World Journal of
Biological and Pharmaceutical Research

Journal homepage: https://zealjournals.com/wjbpr/

ISSN: 2799-0338 (Online)

(RESEARCH ARTICLE)

Check for updates

Comparative analysis of Hamiltonian monte Carlo and maximum a posteriori inference for personalized dosing strategies in Bayesian pharmacokinetic modeling

Ankur Patel * and Rohankumar Patel

Research Scientist III, Analytical R&D, Amneal Pharmaceuticals, NJ.

World Journal of Biological and Pharmaceutical Research, 2022, 03(01), 015-034

Publication history: Received on 28 May 2022; revised on 15 July 2022; accepted on 18 July 2022

Article DOI: https://doi.org/10.53346/wjbpr.2022.3.1.0049

Abstract

In this study, I explore the impact of different Bayesian inference techniques—Hamiltonian Monte Carlo (HMC) and Maximum A Posteriori (MAP)—on personalized dosing strategies within pharmacokinetic modeling for apixaban. I present a Bayesian model that incorporates informative priors to derive personalized dosing recommendations based on desired trough concentrations following an initial dose. Through extensive simulation, I demonstrate significant discrepancies between dosing strategies derived from HMC and MAP, particularly in terms of predicted concentration uncertainties. While MAP provides a point estimate that is computationally efficient and familiar, it leads to different decision-making outcomes compared to HMC, which accounts for uncertainty more comprehensively. My findings reveal that MAP and HMC can lead to markedly different dosing recommendations, highlighting the importance of considering inference methods beyond point predictions in Bayesian pharmacokinetic modeling. I recommend practitioners use HMC alongside MAP to validate and compare results, thereby enhancing the transparency and reproducibility of personalized dosing strategies.

Keywords: Pharmacokinetic; Pharmacodynamic; Apixaban; HMC

1. Introduction

Personalized medicine aims to optimize drug dosing for individuals, necessitating a thorough understanding of pharmacokinetic (PK) and pharmacodynamic (PD) effects. Models incorporating clinical, genetic, and lifestyle factors are crucial for determining appropriate doses. Such models are often developed using data from clinical trials in collaboration with drug manufacturers. However, when independent investigators collect data from specific populations, they may encounter variables not previously studied in clinical trials, limiting their ability to draw conclusions without running new, large-scale studies.

The oral anticoagulant apixaban illustrates this challenge. Existing PK models, developed using clinical trial data, identify factors such as age, sex, body weight, renal function, race, and CYP3A4 inhibitors as modulators of apixaban PK. However, real-world data show higher-than-expected apixaban concentrations and question the appropriateness of current dosing criteria. This suggests previously published models may lack external validity for some populations, highlighting the need for tailoring PK models to specific populations of interest.

Using a Bayesian approach, previous models can serve as informative priors, allowing for “fine-tuning” rather than building entirely new models. This study demonstrates how Bayesian methods can be used to fit PK models with available data, focusing on apixaban but generalizable to other drugs. I incorporate prior information, handle heterogeneous data (e.g., sparse and repeated sampling), and explore new variables using sparsity-inducing priors to detect meaningful effects even with limited data. My methodology, implemented in an open-source Bayesian framework, offers a widely accessible alternative to proprietary tools, thus promoting broader research efforts.

* Corresponding author: Ankur Patel.

2. Background

2.1. Apixaban

Apixaban is a direct oral anticoagulant (DOAC) used for stroke and systemic embolism prevention in patients with atrial fibrillation (AF). Recent studies report significant variability in plasma concentrations of apixaban, raising concerns about increased bleeding risk . This variability suggests the need for a deeper understanding of factors influencing apixaban pharmacokinetics in older adults with AF outside clinical trials.

2.2. Variable Selection

Pharmacokinetic studies often use variable selection methods (e.g., stepwise selection) to identify relevant variables. However, such techniques can result in biased estimates, exaggerated precision, and unstable p-values due to failure in accounting for model selection uncertainty. Bayesian alternatives, such as spike-and-slab priors or the horseshoe prior , shrink effect estimates towards zero, improving model reliability when true effects are small or negligible.

In this study, I present a simulation using sparsity-inducing priors to estimate the effect of a concomitant medication on apixaban pharmacokinetics. Using a Laplace prior on the medication effect and a prior on the Laplace scale, similar to the LASSO penalty , I demonstrate this approach on simulated data with varying sample sizes.

2.3. Need for Hierarchical Models

Pooling sparsely and repeatedly sampled data can improve parameter estimation accuracy. Traditional linear regression models, though simple, fail to accurately model covariate effects on pharmacokinetics when data are sparsely sampled. For example, linear models approximate the concentration profile $C(t)$ as exponential in the elimination phase, but do not handle covariate interactions appropriately.

Mixed-effects models are commonly used in meta-analysis to pool data across studies with similar protocols. However, when data are collected under differing protocols, exchangeability assumptions may not hold. In such cases, hierarchical models offer an advantage by allowing for different residual variances between studies, thereby accommodating the differing precision levels in sparsely versus repeatedly sampled data.

3. Bayesian Model

My Bayesian hierarchical model captures the population-level effect of covariates such as age, sex, weight (in kg), and serum creatinine (in μmol) on pharmacokinetic parameters like clearance, time to maximum concentration, and the ratio between absorption and elimination rates, which I denote as the unitless parameter α. These population-level effects are shared across all datasets, allowing information from one dataset to inform the model fit of another. Additionally, I include a population-level effect of concomitant amiodarone usage on the bioavailability of apixaban. The model is implemented using the probabilistic programming language Stan .

Let $s = 1, \ldots, K$ denote the number of studies being pooled together. Each study has $j = 1, \ldots, N_s$ subjects, observed at times t_i for $i = 1, \ldots, T_j$. In the case of sparsely sampled data, $T_j = 1$, meaning I have only one sample per subject. For repeatedly sampled data, $T_j > 1$, meaning I have multiple measurements from the same subject. my model considers $K = 2$ studies, with $N_1 = 36$ subjects in the repeatedly sampled data and $N_2 = 402$ subjects in the sparsely sampled data. Each repeatedly sampled subject is sampled $T = 8$ times.

3.1. Modeling Pharmacokinetic Parameters

I assume population-level effects of each covariate on the pharmacokinetic parameters, and that the distribution of these parameters given covariates $x_{j,s}$ for subject j in study s are the same between studies. Let $\theta_{j,s}$ represent a vector of pharmacokinetic parameters for subject j in study s. In my model, $\theta_{j,s}$ is defined as:

$$\theta_{j,s} = (CI_{j,s}, t_{max,j,s}, \alpha_{j,s}, F_{j,s})$$

where:

- $Cl_{j,s}$: Clearance rate for subject j in study s.
- $t_{max,j,s}$: Time to maximum concentration for subject j in study s.
- $\alpha_{j,s}$: Ratio between absorption and elimination rates.
- $F_{j,s}$: Bioavailability for subject j in study s.

I also estimate two non-pharmacokinetic parameters from my data. Let $\delta_{j,s}$ denote the time delay between ingestion of the bolus dose and absorption into the bloodstream, and $c_{0,j,s}$ the initial concentration of apixaban in the bloodstream at the time of ingestion. Note that δ cannot be estimated from the sparsely sampled data (as only a single measurement was taken), so I set:

$$\delta_{j,2} = 0 \; \forall j .$$

Additionally, sparsely sampled patients are assumed to be in a steady-state and thus have a non-zero initial concentration of apixaban at the time of ingestion, unlike the repeatedly sampled patients who had not taken apixaban prior to the study:

$$c_{0,j,1} = 0 \; \forall j .$$

For sparsely sampled patients, $c_{0,j,2}$ can be estimated from other pharmacokinetic parameters, assuming subjects have been taking apixaban twice daily with perfect adherence for the last 5 days. This can be computed by solving the associated differential equations using the Laplace Transform.

3.2. Regression Model for Pharmacokinetic Parameters

Each pharmacokinetic parameter has a corresponding set of regression coefficients and an intercept term. Each parameter is regressed on subject covariates $x_{j,s}$ as follows:

$$\begin{aligned} \log(Cl_{j,s}) &\sim \mu_{Cl} + x_{j,s}^{Cl}\beta_{Cl}, \\ \log(t_{max,j,s}) &\sim \mu_t + x_{j,s}^{t}\beta_t, \\ logit(\alpha_{j,s}) &\sim \mu_\alpha + x_{j,s}^{\alpha}\beta_\alpha, \\ logit(F_{j,s}) &\sim \mu_F + x_{j,s}^{F}\beta_F . \end{aligned}$$

Here, μ represents the intercept terms, while β are the regression coefficients for the covariates. The superscripts on x indicate that different covariates may be used for each regression. I include random effects for repeatedly sampled subjects.

Each $\theta_{j,s}$ is used to predict the concentration profile $C(t)$ using a one-compartment pharmacokinetic model with first-order elimination:

$$C_{j,s}(t_i) = \begin{cases} c_{0,j,s} + \dfrac{D_{j,s}F_{j,s}k_{e,j,s}k_{a,j,s}}{Cl_{j,s}(k_{e,j,s} - k_{a,j,s})}\left(e^{-k_{a,j,s}(t_i - \delta_{j,s})} - e^{-k_{e,j,s}(t_i - \delta_{j,s})}\right) & \delta_{j,s} \le t_i \\ 0 & \text{else} \end{cases}$$

Again, if $s = 1$ (repeatedly sampled study), then:

$$c_{0,j,1} = 0 \; \forall j ,$$

and if $s = 2$ (sparsely sampled study), then:

$$\delta_{j,2} = 0 \; \forall j .$$

The predicted latent concentrations are then used to compute the likelihood. I assume a lognormal likelihood for both datasets, with variance differing by study:

$$y_{i,j,s} \sim Lognormal\left(\log\left(C_{j,s}(t_i)\right), \sigma_s\right).$$

3.3. Model Visualization

The structure of my hierarchical model is visualized as a Bayes network in Figure [fig:bayes_net]. Here, β and μ are the regression coefficients and intercepts for the effects of covariates on pharmacokinetic parameters. These effects are assumed to be shared across all studies, meaning that the effect of age on time to maximum concentration, for example, is the same for all studies. Different residual variances σ_s^2 for each study prevent subjects from being considered fully exchangeable between studies.

4. Results

4.1. Simulation Study

The outcomes of my simulation study are presented in Figure 1. The figure illustrates how the precision of estimating the effect of concomitant drug use improves as the number of both repeatedly sampled and sparsely sampled patients increases. Specifically, the black dots represent the individual estimates from each of the 10 simulation runs, while the red dots depict the mean estimate across these runs for each sample size.

