvement was seen by adding other covariates including either form of $UGT1A4$ or $UGT2B10$ or predicted $CYP2A6$ activity score (Tables 4.4, 4.5, and 4.6) to the clearance submodel. No strong covariate effects were seen for V_1, Q, or $V2$ based on graphical goodness-of-fit plots and adding covariates to the V_1 or $V2$ submodels did not meaningfully improve the model. The final

covariate model is described as follows:

$$
\begin{aligned}
CL_i &= \theta_1 \times \left(\frac{WT_i}{70}\right)^{0.75} \times \frac{1}{1+\left(\frac{TM_{50}}{PMA_i}\right)^{Hill}} \times \exp(\eta_i^{CL}) \\
V1_i &= \theta_2 \times \left(\frac{WT_i}{70}\right) \times \exp(\eta_i^{V1}) \\
Q_i &= \theta_3 \times \left(\frac{WT_i}{70}\right)^{0.75} \times \exp(\eta_i^{Q}) \\
V2_i &= \theta_4 \times \left(\frac{WT_i}{70}\right) \times \exp(\eta_i^{V2})
\end{aligned}
$$

where CL_i, $V1_i$, Q_i, and $V2_i$ are the individual-specific PK parameters corresponding to CL, $V1$, Q, and $V2$, WT_i is subject weight in kilograms (kg); and PMA_i is subject postmenstrual age in weeks. The η_i^{CL}, η_i^{V1}, η_i^{Q}, and η_i^{V2} are random effects explaining inter-individual variability for the PK parameters which follow a normal distribution with mean zero and variance of ω_{CL}^2, ω_{V1}^2, ω_Q^2, and ω_{V2}^2, respectively. The θs are estimated model parameters. Diagnostic plots for the final model are shown in Figures 4.2 - 4.7. The plot of observed dexmedetomidine concentrations vs. population predictions reflects the relatively large inter-individual variability and potential misspecification for small concentration values, however no major deficiencies in the structural or error models are seen when comparing observed and individual predicted concentrations. No trends were detected in plots of individual weighted residuals vs. predicted concentration or time. The prediction-corrected visual predictive check showed good agreement between the observed and theoretical median and 90th percentiles; model misspecification is seen for the 10th percentile where many values are below the LLOQ and were simulated from the estimated model and where data are sparse (e.g., times greater than 5 days after first dose).

The estimates of CL, $V1$, Q and $V2$ in terms of a standard weight of 70 kg are shown in Table 4.3: CL (θ_1) = 27.3 L/h, $V1$ (θ_2) = 161 L, Q (θ_3) = 26.0 L/h, and $V2$ (θ_4) = 7903 L. We estimate CL, $V1$, Q and $V2$ as 6.04 L/h, 21.6 L, 5.7 L/h, and 1061.26 L for a child at the median weight of 9.4 kg and median postmenstrual age of 104.6 weeks; After including covariates, the CV for CL was substantially reduced from 201% estimated in the base model to 123% in the weight only model to 103% in the final model. CV remains large for some parameters, especially $V2$ (624%) and $V1$ (138%), likely due to lack of data to accurately estimate them. Despite this, allowing random effects for all the main PK parameters results in a significantly better model fit than assuming no variability for these parameters.

Table 4.3: Estimates from Final Model without Genetic Covariates

Weight and Maturation with simplified variance structure (Obj = -1915.1, BICc = -1835.0)		
Parameters	Estimates (SE) [95% CI][a]	Relative SE (%)
$Cl = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}} \right)$		
θ_1	$27.3\,(1.82)\,[24.0, 31.1]$	6.64
TM_{50}	$41.9\,(0.28)\,[41.4, 42.5]$	0.65
$Hill$	$7.04\,(0.022)\,[6.99, 7.08]$	0.30
$V_1 = \theta_2 (WT/70)$		
θ_2	$161\,(12.1)\,[139, 187]$	7.47
$Q = \theta_3 (WT/70)^{0.75}$		
θ_3	$26.0\,(1.90)\,[22.5, 30.0]$	7.29
$V_2 = \theta_4 (WT/70)$		
θ_4	$7903\,(1408)\,[5617, 11119]$	17.8
$\omega_{Cl}\,(\%CV)$	$103\,(8)\,[88, 120]$	7.77
$\omega_{V1}\,(\%CV)$	$138\,(13)\,[114, 166]$	9.48
$\omega_{Q}\,(\%CV)$	$82\,(9)\,[65, 102]$	11.6
$\omega_{V2}\,(\%CV)$	$624\,(157)\,[391, 1048]$	25.1
$\rho_{Cl,V1}$	$0.923\,(0.027)\,[0.871, 0.975]$	2.92
other ρ terms	0 (fixed)	-
$\sigma_{add}\,(ng/mL)$	0 (fixed)	-
$\sigma_{prop}\,(\%CV)$	$50.5\,(1.6)\,[47.5, 53.5]$	3.16

Table 4.4: Estimates from Categorical Gene Models for $UGT1A4$ or $UGT2B10$

$UGT1A4$ **categorical gene model** (Obj = -1918.3, BICc = -1832.4)			$UGT2B10$ **categorical gene model** (Obj = -1913.8, BICc = -1827.9)		
Parameters	Estimates (SE) [95% CI][a]	Relative SE (%)	Parameters	Estimates (SE) [95% CI]	Relative SE (%)
$Cl = \theta_1 (WT/70)^{0.75} \times \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}} \right) \times \exp(\theta_5 I[UGT1A4 > 0])$			$Cl = \theta_1 (WT/70)^{0.75} \times \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}} \right) \times \exp(\theta_5 I[UGT2B10 > 0])$		
θ_1	24.0 (2.19) [20.1, 28.7]	9.14	θ_1	20.4 (2.07) [16.8, 24.9]	10.1
TM_{50}	42.2 (0.041) [42.1, 42.3]	0.096	TM_{50}	45.7 (0.078) [45.6, 45.9]	0.16
$Hill$	4.17 (0.0016) [4.16, 4.17]	0.037	$Hill$	4.96 (0.0073) [4.95, 4.97]	0.14
θ_5	-0.221 (0.16) $[-0.54, 0.09]$	-72.5	θ_5	-0.104 (0.11) $[-0.32, 0.11]$	-106.4
$V_1 = \theta_2 (WT/70)$			$V_1 = \theta_2 (WT/70)$		
θ_2	169 (20.4) [133, 213]	12.0	θ_2	178 (12.6) [155, 204]	7.11
$Q = \theta_3 (WT/70)^{0.75}$			$Q = \theta_3 (WT/70)^{0.75}$		
θ_3	31.2 (2.74) [26.3, 37.1]	8.78	θ_3	33.2 (1.84) [29.8, 37.0]	5.52
$V_2 = \theta_4 (WT/70)$			$V_2 = \theta_4 (WT/70)$		
θ_4	14346 (2782) [9908, 20770]	19.3	θ_4	19655 (2804) [14921, 25889]	14.3
ω_{Cl} (%CV)	113 (13) [90, 140]	11.1	ω_{Cl} (%CV)	131 (13) [108, 158]	9.79
ω_{V1} (%CV)	139 (31) [88, 216]	22.5	ω_{V1} (%CV)	126 (12) [104, 152]	9.53
ω_{Q} (%CV)	67 (7) [54, 81]	10.4	ω_{Q} (%CV)	67 (6) [56, 79]	8.63
ω_{V2} (%CV)	640 (158) [404, 1067]	24.7	ω_{V2} (%CV)	544 (114) [367, 837]	21.0

Table 4.4: Estimates from Categorical Gene Models for $UGT1A4$ or $UGT2B10$ *(continued)*

Parameters	Estimates (SE) [95% CI][a]	Relative SE (%)	Parameters	Estimates (SE) [95% CI]	Relative SE (%)
$\rho_{Cl,V1}$	0.953 (0.043) [0.869, 1.0]	4.51	$\rho_{Cl,V1}$	0.939 (0.018) [0.904, 0.974]	1.91
other ρ terms	0 (fixed)	-	other ρ terms	0 (fixed)	-
$\sigma_{add}\,(ng/mL)$	0 (fixed)	-	$\sigma_{add}\,(ng/mL)$	0 (fixed)	-
$\sigma_{prop}\,(\%CV)$	50.6 (1.6) [47.5, 53.8]	3.16	$\sigma_{prop}\,(\%CV)$	50.7 (1.6) [47.6, 53.8]	3.15

Table 4.5: Estimates from Additive Gene Models for $UGT1A4$ or $UGT2B10$

$UGT1A4$ **additive gene model** (Obj = -1917.6, BICc = -1831.7)			$UGT2B10$ **additive gene model** (Obj = -1917.3, BICc = -1831.4)		
Parameters	Estimates (SE) [95% CI][a]	Relative SE (%)	Parameters	Estimates (SE) [95% CI]	Relative SE (%)
$Cl = \theta_1 (WT/70)^{0.75} \times \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \times \exp(\theta_5 UGT1A4)$			$Cl = \theta_1 (WT/70)^{0.75} \times \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \times \exp(\theta_5 UGT2B10)$		
θ_1	23.8 (2.47) [19.5, 29.1]	10.3	θ_1	23.5 (2.43) [19.2, 28.7]	10.3
TM_{50}	47.6 (0.136) [47.3, 47.9]	0.28	TM_{50}	41.6 (0.206) [41.2, 42.0]	0.49
$Hill$	5.55 (0.013) [5.53, 5.58]	0.23	$Hill$	17.6 (0.052) [17.5, 17.7]	0.29
θ_5	−0.166 (0.12) [−0.406, 0.073]	-73.4	θ_5	−0.108 (0.068) [−0.241, 0.025]	-63.1
$V_1 = \theta_2 (WT/70)$			$V_1 = \theta_2 (WT/70)$		
θ_2	171 (16.2) [142, 206]	9.48	θ_2	161 (14.2) [136, 191]	8.78
$Q = \theta_3 (WT/70)^{0.75}$			$Q = \theta_3 (WT/70)^{0.75}$		

Table 4.5: Estimates from Additive Gene Models for $UGT1A4$ or $UGT2B10$ *(continued)*

Parameters	Estimates (SE) [95% CI][a]	Relative SE (%)	Parameters	Estimates (SE) [95% CI]	Relative SE (%)
θ_3	31.6 (2.1) [27.7, 36.0]	6.67	θ_3	31.5 (2.2) [27.5, 36.0]	6.95
$V_2 = \theta_4(WT/70)$			$V_2 = \theta_4(WT/70)$		
θ_4	16576 (3229) [11431, 24038]	19.4	θ_4	12588 (2603) [8495, 18654]	20.6
$\omega_{Cl}\,(\%CV)$	110 (10) [91, 131]	9.28	$\omega_{Cl}\,(\%CV)$	109 (12) [88, 135]	10.8
$\omega_{V1}\,(\%CV)$	136 (18) [104, 177]	13.2	$\omega_{V1}\,(\%CV)$	139 (14) [113, 170]	10.4
$\omega_{Q}\,(\%CV)$	69 (7) [56, 84]	10.1	$\omega_{Q}\,(\%CV)$	64 (7) [52, 78]	10.2
$\omega_{V2}\,(\%CV)$	548 (120) [364, 858]	21.8	$\omega_{V2}\,(\%CV)$	792 (212) [482, 1377]	26.7
$\rho_{Cl,V1}$	0.933 (0.03) [0.88, 0.99]	3.10	$\rho_{Cl,V1}$	0.943 (0.04) [0.87, 1.0]	3.92
other ρ terms	0 (fixed)	-	other ρ terms	0 (fixed)	-
$\sigma_{add}\,(ng/mL)$	0 (fixed)	-	$\sigma_{add}\,(ng/mL)$	0 (fixed)	-
$\sigma_{prop}\,(\%CV)$	50.4 (1.6) [47.4, 53.5]	3.17	$\sigma_{prop}\,(\%CV)$	50.6 (1.6) [47.5, 53.7]	3.16

Table 4.6: Estimates from Models using Weight and Maturation with Simplified Variance Structure in $CYP2A6$ Subset

Without $CYP2A6$ score (Obj = -1889.2, BICc = -1809.3)			**With $CYP2A6$ score** (Obj = -1889.9, BICc = -1804.2)		
Parameters	Estimates (SE) [95% CI][a]	Relative SE (%)	Parameters	Estimates (SE) [95% CI]	Relative SE (%)
$Cl = \theta_1 (WT/70)^{0.75} \times \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}} \right)$			$Cl = \theta_1 (WT/70)^{0.75} \times \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}} \right) \times \exp(\theta_5 (CYP2A6score))$		
θ_1	30.5 (2.16) [26.5, 35.0]	7.09	θ_1	24.8 (14.7) [9.42, 65.2]	59.3
TM_{50}	41.5 (0.046) [41.4, 41.6]	0.11	TM_{50}	42.6 (0.098) [42.4, 42.8]	0.23
$Hill$	4.52 (0.011) [4.50, 4.54]	0.24	$Hill$	7.45 (0.015) [7.42, 7.48]	0.19
			θ_5	0.0885 (0.28) [−0.46, 0.64]	319.2
$V_1 = \theta_2 (WT/70)$			$V_1 = \theta_2 (WT/70)$		
θ_2	157 (16.1) [128, 191]	10.2	θ_2	152 (12.1) [130, 178]	7.97
$Q = \theta_3 (WT/70)^{0.75}$			$Q = \theta_3 (WT/70)^{0.75}$		
θ_3	24.8 (1.99) [21.2, 29.0]	8.00	θ_3	24.5 (1.78) [21.2, 28.2]	7.25
$V_2 = \theta_4 (WT/70)$			$V_2 = \theta_4 (WT/70)$		
θ_4	5720 (1129) [3925, 8335]	19.7	θ_4	5756 (1017) [4102, 8077]	17.6
$\omega_{Cl}\,(\%CV)$	103 (10) [86, 124]	9.38	$\omega_{Cl}\,(\%CV)$	100 (8) [86, 116]	7.70
$\omega_{V1}\,(\%CV)$	129 (20) [95, 173]	15.0	$\omega_{V1}\,(\%CV)$	138 (15) [111, 171]	10.9
$\omega_{Q}\,(\%CV)$	87 (10) [69, 109]	11.5	$\omega_{Q}\,(\%CV)$	86 (9) [69, 106]	10.8
$\omega_{V2}\,(\%CV)$	514 (118) [334, 825]	23.0	$\omega_{V2}\,(\%CV)$	549 (133) [349, 906]	24.2

Table 4.6: Estimates from Models using Weight and Maturation with Simplified Variance Structure in *CYP2A6* Subset *(continued)*

Parameters	Estimates (SE) [95% CI][a]	Relative SE (%)	Parameters	Estimates (SE) [95% CI]	Relative SE (%)
$\rho_{Cl,V1}$	0.945 (0.03) [0.88, 1.0]	3.28	$\rho_{Cl,V1}$	0.942 (0.02) [0.91, 0.98]	1.80
other ρ terms	0 (fixed)	-	other ρ terms	0 (fixed)	-
$\sigma_{add}\,(ng/mL)$	0 (fixed)	-	$\sigma_{add}\,(ng/mL)$	0 (fixed)	-
$\sigma_{prop}\,(\%CV)$	50.8 (1.6) [47.6, 54.0]	3.14	$\sigma_{prop}\,(\%CV)$	50.7 (1.6) [47.7, 53.8]	3.15

Notes for Tables 4.3 - 4.6:

SE, standard error; Obj, objective function value; BICc, corrected Bayesian information criteria; CL, total clearance (L/hr); Q, intercompartmental clearance (L/hr); V_1, volume of distribution for the central compartment (L); V_2, volume of distribution for the peripheral compartment (L); TM_{50} postmenstrual age at which clearance is 50% of adult value; $Hill$, maturation factor slope coefficient; CV, coefficient of variation; WT, body weight in kg; PMA, postmenstrual age in weeks; ω_{CL}, ω_{V1}, ω_Q, ω_{V2}, the standard deviation for η_i^{CL}, η_i^{V1}, η_i^Q, η_i^{V2}, respectively; For the standard deviation of random effects, ω, coefficient of variation was calculated as CV% = $100 \times \sqrt{\exp(\omega^2) - 1}$; ρ are correlation terms between random effects; σ_{prop} and σ_{add} are proportional and additive residual error terms.

[a] 95% asymptotic confidence intervals (CIs)

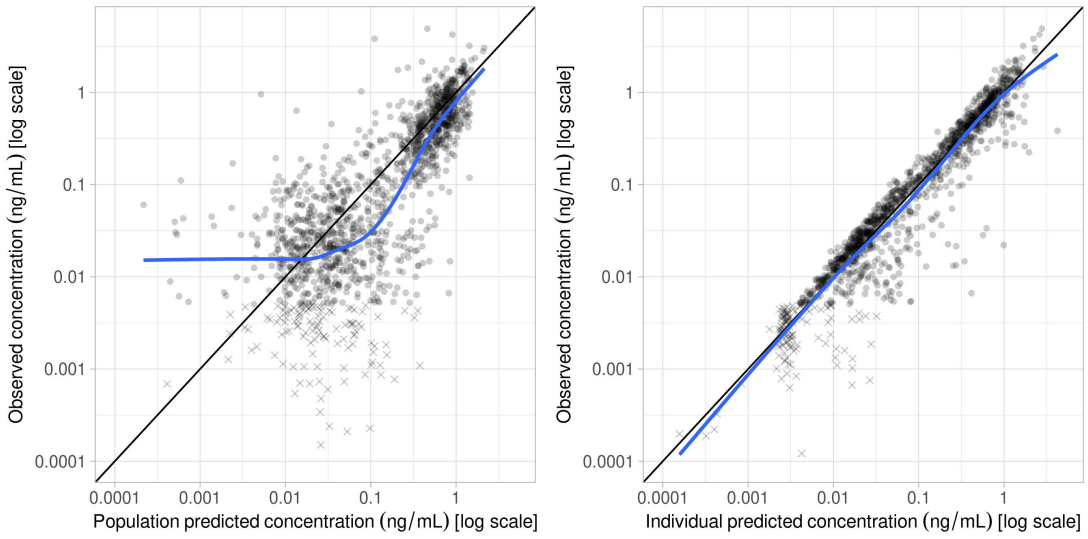

Figure 4.2: Diagnostic plots for the final population pharmacokinetic model. (A) Observed dexmedetomidine concentrations vs. population predicted and (B) individual predicted concentrations. Filled circles indicate observed values, x indicates simulated values based on the estimated model for observations below the lower limit of quantification (0.005 ng/mL), blue lines are loess smoothers and the solid black lines represent the line of identity.

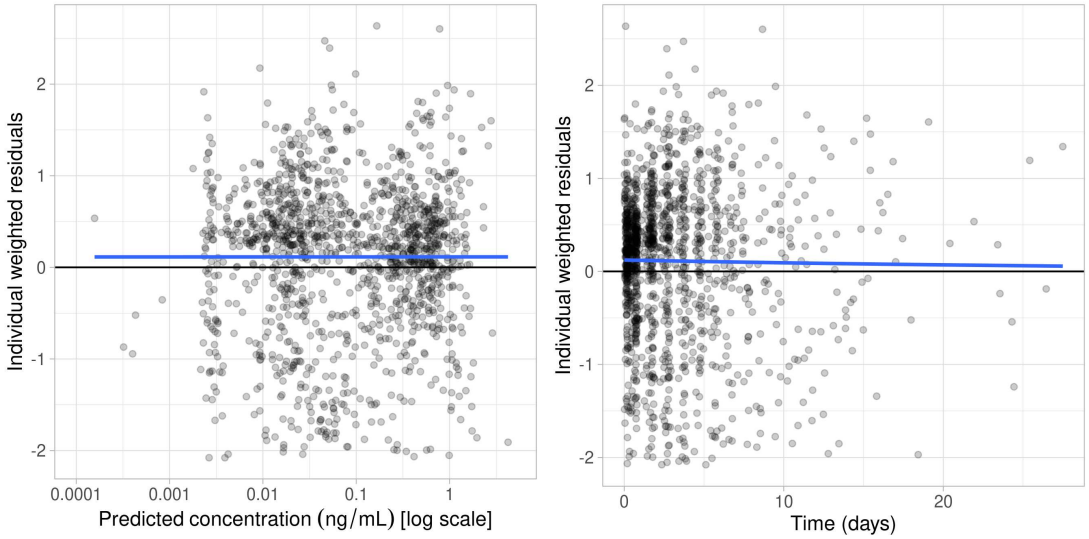

Figure 4.3: Diagnostic plots for the final population pharmacokinetic model. (A) Individual weighted residuals vs. predicted concentrations and (B) time. Blue lines are loess smoothers and the black horizontal lines at zero represent no trend.

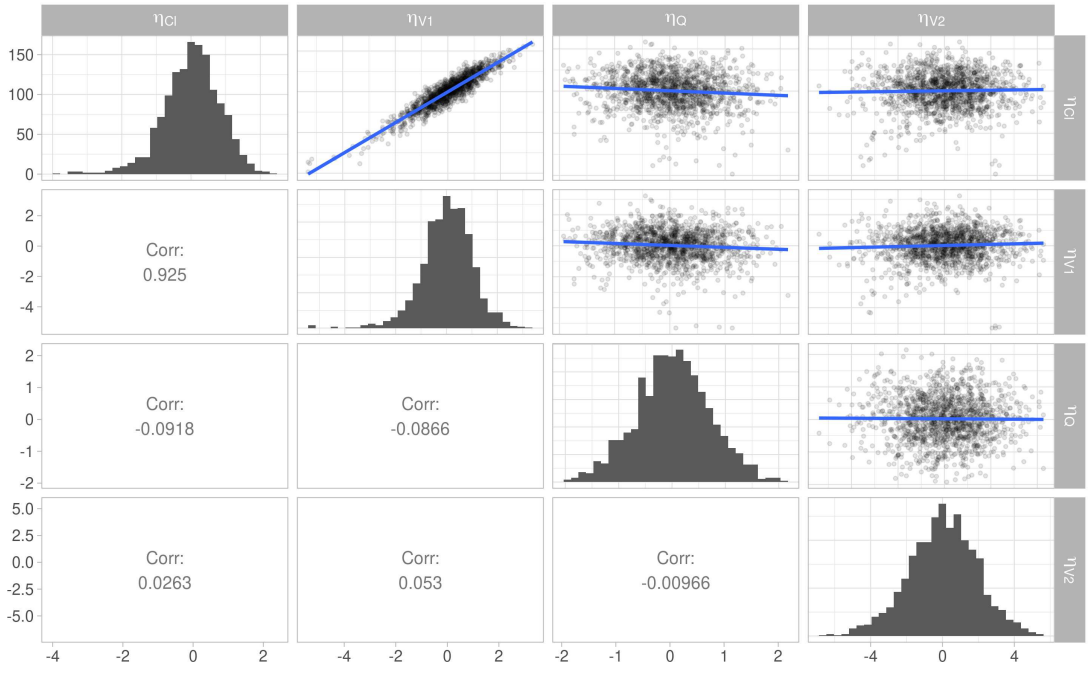

Figure 4.4: Diagnostic plots for the final population pharmacokinetic model. Correlation between random effects. Blue lines are least-squares fits.

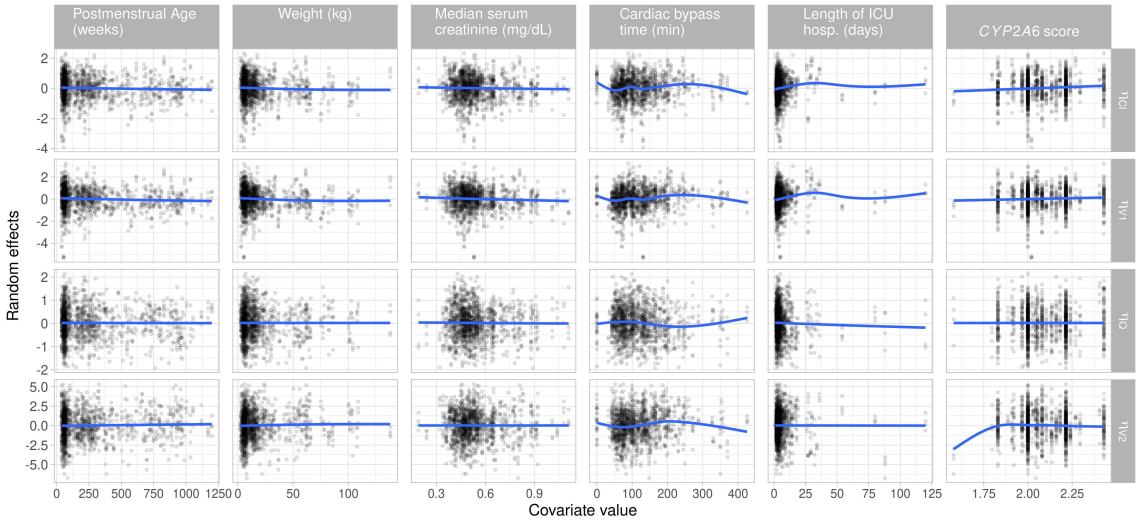

Figure 4.5: Diagnostic plots for the final population pharmacokinetic model. Random effects vs. continuous covariates. Blue lines are loess smoothers.

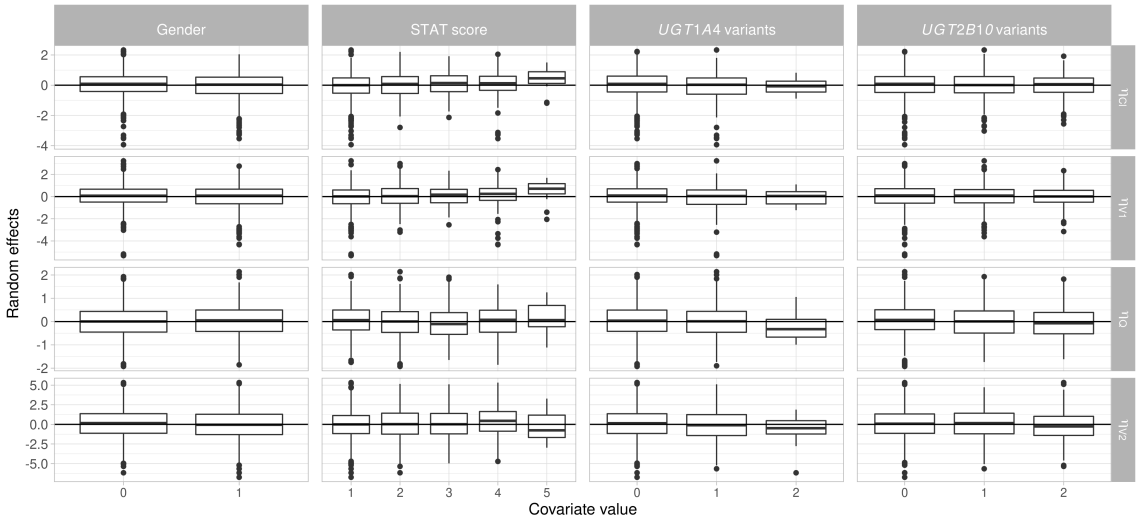

Figure 4.6: Diagnostic plots for the final population pharmacokinetic model. Random effects vs. categorical covariates.

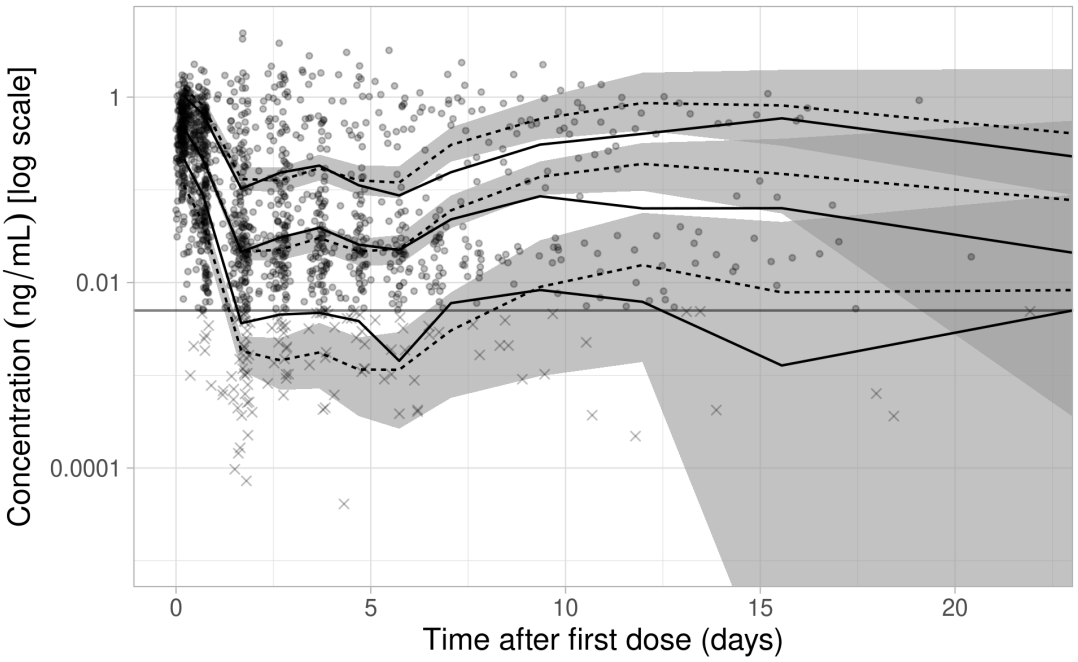

Figure 4.7: Diagnostic plots for the final population pharmacokinetic model. Prediction-corrected visual predictive check with 10th, 50th and 90th percentile of observed values (solid lines) and theoretical values (dashed lines) along with 90% prediction interval for theoretical percentiles (shaded region). Filled circles indicate observed values, x indicates simulated values based on the estimated model for observations below the lower limit of quantification of 0.005 ng/mL (represented by a horizontal grey line); time was binned using the least-squares criteria.

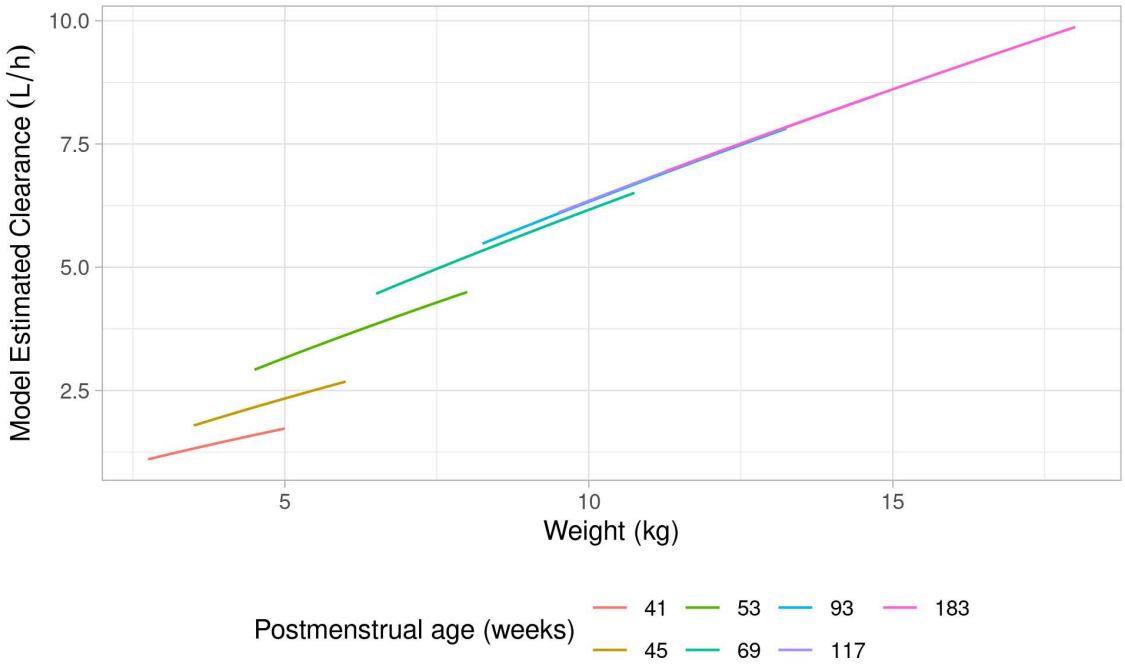

Figure 4.8: Predicted clearance by weight for selected ages from the final weight and age maturation model. Plausible weight ranges for each age group are: 41 weeks (2.7–5.1 kg), 45 weeks (3.3–6.1 kg), 53 weeks (4.5–8.0 kg), 69 weeks (6.3–10.8 kg), 93 weeks (8.1–13.4 kg), 117 weeks (9.3–14.9 kg), 183 weeks (11.2–18.1 kg). Overlapping lines between different age categories represent weights that are plausible for multiple age groups.

Model estimated clearance from the final model for seven age groups across a range of plausible weights is shown in Figure 4.8 (overlap between lines indicates weights that are plausible for multiple age groups). Weight impacts mean estimated *CL* for all ages while postmenstrual age has a large impact only for the youngest age groups. For those over 93 weeks postmenstrual age, maturation is near adult level and mean estimated *CL* is primarily determined by weight.

4.3.3 Genetic Effects on Clearance and Concentration

$UGT1A4$, $UGT2B10$, and $CYP2A6$ were not significant at the $\alpha = 0.05$ level. For the $UGT1A4$ categorical gene model the estimated effect of any variants vs. no variants was -0.221 (95% CI: -0.54 to 0.09) for a 20% decrease [exp(-0.221) $\approx$ 0.80] in *CL* on average for those with any $UGT1A4$ variants holding age and weight constant. In the $UGT1A4$ additive gene model, the estimated effect of each additional variant was -0.166 (95% CI: -0.406 to 0.073). For the $UGT2B10$ categorical model, the estimated genotype effect was -0.104 (95% CI: -0.32 to 0.11), indicating a 10% decrease on average. The $UGT2B10$ additive model estimated the effect of additional variants as -0.108 (95% CI: -0.241 to 0.025). For the $CYP2A6$ model, the estimated effect of a unit increase in risk score was 0.0885 (95% CI: -0.46 to 0.64).

Although these effects are not statistically significant using the $\alpha = 0.05$ threshold, we perform simulations to assess the hypothetical impact on total clearance and clinical dosing if the categorical $UGT1A4$ or $UGT2B10$ model estimates were utilized. Results are included in Figures 4.37 – 4.42. Including these effects in the PK model has a negligible impact on dosing.

4.4 Discussion

Using remnant specimens along with dosing, clinical, and demographic information from an EHR system we were able to develop a dexmedetomidine population PK model for a large pediatric cohort of 354 patients. We identified patient characteristics that alter the PK profile. This study is one of the largest pediatric dexmedetomidine population PK studies reported.

We confirmed a structural model and covariate relationships which are in line with those previously reported for dexmedetomidine PK. Specifically, our model included both weight and age maturation effects on *CL*. We estimated a weight-standardized *CL* of 27.3 L/h (CV 103%). Our estimated *CL* is somewhat smaller (with larger CV) than those reported in other pediatric PK studies. For a standard weight of 70 kg, Potts et al. (2009) [120] found a population *CL* estimate of 42.1 L/h (CV 30.9%); including a scaling factor of 0.73 for children given infusion (vs. bolus) reduced the *CL* estimate to 30.7 L/h. Zuppa et al. (2019) [127] estimated *CL* of 37.3 L/h (CV 48%) for neonates and infants age 0 – 6 months after cardiac bypass and Su et al. (2016) [123]

estimated CL as 39.4 L/h (CV 28%) for children age 1 – 24 months after open heart surgery. The discrepancy between studies could be related to several factors including study design and study population. For example, our study used sparse and opportunistic sampling and included a more heterogenous population which included older children and teenagers while the other studies used densely measured drug levels and were performed in a well-controlled clinical setting with a younger and more homogeneous population.

After controlling for weight and age maturation, we found little evidence to support the importance of variants in $UGT1A4$, $UGT2B10$, or predicted $CYP2A6$ activity in explaining variability of CL between subjects. There were no significant improvements to the model when genetic data were included. The largest potential effect was a 20% decrease in CL due to variants in $UGT1A4$. While this may appear to be a substantial change, modeling of the effects of this difference in clearance demonstrated this does not translate to a clinically relevant difference in dosing of dexmedetomidine. Using population PK models derived from EHR data and remnant specimens offers the possibility of more accurate prediction of individual dosing requirements in a real-life setting, especially in populations where large, intensive-sampling PK clinical trials are difficult to perform due to ethical or logistical considerations. The results from such model-informed precision dosing could also be integrated into EHR-embedded decision support tools; the development and implementation of several of these tools has been recently described by Mizuno et al. (2020) [151] and Vinks et al. (2020) [152].

There are several limitations related to the use of EHR and remnant specimens for our study. Although our data were generated using a standardized system to construct the PK data [133], there may be some errors due to inherent limitations of EHR data, which is not primarily collected for research use. First, data collected for clinical purposes may be subject to errors related to data entry or missingness. Further, real-world dosing data are not standardized with large heterogeneity in the frequency, duration, and timing of administered infusion and bolus doses. In addition, the specimens are very sparse for some subjects and their collection is not timed to facilitate optimal PK estimation. These limitations may be related to the imprecision in estimates for some PK parameters, notably $V2$. Future studies could address some of these limitations by incorporating prior information from previous or smaller pilot studies with more densely sampled data.

Despite these limitations, our study provides further evidence for the feasibility of using EHR data and remnant specimens for population PK analysis. Our study findings, such as weight effects on CL, could be helpful to develop a model-based dosing that may be superior to the current fixed weight-based dosing scheme. However, this should be tested in a future study for its clinical utility in the pediatric population. Because dexmedetomidine is used to achieve specific sedation goals, it would also be of interest to incorporate the current study results into a joint pharmacokinetic-pharmacodynamic model using sedation outcomes also derived from the EHR. These models are an important step toward the ultimate goal of precision dosing.

4.5 Appendix

4.5.1 SAEM Algorithm Settings

The SAEM algorithm was used with the following settings:

Burn-in phase (prior SAEM): 5

Exploratory phase with auto-stop (max iterations: 2000, min iterations 200, stepsize exponent: 0) and simulated annealing (decreasing rate for variance of residual errors: 0.95, decreasing rate for variance of individual parameters: 0.95)

Smoothing phase with auto-stop (max iterations: 400, min iterations 200, stepsize exponent: 0.7)

All other settings used `Monolix` defaults.

4.5.2 Supplemental Tables

Table 4.7: Tests of Deviation from Hardy Weinberg Equilibrium

rs ID	p-value
rs2942857	0.518922
rs112561475	0.762602
rs61750900	0.825786
rs2011425	0.586849
rs3892221	0.500481
rs6755571	0.534690
rs56113850	0.466614
rs2316204	0.762141
rs113288603	0.656682
rs28399442	0.122089
rs1801272	0.312122
rs28399433	0.615292

Table 4.8: Postnatal Age Categories

	Entire Cohort
n	354
Postnatal Age	
<6 mo	111 (31%)
6-24 mo	75 (21%)
2-13 yr	134 (38%)
14-18 yr	25 (7%)
>18 yr	9 (3%)

Table 4.9: Concomitant Medications Administered to at least 5% of Participants

Medication Name	Count	Percent
acetaminophen	322	92.5
cefazolin	320	92.0
famotidine	310	89.1
morphine	308	88.5
furosemide	280	80.5
fentanyl	271	77.9
rocuronium	242	69.5
oxycodone	208	59.8
heparin	203	58.3
lorazepam	179	51.4
ketorolac	162	46.6
midazolam	160	46.0
nicardipine	155	44.5
chlorothiazide	151	43.4
milrinone	129	37.1
epinephrine	119	34.2
ondansetron	116	33.3
hydromorphone	114	32.8
docusate	93	26.7
aminocaprioc	92	26.4
aspirin	85	24.4
dexamethasone	83	23.9
phenylephrine	77	22.1
vecuronium	74	21.3
vancomycin	59	17.0
aminocaproic	53	15.2
hydrocortisone	52	14.9
nitroprusside	52	14.9
spironolactone	47	13.5
enalapril	43	12.4
ketamine	43	12.4
omeprazole	43	12.4
diphenhydramine	39	11.2
cefepime	35	10.1
ephedrine	35	10.1
albuterol	33	9.5
propofol	32	9.2
dopamine	30	8.6
glycopyrrolate	30	8.6
vasopressin	27	7.8
ibuprofen	22	6.3
neostigmine	22	6.3
lidocaine	21	6.0
lisinopril	19	5.5
captopril	18	5.2

Note:
Concomitant medications unavailable for 6 participants in final study cohort.

