Fall 2021

Bayesian Calibration of the ICRP Zirconium Biokinetic Model and Use of Canned Priors for the Evaluation of Bioassay

Thomas Raymond LaBone

Bayesian Calibration of the ICRP Zirconium Biokinetic Model and Use of Canned Priors for the Evaluation of Bioassay

by

Thomas Raymond LaBone

Bachelor of Arts
The College of New Jersey 1979

Master of Science
Rutgers University 1981

Master of Industrial Statistics
University of South Carolina 2010

Submitted in Partial Fulfillment of the Requirements

for the Degree of Doctor of Philosophy in

Biostatistics

Arnold School of Public Health

University of South Carolina

2022

Accepted by:

Alexander McLain, Major Professor

Bo Cai, Committee Member

James Hardin, Committee Member

Derek Jokisch, Committee Member

Tracey L. Weldon, Interim Vice Provost and Dean of the Graduate School

Acknowledgments

With great pleasure I will take this opportunity to acknowledge the people in my life who have made this dissertation possible. First, I would like to thank those who saw something in me worth nurturing and took a personal interest in mentoring me. They have positively impacted my life more than they will ever know. In more or less chronological order, I would like to thank Gerry Nicholls, Ken Skrable, Len Smith, Roscoe Hall, and Don Edwards for taking me under their wing and showing me the ropes. I can't imagine anyone successfully running the PhD gauntlet without the support of their friends and family, so I would like to thank my daughters Beth LaBone and Randa Johnson, and my son-in-law Chris Johnson for their understanding, support, and constant cheerleading. I worked full time while pursuing my PhD, and I would be remiss if I did not acknowledge the frequent pep talks and therapy sessions provided by my colleagues Nancy Chalmers and Mitch Findley. I thank you both many times over. Moving on to my dissertation committee, I would like to thank Dr. Bo Cai, Dr. James Hardin, Dr. Derrick Jokisch, and my advisor Dr. Alexander McLain. I especially thank my dissertation advisor for accepting the challenge of teaching an old dog new tricks and not giving up on me. I pursued a PhD in an effort to take my game to the next level, and Alex has been a marvelous coach. Finally, pursing a PhD would not have been even remotely possible without my wife Joan holding down the fort, getting me through the dark times, helping me to believe in myself, and loving me unconditionally. Thank you, dear.

Abstract

The International Commission on Radiological Protection (ICRP) publishes biokinetic models that relate measurements of radioactive material in the body and excreta (bioassay) to the amount of the material taken into the body (intake). Given the intake and the biokinetic model, radiation dose to organs and tissues can be calculated. The ICRP approximates the biokinetics of radioactive materials in the body with compartmental models expressed mathematically as a system of ordinary differential equations, for which they provide point estimates of the rate constants. Inaccurate estimates of intake and radiation dose can result in cases where the biokinetics of an individual differ from the ICRP model (i.e., model misspecification), and currently there is no disciplined way to correct for this problem. In addition, the ICRP models don't allow for the estimation of realistic uncertainties in the intakes and radiation doses. To address these limitations, proper probability distributions must be assigned to the rate constants. In this dissertation a proof-of-principle example is given showing how the ICRP 134 zirconium model can be calibrated with bioassay data from an experimental study and Bayesian hierarchical methods, which give the desired probability distributions for the rate constants. Once the ICRP 134 zirconium biokinetic model is calibrated, its posterior distributions are expressed as probability distributions that are used as "canned" priors in subsequent Bayesian evaluations of bioassay data from other individuals. Examples of canned prior Bayesian calculations are given and compared to the results obtained by 3 other methods: standard regression method, add-one-in Bayesian analysis, and MAP analysis of the canned prior Bayesian model. Both Bayesian analyses gave intake estimates with

higher accuracy and more realistic uncertainties than the standard method, but the canned prior Bayesian analysis was approximately 60 times faster than the add-one-in Bayesian analysis. In addition, the canned priors are concise and can be provided along with the biokinetic model in ICRP publications. The MAP analysis gave some of the benefits of the Bayesian analysis with the speed of the standard method, and automatically comes along with the canned Bayesian analysis.

Preface

"What stands in the way becomes the way."

– Marcus Aurelius

Internal dosimetry is a branch of radiation protection that is concerned with the evaluation of radiation doses to individuals resulting from the assimilation of radioactive materials. I have been an internal dosimetrist for over 40 years, specializing in evaluating occupational exposure of individuals to radioactive materials in the workplace. Starting in the late 1990's I had numerous communications with Dr. Guthrie Miller at Los Alamos National Laboratory about the application of Bayesian statistics to problems in internal dosimetry. I was intrigued by the work he was doing, but lacked the background in statistics to fully understand and take advantage of it. I wanted to go back to school to get the background in statistics needed to understand what Guthrie was doing, but the demands of a full-time job and life prevented me from doing so. In 2006 I was blessed to get a new job that offered the flexibility to attend school and work full time (thank you Dave Dooley), so I enrolled in the Master of Industrial Statistics program at the University of South Carolina. It was a wonderful program that got me going in the right direction, but I needed more. In 2014, I spoke to Dr. James Hussey, then chairman of the Biostatistics and Epidemiology Department at the University of South Carolina about the feasibility of pursuing a PhD in his department, and the rest is history. Dr. Hussey was instrumental in starting me on a journey that is nearing its end, and I thank him.

The research detailed in this dissertation uses the Bayesian methods first introduced to me by Guthrie 25 years ago and perfected (more or less) over the last 7

years. My ultimate goal was to offer innovative, practical solutions to some of the problems we face in internal dosimetry. Thus, the researchers who develop and use the biokinetic models in internal dosimetry are my target audience. This means that I have an interdisciplinary audience consisting of internal dosimetrists and statisticians, so I have provided ample background material for both. The final product you see here is probably a bit longer than it needed to be, but only represents a small fraction of what I have learned since 2006.

Thomas LaBone

Aiken, South Carolina

November 22, 2021

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Internal Dosimetry

Radioactive materials spontaneously emit ionizing radiation, which makes them very useful for many industrial and medical applications. However, this property also makes them carcinogenic. To limit the risk of detrimental health effects in individuals, federal and state governments in the United States set regulatory limits on the quantity of radioactive materials that can be taken into the body as a result of occupational exposures, and the concentrations of radioactive materials in the air and water that the public can access, which is intended to limit the quantity of radioactive material a member of the public takes into the body.

The quantity of material taken into the body that is the subject of these regulations is referred to as the *intake*. Intakes can occur through various pathways such as inhalation, ingestion, absorption through intact skin, and through wounds (which includes intravenous injections). An inhalation intake is defined to be the quantity of radioactive material that passes through the nares and an ingestion intake the quantity of material that passes through the mouth. It is not possible in practice to measure an intake directly. Instead, for occupationally exposed individuals the magnitude of the intake is inferred from measurements of the quantity of the radioactive material present in various regions of the body or excreta (e.g., urine and feces) at various times after the intake. Such measurements are referred to as *radiobioassay* or just *bioassay* for short. Once an intake is calculated, the time-integrated retention of

the radioactive material is determined for all significant organs and tissues. Given the time integrated retention and factors that describe the transport of radiation between the organs that contain the radioactive material and the organs of interest, we can calculate the *radiation dose* to the organ of interest. Radiation dose is the amount of ionizing radiation absorbed in tissues and organs and is used as an index of harm (e.g., the risk of developing cancer).

1.2 Biokinetic Models

The process of using bioassay to calculate an intake and using this intake to calculate radiation dose is called *internal dosimetry*[1], and it requires an appropriate *biokinetic model*. A biokinetic model is an idealized mathematical representation of how a material is deposited in the body, subsequently translocated to various organs and tissues, and ultimately excreted ((Boecker, 1998); Harrison, 2009). Thus, the biokinetic model provides a useful mathematical relationship between the intake and the bioassay data, and between the intake and the radiation dose delivered to the organs and tissues of the body by the intake.

The term *kinetic model* has it origins in physiological and biomedical research where a radioactive or relatively rare stable isotope of an element is used to trace the metabolic pathways in living organisms (Robertson, 1983, Chapter 1). When applied to occupational internal dosimetry, kinetic models are usually referred to as *biokinetic models*. Although the applications differ, biokinetic models are basically the same as pharmacokinetic models used in the development of drugs (Bonate, 2011) and toxicokinetic models used to assess the risk from intakes of toxic chemicals (Hubal et al., 2019).

[1]See Li (2018) for a detailed review of internal dosimetry.

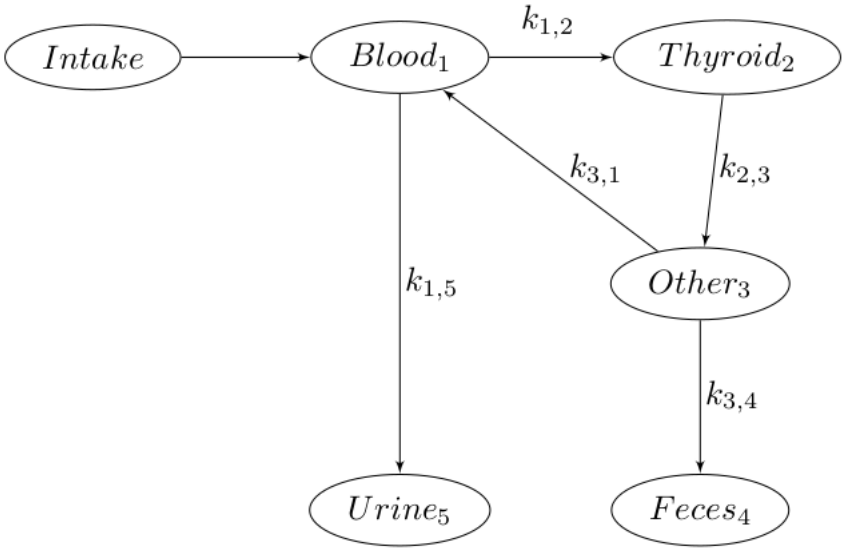

Figure 1.1 Riggs iodine compartmental model.