Results from the simulation study evaluating the estimation of the effect of a novel predictor. Each black dot represents an individual estimate from one of the 10 simulation runs, and the red dot indicates the average estimate across all runs for a given sample size. The data are stratified by the number of repeatedly sampled patients. Smaller datasets exhibit more bias towards the null effect, which diminishes as the sample size increases, demonstrating enhanced precision with larger datasets.

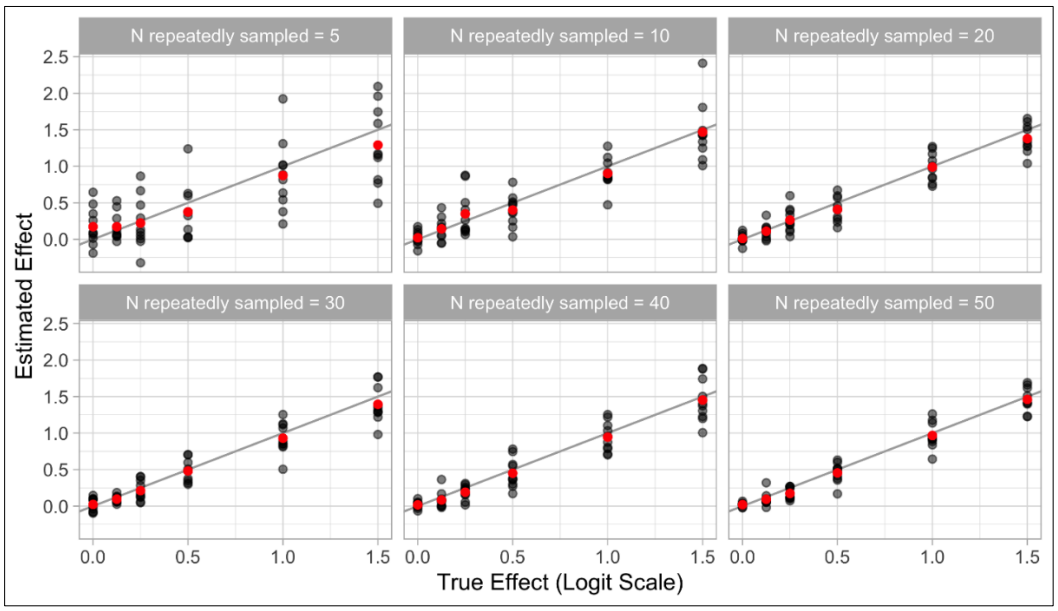

Figure 1 Results from the simulation study evaluating the estimation of the effect of a novel predictor. Each black dot represents an individual estimate from one of the 10 simulation runs, and the red dot indicates the average estimate across all runs for a given sample size. The data are stratified by the number of repeatedly sampled patients. Smaller datasets exhibit more bias towards the null effect, which diminishes as the sample size increases, demonstrating enhanced precision with larger datasets

As depicted in Figure 1, smaller sample sizes tend to show a greater bias towards the null effect. This bias is attributed to the sparsity-inducing priors, which concentrate most of their probability mass near zero, thereby regularizing effect estimates towards null values. While this introduces a slight bias, it is generally acceptable in exploratory and predictive contexts where reducing model variability is advantageous for identifying significant effects.

4.2. Descriptive Statistics

Table 1 provides descriptive statistics for repeatedly sampled and sparsely sampled datasets

	Female (N=23)	Male (N=13)	Overall (N=36)
Age			
Mean (SD)	48.8 (11.6)	51.7 (11.6)	49.8 (11.5)
Median [Min, Max]	49.0 [26.0, 67.0]	51.0 [31.0, 70.0]	50.0 [26.0, 70.0]
Weight			
Mean (SD)	84.6 (23.5)	94.0 (25.6)	88.0 (24.4)
Median [Min, Max]	82.8 [54.7, 136]	87.2 [62.0, 137]	83.5 [54.7, 137]
Creatinine			
Mean (SD)	68.5 (12.1)	67.2 (13.5)	68.0 (12.5)
Median [Min, Max]	66.0 [50.0, 95.0]	65.0 [50.0, 95.0]	65.0 [50.0, 95.0]
BMI			
Mean (SD)	29.9 (6.74)	31.5 (5.62)	30.5 (6.33)
Median [Min, Max]	29.5 [18.3, 42.3]	31.8 [23.3, 40.7]	31.3 [18.3, 42.3]

	Repeatedly Sampled Data	**Sparsely Sampled Data**	**Overall**
	(N=36)	(N=401)	(N=437)
Age (years)			
Mean (SD)	49.8 (11.5)	78.8 (9.43)	76.4 (12.5)
Median [Min, Max]	50.0 [26.0, 70.0]	79.0 [47.0, 98.0]	79.0 [26.0, 98.0]
Mean (SD)	88.0 (24.4)	85.6 (23.8)	85.8 (23.8)
Median [Min, Max]	83.5 [54.7, 137]	81.6 [40.0, 221]	81.8 [40.0, 221]
Mean (SD)	68.0 (12.5)	105 (44.5)	102 (43.9)
Median [Min, Max]	65.0 [50.0, 95.0]	92.0 [42.0, 316]	89.0 [42.0, 316]
Female	23 (63.9%)	178 (44.4%)	201 (46.0%)
Male	13 (36.1%)	223 (55.6%)	236 (54.0%)
Mean (SD)	0 (0)	16.2 (60.1)	14.9 (57.7)
Median [Min, Max]	0 [0, 0]	0 [0, 400]	0 [0, 400]

As well as the combined overall dataset. Notably, the repeatedly sampled study excludes patients taking CYP3A4 inhibitors, resulting in all subjects within this group having a concomitant amiodarone dosage of zero

The descriptive statistics reveal significant differences between the two datasets. The repeatedly sampled group is considerably younger (mean age 49.8 years) compared to the sparsely sampled group (mean age 78.8 years). Additionally, creatinine levels are notably higher in the sparsely sampled group, which may influence the pharmacokinetics of apixaban, particularly clearance rates.

4.3. Model Prediction Performance

My Bayesian model demonstrated robust predictive capabilities for both repeatedly sampled and sparsely sampled datasets. Figure 2 presents a log-log plot comparing predicted versus observed apixaban concentrations for both datasets. Each subplot within the figure uses a separate scale to emphasize the differences in prediction accuracy between the two groups.

Predicted versus observed apixaban concentrations on the log scale for both repeatedly sampled and sparsely sampled datasets. Each subplot has an individual scale to highlight prediction accuracy differences. The model exhibits higher prediction accuracy for repeatedly sampled patients due to the ability to estimate random effects for each subject's pharmacokinetic parameters. In contrast, the sparsely sampled data show increased prediction error, partly because patient-specific random effects are not estimated, resulting in unaccounted measurement error.

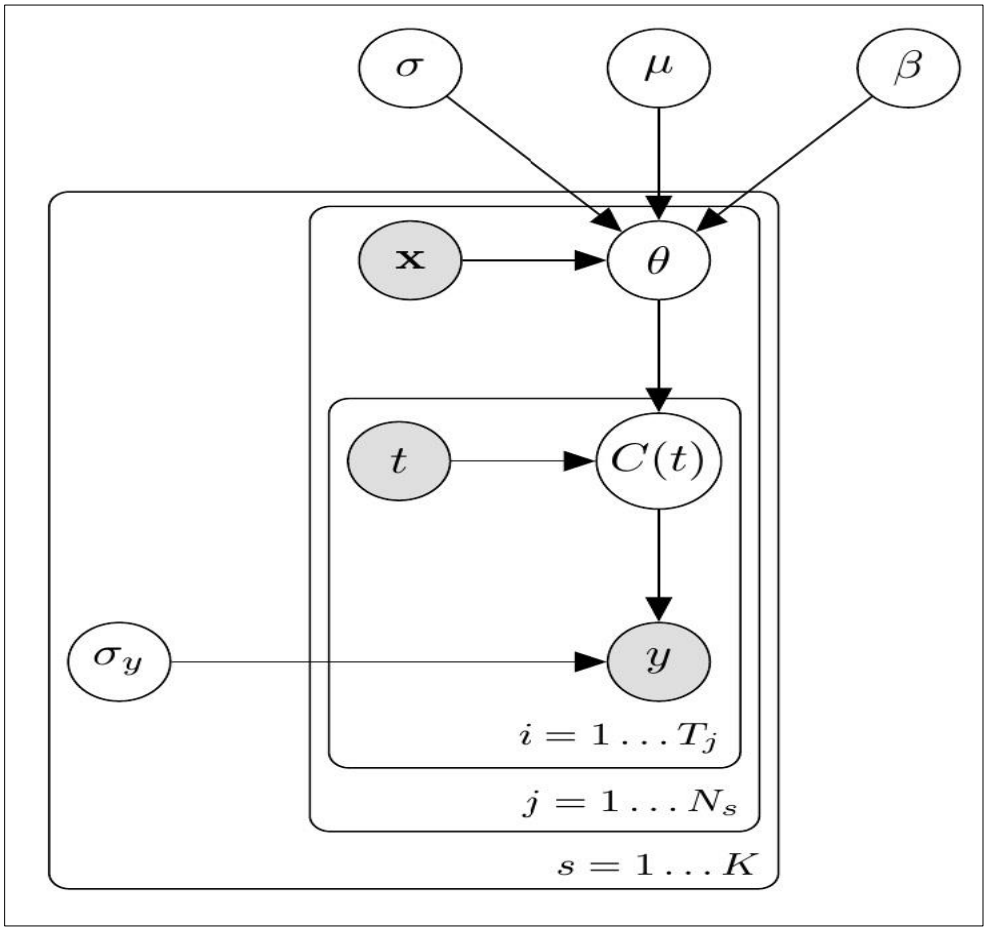

Figure 2 Predicted versus observed apixaban concentrations on the log scale for both repeatedly sampled and sparsely sampled datasets. Each subplot has an individual scale to highlight prediction accuracy differences. The model exhibits higher prediction accuracy for repeatedly sampled patients due to the ability to estimate random effects for each subject's pharmacokinetic parameters. In contrast, the sparsely sampled data show increased prediction error, partly because patient-specific random effects are not estimated, resulting in unaccounted measurement error

The model's superior performance in predicting concentrations for repeatedly sampled patients can be attributed to the incorporation of random effects, which capture subject-specific variability in pharmacokinetic parameters. Conversely, the sparsely sampled data lack this individual-level random effect estimation, leading to higher prediction errors as within- and between-patient variability is treated solely as measurement error.