4.5.2.1 Stage 1 Models

The goal of the first stage of modeling was to find an adequate non-genotype covariate PK model

Table 4.10 Base Models Description

1. One-compartment model with combined additive and proportional error
2. Two-compartment model with combined additive and proportional error
3. Two-compartment model with additive error
4. Two-compartment model with proportional error

Table 4.10: Base Models

Model	Structural Model	Residual variability	Submodels	Between subject variability	OFV	BICc
1	One-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon_1)+\varepsilon_2$ $\varepsilon_1 \sim N(0,\sigma^2_{prop})$ $\varepsilon_2 \sim N(0,\sigma^2_{add})$	$\psi_i = \{Cl_i, V_i\}$ $Cl_i = \theta_1 \exp(\eta_{Cl_i})$ $V_i = \theta_2 \exp(\eta_{V_i})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl})$; $\eta_{V_i} \sim N(0,\omega^2_V)$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V} \\ & \omega^2_V \end{bmatrix}$	-961.3	-914.7
2	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon_1)+\varepsilon_2$ $\varepsilon_1 \sim N(0,\sigma^2_{prop})$ $\varepsilon_2 \sim N(0,\sigma^2_{add})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 \exp(\eta_{Cl_i})$ $V_{1i} = \theta_2 \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl})$; $\eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_Q)$; $\eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & \omega_{Cl,Q} & \omega_{Cl,V2} \\ & \omega^2_{V1} & \omega_{V1,Q} & \omega_{V1,V2} \\ & & \omega^2_Q & \omega_{Q,V2} \\ & & & \omega^2_{V2} \end{bmatrix}$	-1517.3	-1415.1
3	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)+\varepsilon$ $\varepsilon \sim N(0,\sigma^2_{add})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 \exp(\eta_{Cl_i})$ $V_{1i} = \theta_2 \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl})$; $\eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_Q)$; $\eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & \omega_{Cl,Q} & \omega_{Cl,V2} \\ & \omega^2_{V1} & \omega_{V1,Q} & \omega_{V1,V2} \\ & & \omega^2_Q & \omega_{Q,V2} \\ & & & \omega^2_{V2} \end{bmatrix}$	1018.7	1113.6
4	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 \exp(\eta_{Cl_i})$ $V_{1i} = \theta_2 \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl})$; $\eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_Q)$; $\eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & \omega_{Cl,Q} & \omega_{Cl,V2} \\ & \omega^2_{V1} & \omega_{V1,Q} & \omega_{V1,V2} \\ & & \omega^2_Q & \omega_{Q,V2} \\ & & & \omega^2_{V2} \end{bmatrix}$	-1352.5	-1257.6

Table 4.11 Allometric Scaling Models Description

5. Two-compartment model with combined error and fixed theory-based allometric scaling parameters
6. Two-compartment model with combined error and estimated allometric scaling parameters

Table 4.11: Allometric Scaling Models

Model	Structural Model	Residual variability	Submodels	Between subject variability	OFV	BICc
5	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon_1)+\varepsilon_2$ $\varepsilon_1 \sim N(0,\sigma^2_{prop})$ $\varepsilon_2 \sim N(0,\sigma^2_{add})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \exp(\eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl})$; $\eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q})$; $\eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & \omega_{Cl,Q} & \omega_{Cl,V2} \\ & \omega^2_{V1} & \omega_{V1,Q} & \omega_{V1,V2} \\ & & \omega^2_{Q} & \omega_{Q,V2} \\ & & & \omega^2_{V2} \end{bmatrix}$	-1923.3	-1821.1
6	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon_1)+\varepsilon_2$ $\varepsilon_1 \sim N(0,\sigma^2_{prop})$ $\varepsilon_2 \sim N(0,\sigma^2_{add})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{\beta_1} \exp(\eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70)^{\beta_2} \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{\beta_3} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70)^{\beta_4} \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl})$; $\eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q})$; $\eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & \omega_{Cl,Q} & \omega_{Cl,V2} \\ & \omega^2_{V1} & \omega_{V1,Q} & \omega_{V1,V2} \\ & & \omega^2_{Q} & \omega_{Q,V2} \\ & & & \omega^2_{V2} \end{bmatrix}$	-1735.5	-1604.3

Table 4.12 Allometric scaling and Maturation Models Description

7. Two-compartment model with combined error and fixed theory-based allometric scaling parameters and exponential age
8. Two-compartment model with combined error and fixed theory-based allometric scaling parameters and sigmoid (Hill) maturation
9. Two-compartment model with combined error and bodyweight-dependent allometric scaling parameter for total clearance
10. Two-compartment model with combined error and age-dependent allometric scaling parameter for total clearance

Table 4.12: Allometric Scaling and Maturation Models

Model	Structural Model	Residual variability	Submodels	Between subject variability	OFV	BICc
7	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon_1)+\varepsilon_2$ $\varepsilon_1 \sim N(0,\sigma^2_{prop})$ $\varepsilon_2 \sim N(0,\sigma^2_{add})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \exp(\theta_5 AGE + \eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q}); \eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & \omega_{Cl,Q} & \omega_{Cl,V2} \\ & \omega^2_{V1} & \omega_{V1,Q} & \omega_{V1,V2} \\ & & \omega^2_{Q} & \omega_{Q,V2} \\ & & & \omega^2_{V2} \end{bmatrix}$	-1932.3	-1824.3
8	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon_1)+\varepsilon_2$ $\varepsilon_1 \sim N(0,\sigma^2_{prop})$ $\varepsilon_2 \sim N(0,\sigma^2_{add})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \exp(\eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q}); \eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & \omega_{Cl,Q} & \omega_{Cl,V2} \\ & \omega^2_{V1} & \omega_{V1,Q} & \omega_{V1,V2} \\ & & \omega^2_{Q} & \omega_{Q,V2} \\ & & & \omega^2_{V2} \end{bmatrix}$	-1945.9	-1829.3
9	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon_1)+\varepsilon_2$ $\varepsilon_1 \sim N(0,\sigma^2_{prop})$ $\varepsilon_2 \sim N(0,\sigma^2_{add})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{k_1} \exp(\eta_{Cl_i})$ $k_1 = k_0 - \frac{k_{max} WT^{Hill}}{k_{50}^{Hill} + WT^{Hill}}$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q}); \eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & \omega_{Cl,Q} & \omega_{Cl,V2} \\ & \omega^2_{V1} & \omega_{V1,Q} & \omega_{V1,V2} \\ & & \omega^2_{Q} & \omega_{Q,V2} \\ & & & \omega^2_{V2} \end{bmatrix}$	-1946.9	-1815.8
10	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon_1)+\varepsilon_2$ $\varepsilon_1 \sim N(0,\sigma^2_{prop})$ $\varepsilon_2 \sim N(0,\sigma^2_{add})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{k_1} \exp(\eta_{Cl_i})$ $k_1 = k_0 - \frac{k_{max} AGE^{Hill}}{k_{50}^{Hill} + AGE^{Hill}}$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q}); \eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & \omega_{Cl,Q} & \omega_{Cl,V2} \\ & \omega^2_{V1} & \omega_{V1,Q} & \omega_{V1,V2} \\ & & \omega^2_{Q} & \omega_{Q,V2} \\ & & & \omega^2_{V2} \end{bmatrix}$	-1939.1	-1808.0

Table 4.13 Simplified Variance Component Models Description

11. Two-compartment model with combined error and fixed theory-based allometric scaling parameters and sigmoid (Hill) maturation, no $V2$ random effects
12. Two-compartment model with combined error and fixed theory-based allometric scaling parameters and sigmoid (Hill) maturation, no $V2$ or Q random effects
13. Two-compartment model with proportional error and fixed theory-based allometric scaling parameters and sigmoid (Hill) maturation, all random effects
14. Two-compartment model with proportional error and fixed theory-based allometric scaling parameters and sigmoid (Hill) maturation, no $V2$ random effects
15. Two-compartment model with proportional error and fixed theory-based allometric scaling parameters and sigmoid (Hill) maturation, block correlation 1
16. Two-compartment model with proportional error and fixed theory-based allometric scaling parameters and sigmoid (Hill) maturation, block correlation 2

Table 4.13: Simplified Variance Component Models

Model	Structural Model	Residual variability	Submodels	Between subject variability	OFV	BICc
11	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon_1)+\varepsilon_2$ $\varepsilon_1 \sim N(0,\sigma^2_{prop})$ $\varepsilon_2 \sim N(0,\sigma^2_{add})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \exp(\eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70)$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & \omega_{Cl,Q} \\ & \omega^2_{V1} & \omega_{V1,Q} \\ & & \omega^2_{Q} \end{bmatrix}$	-1504.7	-1411.6
12	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon_1)+\varepsilon_2$ $\varepsilon_1 \sim N(0,\sigma^2_{prop})$ $\varepsilon_2 \sim N(0,\sigma^2_{add})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \exp(\eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75}$ $V_{2i} = \theta_4 (WT/70)$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} \\ & \omega^2_{V1} \end{bmatrix}$	-1408.4	-1332.9

Table 4.13: Simplified Variance Component Models *(continued)*

Model	Structural Model	Residual variability	Submodels	Between subject variability	OFV	BICc
13	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \exp(\eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q}); \eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & \omega_{Cl,Q} & \omega_{Cl,V2} \\ & \omega^2_{V1} & \omega_{V1,Q} & \omega_{V1,V2} \\ & & \omega^2_{Q} & \omega_{Q,V2} \\ & & & \omega^2_{V2} \end{bmatrix}$	-1932.9	-1823.5
14	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \exp(\eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70)$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & \omega_{Cl,Q} \\ & \omega^2_{V1} & \omega_{V1,Q} \\ & & \omega^2_{Q} \end{bmatrix}$	-1479.6	-1393.6
15	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \exp(\eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q}); \eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & \omega_{Q,V2} \\ & & & \omega^2_{V2} \end{bmatrix}$	-1915.2	-1829.3
16	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \exp(\eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q}); \eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & 0 \\ & & & \omega^2_{V2} \end{bmatrix}$	-1915.1	-1835.0

Table 4.14 Additional Simplified Variance Components Models Description

17. Two-compartment model with proportional error and estimated allometric scaling parameters, block correlation 2
18. Two-compartment model with proportional error and estimated allometric scaling parameters and exponential age, block correlation 2
19. Two-compartment model with proportional error and fixed theory-based allometric scaling parameters and exponential age, block correlation 2
20. Two-compartment model with proportional error and estimate allometric scaling parameters and sigmoid (Hill) maturation, block correlation 2

Table 4.14: Additional Simplified Variance Components Models

Model	Structural Model	Residual variability	Submodels	Between subject variability	OFV	BICc
17	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{\beta_1} \exp(\eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70)^{\beta_2} \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{\beta_3} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70)^{\beta_4} \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl})$; $\eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q})$; $\eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & 0 \\ & & & \omega^2_{V2} \end{bmatrix}$	-1910.1	-1815.5
18	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{\beta_1} \exp(\theta_5 AGE + \eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70)^{\beta_2} \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{\beta_3} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70)^{\beta_4} \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl})$; $\eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q})$; $\eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & 0 \\ & & & \omega^2_{V2} \end{bmatrix}$	-1904.0	-1803.6
19	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \exp(\theta_5 AGE + \eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl})$; $\eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q})$; $\eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & 0 \\ & & & \omega^2_{V2} \end{bmatrix}$	-1899.1	-1827.6
20	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{\beta_1} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}} \right) \exp(\eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70)^{\beta_2} \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{\beta_3} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70)^{\beta_4} \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl})$; $\eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q})$; $\eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & 0 \\ & & & \omega^2_{V2} \end{bmatrix}$	-1927.4	-1818.4

Table 4.15 Additional Non-genotype Covariate Models Description

21. Add exponential gender effect on CL to Model 16
22. Add exponential serum creatinine effect on CL to Model 16
23. Add exponential cardiac bypass time effect on CL to Model 16
24. Add exponential STAT score effect on CL to Model 16
25. Add exponential ICU hospitalization length effect on CL to Model 16
26. Add exponential gender effect on $V1$ to Model 16
27. Add exponential serum creatinine effect on $V1$ to Model 16
28. Add exponential cardiac bypass time effect on $V1$ to Model 16
29. Add exponential STAT score effect on $V1$ to Model 16
30. Add exponential ICU hospitalization length effect on $V1$ to Model 16
31. Add exponential gender effect on $V2$ to Model 16
32. Add exponential serum creatinine effect on $V2$ to Model 16
33. Add exponential cardiac bypass time effect on $V2$ to Model 16
34. Add exponential STAT score effect on $V2$ to Model 16
35. Add exponential ICU hospitalization length effect on $V2$ to Model 16

Table 4.15: Additional non-Genotype Covariate Models

Model	Structural Model	Residual variability	Submodels	Between subject variability	OFV	BICc
21	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \exp(\theta_{Cl,gender} I[female] + \eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q}); \eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & 0 \\ & & & \omega^2_{V2} \end{bmatrix}$	-1914.1	-1828.2
22	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \exp(\theta_{Cl,creat} creat + \eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q}); \eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & 0 \\ & & & \omega^2_{V2} \end{bmatrix}$	-1913.1	-1827.2

Table 4.15: Additional non-Genotype Covariate Models *(continued)*

Model	Structural Model	Residual variability	Submodels	Between subject variability	OFV	BICc
23	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \exp(\theta_{Cl,cpb}(bypasstime) + \eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q}); \eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & 0 \\ & & & \omega^2_{V2} \end{bmatrix}$	-1913.4	-1827.5
24	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \exp(\theta_{Cl,STAT}(STATscore) + \eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q}); \eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & 0 \\ & & & \omega^2_{V2} \end{bmatrix}$	-1914.9	-1829.0
25	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \exp(\theta_{Cl,loi}(ICUtime) + \eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q}); \eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & 0 \\ & & & \omega^2_{V2} \end{bmatrix}$	-1913.9	-1827.9
26	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \exp(\eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\theta_{V1,gender} I[female] + \eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q}); \eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & 0 \\ & & & \omega^2_{V2} \end{bmatrix}$	-1914.8	-1828.8
27	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \exp(\eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\theta_{V1,creat} creat + \eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q}); \eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & 0 \\ & & & \omega^2_{V2} \end{bmatrix}$	-1921.7	-1835.8

Table 4.15: Additional non-Genotype Covariate Models *(continued)*

Model	Structural Model	Residual variability	Submodels	Between subject variability	OFV	BICc
28	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \exp(\eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\theta_{V1,cpb}(bypasstime) + \eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q}); \eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & 0 \\ & & & \omega^2_{V2} \end{bmatrix}$	-1901.9	-1816.0
29	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \exp(\eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\theta_{V1,STAT}(STATscore) + \eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q}); \eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & 0 \\ & & & \omega^2_{V2} \end{bmatrix}$	-1919.4	-1833.5
30	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \exp(\eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\theta_{V1,loi}(ICUtime) + \eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q}); \eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & 0 \\ & & & \omega^2_{V2} \end{bmatrix}$	-1915.3	-1829.4
31	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \exp(\eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\theta_{V2,gender} I[female] + \eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q}); \eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & 0 \\ & & & \omega^2_{V2} \end{bmatrix}$	-1914.8	-1828.8
32	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \exp(\eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\theta_{V2,creat} creat + \eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q}); \eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & 0 \\ & & & \omega^2_{V2} \end{bmatrix}$	-1915.4	-1829.5

Table 4.15: Additional non-Genotype Covariate Models *(continued)*

Model	Structural Model	Residual variability	Submodels	Between subject variability	OFV	BICc
33	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}} \right) \exp(\eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\theta_{V2,cpb}(bypasstime) + \eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl})$; $\eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q})$; $\eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & 0 \\ & & & \omega^2_{V2} \end{bmatrix}$	-1920.7	-1834.7
34	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}} \right) \exp(\eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\theta_{V2,STAT}(STAT\,score) + \eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl})$; $\eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q})$; $\eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & 0 \\ & & & \omega^2_{V2} \end{bmatrix}$	-1915.5	-1829.6
35	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}} \right) \exp(\eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\theta_{V2,loi}(ICUtime) + \eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl})$; $\eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q})$; $\eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & 0 \\ & & & \omega^2_{V2} \end{bmatrix}$	-1920.2	-1834.3

4.5.2.2 Stage 2 Models

The goal of the second stage of modeling was to test for improvement in the best covariate model by adding genotype effects

Table 4.16 Best Covariate Model and UGT* Genotype Effects Models Description

36. Add exponential $UGT2B10$ categorical effect (no variants vs. any variants) to Model 16
37. Add exponential $UGT2B10$ additive effect to Model 16
38. Add exponential $UGT1A4$ categorical effect (no variants vs. any variants) to Model 16
39. Add exponential $UGT1A4$ additive effect to Model 16

Table 4.16: Best Covariate Model and UGT* Genotype Effects

Model	Structural Model	Residual variability	Submodels	Between subject variability	OFV	BICc
36	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \exp(\theta_{Cl,UGT2B10} I[UGT2B10>0] + \eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q}); \eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & 0 \\ & & & \omega^2_{V2} \end{bmatrix}$	-1913.8	-1827.9
37	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \exp(\theta_{Cl,UGT2B10} UGT2B10 + \eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q}); \eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & 0 \\ & & & \omega^2_{V2} \end{bmatrix}$	-1917.3	-1831.4
38	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \exp(\theta_{Cl,UGT1A4} I[UGT1A4>0] + \eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q}); \eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & 0 \\ & & & \omega^2_{V2} \end{bmatrix}$	-1918.3	-1832.4
39	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \exp(\theta_{Cl,UGT1A4} UGT1A4 + \eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q}); \eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & 0 \\ & & & \omega^2_{V2} \end{bmatrix}$	-1917.6	-1831.7

Table 4.17 Best Covariate Model and *CYP2A6* PRS Score in Subset with Complete PRS Data Models Description

40. Two-compartment model with proportional error and fixed theory-based allometric scaling parameters and sigmoid (Hill) maturation, block correlation 2

41. Add exponential *CYP2A6* additive effect to Model 30

Table 4.17: Best Covariate Model and *CYP2A6* PRS Score in subset with Complete PRS data (n=350)

Model	Structural Model	Residual variability	Submodels	Between subject variability	OFV	BICc
40	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \exp(\eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q}); \eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & 0 \\ & & & \omega^2_{V2} \end{bmatrix}$	-1889.2	-1809.3
41	Two-compartment multidose IV infusion/bolus with linear elimination	$y = C(\psi;x)(1+\varepsilon)$ $\varepsilon \sim N(0,\sigma^2_{prop})$	$\psi_i = \{Cl_i, V_{1i}, Q_i, V_{2i}\}$ $Cl_i = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}}\right) \exp(\theta_{Cl,CYP2A6}(CYP2A6score) + \eta_{Cl_i})$ $V_{1i} = \theta_2 (WT/70) \exp(\eta_{V_{1i}})$ $Q_i = \theta_3 (WT/70)^{0.75} \exp(\eta_{Q_i})$ $V_{2i} = \theta_4 (WT/70) \exp(\eta_{V_{2i}})$ $\eta_{Cl_i} \sim N(0,\omega^2_{Cl}); \eta_{V1_i} \sim N(0,\omega^2_{V1})$ $\eta_{Q_i} \sim N(0,\omega^2_{Q}); \eta_{V2_i} \sim N(0,\omega^2_{V2})$	$\Omega = \begin{bmatrix} \omega^2_{Cl} & \omega_{Cl,V1} & 0 & 0 \\ & \omega^2_{V1} & 0 & 0 \\ & & \omega^2_{Q} & 0 \\ & & & \omega^2_{V2} \end{bmatrix}$	-1889.9	-1804.2

Abbreviations for Tables 4.10 – 4.17: IV, intravenous; C, concentration; σ_{prop} and σ_{add} are proportional and additive residual error terms; OFV, objective function value; BICc, corrected Bayesian information criteria; *CL*, total clearance (L/hr); V, volume of distribution for the central compartment (L) in one-compartment model; *Q*, intercompartmental clearance (L/hr); $V1$, volume of distribution for the central compartment (L) in two-compartment model; $V2$, volume of distribution for the peripheral compartment (L); ω_{Cl}, ω_{V1}, ω_{Q}, ω_{V2}, the standard deviation for η_i^{Cl}, η_i^{V1}, η_i^{Q}, η_i^{V2}, respectively; Ω is intra-individual variance-covariance matrix; TM_{50} postmenstrual age at which clearance is 50% of adult value; Hill, maturation factor slope coefficient; CV, coefficient of variation; WT, body weight in kg; PMA, postmenstrual age in weeks

Table 4.18: Estimates of Parameters for Population Pharmacokinetic Models

Base Model (Obj = -1517.3, BICc = -1415.1)		**Weight Only Model** (Obj = -1923.3, BICc = -1821.1)		**Weight and Maturation with simplified variance structure** (Obj = -1915.1, BICc = -1835.0)	
Parameters	Estimates (SE) [95% CI][a]	Parameters	Estimates (SE) [95% CI]	Parameters	Estimates (SE) [95% CI]
Cl		$Cl = \theta_1 (WT/70)^{0.75}$		$Cl = \theta_1 (WT/70)^{0.75} \left(\frac{1}{1+\left(\frac{TM_{50}}{PMA}\right)^{Hill}} \right)$	
	6.27 (0.52) [5.33, 7.37]	θ_1	22.3 (1.85) [19.0, 26.2]	θ_1	27.3 (1.82) [24.0, 31.1]
				TM_{50}	41.9 (0.28) [41.4, 42.5]
				$Hill$	7.04 (0.022) [6.99, 7.08]
V_1		$V_1 = \theta_2 (WT/70)$		$V_1 = \theta_2 (WT/70)$	
	19.5 (1.37) [17.0, 22.4]	θ_2	123 (11.5) [102, 148]	θ_2	161 (12.1) [139, 187]
Q		$Q = \theta_3 (WT/70)^{0.75}$		$Q = \theta_3 (WT/70)^{0.75}$	
	5.26 (0.44) [4.47, 6.20]	θ_3	26.6 (2.23) [22.6, 31.3]	θ_3	26.0 (1.90) [22.5, 30.0]
V_2		$V_2 = \theta_4 (WT/70)$		$V_2 = \theta_4 (WT/70)$	
	763 (126) [555, 1048]	θ_4	6674 (1499) [4363, 10211]	θ_4	7903 (1408) [5617, 11119]
ω_{Cl} (%CV)	201 (12) [178, 226]	ω_{Cl} (%CV)	123 (13) [100, 150]	ω_{Cl} (%CV)	103 (8) [88, 120]
ω_{V1} (%CV)	161 (9) [145, 179]	ω_{V1} (%CV)	168 (22) [130, 217]	ω_{V1} (%CV)	138 (13) [114, 166]
ω_{Q} (%CV)	146 (8) [132, 162]	ω_{Q} (%CV)	91 (11) [71, 115]	ω_{Q} (%CV)	82 (9) [65, 102]
ω_{V2} (%CV)	672 (81) [534, 855]	ω_{V2} (%CV)	857 (256) [494, 1595]	ω_{V2} (%CV)	624 (157) [391, 1048]
$\rho_{Cl,V1}$	0.964 (0.003) [0.959, 0.97]	$\rho_{Cl,V1}$	0.849 (0.055) [0.741, 0.956]	$\rho_{Cl,V1}$	0.923 (0.027) [0.871, 0.975]
$\rho_{Cl,V2}$	−0.032 (0.043) [−0.116, 0.053]	$\rho_{Cl,V2}$	−0.174 (0.101) [−0.372, 0.023]	$\rho_{Cl,V2}$	0 (fixed)

Table 4.18: Estimates of Parameters for Population Pharmacokinetic Models *(continued)*

Parameters	Estimates (SE) [95% CI][a]	Parameters	Estimates (SE) [95% CI]	Parameters	Estimates (SE) [95% CI]
$\rho_{Cl,Q}$	0.043 (0.076) [−0.106, 0.192]	$\rho_{Cl,Q}$	−0.509 (0.110) [−0.726, −0.293]	$\rho_{Cl,Q}$	0 (fixed)
$\rho_{V1,V2}$	0.195 (0.044) [0.109, 0.28]	$\rho_{V1,V2}$	0.277 (0.107) [0.068, 0.486]	$\rho_{V1,V2}$	0 (fixed)
$\rho_{V1,Q}$	−0.076 (0.073) [−0.22, 0.067]	$\rho_{V1,Q}$	−0.410 (0.125) [−0.656, −0.165]	$\rho_{V1,Q}$	0 (fixed)
$\rho_{V2,Q}$	0.089 (0.083) [−0.074, 0.252]	$\rho_{V2,Q}$	0.21 (0.121) [−0.028, 0.448]	$\rho_{V2,Q}$	0 (fixed)
$\sigma_{add}\,(ng/mL)$	$2.22e-16$ ($5.65e-13$) [0, $1.11e-12$]	$\sigma_{add}\,(ng/mL)$	$1.9e-08$ ($3.73e-07$) [0, $7.51e-07$]	$\sigma_{add}\,(ng/mL)$	0 (fixed)
$\sigma_{prop}\,(\%CV)$	50.3 (1.4) [47.6, 53.0]	$\sigma_{prop}\,(\%CV)$	50.3 (1.6) [47.2, 53.4]	$\sigma_{prop}\,(\%CV)$	50.5 (1.6) [47.5, 53.5]

Abbreviations for Table 4.18: [a] 95% Asymptotic confidence intervals (CIs); SE, standard error; Obj, objective function value; BICc, corrected Bayesian information criteria; CL, total clearance (L/hr); Q, intercompartmental clearance (L/hr); V_1, volume of distribution for the central compartment (L); V_2, volume of distribution for the peripheral compartment (L); TM_{50} postmenstrual age at which clearance is 50% of adult value; $Hill$, maturation factor slope coefficient; CV, coefficient of variation; WT, body weight in kg; PMA, postmenstrual age in weeks; ω_{CL}, ω_{V1}, ω_Q, ω_{V2}, the standard deviation for η_i^{CL}, η_i^{V1}, η_i^Q, η_i^{V2}, respectively; For the standard deviation of random effects, ω, coefficient of variation was calculated as CV% = $100 \times \sqrt{\exp(\omega^2) - 1}$; ρ are correlation terms between random effects; σ_{prop} and σ_{add} are proportional and additive residual error terms.

4.5.3 Supplemental Figures

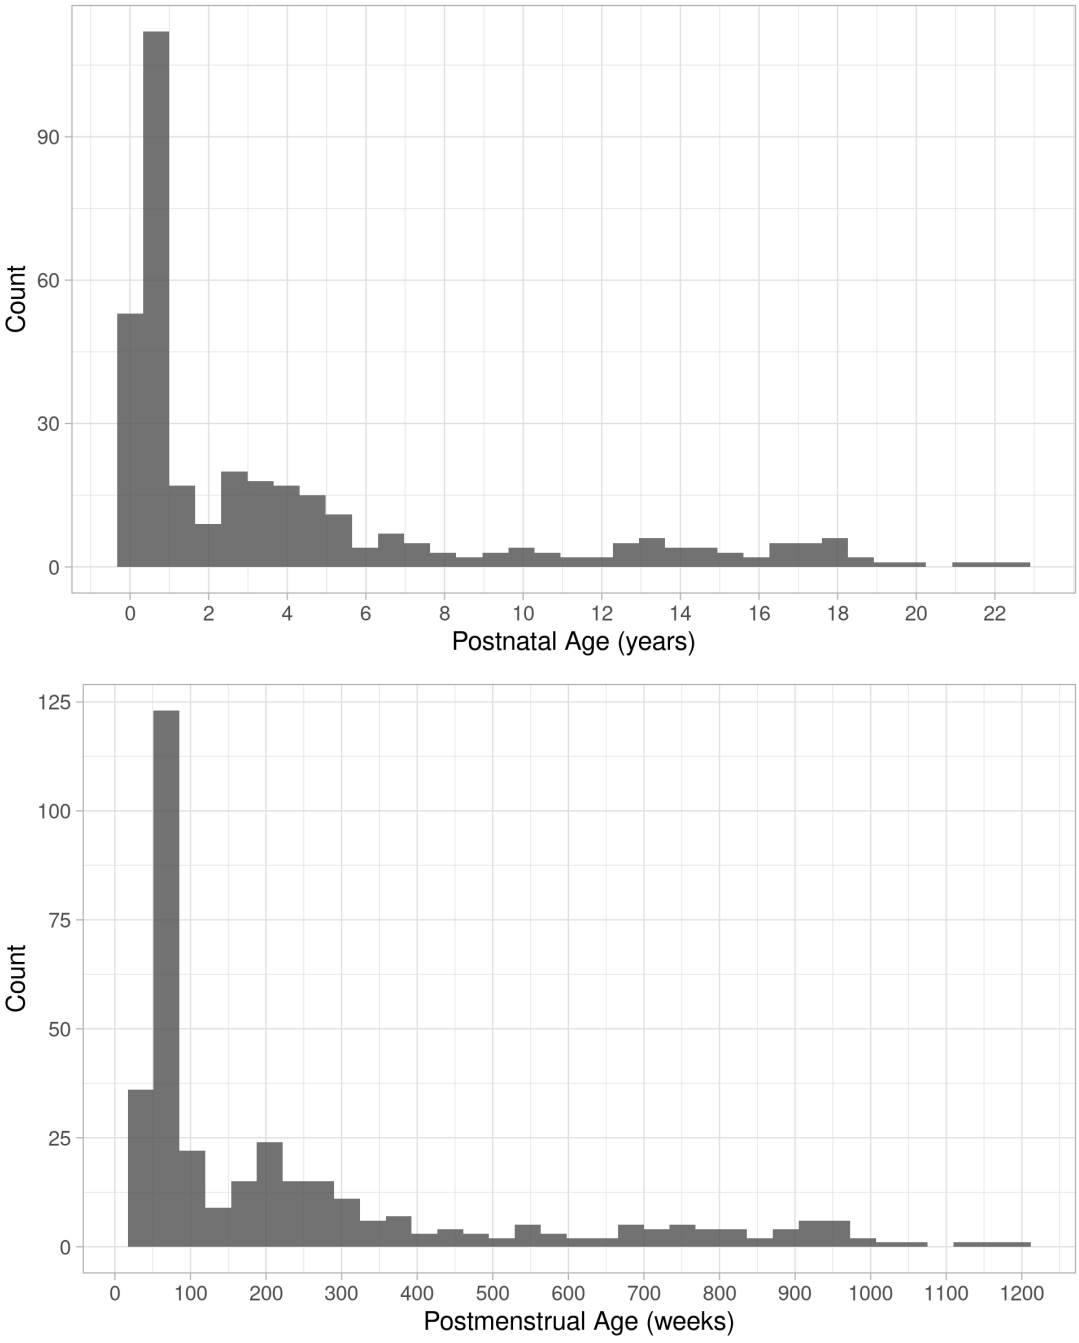

Figure 4.9: Distribution of (A) postnatal age and (B) postmenstrual age in final study population.

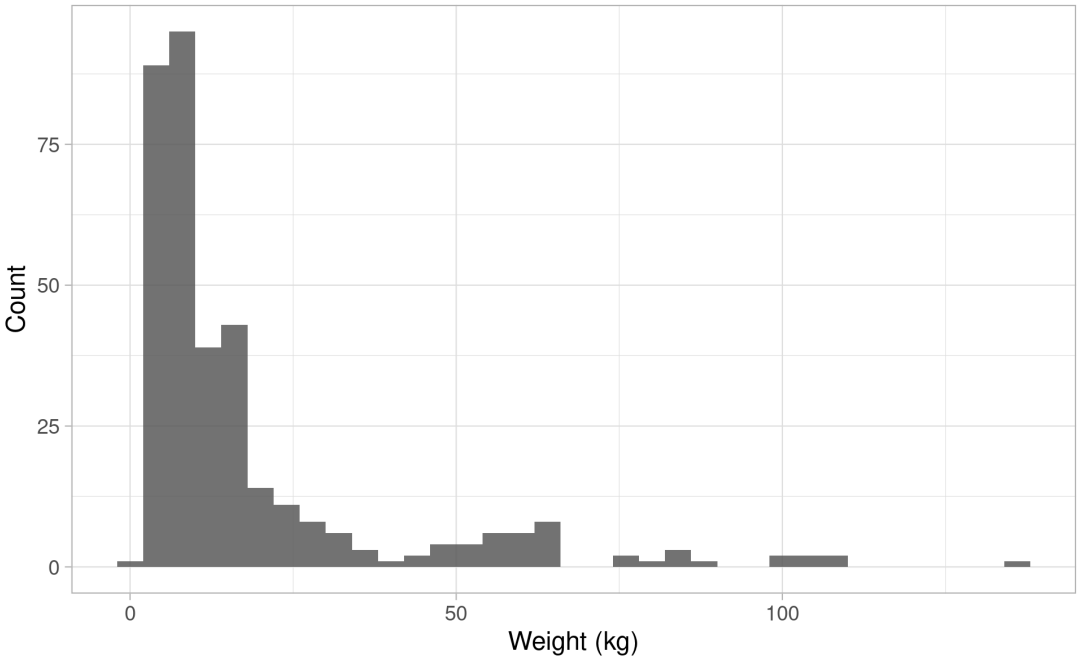

Figure 4.10: Distribution of weight in final study population.

4.5.3.1 Model 2 Goodness-of-fit Plots

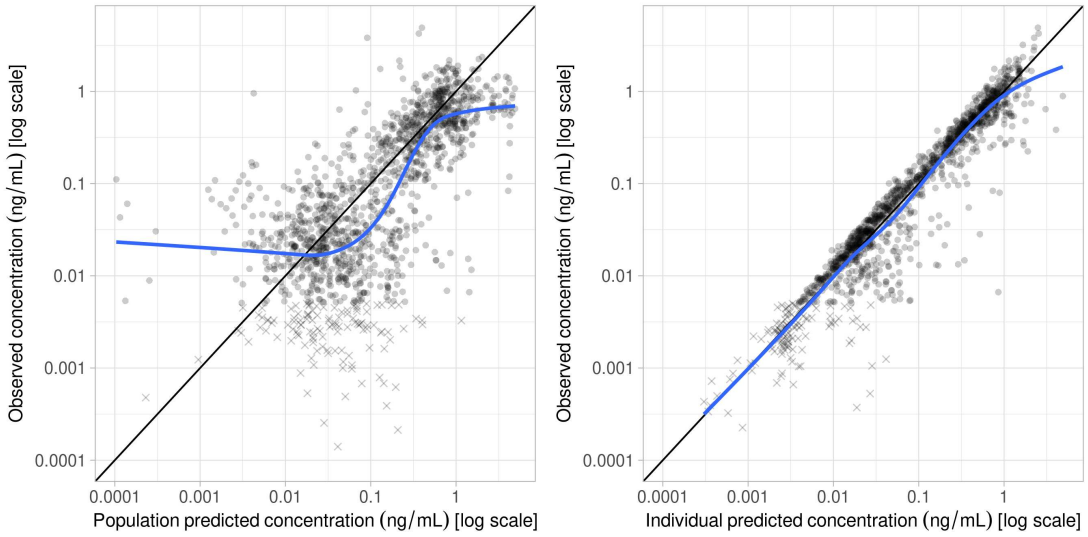

Figure 4.11: (A) Observed vs. population predicted concentrations and (B) observed vs. individual predicted concentrations.

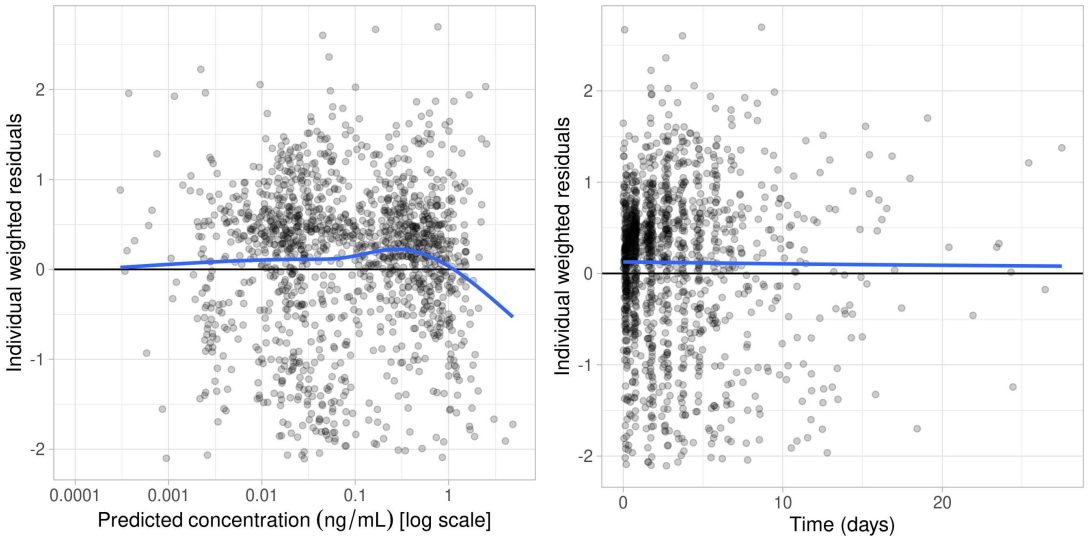

Figure 4.12: (A) Individual weighted residuals vs. predicted concentration and (B) individual weighted residuals vs. time.