Biokinetic models for use in occupational internal dosimetry are issued in International Commission on Radiological Protection (ICRP)[2] publications that are subsequently used by regulatory agencies to derive limits for workers and the public.

Biokinetic models are represented by a system of compartments between which the transfer of material is described by first order differential equations. For example, the Riggs iodine biokinetic model for an injection intake of iodine into the bloodstream was first proposed in the early 1950's (Riggs, 1952). This model, which is shown in Figure 1.1, was used with minor modifications in the ICRP Publication 30 (ICRP, 1979) and ICRP Publication 68 (ICRP, 1995), only recently being superseded by the more complex Leggett iodine model in ICRP Publication 137 (ICRP, 2017).

The compartments in a biokinetic model are mathematical constructs used to describe the retention of material in the body and do not necessarily match up with

[2]The *International Commission on Radiological Protection* is an independent organization chartered in 1928 that has developed, maintained, and elaborated the International System of Radiological Protection used world-wide as the common basis for radiological protection standards, legislation, guidelines, programs, and practice. The Commission has published over 100 reports (see www.icrp.org).

specific tissues and organs in the body. In the iodine biokinetic model *blood* represents the pool of inorganic iodine in the body, *thyroid* represents the pool of organic iodine (thyroid hormones) in the overall thyroid, *other* represents the pool of organic iodine in the body other than the thyroid, *urine* represents excretion to the urine, and *feces* represents excretion to the feces. *Physiologically based biokinetic models* (PBBK) incorporate realistic physiological information like blood flow through organs and tissues (Leggett and Williams, 1995), which allows for more accurate modeling of biokinetics across individuals, populations, and species.

1.3 Forward and Inverse Problems

Compartment models are used to solve two different types of problems: the *Forward Problem* and the *Inverse Problem* (Jacquez, 1985, Sections 1 & 3) shown in Figure 1.2. In the Forward Problem we are given the compartmental structure of the model and its rate constants that are used to solve the system of ODE for the content $q_x(t)$ of each compartment x at time t after intake. For example, given measurements of iodine in the thyroid $q_{thy}(t)$ at various times the quantity of iodine initially injected into the blood can be calculated, or given the quantity of iodine injected into the blood the content of the thyroid at any later time can be predicted. The internal dosimetry practitioner uses the ICRP biokinetic models to solve the Forward Problem, i.e., to evaluate bioassay measurements from individuals who were occupationally exposed to radioactive materials (Skrable, Chabot, et al., 1988).

In the Inverse Problem the compartment structure and its transfer rate constants are treated as unknown parameters to be estimated from the measurements of the contents of the compartments and prior knowledge. The modeler solves the Inverse Problem, developing the biokinetic models that are eventually adopted by the ICRP and used by the practitioner. Thus, the modeler takes available data and biological knowledge, applies simplifying assumptions, and develops an idealized compartmental

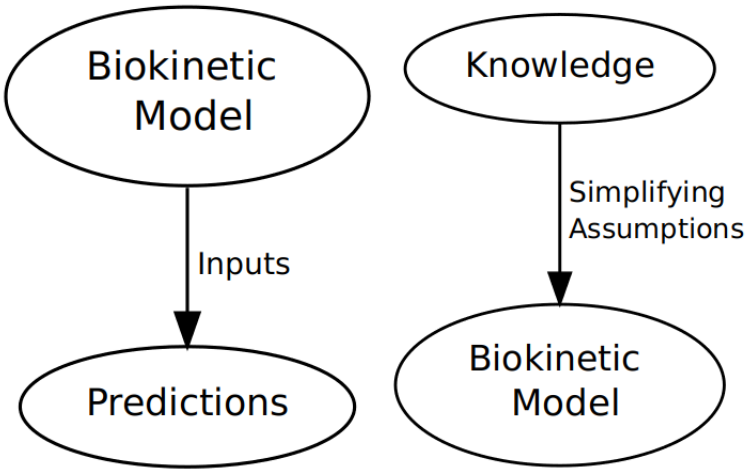

Figure 1.2 Forward Problem (left) and Inverse Problem (right).

model that is *adequate for the task at hand*[3]. This is a complicated process that requires considerable expertise and judgment.

1.4 Problems with the Forward Problem

At first glance Forward Problem appears to be a relatively simple math problem, but there are issues that complicate the task. First of all, frequently the predictions of the compartment contents made with the ICRP biokinetic models don't agree with the observed bioassay measurements. Such disagreements can be the result of measurement error, but often the biokinetic model is not a good approximation to the biokinetics of the individual, i.e., there is model misspecification. This is a consequence of ICRP biokinetic models being deterministic, with all model parameters provided as exact point values (Paquet, Bailey, et al., 2016). Thus, the models don't account for uncertainty in the model structure or parameters, and can't accommodate variability between people. When confronted with this lack-of-fit problem, the practitioner can do one of three things, only the last of which is of interest here: (i) just continue on, stating that the conclusions of the analysis are valid despite the

[3]To paraphrase G. E. P. Box, all biokinetic models are wrong but some are useful (Box, 1979)

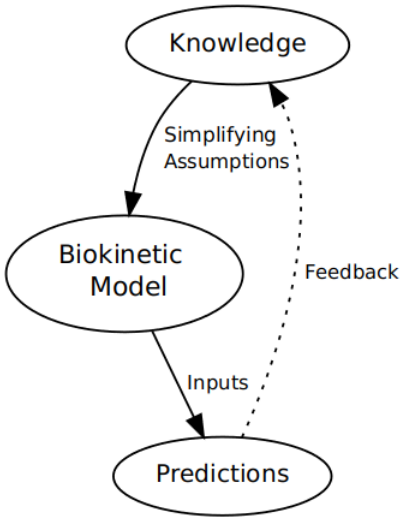

Figure 1.3 Combination of Forward and Inverse Problems.

poor agreement between the predicted and observed measurements (most common in regulatory programs); (ii) use some sort of maximizing technique that will give an overestimate of the intake and dose (most common in compensation programs); (iii) tweak parameters in the biokinetic model to get a better fit to the data.

Model parameters are tweaked in an effort to reduce the misspecification error in the biokinetic model, thus improving the agreement between the predictions of the model and the observed bioassay measurements. This is the Inverse Problem, and attempts to tweak a model must be exercised with caution because: we really don't know which parameters to tweak and by how much – a good fit to the data does not necessarily mean good choices have been made; parameters are correlated, so we can't just change one without the risk of creating an unrealistic biokinetic model; we risk over fitting, i.e., changing parameters in the model to fit noise in the observed measurements.

A disciplined mechanism is needed for creating acceptable agreement between the observed and predicted bioassay measurements while at the same time not making unrealistic changes to the biokinetic model (Figure 1.3).

Closely related to the problem of model misspecification is the desire to estimate and report the uncertainty in assigned intakes and doses. For example, the NCRP[4] (NCRP, 2010) stated that

"*... it is scientifically and ethically necessary to assess the possibility that persons with assigned estimates of internal dose did not in fact receive much larger doses. This is the reason to evaluate the uncertainties in assigned dose.*"

A motivating question behind this dissertation is *how exactly does one go about making a realistic estimate of the uncertainty for an internal dose*? This is a complicated question (Paquet, Bailey, et al., 2016) that the NCRP did not fully address.

In summary, in this dissertation I will focus on how to reduce the inaccuracy of internal dose estimates caused by standard ICRP biokinetic models that do not adequately describe the kinetics of an individual, and estimate of uncertainty in intake and radiation dose caused by variability in the parameters in the biokinetic model at the level of the population and at the level of the individual. The uncertainties in internal dose associated with the transport of radiation within and between organs are not addressed (Fairlie, 2005).

1.5 Path Forward

To deal with the lack-of-fit problem while at the same time assessing the uncertainties in estimated intakes the modeler has to provide more informative biokinetic models. More specifically, realistic probability distributions that describe the uncertainty and variability in the rate constants are needed in place of the point estimates. In addition, the practitioner needs more advanced computational methods that can use these informative biokinetic models.

[4]The National Council on Radiation Protection and Measurements (NCRP) was chartered by the U.S. Congress in 1964 to develop basic concepts about radiation quantities, units and measurements, about the application of these concepts, and about radiation protection.

The practitioner and modeler both face some significant challenges. However, a lot of work has been done in the last 30 years in the field of setting chemical exposure limits for populations and evaluating individual chemical exposures using PBBK biokinetic models and *Bayesian population modeling* (for example, see Hack (2006), Allen, Hack, and Clewell (2007), Bois et al. (1996)). To illustrate the methodology, let's assume that the modeler wants to create a more informative version of the Riggs iodine model. The modeler will:

- Assemble measurements of iodine kinetics in a number of human subjects (e.g., thyroid bioassay, urine bioassay).
- Propose distributions for all rate constants based on all prior knowledge (e.g., previous human and animal studies, SWAGs[5], etc.).
- Update the prior knowledge with the observed measurements using *Bayes Rule* to create posterior distributions of the rate constants. This is referred to as a *Bayesian calibration* of the biokinetic model ((Kennedy and O'Hagan, 2002); Hack, 2006).
- Summarize and publish the posterior distributions for the iodine model in a form that is usable by the practitioner.

This is a fairly complex process, but it produces an iodine model that accounts for variability between individuals and lack of knowledge about the biokinetics of iodine – it has all the properties we need. Once the practitioner has access to the new informative iodine model, it becomes the prior information for the Bayesian analysis of the iodine bioassay data from an out-of-study individual. Such an evaluation allows the individual's bioassay data to influence the estimates of rate constants in the biokinetic model while ensuring that the model is identifiable and physiologically realistic (see

[5]"Scientific wild-ass guess (SWAG) is an American English slang term meaning a rough estimate made by an expert in the field, based on experience and intuition." -Wikipedia.