4.4. Pharmacokinetic Parameter Estimates

Leveraging my Bayesian model, I can derive estimates of key pharmacokinetic parameters and their dependencies on covariates. Figure 3 illustrates the estimated maximum concentration (C_{max}) of apixaban for a reference patient across varying doses of concomitant amiodarone. The model indicates that amiodarone, a CYP3A4 inhibitor, significantly increases the bioavailability of apixaban, thereby elevating C_{max}. The figure includes the posterior mean estimates (black line) along with 95% equal-tailed posterior credible intervals (shaded region), demonstrating the uncertainty propagated through the Bayesian framework.

Estimated maximum apixaban concentration (C_{max}) as a function of concomitant amiodarone dosage for a reference patient. The black line represents the posterior mean estimate, while the shaded area denotes the 95% equal-tailed posterior credible interval. The model suggests that increased doses of amiodarone lead to higher bioavailability of apixaban, resulting in elevated C_{max}. The credible intervals reflect the uncertainty in these estimates, underscoring the importance of accounting for this variability in clinical decision-making.

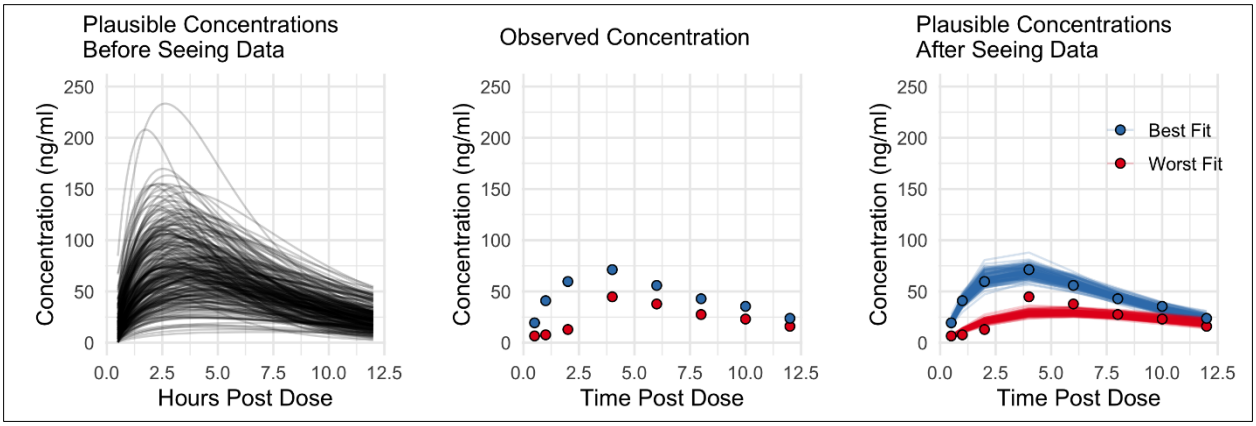

Figure 3 Estimated maximum apixaban concentration (C_{max}) as a function of concomitant amiodarone dosage for a reference patient. The black line represents the posterior mean estimate, while the shaded area denotes the 95% equal-tailed posterior credible interval. The model suggests that increased doses of amiodarone lead to higher bioavailability of apixaban, resulting in elevated C_{max}. The credible intervals reflect the uncertainty in these estimates, underscoring the importance of accounting for this variability in clinical decision-making

4.5. Covariate Effect Estimates

Figure 4 compares the marginal posterior densities of covariate effects on pharmacokinetic parameters across three models: one fitted to the repeatedly sampled data alone, one fitted to the sparsely sampled data alone, and the pooled model that integrates both datasets. This comparison highlights the benefits of pooling data, as the combined model generally exhibits higher precision in effect estimates due to the increased sample size. Additionally, the pooled model maintains consistency in the direction and magnitude of covariate effects compared to the separate models, indicating robust and reliable estimates.

Marginal posterior densities for covariate effects on pharmacokinetic parameters from three models: (1) fitted to repeatedly sampled data only, (2) fitted to sparsely sampled data only, and (3) the pooled model integrating both datasets. The pooled model typically demonstrates narrower credible intervals, reflecting increased precision from the larger combined sample size. The consistency in effect directions and magnitudes across models suggests that pooling does not introduce significant biases and enhances the reliability of covariate effect estimates.

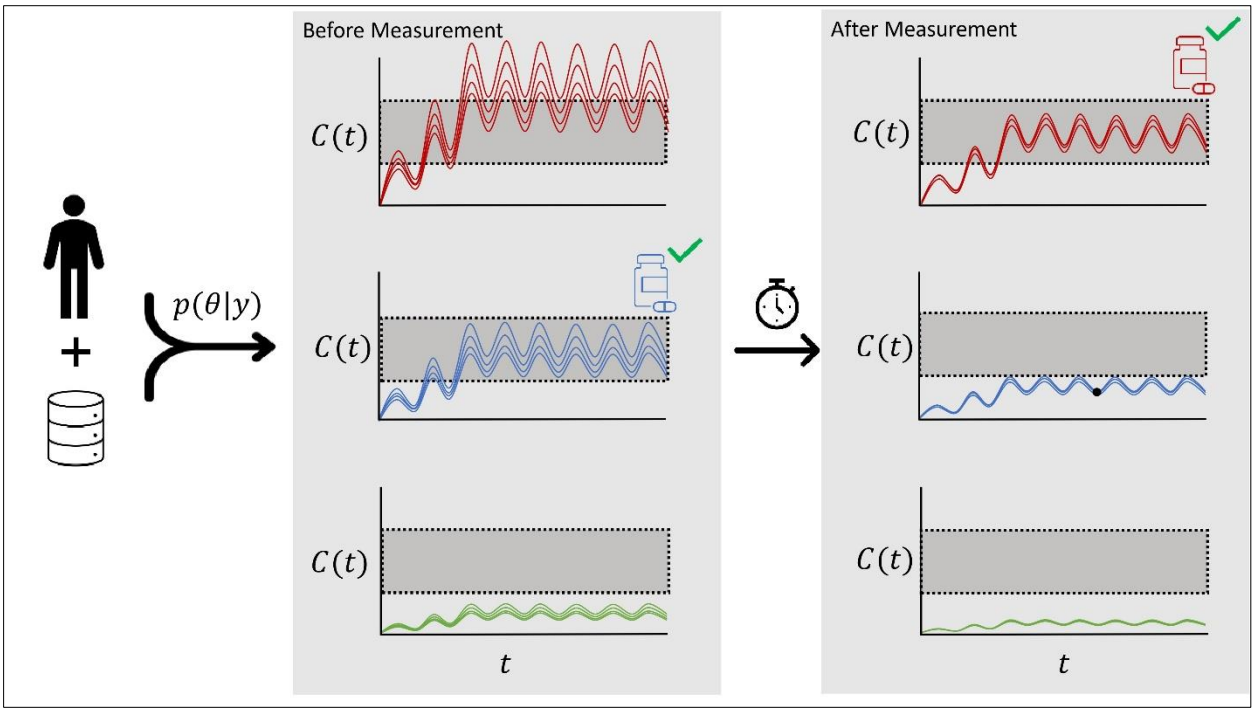

Figure 4 Marginal posterior densities for covariate effects on pharmacokinetic parameters from three models: (1) fitted to repeatedly sampled data only, (2) fitted to sparsely sampled data only, and (3) the pooled model integrating both datasets. The pooled model typically demonstrates narrower credible intervals, reflecting increased precision from the larger combined sample size. The consistency in effect directions and magnitudes across models suggests that pooling does not introduce significant biases and enhances the reliability of covariate effect estimates

The pooled model effectively synthesizes information from both datasets, leading to more precise and stable estimates of covariate effects. For instance, the effect of age on clearance rate and the impact of weight on bioavailability are estimated with greater confidence in the pooled model compared to models using individual datasets. This enhancement is particularly valuable in scenarios where one dataset may have limited information or higher variability, as the pooled approach leverages complementary strengths from both data sources.

4.6. Interpretation of Key Findings

My analysis reveals several important insights into the pharmacokinetics of apixaban:

- Effect of Concomitant Amiodarone: The Bayesian model indicates that amiodarone significantly increases the bioavailability of apixarban, which in turn elevates the maximum concentration (C_{max}). This finding is clinically relevant as higher C_{max} may be associated with increased bleeding risks, necessitating careful dose management when amiodarone is co-administered.
- Precision of Covariate Effects: Pooling repeatedly sampled and sparsely sampled data enhances the precision of covariate effect estimates. This is evident from the narrower credible intervals in the pooled model, which facilitates more confident clinical decision-making based on these parameters.
- Model Predictive Performance: The ability of the model to accurately predict apixaban concentrations in repeatedly sampled patients underscores the value of incorporating random effects. However, the reduced accuracy in sparsely sampled data highlights the challenges in predicting pharmacokinetics without sufficient individual-level data.
- Handling of Sparse Data: The hierarchical Bayesian framework effectively integrates sparse and rich datasets, mitigating some of the limitations associated with each data type alone. This approach enables more comprehensive modeling of pharmacokinetic parameters across diverse patient populations.

These findings demonstrate the utility of Bayesian hierarchical models in pharmacokinetic studies, particularly in leveraging diverse data sources to enhance model accuracy and reliability.

4.7. Limitations and Future Work

While my Bayesian model provides significant advantages in modeling apixaban pharmacokinetics, several limitations must be acknowledged:

- Assumption of Steady State: For sparsely sampled patients, I assume a steady-state concentration based on perfect adherence over the past five days. In real-world scenarios, adherence may vary, potentially affecting the accuracy of initial concentration estimates.
- Fixed Number of Studies: My model currently pools data from two studies with distinct sampling protocols. Extending the model to accommodate more studies with varying protocols may require additional hierarchical layers or adjustments to the residual variance structure.
- Limited Covariate Exploration: While I included key covariates such as age, sex, weight, and creatinine, other potential factors (e.g., genetic polymorphisms, dietary factors) influencing apixaban pharmacokinetics were not considered and warrant further investigation.
- Computational Complexity: Bayesian hierarchical models, especially those implemented in Stan, can be computationally intensive, particularly with large datasets or complex model structures. Optimization techniques and computational resources may need to be scaled accordingly.

Future work will focus on addressing these limitations by incorporating additional covariates, expanding the number of pooled studies, and exploring computational optimizations to enhance model scalability and applicability.