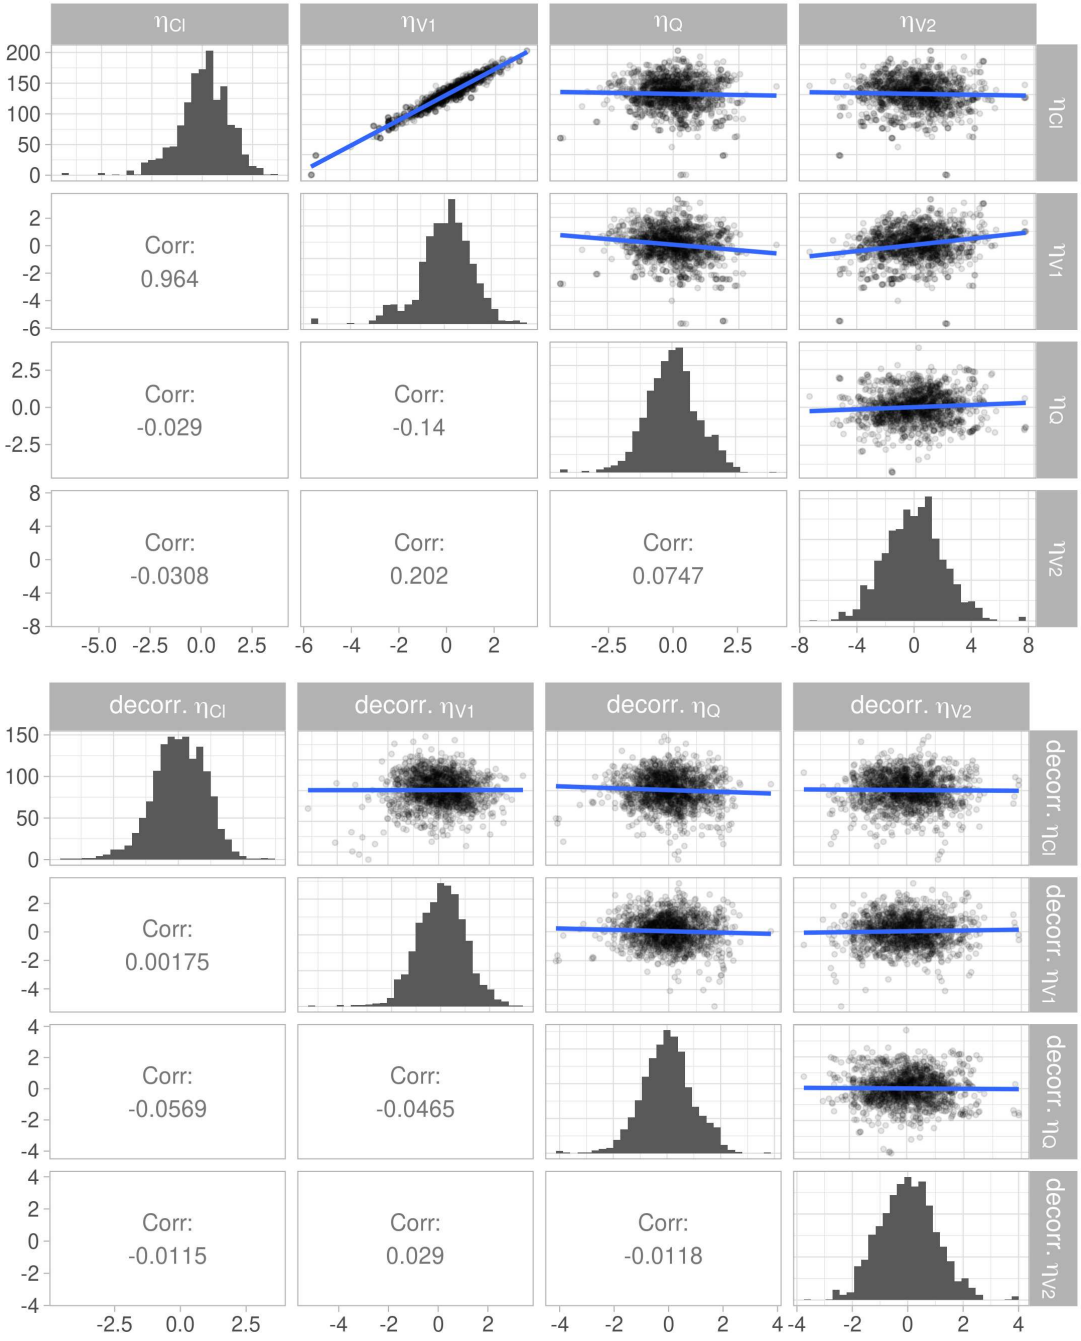

Figure 4.13: (A) Random effects correlations and (B) decorrelated random effects correlations.

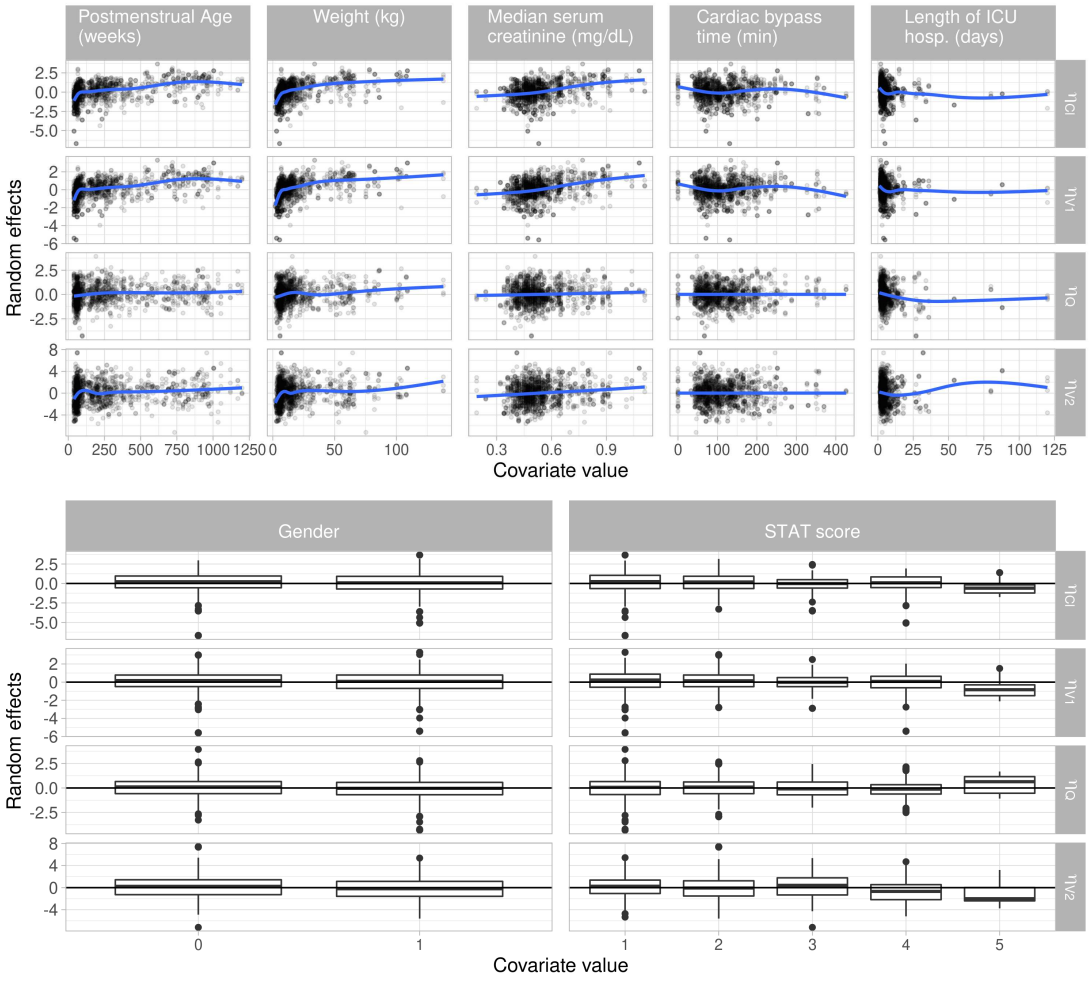

Figure 4.14: (A) Random effects vs. continuous covariates and (B) random effects vs. categorical covariates.

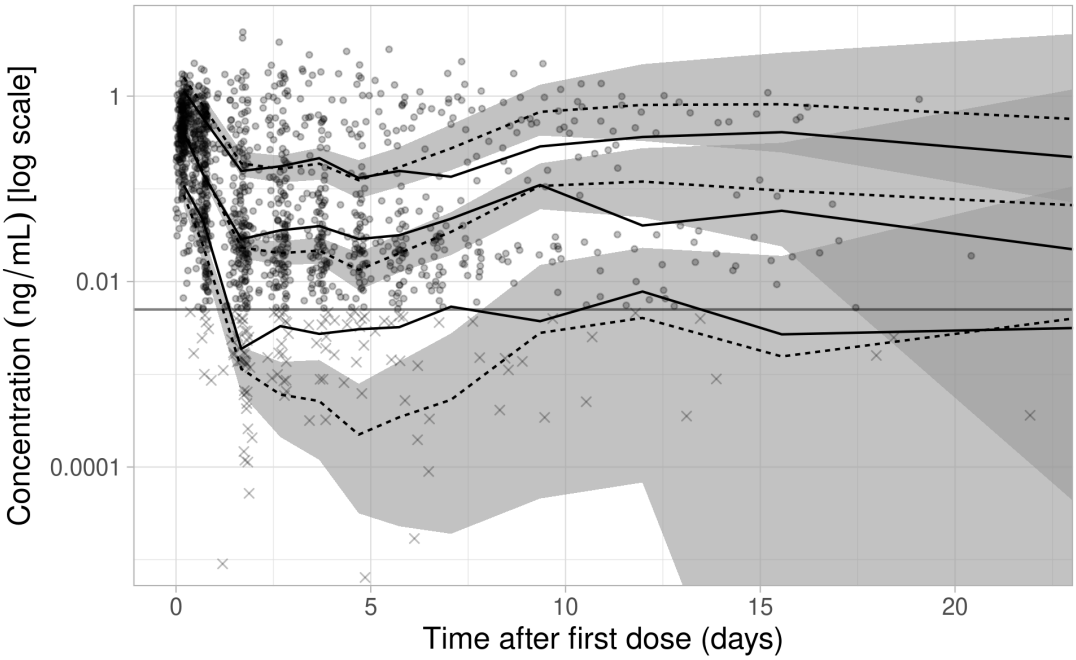

Figure 4.15: Prediction corrected visual predictive check.

4.5.3.2 Model 5 Goodness-of-fit Plots

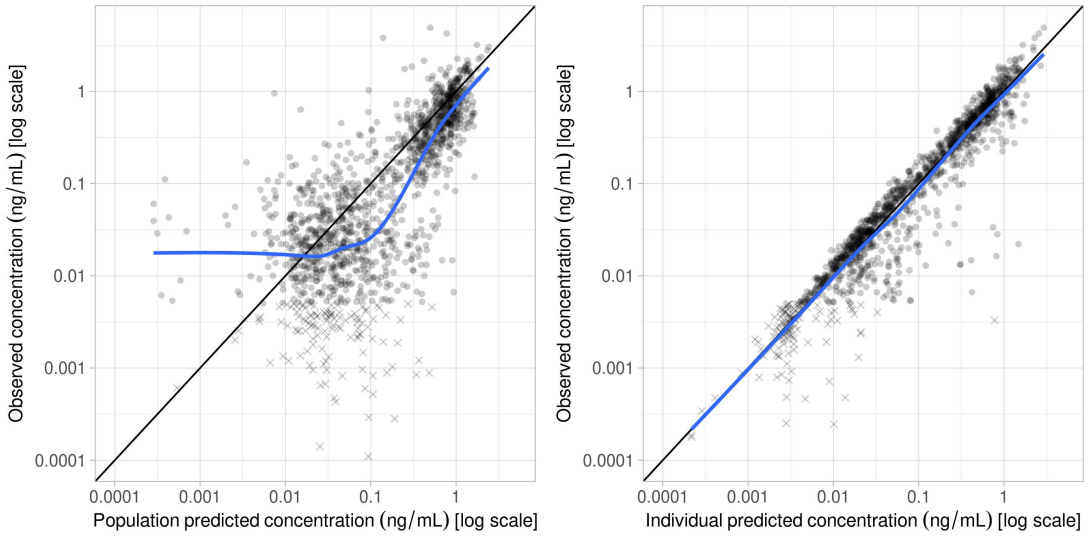

Figure 4.16: (A) Observed vs. population predicted concentrations and (B) observed vs. individual predicted concentrations.

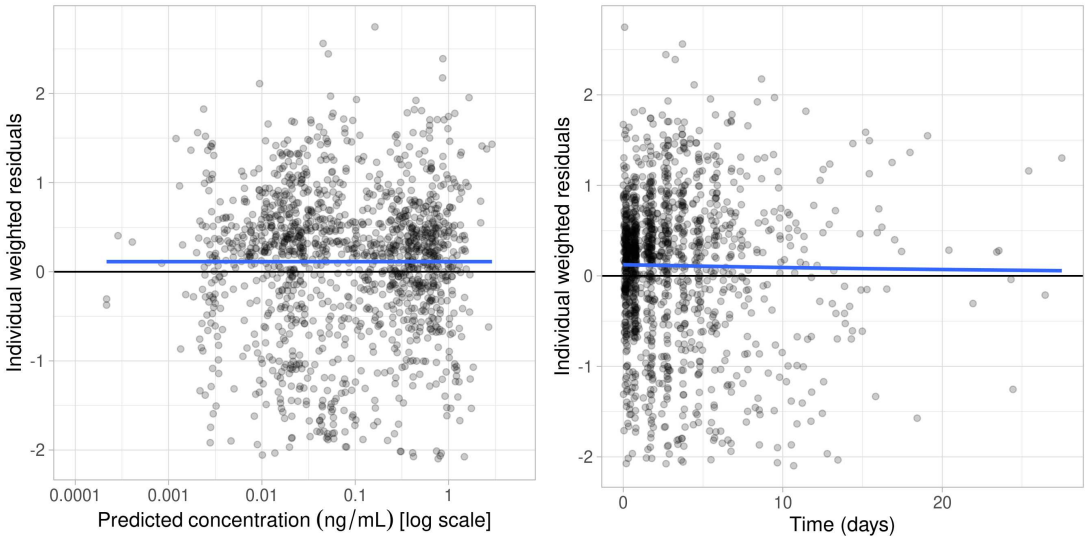

Figure 4.17: (A) Individual weighted residuals vs. predicted concentration and (B) individual weighted residuals vs. time.

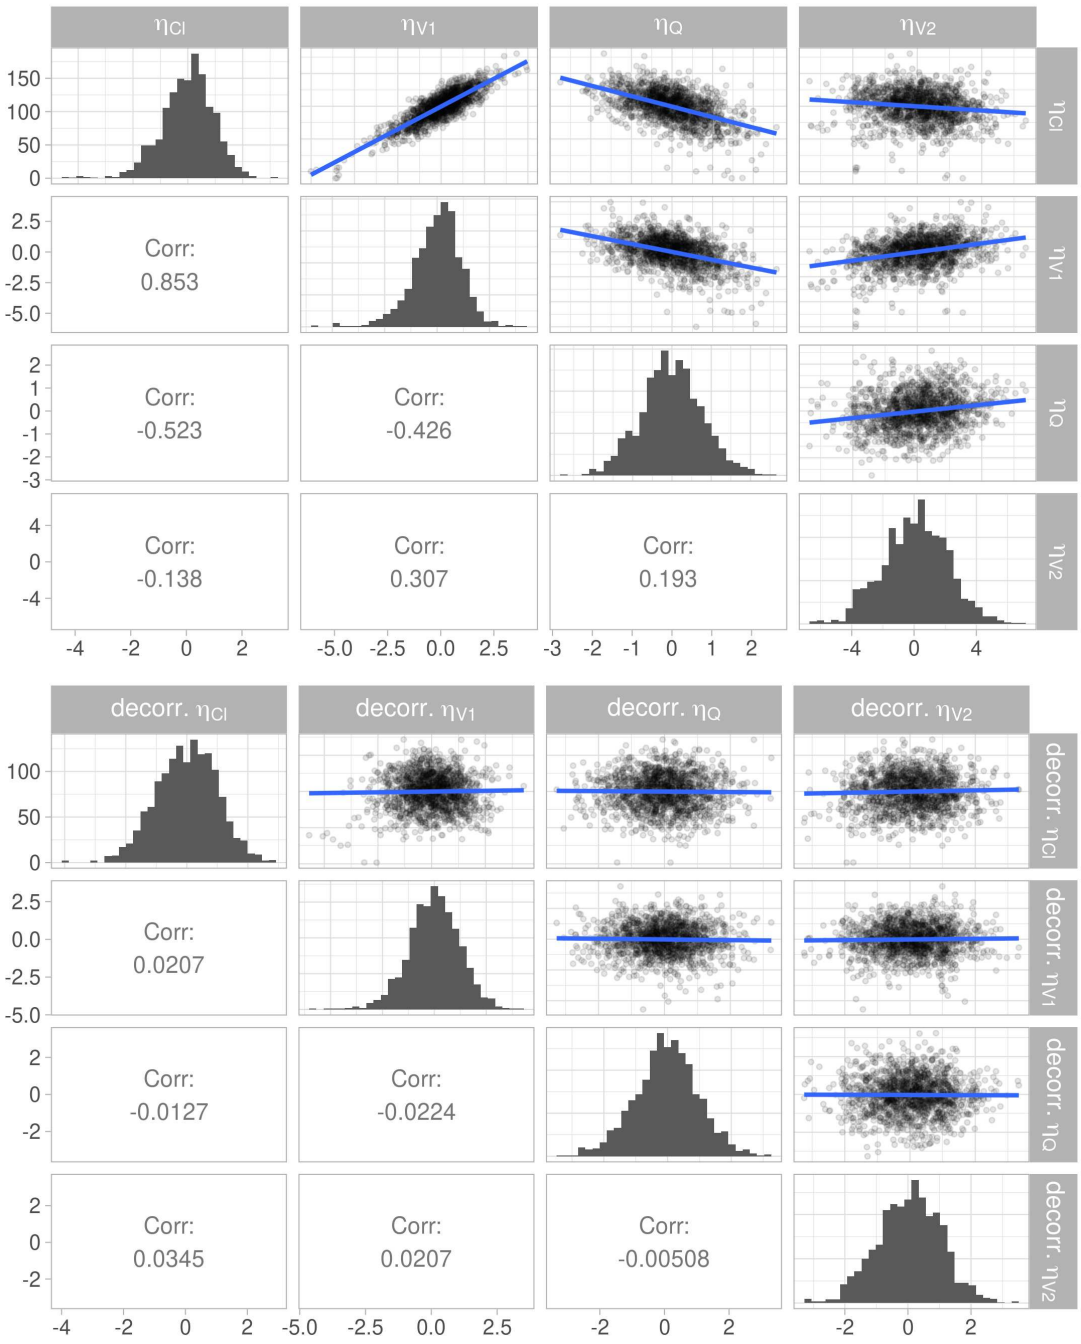

Figure 4.18: (A) Random effects correlations and (B) decorrelated random effects correlations.

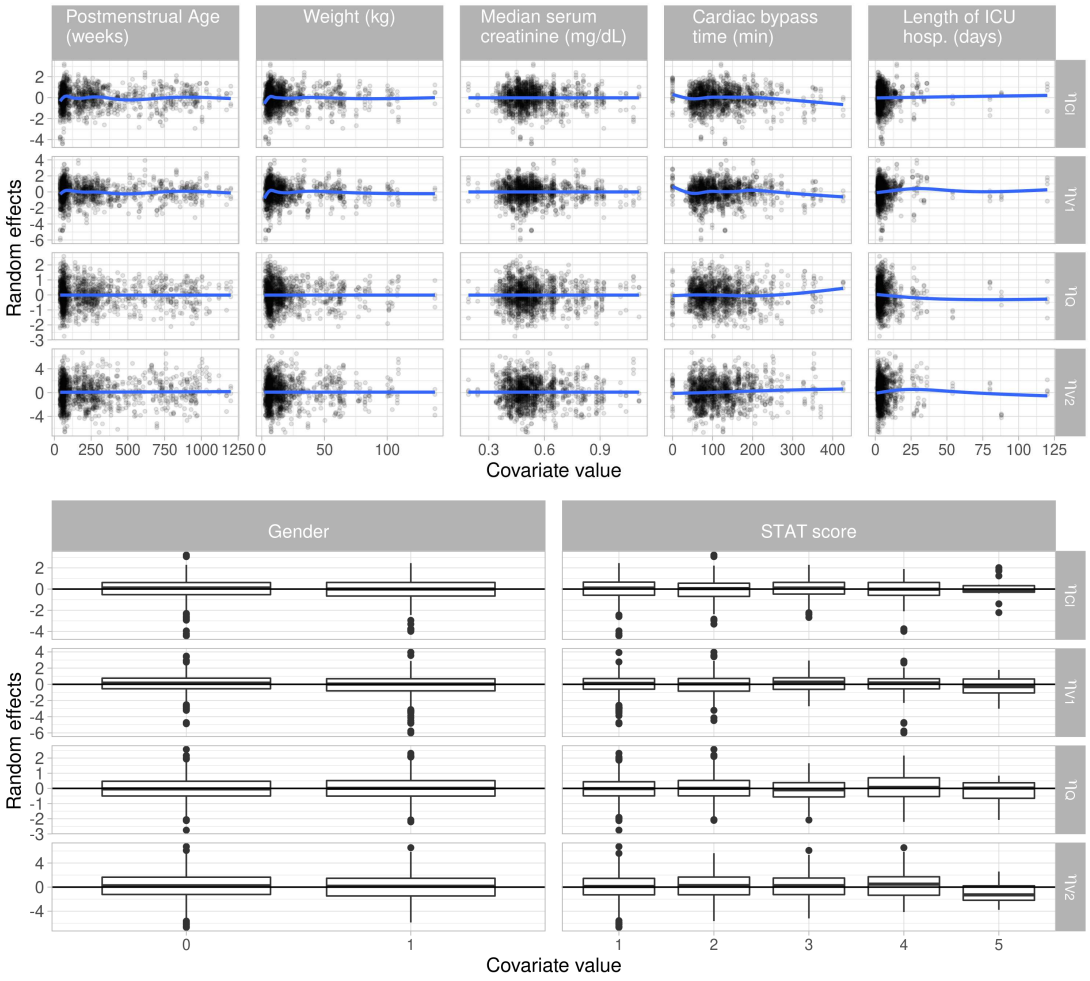

Figure 4.19: (A) Random effects vs. continuous covariates and (B) random effects vs. categorical covariates.

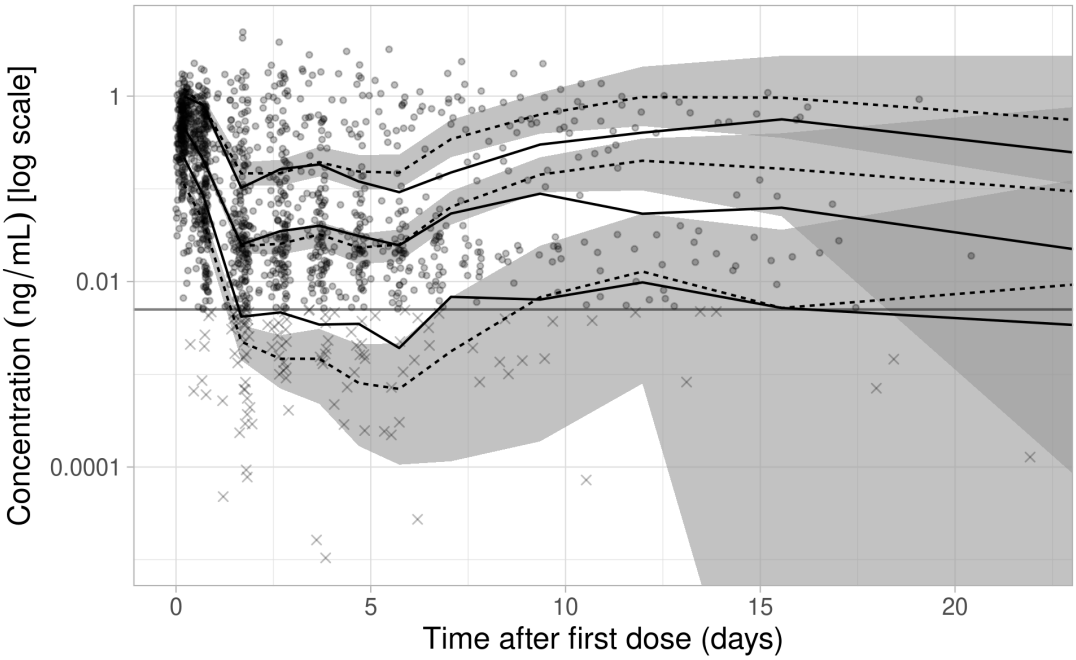

Figure 4.20: Prediction corrected visual predictive check.

4.5.3.3 Model 8 Goodness-of-fit Plots

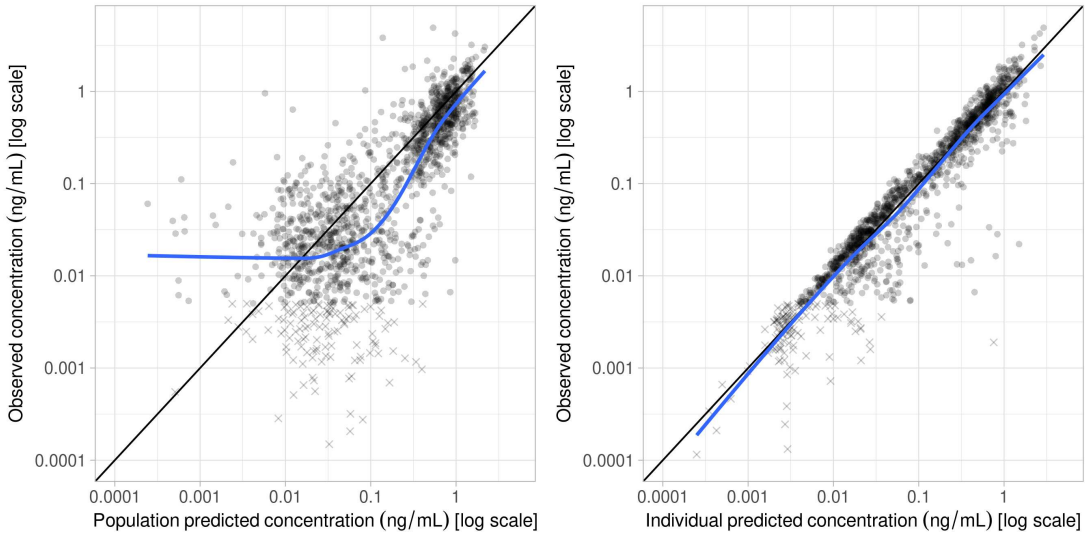

Figure 4.21: (A) Observed vs. population predicted concentrations and (B) observed vs. individual predicted concentrations.

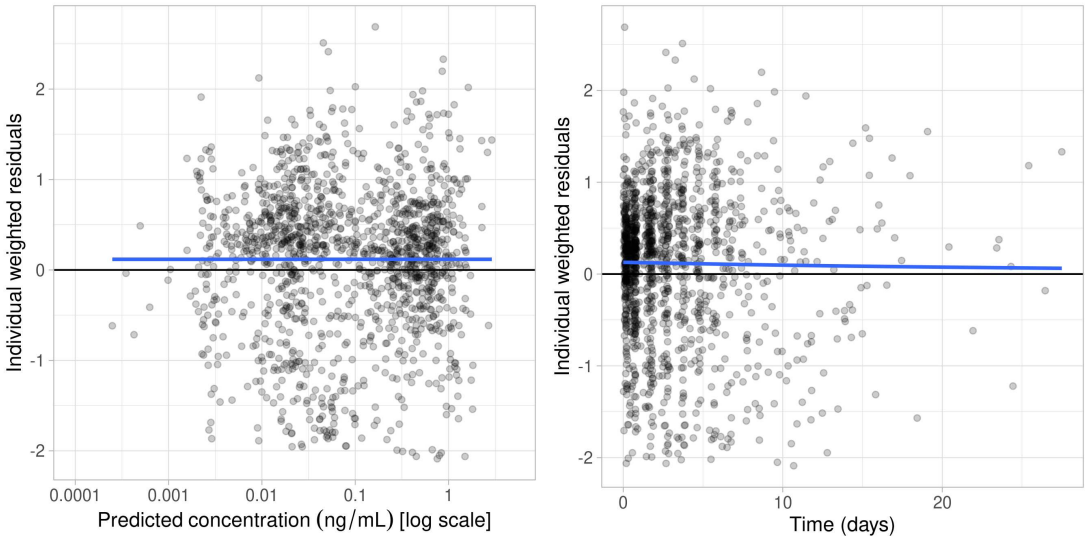

Figure 4.22: (A) Individual weighted residuals vs. predicted concentration and (B) individual weighted residuals vs. time.

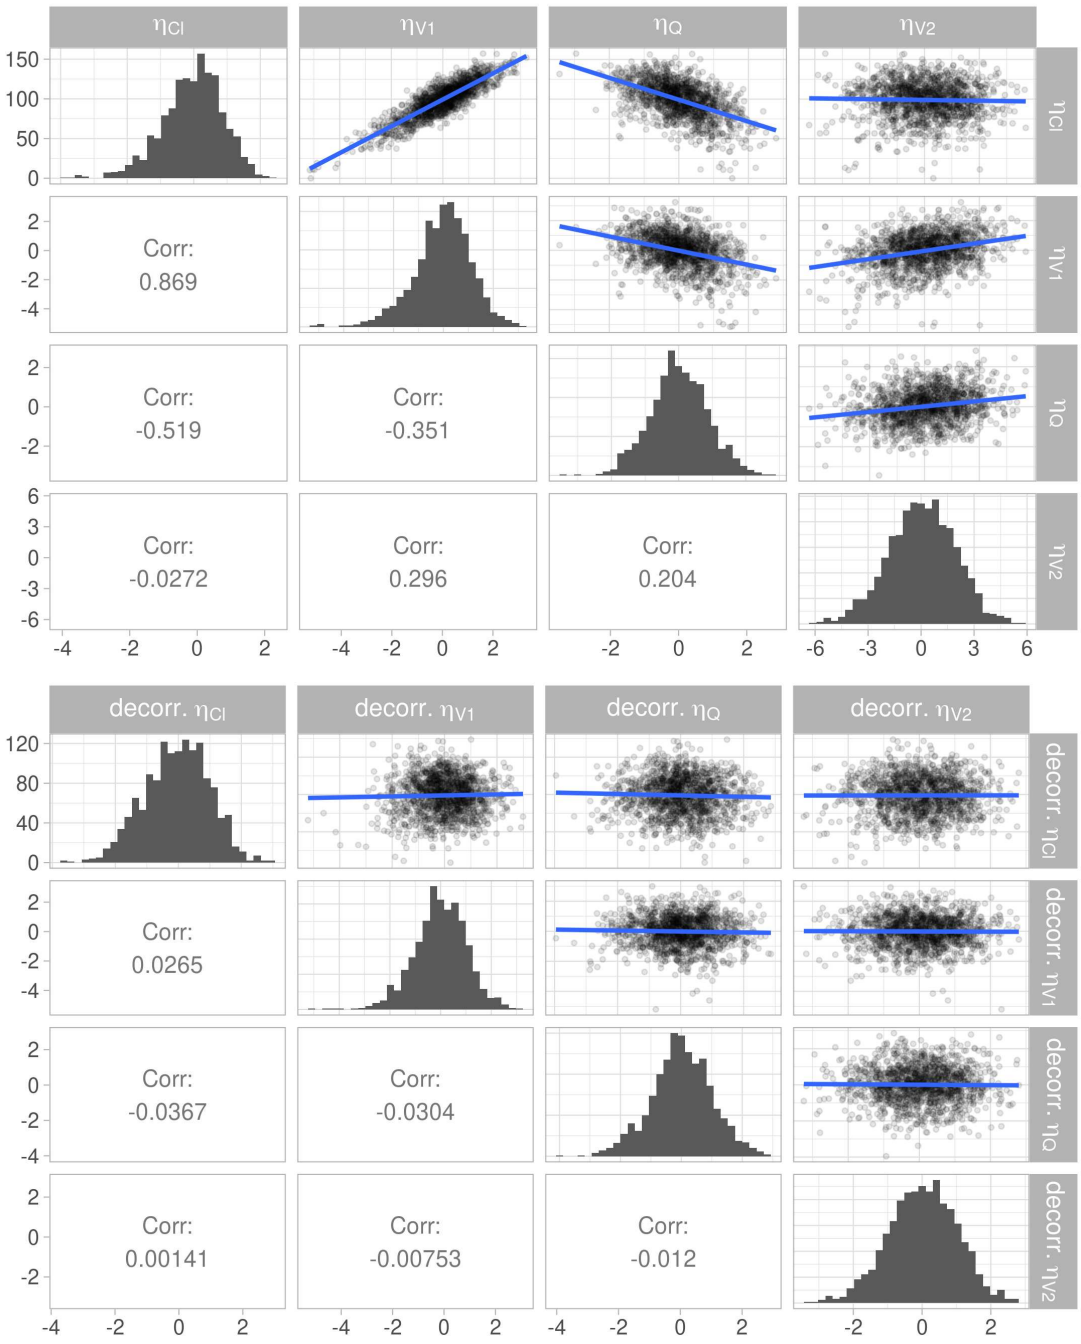

Figure 4.23: (A) Random effects correlations and (B) decorrelated random effects correlations.

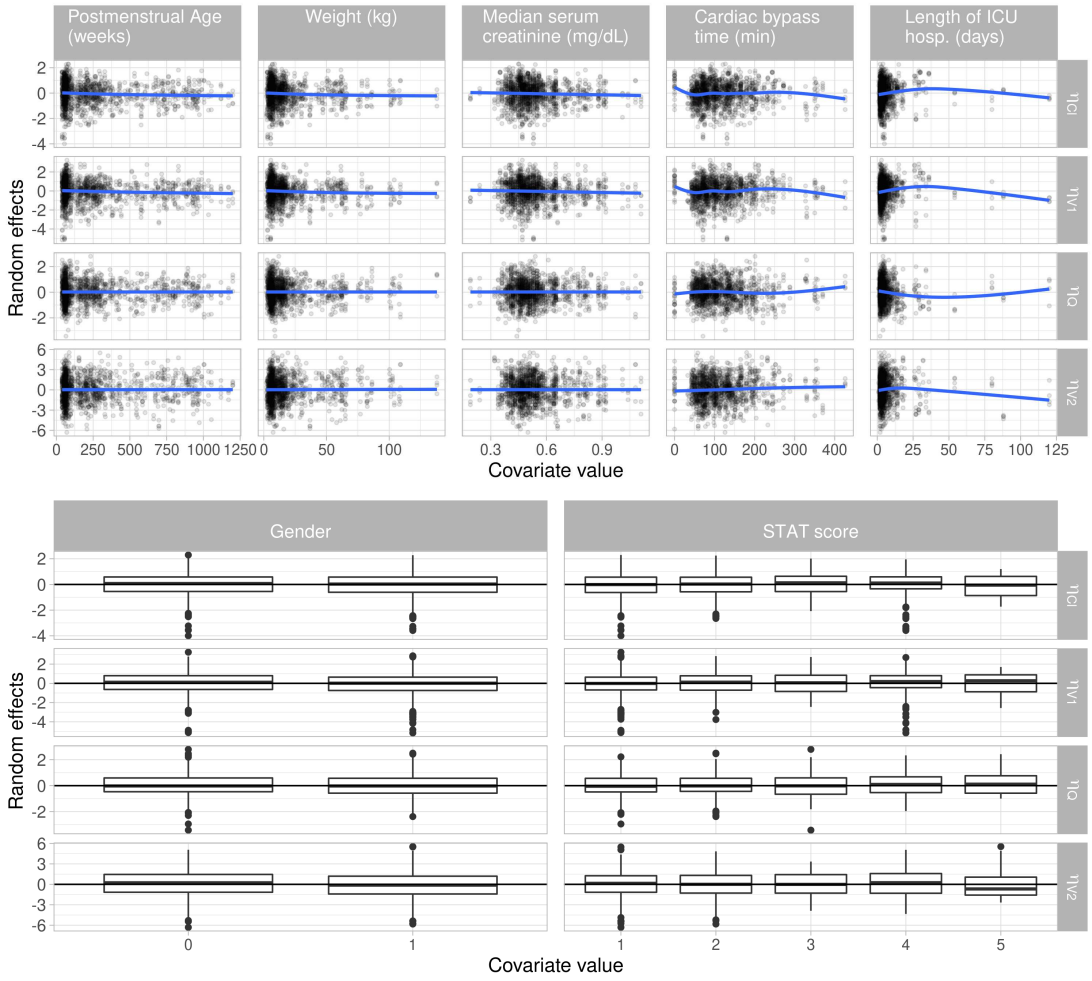

Figure 4.24: (A) Random effects vs. continuous covariates and (B) random effects vs. categorical covariates.

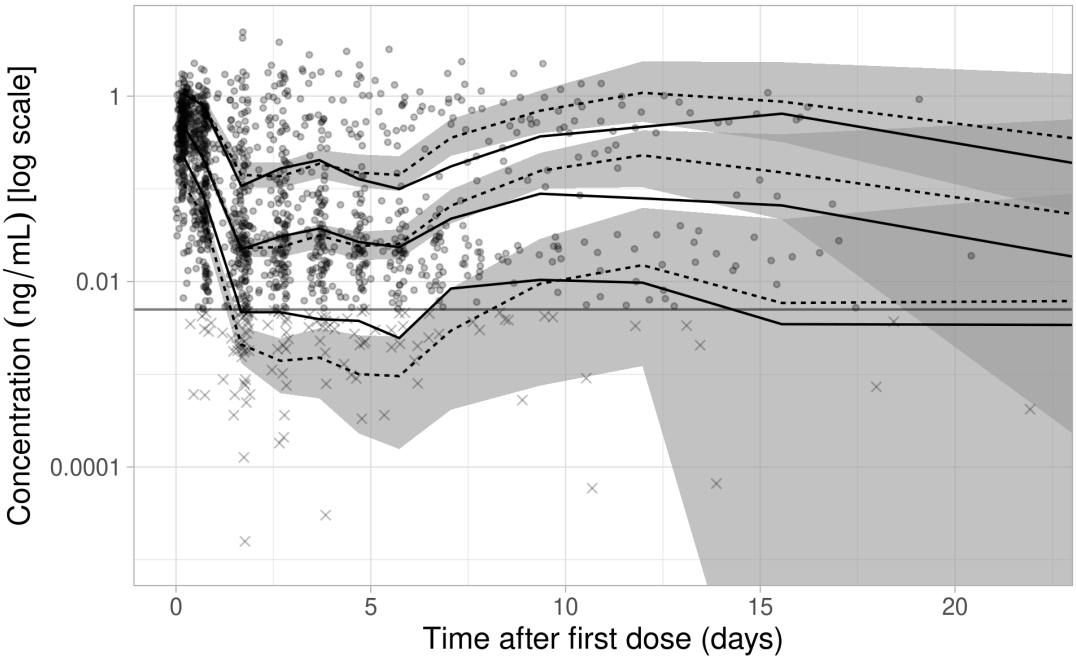

Figure 4.25: Prediction corrected visual predictive check.

4.5.3.4 Model 13 Goodness-of-fit Plots

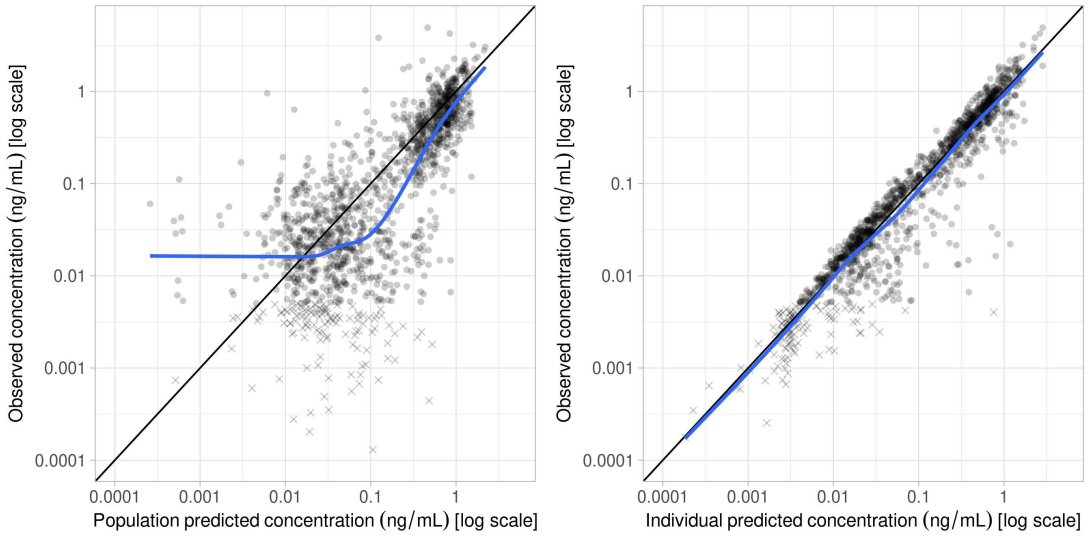

Figure 4.26: (A) Observed vs. population predicted concentrations and (B) observed vs. individual predicted concentrations.