Section A). In addition, the posterior distributions of the rate constants calculated for the individual can be used to estimate the uncertainty in the individual's iodine intake and organ doses.

The objective of this dissertation is to flesh out the details of how one might implement the procedure discussed in the previous paragraph. A proof-of-principle analysis will be performed with measurements of zirconium[6] in the blood plasma and urine of human subjects after injection intakes. Methods that allow us to generate informative distributions for rate constants are proposed in Aim 1 of this dissertation:

1. Assemble an appropriate zirconium bioassay dataset and perform population-level Bayesian calibration of ICRP 134 zirconium biokinetic model (Chapter 3).

Methods that allow us to use the informative posterior distributions from are given in Aim 2:

2. Fit distributions to the population-level posteriors and use them as priors in the evaluation of out-of-study datasets (Chapter 4).

Aim 3 is to fully document the computational routines and provide simulated bioassay datasets that allow interested individuals to reproduce the calculations in Aims 1 and 2:

3. Create a vignette demonstrating how to calibrate the zirconium model and use the resulting informative priors to evaluate the bioassay data from an individual and make it available on GitHub (Chapter 5).

This is an interdisciplinary study, using Bayesian statistical methods to attack applied problems in biokinetic modeling. The meat of this dissertation lies in Chapters 3 and

[6]Zirconium (Zr) is a metal with an atomic number of 40 whose radioisotopes are common around nuclear power reactors.

4, which will be submitted for publication. Out of necessity those chapters jump right into the details with little background exposition, so to ease statisticians and internal dosimetrists into the discussion some background material on statistical methods and biokinetic modeling is given in Chapter 2. Supporting information that expands on topics discussed in the main chapters of the dissertation are given in the appendix.

Chapter 2

Primer on Biokinetic Modeling

2.1 Biokinetic Models

The Riggs iodine biokinetic model (Riggs, 1952) shown in Figure 1.1 is typical of all ICRP biokinetic models, consisting of compartments between which the movement of material is described by a system of first-order linear ordinary differential equations (ODE). The system of ODE can be solved to give the content of each compartment (and hence any combination of compartments) at any time after the intake. The Riggs iodine model that describes the rate of change in each compartment can be represented as the following system of ODE

$$\begin{aligned}
\frac{dq_1(t)}{dt} &= -k_{11}q_1(t) + k_{21}q_2(t) + k_{31}q_3(t) + k_{41}q_4(t) + k_{51}q_5(t) \\
\frac{dq_2(t)}{dt} &= k_{12}q_1(t) - k_{22}q_2(t) + k_{32}q_3(t) + k_{42}q_4(t) + k_{52}q_5(t) \\
\frac{dq_3(t)}{dt} &= k_{13}q_1(t) + k_{23}q_2(t) - k_{33}q_3(t) + k_{43}q_4(t) + k_{53}q_5(t) \\
\frac{dq_4(t)}{dt} &= k_{14}q_1(t) + k_{24}q_2(t) + k_{34}q_3(t) - k_{44}q_4(t) + k_{54}q_5(t) \\
\frac{dq_5(t)}{dt} &= k_{15}q_1(t) + k_{25}q_2(t) + k_{35}q_3(t) + k_{45}q_4(t) - k_{55}q_5(t).
\end{aligned} \tag{2.1}$$

For example, $q_1(t)$ describes the quantity of iodine in compartment 1 (the blood) at time t after injection of the iodine and $\frac{dq_1(t)}{dt}$ describes the instantaneous rate of change of iodine the blood at that time. A rate constant like k_{12} describes the instantaneous rate at which iodine moves from compartment 1 (the blood) to compartment

2 (the thyroid). Note that rate constants are 0 between compartments with no connections. For example, $k_{51} = 0$ because according to the model there is no transfer of iodine from the urine back to the blood. Rate constants on the diagonal of the rate matrix describe the removal of iodine from the compartment by all pathways. For example, k_{11} describes the rate at which stable iodine moves from the blood to the thyroid and from the blood to the urine.

The rate constants for the Riggs iodine model are given in Table 2.1. Riggs derived the rate constants by looking at the mass balance of iodine in a compartment system that was in equilibrium, i.e., when the rate at which iodine enters the system equals the rate at which it leaves. Such analyses are possible for essential elements like iodine where the material is present in the food we eat. To get a better idea of how this is done let's look at the process in more detail. Riggs assumed that there is 150 $\mu g/day$ of inorganic iodine that comes into the blood from outside sources (the intake) and 150 $\mu g/day$ leaves by the urinary and fecal pathways. Of the 150 $\mu g/day$ of iodine that enters the blood, 70 $\mu g/day$ is taken up by the thyroid and 70 $\mu g/day$ leaves the thyroid since the system is in equilibrium. The metabolic pool of organic iodine in the thyroid under these conditions is 8000 μg, which gives a transfer rate constant of

$$k_{23} = \frac{70\ \mu g/day}{8000\ \mu g} = 0.008736\ day^{-1}. \tag{2.2}$$

Since 70 $\mu g/day$ leaves the thyroid, 70 $\mu g/day$ also enters and leaves the other compartment, which has a 1200 μg iodine pool. The total removal rate constant for other is thus

$$k_{33} = k_{31} + k_{34} = \frac{70\ \mu g/day}{1200\ \mu g} = 0.05833\ day^{-1}. \tag{2.3}$$

Of the 70 $\mu g/day$ leaving Other, 6 $\mu g/day$ goes to feces and 64 $\mu g/day$ is recycled back to the blood. This gives a fraction going to feces f_f of

$$f_f = \frac{6\ \mu g/day}{70\ \mu g/day} = 0.08571. \tag{2.4}$$

The rate constants for Other to feces and Other to blood are

$$k_{34} = f_f \left(0.058272\ day^{-1}\right) = 0.004995\ day^{-1}, \tag{2.5}$$

$$k_{31} = (1 - f_f) \left(0.058272\ day^{-1}\right) = 0.05328\ day^{-1}. \tag{2.6}$$

Finally, if there is 6 $\mu g/day$ going to feces then under equilibrium conditions there must be 144 $\mu g/day$ going to the urine, which gives a transfer rate constant for blood to urine of

$$k_{15} = \frac{144\ \mu g/day}{75\ \mu g} = 1.92\ day^{-1}. \tag{2.7}$$

The system of ODE in (2.1) can be presented in the following matrix notation

$$\begin{pmatrix} \dot{q_1}(t) \\ \dot{q_2}(t) \\ \vdots \\ \dot{q_{n_c}}(t) \end{pmatrix} = \begin{pmatrix} -k_{11} & k_{21} & \cdots & k_{n_c 1} \\ k_{12} & -k_{22} & \cdots & k_{n_c 2} \\ \vdots & \vdots & \vdots & \vdots \\ k_{1n_c} & k_{2n_c} & \cdots & -k_{n_c n_c} \end{pmatrix} \begin{pmatrix} q_1(t) \\ q_2(t) \\ \cdots \\ q_{n_c}(t) \end{pmatrix} \tag{2.8}$$

where $n_c = 5$ is the number of compartments in the iodine model. An even more concise expression is $\dot{\mathbf{q}} = \boldsymbol{kq}$, where $\boldsymbol{k}$ is referred to as the rate matrix. The goal is to solve this system of ODE for the quantity of iodine $q_i(t)$ in each compartment i at time t.

There are a number of different ways to solve this system of linear first-order differential equations, including Laplace transforms (Jacquez, 1985, App 3), Runge–Kutta family of methods (Spiegel, 1981, Ch 9), eigenvalue/eigenvector method (Jacquez, 1985, App 2), and the matrix exponential (Polig, 2001).

Laplace transforms are tedious to use for all but the simplest models, e.g., one can use this method to solve the iodine model but it is not a pleasant experience because of the algebra involved. The other three methods are numerical, so the primary factor in choosing one is the availability of suitable routines that implement the method. The software used in this dissertation for the Bayesian calculations (Stan) implements the

Table 2.1 Rate constants for Riggs iodine model. Rate constants on the diagonal of the rate matrix (e.g., k_{22}) are the sum of all rate constants for material leaving that compartment.

From	To		Rate Constant (1/d)	Half Life (days)
blood	thyroid	k_{12}	0.9336	0.7424
blood	urine	k_{15}	1.92	0.3610
other	feces	k_{34}	0.004992	138.8516
other	blood	k_{31}	0.05328	13.0095
thyroid	other	k_{23}	0.008736	79.3438
other		k_{33}	0.058272	11.8950
blood		k_{11}	2.8536	0.2429
urine		k_{55}	0	
feces		k_{44}	0	
thyroid		k_{22}	0.008736	79.3438

Runge–Kutta methods and the matrix exponential. The matrix exponential method is used here per the recommendations of the software developers (Stan Development Team, 2021b, pg 181) and it has been found to be an efficient method for solving systems of ODE like those associated with biokinetic models. The matrix exponential solution to the system of ODE is

$$\mathbf{q_t} = \exp\left(\mathbf{k}\boldsymbol{t}\right)\mathbf{q_0}, \tag{2.9}$$

where $\mathbf{q_0}$ is the quantity of iodine in each compartment at $t = 0$ and $\exp\left(\mathbf{k}t\right)$ is defined by the Taylor series expansion

$$\exp\left(\mathbf{k}\boldsymbol{t}\right) = \sum_{j=0}^{\infty} \frac{(\mathbf{kt})^j}{j!}. \tag{2.10}$$

Although (2.10) defines the matrix exponential it is not calculated directly using a truncated version (2.10) because of numerical instability of this method (especially when j is large). There are a variety of numerical methods used to solve the matrix exponential (Moler and Van Loan, 2003) but the exact method used in Stan is not specified in its reference documentation.