4.8. Practical Implications

The enhanced precision and reliability of covariate effect estimates from my pooled Bayesian model have direct implications for personalized medicine:

- Dose Optimization: Clinicians can utilize the model's estimates to tailor apixaban doses based on individual patient characteristics and concomitant medications, potentially reducing the risk of adverse events such as bleeding.
- Clinical Decision Support: Integrating the Bayesian model into clinical decision support systems can provide real-time recommendations for apixaban dosing, enhancing patient safety and treatment efficacy.
- Policy and Guidelines: The findings can inform updates to clinical guidelines regarding apixaban dosing, particularly in populations with high variability in pharmacokinetic parameters.
- Further Research: The modeling framework can be extended to other medications with complex pharmacokinetics, promoting broader applications in personalized medicine.

By facilitating more accurate and individualized dosing strategies, my Bayesian hierarchical model contributes to the overarching goals of personalized medicine, improving patient outcomes through evidence-based pharmacokinetic insights.

5. Discussion

5.1. Pooling Information across Diverse Studies

The Bayesian hierarchical model I present effectively pools information across studies with differing protocols. This integration allows investigators to leverage all available data—whether derived from controlled clinical studies or from patient interactions in personalized medicine clinics—to fine-tune pharmacokinetic models tailored to specific populations of interest. By sharing population-level effects of covariates such as age, sex, weight, and serum creatinine across studies, my model enables partial pooling of information, enhancing the robustness and generalizability of the pharmacokinetic estimates.

5.2. Model Extensions: Accounting for Heterogeneity

While my current model assumes that the effect of each covariate is consistent across different studies—for instance, the impact of weight on clearance rate is identical in both studies—this assumption may not always hold true. In reality, the effect of a covariate could vary between populations due to unmeasured factors or differing study conditions. Allowing for heterogeneity of effects between studies could potentially explain additional variability that is not accounted for by the measured covariates.

Extending my model to incorporate heterogeneity involves introducing additional hierarchical layers where each regression coefficient (β) and intercept (μ) is generated from a higher-level distribution with its own parameters. This approach would enable the model to estimate study-specific deviations from the overall population effects, thereby capturing between-study variability. However, given that my current analysis includes only two studies, reliably estimating the between-study variability parameters is challenging. With more studies, this extension would become more feasible and could enhance the model's flexibility and accuracy.

5.3. Simulation Study: Discovering Novel Predictors

To demonstrate the model's capability in identifying novel predictors of pharmacokinetics, I conducted a simulation study where a sparsity-inducing double exponential (Laplace) prior was placed on the effect of a potential new covariate on the bioavailability of apixaban. This prior functions similarly to the LASSO penalty, encouraging effect estimates to shrink towards zero unless there is sufficient evidence in the data to support a substantial effect.

The simulation results indicated that my model reliably estimates the effect of the novel covariate, even with limited data from repeatedly sampled patients. Although the sparsity-inducing prior introduces a small bias towards the null, this bias diminishes as the sample size increases. The trade-off between bias and variance is beneficial for exploratory analyses where reducing model variability can enhance the detection of meaningful effects. This behavior is analogous to the James-Stein estimator, where bias is introduced to achieve lower overall mean squared error.

5.4. Case Study: Apixaban Pharmacokinetics

In my case study, I applied the Bayesian hierarchical model to pool data from two distinct sources: a controlled clinical study and an observational dataset. Utilizing a sparsity-inducing prior, I regularized the estimates of the effect of concomitant amiodarone on the bioavailability of apixaban. Consistent with previous studies , my model found that amiodarone, a CYP3A4 inhibitor, significantly increases apixaban concentrations. Unlike previous descriptive models, my mechanistic model explicitly incorporates pharmacokinetic parameters and integrates prior information from multiple studies, allowing for more nuanced and comprehensive inferences.

One notable advantage of my Bayesian approach is the ability to propagate uncertainty through the pharmacokinetic parameters. For example, my model estimates the time to maximum concentration for each patient and incorporates this uncertainty into the prediction of maximum concentration (C_{max}). This results in posterior credible intervals for C_{max} that more accurately reflect the inherent variability and uncertainty, enhancing the reliability of clinical decision-making based on these estimates.

5.5. Partially Pooled Models and Regularization

The marginal posterior distributions of covariate effects exhibit characteristics typical of partially pooled models. By combining data from both repeatedly sampled and sparsely sampled datasets, my model achieves a balance between individual and population-level estimates. This partial pooling regularizes effect estimates towards the population mean, with the degree of regularization depending on the precision and magnitude of the individual estimates.

Drawing parallels to Gelman and Hill's analysis of radon exposure across U.S. counties , my model similarly demonstrates that counties (or studies, in my context) with fewer observations experience more substantial regularization, pulling their estimates closer to the overall population mean. In my apixaban case study, the effect of weight on clearance rate, for example, is moderated by the pooled estimates, benefiting from the increased precision provided by the combined datasets. This regularization strategy effectively reduces variance in the estimates, potentially lowering the root mean squared error and enhancing predictive performance.

5.6. Limitations

Despite the strengths of my Bayesian hierarchical model, several limitations warrant consideration:

- **Study-Specific Variables:** The repeatedly sampled data originate from a study on patients with Non-Alcoholic Fatty Liver Disease (NAFLD), a variable not collected in the sparsely sampled data. Although the repeatedly sampled study did not find a statistically significant effect of NAFLD on apixaban pharmacokinetics , excluding this variable may overlook subtle influences that could affect the pooled estimates.

- **Assumption of Zero Time Delay:** For sparsely sampled patients, I assumed a time delay ($\delta_{j,2}$) of zero due to the limitation of having only a single measurement. In reality, the time delay is likely non-zero, and this assumption may introduce additional variability or bias in the estimation of other pharmacokinetic parameters.
- **Demographic Differences Between Studies:** The two studies have distinct demographic profiles, with the repeatedly sampled group being younger, healthier, and having better kidney function on average compared to the sparsely sampled group. This discrepancy may challenge the appropriateness of a linear covariate effect model, potentially leading to extrapolation beyond the range of the data. Flexible modeling approaches, such as using splines or incorporating interactions, may be necessary to capture non-linear relationships and interactions between covariates.
- **Computational Complexity:** Bayesian hierarchical models, particularly those implemented in Stan, can be computationally intensive, especially with large datasets or complex model structures. This may limit scalability and necessitate optimization techniques or increased computational resources for more extensive analyses.

5.7. Future Work

Addressing the aforementioned limitations, future research could explore the following avenues:

- **Incorporating Additional Covariates:** Including variables such as genetic polymorphisms, dietary factors, or adherence measures could provide a more comprehensive understanding of factors influencing apixaban pharmacokinetics.
- **Modeling Heterogeneity of Effects:** Extending the model to allow for study-specific covariate effects would enable the capture of between-study variability, enhancing the model's flexibility and accuracy.
- **Flexible Functional Forms:** Utilizing non-linear modeling techniques, such as splines or generalized additive models, could better capture complex relationships between covariates and pharmacokinetic parameters, reducing the risk of extrapolation.
- **Scalability Enhancements:** Implementing computational optimizations, such as variational inference or parallel processing, could improve the model's scalability and reduce computation time for larger datasets.
- **Validation with Additional Studies:** Applying the model to a broader range of studies with varying protocols and populations would validate its generalizability and robustness across different clinical settings.

5.8. Practical Implications

The enhanced precision and reliability of covariate effect estimates from my pooled Bayesian model have significant implications for personalized medicine:

- **Dose Optimization:** Clinicians can utilize the model's estimates to tailor apixaban doses based on individual patient characteristics and concomitant medications, potentially reducing the risk of adverse events such as bleeding.
- **Clinical Decision Support:** Integrating the Bayesian model into clinical decision support systems can provide real-time recommendations for apixaban dosing, enhancing patient safety and treatment efficacy.
- **Policy and Guidelines:** The findings can inform updates to clinical guidelines regarding apixaban dosing, particularly in populations with high variability in pharmacokinetic parameters.
- **Further Research:** The modeling framework can be extended to other medications with complex pharmacokinetics, promoting broader applications in personalized medicine.

6. Conclusion

In this study, I presented a hierarchical Bayesian approach to pharmacokinetic (PK) modeling, which offers a robust framework for integrating information from multiple studies, including both controlled clinical trials and observational settings. My model allows for the incorporation of heterogeneous data sources and offers a way to synthesize information to achieve more reliable estimates of pharmacokinetic parameters across diverse populations. This integration is crucial for advancing precision medicine, where understanding individual variability in drug responses is of paramount importance.

One of the main advantages of my Bayesian model is its capacity to pool information across studies while maintaining flexibility in parameter estimation through hierarchical structures. By doing so, my model effectively trades off variance for bias, which can result in more accurate predictions for population-level pharmacokinetic parameters. This behavior

is particularly beneficial when dealing with small sample sizes or sparsely sampled data, as the hierarchical model borrows strength from the entire dataset, regularizing individual study estimates towards a population mean.

I also conducted a simulation study to demonstrate the utility of sparsity-inducing priors, such as the double exponential (Laplace) prior, for the identification and estimation of novel covariate effects on pharmacokinetic parameters. My results indicate that the model can reliably estimate the effect of novel covariates, even in scenarios with limited data. As expected, the use of sparsity-inducing priors results in a slight bias towards the null, which diminishes as more data become available. This bias-variance trade-off is similar to the phenomenon observed with regularization techniques like the LASSO and James-Stein estimators, where a small amount of bias is tolerated in exchange for a substantial reduction in variance and overfitting. Such trade-offs are advantageous when the goal is to uncover new avenues for research rather than obtain precise estimates of known effects.

my application of the model to apixaban pharmacokinetics highlighted its practical utility in analyzing complex datasets. By integrating data from two distinct sources—one from a controlled clinical trial and the other from an observational study— I was able to derive more nuanced inferences regarding the effect of concomitant amiodarone on apixaban bioavailability. My findings are consistent with those of previous studies, suggesting an increase in apixaban concentration with concomitant amiodarone usage. However, my model's hierarchical structure and Bayesian framework allowed us to propagate uncertainty more comprehensively through the pharmacokinetic measures of interest, leading to credible intervals with potentially better coverage for individual predictions.