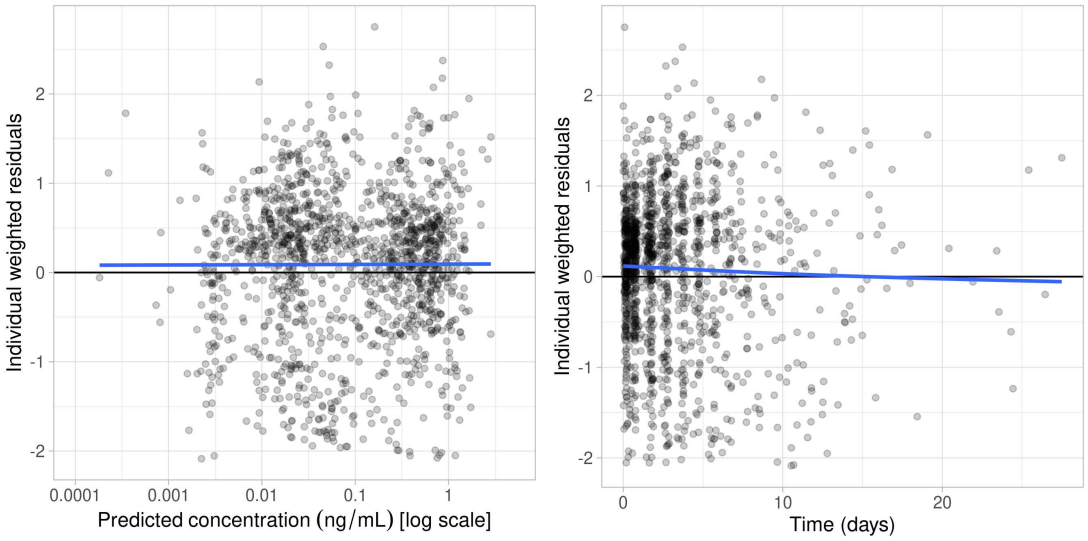

Figure 4.27: (A) Individual weighted residuals vs. predicted concentration and (B) individual weighted residuals vs. time.

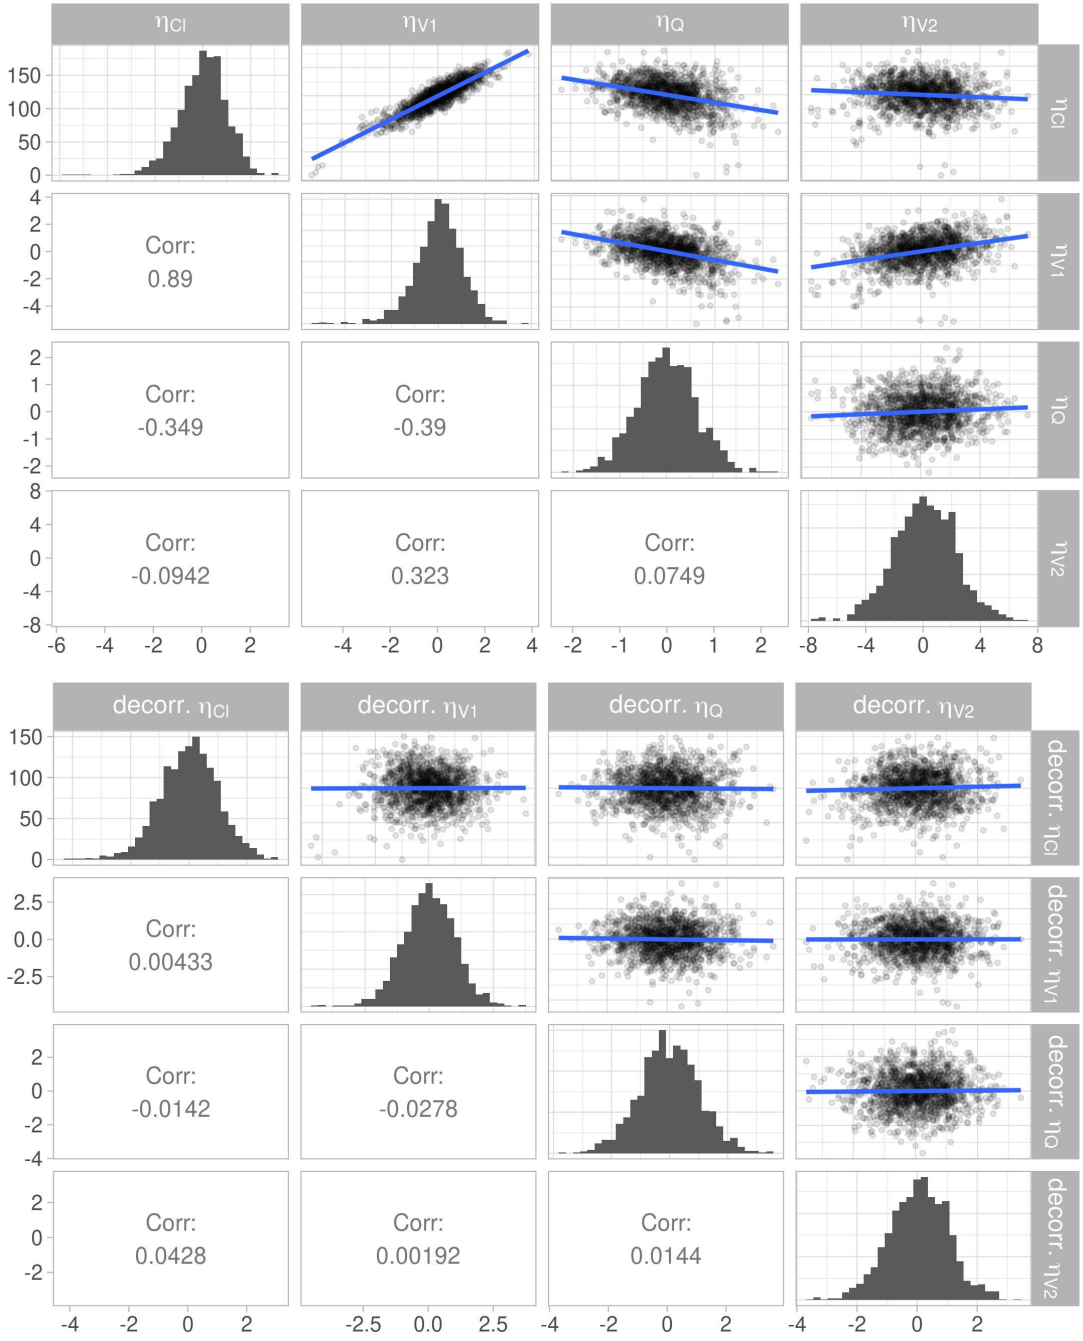

Figure 4.28: (A) Random effects correlations and (B) decorrelated random effects correlations.

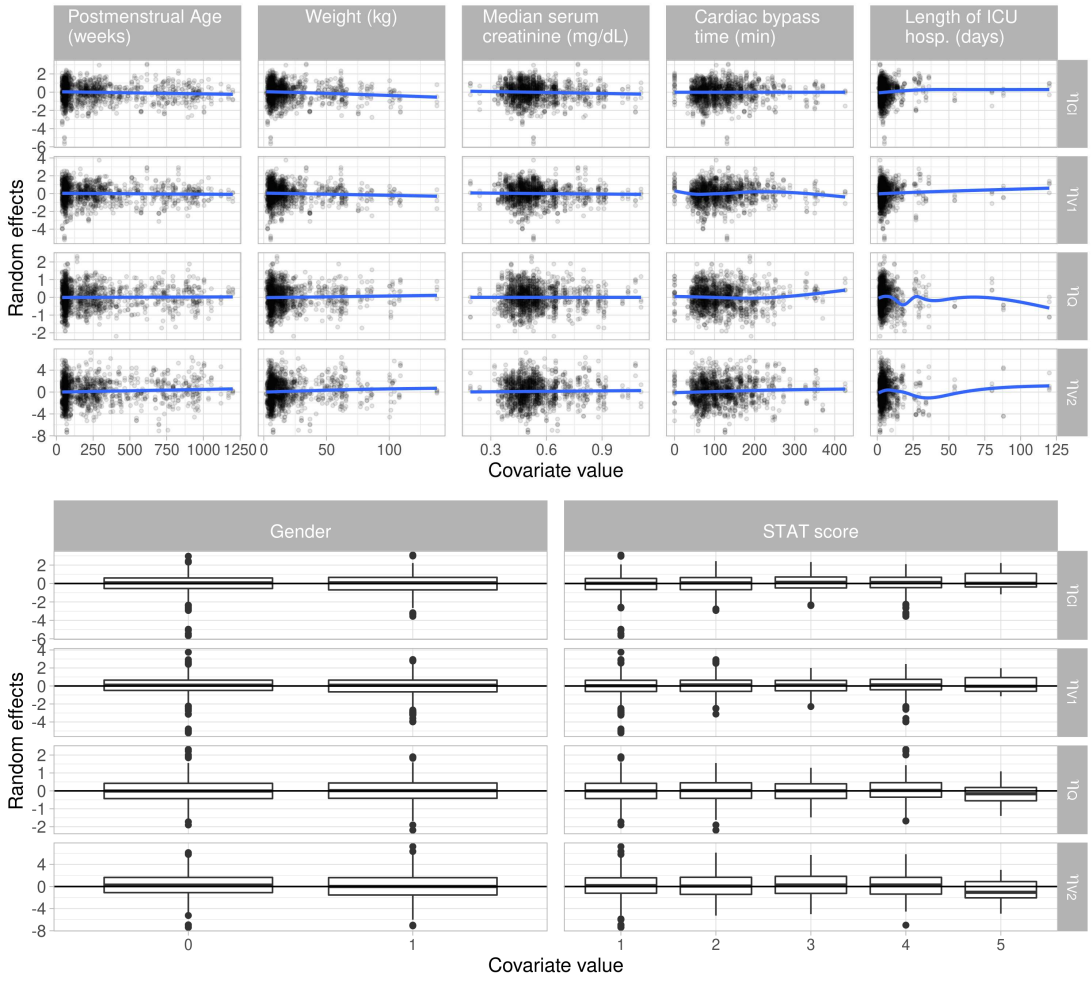

Figure 4.29: (A) Random effects vs. continuous covariates and (B) random effects vs. categorical covariates.

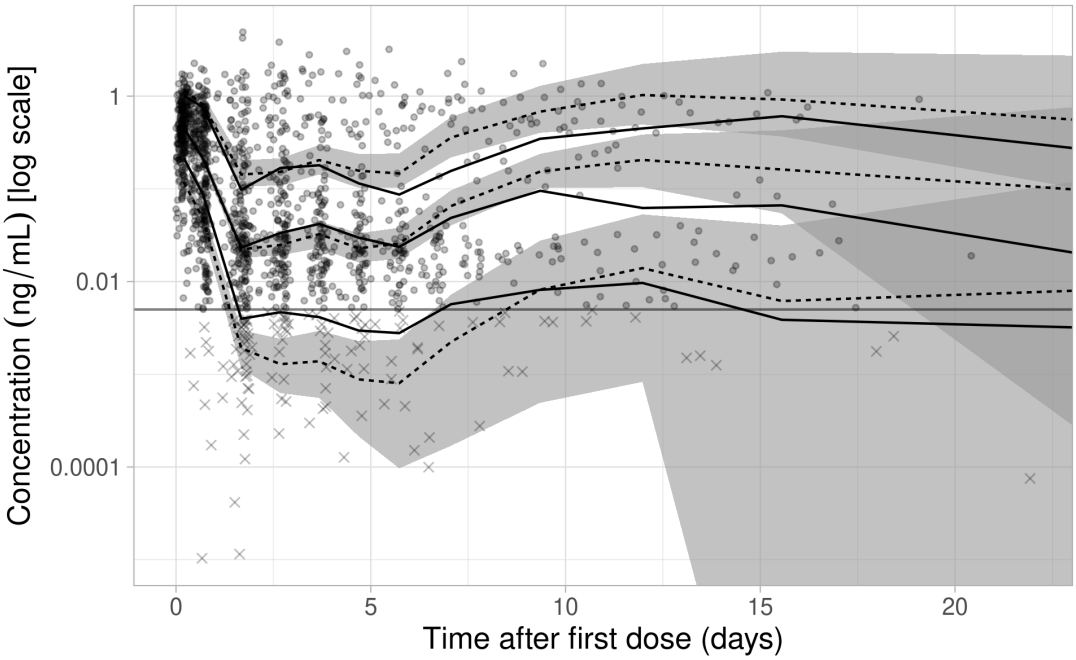

Figure 4.30: Prediction corrected visual predictive check.

4.5.3.5 Model 15 Goodness-of-fit Plots

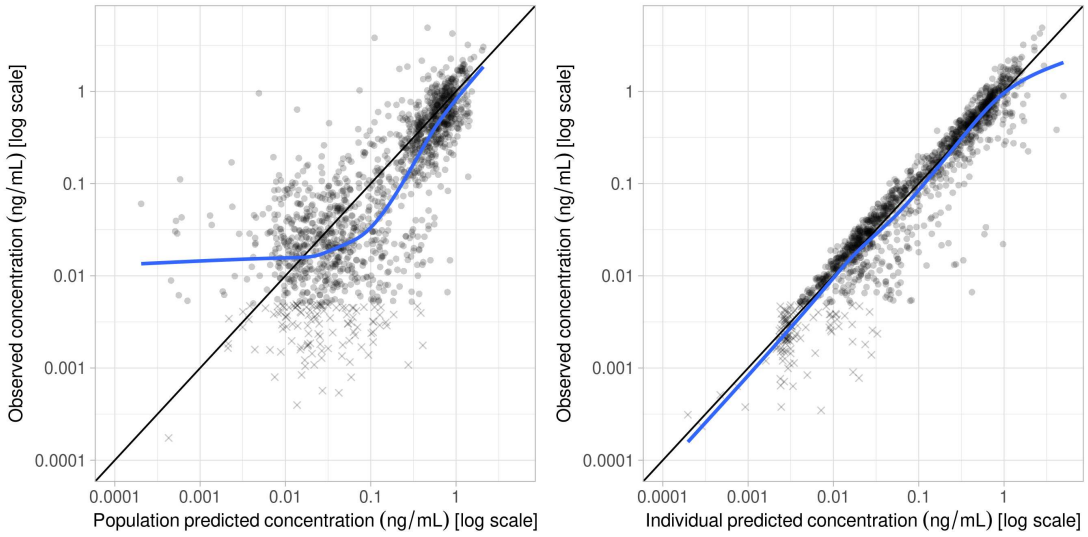

Figure 4.31: (A) Observed vs. population predicted concentrations and (B) observed vs. individual predicted concentrations.

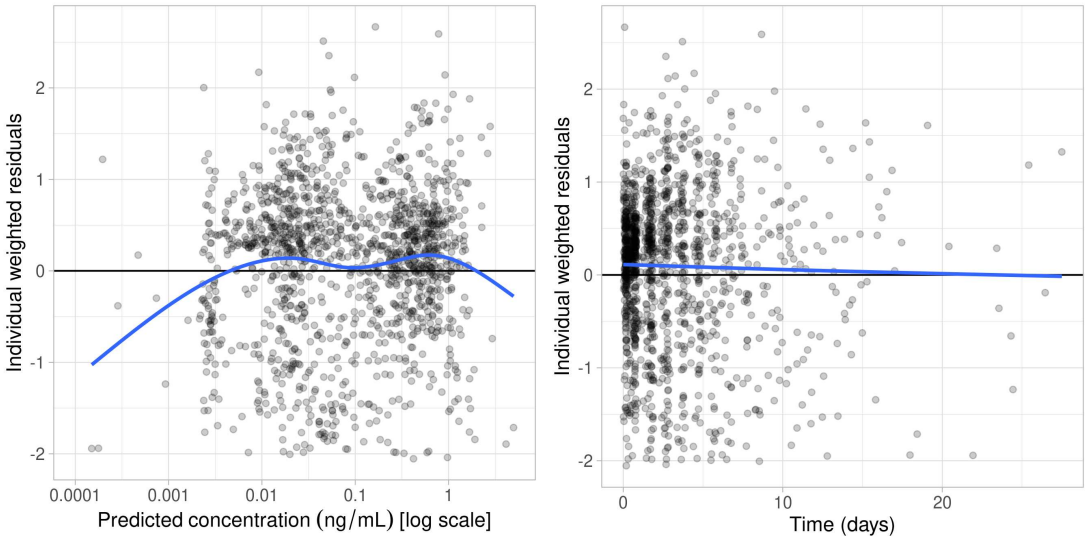

Figure 4.32: (A) Individual weighted residuals vs. predicted concentration and (B) individual weighted residuals vs. time.

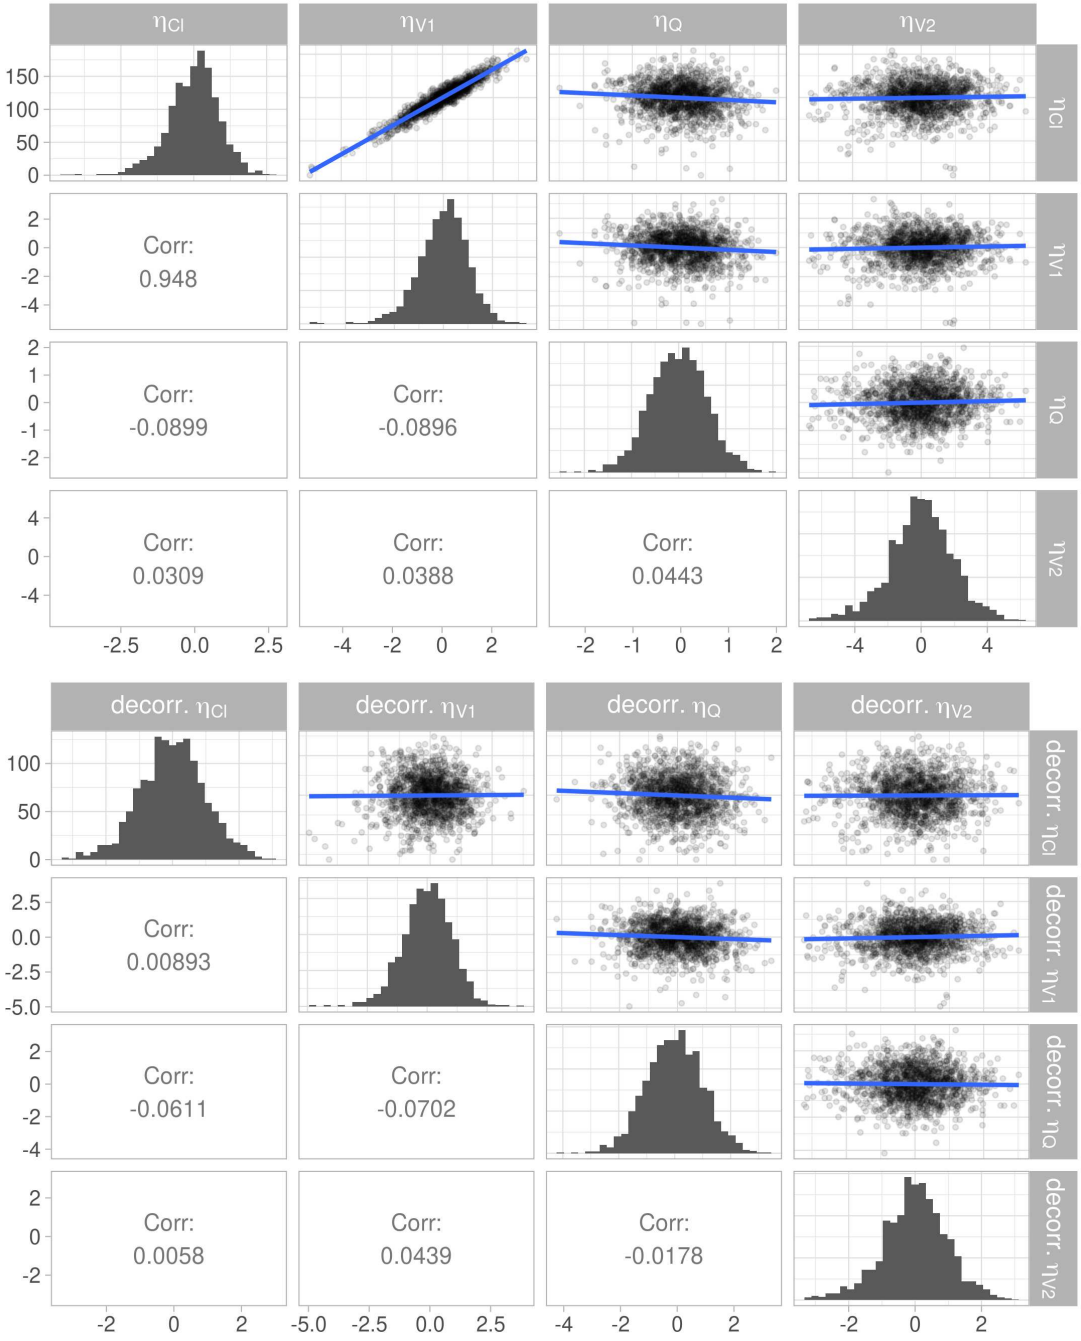

Figure 4.33: (A) Random effects correlations and (B) decorrelated random effects correlations.

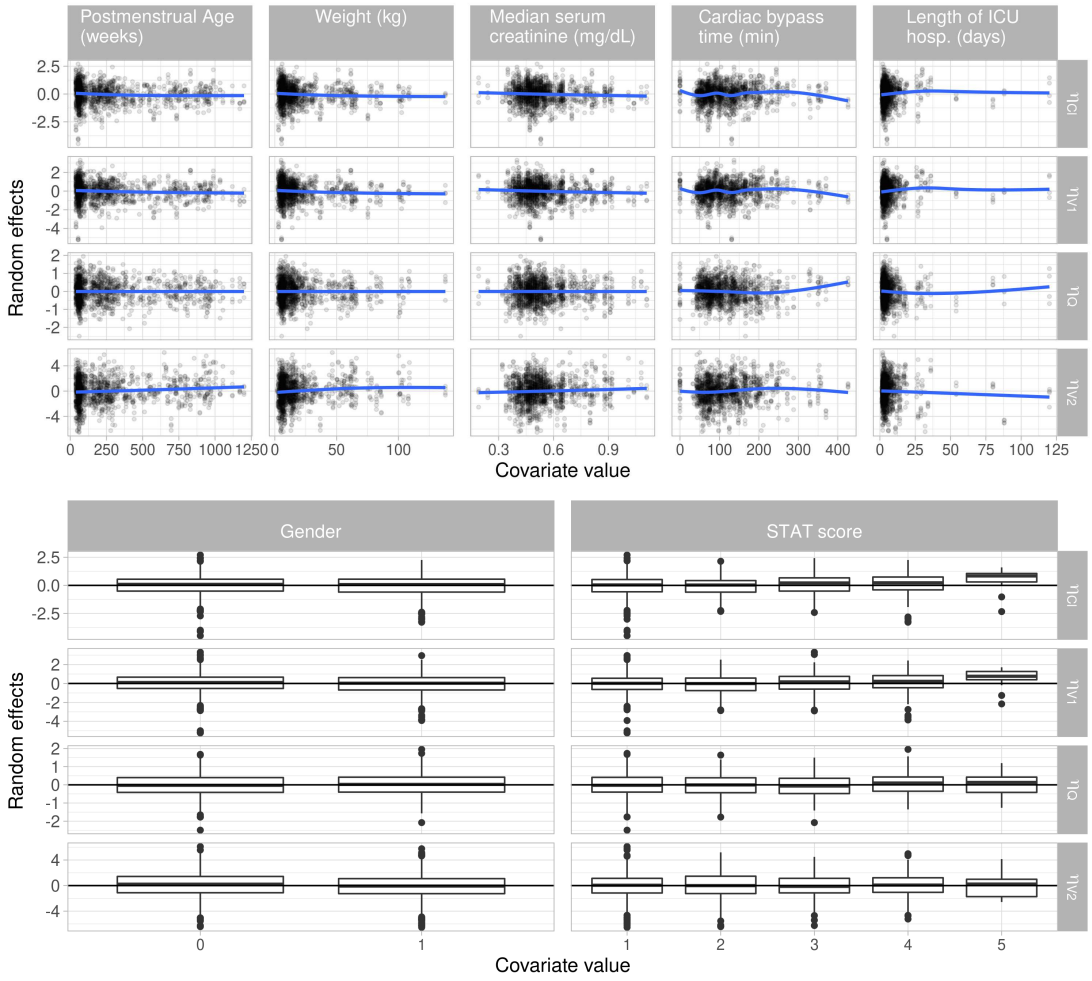

Figure 4.34: (A) Random effects vs. continuous covariates and (B) random effects vs. categorical covariates.

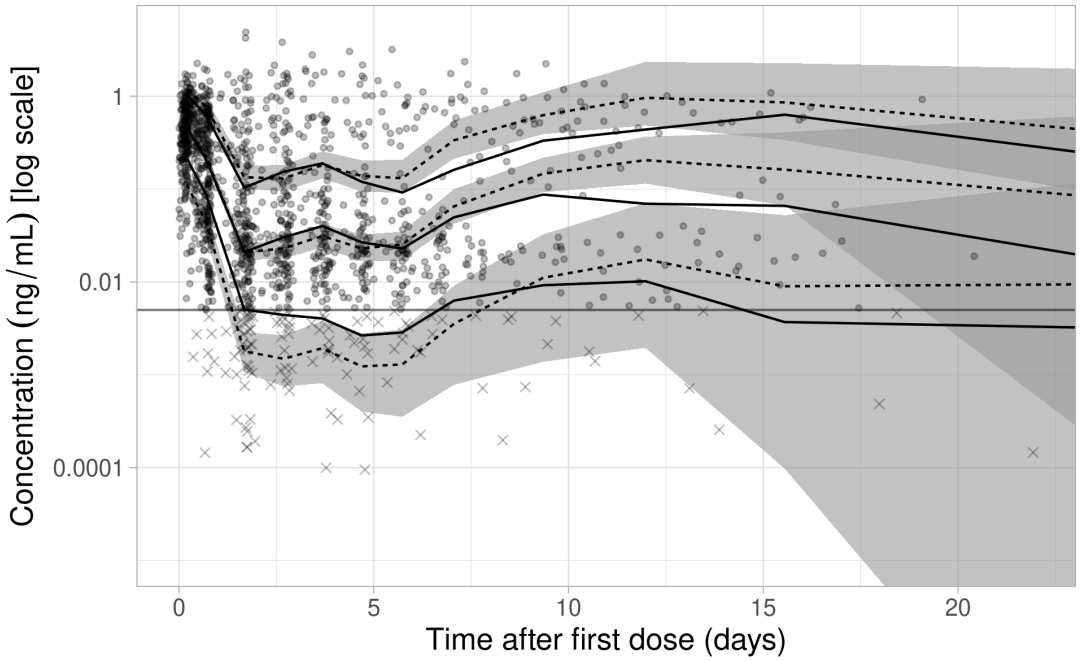

Figure 4.35: Prediction corrected visual predictive check.

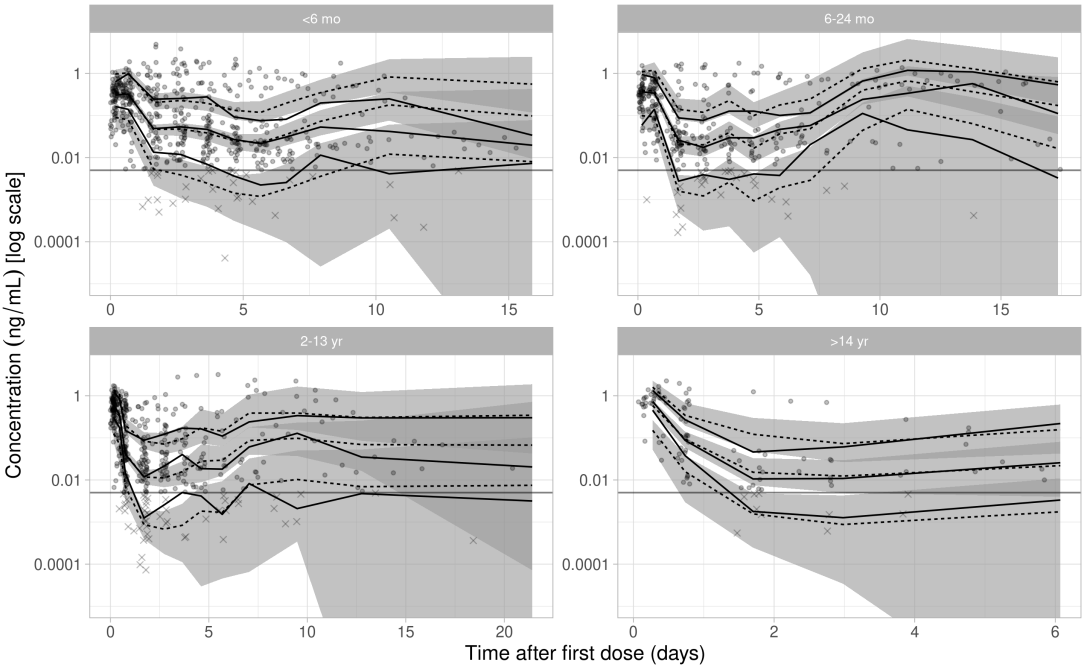

Figure 4.36: Prediction corrected visual predictive check for final model (Model 16) stratified by postnatal age group.

4.5.3.6 Genetic Effects on Clearance and Concentration

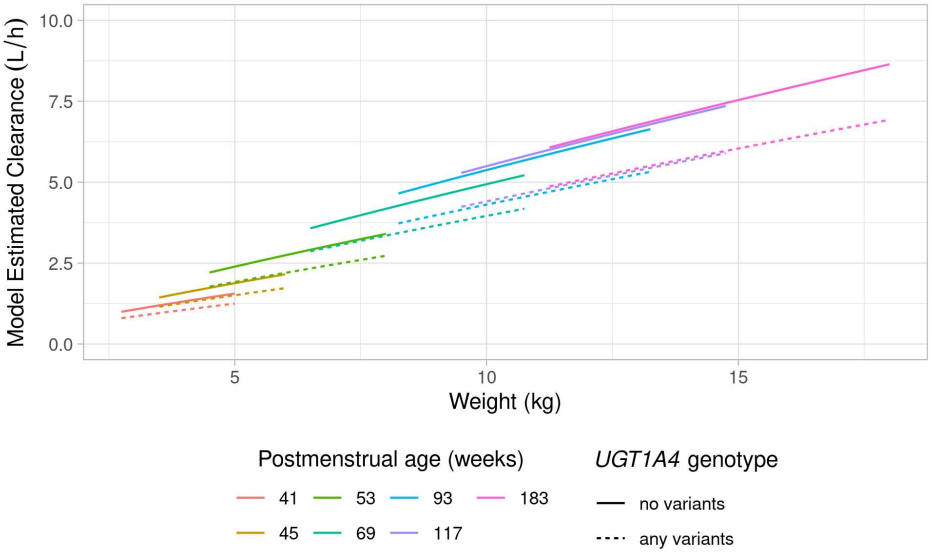

Figure 4.37: Model estimated clearance by $UGT1A4$ genotype.

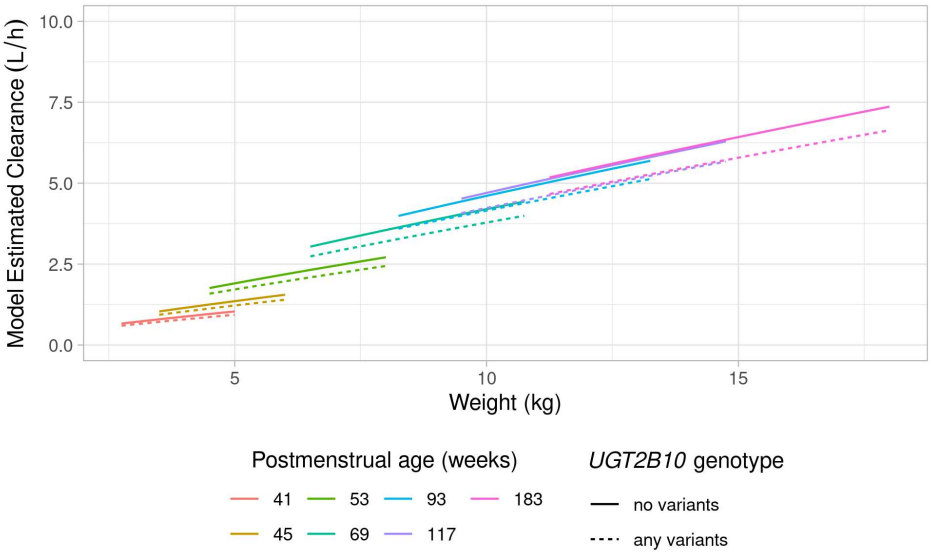

Figure 4.38: Model estimated clearance by $UGT2B10$ genotype.

4.5.3.6.1 Simulations of Concentration

To demonstrate the hypothetical effect of genotype covariates on predicted dexmedetomidine concentration, we estimated predicted concentration-time profiles using PK parameter estimates from the $UGT1A4$ and $UGT2B10$ categorical models. 90% confidence bounds were produced using quantiles from 200 simulated subjects with residual error (proportional component standard deviation = 0.478 ng/mL) for the following combination of covariates and dosing groups: (1) the 5th, median, or 95th percentile of weight and postmen-

strual age; (2) with or without $UGT1A4$ or $UGT2B10$ genotype variants; (3) fixed infusion rates of 0.4 or 0.6 mcg/kg/h for a 12-hour infusion. Because we are primarily interested in differences in concentration for the 'same subject' with and without variants, inter-individual variability was not included.

Results are shown in Figures 4.39 and 4.40. As expected, subjects with $UGT1A4$ variants have higher predicted concentration than those without variants, however the difference between genotypes is much smaller than the residual variability within subjects represented by the colored regions. For the 0.4 mcg/kg/h dosing rate, the predicted concentration levels are below the target range of 0.4 to 0.8 ng/mL for subjects at the 5th and median of weight and age both with and without variants. In contrast, the 0.6 mcg/kg/h rate yields predicted concentration levels within the target range for both genotype groups. Similar patterns are also seen for simulated concentrations from the $UGT2B10$ model.

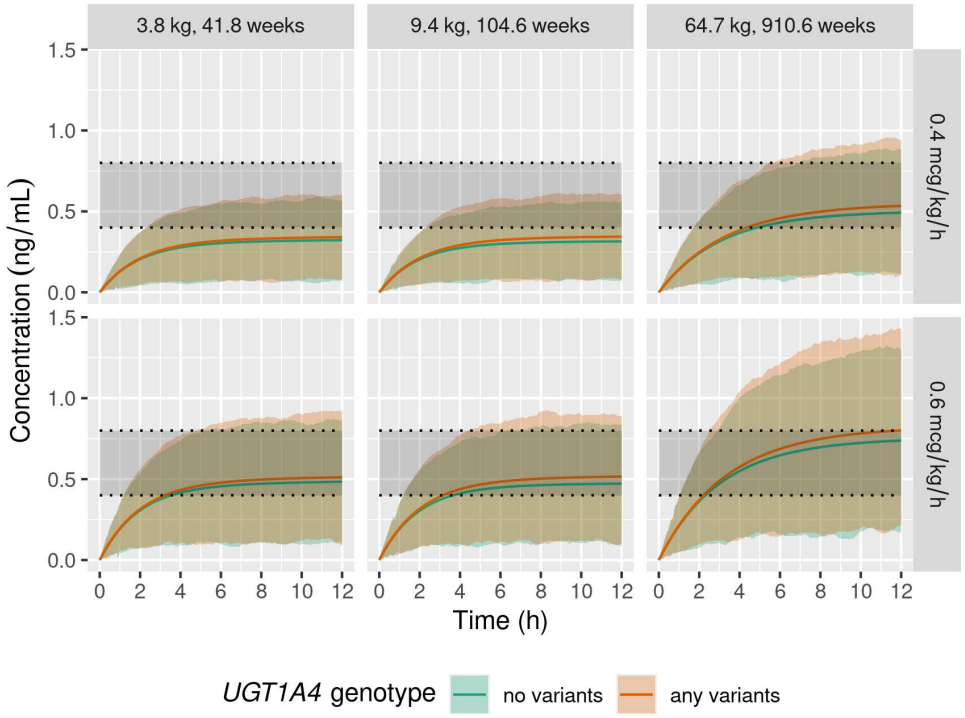

Figure 4.39: Predicted concentration simulations from $UGT1A4$ categorical model.

We also performed simulations to find the infusion rates that would yield similar concentrations for subjects of the same age and weight with and without genetic variants after a 12-hour infusion. First, concentration-time profiles were generated using identical dosing rate (0.6 mcg/kg/h), weight (50kg), and postmenstrual age (520 weeks) for a simulated subject with and without $UGT1A4$ or $UGT2B10$ variants, respectively. Then the infusion rate for the simulated subject with variants was adjusted in increments of 0.01 mcg/kg/h until the concentration achieved at the end of the 12-hour infusion most closely matched the concentration of

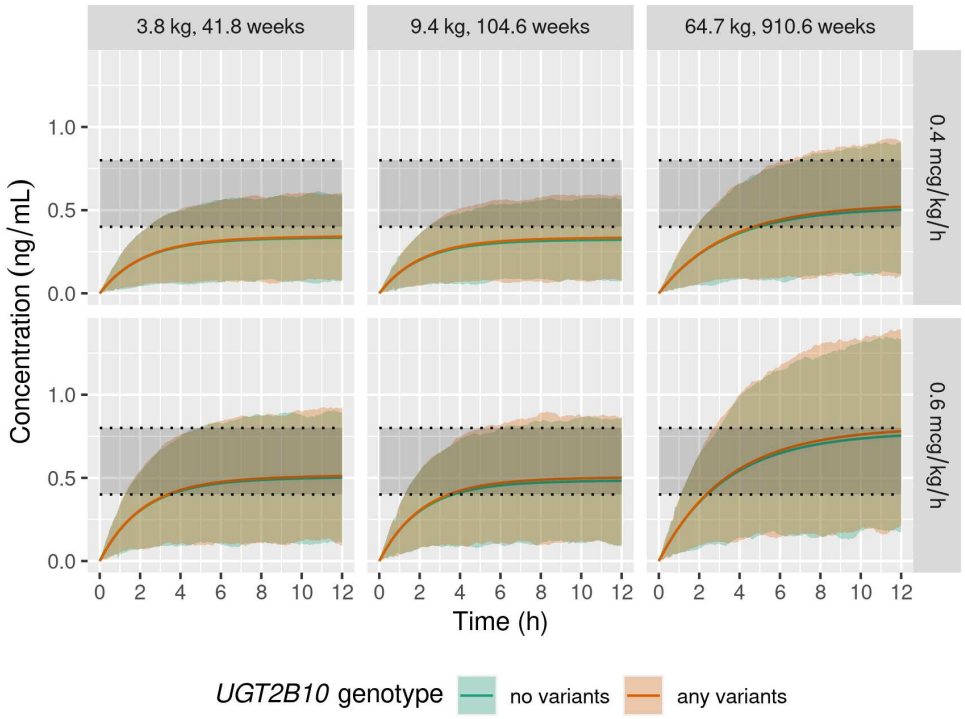

Figure 4.40: Predicted concentration simulations from $UGT2B10$ categorical model.

the simulated subject without variants. The `R` package `mrgsolve` [153] was used for all simulations. The simulated patient with $UGT1A4$ variants required a rate of 0.56 mcg/kg/h to approximate the concentration of those without $UGT1A4$ variants at a rate of 0.6 mcg/kg/h. The simulated patient with $UGT2B10$ variants required a rate of 0.58 mcg/kg/h to approximate the concentration of those without $UGT1A4$ variants at a rate of 0.6 mcg/kg/h. These changes in dose are not large enough to impact clinical dosing as increments for dexmedetomidine titration are typically 0.1 mcg/kg/h.