As mentioned previously, the compartments in a biokinetic model are mathematical constructs that do not necessarily correspond directly to the organs and tissues in

the body or excreta. To improve the agreement between the predictions of the biokinetic model and what is actually observed in bioassay, a collection of compartments in the biokinetic model can considered as a group. This collection of compartments in the biokinetics model that are considered to correspond to what was measured in bioassay is referred to a *reference bioassay function* $m(t)$ (Paquet, Leggett, et al., 2016). For example, if we want to predict the quantity of iodine in the thyroid (compartment 2) the reference bioassay function is

$$m_{thy}(t) = q_2(t), \tag{2.11}$$

if we are interested in predicting the quantity of iodine in the total body, noting $q_4(t)$ and $q_5(t)$ are not in the body, the reference bioassay function is

$$m_{tot}(t) = q_1(t) + q_2(t) + q_3(t), \tag{2.12}$$

and if we are interested in predicting the cumulative quantity of iodine excreted in the cumulative urine the reference bioassay function is

$$m_{urn}(t) = q_5(t). \tag{2.13}$$

Note that $m(t)$ is also a function of $\boldsymbol{k}$ and $\boldsymbol{q}_0$, but they are suppressed here to make the notation more concise. The reference bioassay function is defined for a specific mode of administration. For example, the reference bioassay functions for iodine discussed thus far are for intravenous injection of iodine. Other common modes of administration are inhalation and ingestion. By convention, in a reference bioassay function a unit quantity of material is administered so that the initial contents $\boldsymbol{q}_0$ of all the compartments sum to 1. Thus, the reference bioassay function for an injection of iodine assumes that the initial content of the blood compartment is 1 and the contents of all other compartments are 0.

2.2 The Forward Problem

Predicting the quantity of iodine in the compartments of the iodine biokinetic model given the reference bioassay functions is an example of what is referred to as the Forward Problem. Because the reference bioassay function assumes a unit injection of iodine, if we are interested in predicting the quantity of iodine in the thyroid after injecting 2 units of iodine the desired result is $2m_{thy}(t)$. The multiplier of 2 scales a reference bioassay function to accommodate cases where the administered quantity of iodine is equal to 2. This multiplier that represents the quantity of material administered is referred to as the *intake* β.

An applied problem of interest is to calculate the intake given the the reference bioassay functions and the quantity M of material actually measured in specific organs and tissues or excreta. For example, given a single measurement $M_{thy}(10)$ of iodine in the thyroid at $t = 10$ days after an injection of iodine, the estimated intake $\hat{\beta}$ is

$$\hat{\beta} = \frac{M_{thy}(10)}{m_{thy}(10)} \tag{2.14}$$

and the predicted quantity of iodine in the thyroid is

$$\hat{M}_{thy}(10) = \hat{\beta} m_{thy}(10) \tag{2.15}$$

Such calculations make the implicit assumption that the biokinetics of the individual in whom we measured $M_{thy}(10)$ are identical to the biokinetics of specified by the Riggs biokinetic model.

If there are $n > 1$ thyroid bioassay measurements from a single individual following a single intake at $t = 0$ we have to calculate the intake that gives the "best" agreement between the observed and predicted bioassay. One way to do this is to define the sum of the squares of the differences between the observed and predicted values (the residual sum of squares)

$$RSS = \sum_{i=1}^{n} \left[M_{thy}(t_i) - \beta m_{thy}(t_i) \right]^2 \tag{2.16}$$

and select the value of β that minimizes the sum. This is a regression of $\boldsymbol{M}$ on $\boldsymbol{m}$ through the origin, the formal equation of which is

$$M(t_i) = \beta \cdot m(t_i) + \varepsilon, \quad i = 1 \ldots, n \tag{2.17}$$

where the intake β is the slope of the line and ε is the homoscedastic regression error. Bioassay measurement errors are usually heteroscedastic, so a refinement is to perform a weighted regression, minimizing the weighted residual sum of squares (Claudio Cobelli, Foster, and Toffolo, 2007, Chapter 10; Skrable, Chabot, et al., 1994)

$$WRSS = \sum_{i=1}^{n} w_i \left[M_{thy}(t_i) - \beta m_{thy}(t_i)\right]^2, \tag{2.18}$$

where w_i are the weighting factors. Commonly used weighting factors are analytical measurement uncertainties u (LaBone, 2010) or functions of the observed or predicted bioassay measurements (Skrable, French, et al., 2002; Wakefield, 1996; Wakefield, Aarons, and Racine-Poon, 1999).

Compartmental models like the Riggs iodine model are referred to as *models of systems* (DiStefano and Landaw, 1984; Carson, Finkelstein, and C. Cobelli, 1983, pg 39). In system models the form of the regression given in (2.17) is dictated by the biokinetic model. Arbitrary regression functions fit to bioassay measurements are referred to as *models of data*. Relatively simple data models can provide answers based on the observed data. For example, measurements of the thyroid content alone over time can't be used to estimate the content of the Other compartment or the content of the thyroid compartment at times past the last measurement. On the other hand, system models can be used to incorporate all available knowledge into a single model that can give answers to questions concerning compartments under conditions not directly observed. For example, following an injection of radioactive iodine we might want to know the radiation dose to organs and tissues other than the thyroid, i.e., the combination of the blood and Other compartments, based on measurements of iodine in the thyroid. To do this requires a system model (the Riggs

iodine model) that relates the quantity of iodine in the thyroid to the quantity of iodine in other compartments as a function of time.

As an example, reference bioassay functions were calculated for the thyroid by solving the Riggs iodine model and the intake of iodine calculated from measurements of iodine in the thyroid using weighted least squares regression for two individuals. In the regression the weights were equal to $w_i = 1/u_i^2$, where u_i is the measurement uncertainty. The results for Subject 4 and Subject 10 are shown in Figure 2.1 (a) and (b), respectively. The results of the weighted least squares regressions are given in Table 2.2. The plot for Subject 4 exhibits lack of fit, indicating that biokinetics of Subject 4 are not adequately approximated by the Riggs biokinetic model.

These linear regressions are examples of the Forward Problem where a proportionality constant (the intake) is estimated. Note that when we perform regression to estimate the intake it is assumed that we have the "correct" biokinetic model and well behaved bioassay data. In this case the data points will lie neatly around a straight line on a plot of observed bioassay data versus reference bioassay functions. To obtain a better fit to the bioassay data like those from Subject 4 we have to adjust the parameters in the Riggs iodine model based on the bioassay results of the individual while performing the regression. Maximum likelihood methods like those discussed in the next section are better suited to such tasks than is linear regression.

2.3 Maximum Likelihood

Assume that the errors for the regression are independent, homoscedastic, and normally distributed. Given this, the sampling distribution of the data $\boldsymbol{M}$ is

$$\boldsymbol{M} \sim N\left(\beta \boldsymbol{m}, \boldsymbol{\Sigma}\right), \tag{2.19}$$