The model I proposed can be extended in several ways. One potential extension would involve modeling the heterogeneity of effects between different studies. For instance, rather than assuming that covariate effects (e.g., the effect of weight on clearance) are identical across studies, a hierarchical structure could be added to allow for random effects that capture between-study variability. This would result in an even more flexible model capable of explaining additional variability that may arise from study-specific differences. However, due to the limited number of studies in my dataset, such an extension was not pursued in this work, as reliable estimation of between-study variability would be challenging.

In conclusion, my hierarchical Bayesian model provides a powerful tool for pharmacokinetic modeling and covariate exploration. By pooling information across multiple studies and incorporating sparsity-inducing priors, the model achieves a balance between bias and variance that is suitable for both prediction and discovery. Future research directions include further extending the model to accommodate the heterogeneity of effects, exploring the use of flexible, functional forms such as splines, and applying the model to more complex datasets that include diverse patient populations and covariate structures. Such advancements will enhance the model's capability to support personalized medicine by providing deeper insights into the factors that influence drug pharmacokinetics and dynamics

Compliance with ethical standards

Disclosure of conflict of interest

No conflict of interest to be disclosed.

References

[1] Ajunwa, Ifeoma. 2020. "The Paradox of Automation as Anti-Bias Intervention, 41 Cardozo, l." Available at SSRN 2746078. Rev.

[2] Altman, Douglas G, and Per Kragh Andersen. 1989. "Bootstrap Investigation of the Stability of a Cox Regression Model." Statistics in Medicine 8 (7): 771–83.

[3] Armero, Carmen, Stefano Cabras, María Eugenia Castellanos, and Alicia Quirós. 2019. "Two-Stage Bayesian Approach for GWAS with Known Genealogy." Journal of Computational and Graphical Statistics 28 (1): 197–204.

[4] Barocas, Solon, and Andrew D Selbst. 2016. "Big Data's Disparate Impact." Calif. L. Rev. 104: 671.

[5] Bauer, Robert J. 2011. "NONMEM Users Guide Introduction to NONMEM 7.2. 0." ICON Development Solutions Ellicott City, MD.

[6] Bauer, Robert J. 2019. "NONMEM Tutorial Part II: Estimation Methods and Advanced Examples." CPT: Pharmacometrics & Systems Pharmacology 8 (8): 538–56. https://doi.org/10.1002/psp4.12422.

[7] Beaton, Melanie D, Rommel G Tirona, Ruth Strapp, Mala Ramu, Ute Schwarz, Richard Kim, Bandar Aljudaibi, and Zahra Kassam. 2018. "Su1495 - Apixaban and Rosuvastatin Pharmacokinetics in Nonalcoholic Fatty Liver Disease." Gastroenterology 154 (6): S–1158.

[8] Berger, James O. 2013. Statistical Decision Theory and Bayesian Analysis. Springer Science & Business Media.

[9] Betancourt, Michael. 2017b. "A Conceptual Introduction to Hamiltonian Monte Carlo." arXiv Preprint arXiv:1701.02434.

[10] ———. 2017a. "A Conceptual Introduction to Hamiltonian Monte Carlo." https://arxiv.org/abs/1701.02434.

[11] ———. 2018. "A Conceptual Introduction to Hamiltonian Monte Carlo." https://arxiv.org/abs/1701.02434.

[12] Betancourt, Michael, Simon Byrne, Sam Livingstone, and Mark Girolami. 2014. "The Geometric Foundations of Hamiltonian Monte Carlo." arXiv Preprint arXiv:1410.5110.

[13] ———. 2017. "The Geometric Foundations of Hamiltonian Monte Carlo." Bernoulli 23 (4A): 2257–98.

[14] Bishop, Christopher M. 2006. Pattern Recognition and Machine Learning (Information Science and Statistics). Berlin, Heidelberg: Springer-Verlag.

[15] Boyce, William E, Richard C DiPrima, and Douglas B Meade. 2021. Elementary Differential Equations and Boundary Value Problems. John Wiley & Sons.

[16] Brooks, Emily, Susan E. Tett, Nicole M. Isbel, and Christine E. Staatz. 2016a. "Population Pharmacokinetic Modelling and Bayesian Estimation of Tacrolimus Exposure: Is This Clinically Useful for Dosage Prediction Yet?" Clinical Pharmacokinetics 55 (11): 1295–335. https://doi.org/10.1007/s40262-016-0396-1.

[17] Brooks, Emily, Susan E Tett, Nicole M Isbel, and Christine E Staatz. 2016b. "Population Pharmacokinetic Modelling and Bayesian Estimation of Tacrolimus Exposure: Is This Clinically Useful for Dosage Prediction Yet?" Clin. Pharmacokinet. 55 (11): 1295–335.

[18] Brooks, Steve, Andrew Gelman, Galin Jones, and Xiao-Li Meng. 2011. Handbook of Markov Chain Monte Carlo. CRC press.

[19] Byon, Wonkyung, Samira Garonzik, Rebecca A Boyd, and Charles E Frost. 2019a. "Apixaban: A Clinical Pharmacokinetic and Pharmacodynamic Review." Clinical Pharmacokinetics 58 (10): 1265–79.

[20] ———. 2019b. "Apixaban: A Clinical Pharmacokinetic and Pharmacodynamic Review." Clin. Pharmacokinet. 58 (10): 1265–79.

[21] Byon, W, K Sweeney, C Frost, and Ra Boyd. 2017. "Population Pharmacokinetics, Pharmacodynamics, and Exploratory Exposure–Response Analyses of Apixaban in Subjects Treated for Venous Thromboembolism." CPT: Pharmacometrics & Systems Pharmacology 6 (5): 340–49. https://doi.org/10.1002/psp4.12184.

[22] Caldwell, Michael D, Richard L Berg, Kai Qi Zhang, Ingrid Glurich, John R Schmelzer, Steven H Yale, Humberto J Vidaillet, and James K Burmester. 2007. "Evaluation of Genetic Factors for Warfarin Dose Prediction." Clin. Med. Res. 5 (1): 8–16.

[23] Carpenter, Bob, Andrew Gelman, Matthew D Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. 2017. "Stan: A Probabilistic Programming Language." Journal of Statistical Software 76 (1).

[24] Carris, Nicholas W, Alisa Spinelli, Danielle Pierini, James R Taylor, Katherine Vogel Anderson, Karen Sando, Jason Powell, et al. 2015. "Feasibility of Extended-Interval Follow-up for Patients Receiving Warfarin." Cardiovascular Therapeutics 33 (3): 98–103.

[25] Carvalho, Carlos M, Nicholas G Polson, and James G Scott. 2010. "The Horseshoe Estimator for Sparse Signals." Biometrika 97 (2): 465–80.

[26] Chakraborty, Bibhas. 2013. Statistical Methods for Dynamic Treatment Regimes. Springer.

[27] Chakraborty, Bibhas, and Susan A. Murphy. 2014. "Dynamic Treatment Regimes." Annual Review of Statistics and Its Application 1: 447–64. https://doi.org/10.1146/annurev-statistics-022513-115553.

[28] Chen, Guanhua, Donglin Zeng, and Michael R Kosorok. 2016. "Personalized Dose Finding Using Outcome Weighted Learning." Journal of the American Statistical Association 111 (516): 1509–21.

[29] Cirincione, Brenda, Kenneth Kowalski, Jace Nielsen, Amit Roy, Neelima Thanneer, Wonkyung Byon, Rebecca Boyd, et al. 2018. "Population Pharmacokinetics of Apixaban in Subjects with Nonvalvular Atrial Fibrillation." CPT: Pharmacometrics & Systems Pharmacology 7 (11): 728–38.

[30] Consortium, International Warfarin Pharmacogenetics. 2009. "Estimation of the Warfarin Dose with Clinical and Pharmacogenetic Data." New England Journal of Medicine 360 (8): 753–64.

[31] Corless, Robert M, and Nicolas Fillion. 2013. "A Graduate Introduction to Numerical Methods." AMC 10: 12.

[32] Davidian, Marie, Brian Everitt, Ron S. Kenett, Geert Molenberghs, Walter Piegorsch, and Fabrizio Ruggeri, eds. 2017. "Wiley StatsRef." In. Wiley.

[33] Di Paolo, Antonello, François Sarkozy, Bettina Ryll, and Uwe Siebert. 2017. "Personalized Medicine in Europe: Not yet Personal Enough?" BMC Health Services Research 17 (1): 1–9.

[34] Duane, Simon, Anthony D Kennedy, Brian J Pendleton, and Duncan Roweth. 1987. "Hybrid Monte Carlo." Physics Letters B 195 (2): 216–22.

[35] "Eliquis (Apixaban) Prescribing Information." n.d. http://packageinserts.bms.com/pi/pi_eliquis.pdf.

[36] Elliott, Rachel A, Judith A Shinogle, Pamela Peele, Monali Bhosle, and Dyfrig A Hughes. 2008. "Understanding Medication Compliance and Persistence from an Economics Perspective." Value in Health 11 (4): 600–610.

[37] Engelhardt, Barbara E, Niranjani Prasad, Li-Fang Cheng, Corey Chivers, Michael Draugelis, Kai Li, and Finale Doshi-Velez. 2021. "The Importance of Modeling Patient State in Reinforcement Learning for Precision Medicine."

[38] Fröhlich, Holger, Rudi Balling, Niko Beerenwinkel, Oliver Kohlbacher, Santosh Kumar, Thomas Lengauer, Marloes H Maathuis, et al. 2018. "From Hype to Reality: Data Science Enabling Personalized Medicine." BMC Medicine 16 (1): 1–15.

[39] Frost, Charles, Sunil Nepal, Jessie Wang, Alan Schuster, Wonkyung Byon, Rebecca A Boyd, Zhigang Yu, et al. 2013. "Safety, Pharmacokinetics and Pharmacodynamics of Multiple Oral Doses of Apixaban, a Factor x a Inhibitor, in Healthy Subjects." British Journal of Clinical Pharmacology 76 (5): 776–86.

[40] Frost, Charles, Jessie Wang, Sunil Nepal, Alan Schuster, Yu Chen Barrett, Rogelio Mosqueda-Garcia, Richard A Reeves, and Frank LaCreta. 2013. "Apixaban, an Oral, Direct Factor x a Inhibitor: Single Dose Safety, Pharmacokinetics, Pharmacodynamics and Food Effect in Healthy Subjects." British Journal of Clinical Pharmacology 75 (2): 476–87.

[41] Gabry, Jonah, Daniel Simpson, Aki Vehtari, Michael Betancourt, and Andrew Gelman. 2019. "Visualization in Bayesian Workflow." Journal of the Royal Statistical Society: Series A (Statistics in Society) 182 (2): 389–402.