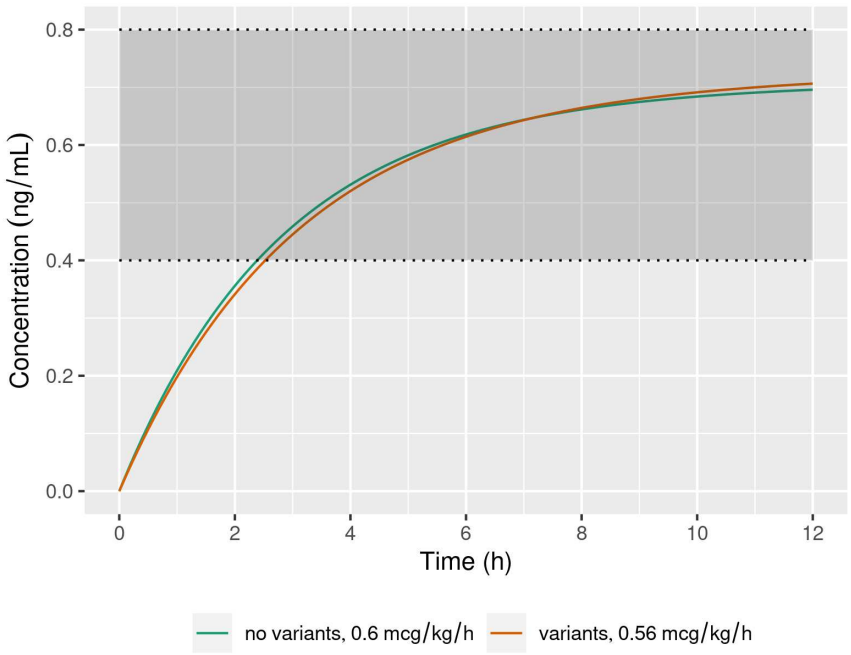

Figure 4.41: Simulated Dose needed to achieve same concentration for patients of same postmenstrual age (520 weeks) and weight (50 kg) with and without variants from $UGT1A4$ model.

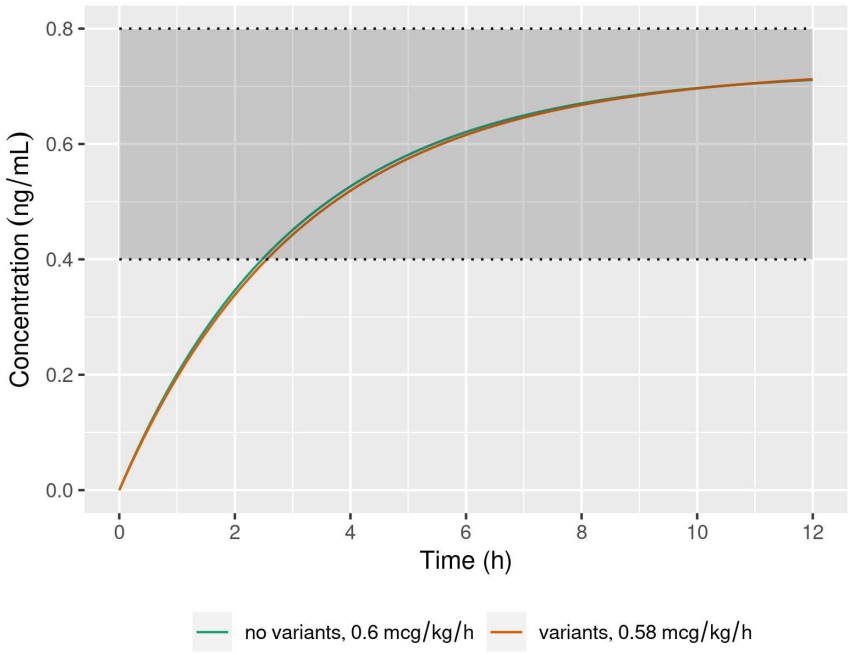

Figure 4.42: Simulated Dose needed to achieve same concentration for patients of same postmenstrual age (520 weeks) and weight (50 kg) with and without variants from $UGT2B10$ model.

4.5.4 Recommended Protocols for Analgesia and Sedation Drug and Dose Selection

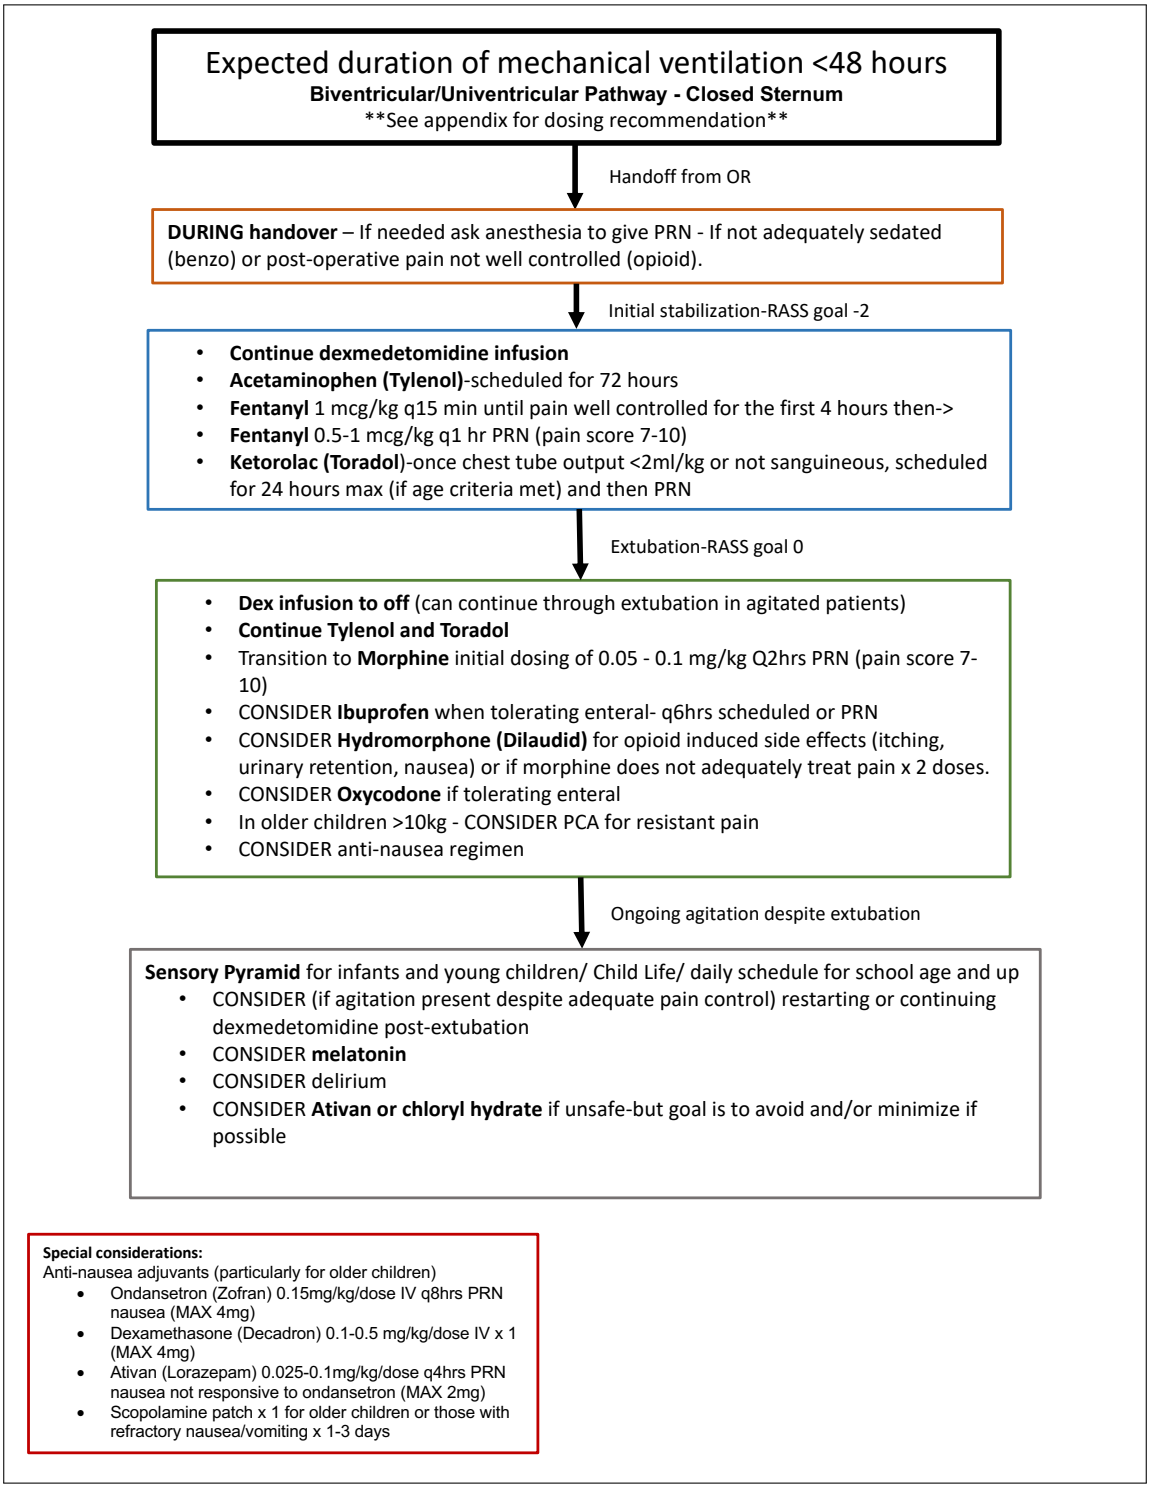

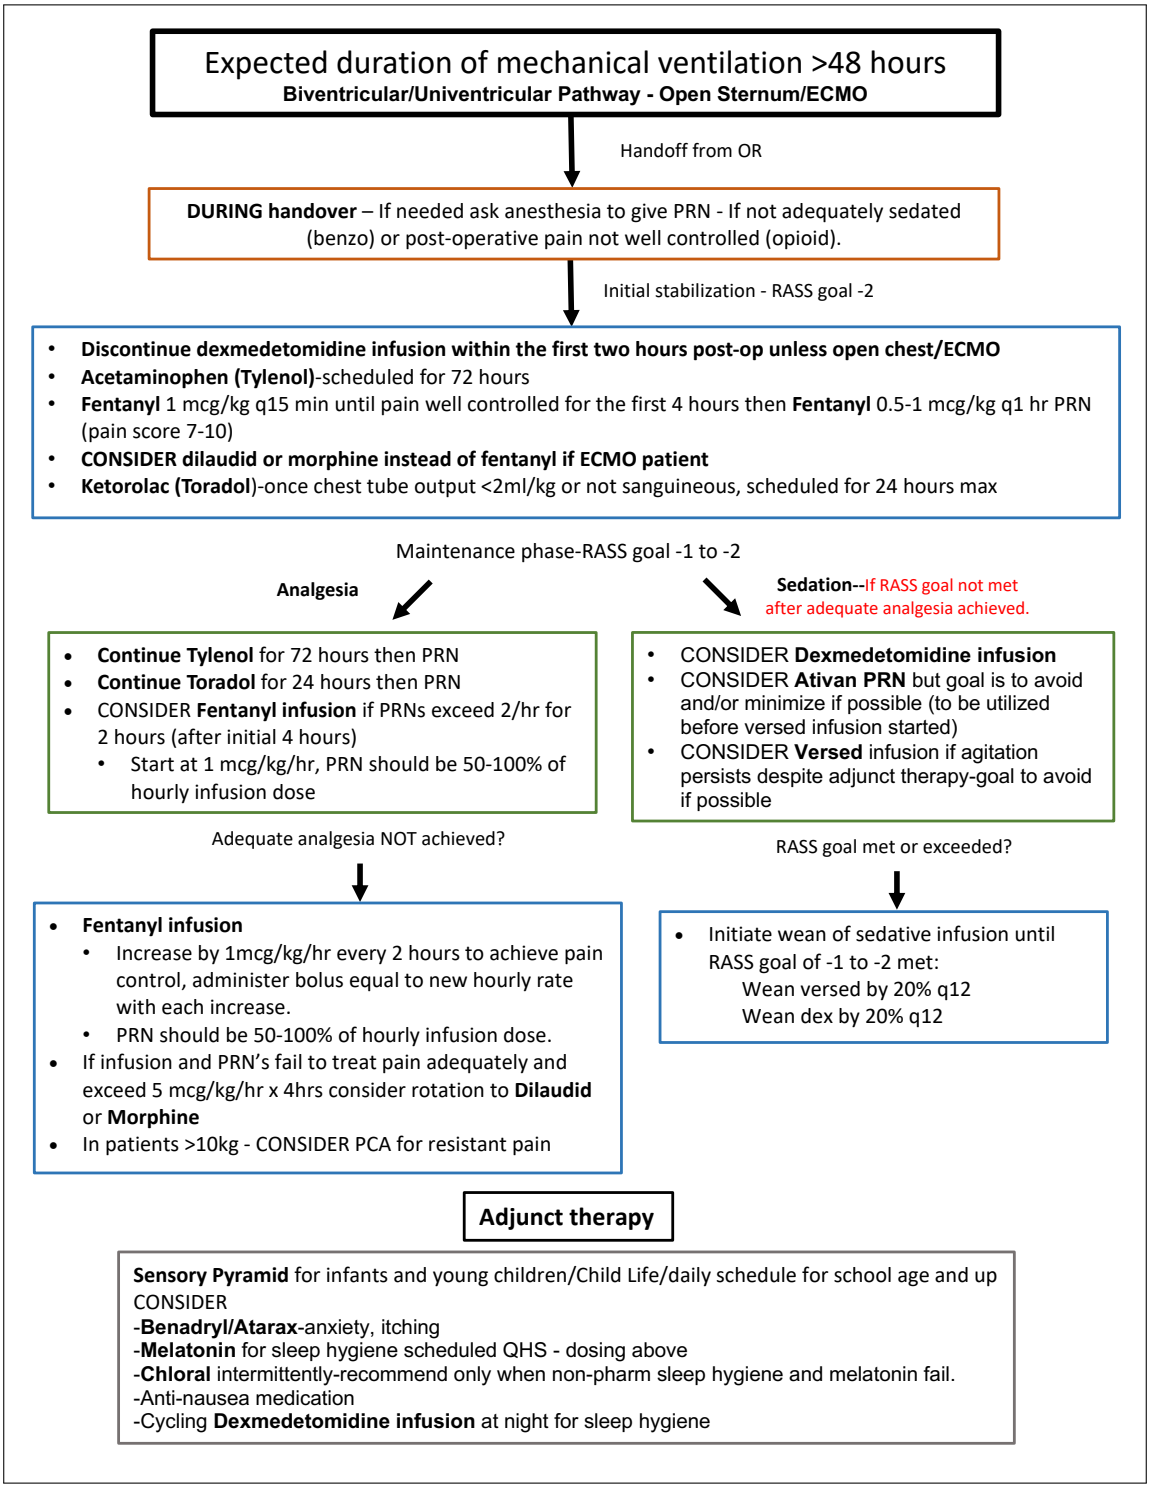
Expected duration of mechanical ventilation >48 hours
Biventricular/Univentricular Pathway - Open Sternum/ECMO
Handoff from OR
DURING handover – If needed ask anesthesia to give PRN - If not adequately sedated (benzo) or post-operative pain not well controlled (opioid).
Initial stabilization - RASS goal -2
Discontinue dexmedetomidine infusion within the first two hours post-op unless open chest/ECMO
Acetaminophen (Tylenol)-scheduled for 72 hours
Fentanyl 1 mcg/kg q15 min until pain well controlled for the first 4 hours then Fentanyl 0.5-1 mcg/kg q1 hr PRN (pain score 7-10)
CONSIDER dilaudid or morphine instead of fentanyl if ECMO patient
Ketorolac (Toradol)-once chest tube output <2ml/kg or not sanguineous, scheduled for 24 hours max
Maintenance phase-RASS goal -1 to -2
Analgesia
Sedation--If RASS goal not met after adequate analgesia achieved.
Continue Tylenol for 72 hours then PRN
Continue Toradol for 24 hours then PRN
CONSIDER Fentanyl infusion if PRNs exceed 2/hr for 2 hours (after initial 4 hours)
Start at 1 mcg/kg/hr, PRN should be 50-100% of hourly infusion dose
CONSIDER Dexmedetomidine infusion
CONSIDER Ativan PRN but goal is to avoid and/or minimize if possible (to be utilized before versed infusion started)
CONSIDER Versed infusion if agitation persists despite adjunct therapy-goal to avoid if possible
Adequate analgesia NOT achieved?
RASS goal met or exceeded?
Fentanyl infusion
Increase by 1mcg/kg/hr every 2 hours to achieve pain control, administer bolus equal to new hourly rate with each increase.
PRN should be 50-100% of hourly infusion dose.
If infusion and PRN's fail to treat pain adequately and exceed 5 mcg/kg/hr x 4hrs consider rotation to Dilaudid or Morphine
In patients >10kg - CONSIDER PCA for resistant pain
Initiate wean of sedative infusion until RASS goal of -1 to -2 met:
Wean versed by 20% q12
Wean dex by 20% q12
Adjunct therapy
Sensory Pyramid for infants and young children/Child Life/daily schedule for school age and up
CONSIDER
-Benadryl/Atarax-anxiety, itching
-Melatonin for sleep hygiene scheduled QHS - dosing above
-Chloral intermittently-recommend only when non-pharm sleep hygiene and melatonin fail.
-Anti-nausea medication
-Cycling Dexmedetomidine infusion at night for sleep hygiene

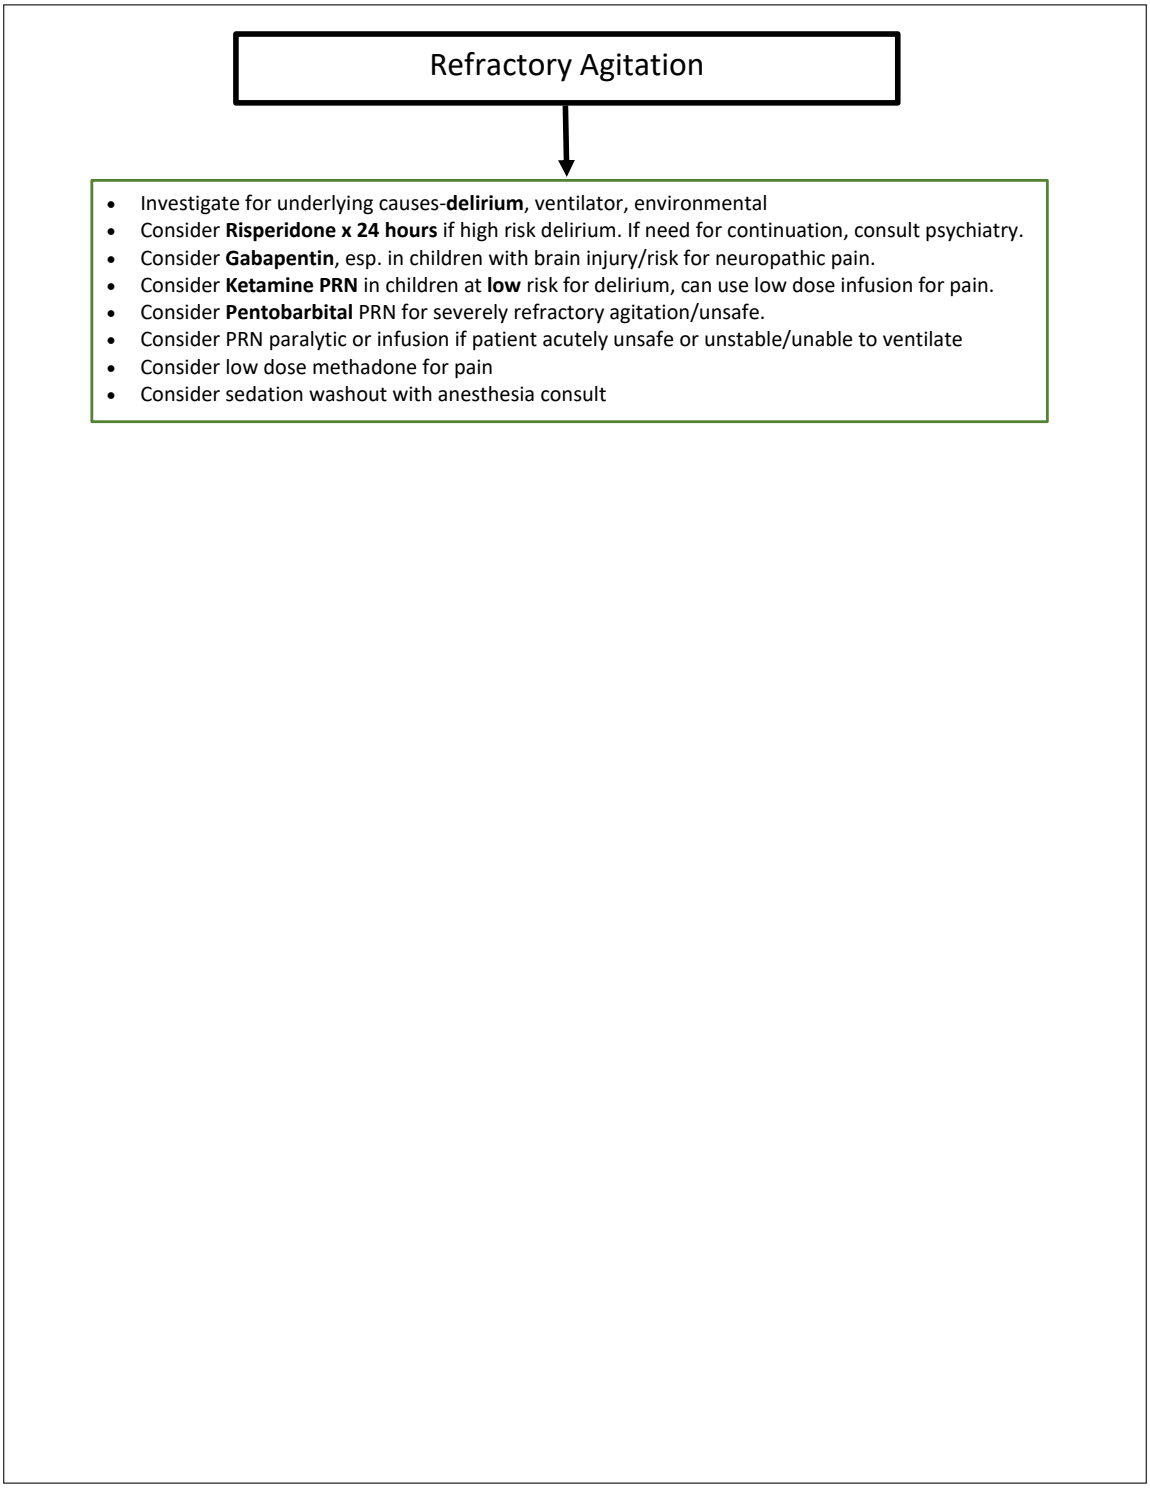
Refractory Agitation
Investigate for underlying causes-**delirium**, ventilator, environmental
Consider **Risperidone x 24 hours** if high risk delirium. If need for continuation, consult psychiatry.
Consider **Gabapentin**, esp. in children with brain injury/risk for neuropathic pain.
Consider **Ketamine PRN** in children at **low** risk for delirium, can use low dose infusion for pain.
Consider **Pentobarbital** PRN for severely refractory agitation/unsafe.
Consider PRN paralytic or infusion if patient acutely unsafe or unstable/unable to ventilate
Consider low dose methadone for pain
Consider sedation washout with anesthesia consult

CHAPTER 5

Bayesian Population Pharmacokinetic Analysis of Dexmedetomidine in Children using Real World Data and Informative Priors

5.1 Introduction

An analysis by James et al. (2022) [110], included as Chapter 4, explored the impact of patient characteristics and pharmacogenetics on dexmedetomidine pharmacokinetics (PK) using electronic health record (EHR) data and remnant samples under the classical 'frequentist' statistical paradigm. The analysis confirmed the importance of weight and age as predictors of total clearance in the pediatric population, but did not find evidence for pharmacogenetic effects of $UGT1A4$ or $UGT2B10$ genotype or $CYP2A6$ risk score. However, that analysis exhibited several limitations. First, the observed drug concentration outcome data were quite sparse with a median of only 1.2 concentration samples per subject per day[1]. Relatedly, because the outcome data were not collected with the goal of estimating or monitoring PK characteristics such as trough concentration or C_{max}, as in an optimally designed experiment, they contain less information to estimate some parameters, even using a population PK analysis which borrows information across subjects. In addition, real-world dosing data are much more complex than the simple dosing scenarios (often just a single dose) used for designed experiments with large variations in the amount, frequency, and timing of doses from subject to subject. Finally, observational data sourced from EHRs are more prone to missingness and data entry errors than those from designed PK studies. The stochastic approximation expectation-maximization (SAEM) algorithm used in frequentist analysis encountered occasional convergence problems and included simplifications such as fixing most random effects correlations at zero to improve estimation and model fit. The analysis also produced population PK estimates which substantially differed from previous studies for some parameters. For instance, the peripheral compartment volume of distribution for the two-compartment model was 100 times larger than estimates reported in other pediatric dexmedetomidine PK studies while estimated total clearance was 10 to 15 L/h smaller (with larger coefficient of variation).

To address these limitations, we reanalyze the real-world EHR and remnant specimen data using Bayesian methodology to incorporate prior information derived from several recent pediatric dexmedetomidine PK studies. The Bayesian analysis provides a coherent framework for combining prior knowledge (or external

[1]For comparison, Ber et al. (2020) collected 1 to 3 samples per subject per day during infusion and 8 samples within 6 hours after end of infusion [115]; Zuppa et al. (2019) collected at least 4 samples per subject over 12 hours during infusion and at least 8 samples within 24 hours after end of infusion [127]; Zimmerman et al. (2019) collected at least 5 samples per subject during cardiac bypass procedures averaging 180 minutes [126]; Song et. al (2019) collected 3 samples per subject within 60 minutes of start of infusion and 6 samples within 8 hours after end of infusion [122].

evidence) and current data and helps stabilize estimation by borrowing information about parameters from smaller, more densely sampled previous studies. The Bayesian population PK method also has advantages beyond incorporating prior information. First, it produces parameter estimates and intervals that are easier to interpret than frequentist estimates. For example, a 90% posterior credible interval of 35 to 45 L/h for population total clearance has the intuitive interpretation that given the prior information and current data, there is a 90% probability that clearance is between 35 and 45 L/h. In contrast, a frequentist 90% confidence interval is a statement about how the interval is built; it is calculated using a procedure such that 90% of all constructed intervals will contain the true value of clearance. In addition, inference for Bayesian population PK analysis uses samples from the posterior distribution (or an approximation to the posterior) which has several benefits. For instance, parameter uncertainty is quantified without using asymptotic arguments which may be difficult to assess or post hoc procedures such as bootstrapping (although the accuracy of the Bayesian approximation must also be considered) [61,62]. It is also straightforward to derive point and interval estimates for functions of parameters (such as maximum concentration).

Despite these advantages, fitting Bayesian models is more time-consuming in two ways. First, more time must be devoted to specifying the complete model, namely the prior distribution, which is absent from the frequentist paradigm. Second, fitting Bayesian population PK models using Markov chain Monte Carlo (MCMC) takes more computing time than frequentist estimation [154,155]. Therefore, we also demonstrate the use of automatic differentiation variational inference (ADVI) to more quickly provide approximate solutions during Bayesian model development and selection, as described in Chapter 3. After using ADVI for model development, the final models are refit using more accurate MCMC methods, specifically Hamiltonian Monte Carlo (HMC).

The major goals of the reanalysis are to provide estimates which are informed by both current and previous data to summarize our knowledge and uncertainty about dexmedetomidine PK and to quantify genetic effects that were selected based on previous studies and known metabolic pathways. Ultimately, better understanding the factors affecting dexmedetomidine PK can facilitate more precise, personalized dosing.

5.2 Methods

The data for this reanalysis are identical to those in the original work by James et al. (2022). The study design, data collection, drug concentration measurement and laboratory analysis, genotyping and *CYP*2*A*6 activity score prediction, and data processing are described in detail in that paper and Van Driest et al. (2016) [110,136]. In short, observational data, including remnant clinical testing specimens for drug concentration measurements, peripheral blood for genetic analysis, dexmedetomidine dosing from the EHR, and other demographic, surgical,

and clinical information were collected for pediatric patients undergoing surgery for congenital heart disease. Specimens were not collected in connection with dose administration or to monitor PK characteristics. Data from all sources were combined and processed using the `R` software [138] package `EHR` [135].

5.2.1 Bayesian Population Pharmacokinetic Analysis

We performed population PK analysis using the three stage hierarchical Bayesian nonlinear modeling framework of Wakefield, Aarons, and Racine-Poon (1999) described in detail in Chapter 1 (section 1.3.3) [62]. Concentrations below the LLOQ were considered to be censored between 0 and 0.005 ng/mL and were handled in the modeling by integrating the first-stage conditional distribution between these bounds[2] [63]. For model development and exploration, models were compared numerically using out-of-sample predictive accuracy estimated by 10-fold leave-subject-out cross-validation log pointwise predictive density ($lppd_{cv}$) with ADVI estimation. In simulations of Bayesian population PK models estimated with ADVI (described in Chapter 3), $lppd_{cv}$ had better performance for model selection than information criteria such as DIC or WAIC which are based on observation-level partitioning of data. Specifically, we create $K = 10$ groups each containing approximately $\frac{1}{10}$ of the subjects. Each group contains all observations for subjects assigned to that group. The cross-validation log pointwise predictive density is $lppd_{cv} = \sum_{k=1}^{K} \sum_{i \in k} \log p_{post(-k)}(y_i)$ where $p_{post(-k)}$ is the posterior distribution for a model fit *without* the observations in fold k and the inner sum evaluates the log predictive density over the observations y_i in fold k. Summarizing $p_{post(-k)}$ by S simulation draws $\theta^{-k,s}$, the $lppd_{cv}$ can be computed as $\sum_{k=1}^{K} \sum_{i \in k} \log \left(\frac{1}{S} \sum_{s=1}^{S} p(y_i|\theta^{-k,s}) \right)$. Models with higher $lppd_{cv}$ have better out-of-sample predictive performance.

Models selected at each stage of development were refit using HMC to obtain final parameter estimates and perform posterior model checks. We used several graphical methods for model evaluation including plots of observed vs. population and individual concentration predictions, plots of samples from the conditional total clearance random effects distributions vs. covariates, and visual predictive checks [144,145]. Settings for both estimation methods are included in the Appendix. The models were implemented using the `CmdStanR` interface (version 0.4.0) to the `CmdStan` probabilistic programming language (version 2.25.0) along with the PK library `Torsten` (version 0.88) [65,66,87].

5.2.1.1 Model Development

Based on previous dexmedetomidine studies, we began with a two-compartment model with additive and proportional residual error and fixed allometric weight scaling. The main PK parameters for this model are

[2] E.g., the contribution to the likelihood for a censored observation is $\int_0^{0.005} p_{y_{ij}|\psi_i}(u|\psi_i)\,du$ compared to $p(y_{ij}|\psi_i)$ for an uncensored observation where the distribution $p(\cdot)$ is determined by the residual error model.

total clearance (CL, L/h), volume of distribution for the central compartment ($V1$, L), inter-compartmental clearance (Q, L/h) and volume of distribution for the peripheral compartment ($V2$, L). For inter-individual variability (IIV), we assumed a full (4×4) covariance matrix allowing correlations for all main PK parameters. Weight varied little within each subject during the time frame when concentration data were available, so baseline demographic weight was used for scaling.

Models including additional covariates were developed in several stages. First, we considered adding age maturation effects to the initial model since both size and maturation have been shown to be important factors in pediatric PK models with a large age range [115,118,120,125,147–149]. For maturation, we explored a sigmoid Hill model and an exponential model using postmenstrual age. Next, we considered adding non-genotype covariate effects on total clearance with the goal of building a model that adequately describes dexmedetomidine PK while accounting for IIV. Based on previous research and biological plausibility, we used the covariates sex, Society of Thoracic Surgery–European Association for Cardio-Thoracic Surgery (STAT) Congenital Heart Surgery Mortality score [146], cardiac bypass time, length of ICU stay, and serum creatinine. STAT score was included in the model as a continuous exponential term. Cardiac bypass time, length of ICU stay and serum creatinine were first standardized by subtracting the sample mean and dividing by the sample standard deviation, then included using an exponential term. Finally, we examined the hypothesized association between the genotype variables ($UGT1A4$ and $UGT2B10$ variants and $CYP2A6$ score) and total clearance. For $UGT1A4$ and $UGT2B10$, dichotomous models (coding individuals as having a loss-of-function variant or not) and additive models (counting the number of variants) were considered. For $CYP2A6$, the predicted enzyme activity score was standardized and included as an exponential term [132].

5.2.1.2 Prior Specification

We developed an informative prior using estimates from eight recently published pediatric dexmedetomidine PK studies (Table 5.8). The goal of the informative prior was to provide plausible parameter values based on previous literature in similar populations to supplement the current data and allow more precise estimation. We compared our prior to one developed by Wiczling et al. (2016) for an informative Bayesian analysis of dexmedetomidine PK in 38 critically ill pediatric patients using intensive sampling[3] [125]. To begin, we formed a weighted average of estimates from the previous studies for each of the main PK population parameters, total clearance (CL_{pop}), inter-compartmental clearance (Q_{pop}), central compartment volume of distribution ($V1_{pop}$) and peripheral compartment volume of distribution ($V2_{pop}$). The priors were specified using the same log-Normal distributional family as Wiczling et al., but with the most probable value (mode) equal to the weighted average; the variance was adjusted so the 90% highest density interval (HDI) included

[3]Eight samples over 24 hours during infusion and eight samples over six hours at the cessation of infusion.

all the previous study estimates. More informative (lower variance) priors were given to parameters that have been historically difficult to estimate with sparse data such as $V2_{pop}$. For IIV standard deviation parameters (ω_{CL}, ω_{Q}, ω_{V1}, and ω_{V2}) we used the same distributional form as Wiczling et al., but with 90% HDI widths between 5 to 10 times larger to account for additional uncertainty. For the residual (intra-individual) standard deviation parameters, we used half-Cauchy priors with mode at 0. Model parameters for non-genotype and genotype covariate effects on total clearance were given generic weakly informative skeptical standard normal priors, with continuous covariates standardized so a one-unit change in the standardized variable is equal to a standard deviation change on the original scale.

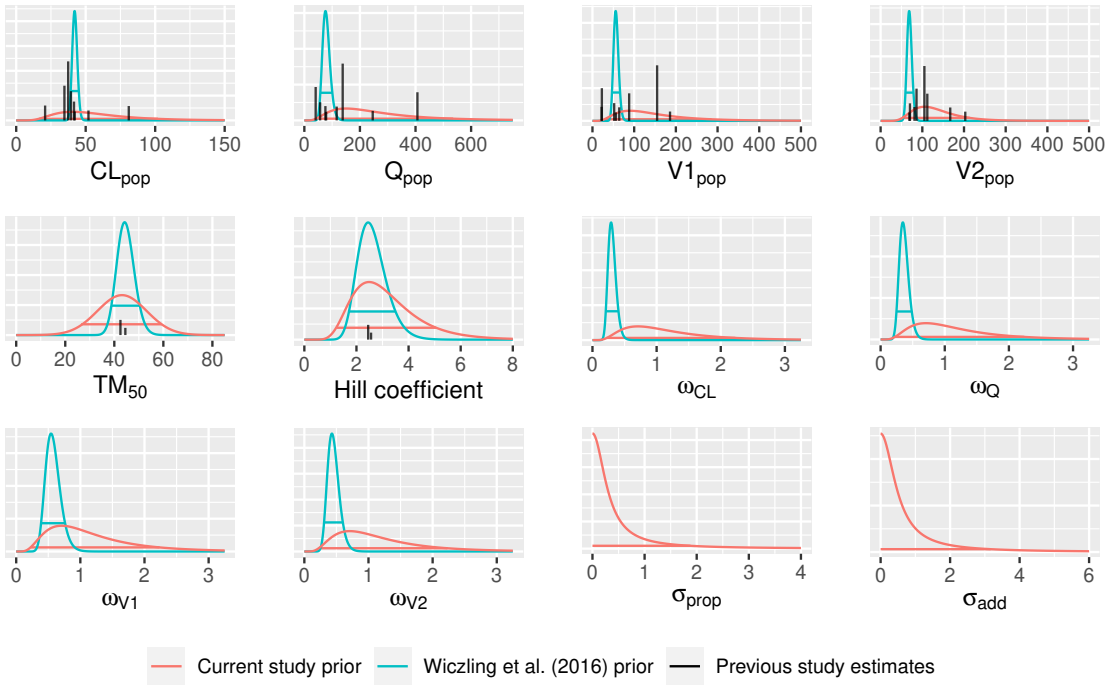

Figure 5.1: Prior distributions for two-compartment pharmacokinetic model. Red curves are current study prior densities, teal curves are Wiczling et al. prior densities, and vertical black lines indicate the values of previous dexmedotomidine study estimates with the line height scaled to the sample size for each study. Horizontal lines indicate the 90% highest density interval for each prior.

Figure 5.1 shows the current study priors compared to those used in the informative Bayesian analysis of Wiczling et al. (see Table 5.9 for specific mathematical forms). The priors used in the current study are much less informative than in the previous Bayesian analysis which had a smaller sample size but more regular, intensive sampling. It is convenient to specify the joint prior using marginal components as displayed in Figure 5.1; however, it is easier to compare the overall joint priors using prior predictive concentrations. Figure 5.2 shows the median and 80% credible intervals for prior predictive distributions of concentration over time using the two-compartment allometric weight scaling model with no other covariates. The plots use a hypothetical

reference subject with the median weight of 9.4 kg and a loading dose of 0.4 mcg/kg delivered over 5 minutes followed by a 6 hour infusion at 0.6 mcg/kg/h and a 2 hour infusion at 0.5 mcg/kg/h. Because they are based solely on the specified PK model structure and prior information, these plots represent the assumed state of knowledge before incorporating any observed data. The parameters for the current study prior contain more uncertainty (larger variance), so the range of feasible concentration profiles is much wider.