$$\boldsymbol{\Sigma} = \sigma^2 \boldsymbol{I}, \tag{2.20}$$

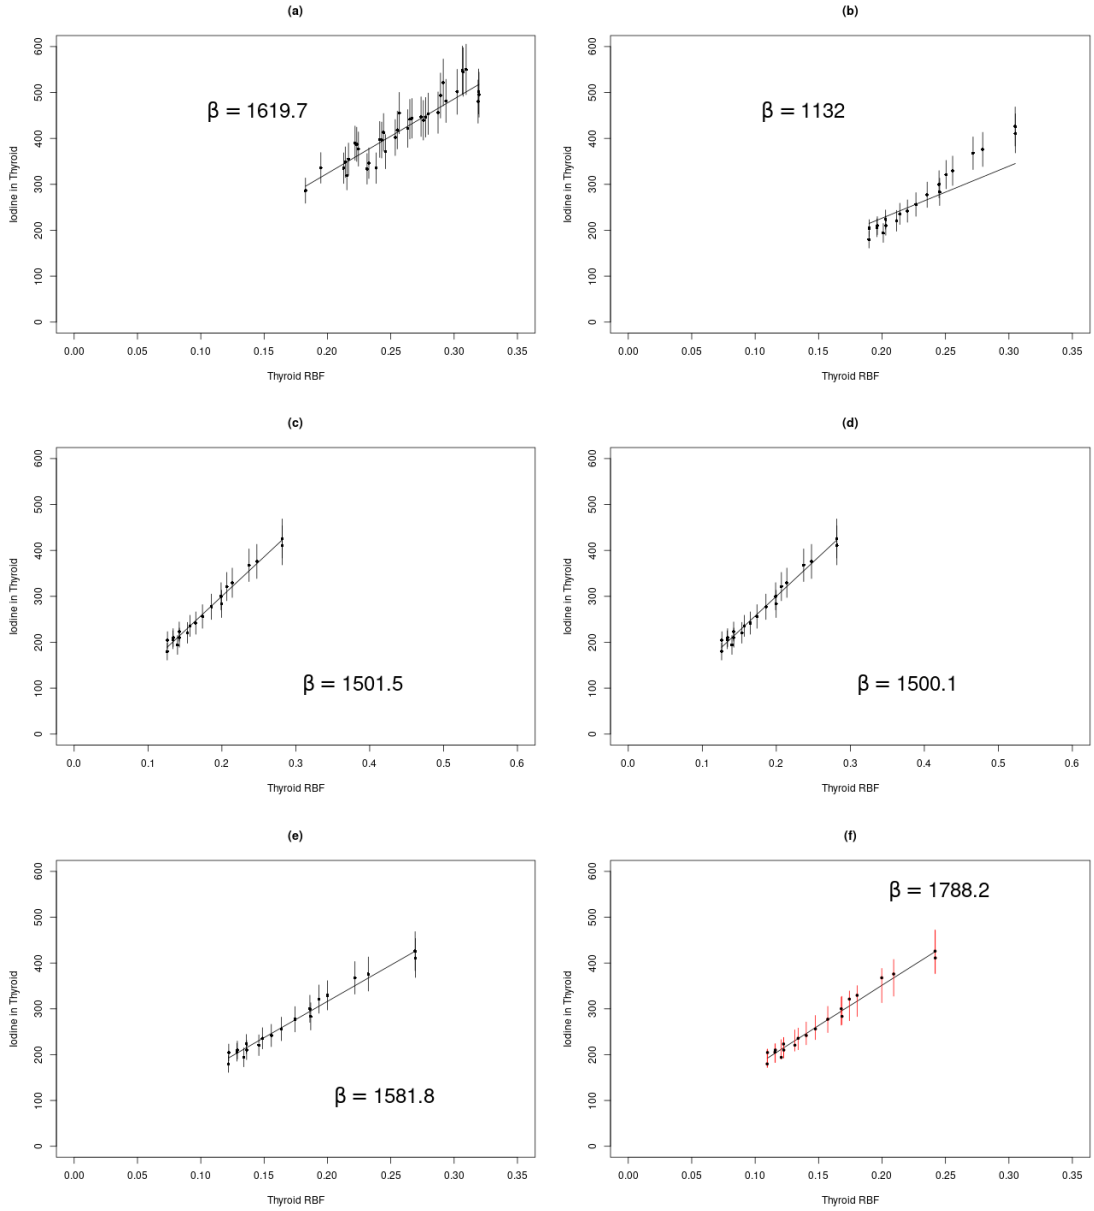

Figure 2.1 Scatter plots of observed thyroid burden versus thyroid reference bioassay function for Subject 10 (a) and Subject 4 (b). These fits were calculated using weighted least squares estimates of the intakes (the errors bars are 2σ measurement uncertainties). The fits obtained for Subject 4 using unrestricted maximum likelihood (c), restricted maximum likelihood (d), MAP (e), and Bayesian (f) estimates of intake and the rate matrix. In plot (f) the error bars are the 95% credible intervals for each measurement and the intake is the geometric mean (see Appendix D for definition) of the intake posterior distribution.

Table 2.2 The results of weighted least squares (WLS), unrestricted maximum likelihood (MLU), restricted maximum likelihood (MLR), and maximum a posteriori (MAP) estimates of intake and the rate constants (all except WLS, which estimates intake only). The values in red are problematic because they are either physiologically improbable or physically impossible.

	Sub	Intake	k[bld,thy]	k[bld,urn]	k[oth,fec]	k[oth,bld]	k[thy,oth]
WLS	4	1132	0.9336	0.0087	0.0533	0.0050	1.9200
MLU	4	1502	0.9408	0.0106	-1.7086	146.202	2.1100
MLR	4	1500	0.9418	0.0106	0.6862	146.576	2.1092
MAP	4	1582	0.9512	0.0144	0.0486	0.0049	2.1905
WLS	10	1620	0.9336	0.0087	0.0533	0.0050	1.9200
MLU	10	418	1.4234	0.0322	1023.57	692.644	-0.3293
MLR	10	956	0.4900	0.0120	909.23	413.286	0.3799
MAP	10	1667	0.9347	0.0092	0.0483	0.0048	1.9566

where σ is the standard deviation of the regression error and $\boldsymbol{I}$ is an $n \times n$ identity matrix. The *likelihood* $L(\beta)$ of the intake given the data, its covariance matrix, and the reference bioassay functions is

$$L(\beta) = \prod_{i=1}^{n} f\left(M(t_i), \beta m(t_i), \Sigma_{i,i}\right) \tag{2.21}$$

where $f(\cdot)$ is the normal pdf with conditional mean $\beta m(t_i)$ and variance $\Sigma_{i,i}$ for observed bioassay result $M(t_i)$. As discussed in Appendix G the likelihood is frequently expressed as $p_\beta(\boldsymbol{M}|\mathbf{m}, \beta) = L(\beta)$. The *maximum likelihood estimator* (MLE) for the intake β is that value of β that maximizes the likelihood function, i.e., it is the value of β that makes the observed data the most likely to have been observed

$$\hat{\beta} = \arg\max L(\beta). \tag{2.22}$$

Numerically, it is more convenient to work with the log of the likelihood function

$$\mathcal{L}(\beta) = \log\left(L(\beta)\right), \tag{2.23}$$

which has the same maximum as $L(\beta)$ because the maximum of a function is invariant under a monotonic transformation like the log.

Finding the MLE of single parameter like β is straightforward (it can be accomplished by hand), but more sophisticated optimizers are required for the multidimensional optimizations discussed in the next section. Note that the MLE of β is the same as the weighted least squares estimate because of the assumption of normally distributed regression errors.

Errors in the parameter estimates (β in this case) are based on quadratic approximation to the curvature of the log likelihood function evaluated at the maximum likelihood estimate (Pawitan, 2001, page 31), i.e., the observed Fisher information $I\left(\hat{\beta}\right)$

$$I\left(\hat{\beta}\right) = -\frac{d^2}{d\beta^2}\mathcal{L}\left(\beta\right). \tag{2.24}$$

A high curvature in the likelihood function results in a narrow likelihood function that is well defined and has a small estimated standard error. A quadratic curvature corresponds to a normal distribution, but this approximation is not accurate if the log likelihood function is not quadratic. The normal approximations of standard errors in the parameters are typically given in the output of optimization routines from the *bbmle* library (Bolker, 2020) and can be used to construct normal Wald confidence intervals for the uncertainty in the parameters (Pawitan, 2001, pg 42).

Heterscedasticity in bioassay measurement errors is addressed with a covariance matrix $\mathbf{\Sigma}$ that contains the squared measurement uncertainties $\boldsymbol{u}^2$ on its diagonal and zeros elsewhere (i.e. the analytical uncertainties in the measurements are uncorrelated):

$$\mathbf{\Sigma}_{i=i} = u_i^2, \quad \mathbf{\Sigma}_{i\neq j} = 0. \tag{2.25}$$

2.4 The Inverse Problem

Fits of the bioassay data to reference bioassay functions often show excessive lack of fit like that seen for Subject 4 because the standard biokinetic model does not adequately describe the biokinetics of the individual. As a remedy we can attempt to

estimate the reference bioassay functions and the intake using maximum likelihood. To accomplish this the reference bioassay function in the likelihood must be expressed as an explicit function of the rate matrix

$$p_\beta\left(\boldsymbol{M}|\mathbf{m},\beta\right)=L\left(\beta,\boldsymbol{k}\right)=\prod_{i=1}^{n}f\left(M\left(t_i\right),\beta m\left(t_i,\boldsymbol{k}\right),\Sigma_{i,i}\right). \tag{2.26}$$

Note that the intake β is a physical quantity that could conceivable be known, but the rate constants $\boldsymbol{k}$ are latent variables that are not directly observable.

In theory, the maxima of the likelihood in (2.26) with respect to β and $\boldsymbol{k}$ give the maximum likelihood estimators of $(\beta,\boldsymbol{k})$. However, there are significant problems encountered if we try to do this because the ML estimates may be physiologically unrealistic, e.g., negative rate constants or intakes may be obtained if no restrictions are imposed on the estimated values of the parameters (unrestricted maximum likelihood). The plot for Subject 4 shown in Figure 2.1(c) and looks quite reasonable. However, inspection of the results (rows 3 and 4 of Table 2.2) shows physiologically unrealistic values for $k_{oth,fec}$ and $k_{oth,bld}$, including negative rate constants that are not physically possible. The likelihood function can be constrained to produce positive parameter estimates, but physiologically unrealistic rate constants like $k_{oth,fec}=909.2\ day^{-1}$ in Table 2.2 are still possible. Such problems can't be identified by examination of the fit of the model predictions to the data. For example, the extreme restricted likelihood estimates in rows 5 and 6 of Table 2.2 appear to fit the data well in Figure 2.1(d). The problem is that the Riggs iodine biokinetic model is not *identifiable* given only thyroid bioassay data. In a non-identifiable biokinetic model there are multiple sets (perhaps infinite number) of rate constants that are solutions to the model. In other words, a non-identifiable biokinetic model has a likelihood function that is essentially flat for many combinations of parameters. As discussed in Appendix A, non-identifiability can occur even if we have an infinite amount of error-free data, and ICRP biokinetic models should be assumed to be non-identifiable

in all practical applications where there are limited quantities and types of bioassay data.

2.4.1 Maximum A Posteriori Estimates

One way to constrain parameter estimates to physiologically realistic values and ensure that the model is identifiable is to bring additional information about the parameters into the calculation — information that is not derived from the current dataset. This prior information is expressed as probability distributions for the parameters that are referred to as *prior probability distributions*. For example, based on previous knowledge a multivariate lognormal prior $p(\boldsymbol{k})$ can be assigned to the rate matrix $\boldsymbol{k}$ where the log means of the distribution are the logs of the Riggs rate constants, the log standard deviations of the distribution are are log(1.25), and the rate constants are not correlated. Further, based on knowledge of the incident that caused the intake of iodine, a lognormal distribution is assigned to the prior for the intake $p(\beta)$ where the log mean is log(1600) and the log standard deviation log(3). The maxima of the resulting product $L(\beta, \boldsymbol{k})\, p(\boldsymbol{k})\, p(\beta)$ are the a posteriori maximum likelihood (MAP) estimates of the parameters, which are given in rows 7 and 8 of Table 2.2. The fit to the data using the MAP parameter estimates is shown in Figure 2.1(e).

2.5 Bayesian Methods

All the methods discussed thus far to solve the Inverse Problem provide point estimates of the model parameters and estimates of the uncertainty in those point estimates by treating the data as random and the parameters fixed. This paradigm can be referred to as *frequentist statistics*. In a reversal of paradigm, we can treat the data as fixed and the parameters as random, which is referred to as *Bayesian statistics*. A brief introduction to Bayesian analysis is given here, but there is a large body of literature discussing the theory and application of Bayes Law. In order of

increasing complexity, I recommend Lambert (2018), McElreath (2018), and Gelman, Carlin, et al. (2013) to those unfamiliar with Bayesian methods. The application of Bayesian methods to biokinetic models is discussed in Bonate (2011, Ch 10).