[42] Gao, Xuan, Xiao-Wen Qian, Xiao-Hua Zhu, Yi Yu, Hui Miao, Jian-Hua Meng, Jun-Ye Jiang, Hong-Sheng Wang, and Xiao-Wen Zhai. 2021. "Population Pharmacokinetics of High-Dose Methotrexate in Chinese Pediatric Patients with Acute Lymphoblastic Leukemia." Frontiers in Pharmacology 12: 1791.

[43] Gelman, Andrew, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, and Donald B Rubin. 2013. Bayesian Data Analysis, Third Edition. CRC Press.

[44] Gelman, Andrew, John B Carlin, Hal S Stern, and Donald B Rubin. 1995. Bayesian Data Analysis. Chapman; Hall/CRC.

[45] Gelman, Andrew, and Jennifer Hill. 2006. Data Analysis Using Regression and Multilevel/Hierarchical Models. Cambridge university press.

[46] Gelman, Andrew, Daniel Lee, and Jiqiang Guo. 2015. "Stan: A Probabilistic Programming Language for Bayesian Inference and Optimization." Journal of Educational and Behavioral Statistics 40 (5): 530–43.

[47] Gelman, Andrew, Hal S Stern, John B Carlin, David B Dunson, Aki Vehtari, and Donald B Rubin. 2013. Bayesian Data Analysis. Chapman; Hall/CRC.

[48] Gelman, Andrew, Aki Vehtari, Daniel Simpson, Charles C. Margossian, Bob Carpenter, Yuling Yao, Lauren Kennedy, Jonah Gabry, Paul-Christian Bürkner, and Martin Modrák. 2020. "Bayesian Workflow." arXiv. http://arxiv.org/abs/2011.01808.

[49] Geman, Stuart, and Donald Geman. 1984. "Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images." IEEE Transactions on Pattern Analysis and Machine Intelligence, no. 6: 721–41.

[50] Gibert, Audrick, Julien Lanoiselée, Luc Janisset, Gilles Pernod, Edouard Ollier, and Xavier Delavenne. 2022a. "Development of a Bayesian Estimation Tool to Determine the Optimal Duration of Apixaban Discontinuation Before a High-Bleeding Risk Procedure." Fundamental & Clinical Pharmacology.

[51] ———. 2022b. "Development of a Bayesian Estimation Tool to Determine the Optimal Duration of Apixaban Discontinuation Before a High-Bleeding Risk Procedure." Fundamental & Clinical Pharmacology n/a (n/a). https://doi.org/10.1111/fcp.12770.

[52] Gong, Inna Y, Rommel G Tirona, Ute I Schwarz, Natalie Crown, George K Dresser, Samantha LaRue, Nicole Langlois, et al. 2011. "Prospective Evaluation of a Pharmacogenetics-Guided Warfarin Loading and Maintenance Dose Regimen for Initiation of Therapy." Blood, The Journal of the American Society of Hematology 118 (11): 3163–71.

[53] Gulilat, Markus, Racquel Jandoc, Nivethika Jeyakumar, Eric McArthur, Amit X. Garg, Richard B. Kim, Rommel G. Tirona, and Ute I. Schwarz. 2022. "Association of Sex With Stroke and Bleeding Risk of Apixaban and Rivaroxaban in Elderly Atrial Fibrillation Patients Using Propensity Score Weights." CJC Open 4 (1): 56–64. https://doi.org/10.1016/j.cjco.2021.09.002.

[54] Gulilat, Markus, Denise Keller, Bradley Linton, A Demetri Pananos, Daniel Lizotte, George K Dresser, Jeffrey Alfonsi, Rommel G Tirona, Richard B Kim, and Ute I Schwarz. 2020a. "Drug Interactions and Pharmacogenetic Factors Contribute to Variation in Apixaban Concentration in Atrial Fibrillation Patients in Routine Care." Journal of Thrombosis and Thrombolysis 49 (2): 294–303.

[55] Gulilat, Markus, Denise Keller, Bradley Linton, A. Demetri Pananos, Daniel Lizotte, George K. Dresser, Jeffrey Alfonsi, Rommel G. Tirona, Richard B. Kim, and Ute I. Schwarz. 2020b. "Drug Interactions and Pharmacogenetic Factors Contribute to Variation in Apixaban Concentration in Atrial Fibrillation Patients in Routine Care." Journal of Thrombosis and Thrombolysis 49 (2): 294–303. https://doi.org/10.1007/s11239-019-01962-2.

[56] Harrell, Frank E et al. 2015. Regression Modeling Strategies: With Applications to Linear Models, Logistic and Ordinal Regression, and Survival Analysis. Vol. 3. Springer.

[57] Hoffman, Mathew, and Andrew Gelman. 2014. "The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo." Journal of Machine Learning Research 15: 1593–623.

[58] Hoffman, Matthew D, and Andrew Gelman. 2014. "The No-U-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo." Journal of Machine Learning Research 15 (1): 1593–623.

[59] Holbrook, Anne, Sam Schulman, Daniel M Witt, Per Olav Vandvik, Jason Fish, Michael J Kovacs, Peter J Svensson, David L Veenstra, Mark Crowther, and Gordon H Guyatt. 2012. "Evidence-Based Management of Anticoagulant Therapy: Antithrombotic Therapy and Prevention of Thrombosis: American College of Chest Physicians Evidence-Based Clinical Practice Guidelines." Chest 141 (2): e152S–e184S.

[60] Holmes, Mark H. 2009. "Introduction to the Foundations of Applied Mathematics."

[61] Houts, Peter S, Allan Lipton, Harold A Harvey, Barbara Martin, Mary A Simmonds, Richard H Dixon, Santo Longo, et al. 1984. "Nonmedical Costs to Patients and Their Families Associated with Outpatient Chemotherapy." Cancer 53 (11): 2388–92.

[62] Ishwaran, Hemant. 1999. "Applications of Hybrid Monte Carlo to Bayesian Generalized Linear Models: Quasicomplete Separation and Neural Networks." Journal of Computational and Graphical Statistics 8 (4): 779–99.

[63] James, William, and Charles Stein. 1992. "Estimation with Quadratic Loss." In Breakthroughs in Statistics, 443–60. Springer.

[64] Juurlink, David N. 2007. "Drug Interactions with Warfarin: What Clinicians Need to Know." Cmaj 177 (4): 369–71.

[65] Kass, Robert E, Bradley P Carlin, Andrew Gelman, and Radford M Neal. 1998. "Markov Chain Monte Carlo in Practice: A Roundtable Discussion." The American Statistician 52 (2): 93–100.

[66] Kasztura, Miriam, Aude Richard, Nefti-Eboni Bempong, Dejan Loncar, and Antoine Flahault. 2019. "Cost-Effectiveness of Precision Medicine: A Scoping Review." International Journal of Public Health 64 (9): 1261–71.

[67] Koller, Daphne, Nir Friedman, and Francis Bach. 2009. Probabilistic Graphical Models: Principles and Techniques. MIT press.

[68] Kruizinga, Matthijs D, Rob GJA Zuiker, Kirsten R Bergmann, Annelies C Egas, Adam F Cohen, Gijs WE Santen, and Michiel J van Esdonk. 2022. "Population Pharmacokinetics of Clonazepam in Saliva and Plasma: Steps Towards Noninvasive Pharmacokinetic Studies in Vulnerable Populations." British Journal of Clinical Pharmacology 88 (5): 2236–45.

[69] Laber, Eric B, and Ying-Qi Zhao. 2015. "Tree-Based Methods for Individualized Treatment Regimes." Biometrika 102 (3): 501–14.

[70] Lee, Anna, Kanan Shah, and Fumiko Chino. 2020. "Assessment of Parking Fees at National Cancer Institute–Designated Cancer Treatment Centers." JAMA Oncology 6 (8): 1295–97.

[71] Lee, Juhee, Peter F Thall, and Katy Rezvani. 2019. "Optimizing Natural Killer Cell Doses for Heterogeneous Cancer Patients on the Basis of Multiple Event Times." Journal of the Royal Statistical Society: Series C (Applied Statistics) 68 (2): 461–74.

[72] Li, Pin, Jeremy MG Taylor, Spring Kong, Shruti Jolly, and Matthew J Schipper. 2020. "A Utility Approach to Individualized Optimal Dose Selection Using Biomarkers." Biometrical Journal 62 (2): 386–97.

[73] Lin, C C, and L A Segel. 1988. Mathematics Applied to Deterministic Problems in the Natural Sciences. Philadelphia: SIAM.

[74] Livingstone, Samuel, Michael Betancourt, Simon Byrne, Mark Girolami, et al. 2019. "On the Geometric Ergodicity of Hamiltonian Monte Carlo." Bernoulli 25 (4A): 3109–38.

[75] Lizotte, Daniel J. 2017. "Wiley StatsRef: Statistics Reference Online." In. Wiley.

[76] Looff, Margot de, Bob Wilffert, Cornelis Boersma, Lieven Annemans, Stefan Vegter, Job FM van Boven, and Maarten J Postma. 2016. "Economic Evaluations of Pharmacogenetic and Pharmacogenomic Screening Tests: A Systematic Review. Second Update of the Literature." PloS One 11 (1): e0146262.

[77] Lum, Kristian, and William Isaac. 2016. "To Predict and Serve?" Significance 13 (5): 14–19.

[78] Lunn, David J, Andrew Thomas, Nicky Best, and David Spiegelhalter. 2000. "WinBUGS-a Bayesian Modelling Framework: Concepts, Structure, and Extensibility." Statistics and Computing 10 (4): 325–37.

[79] Luque, S, W Hope, N Campillo, R Muñoz-Bermúdez, L Sorli, J Barceló-Vidal, E González-Colominas, et al. 2019. "Population Pharmacokinetics of Anidulafungin in Critically Ill Patients." Antimicrobial Agents and Chemotherapy 63 (7): e00378–19.

[80] MacKay, David JC, David JC Mac Kay, et al. 2003. Information Theory, Inference and Learning Algorithms. Cambridge university press.

[81] Martin, Janet, Avtar Lal, Jessica Moodie, Fang Zhu, and Davy Cheng. 2016a. "Hospital-Based HTA and Know4Go at MEDICI in London, Ontario, Canada." In Hospital-Based Health Technology Assessment: The Next Frontier for Health Technology Assessment, edited by Laura Sampietro-Colom and Janet Martin, 127–52. Cham: Springer International Publishing.