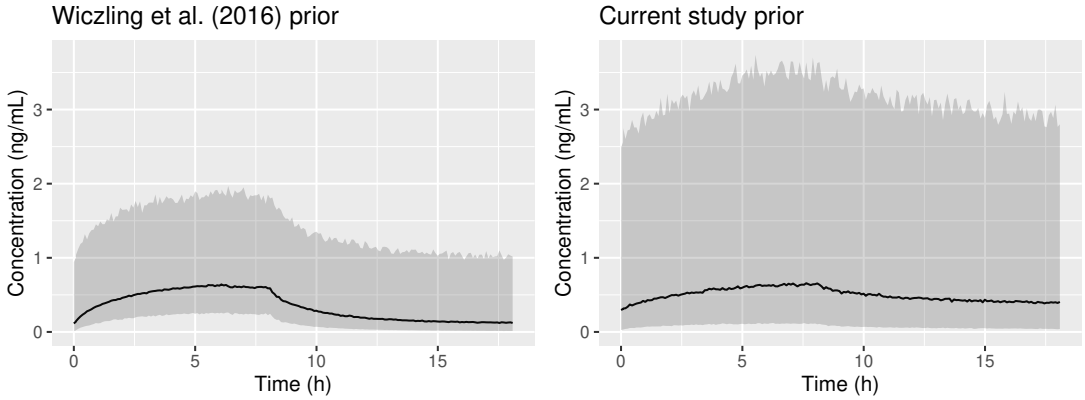

Figure 5.2: Prior predictive distributions for two-compartment allometric scaling model for reference subject with median weight of 9.4 kg. Median (solid black line) and 80% credible intervals (shaded region) are shown.

In addition to the informative prior, we also consider sensitivity analyses using an improper non-informative prior and a more weakly informative prior. The simulation study in Chapter 3 showed poor performance can result when using non-informative priors and inclusion of prior information is the primary motivation for performing the Bayesian reanalysis, however it may be of interest to compare a Bayesian model without informative priors to the frequentist analysis; to "let the data speak" without any external information. However, as noted by Wakefield et al. (1999) [62], "for those elements ... that correspond to nonlinear parameters at the first stage, proper priors are required in order to guarantee propriety of the posterior distribution." Given the above and the limitations of the observational data in this study, a non-informative prior may not produce reasonable results. Therefore, we also include a sensitivity analysis using less informative proper priors which are in the same distributional families as the main analysis prior, but with the log-scale standard deviation or scale increased by 50%[4] (see Table 5.10 for specific mathematical forms).

[4]e.g., for the CL_{pop} parameter the main analysis prior is lognormal; the logarithm of the prior distribution has a mode of 40.6 and standard deviation of 0.5 and the 90% highest density interval (HDI) is 16.2 - 101.9. The more weakly informative prior has the same mode of 40.6, but standard deviation (on the log-scale) of $0.75 = 0.5 \times 1.5$ and the 90% HDI is 9.4 - 189.6. For comparison, the minimum and maximum estimates of CL_{pop} among the recent studies contributing to the main analysis prior were 20.8 and 81, respectively, so the more weakly informative prior includes plausible values far beyond the range of the observed estimates.

5.3 Results

5.3.1 Study Population and Specimens

As detailed in Chapter 4, the final cohort contained 354 subjects with 2,386 dexmedetomidine dosing events (2,351 intravenous infusions and 35 bolus administrations) and 1,400 specimens with dexmedetomidine concentration measurements. The *CYP2A6* predicted activity score was available for 350 of the 354 subjects and all models including *CYP2A6* score were estimated using this subset of subjects.

5.3.2 Bayesian Population Pharmacokinetic Model

Table 5.1 shows the estimated out-of-sample predictive accuracy of the initial allometric scaling, age maturation, and non-genotype covariate models. Both the Hill maturation model and exponential age model had higher $lppd_{cv}$ than the model with only allometric scaling confirming the importance of including both weight and age covariates for this population. In addition, the $lppd_{cv}$ for the model including a Hill maturation factor for total clearance indicated better predictive performance than the exponential age model. Moving to the next stage, there were no improvements in predictive accuracy for models which added non-genotype covariate effects on total clearance to the model with allometric scaling and Hill maturation. Further, in plots of the total clearance random effects versus covariates no strong trends were seen for the non-genotype covariates (Figure 5.9), indicating little association between these variables and total clearance after adjusting for weight and age. Additional diagnostic plots are shown in Figures 5.10, 5.11, and 5.14.

Table 5.1: Log Pointwise Predictive Density for Allometric Scaling, Age Maturation, and non-Genotype Covariate Models

Model	$lppd_{cv}$
Two-compartment allometric scaling	706.1
+ Hill maturation	856.6
+ exponential postmenstrual age	813.3
Two-compartment allometric scaling + Hill maturation	856.6
+ serum creatinine	735.3
+ sex	805.2
+ STAT score	719.3
+ cardiac bypass time	834.8
+ length of ICU stay	775.3

The final covariate model before including genotype effects was:

$$
\begin{aligned}
CL_i &= CL_{pop} \times \left(\frac{WT_i}{70}\right)^{0.75} \times \frac{1}{1+\left(\frac{TM_{50}}{PMA_i}\right)^{Hill}} \times \exp(\eta_{CL,i}) \\
V1_i &= V1_{pop} \times \left(\frac{WT_i}{70}\right) \times \exp(\eta_{V1,i}) \\
Q_i &= Q_{pop} \times \left(\frac{WT_i}{70}\right)^{0.75} \times \exp(\eta_{Q,i}) \\
V2_i &= V2_{pop} \times \left(\frac{WT_i}{70}\right) \times \exp(\eta_{V2,i})
\end{aligned}
$$

where CL_i, $V1_i$, Q_i, and $V2_i$ are the individual-specific PK parameters corresponding to CL, $V1$, Q, and $V2$; CL_{pop}, $V1_{pop}$, Q_{pop}, and $V2_{pop}$ are population parameters; WT_i is subject weight in kilograms (kg); and PMA_i is subject postmenstrual age in weeks. TM_{50} is a parameter estimating the postmenstrual age when 50% of adult clearance is achieved and the $Hill$ parameter controls the shape of the sigmoid function. The $\eta_{CL,i}$, $\eta_{V1,i}$, $\eta_{Q,i}$, and $\eta_{V2,i}$ are random effects explaining IIV for the PK parameters which follow a normal distribution with mean zero and variance of ω^2_{CL}, ω^2_{V1}, ω^2_{Q}, and ω^2_{V2}, respectively. Table 5.2 summarizes the HMC-estimated posterior parameters for the model with allometric scaling and Hill maturation with estimates given in terms of a standard subject weight of 70kg: $CL_{pop} = 39.9$ L/h, $V1_{pop} = 141$ L, $Q_{pop} = 14.4$ L/h, and $V2_{pop} = 1,270$ L. The posterior predictive distribution from this model for a reference subject with median weight of 9.4 kg, postmenstrual age of 104.6 weeks, and the same dosing used for the prior predictive plots is shown in Figure 5.3. Updating the prior information with the observed data and including postmenstrual age in the structural model resulted in a substantial reduction in the 80% credible interval width.

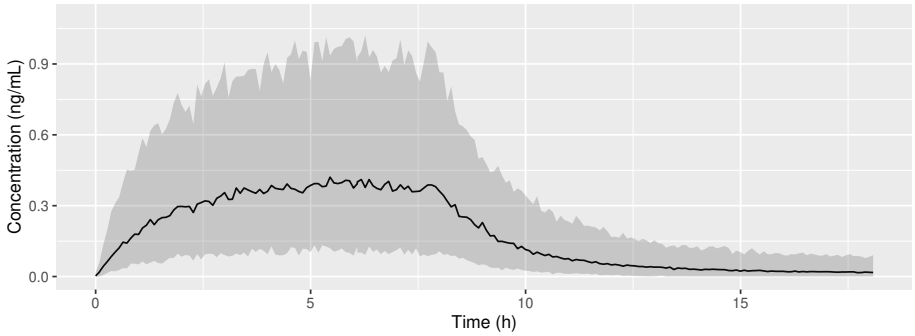

Figure 5.3: Posterior predictive distribution for model with allometric scaling and Hill maturation for reference subject with median weight of 9.4 kg and postmenstrual age of 104.6 weeks. Median (solid black line) and 80% credible intervals (shaded region) are shown.

Table 5.2: Summary of Posterior Distribution Estimated by HMC for Parameters in Allometric Scaling and Hill Maturation Model

			90% Highest Density Interval	
	Mean	Median	Lower Bound	Upper Bound
CL_{pop}	39.9	39.8	35	44.9
Q_{pop}	14.4	14.2	10.3	18.8
$V1_{pop}$	141	140	121	161
$V2_{pop}$	1270	1214	690	1786
TM_{50}	40.5	40.9	35.3	46.8
$Hill$	4.27	4.11	2.28	6.36
ω^2_{CL}	0.665	0.658	0.537	0.824
ω^2_{Q}	1.02	1.01	0.573	1.48
ω^2_{V1}	0.788	0.775	0.511	1.08
ω^2_{V2}	4.07	3.9	2.4	5.89
σ^2_{add}	0.00021	0.000206	0.000153	0.000269
σ^2_{prop}	0.185	0.184	0.161	0.209

5.3.3 Genetic Effects on Clearance and Concentration

Starting with the model including allometric scaling and Hill maturation, we added genotype effects individually to the total clearance model. The estimated predictive accuracy measured by $lppd_{cv}$ for the $UGT1A4$ and $UGT2B10$ models is shown in Table 5.3. Adding the UGT variants with an additive or dichotomous functional form to the model did not improve estimated predictive accuracy.

Table 5.3: Log Pointwise Predictive Density for $UGT1A4$ and $UGT2B10$ Models

Model	$lppd_{cv}$
Two-compartment allometric scaling + Hill maturation	856.6
+ $UGT1A4$ additive	772.9
+ $UGT2B10$ additive	783.8
+ $UGT1A4$ dichotomous	746.3
+ $UGT2B10$ dichotomous	809.8

Table 5.4 compares the predictive accuracy of the allometric scaling and Hill maturation model to the model that also includes $CYP2A6$ risk score among the 350 subjects with an available score. In contrast to the UGT models, adding a $CYP2A6$ score effect for the total clearance parameter improved the model with $lppd_{cv}$ increasing from 780.3 to 803.3. The plots of observed vs. population predicted concentration and observed vs. individual predicted concentration in Figure 5.4 show a reasonably good fit. Additional diagnostic plots are shown in Figures 5.12, 5.13, and 5.15.

Table 5.4: Log Pointwise Predictive Density for $CYP2A6$ score Model*

Model	$lppd_{cv}$
Two-compartment allometric scaling + Hill maturation	780.3
+ $CYP2A6$ score	803.3

* Among n = 350 subjects with available score

For the model with $CYP2A6$ score, total clearance is given by:

$$CL_i = CL_{pop} \times \left(\frac{WT_i}{70}\right)^{0.75} \times \frac{1}{1+\left(\frac{TM_{50}}{PMA_i}\right)^{Hill}} \times \exp(\beta_{CL,CYP2A6} \cdot \text{standardized } CYP2A6 \text{ score}) \times \exp(\eta_{CL,i})$$

The posterior distribution of parameters for this model are summarized in Table 5.5. The mean effect of the standardized $CYP2A6$ score was 0.0627 with 90% probability that the effect is between -3.51×10^{-4} and 0.123 and only 4.8% probability that $\beta_{CL,CYP2A6} \leq 0$. A one standard deviation increase in $CYP2A6$ score was associated with approximately $\exp(0.0627) \approx 1.06$ times higher total clearance on average. Because weight is

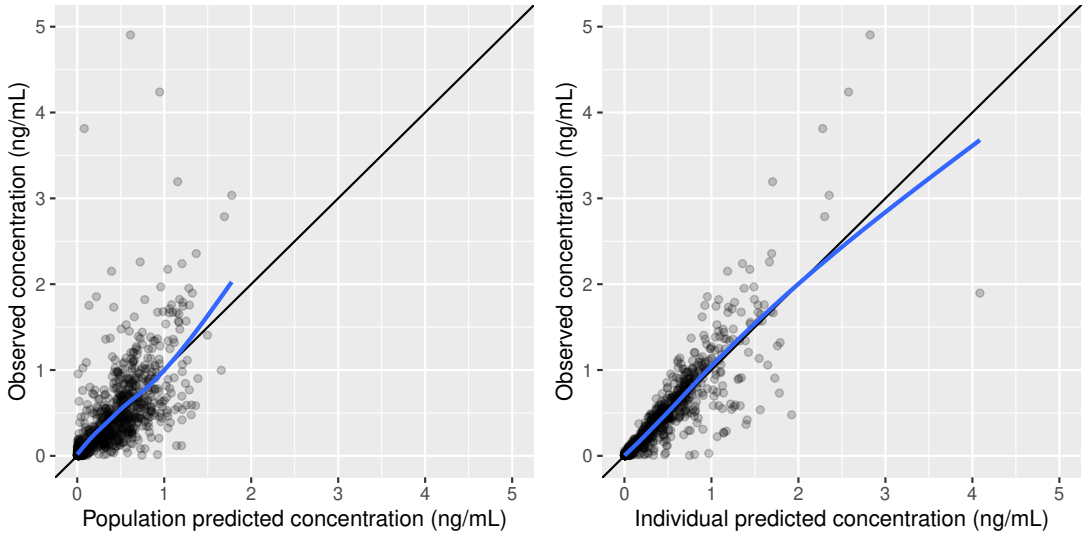

Figure 5.4: Observed vs. predicted posterior median concentration for model with allometric scaling, Hill maturation, and *CYP2A6* score. Blue lines are loess smoothers.

the most important factor in determining total clearance for a pediatric population, we can also explain the *CYP2A6* effect in terms of the weight change needed to achieve an equivalent increase in clearance. Increasing *CYP2A6* score by one standard deviation (0.14) increases clearance by the same amount as a weight increase of about 9% [$\exp\left(\frac{0.0627}{0.75}\right) \approx 1.09$] for a subject without the standard deviation increase in *CYP2A6* score. A detailed calculation is provided in the Appendix.

The model including *CYP2A6* score implies that a subject with a high *CYP2A6* score will require a larger dose to achieve the same concentration as a subject with lower *CYP2A6* score, holding other covariates constant. Figure 5.5 illustrates this for hypothetical subjects with median weight (9.4kg), median postmenstrual age (104.6 weeks) and *CYP2A6* scores of 2.0 and 2.43, around the median and maximum, respectively. In Figure 5.5(a) dosing was the same for both subjects (loading dose of 0.4 mcg/kg delivered over 5 minutes followed by a 6 hour infusion at 0.6 mcg/kg/h and a 2 hour infusion at 0.5 mcg/kg/h) and the peak concentration for the subject with *CYP2A6* score of 2.43 is around 0.1 ng/mL lower than the equivalent subject with *CYP2A6* score of 2.0. In Figure 5.5(b), the subject with *CYP2A6* score of 2.43 was given a 0.1 mcg/kg higher dose (loading dose of 0.5 mcg/kg over 5 minutes followed by a 6 hour infusion at 0.7 mcg/kg/h and then a 2 hour infusion at 0.6 mcg/kg/h). The higher dose results in a concentration profile nearly identical to the subject with *CYP2A6* score of 2.0.

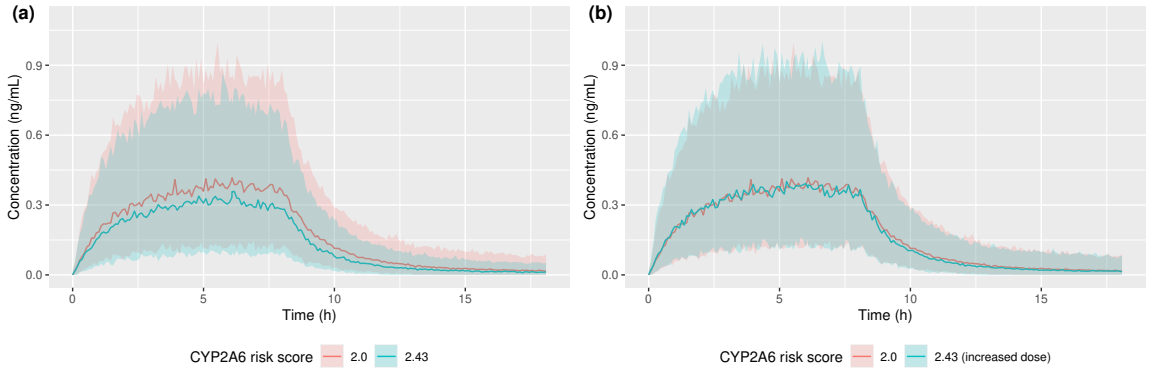

Figure 5.5: Posterior predictive distribution for model with allometric scaling, Hill maturation and $CYP2A6$ score for (a) reference subjects with median weight of 9.4 kg, postmenstrual age of 104.6 weeks, loading dose of 0.4 mcg/kg delivered over 5 minutes followed by a 6 hour infusion at 0.6 mcg/kg/h and a 2 hour infusion at 0.5 mcg/kg/h, and $CYP2A6$ score of 2.0 and 2.43 and (b) same covariates as (a) but with dose for subject with $CYP2A6$ score of 2.43 increased by 0.1 mcg/kg (loading dose of 0.5 mcg/kg over 5 minutes followed by a 6 hour infusion at 0.7 mcg/kg/h and then a 2 hour infusion at 0.6 mcg/kg/h). Medians (solid lines) and 80% credible intervals (shaded regions) are shown.

Table 5.5: Summary of Posterior Distribution Estimated by HMC for Parameters in Allometric Scaling, Hill Maturation, and $CYP2A6$ Score Model

			90% Highest Density Interval	
	Mean	Median	Lower Bound	Upper Bound
CL_{pop}	40.9	40.8	36.1	45.6
$\beta_{CL,CYP2A6}$	0.0627	0.0634	-0.000351	0.123
Q_{pop}	13.7	13.5	9.79	17.6
$V1_{pop}$	142	142	122	164
$V2_{pop}$	1170	1136	636	1661
TM_{50}	41	41.3	35.7	46.2
$Hill$	4.34	4.31	2.65	6.06
ω^2_{CL}	0.628	0.62	0.494	0.761
ω^2_{Q}	1.03	1.01	0.647	1.41
ω^2_{V1}	0.842	0.827	0.52	1.11
ω^2_{V2}	3.86	3.75	2.06	5.58
σ^2_{add}	0.000215	0.00021	0.000153	0.000267
σ^2_{prop}	0.188	0.187	0.163	0.208

5.3.4 Prior Sensitivity Analysis

For the sensitivity analysis, the model including allometric scaling and Hill maturation and the model with those effects plus $CYP2A6$ score on total clearance were both fit using a non-informative prior and the more weakly informative prior described above. The HMC sampler settings were modified to include 500 additional warm-up samples since models with less prior information may be slower to converge.

The results for the non-informative prior were very poor with both models failing to reach convergence after 3000 warm-up samples. Examining the trace plots for these models in Figures 5.16 and 5.17 reveals serious problems with the sampler with multiple chains 'stuck' around the same value and unrealistically large values for several parameters, such as CL_{pop} and TM_{50}, despite the chains being initialized to values around the modes of the main analysis informative prior.

The more weakly informative prior produced better results than the non-informative prior. The trace plots in Figures 5.18 and 5.19 show much better mixing and convergence of the chains. Table 5.6 shows the results from the model with only allometric scaling and Hill maturation using the weakly informative prior. There are some differences compared to the main analysis prior, but the posterior distributions are close for most parameters. The mean estimates using the more weakly informative prior are: $CL_{pop} = 34.8$, $V1_{pop} = 149$, $Q_{pop} = 18.8$, $V2_{pop} = 3328$, $TM_{50} = 41.5$, and $Hill = 5.04$ compared to $CL_{pop} = 39.9$, $V1_{pop} = 141$, $Q_{pop} = 14.4$, and $V2_{pop} = 1270$, $TM_{50} = 40.5$, and $Hill = 4.27$ for the main analysis prior. The largest difference is for the $V2_{pop}$ parameter; the observed data does not contain much information to estimate this parameter precisely, as evidenced by the frequentist analysis, and the prior has a larger influence on the posterior.

Results for the model including allometric scaling, Hill maturation, and $CYP2A6$ score estimated with the more weakly informative prior are shown in Table 5.7. The mean estimates for the sensitivity analysis are: $CL_{pop} = 35.8$, $\beta_{CL,CYP2A6} = 0.0843$, $V1_{pop} = 146$, $Q_{pop} = 17.9$, $V2_{pop} = 2841$, $TM_{50} = 41$, and $Hill = 5.23$ vs. $CL_{pop} = 40.9$, $\beta_{CL,CYP2A6} = 0.0627$, $V1_{pop} = 142$, $Q_{pop} = 13.7$, $V2_{pop} = 1170$, $TM_{50} = 41$, and $Hill = 4.34$ for the main analysis prior. Again, the largest differences are seen for parameters Q_{pop} and $V2_{pop}$. Figure 5.6 shows the posterior predictive distributions used to illustrate the effects of dose change using the more weakly informative prior. Comparing to the same plots from the main analysis (Figure 5.5) shows that using the weaker prior has little impact on predicted concentrations and inference, which are nearly identical.

5.3.5 Frequentist Analysis Comparison

Figures 5.7 and 5.8 show the prior and posterior for the Bayesian reanalysis along with the frequentist mean estimate and 95% confidence intervals for the model including allometric scaling and Hill maturation and the model including allometric scaling, Hill maturation and $CYP2A6$ score. Note that the frequentist

Table 5.6: Summary of Posterior Distribution Estimated by HMC for Parameters in Allometric Scaling and Hill Maturation Model for Sensitivity Analysis

			90% Highest Density Interval	
	Mean	Median	Lower Bound	Upper Bound
CL_{pop}	34.8	35	27.9	40
Q_{pop}	18.8	18.2	12.4	25.2
$V1_{pop}$	149	149	128	168
$V2_{pop}$	3328	3067	1418	5480
TM_{50}	41.5	41.9	34.5	48.6
$Hill$	5.04	4.71	1.85	7.7
ω^2_{CL}	0.708	0.69	0.517	0.872
ω^2_{Q}	0.764	0.753	0.291	1.22
ω^2_{V1}	0.72	0.703	0.459	0.951
ω^2_{V2}	4.86	4.61	2.66	7.1
σ^2_{add}	0.000191	0.000189	0.000141	0.000246
σ^2_{prop}	0.188	0.188	0.163	0.212

Table 5.7: Summary of Posterior Distribution Estimated by HMC for Parameters in Allometric Scaling, Hill Maturation, and $CYP2A6$ Score Model for Sensitivity Analysis

			90% Highest Density Interval	
	Mean	Median	Lower Bound	Upper Bound
CL_{pop}	35.8	35.9	29.3	42.6
$\beta_{CL,CYP2A6}$	0.0843	0.083	0.0125	0.158
Q_{pop}	17.9	17	12.1	25.3
$V1_{pop}$	146	146	126	167
$V2_{pop}$	2841	2582	1280	4490
TM_{50}	41	41.6	35.7	47.1
$Hill$	5.23	4.91	2.32	7.86
ω^2_{CL}	0.674	0.66	0.528	0.853
ω^2_{Q}	0.827	0.818	0.319	1.32
ω^2_{V1}	0.73	0.709	0.438	0.989
ω^2_{V2}	4.84	4.56	2.42	7.44
σ^2_{add}	0.000197	0.000194	0.000145	0.000253
σ^2_{prop}	0.188	0.188	0.164	0.211

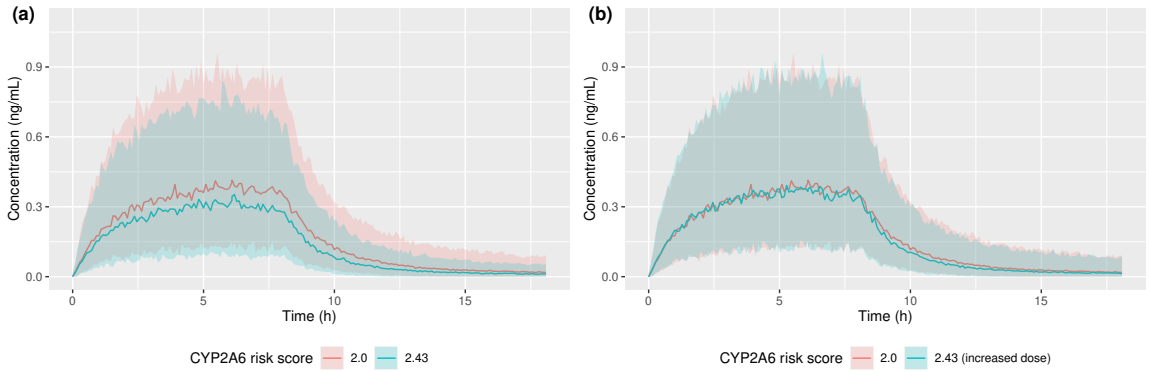

Figure 5.6: Posterior predictive distribution for model with allometric scaling, Hill maturation and *CYP2A6* score for sensitivity analysis using more weakly informative priors for (a) reference subjects with median weight of 9.4 kg, postmenstrual age of 104.6 weeks, loading dose of 0.4 mcg/kg delivered over 5 minutes followed by a 6 hour infusion at 0.6 mcg/kg/h and a 2 hour infusion at 0.5 mcg/kg/h, and *CYP2A6* score of 2.0 and 2.43 and (b) same covariates as (a) but with dose for subject with *CYP2A6* score of 2.43 increased by 0.1 mcg/kg (loading dose of 0.5 mcg/kg over 5 minutes followed by a 6 hour infusion at 0.7 mcg/kg/h and then a 2 hour infusion at 0.6 mcg/kg/h). Medians (solid lines) and 80% credible intervals (shaded regions) are shown.

and Bayesian specifications for the subject level (Stage one) and population level (Stage two) models are not identical; the frequentist models used a time-varying weight covariate and non-standardized *CYP2A6* score, and had a simplified variance structure with most of the correlations between random effects (ρ) and the additive component of the residual variance (σ_{add}) fixed at zero; the Bayesian models used baseline weight, standardized *CYP2A6*, and include both additive and proportional residual error and all random effect correlations.

In the final Bayesian model with weight, age maturation, and *CYP2A6* score we estimated a weight-standardized mean CL_{pop} of 40.9 L/h (CV 93.5%) while the frequentist estimate for the most similar model was 24.8 L/h (CV 100%); the Bayesian mean Q_{pop} estimate was 13.7 L/h (CV 134.2%) vs. 24.5 L/h (CV 86%) for the frequentist model; $V1_{pop}$ was estimated as 142 L (CV 114.9%) vs. 152 L (CV 138%) for the Bayesian and frequentist models, respectively; similarly, the Bayesian vs. frequentist estimates for TM_{50} are 41 vs. 42.6 weeks; for the $Hill$ coefficient 4.34 vs. 7.45 and for $\beta_{CL,CYP2A6}$ 0.0627 vs. 0.0885. The largest difference was seen for $V2_{pop}$ with a mean of 1170 L (CV 681.7%) for the Bayesian model and 5756 L (CV 549%) for the frequentist model. For most parameters, the confidence intervals from the frequentist analysis are much narrower than the corresponding Bayesian credible intervals. Further, the final model from the frequentist analysis included only allometric scaling and age maturation and pharmacogenetic effects were not statistically or clinically significant.

There are several potential explanations for the different conclusions reached using the two statistical paradigms. First, the Bayesian models include both current study data and external evidence (encoded in the prior

distribution). Conceptually this is equivalent including data from additional subjects in the analysis which stabilizes the estimation. In addition, the frequentist and Bayesian analyses used different information criteria. The frequentist model comparison used the somewhat misleadingly named 'corrected Bayesian information criteria' (BICc) [143] while the Bayesian model comparison used cross-validation log pointwise predictive density ($lppd_{cv}$). While the goal of $lppd_{cv}$ is to estimate out-of-sample predictive fit, BICc attempts to approximate the marginal probability density (or evidence). In practice, BICc favors more parsimonious models compared to other common criteria such as AIC and WAIC. Further, as stated by Gelman et al. (2014) [4], "it is completely possible for a complicated model to predict well and have a low AIC, DIC, and WAIC, but, because of the penalty function, to have a relatively high (that is, poor) BIC." For these two analyses the tradeoff between parsimony (lower variance) and predictive accuracy (lower bias) excludes the *CYP2A6* effect when selecting for parsimony (using BICc) in the frequentist model with less information, but includes the effect when selecting for predictive accuracy (using $lppd_{cv}$) in the Bayesian model adding prior information from more 'subjects.'

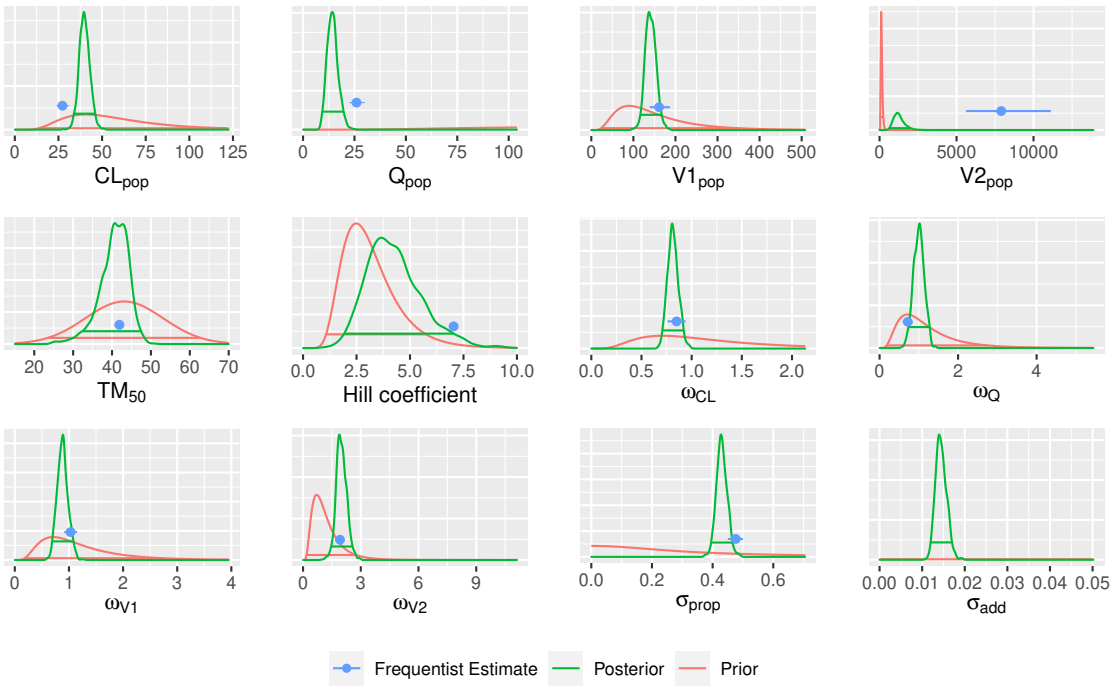

Figure 5.7: Prior and posterior distributions vs. frequentist estimates from James et al. (2022) for model with allometric scaling and Hill maturation. Frequentist maximum likelihood estimates and 95% confidence intervals (blue circle and interval) and Bayesian prior (red curve) and posterior (green curve) distributions are shown. Horizontal lines indicate the 90% highest density intervals for the Bayesian distributions.

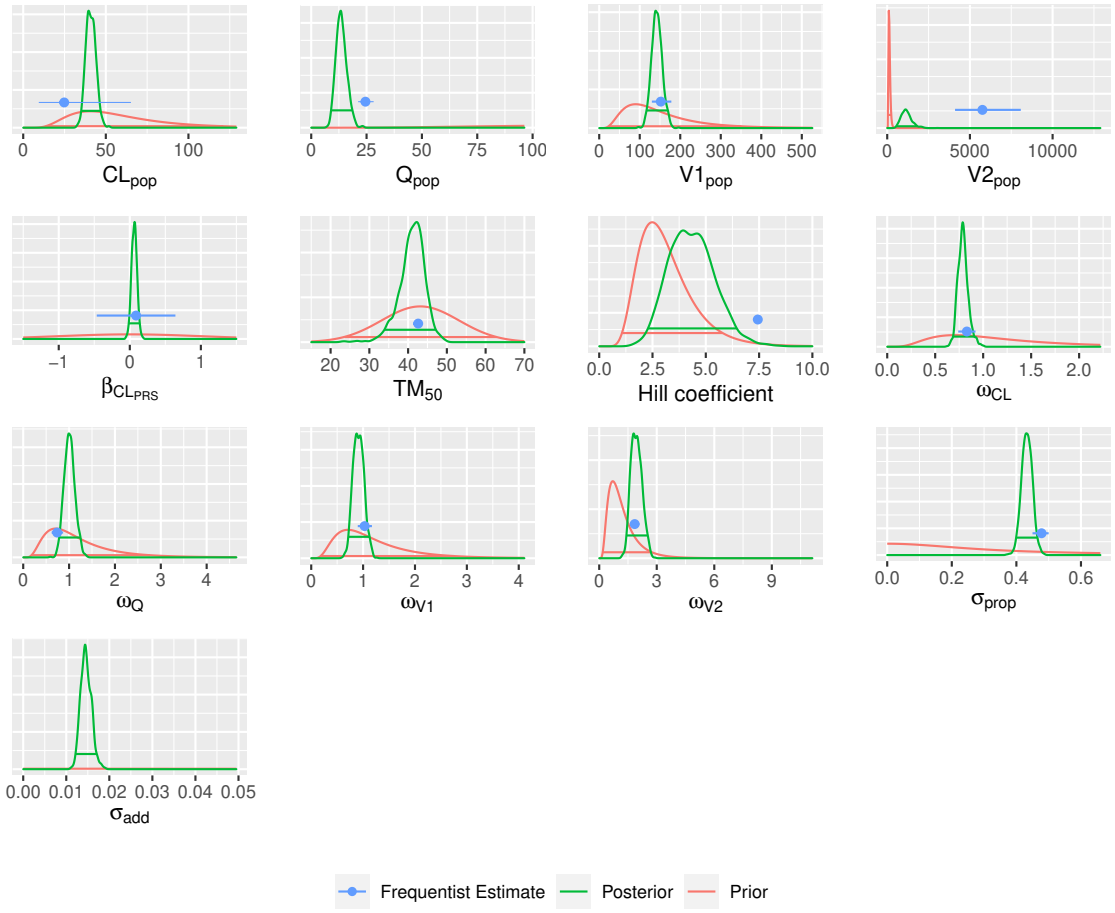

Figure 5.8: Prior and posterior distributions vs. frequentist estimates from James et al. (2022) for model with allometric scaling, Hill maturation and *CYP2A6* score. Frequentist maximum likelihood estimates and 95% confidence intervals (blue circle and interval) and Bayesian prior (red curve) and posterior (green curve) distributions are shown. Horizontal lines indicate the 90% highest density intervals for the Bayesian distributions.

5.4 Discussion

Using remnant specimens along with dosing, clinical, and demographic information from an EHR system we developed a dexmedetomidine population PK model for a large pediatric cohort of 354 patients. Incorporating prior information using the Bayesian paradigm stabilized the estimation and produced results much closer to the previous literature for most parameters compared to the frequentist analysis in James et al. (2022) [110]. In addition, we found that including predicted *CYP*2*A*6 enzyme activity using a novel polygenic risk score improved out-of-sample predictive accuracy and can potentially make a clinically meaningful difference in dosing.

Using Bayesian analysis for population PK models involves several tradeoffs. Bayesian estimation has a higher computational burden than frequentist estimation. For this analysis, the ADVI approximation was used to reduce computing time during model development, but fitting the models still took substantially longer (e.g., the final model took around 57 minutes to fit using frequentist SAEM estimation, 6 minutes to fit using ADVI and 595 minutes, or nearly 10 hours, to fit using HMC). Also, there may be disagreement about the appropriateness or subjectivity of the prior. Although Bayesian estimation failed with non-informative priors, comparing two different informative priors with plausible parameter values produced roughly similar Bayesian posterior distributions for most parameters and did not have a major impact on predicted concentrations. Moreover, the informative priors used for this analysis were derived from previous literature and encoded much greater uncertainty than those used for the previous informative Bayesian analysis of Wiczling et al. An alternative approach such as formal expert elicitation could be used to arrive at a consensus prior.

Using population PK models with EHR data and remnant specimens allows collection of much larger sample sizes while supplementing this observational data with prior information from previous studies helps address limitations related to irregular, sparse sampling. By utilizing both the current data and external evidence, we can better understand what factors affect dexmedetomidine PK to enable improved individualized dosing decisions.

5.5 Appendix

5.5.1 Model Parameterization

The variance-covariance matrix for the parameter random effects, Ω, is parameterized using a decomposition into standard deviations and a correlation matrix [109]; $\Omega = \begin{bmatrix} \omega_{Cl} & 0 \\ 0 & \omega_V \end{bmatrix} \times \rho \times \begin{bmatrix} \omega_{Cl} & 0 \\ 0 & \omega_V \end{bmatrix}$. This decomposition allows prior information about the variability of population PK parameters to be specified on the same scale as the parameter. In addition, a non-centered parameterization is used for efficiency and numerical stability [102]. In the `Stan` model implementation, a prior is specified on the Cholesky factor of the correlation matrix L_ρ where $\rho = L_\rho \times L'_\rho$ instead of directly on ρ. Note that for the LKJ prior $\eta = 1$ is a uniform distribution over correlation matrices.