If we express the likelihood as $p(\boldsymbol{M}|\mathbf{k}, \beta)$ the laws of conditional probability tell us that the joint distribution $p(\boldsymbol{k}, \boldsymbol{\beta}, \boldsymbol{M})$ of the parameters and the data is

$$p(\boldsymbol{k}, \boldsymbol{\beta}, \boldsymbol{M}) = p(\boldsymbol{k}, \boldsymbol{\beta}|\boldsymbol{M})\, p(\boldsymbol{M}) = p(M|\boldsymbol{k}, \beta)\, p(\boldsymbol{k})\, p(\beta), \tag{2.27}$$

where $p(\boldsymbol{M})$ is the marginal probability density of the bioassay data and $p(\boldsymbol{k}, \boldsymbol{\beta}|\boldsymbol{M})$ is the probability of observing the parameters given the data, which is called the *posterior distribution*. A minor rearrangement of (2.28) gives the immensely useful Bayes Law

$$p(\boldsymbol{k}, \boldsymbol{\beta}|\boldsymbol{M}) = \frac{p(M|\boldsymbol{k}, \beta)\, p(\boldsymbol{k})\, p(\beta)}{p(\boldsymbol{M})}. \tag{2.28}$$

In Bayesian methods the posterior distribution contains all that is presently known about the parameters, combining what the data are telling us with our prior knowledge. All inferences about the fate of iodine in this system are calculated with the posterior distributions. Like MAP methods, Bayesian methods make the parameters in biokinetic models identifiable through the incorporation of prior knowledge about the biokinetic model and its parameters. Unlike MAP methods, Bayesian methods provide estimates of the probability distribution of parameters that are interpreted as the uncertainty in the parameters of the biokinetic model.

The marginal probability density of the data $p(\boldsymbol{M})$ is a normalization constant that ensures the posterior distribution is a proper probability distribution and is

$$p(\boldsymbol{M}) = \int\int p(M|\boldsymbol{k}, \beta)\, p(\boldsymbol{k})\, p(\beta)\, d\boldsymbol{k} d\beta. \tag{2.29}$$

This integration is intractable in most real-life problems so it is side stepped by using a simulation that takes correlated random samples directly from the posterior distribution. This type of sampling is referred to as Markov Chain Monte Carlo

(MCMC) to help differentiate it from regular Monte Carlo sampling where the samples are independent of each other. The MCMC samples are used to estimate the posterior distributions and their summary statistics.

The results of a Bayesian analysis are expressed as the posterior distributions of the parameters. For example, the data for Subject 4 were analyzed using the lognormal prior distributions on the rate constants and the intake discussed above. The posterior distributions of the rate constants are shown in Figure 2.2(a) and the posterior distribution of the intake in 2.2(b) . The rate constant for $k_{thy,oth}$ is noticeably more precisely than the other rate constants (i.e., it has a narrower peak) because of the large influence the thyroid data had on its posterior. The geometric mean (gm) and geometric standard deviation (gsd) are used to as summary statistics for the intake. The posterior distributions are proper probability density functions and it is appropriate to make statements like there is a 95% probability that the *true* iodine intake by Subject 4 is in the *credible* interval of $(1271, 2691)$, which is denoted by the vertical dashes lines in Figure 2.2(b). The *confidence* intervals on parameters generated by the other methods like maximum likelihood can not be interpreted in this way (Lambert, 2018).

The fit to the bioassay data using Bayesian methods is shown in Figure 2.1(f). The error bars are pointwise 95% credible intervals on the predicted quantity of iodine in the thyroid, i.e., on the regression line. Note that the error bars in the other plots in Figure 2.1 are 95% coverage intervals on the measurements and should not be compared directly to the credible intervals.

2.6 Summary

The Inverse Problem in biokinetic modeling, which is the estimation of parameters in the biokinetic model based on observed bioassay data, was discussed in this chapter. This process is referred to as calibrating the biokinetic model. The relatively

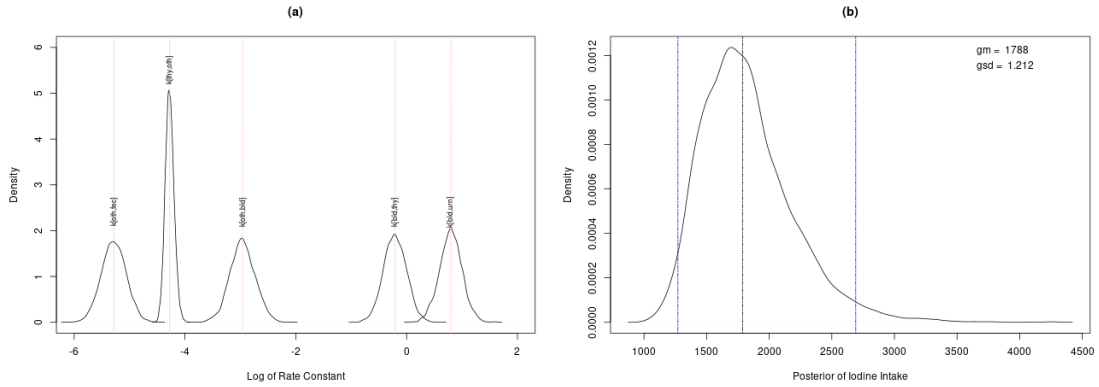

Figure 2.2 Marginal posterior distributions for rate constants (left) and intake (right) for Subject 4.

simple Riggs iodine biokinetic model was used to introduce the nomenclature and mathematics of biokinetic models and to illustrate several statistical methods used to calibrate the biokinetic model with observed bioassay data. The Bayes method for calibrating a biokinetic model was shown to have the advantages of allowing us to incorporate prior knowledge, which makes the biokinetic model identifiable, and expressing results as probability distributions, which allows us to properly assess the uncertainty in the results. The general workflow of a Bayesian analysis is often referred to as "Turning the Bayesian Crank" because a similar workflow is used in all Bayesian problems. Adapted from Kruschke (2015, pg 25), the Bayesian Crank used in this dissertation can be described as follows:

1. Identify the bioassay data.
2. Identify the biokinetic model.
3. Specify prior distributions for all the parameters in the biokinetic model.
4. Use Bayesian inference to meld the data and prior distributions of the parameters into the posterior distributions of the parameters.
5. Check to see whether the predictions of the biokinetic model using the posterior distributions agrees with the observed bioassay data.

The primary difficulties associated with turning the crank are Step 3, fabricating prior distributions that are realistic and informative, and Step 4, dealing with the integral in (2.29). These problems are discussed in more detail in the next chapter where the more complex zirconium biokinetic model is calibrated.

Chapter 3

Population Zirconium Biokinetic Model

3.1 Introduction

Radioactive materials spontaneously emit radiation, which makes them very useful for many industrial and medical applications. However, this property also makes them potentially carcinogenic (National Research Council, 1988). In order to limit the risk of detrimental health effects in individuals, limits are set on the quantity of radioactive materials that can be taken into the body as a result of occupational exposures. The quantity of material taken into the body that is the subject of these limits is referred to as the *intake*. Intakes can occur through various pathways such as inhalation, ingestion, absorption through intact skin, and through wounds. In practice it is not possible to measure an intake directly. Instead, for occupationally exposed individuals the magnitude of the intake is inferred from measurements of the quantity of the radioactive material present in regions of the body or excreta (e.g., urine and feces) at various times after the intake occurred. Such measurements are referred to as *radiobioassay* or just *bioassay* for short.

Using bioassay to calculate an intake requires an appropriate *biokinetic model*. A biokinetic model is an idealized mathematical representation of how a material is deposited in the body, subsequently translocated to various organs and tissues, and ultimately excreted (Boecker, 1998; Harrison, 2009). The biokinetic model provides a useful mathematical relationship between the intake and the bioassay data, and between the intake and the *radiation dose* delivered to the various organs and tissues

of the body. Radiation dose is the amount of ionizing radiation absorbed in tissues and organs and is used as an index of harm, relating the intake to the risk of developing cancer, and is calculated by applying factors that describe the transport of radiation between organs to the time-integrated retention of the radioactive material in the organs.

The International Commission on Radiological Protection (ICRP) publishes biokinetic models for all radionuclides of practical concern. These models are the international standard, being used to calculate intakes from bioassay data and to calculate permissible concentrations of radionuclides in air and water (Clarke and Valentin, 2009). Two problems with the application of these models are: (i) the ICRP biokinetic model may not adequately describing the biokinetics of a particular individual, and (ii) ICRP models don't give a straightforward way to estimate the uncertainty in estimated intakes and radiation doses. Practical methods for dealing with these two problems are the main focus of this dissertation

My exposition will revolve around the biokinetics of *zirconium* (Zr), which is a metallic element whose radioisotopes like Zr^{95} are commonly associated with nuclear power reactors. The ICRP defined the compartmental structure of the ICRP 134 zirconium biokinetic model (Paquet, Leggett, et al., 2016) shown in Figure 3.2 and provided point values of the rate constants in the system of ordinary differential equations (ODE) that define the biokinetics of zirconium in the human body.

Greiter et al. (2011) presented data from a zirconium stable isotope study (discussed further in Section 3.2) that were used to construct a zirconium biokinetic model that will be referred to as the Helmholtz Zentrum Center for Radiation Research (HMGU) model. Bayesian model comparison methods were used by Schmidl et al. (2012) to show that the HMGU model had better predictive power than the ICRP Publication 56 zirconium biokinetic model (ICRP, 1990), which was the predecessor for the ICRP 134 zirconium model. Schmidl et al. (2012) also discussed the

use of Bayesian methods to estimate credible intervals on radiation doses assigned to an individual. Finally, Schmidl's work is referenced in the discussion in Section 3.4.4 on prior selection. Li, Greiter, et al. (2011a) performed an uncertainty analysis of the HMGU zirconium biokinetic model and Li, Greiter, et al. (2011b) performed a sensitivity analysis of the model using Monte Carlo methods. These studies were concerned with selecting an optimal compartmental structure (i.e., the number of compartments and their interconnections) for a zirconium biokinetic model and comparing it to the ICRP biokinetic model that was in use at the time.