[82] ———. 2016b. "Hospital-Based HTA and Know4Go at MEDICI in London, Ontario, Canada." In Hospital-Based Health Technology Assessment, 127–52. Springer.

[83] McElreath, Richard. 2016. Statistical Rethinking: A Bayesian Course with Examples in r and Stan. Vol. 122. CRC Press.

[84] Merli, Geno J, and George Tzanis. 2009. "Warfarin: What Are the Clinical Implications of an Out-of-Range-Therapeutic International Normalized Ratio?" Journal of Thrombosis and Thrombolysis 27 (3): 293–99.

[85] Metropolis, Nicholas, Arianna W Rosenbluth, Marshall N Rosenbluth, Augusta H Teller, and Edward Teller. 1953. "Equation of State Calculations by Fast Computing Machines." The Journal of Chemical Physics 21 (6): 1087–92.

[86] Metropolis, Nicholas, and Stanislaw Ulam. 1949. "The Monte Carlo Method." Journal of the American Statistical Association 44 (247): 335–41.

[87] Miller, Richard K, and Anthony N Michel. 1982. Ordinary Differential Equations. Academic Press.

[88] Mitchell, Toby J, and John J Beauchamp. 1988. "Bayesian Variable Selection in Linear Regression." Journal of the American Statistical Association 83 (404): 1023–32.

[89] "Monolix 2018R1 User Guide." n.d. Accessed June 4, 2022. https://monolix.lixoft.com/single-page/.

[90] Morse, Bridget L, and Richard B Kim. 2015. "Is Personalized Medicine a Dream or a Reality?" Critical Reviews in Clinical Laboratory Sciences 52 (1): 1–11.

[91] Mould, Diane R, and Richard Neil Upton. 2013. "Basic Concepts in Population Modeling, Simulation, and Model-Based Drug Development—Part 2: Introduction to Pharmacokinetic Modeling Methods." CPT: Pharmacometrics & Systems Pharmacology 2 (4): 1–14.

[92] Murphy, Kevin P. 2012. Machine Learning: A Probabilistic Perspective. MIT press.

[93] Naimi, Ashley I, Stephen R Cole, and Edward H Kennedy. 2017. "An Introduction to g Methods." International Journal of Epidemiology 46 (2): 756–62.

[94] Neal, Radford. 2011. "MCMC Using Hamiltonian Dynamics." Handbook of Markov Chain Monte Carlo 2 (11): 2.

[95] Neal, Radford M. 1996. Bayesian Learning for Neural Networks. Lecture Notes in Statistics. New York: Springer.

[96] Nguyen, Thu Thuy, Henri Bénech, Alain Pruvost, and Natacha Lenuzza. 2016. "A Limited Sampling Strategy Based on Maximum a Posteriori Bayesian Estimation for a Five-Probe Phenotyping Cocktail." Eur. J. Clin. Pharmacol. 72 (1): 39–51.

[97] Ni, Yang, Francesco C Stingo, Min Jin Ha, Rehan Akbani, and Veerabhadran Baladandayuthapani. 2019. "Bayesian Hierarchical Varying-Sparsity Regression Models with Application to Cancer Proteogenomics." Journal of the American Statistical Association 114 (525): 48–60.

[98] Nyberg, Joakim, Caroline Bazzoli, Kay Ogungbenro, Alexander Aliev, Sergei Leonov, Stephen Duffull, Andrew C. Hooker, and France Mentré. 2015. "Methods and Software Tools for Design Evaluation in Population Pharmacokinetics–Pharmacodynamics Studies." British Journal of Clinical Pharmacology 79 (1): 6–17. https://doi.org/10.1111/bcp.12352.

[99] Obermeyer, Ziad, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. 2019. "Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations." Science 366 (6464): 447–53.

[100] Ohara, Minami, Yasuhiko Suzuki, Saki Shinohara, Inna Y Gong, Crystal L Schmerk, Rommel G Tirona, Ute I Schwarz, et al. 2019. "Differences in Warfarin Pharmacodynamics and Predictors of Response Among Three Racial Populations." Clinical Pharmacokinetics 58 (8): 1077–89.

[101] Owen, Joel S, and Jill Fiedler-Kelly. 2014. Introduction to Population Pharmacokinetic/Pharmacodynamic Analysis with Nonlinear Mixed Effects Models. John Wiley & Sons.

[102] Pananos, A Demetri, and Daniel J Lizotte. 2020. "Comparisons Between Hamiltonian Monte Carlo and Maximum a Posteriori for a Bayesian Model for Apixaban Induction Dose & Dose Personalization." In Machine Learning for Healthcare Conference, 397–417. PMLR.

[103] Park, Hyung, Eva Petkova, Thaddeus Tarpey, and R Todd Ogden. 2021. "A Single-Index Model with a Surface-Link for Optimizing Individualized Dose Rules." Journal of Computational and Graphical Statistics, 1–10.

[104] Peretz, David, Ngoc H. On, Donald Miller, Richard Kim, Carla Franklin, Roman Dascal, Cori Knowles, and Gerald Y. Minuk. 2021. "Pharmacokinetics of a Once-Daily Tacrolimus Formulation in First Nations and Caucasian Liver Transplant Recipients." Transplant International 34 (11): 2266–73. https://doi.org/10.1111/tri.13997.

[105] Piironen, Juho, and Aki Vehtari. 2017. "Sparsity Information and Regularization in the Horseshoe and Other Shrinkage Priors." Electronic Journal of Statistics 11 (2): 5018–51.

[106] Plummer, Martyn et al. 2003. "JAGS: A Program for Analysis of Bayesian Graphical Models Using Gibbs Sampling." In Proceedings of the 3rd International Workshop on Distributed Statistical Computing, 124:1–10. 125.10. Vienna, Austria.

[107] Preijers, Tim, Britta A P Laros-vanGorkom, Ron A A Mathôt, and Marjon H Cnossen. 2019. "Pharmacokinetic-Guided Dosing of Factor VIII Concentrate in a Morbidly Obese Severe Haemophilia a Patient Undergoing Orthopaedic Surgery." BMJ Case Reports 12 (1): bcr–2018.

[108] Rackauckas, Chris, Yingbo Ma, Andreas Noack, Vaibhav Dixit, Patrick Kofod Mogensen, Simon Byrne, Shubham Maddhashiya, et al. 2020. "Accelerated Predictive Healthcare Analytics with Pumas, a High Performance Pharmaceutical Modeling and Simulation Platform."

[109] Rich, Benjamin, Erica EM Moodie, and David A Stephens. 2014. "Simulating Sequential Multiple Assignment Randomized Trials to Generate Optimal Personalized Warfarin Dosing Strategies." Clinical Trials 11 (4): 435–44.

[110] ———. 2016. “Optimal Individualized Dosing Strategies: A Pharmacologic Approach to Developing Dynamic Treatment Regimens for Continuous-Valued Treatments.” Biometrical Journal 58 (3): 502–17.

[111] Rich, Benjamin, Erica EM Moodie, David A Stephens, and Robert W Platt. 2010. “Model Checking with Residuals for g-Estimation of Optimal Dynamic Treatment Regimes.” The International Journal of Biostatistics 6 (2).

[112] Robins, James M. 2004. “Optimal Structural Nested Models for Optimal Sequential Decisions.” In Proceedings of the Second Seattle Symposium in Biostatistics, 189–326. Springer.

[113] Robins, James M, Miguel Ángel Hernán, and Babette Brumback. 2000. “Marginal Structural Models and Causal Inference in Epidemiology.” Epidemiology. Wolters Kluwer.

[114] “Robust Statistical Workflow with Rstan.” n.d. https://betanalpha.github.io/assets/case_studies/rstan_workflow.html#25_validating_a_fit_in_stan.

[115] Rogowski, Wolf, Katherine Payne, Petra Schnell-Inderst, Andrea Manca, Ursula Rochau, Beate Jahn, Oguzhan Alagoz, Reiner Leidl, and Uwe Siebert. 2015. “Concepts of ‘Personalization’in Personalized Medicine: Implications for Economic Evaluation.” Pharmacoeconomics 33 (1): 49–59.

[116] Rosenbaum, Sara E. 2016. Basic Pharmacokinetics and Pharmacodynamics: An Integrated Textbook and Computer Simulations, Secpnd Edition. John Wiley & Sons.

[117] Salvatier, John, Thomas V Wiecki, and Christopher Fonnesbeck. 2016b. “Probabilistic Programming in Python Using PyMC3.” PeerJ Computer Science 2: e55.

[118] ———. 2016a. “Probabilistic Programming in Python Using PyMC3.” PeerJ Computer Science 2: e55.

[119] Salway, Ruth, and Jon Wakefield. 2008a. “Gamma Generalized Linear Models for Pharmacokinetic Data.” Biometrics 64 (2): 620–26.

[120] ———. 2008b. “Gamma Generalized Linear Models for Pharmacokinetic Data.” Biometrics 64 (2): 620–26.

[121] Sarma, Abhraneel, and Matthew Kay. 2020. “Prior Setting in Practice: Strategies and Rationales Used in Choosing Prior Distributions for Bayesian Analysis.” In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, 1–12.

[122] Schmidt, Mikkel N. 2009. “Function Factorization Using Warped Gaussian Processes.” In Proceedings of the 26th Annual International Conference on Machine Learning, 921–28.

[123] Schulz, Juliana, and Erica EM Moodie. 2021. “Doubly Robust Estimation of Optimal Dosing Strategies.” Journal of the American Statistical Association 116 (533): 256–68.

[124] Schwarz, Ute I, Marylyn D Ritchie, Yuki Bradford, Chun Li, Scott M Dudek, Amy Frye-Anderson, Richard B Kim, Dan M Roden, and C Michael Stein. 2008. “Genetic Determinants of Response to Warfarin During Initial Anticoagulation.” New England Journal of Medicine 358 (10): 999–1008.

[125] Shabaruddin, Fatiha H, Nigel D Fleeman, and Katherine Payne. 2015. “Economic Evaluations of Personalized Medicine: Existing Challenges and Current Developments.” Pharmacogenomics and Personalized Medicine 8: 115.

[126] Smith, Gary. 2018. “Step Away from Stepwise.” Journal of Big Data 5 (1): 1–12.

[127] Sohrabi, Mohammad Karim, and Alireza Tajik. 2017. “Multi-Objective Feature Selection for Warfarin Dose Prediction.” Comput. Biol. Chem. 69 (August): 126–33.

[128] Stan Development Team. 2012. “Stan Modeling Language Users Guide and Reference Manual, Version 1.0.0.” http://mc-stan.org/.