Table 5.8: Prior Pediatric Dexmedetomidine Studies

Title	Authors	Year	N	CL_{pop}	Q_{pop}	$V1_{pop}$	$V2_{pop}$	TM_{50}	Hill
Pharmacokinetics of Dexmedetomidine in Infants and Children After Orthotopic Liver Transplantation	Damian, M. et al.	2020	20	52	246	186	203	-	-
Population Pharmacokinetic Model of Dexmedetomidine in a Heterogeneous Group of Patients	Ber, J. et al.	2020	70	34.7	40.8	22.5	86.1	-	-
Results of a phase I multicentre investigation of dexmedetomidine bolus and infusion in corrective infant cardiac surgery	Zuppa, A. et al.	2019	119	37.4	138	155	105	-	-
Dexmedetomidine Pharmacokinetics and a New Dosing Paradigm in Infants Supported With Cardiopulmonary Bypass	Zimmerman, K. O. et al.	2019	18	42.1	78.3	56.3	69	44.5	2.56
A Population Pharmacokinetic Model of Intravenous Dexmedetomidine for Mechanically Ventilated Children after Neurosurgery	Song, I. et al.	2019	29	81	116.4	64.2	167	-	-
Population Pharmacokinetics and Pharmacodynamics of Dexmedetomidine in Children Undergoing Ambulatory Surgery	Pérez-Guillé, M.-G. et al.	2018	30	20.8	75.8	21.9	81.2	-	-
Dexmedetomidine Pharmacology in Neonates and Infants After Open Heart Surgery	Su, F. et al.	2016	59	39.4	406.8	88	112	-	-
The pharmacokinetics of dexmedetomidine during long-term infusion in critically ill pediatric patients. A Bayesian approach with informative priors	Wiczling, P. et al.	2016	38	41.6	56.8	52	70.4	42.5	2.45

Table 5.9: Dexmedetomidine Informative Priors

Parameter	Derived from Wiczling et al. (2016)	Current study informative prior
CL_{pop}	$logNormal(\mu = \log(42.1), \sigma = \sqrt{0.00185})$	$logNormal(\mu = \log(40.6) + 0.5^2, \sigma = 0.5)$
Q_{pop}	$logNormal(\mu = \log(78.3), \sigma = \sqrt{0.02862})$	$logNormal(\mu = \log(149.9) + 0.65^2, \sigma = 0.65)$
$V1_{pop}$	$logNormal(\mu = \log(56.3), \sigma = \sqrt{0.01122})$	$logNormal(\mu = \log(89.9) + 0.6^2, \sigma = 0.6)$
$V2_{pop}$	$logNormal(\mu = \log(69), \sigma = \sqrt{0.00726})$	$logNormal(\mu = \log(105.5) + 0.4^2, \sigma = 0.4)$
TM_{50}	$logNormal(\mu = \log(44.5), \sigma = \sqrt{0.00636})$	$Normal(\mu = 43.1, \sigma = 10)$
$Hill$	$logNormal(\mu = \log(2.56), \sigma = \sqrt{0.04472})$	$logNormal(\mu = \log(2.48) + 0.4^2, \sigma = 0.4)$
ω_{Cl}	$logNormal(\mu = \log(\sqrt{0.091}), \sigma = 0.2)$	$logNormal(\mu = \log(\sqrt{0.5}) + 0.6^2, \sigma = 0.6)$
ω_{Q}	$logNormal(\mu = \log(\sqrt{0.13}), \sigma = 0.2)$	$logNormal(\mu = \log(\sqrt{0.5}) + 0.6^2, \sigma = 0.6)$
ω_{V1}	$logNormal(\mu = \log(\sqrt{0.32}), \sigma = 0.2)$	$logNormal(\mu = \log(\sqrt{0.5}) + 0.6^2, \sigma = 0.6)$
ω_{V2}	$logNormal(\mu = \log(\sqrt{0.20}), \sigma = 0.2)$	$logNormal(\mu = \log(\sqrt{0.5}) + 0.6^2, \sigma = 0.6)$
ρ	$LKJ(\eta = 1)$	$LKJ(\eta = 1)$
σ_{prop}	$Half-Cauchy(\mu = 0, \sigma = 0.2)$	$Half-Cauchy(\mu = 0, \sigma = 0.3)$
σ_{add}	$Half-Cauchy(\mu = 0, \sigma = 0.2)$	$Half-Cauchy(\mu = 0, \sigma = 0.5)$

Table 5.10: Dexmedetomidine Sensitivity Analysis Priors

Parameter	More weakly informative prior	Non-informative prior
CL_{pop}	$logNormal(\mu = \log(40.6) + 0.75^2, \sigma = 0.75)$	$Uniform(0, \infty)$
Q_{pop}	$logNormal(\mu = \log(149.9) + 0.975^2, \sigma = 0.975)$	$Uniform(0, \infty)$
$V1_{pop}$	$logNormal(\mu = \log(89.9) + 0.9^2, \sigma = 0.9)$	$Uniform(0, \infty)$
$V2_{pop}$	$logNormal(\mu = \log(105.5) + 0.6^2, \sigma = 0.6)$	$Uniform(0, \infty)$
TM_{50}	$trunc-Normal(min = 0, \mu = 43.1, \sigma = 15)$	$Uniform(0, \infty)$
$Hill$	$logNormal(\mu = \log(2.48) + 0.6^2, \sigma = 0.6)$	$Uniform(0, \infty)$
ω_{Cl}	$logNormal(\mu = \log(\sqrt{0.5}) + 0.9^2, \sigma = 0.9)$	$Uniform(0, \infty)$
ω_{Q}	$logNormal(\mu = \log(\sqrt{0.5}) + 0.9^2, \sigma = 0.9)$	$Uniform(0, \infty)$
ω_{V1}	$logNormal(\mu = \log(\sqrt{0.5}) + 0.9^2, \sigma = 0.9)$	$Uniform(0, \infty)$
ω_{V2}	$logNormal(\mu = \log(\sqrt{0.5}) + 0.9^2, \sigma = 0.9)$	$Uniform(0, \infty)$
ρ	$LKJ(\eta = 1)$	$LKJ(\eta = 1)$
σ_{prop}	$Half-Cauchy(\mu = 0, \sigma = 0.45)$	$Uniform(0, \infty)$
σ_{add}	$Half-Cauchy(\mu = 0, \sigma = 0.75)$	$Uniform(0, \infty)$

5.5.2 Estimation Algorithm Settings

5.5.2.1 Automatic Differentiation Variational Inference Settings

The `cmdstanr $variational()` function was used for ADVI estimation. The following settings were used:

- algorithm: "meanfield"
- init: an initialization function based on the prior distribution

```
# Wiczling prior
init0 <- function(){
list(CL_pop = exp(rnorm(1, log(40), 0.1)),
     Q_pop = exp(rnorm(1, log(80), 0.1)),
     V1_pop = exp(rnorm(1, log(55), 0.1)),
     V2_pop = exp(rnorm(1, log(70), 0.1)))
```

```
}
# informative prior
init0 <- function(){
list(CL_pop = exp(rnorm(1, log(40), 0.1)),
     Q_pop = exp(rnorm(1, log(150), 0.1)),
     V1_pop = exp(rnorm(1, log(90), 0.1)),
     V2_pop = exp(rnorm(1, log(105), 0.1)))
}
```

- convergence tolerance on the relative norm of the objective (tol_rel_obj): 0.003 (for allometric scaling and allometric scaling + maturation models); 0.005 (for allometric scaling + Hill maturation + additional covariate models); 0.01 (for cross-validation models)
- maximum number of iterations (iter): 15000 (17500 for cross-validation models)
- number of samples for Monte Carlo estimate of gradients (grad_samples): 4
- number of samples for Monte Carlo estimate of ELBO (elbo_samples): 150 (50 for cross-validation models)
- evaluate ELBO every Nth iteration (eval_elbo): 100 (Default)
- maximum tries to refit: 6 for cross-validation models
- all other parameters used `$variational()` defaults

5.5.2.2 No U Turns Hamiltonian Monte Carlo Simulation Settings

The `cmdstanr` `$sample()` function was used for HMC estimation. The following settings were used:

- algorithm: "hmc" (Default)
- engine: "nuts" (Default)
- init: an initialization function based on the true population parameters
- init: an initialization function based on the prior distribution

```
# Wiczling prior
init0 <- function(){
list(CL_pop = exp(rnorm(1, log(40), 0.1)),
     Q_pop = exp(rnorm(1, log(80), 0.1)),
     V1_pop = exp(rnorm(1, log(55), 0.1)),
     V2_pop = exp(rnorm(1, log(70), 0.1)))
}
# informative and more weakly informative prior
```

```
init0 <- function(){
list(CL_pop = exp(rnorm(1, log(40), 0.1)),
     Q_pop = exp(rnorm(1, log(150), 0.1)),
     V1_pop = exp(rnorm(1, log(90), 0.1)),
     V2_pop = exp(rnorm(1, log(105), 0.1)))
}
```

- number of warmup iterations to run per chain (iter_warmup): 2500 (3000 for sensitivity analysis with more weakly informative prior)
- number of post-warmup iterations to run per chain (iter_sampling): 250
- number of Markov chains to run (chains): 4
- maximum number of MCMC chains to run in parallel (parallel_chains): 4
- The adaptation target acceptance statistic (adapt_delta): 0.85
- The number of iterations between printed screen updates (refresh): 500
- all other parameters used `$sample()` defaults

5.5.3 Diagnostic Plots

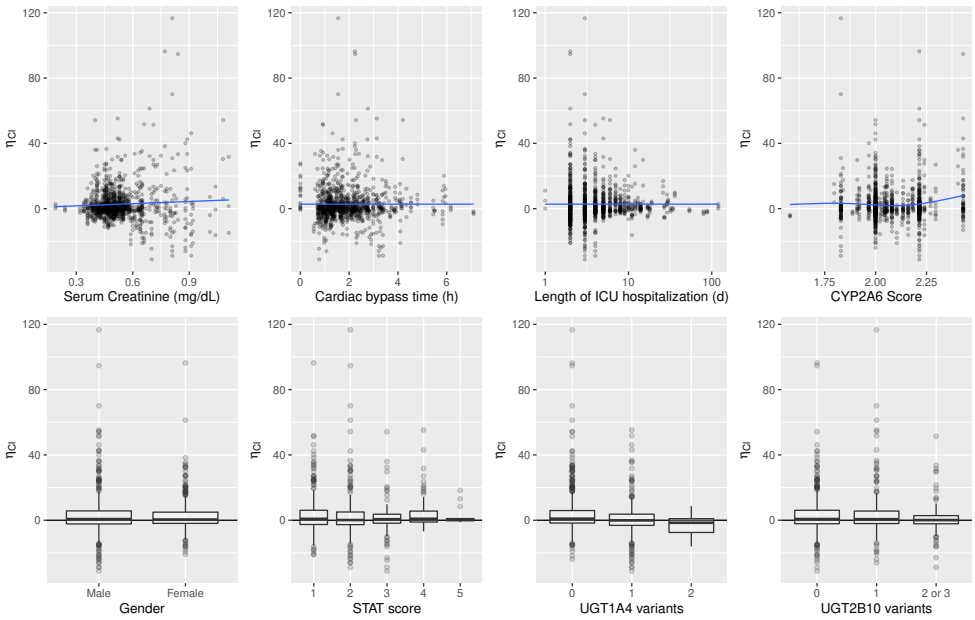

Figure 5.9: Total clearance random effects versus non-genotype covariates for allometric scaling and Hill maturation model. Blue lines are loess smoothers.

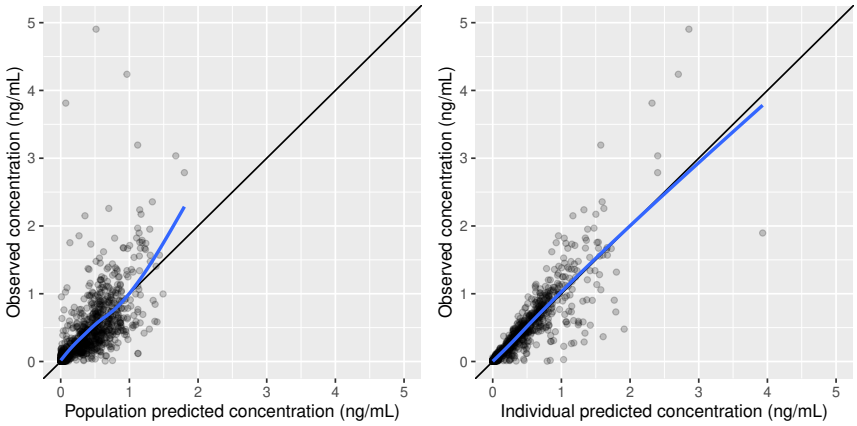

Figure 5.10: Observed vs. predicted concentration for allometric scaling and Hill maturation model. Blue lines are loess smoothers.

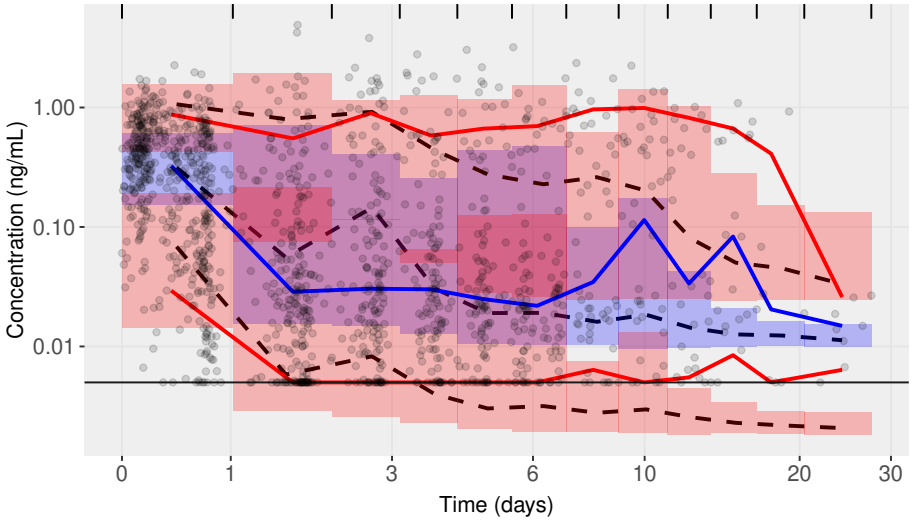

Figure 5.11: Visual predictive check for allometric scaling and Hill maturation model with 10th, 50th and 90th percentile of observed values (solid lines) and theoretical values (dashed lines) along with 90% prediction interval for theoretical percentiles (shaded regions). Filled circles indicate observed values and the lower limit of quantification of 0.005 ng/mL is represented by a horizontal grey line; time was binned using the Jenks natural breaks classification method.

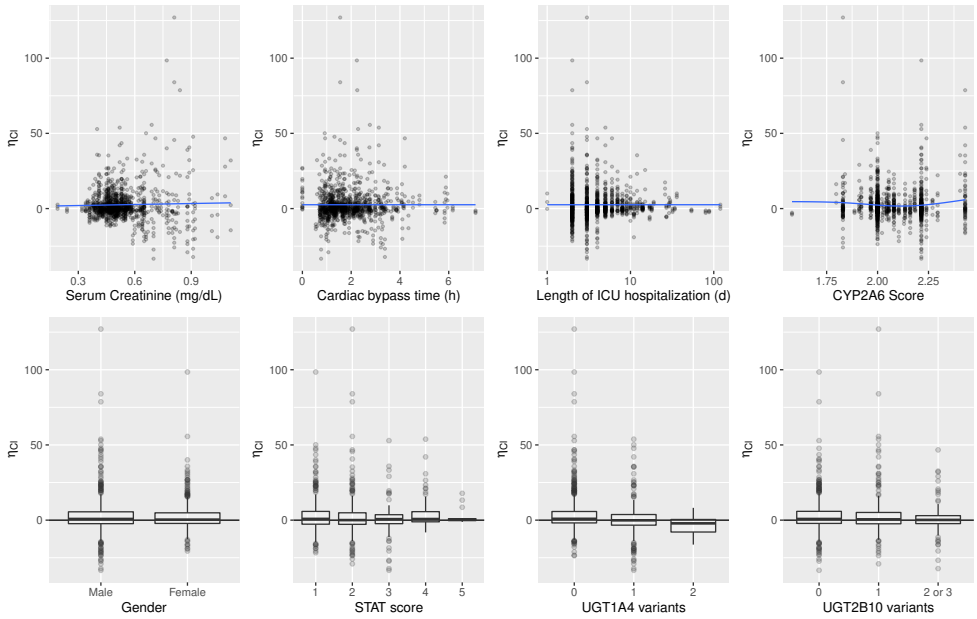

Figure 5.12: Total clearance random effects versus non-genotype covariates for allometric scaling, Hill maturation, and *CYP2A6* score model. Blue lines are loess smoothers.

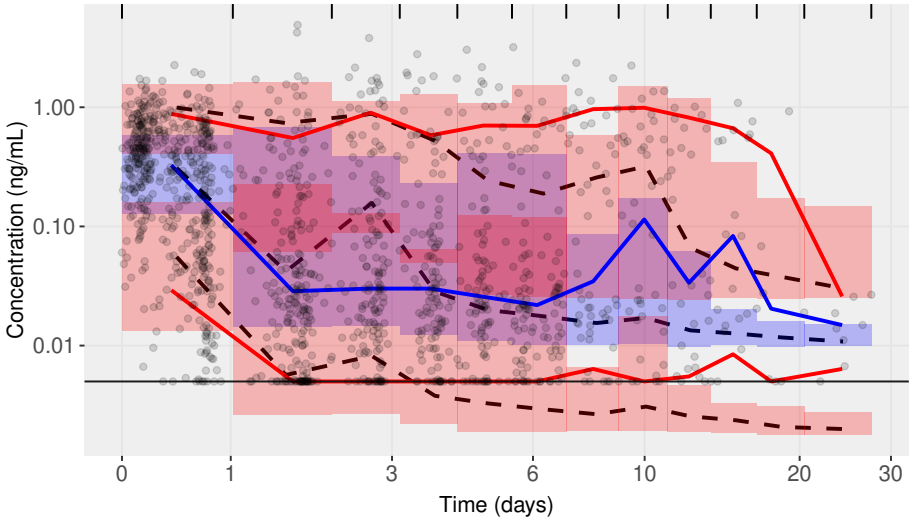

Figure 5.13: Visual predictive check for allometric scaling, Hill maturation, and *CYP2A6* score model with 10th, 50th and 90th percentile of observed values (solid lines) and theoretical values (dashed lines) along with 90% prediction interval for theoretical percentiles (shaded regions). Filled circles indicate observed values and the lower limit of quantification of 0.005 ng/mL is represented by a horizontal grey line; time was binned using the Jenks natural breaks classification method.

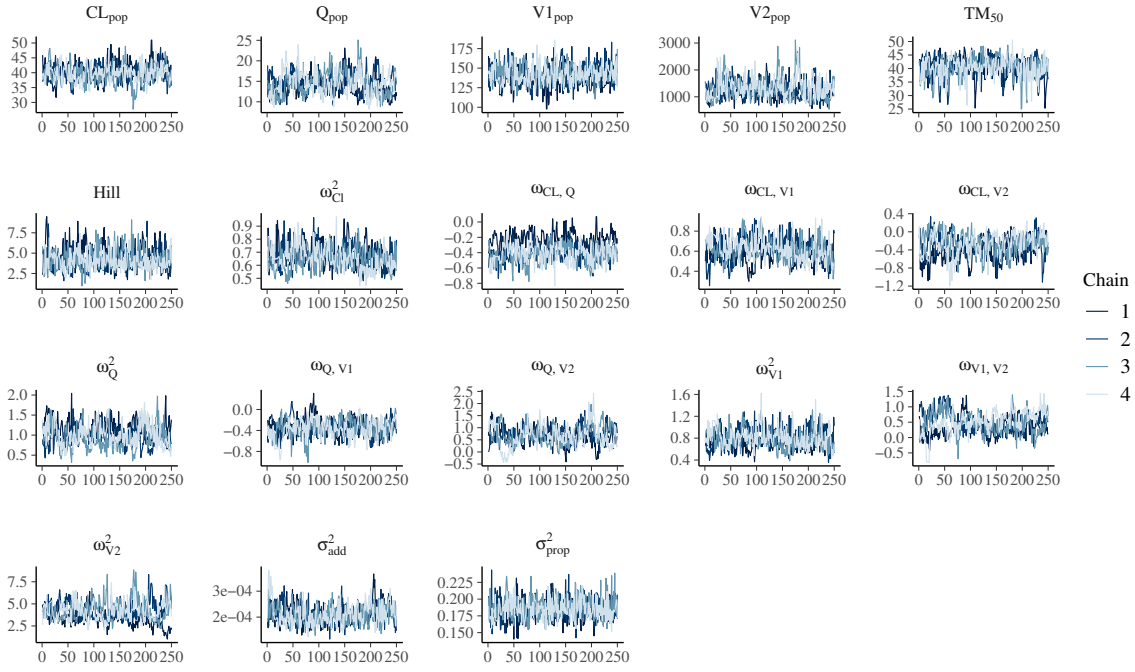

Figure 5.14: HMC trace plots for model with allometric scaling and Hill maturation using main analysis informative prior

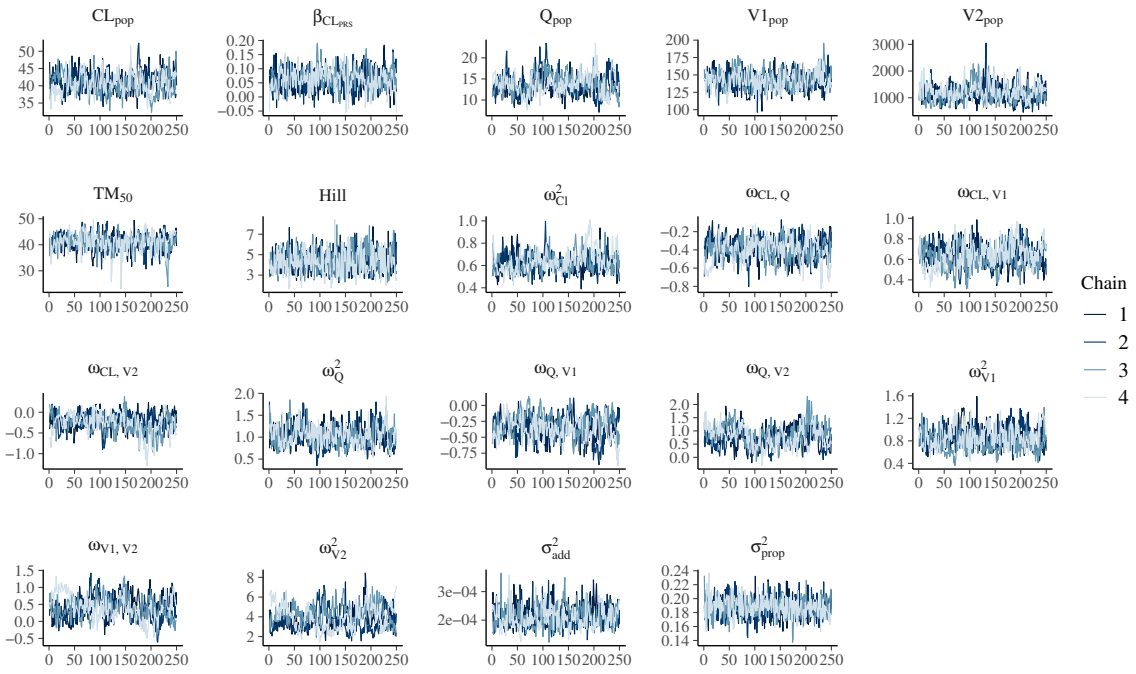

Figure 5.15: HMC trace plots for model with allometric scaling, Hill maturation, and *CYP2A6* score using main analysis informative prior

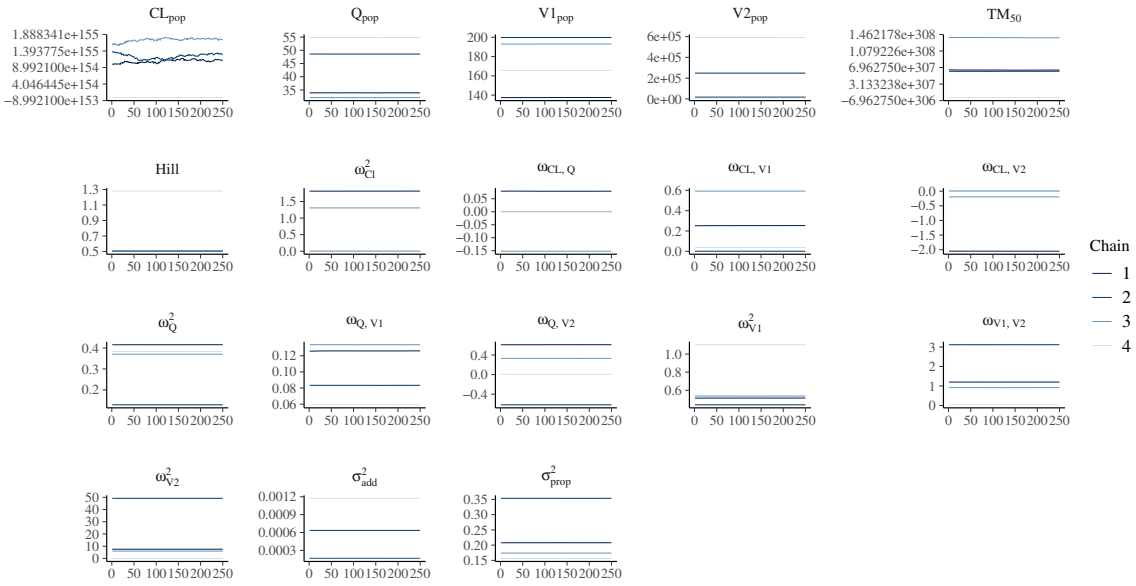

Figure 5.16: HMC trace plots for model with allometric scaling and Hill maturation using sensitivity analysis non-informative prior

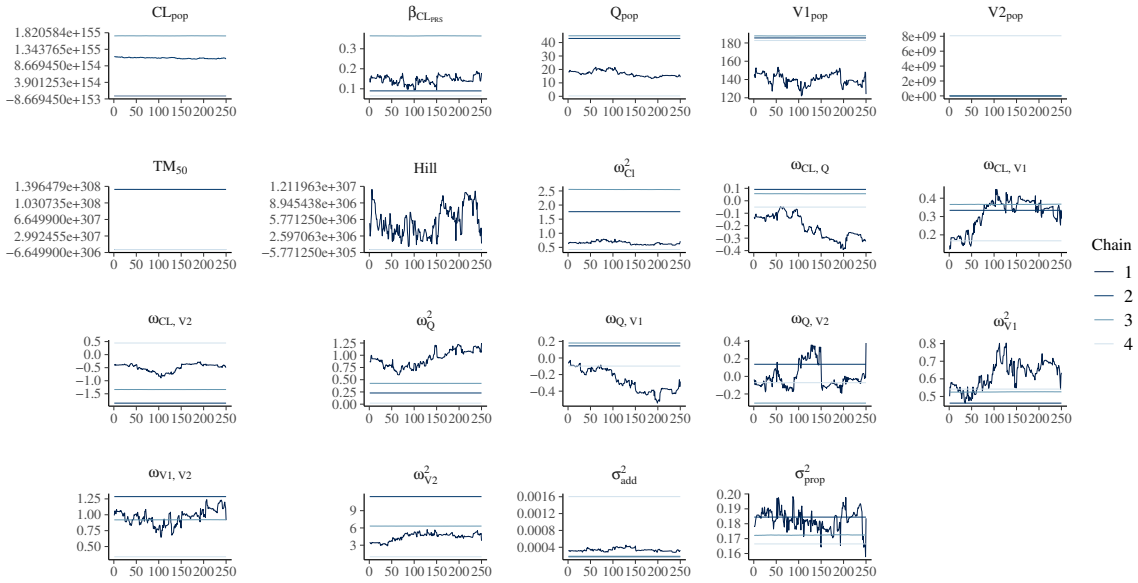

Figure 5.17: HMC trace plots for model with allometric scaling, Hill maturation, and *CYP2A6* score using sensitivity analysis non-informative prior

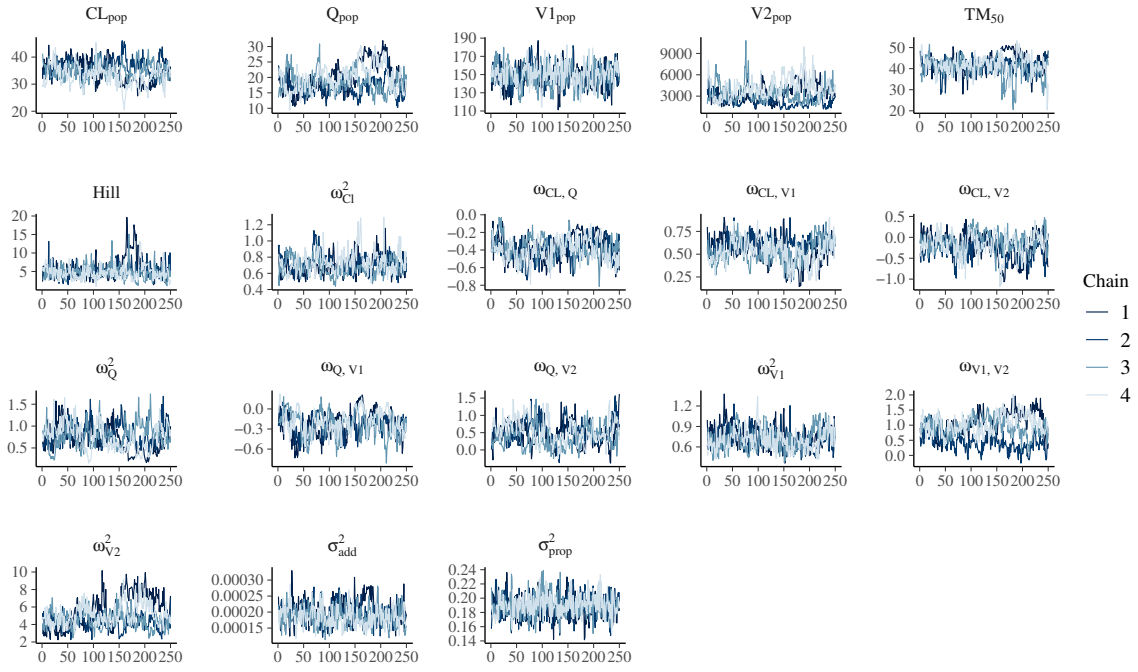

Figure 5.18: HMC trace plots for model with allometric scaling and Hill maturation using sensitivity analysis more weakly informative prior

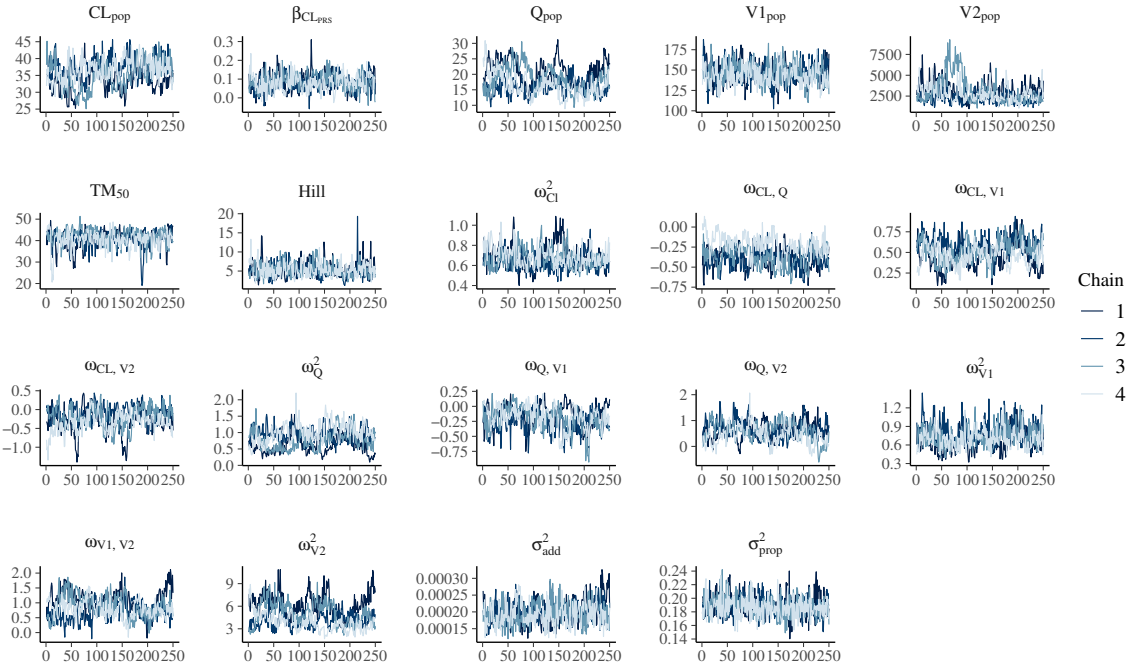

Figure 5.19: HMC trace plots for model with allometric scaling, Hill maturation, and *CYP2A6* score using sensitivity analysis more weakly informative prior

5.5.4 Calculation of *CYP2A6* Effect in terms of Equivalent Change in Weight

From the *CYP2A6* model average $CL = CL_{pop} \times \left(\frac{WT_i}{70}\right)^{0.75} \times \frac{1}{1+\left(\frac{TM_{50}}{PMA_i}\right)^{Hill}} \times \exp(\beta_{CL,CYP2A6})$ for a subject with standardized *CYP2A6* score of 1 (around 2.22 on original scale). A subject with the equivalent postmenstrual age, standardized *CYP2A6* score of 0 (around 2.08 on original scale), and weight changed by a factor Δ will have average $CL = CL_{pop} \times \left(\frac{WT_i \times \Delta}{70}\right)^{0.75} \times \frac{1}{1+\left(\frac{TM_{50}}{PMA_i}\right)^{Hill}}$. Setting these equations equal and solving for Δ will give the proportional weight change that has the same impact on CL as a one-unit increase in standardized *CYP2A6* score.

$$
\begin{aligned}
CL_{pop} \times \left(\frac{WT_i}{70}\right)^{0.75} \times \frac{1}{1+\left(\frac{TM_{50}}{PMA_i}\right)^{Hill}} \times \exp(\beta_{CL,CYP2A6}) &= CL_{pop} \times \left(\frac{WT_i \times \Delta}{70}\right)^{0.75} \times \frac{1}{1+\left(\frac{TM_{50}}{PMA_i}\right)^{Hill}} \\
\left(\frac{WT_i}{70}\right)^{0.75} \times \exp(\beta_{CL,CYP2A6}) &= \left(\frac{WT_i}{70}\right)^{0.75} \times \Delta^{0.75} \\
\beta_{CL,CYP2A6} &= 0.75 \log(\Delta) \\
\exp\left(\frac{\beta_{CL,CYP2A6}}{0.75}\right) &= \Delta
\end{aligned}
$$

For the *CYP2A6* model, $\beta_{CL,CYP2A6} = 0.0627$ so the equivalent proportional weight change is $\Delta = \exp\left(\frac{0.0627}{0.75}\right) \approx 1.09$

CHAPTER 6

Conclusion

The best approach to encourage adoption of Bayesian methods is by showing how these methods can provide practical advantages to solve real problems. In this dissertation, we have demonstrated this methodology in two application areas, semi-parametric regression modeling with CPMs and population PK analysis.

In Chapter 2, we described how the CPM model can be reparameterized to handle a large number of ordinal categories. Using a Dirichlet prior anchored to the conditional probabilities of category membership for a specific covariate vector along with the inverse of the assumed link function we induce a prior for the ordered intercept parameters. The Bayesian CPM is a flexible model which can handle both continuous and discrete ordered outcomes and estimation of the full conditional CDF, along with quantiles and other functionals using a single model fit. The model performed reasonably well for the simulations studied, although it is best suited for cases with fairly dense data that are sufficient to describe the posterior CDF.

In Chapter 3, the ADVI approximation was compared to MCMC using extensive simulations encompassing two PK models (one and two compartment IV infusion), two data sampling schemes (dense and sparse), and four prior specifications (strong, weak, and misspecified informative priors, and a non-informative prior). The estimation approaches were compared for population and individual parameter estimation, population and individual posterior concentration prediction, fit time, and model selection using a correctly specified and two incorrectly specified models. ADVI is best used when fitting many models, such as during model development and refinement or for model selection using cross-validation, or when the goal is to evaluate models quickly and approximate results are acceptable. A more accurate method such as MCMC should be used for final estimates.

Chapter 4 detailed a frequentist population PK analysis of dexmedetomidine using real-world EHR data. While this analysis demonstrated the benefit of using real-world data for population PK analysis by confirming the importance of weight and age for dexmedetomidine PK it did not find evidence for pharmacogenetic effects of $UGT1A4$ or $UGT2B10$ variants or $CYP2A6$ risk score. Reanalyzing the dexmedetomdine PK data using a Bayesian approach in Chapter 5 addressed a limitation of the frequentist analysis by stabilizing estimation using prior information from previously published studies. In addition, this analysis demonstrated the combined estimation approach advocated in Chapter 3 by using ADVI for model development and selection and MCMC for final model estimates and posterior predictive checks.

Throughout, we have described the advantages of the Bayesian approach. All parameter estimates and intervals can be interpreted intuitively. For example, in the HIV inflammatory biomarker case study model from Chapter 2 there is a 95% probability that a white, male, nonsmoker with average age and CD4 count in the LiNC study and BMI of 30.3 has a median IL-6 between 2.02 and 2.68; for the same subject there is a 75% probability that the 90th percentile of IL-6 is greater than 5.75. From the Bayesian dexmedetomdine PK reanalysis, the probability that total clearance is between 36.1 and 45.6 L/h for a subject with standardized 70kg weight, full adult maturity, and mean $CYP2A6$ score is 90%. The dexmedetomdine reanalysis also showed the benefit of informative priors to stabilize estimates. The mean $V2_{pop}$ for the model with allometric scaling, Hill maturation, and $CYP2A6$ score was estimated as 5756 L in the frequentist analysis, far larger than any value previously reported in the literature. The Bayesian estimate of 1170 L, while still large, is much more plausible.

We have also attempted to address what we view as the main drawbacks of Bayesian modeling, the additional time required to both define and fit models. While attention to the prior will remain a prerequisite for all Bayesian analysis, the reparameterization and prior choices explored for the Bayesian CPM represent an example of a reasonable starting point that requires the analyst to specify fewer hyperparameters than alternatives that parameterize the intercepts directly. The ADVI approximation addresses the computational burden of fitting many Bayesian population PK analysis.

There are several avenues for additional research related to the Bayesian CPM model. First, a direct comparison to other Bayesian semiparametric modeling approaches would help clarify the connections between approaches and highlight benefits and drawbacks. Second, additional flexibility can potentially be gained by substituting the Dirichlet prior for a infinite-dimensional Bayesian nonparametric prior, such as a Dirichlet process prior, which does not require the number of categories to be known *a priori* and could better accommodate Bayesian updating of the posterior with data containing new distinct continuous values or categories. Third, if inference for the parameters is not of primary interest, specification of the link function could be avoided by either estimating it nonparametrically or using a more flexible mixture link function.

For Bayesian PK modeling, the results from ADVI are not identical to those using the MCMC gold standard, but the decrease in fit time can justify the increase in bias and misestimated precision depending on the goals of the modeling procedure. Two directions for future work in this area are evaluation of extensions to address the inaccuracies of mean-field ADVI such as full-rank ADVI or LR-ADVI and application to more complex pharmacometric models, such as pharmacokinetic-pharmacodynamic or physiologically based pharmacokinetic models, which could potentially benefit from the hybrid approach of combining large observational datasets with prior information from small, well-designed studies.

References

1. McElreath R. Statistical rethinking: A Bayesian course with examples in R and Stan. Boca Raton: CRC Press/Taylor & Francis Group; 2016.

2. Kruschke JK. Doing Bayesian data analysis: A tutorial with R, JAGS, and Stan. Edition 2. Boston: Academic Press; 2015.

3. Lunn D. The BUGS book: A practical introduction to Bayesian analysis. Boca Raton, FL: CRC Press, Taylor & Francis Group; 2013.

4. Gelman A, Carlin J, Stern H, Dunson D, Vehtari A, Rubin D. Bayesian Data Analysis. Third edition. Boca Raton: CRC Press; 2014.

5. Wasserman L. All of statistics: A concise course in statistical inference. New York: Springer; 2004.

6. Brooks S. Handbook for Markov chain Monte Carlo [Internet]. Boca Raton, Fla.: Taylor & Francis; 2011 [cited 2022 Mar 27]. Available from: http://grail.eblib.com.au/patron/FullRecord.aspx?p=762505

7. Turkman MAA, Paulino CD, Müller P. Computational Bayesian statistics: An introduction. Cambridge ; New York, NY: Cambridge University Press; 2019.

8. Beal MJ, Ghahramani Z. The Variational Bayesian EM Algorithm for Incomplete Data: With Application to Scoring Graphical Model Structures. In: Bernardo JM, Bayarri MJ, Dawid AP, Berger JO, Heckerman D, Smith AFM, et al., editors. Proceedings of the Seventh Valencia International Meeting. Oxford University Press; 2003. page 10.

9. Winn J, Bishop CM. Variational Message Passing. Journal of Machine Learning Research 2005;6:661–94.

10. Beal MJ, Ghahramani Z. Variational Bayesian learning of directed graphical models with hidden variables. Bayesian Analysis [Internet] 2006 [cited 2021 Oct 11];1. Available from: https://projecteuclid.org/journals/bayesian-analysis/volume-1/issue-4/Variational-Bayesian-learning-of-directed-graphical-models-with-hidden-variables/10.1214/06-BA126.full

11. Wainwright MJ, Jordan MI. Graphical Models, Exponential Families, and Variational Inference. Foundations and Trends in Machine Learning [Internet] 2008 [cited 2021 Oct 27];1:1–305. Available from: http://www.nowpublishers.com/article/Details/MAL-001

12. Chappell MA, Groves AR, Whitcher B, Woolrich MW. Variational Bayesian Inference for a Nonlinear Forward Model. IEEE Transactions on Signal Processing 2009;57:223–36.