In the calculations performed in this dissertation, the system of ODE must be solved for each MCMC iteration, which is extremely computationally intensive. Schmidl et al. (2012) used the matrix exponential to solve the system of ODE as was done here (see Section 3.3). This is an expensive calculation that can be avoided by generating a large number of candidate combinations of rate constants, solving the systems of ODE once, and storing the solutions in a look-up table (Miller, 2008; Miller, 2017; Poudel et al., 2018; NCRP 2010, Section 5.3.2.3). Another option for this computationally intensive task is *weighted likelihood Monte Carlo sampling* (WELMOS) (Puncher, Birchall, and Bull, 2012; Puncher, Birchall, and Bull, 2014; Puncher and Riddell, 2016; NCRP 2010, Section 5.5.2). In WELMOS, random draws are taken from prior distributions for the the parameters in the biokinetic rate matrix and a likelihood calculated for a discretized intake prior. The likelihoods are summed to give an approximation to the posterior distribution of the intake. This is not a true Bayesian approach because it samples from the prior rather than the posterior distribution, but if the data are not highly informative it will give results very similar to a true Bayesian approach (NCRP, 2010, pg 169).

The goals of this dissertation are fundamentally different the works cited above. In this chapter a Bayesian *calibration* (Hack, 2006) is performed on the ICRP 134 biokinetic model using bioassay data from the Greiter stable isotope study. This

calibration incorporates the prior distributions of the model parameters and results in posterior probability distributions of the rate constants in the model given the observed data. This stochastic biokinetic model is designed to achieve the best agreement between the observations and the model predictions. A hierarchical Bayesian calibration is performed, which accounts for parameter variability in an individual and between individuals in the study group. In principle, knowledge of the posterior probability distributions of the model rate constants allows one to calculate the uncertainty in all quantities derived from the biokinetic model. For example, it can provide estimates and credible intervals for the intake of zirconium and the radiation doses to organs and tissues. This population-level Bayesian calibration is a proof-of-principle exercise that shows how researchers who develop biokinetic models for the ICRP can provide probability distributions for the parameters in their models that are consistent with observed bioassay and prior knowledge. In an actual model development, prior information from multiple experimental studies and accidental exposures can be combined with expert elicitation (i.e., high quality SWAGs) to create highly informative models that encapsulate all prior knowledge.

In the following Section we give a presentation of the Greiter zirconium data followed by a summary of the ICRP 134 zirconium biokinetic model in Section 3.3. Then, in Section 3.3.1 reference bioassay functions and intakes are defined in terms of the biokinetic model. At this point population and individual Bayesian statistical models are defined in Section 3.4, which include the specification of the likelihood and prior distributions. The chapter concludes with the presentation of the results in Section 3.5 and the discussion of results in Section 3.6.

3.2 Data

The data from the Greiter study were provided by Dr. Augusto Giussani[1] of the German Bundesamt fur Strahlenschutz. The studies were originally performed at the German Helmholtz Zentrum Center for Radiation Research (HMGU). The data from the $N = 16$ study subjects used in previous studies (Greiter, 2008; Greiter et al., 2011) and data from 2 additional out-of-study subjects were provided. Bioassay measurements consisted of the concentration of zirconium in the blood plasma and the quantity of zirconium in incremental urine samples at various times after intravenous administration of a known quantity of stable zirconium. The individuals in the study each have n_j bioassay measurements, $j = 1, 2, \ldots N$. Each individual has n_p measurements of the concentration of zirconium in the blood plasma and n_u measurements of the zirconium in incremental urine samples, with $n_j = n_{u_j} + n_{p_j}$. For the j^{th} individual, the response variable $\boldsymbol{M}_j$ is a $n_j \times 1$ vector of the bioassay measurements that consists of n_{p_j} plasma measurements stacked on top of n_{u_j}urine measurements. These measurements were made at times $\boldsymbol{t}_j$ (a $n_j \times 1$ vector) after injection of the zirconium. Urine samples are collected over a $n_{u_j} \times 1$ vector of time increments denoted by $\boldsymbol{\Delta t}_j$. Urine samples tend to have a Δt equal to either 12 or 24 hours.

The combined standard uncertainty u of the measurement calculated according to guidance given in the *ISO Guide to the Expression of Uncertainty in Measurement (GUM)* (ISO, 2008) was provided for each measurement. To facilitate comparison of data from different individuals in the study, all bioassay measurements and associated measurement uncertainties were normalized by dividing them by the quantity of zirconium injected. Thus, in the original dataset the units of the blood plasma bioassay are fraction of injected zirconium per kg of blood plasma and the units of the

[1]Personal communication on 4/14/2020.

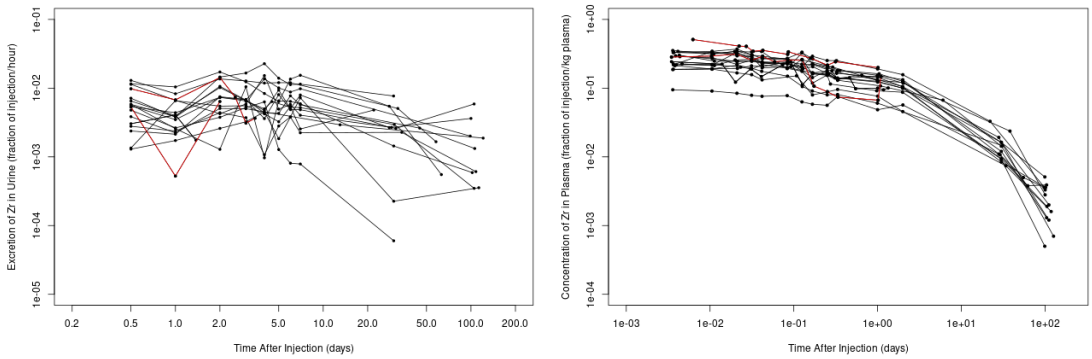

Figure 3.1 Urine and blood plasma data for 16 subjects from the Greiter zirconium study used in the population model and 2 subjects (in red) that were not used in the model. These two out-of-sample individuals are discussed in Chapter 4.

incremental urine bioassay are fraction of injected zirconium excreted per hour. Some bioassay results were declared to be below the detection level (BDL) of the analytical method and were not used in previous studies. The values of BDL measurements and their associated uncertainties are available and are used in this analysis, i.e., no data are excluded. Plots of the urine and blood plasma data are shown in Figure 3.1. The urine data have more scatter that the blood plasma data because there is more zirconium in the blood plasma than in the urine and because the urine samples can have different collection intervals.

3.3 Zirconium Biokinetic Model

The biokinetic model for zirconium shown in Figure 3.2 is given in ICRP Publication 134 (Paquet, Leggett, et al., 2016, Page 273). This model consist of compartments between which the movement of material is described by a system of first-order ODE. The system of ODE can be solved to give the fraction of a unit intake[2] (an injection in

[2]In pharmacokinetics literature the *dose* is the amount of material delivered to an individual, which is referred to as an *intake* here to be consistent with the internal dosimetry literature and to avoid confusion with *radiation dose*.

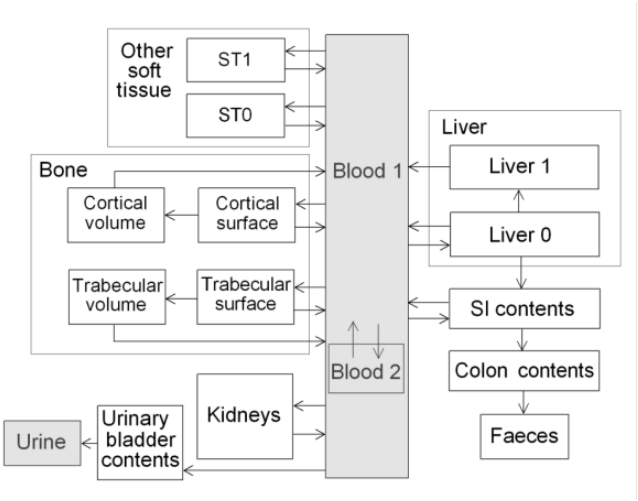

Figure 3.2 ICRP 134 biokinetic model for zirconium.

this case) of the zirconium that is present in each compartment, or any combination of compartments at any time after the intake.

The $n_c = 16$ compartment ICRP zirconium model that describes the rate of change in each compartment can be represented by the following system of n_c ODE

$$\begin{aligned}
\frac{dq_1(t)}{dt} &= \dot{q_1}(t) = -k_{11}q_1(t) + k_{21}q_2(t) + \cdots + k_{n_c1}q_{n_c}(t) \\
\frac{dq_2(t)}{dt} &= \dot{q_2}(t) = k_{12}q_1(t) - k_{22}q_2(t) + \cdots + k_{n_c2}q_{n_c}(t) \\
&\vdots \\
\frac{dq_{n_c}(t)}{dt} &= \dot{q_{n_c}}(t) = k_{1n_c}q_1(t) + k_{2n_c}q_2(t) + \cdots - k_{n_cn_c}q_{n_c}(t),
\end{aligned} \tag{3.1}$$

where $q_i(t)$ describes the content of compartment i at time t, $\frac{dq_i(t)}{dt}$ describes the rate of change in compartment i at time t, and k_{ij} describes the instantaneous rate at which zirconium is transferred from compartment i to compartment j. Note that rate constants for compartments with no connections are 0. For example, $k_{43} = 0$ because there is no transfer of zirconium from the liver-1 compartment to the liver-0 compartment according to the biokinetic model (see Figure 3.2). In fact, for the zirconium model only $n_r = 27$ of the possible 256 of rate constants are non-zero. The rate constants for the ICRP 134 zirconium model are given in Table 3.1.