[129] Standing, Joseph F. 2017. “Understanding and Applying Pharmacometric Modelling and Simulation in Clinical Practice and Research.” British Journal of Clinical Pharmacology 83 (2): 247–54.

[130] Stein, Charles et al. 1956. “Inadmissibility of the Usual Estimator for the Mean of a Multivariate Normal Distribution.” In Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability, 1:197–206. 1.

[131] Stifft, Frank, Franciscus Vandermeer, Cees Neef, Sander van Kuijk, and Maarten H L Christiaans. 2020. “A Limited Sampling Strategy to Estimate Exposure of Once-Daily Modified Release Tacrolimus in Renal Transplant Recipients Using Linear Regression Analysis and Comparison with Bayesian Population Pharmacokinetics in Different Cohorts.” Eur. J. Clin. Pharmacol. 76 (5): 685–93.

[132] Sturkenboom, Marieke G. G., Anne-Grete Märtson, Elin M. Svensson, Derek J. Sloan, Kelly E. Dooley, Simone H. J. van den Elsen, Paolo Denti, Charles A. Peloquin, Rob E. Aarnoutse, and Jan-Willem C. Alffenaar. 2021. "Population Pharmacokinetics and Bayesian Dose Adjustment to Advance TDM of Anti-TB Drugs." Clinical Pharmacokinetics 60 (6): 685–710. https://doi.org/10.1007/s40262-021-00997-0.

[133] Sturkenboom, Marieke GG, Anne-Grete Märtson, Elin M Svensson, Derek J Sloan, Kelly E Dooley, Simone HJ van den Elsen, Paolo Denti, Charles A Peloquin, Rob E Aarnoutse, and Jan-Willem C Alffenaar. 2021. "Population Pharmacokinetics and Bayesian Dose Adjustment to Advance TDM of Anti-TB Drugs." Clinical Pharmacokinetics 60 (6): 685–710.

[134] Sturkenboom, Marieke GG, Leonie W Mulder, Arthur de Jager, Richard van Altena, Rob E Aarnoutse, Wiel CM de Lange, Johannes H Proost, Jos GW Kosterink, Tjip S van der Werf, and Jan-Willem C Alffenaar. 2015. "Pharmacokinetic Modeling and Optimal Sampling Strategies for Therapeutic Drug Monitoring of Rifampin in Patients with Tuberculosis." Antimicrobial Agents and Chemotherapy 59 (8): 4907–13.

[135] Sukumar, Smrithi, Markus Gulilat, Bradley Linton, Steven E Gryn, George K Dresser, Jeffrey E Alfonsi, Ute I Schwarz, Richard B Kim, and Janice B Schwartz. 2019. "Apixaban Concentrations with Lower Than Recommended Dosing in Older Adults with Atrial Fibrillation." Journal of the American Geriatrics Society 67 (9): 1902–6.

[136] Sutton, Richard S, and Andrew G Barto. 2011. Reinforcement Learning: An Introduction. Cambridge, MA: MIT Press.

[137] Tenenbaum, Morris., and Harry Pollard. 1963. Ordinary Differential Equations: An Elementary Textbook for Students of Mathematics, Engineering, and the Sciences. Dover Publications.

[138] Tibshirani, Robert. 1996. "Regression Shrinkage and Selection via the Lasso." Journal of the Royal Statistical Society: Series B (Methodological) 58 (1): 267–88.

[139] Tirona, Rommel G, Zahra Kassam, Ruth Strapp, Mala Ramu, Catherine Zhu, Melissa Liu, Ute I Schwarz, Richard B Kim, Bandar Al-Judaibi, and Melanie D Beaton. 2018. "Apixaban and Rosuvastatin Pharmacokinetics in Nonalcoholic Fatty Liver Disease." Drug Metabolism and Disposition 46 (5): 485–92.

[140] Touma, Lahoud, Kristian B Filion, Renee Atallah, Maria Eberg, and Mark J Eisenberg. 2015. "A Meta-Analysis of Randomized Controlled Trials of the Risk of Bleeding with Apixaban Versus Vitamin k Antagonists." The American Journal of Cardiology 115 (4): 533–41.

[141] Tripuraneni, Nilesh, Mark Rowland, Zoubin Ghahramani, and Richard Turner. 2017. "Magnetic Hamiltonian Monte Carlo." In Proceedings of the 34th International Conference on Machine Learning - Volume 70, 3453–61. ICML'17. Sydney, NSW, Australia: JMLR.org.

[142] Tsiatis, A. A., M. Davidian, S. T. Holloway, and E. B. Laber. 2019. Dynamic Treatment Regimes: Statistical Methods for Precision Medicine. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. CRC Press. https://books.google.ca/books?id=FD0PEAAAQBAJ.

[143] Ueshima, Satoshi, Daiki Hira, Yuuma Kimura, Ryo Fujii, Chiho Tomitsuka, Takuya Yamane, Yohei Tabuchi, et al. 2018. "Population Pharmacokinetics and Pharmacogenomics of Apixaban in Japanese Adult Patients with Atrial Fibrillation." British Journal of Clinical Pharmacology 84 (6): 1301–12.

[144] Upreti, Vijay V., Jessie Wang, Yu Chen Barrett, Wonkyung Byon, Rebecca A. Boyd, Janice Pursley, Frank P. LaCreta, and Charles E. Frost. 2013a. "Effect of Extremes of Body Weight on the Pharmacokinetics, Pharmacodynamics, Safety and Tolerability of Apixaban in Healthy Subjects." British Journal of Clinical Pharmacology 76 (6): 908–16. https://doi.org/10.1111/bcp.12114.

[145] Upreti, Vijay V, Jessie Wang, Yu Chen Barrett, Wonkyung Byon, Rebecca A Boyd, Janice Pursley, Frank P LaCreta, and Charles E Frost. 2013b. "Effect of Extremes of Body Weight on the Pharmacokinetics, Pharmacodynamics, Safety and Tolerability of Apixaban in Healthy Subjects." British Journal of Clinical Pharmacology 76 (6): 908–16.

[146] Vakkalagadda, Blisse, Charles Frost, Wonkyung Byon, Rebecca A Boyd, Jessie Wang, Donglu Zhang, Zhigang Yu, Clapton Dias, Andrew Shenker, and Frank LaCreta. 2016. "Effect of Rifampin on the Pharmacokinetics of Apixaban, an Oral Direct Inhibitor of Factor Xa." American Journal of Cardiovascular Drugs 16 (2): 119–27.

[147] Vehtari, Aki, Andrew Gelman, Daniel Simpson, Bob Carpenter, and Paul-Christian Burkner. 2019. "Rank-Normalization, Folding, and Localization: An Improved r-Hat for Assessing Convergence of MCMC." arXiv Preprint arXiv:1903.08008.

[148] Vinogradova, Yana, Carol Coupland, Trevor Hill, and Julia Hippisley-Cox. 2018. "Risks and Benefits of Direct Oral Anticoagulants Versus Warfarin in a Real World Setting: Cohort Study in Primary Care." BMJ 362: k2505.

[149] Vu, Anh, Tao T Qu, Rachel Ryu, Shuktika Nandkeolyar, Alan Jacobson, and Lisa T Hong. 2021. "Critical Analysis of Apixaban Dose Adjustment Criteria." Clinical and Applied Thrombosis/Hemostasis 27: 10760296211021158.

[150] Wakefield, Jon. 1996b. "The Bayesian Analysis of Population Pharmacokinetic Models." Journal of the American Statistical Association 91 (433): 62–75. https://doi.org/10.1080/01621459.1996.10476664.

[151] ———. 1996a. "The Bayesian Analysis of Population Pharmacokinetic Models." Journal of the American Statistical Association 91 (433): 62–75.

[152] Whittingham, Mark J, Philip A Stephens, Richard B Bradbury, and Robert P Freckleton. 2006. "Why Do We Still Use Stepwise Modelling in Ecology and Behaviour?" Journal of Animal Ecology 75 (5): 1182–89.

[153] Wiens, Jenna, Suchi Saria, Mark Sendak, Marzyeh Ghassemi, Vincent X Liu, Finale Doshi-Velez, Kenneth Jung, et al. 2019. "Do No Harm: A Roadmap for Responsible Machine Learning for Health Care." Nature Medicine 25 (9): 1337–40.

[154] Wigle, Theodore John, Brandi Povitz, Wendy Teft, Robin Legan, John Gordon Lenehan, Markus Gulilat, Stephanie Nevison, et al. 2019. "Prospective Cohort Study of the Impact of Hospital-Wide Dihydropyrimidine Dehydrogenase (DPYD) Genotype Testing for Fluoropyrimidine-Based Chemotherapy on Adverse Events and Hospital Costs." American Society of Clinical Oncology.

[155] Xue, Ling, Nick Holford, Xiao-liang Ding, Zhen-ya Shen, Chen-rong Huang, Hua Zhang, Jing-jing Zhang, et al. 2017. "Theory-based Pharmacokinetics and Pharmacodynamics of S- and R-warfarin and Effects on International Normalized Ratio: Influence of Body Size, Composition and Genotype in Cardiac Surgery Patients." British Journal of Clinical Pharmacology 83 (4): 823–35.

[156] Zhang, Kristine, Yuanheng Wang, Jianzhun Du, Brian Chu, Leo Anthony Celi, Ryan Kindle, and Finale Doshi-Velez. 2021. "Identifying Decision Points for Safe and Interpretable Reinforcement Learning in Hypotension Treatment." https://arxiv.org/abs/2101.03309.

[157] Zhao, Yingqi, Donglin Zeng, A John Rush, and Michael R Kosorok. 2012. "Estimating Individualized Treatment Rules Using Outcome Weighted Learning." Journal of the American Statistical Association 107 (499): 1106–18.

[158] Zhou, Xiang, Peter Carbonetto, and Matthew Stephens. 2013. "Polygenic Modeling with Bayesian Sparse Linear Mixed Models." PLoS Genetics 9 (2): e1003264.

[159] Zhu, Yu-Bin, Xian-Hua Hong, Meng Wei, Jing Hu, Xin Chen, Shu-Kui Wang, Jun-Rong Zhu, Feng Yu, and Jian-Guo Sun. 2017. "Development of a Novel Individualized Warfarin Dose Algorithm Based on a Population Pharmacokinetic Model with Improved Prediction Accuracy for Chinese Patients After Heart Valve Replacement." Acta Pharmacol. Sin. 38 (3): 434–42.