13. Wang C. Variational Inference in Nonconjugate Models. Journal of Machine Learning Research 2013;14:1005–31.

14. Kucukelbir A, Tran D, Ranganath R, Gelman A, Blei DM. Automatic Differentiation Variational Inference. Journal of Machine Learning Research 2017;18:430–74.

15. Blei DM, Kucukelbir A, McAuliffe JD. Variational Inference: A Review for Statisticians. Journal of the American Statistical Association [Internet] 2017 [cited 2018 Oct 23];112:859–77. Available from: http://arxiv.org/abs/1601.00670

16. Agresti A. Categorical Data Analysis. 2nd ed. New York: Wiley-Interscience; 2002.

17. Walker SH, Duncan DB. Estimation of the probability of an event as a function of several independent variables. Biometrika [Internet] 1967 [cited 2020 Jul 17];54:167–79. Available from: https://academic.oup.com/biomet/article-lookup/doi/10.1093/biomet/54.1-2.167

18. McCullagh P. Regression Models for Ordinal Data. Journal of the Royal Statistical Society. Series B (Methodological) [Internet] 1980;42:109–42. Available from: http://www.jstor.org/stable/2984952

19. Albert JH, Chib S. Bayesian Analysis of Binary and Polychotomous Response Data. Journal of the American Statistical Association [Internet] 1993 [cited 2018 May 9];88:669. Available from: https://www.jstor.org/stable/2290350?origin=crossref

20. Albert J, Chib S. Bayesian Methods for Cumulative, Sequential and Two-step Ordinal Data Regression Models. 1997;33.

21. Johnson VE, Albert J. Ordinal data modeling. New York: Springer; 1999.

22. Peterson B, Harrell FE. Partial Proportional Odds Models for Ordinal Response Variables. Applied Statistics [Internet] 1990 [cited 2020 Jul 18];39:205. Available from: https://www.jstor.org/stable/10.2307/2347760?origin=crossref

23. Lang JB. Bayesian ordinal and binary regression models with a parametric family of mixture links. Computational Statistics & Data Analysis [Internet] 1999 [cited 2018 Jul 17];31:59–87. Available from: http://linkinghub.elsevier.com/retrieve/pii/S0167947399000079

24. Congdon P. Bayesian models for categorical data. Chichester ; New York: Wiley; 2005.

25. Liu Q, Shepherd BE, Li C, Harrell FE. Modeling continuous response variables using ordinal regression. Statistics in Medicine [Internet] 2017 [cited 2018 Jan 19];36:4316–35. Available from: http://onlinelibrary.wiley.com/doi/10.1002/sim.7433/abstract

26. Harrell FE. Regression modeling strategies: With applications to linear models, logistic and ordinal regression, and survival analysis. Second edition. Cham Heidelberg New York: Springer; 2015.

27. Tian Y, Hothorn T, Li C, Harrell FE, Shepherd BE. An empirical comparison of two novel transformation models. Statistics in Medicine [Internet] 2019 [cited 2020 Feb 7];Available from: https://onlinelibrary.wiley.com/doi/abs/10.1002/sim.8425

28. Zeng D, Lin DY. Maximum likelihood estimation in semiparametric regression models with censored data. Journal of the Royal Statistical Society: Series B (Statistical Methodology) [Internet] 2007 [cited 2020 Aug 4];69:507–64. Available from: https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/j.1369-7412.2007.00606.x

29. Gelfand AE. Approaches for Semiparametric Bayesian Regression. In: Ghosh S, editor. Asymptotics, Nonparametrics, and Time Series. CRC Press; 1999. page 615–38.

30. Brunner LJ. Bayesian linear regression with error terms that have symmetric unimodal densities. Journal of Nonparametric Statistics [Internet] 1995 [cited 2020 Aug 13];4:335–48. Available from: http://www.tandfonline.com/doi/abs/10.1080/10485259508832625

31. Kottas A, Gelfand AE. Bayesian Semiparametric Median Regression Modeling. Journal of the American Statistical Association [Internet] 2001 [cited 2020 Jul 28];96:1458–68. Available from: http://www.tandfonline.com/doi/abs/10.1198/016214501753382363

32. DeYoreo M, Kottas A. Bayesian nonparametric density regression for ordinal responses. In: Flexible Bayesian regression modelling. Academic Press; 2020. page 65–89.

33. Song X-Y, Lu Z-H. Semiparametric transformation models with Bayesian P-splines. Statistics and Computing [Internet] 2012 [cited 2020 Jul 2];22:1085–98. Available from: http://link.springer.com/10.1007/s11222-011-9280-x

34. Tang N, Wu Y, Chen D. Semiparametric Bayesian analysis of transformation linear mixed models. Journal of Multivariate Analysis [Internet] 2018 [cited 2020 Aug 1];166:225–40. Available from: https://linkinghub.elsevier.com/retrieve/pii/S0047259X18300976

35. Mallick BK, Walker S. A Bayesian semiparametric transformation model incorporating frailties. Journal of Statistical Planning and Inference [Internet] 2003 [cited 2020 Jun 30];112:159–74. Available from: https://linkinghub.elsevier.com/retrieve/pii/S0378375802003300

36. Lin J, Sinha D, Lipsitz S, Polpo A. Semiparametric Bayesian Survival Analysis using Models with Log-Linear Median. Biometrics [Internet] 2012 [cited 2020 Aug 4];68:1136–45. Available from: https://www.jstor.org/stable/41806032

37. Hanson TE, Jara A. Surviving fully Bayesian nonparametric regression models [Internet]. In: Damien P, Dellaportas P, Polson NG, Stephens DA, editors. Bayesian Theory and Applications. Oxford University Press; 2013 [cited 2020 Jul 2]. page 593–616.Available from: http://www.oxfordscholarship.com/view/10.1093/acprof:oso/9780199695607.001.0001/acprof-9780199695607-chapter-30

38. Hanson T, Yang M. Bayesian Semiparametric Proportional Odds Models. Biometrics [Internet] 2007 [cited 2020 Jun 30];63:88–95. Available from: http://onlinelibrary.wiley.com/doi/abs/10.1111/j.1541-0420.2006.00671.x

39. Ibrahim JG, Chen M-H, Sinha D. Bayesian survival analysis. Softcover repr. of the hardcover 1st edition 2001, corr. 2nd printing. New York: Springer; 2010.

40. Müller P, Quintana FA, Jara A, Hanson T. Bayesian nonparametric data analysis. Cham: Springer; 2015.

41. Hjort NL, editor. Bayesian nonparametrics. Cambridge, UK ; New York: Cambridge University Press; 2010.

42. Gibaldi M, Perrier D. Pharmacokinetics, Second Edition [Internet]. Taylor & Francis; 1982. Available from: https://books.google.com/books?id=cZa3DlaSTqIC

43. Rosenbaum S. Basic pharmacokinetics and pharmacodynamics: An integrated textbook and computer simulations. Second edition. Hoboken, New Jersey: John Wiley & Sons, Inc; 2017.

44. Sheiner LB, Rosenberg B, Marathe VV. Estimation of population characteristics of pharmacokinetic parameters from routine clinical data. Journal of Pharmacokinetics and Biopharmaceutics [Internet] 1977 [cited 2019 Sep 10];5:445–79. Available from: http://link.springer.com/10.1007/BF01061728

45. Sheiner LB, Beal SL. Bayesian individualization of pharmacokinetics: Simple implementation and comparison with non-Bayesian methods. Journal of Pharmaceutical Sciences [Internet] 1982 [cited 2020 Sep 8];71:1344–8. Available from: http://onlinelibrary.wiley.com/doi/abs/10.1002/jps.2600711209

46. Sheiner LB, Beal SL. Evaluation of methods for estimating population pharmacokinetic parameters. I. Michaelis-menten model: Routine clinical pharmacokinetic data. Journal of Pharmacokinetics and Biopharmaceutics [Internet] 1980 [cited 2020 Aug 20];8:553–71. Available from: http://link.springer.com/10.1007/BF01060053

47. Sheiner LB, Beal SL. Evaluation of methods for estimating population pharmacokinetic parameters II. Biexponential model and experimental pharmacokinetic data. Journal of Pharmacokinetics and Biopharmaceutics [Internet] 1981 [cited 2020 Aug 20];9:635–51. Available from: http://link.springer.com/10.1007/BF01061030

48. Mandema JW, Verotta D, Sheiner LB. Building population pharmacokineticpharmacodynamic models. I. Models for covariate effects. Journal of Pharmacokinetics and Biopharmaceutics [Internet] 1992 [cited 2020 Aug 27];20:511–28. Available from: http://link.springer.com/10.1007/BF01061469

49. Sheiner LB, Ludden TM. Population Pharmacokinetics/Dynamics. Annual Review of Pharmacology and Toxicology 1992;32:185–209.

50. Wade JR, Beal SL, Sambol NC. Interaction between structural, statistical, and covariate models in population pharmacokinetic analysis. Journal of Pharmacokinetics and Biopharmaceutics [Internet] 1994 [cited 2020 Aug 27];22:165–77. Available from: http://link.springer.com/10.1007/BF02353542

51. Wählby U, Jonsson EN, Karlsson MO. Comparison of stepwise covariate model building strategies in population pharmacokinetic-pharmacodynamic analysis. AAPS PharmSci [Internet] 2002 [cited 2020 Sep 4];4:68–79. Available from: http://link.springer.com/10.1208/ps040427

52. Wang Y. Derivation of various NONMEM estimation methods. Journal of Pharmacokinetics and Pharmacodynamics [Internet] 2007 [cited 2020 Aug 13];34:575–93. Available from: http://link.springer.com/10.1007/s10928-007-9060-6

53. Bonate PL. Recommended reading in population pharmacokinetic pharmacodynamics. The AAPS Journal [Internet] 2005 [cited 2019 Sep 10];7:E363–73. Available from: http://link.springer.com/10.1208/aapsj070237

54. Davidian M, Giltinan DM. Nonlinear models for repeated measurement data. Boca Raton, Fla: Chapman & Hall/CRC; 1998.

55. Bonate PL. Pharmacokinetic-pharmacodynamic modeling and simulation. 2. ed. New York: Springer; 2011.

56. Best NG, Tan KKC, Gilks WR, Spiegelhalter DJ. Estimation of population pharmacokinetics using the Gibbs sampler. Journal of Pharmacokinetics and Biopharmaceutics [Internet] 1995 [cited 2020 Aug 13];23:407–35. Available from: http://link.springer.com/10.1007/BF02353641

57. Gelman A, Bois F, Jiang J. Physiological Pharmacokinetic Analysis Using Population Modeling and Informative Prior Distributions. Journal of the American Statistical Association [Internet] 1996 [cited 2019 Aug 16];91:1400–12. Available from: http://www.tandfonline.com/doi/abs/10.1080/01621459.1996.10476708

58. Rosner GL, Muller P. Bayesian Population Pharmacokinetic and Pharmacodynamic Analyses Using Mixture Models. Journal of Pharmacokinetics and Biopharmaceutics 1997;25:25.

59. Wakefield JC, Smith AFM, Racine-Poon A, Gelfand AE. Bayesian Analysis of Linear and Non-Linear Population Models by Using the Gibbs Sampler. Applied Statistics 1994;43:201–21.

60. Wakefield J, Bennett J. The Bayesian Modeling of Covariates for Population Pharmacokinetic Models. Journal of the American Statistical Association 1996;91:917–27.

61. Wakefield J. The Bayesian Analysis of Population Pharmacokinetic Models. Journal of the American Statistical Association [Internet] 1996 [cited 2019 Aug 16];91:62–75. Available from: http://www.tandfonline.com/doi/abs/10.1080/01621459.1996.10476664

62. Wakefield J, Aarons L, Racine-Poon A. The Bayesian approach to Population pharmacokinetic/pharmacodynamic modeling [Internet]. In: Bickel P, Diggle P, Fienberg S, Krickeberg K, Olkin I, Wermuth N, et al., editors. Case Studies in Bayesian Statistics. New York, NY: Springer New York; 1999 [cited 2020 Aug 14]. page 205–65.Available from: http://link.springer.com/10.1007/978-1-4612-1502-8_4

63. Lunn DJ, Best N, Thomas A, Wakefield J, Spiegelhalter D. Bayesian Analysis of Population PK/PD Models: General Concepts and Software. Journal of Pharmacokinetics and Pharmacodynamics 2002;29:37.

64. Lunn DJ. Bayesian Analysis of Population Pharmacokinetic/Pharmacodynamic Models [Internet]. In: Husmeier D, Dybowski R, Roberts S, editors. Probabilistic Modeling in Bioinformatics and Medical Informatics. London: Springer-Verlag; 2005 [cited 2020 Aug 13]. page 351–70.Available from: http://link.springer.com/10.1007/1-84628-119-9_11

65. Margossian CC, Zhang Y, Gillespie WR. Flexible and efficient Bayesian pharmacometrics modeling using Stan and Torsten, Part I. arXiv:2109.10184 [stat] [Internet] 2021 [cited 2021 Nov 9];Available from: http://arxiv.org/abs/2109.10184

66. Torsten: A Pharmacokinetic/Pharmacodynamic Model Library for Stan User Manual [Internet]. 2019 [cited 2019 Dec 1];Available from: https://github.com/metrumresearchgroup/Torsten/blob/master/docs/torsten_manual.pdf

67. Box GEP, Tiao GC. Bayesian inference in statistical analysis. Wiley classics library ed. New York: Wiley; 1992.

68. Gelman A, Vehtari A, Simpson D, Margossian CC, Carpenter B, Yao Y, et al. Bayesian Workflow. arXiv:2011.01808 [stat] [Internet] 2020 [cited 2020 Nov 18];Available from: http://arxiv.org/abs/2011.01808

69. McKinley TJ, Morters M, Wood JLN. Bayesian Model Choice in Cumulative Link Ordinal Regression Models. Bayesian Analysis [Internet] 2015 [cited 2019 Nov 2];10:1–30. Available from: http://arxiv.org/abs/1503.07642

70. Betancourt M. Ordinal Regression [Internet]. 2019 [cited 2020 Jul 3];Available from: https://betanalpha.github.io/assets/case_studies/ordinal_regression.html

71. Berger JO, Bernardo JM, Sun D. Overall Objective Priors. Bayesian Analysis [Internet] 2015 [cited 2020 Jul 2];10:189–221. Available from: http://projecteuclid.org/euclid.ba/1422556416

72. Stan Development Team. RStan: The R interface to Stan [Internet]. 2018;Available from: http://mc-stan.org/

73. Neal R. MCMC using Hamiltonian Dynamics. In: Fitzmaurice G, Brooks S, Gelman A, Jones GL, Meng X-L, editors. Handbook of Markov Chain Monte Carlo. New York: CRC Press, Taylor & Francis Group; 2011. page 113–62.

74. Koethe JR, Bian A, Shintani AK, Boger MS, Mitchell VJ, Erdem H, et al. Serum Leptin Level Mediates the Association of Body Composition and Serum C-Reactive Protein in HIV-Infected Persons on Antiretroviral Therapy. AIDS Research and Human Retroviruses [Internet] 2012 [cited 2020 Aug 13];28:552–7. Available from: http://www.liebertpub.com/doi/10.1089/aid.2011.0232

75. Koethe JR, Grome H, Jenkins CA, Kalams SA, Sterling TR. The metabolic and cardiovascular consequences of obesity in persons with HIV on long-term antiretroviral therapy: AIDS [Internet] 2015 [cited 2020 Aug 13];1. Available from: http://journals.lww.com/00002030-900000000-97959

76. Vehtari A, Gelman A, Gabry J. Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. Statistics and Computing [Internet] 2017 [cited 2019 Jun 25];27:1413–32. Available from: http://link.springer.com/10.1007/s11222-016-9696-4

77. Stern HS, Sinharay S. Bayesian Model Checking and Model Diagnostics [Internet]. In: Handbook of Statistics. Elsevier; 2005 [cited 2020 Aug 10]. page 171–92.Available from: https://linkinghub.elsevier.com/retrieve/pii/S016971610525006X

78. Kucukelbir A, Ranganath R, Gelman A, Blei DM. Automatic Variational Inference in Stan. arXiv:1506.03431 [stat] [Internet] 2015 [cited 2021 Oct 18];Available from: http://arxiv.org/abs/1506.03431

79. Ghahramani Z, Jordan MI. Factorial Hidden Markov Models.: Machine Learning 1997;29:245–73.

80. Attias H. A Variational Baysian Framework for Graphical Models. In: Advances in Neural Information Processing Systems. 2000. page 209–15.

81. Corduneanu A, Bishop CM. Variational Bayesian Model Selection for Mixture Distributions. 2001. page 27–34.

82. Blei DM. Latent Dirichlet Allocation. Journal of Machine Learning Research 2003;3:993–1022.

83. Blei DM, Jordan MI. Variational inference for Dirichlet process mixtures. Bayesian Analysis [Internet] 2006 [cited 2021 Sep 3];1. Available from: https://projecteuclid.org/journals/bayesian-analysis/volume-1/issue-1/Variational-inference-for-Dirichlet-process-mixtures/10.1214/06-BA104.full

84. Braun M, McAuliffe J. Variational Inference for Large-Scale Models of Discrete Choice. Journal of the American Statistical Association [Internet] 2010 [cited 2021 Sep 3];105:324–35. Available from: https://www.tandfonline.com/doi/full/10.1198/jasa.2009.tm08030

85. Yao Y, Vehtari A, Simpson D, Gelman A. Yes, but Did It Work?: Evaluating Variational Inference. arXiv:1802.02538 [stat] [Internet] 2018 [cited 2021 May 25];Available from: http://arxiv.org/abs/1802.02538

86. Vehtari A, Simpson D, Gelman A, Yao Y, Gabry J. Pareto Smoothed Importance Sampling. arXiv:1507.02646 [stat] [Internet] 2021 [cited 2021 Jul 1];Available from: http://arxiv.org/abs/1507.02646

87. Stan Development Team. Stan Modeling Language Users Guide and Reference Manual [Internet]. 2020;Available from: https://mc-stan.org

88. Spiegelhalter DJ, Best NG, Carlin BP, Linde A van der. Bayesian measures of model complexity and fit. Journal of the Royal Statistical Society: Series B (Statistical Methodology) [Internet] 2002 [cited 2018 Jul 17];64:583–639. Available from: http://doi.wiley.com/10.1111/1467-9868.00353

89. Watanabe S. Asymptotic Equivalence of Bayes Cross Validation and Widely Applicable Information Criterion in Singular Learning Theory. Journal of Machine Learning Research 2010;11:3571–94.

90. Friston KJ, Ashburner JT, Kiebel SJ, Nichols TE, Penny WD, editors. Statistical parametric mapping: The analysis of functional brain images. 1st ed. Amsterdam Boston: Elsevier / Academic Press; 2007.

91. Nott DJ, Tan SL, Villani M, Kohn R. Regression Density Estimation With Variational Methods and Stochastic Approximation. Journal of Computational and Graphical Statistics [Internet] 2012 [cited 2021 Aug 2];21:797–820. Available from: http://www.tandfonline.com/doi/abs/10.1080/10618600.2012.679897

92. Gelman A, Hwang J, Vehtari A. Understanding predictive information criteria for Bayesian models. Statistics and Computing [Internet] 2014 [cited 2021 Jul 4];24:997–1016. Available from: http://arxiv.org/abs/1307.5928

93. Duffull S, Waterhouse T, Eccleston J. Some Considerations on the Design of Population Pharmacokinetic Studies. Journal of Pharmacokinetics and Pharmacodynamics [Internet] 2005 [cited 2022 Mar 25];32:441–57. Available from: http://link.springer.com/10.1007/s10928-005-0034-2

94. Retout S, Comets E, Samson A, Mentré F. Design in nonlinear mixed effects models: Optimization using the Fedorov–Wynn algorithm and power of the Wald test for binary covariates. Statistics in Medicine [Internet] 2007 [cited 2022 Mar 25];26:5162–79. Available from: https://onlinelibrary.wiley.com/doi/10.1002/sim.2910

95. Meer AF van der, Touw DJ, Marcus MAE, Neef C, Proost JH. Influence of Erroneous Patient Records on Population Pharmacokinetic Modeling and Individual Bayesian Estimation. Therapeutic Drug Monitoring [Internet] 2012 [cited 2022 Mar 12];34:526–34. Available from: https://journals.lww.com/00007691-201210000-00008

96. Park S, Jo S, Lee W. A variational Bayes method for pharmacokinetic model. The Korean Journal of Applied Statistics [Internet] 2021 [cited 2022 Mar 30];34:9–23. Available from: https://doi.org/10.5351/KJAS.2021.34.1.009

97. Wang B, Titterington DM. Inadequacy of interval estimates corresponding to variational Bayesian approximations. In: Cowell R, Ghahramani Z, editors. Proceedings of the 10th International Workshop on Artifical Intelligence and Statistics. Society for Articial Intelligence & Statistics; 2005. page 373–80.

98. Turner RE, Sahani M. Two problems with variational expectation maximisation for time series models [Internet]. In: Barber D, Cemgil AT, Chiappa S, editors. Bayesian Time Series Models. Cambridge: Cambridge University Press; 2011 [cited 2021 Oct 12]. page 104–24.Available from: https://www.cambridge.org/core/product/identifier/CBO9780511984679A041/type/book_part

99. Giordano R, Broderick T, Jordan MI. Covariances, Robustness, and Variational Bayes. Journal of Machine Learning Research 2018;19:1–49.

100. Bertrand J, Mentre F, Lavielle M, Mesa H, Chatel K. Mathematical Expressions of the Pharmacokinetic and Pharmacodynamic Models implemented in the Monolix software. 2008;

101. Abuhelwa AY, Foster DJR, Upton RN. ADVAN-style analytical solutions for common pharmacokinetic models. Journal of Pharmacological and Toxicological Methods [Internet] 2015 [cited 2019 Aug 16];73:42–8. Available from: https://linkinghub.elsevier.com/retrieve/pii/S1056871915000362

102. Betancourt MJ, Girolami M. Hamiltonian Monte Carlo for Hierarchical Models. arXiv:1312.0906 [stat] [Internet] 2013 [cited 2020 Nov 27];Available from: http://arxiv.org/abs/1312.0906

103. Proost JH. Combined proportional and additive residual error models in population pharmacokinetic modelling. European Journal of Pharmaceutical Sciences [Internet] 2017 [cited 2020 Nov 11];109:S78–82. Available from: https://linkinghub.elsevier.com/retrieve/pii/S092809871730249X

104. Bassett R, Deride J. Maximum a posteriori estimators as a limit of Bayes estimators. Mathematical Programming [Internet] 2019 [cited 2022 Apr 28];174:129–44. Available from: http://link.springer.com/10.1007/s10107-018-1241-0

105. Maier C, Hartung N, Wiljes J de, Kloft C, Huisinga W. Bayesian Data Assimilation to Support Informed Decision Making in Individualized Chemotherapy. CPT: Pharmacometrics & Systems Pharmacology [Internet] 2020 [cited 2021 Apr 1];9:153–64. Available from: http://ascpt.onlinelibrary.wiley.com/doi/abs/10.1002/psp4.12492

106. Opper M, Archambeau C. The Variational Gaussian Approximation Revisited. Neural Computation [Internet] 2009 [cited 2021 Oct 15];21:786–92. Available from: https://direct.mit.edu/neco/article/21/3/786-792/7385

107. Tsamandouras N, Rostami-Hodjegan A, Aarons L. Combining the 'bottom up' and 'top down' approaches in pharmacokinetic modelling: Fitting PBPK models to observed clinical data: Parameter estimation in PBPK models. British Journal of Clinical Pharmacology [Internet] 2015 [cited 2022 May 4];79:48–55. Available from: https://onlinelibrary.wiley.com/doi/10.1111/bcp.12234

108. Langdon G, Gueorguieva I, Aarons L, Karlsson M. Linking preclinical and clinical whole-body physiologically based pharmacokinetic models with prior distributions in NONMEM. European Journal of Clinical Pharmacology [Internet] 2007 [cited 2022 Mar 16];63:485–98. Available from: http://link.springer.com/10.1007/s00228-007-0264-x

109. Barnard J, McCulloch R, Meng X-L. Modeling Covariance Matrices in Terms of Standard Deviations and Correlations, with Application to Shrinkage. Statistica Sinica 2000;10:1281–311.

110. James NT, Breeyear JH, Caprioli R, Edwards T, Hachey B, Kannankeril PJ, et al. Population pharmacokinetic analysis of dexmedetomidine in children using real-world data from electronic health records and remnant specimens. British Journal of Clinical Pharmacology [Internet] 2022 [cited 2022 Jan 31];bcp.15194. Available from: https://onlinelibrary.wiley.com/doi/10.1111/bcp.15194

111. Plambech MZ, Afshari A. Dexmedetomidine in the pediatric population: A review. MINERVA ANESTESIOLOGICA 2015;81:13.

112. Weerink MAS, Struys MMRF, Hannivoort LN, Barends CRM, Absalom AR, Colin P. Clinical Pharmacokinetics and Pharmacodynamics of Dexmedetomidine. Clinical Pharmacokinetics [Internet] 2017 [cited 2019 Dec 5];56:893–913. Available from: http://link.springer.com/10.1007/s40262-017-0507-7

113. Schwartz LI, Twite M, Gulack B, Hill K, Kim S, Vener DF. The Perioperative Use of Dexmedetomidine in Pediatric Patients with Congenital Heart Disease: An Analysis from the Congenital Cardiac Anesthesia Society-Society of Thoracic Surgeons Congenital Heart Disease Database. Anesthesia and Analgesia 2016;123:715–21.

114. Shuplock JM, Smith AH, Owen J, Van Driest SL, Marshall M, Saville B, et al. Association between perioperative dexmedetomidine and arrhythmias after surgery for congenital heart disease. Circulation. Arrhythmia and Electrophysiology 2015;8:643–50.

115. Ber J, Wiczling P, Hołysz M, Klupczyńska A, Bartkowska-Śniatkowska A, Bieda K, et al. Population Pharmacokinetic Model of Dexmedetomidine in a Heterogeneous Group of Patients. The Journal of Clinical Pharmacology [Internet] 2020 [cited 2021 Apr 14];60:1461–73. Available from: http://accp1.onlinelibrary.wiley.com/doi/abs/10.1002/jcph.1647

116. Damian MA, Hammer GB, Elkomy MH, Frymoyer A, Drover DR, Su F. Pharmacokinetics of Dexmedetomidine in Infants and Children After Orthotopic Liver Transplantation: Anesthesia & Analgesia [Internet] 2020 [cited 2021 Jan 23];130:209–16. Available from: http://journals.lww.com/10.1213/ANE.0000000000003761

117. Greenberg RG, Wu H, Laughon M, Capparelli E, Rowe S, Zimmerman KO, et al. Population Pharmacokinetics of Dexmedetomidine in Infants: Journal of Clinical Pharmacology. The Journal of Clinical Pharmacology [Internet] 2017 [cited 2021 Apr 14];57:1174–82. Available from: http://doi.wiley.com/10.1002/jcph.904

118. Morse JD, Cortinez LI, Anderson BJ. A Universal Pharmacokinetic Model for Dexmedetomidine in Children and Adults. Journal of Clinical Medicine [Internet] 2020 [cited 2021 Jan 20];9:3480. Available from: https://www.mdpi.com/2077-0383/9/11/3480

119. Pérez-Guillé M-G, Toledo-López A, Rivera-Espinosa L, Alemon-Medina R, Murata C, Lares-Asseff I, et al. Population Pharmacokinetics and Pharmacodynamics of Dexmedetomidine in Children Undergoing Ambulatory Surgery: Anesthesia & Analgesia [Internet] 2018 [cited 2021 Jan 19];127:716–23. Available from: http://journals.lww.com/00000539-201809000-00022

120. Potts AL, Anderson BJ, Warman GR, Lerman J, Diaz SM, Vilo S. Dexmedetomidine pharmacokinetics in pediatric intensive care – a pooled analysis. Pediatric Anesthesia [Internet] 2009 [cited 2021 Jan 7];19:1119–29. Available from: http://onlinelibrary.wiley.com/doi/abs/10.1111/j.1460-9592.2009.03133.x

121. Smuszkiewicz P, Wiczling P, Ber J, Warzybok J, Małkiewicz T, Matysiak J, et al. Pharmacokinetics of dexmedetomidine during analgosedation in ICU patients. Journal of Pharmacokinetics and Pharmacodynamics [Internet] 2018 [cited 2019 Dec 5];45:277–84. Available from: http://link.springer.com/10.1007/s10928-017-9564-7

122. Song I-K, Yi S, Lim H-S, Lee J-H, Kim E-H, Cho J-Y, et al. A Population Pharmacokinetic Model of Intravenous Dexmedetomidine for Mechanically Ventilated Children after Neurosurgery. Journal of Clinical Medicine [Internet] 2019 [cited 2021 Jan 23];8:1563. Available from: https://www.mdpi.com/2077-0383/8/10/1563

123. Su F, Gastonguay MR, Nicolson SC, DiLiberto M, Ocampo-Pelland A, Zuppa AF. Dexmedetomidine Pharmacology in Neonates and Infants After Open Heart Surgery: Anesthesia & Analgesia [Internet] 2016 [cited 2021 Jan 20];122:1556–66. Available from: http://journals.lww.com/00000539-201605000-00045

124. Takeuchi M, Nemoto S, Suzuki Y, Takahashi N, Takenaka N, Takata A, et al. Age-Specific Dose Regimens of Dexmedetomidine for Pediatric Patients in Intensive Care Following Elective Surgery: A Phase 3, Multicenter, Open-Label Clinical Trial in Japan. Pediatric Critical Care Medicine [Internet] 2021 [cited 2021 Apr 14];Available from: https://journals.lww.com/10.1097/PCC.0000000000002730

125. Wiczling P, Bartkowska-Śniatkowska A, Szerkus O, Siluk D, Rosada-Kurasińska J, Warzybok J, et al. The pharmacokinetics of dexmedetomidine during long-term infusion in critically ill pediatric patients. A Bayesian approach with informative priors. Journal of Pharmacokinetics and Pharmacodynamics [Internet] 2016 [cited 2020 Aug 13];43:315–24. Available from: http://link.springer.com/10.1007/s10928-016-9474-0

126. Zimmerman KO, Wu H, Laughon M, Greenberg RG, Walczak R, Schulman SR, et al. Dexmedetomidine Pharmacokinetics and a New Dosing Paradigm in Infants Supported With Cardiopulmonary Bypass: Anesthesia & Analgesia [Internet] 2019 [cited 2021 Jan 20];129:1519–28. Available from: http://journals.lww.com/00000539-201912000-00016

127. Zuppa AF, Nicolson SC, Wilder NS, Ibla JC, Gottlieb EA, Burns KM, et al. Results of a phase 1 multicentre investigation of dexmedetomidine bolus and infusion in corrective infant cardiac surgery. British Journal of Anaesthesia [Internet] 2019 [cited 2021 Apr 14];123:839–52. Available from: https://linkinghub.elsevier.com/retrieve/pii/S0007091219306452

128. DEXMEDETOMIDINE HCL injection [package insert] [Internet]. Berkeley Heights, NJ: Hikma Pharmaceuticals USA Inc.; 2020. Available from: https://dailymed.nlm.nih.gov/dailymed/drugInfo.cfm?setid=de6f01a2-9f2d-4c68-9c02-06643704acb1

129. Kohli U, Pandharipande P, Muszkat M, Sofowora GG, Friedman EA, Scheinin M, et al. CYP2A6 genetic variation and dexmedetomidine disposition. European Journal of Clinical Pharmacology 2012;68:937–42.

130. Rolle A, Paredes S, Cortínez LI, Anderson BJ, Quezada N, Solari S, et al. Dexmedetomidine metabolic clearance is not affected by fat mass in obese patients. British Journal of Anaesthesia [Internet] 2018 [cited 2021 Jan 19];120:969–77. Available from: https://linkinghub.elsevier.com/retrieve/pii/S0007091218301776

131. Guan Y, Li B, Wei W, Wang S, Yuen V-M, Liu Y, et al. Quantitative ultra-high-performance liquid chromatography-tandem mass spectrometry for determination of dexmedetomidine in pediatric plasma samples: Correlation with genetic polymorphisms. Biomedical chromatography: BMC 2019;33:e4683.

132. El-Boraie A, Taghavi T, Chenoweth MJ, Fukunaga K, Mushiroda T, Kubo M, et al. Evaluation of a weighted genetic risk score for the prediction of biomarkers of CYP2A6 activity. Addiction Biology 2020;25:e12741.

133. Choi L, Beck C, McNeer E, Weeks HL, Williams ML, James NT, et al. Development of a System for Postmarketing Population Pharmacokinetic and Pharmacodynamic Studies Using Real-World Data From Electronic Health Records. Clinical Pharmacology & Therapeutics [Internet] 2020 [cited 2020 Apr 30];107:934–43. Available from: https://onlinelibrary.wiley.com/doi/abs/10.1002/cpt.1787

134. Hagos FT, Horvat CM, Au AK, Conley YP, Li L, Poloyac SM, et al. Factors Contributing to Fentanyl Pharmacokinetic Variability Among Diagnostically Diverse Critically Ill Children. Clinical Pharmacokinetics [Internet] 2019 [cited 2021 Apr 24];58:1567–76. Available from: http://link.springer.com/10.1007/s40262-019-00773-1

135. EHR: Electronic Health Record (EHR) Data Processing and Analysis Tool [Internet]. 2020;Available from: https://cran.r-project.org/package=EHR

136. Van Driest SL, Marshall MD, Hachey B, Beck C, Crum K, Owen J, et al. Pragmatic pharmacology: Population pharmacokinetic analysis of fentanyl using remnant samples from children after cardiac surgery. British Journal of Clinical Pharmacology [Internet] 2016 [cited 2020 Jan 30];81:1165–74. Available from: https://bpspubs.onlinelibrary.wiley.com/doi/abs/10.1111/bcp.12903

137. Harris PA, Taylor R, Thielke R, Payne J, Gonzalez N, Conde JG. Research electronic data capture (REDCap)—A metadata-driven methodology and workflow process for providing translational research informatics support. Journal of Biomedical Informatics [Internet] 2009 [cited 2021 Apr 14];42:377–81. Available from: https://linkinghub.elsevier.com/retrieve/pii/S1532046408001226

138. Team RC. R: A Language and Environment for Statistical Computing [Internet]. 2020;Available from: https://www.R-project.org/

139. Monolix [Internet]. 2020;Available from: https://monolix.lixoft.com/

140. Ahn JE, Karlsson MO, Dunne A, Ludden TM. Likelihood based approaches to handling data below the quantification limit using NONMEM VI. J Pharmacokinet Pharmacodyn 2008;21.

141. Beal SL. Ways to Fit a PK Model with Some Data Below the Quantification Limit. Journal of Pharmacokinetics and Pharmacodynamics 2001;28:481–504.

142. Sukarnjanaset W, Wattanavijitkul T, Jarurattanasirikul S. Evaluation of FOCEI and SAEM Estimation Methods in Population Pharmacokinetic Analysis Using NONMEM® Across Rich, Medium, and Sparse Sampling Data. European Journal of Drug Metabolism and Pharmacokinetics [Internet] 2018 [cited 2021 Mar 5];43:729–36. Available from: http://link.springer.com/10.1007/s13318-018-0484-8

143. Delattre M, Lavielle M, Poursat M-A. A note on BIC in mixed-effects models. Electronic Journal of Statistics [Internet] 2014 [cited 2022 May 1];8. Available from: https://projecteuclid.org/journals/electronic-journal-of-statistics/volume-8/issue-1/A-note-on-BIC-in-mixed-effects-models/10.1214/14-EJS890.full

144. Lavielle M, Ribba B. Enhanced Method for Diagnosing Pharmacometric Models: Random Sampling from Conditional Distributions. Pharmaceutical Research [Internet] 2016 [cited 2021 Aug 20];33:2979–88. Available from: http://link.springer.com/10.1007/s11095-016-2020-3

145. Nguyen THT, Mouksassi M-S, Holford N, Al-Huniti N, Freedman I, Hooker AC, et al. Model Evaluation of Continuous Data Pharmacometric Models: Metrics and Graphics. CPT: Pharmacometrics & Systems Pharmacology [Internet] 2017 [cited 2021 Jan 6];6:87–109. Available from: http://ascpt.onlinelibrary.wiley.com/doi/abs/10.1002/psp4.12161

146. Jacobs ML, O'Brien SM, Jacobs JP, Mavroudis C, Lacour-Gayet F, Pasquali SK, et al. An empirically based tool for analyzing morbidity associated with operations for congenital heart disease. The Journal of Thoracic and Cardiovascular Surgery [Internet] 2013 [cited 2021 Apr 11];145:1046–1057.e1. Available from: https://linkinghub.elsevier.com/retrieve/pii/S0022522312007076

147. Anderson BJ, Holford NHG. Tips and traps analyzing pediatric PK data. Pediatric Anesthesia [Internet] 2011 [cited 2021 Aug 11];21:222–37. Available from: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1460-9592.2011.03536.x

148. Back H, Lee JB, Han N, Goo S, Jung E, Kim J, et al. Application of Size and Maturation Functions to Population Pharmacokinetic Modeling of Pediatric Patients. Pharmaceutics [Internet] 2019 [cited 2021 Aug 11];11:259. Available from: https://www.mdpi.com/1999-4923/11/6/259

149. Holford N, Heo Y-A, Anderson B. A Pharmacokinetic Standard for Babies and Adults. Journal of Pharmaceutical Sciences [Internet] 2013 [cited 2021 Jan 8];102:2941–52. Available from: https://linkinghub.elsevier.com/retrieve/pii/S0022354915309357

150. Ding J, Wang Y, Lin W, Wang C, Zhao L, Li X, et al. A Population Pharmacokinetic Model of Valproic Acid in Pediatric Patients with Epilepsy: A Non-Linear Pharmacokinetic Model Based on Protein-Binding Saturation. Clinical Pharmacokinetics [Internet] 2015 [cited 2021 Aug 12];54:305–17. Available from: http://link.springer.com/10.1007/s40262-014-0212-8

151. Mizuno T, Dong M, Taylor ZL, Ramsey LB, Vinks AA. Clinical implementation of pharmacogenetics and model-informed precision dosing to improve patient care. British Journal of Clinical Pharmacology [Internet] 2020 [cited 2021 Apr 16];Available from: http://bpspubs.onlinelibrary.wiley.com/doi/abs/10.1111/bcp.14426

152. Vinks AA, Punt NC, Menke F, Kirkendall E, Butler D, Duggan TJ, et al. Electronic Health Record–Embedded Decision Support Platform for Morphine Precision Dosing in Neonates. Clinical Pharmacology & Therapeutics [Internet] 2020 [cited 2020 Aug 13];107:186–94. Available from: http://ascpt.onlinelibrary.wiley.com/doi/abs/10.1002/cpt.1684

153. Baron KT. Mrgsolve: Simulate from ODE-Based Models [Internet]. 2020;Available from: https://CRAN.R-project.org/package=mrgsolve

154. Bauer RJ, Guzy S, Ng C. A survey of population analysis methods and software for complex pharmacokinetic and pharmacodynamic models with examples. The AAPS Journal [Internet] 2007 [cited 2020 Mar 4];9:E60–83. Available from: http://link.springer.com/10.1208/aapsj0901007

155. Dartois C, Lemenuel-Diot A, Laveille C, Tranchand B, Tod M, Girard P. Evaluation of Uncertainty Parameters Estimated by Different Population PK Software and Methods. Journal of Pharmacokinetics and Pharmacodynamics [Internet] 2007 [cited 2022 Mar 16];34:289–311. Available from: http://link.springer.com/10.1007/s10928-006-9046-9