The rate of change of the i^{th} compartment can be expressed as

$$\dot{q}_i(t) = \sum_{j \neq i} k_{ji} q_j(t) - k_{ii} q_i(t), \tag{3.2}$$

where k_{ii} is the sum of all the rate constants that govern the transfer of material out of compartment i by all routes, and is given by

$$k_{ii} = \lambda + \sum_{j \neq i} k_{ij}, \tag{3.3}$$

where the constant λ is added to account for removal by radioactive decay (Turner, Downing, and Bogard, 2012, page 15; Claudio Cobelli, Foster, and Toffolo, 2007, Section 2.3). For the stable isotopes of zirconium like those used in the Greiter study $\lambda = 0$.

The ODEs in (3.1) can be presented concisely in matrix notation as

$$\begin{pmatrix} \dot{q_1}(t) \\ \dot{q_2}(t) \\ \vdots \\ \dot{q_{n_c}}(t) \end{pmatrix} = \begin{pmatrix} -k_{11} & k_{21} & \cdots & k_{n_c 1} \\ k_{12} & -k_{22} & \cdots & k_{n_c 2} \\ \vdots & \vdots & \vdots & \vdots \\ k_{1n_c} & k_{2n_c} & \cdots & -k_{n_c n_c} \end{pmatrix} \begin{pmatrix} q_1(t) \\ q_2(t) \\ \cdots \\ q_{n_c}(t) \end{pmatrix} \tag{3.4}$$

or $\dot{\mathbf{q}} = \boldsymbol{k}\boldsymbol{q}$, where $\boldsymbol{k}$ is referred to as the rate matrix. The goal is to solve this system of ODE for the quantity of zirconium $q_i(t)$ of each compartment i at time t. For example, $q_1(t)$ describes how much zirconium is in the blood-1 compartment at time t. There are a number of different ways to accomplish this, including Runge–Kutta methods, eigenvalue decomposition, and the matrix exponential (Jacquez, 1985; Polig, 2001; Stan Development Team, 2021b, Chapter 13). Here we will use the *matrix exponential* which is a computationally simple and fast method for solving linear systems of ODE that is available in the MCMC computational software (Stan Development Team, 2021b, pg 181). The matrix exponential solution to the system of ODE is

$$\mathbf{q_t} = \exp(\mathbf{k}t)\,\mathbf{q_0}, \tag{3.5}$$

where $\mathbf{q_0}$ is the initial content of the compartments at $t = 0$, i.e., the initial conditions needed to solve the ODE. Here $\exp(\mathbf{k}t)$ can be expressed using the Taylor series expansion

$$\exp(\mathbf{k}t) = \sum_{j=0}^{\infty} \frac{(\mathbf{k}t)^j}{j!}. \tag{3.6}$$

To calculate (3.6) we use a truncated version that has been shown to be an accurate approximation that avoids numerical instability of raising $\mathbf{k}t$ to high powers (Moler and Van Loan, 2003). The initial quantities $\boldsymbol{q_0}$ of zirconium in the compartments are 1 for blood plasma-1 and 0 for the other compartments because we are assuming an intravenous injection of a unit quantity of zirconium into blood plasma-1. Thus, for this model the initial content of the blood plasma-1 compartment is the intake.

3.3.1 Reference Bioassay Functions and Intake

In this section, we relate the compartment-specific content q_t at time t to the observed blood and urine data discussed in Section 3.2. To this end, we assume the observed bioassay data $\boldsymbol{M}_j$ for the jth subject is represented as

$$\boldsymbol{M}_j = \beta \cdot \boldsymbol{m}_j + \boldsymbol{\varepsilon}_j, \tag{3.7}$$

where $\boldsymbol{\varepsilon_j}$ is the $n_j \times 1$ vector of heteroscedastic regression error with $E(\boldsymbol{\varepsilon}_j) = 0$. The reference bioassay function $\boldsymbol{m}_j = E(\boldsymbol{M}_j|\beta = 1)$ is a function of the q_t's and relates the measured bioassay to the compartment specific contents predicted by the biokinetic model (discussed more below) and β represents the quantity of zirconium injected into the individual (the intake), which is of interest in practice. The estimate of the intake β can be made using the weighted residual sums of squares.

The collection of compartments in the biokinetic model that corresponds to what was measured in a bioassay is referred to a reference bioassay function $\boldsymbol{m}_j = m(\boldsymbol{t}_j)$ at times $\boldsymbol{t}_j$ (Paquet, Leggett, et al., 2016). For blood plasma, the reference bioassay function will be the sum of blood-1 and blood-2 compartments (see Figure 3.2). In

practice, $m_{j,p}$ must account for the volume of the blood plasma compartment, which is incorporated for the subject via

$$\boldsymbol{m}_{j,p} \equiv m_p\left(\boldsymbol{t}_{p_j}, v_j\right) = \frac{q_1\left(\boldsymbol{t}_{p_j}\right) + q_2\left(\boldsymbol{t}_{p_j}\right)}{v_j}, \tag{3.8}$$

where $\boldsymbol{t}_j$ is a vector of time points and v_j is the volume of blood plasma, respectively, for the jth subject (see Appendix B for more details on the calculation of blood volume). Note that $\boldsymbol{m}_{j,p}$ and $\boldsymbol{m}_{j,u}$ are also a function of the rate matrix $\boldsymbol{k}$ and the initial quantities $\boldsymbol{q}_0$ but these parameters are suppressed for the moment to make the notation more concise. The quantity of zirconium in an incremental urine sample has the reference bioassay function

$$\boldsymbol{m}_{j,u} \equiv m_u\left(\boldsymbol{t}_{u_j}, \Delta\boldsymbol{t}_j\right) = q_{12}\left(t_{u_j}\right) - q_{12}\left(t_{u_j} - \Delta t_{u_j}\right), \tag{3.9}$$

where $\boldsymbol{t}_{u_j}$ is the time at which the collection of urine ended and $\Delta\boldsymbol{t}_{u_j}$ is the time interval over which the urine was collected. In practice, an individual collects an incremental urine sample by completely voiding his bladder at time t_1. All subsequent urine voids are collected until a later time t_2, when the individual completely voids his bladder one last time, completing the incremental urine sample. By convention, in the incremental urine reference bioassay function the time of the sample $t_j = t_2$ and the collection interval $\Delta t_{u_j} = t_2 - t_1$. Collection intervals used in the Greiter zirconium study were either 12 hours or 24 hours. The quantity variable $\boldsymbol{m}_j = \left(\boldsymbol{m}_{p_j}, \boldsymbol{m}_{u_j}\right)$ is a $n_j \times 1$ vector that consists of the n_{p_j} plasma reference bioassay functions stacked on top of n_{u_j}urine reference bioassay functions.

3.4 Bayesian Models

The analysis described in the last section becomes more complicated when there exists heterogeneity in the individual level biokinetic models, which results in them not accurately describe the biokinetics of zirconium on an individual. In this section, we

define a Bayesian population model that, with the data from the zirconium study, is used to calibrate the zirconium biokinetic model. This calibration results in estimated posterior distributions for the rate constants in the ICRP 134 zirconium model. Let's begin by looking at the Bayesian approach to modeling data from an individual, after which the population model will be presented.

3.4.1 Individual-level Model

By Bayes Rule, the joint posterior distribution $p(\beta, \boldsymbol{k}, v|\boldsymbol{M})$ of β, $\boldsymbol{k}$, and v given the bioassay data $\boldsymbol{M}$ is

$$p(\beta, \boldsymbol{k}, v|\boldsymbol{M}) = \frac{p(M|\beta, \boldsymbol{k}, v)\, p(\beta, \boldsymbol{k}, v)}{p(\boldsymbol{M})}, \tag{3.10}$$

(Bernillon and Bois, 2000; Bonate, 2011, Chapter 10) where: $p(\boldsymbol{M})$ is the marginal density of the bioassay data $\boldsymbol{M}$; $\boldsymbol{k}$ is the rate matrix; β is the intake; v is the volume of blood plasma; $p(\beta, \boldsymbol{k}, v)$ is the joint prior distribution of β, $\boldsymbol{k}$, and v; and $p(M|\beta, \boldsymbol{k}, v)$ is the likelihood of the data. Thus, the likelihood is

$$p(M|\beta, \boldsymbol{k}, v) = \prod_{j=1}^{n} f\left(M_j, \beta \cdot m\left(\boldsymbol{t}_j, \boldsymbol{k}, v\right), \boldsymbol{\Sigma}_{j,j}\right), \tag{3.11}$$

where $f(\cdot)$ is the pdf for a scaled and shifted Student's t distribution with location parameter $\beta \cdot m(\boldsymbol{t}_j, \boldsymbol{k}, v)$, scale parameter $\boldsymbol{\Sigma}_{j,j}$, and ν degrees of freedom. Assuming that β, $\boldsymbol{k}$, and v are independent of each other, the posterior distribution of the parameters given the observed data is

$$p(\beta, \boldsymbol{k}, v|\boldsymbol{M}) = \frac{p(\boldsymbol{M}|\beta, \boldsymbol{k}, v)\, p(\beta)\, p(\boldsymbol{k})\, p(v)}{p(\boldsymbol{M})}. \tag{3.12}$$

The normalization constant for the posterior, i.e., the constant that makes the posterior a proper probability density function is

$$p(\boldsymbol{M}) = \int\int\int p(\boldsymbol{M}|\beta, \boldsymbol{k}, v)\, p(\beta)\, p(\boldsymbol{k})\, p(v)\ \partial\beta\partial\boldsymbol{k}\partial v. \tag{3.13}$$

All inferences concerning the unknown parameters are derived from their posterior distributions given by (3.12). For example, from the posterior distribution of β a

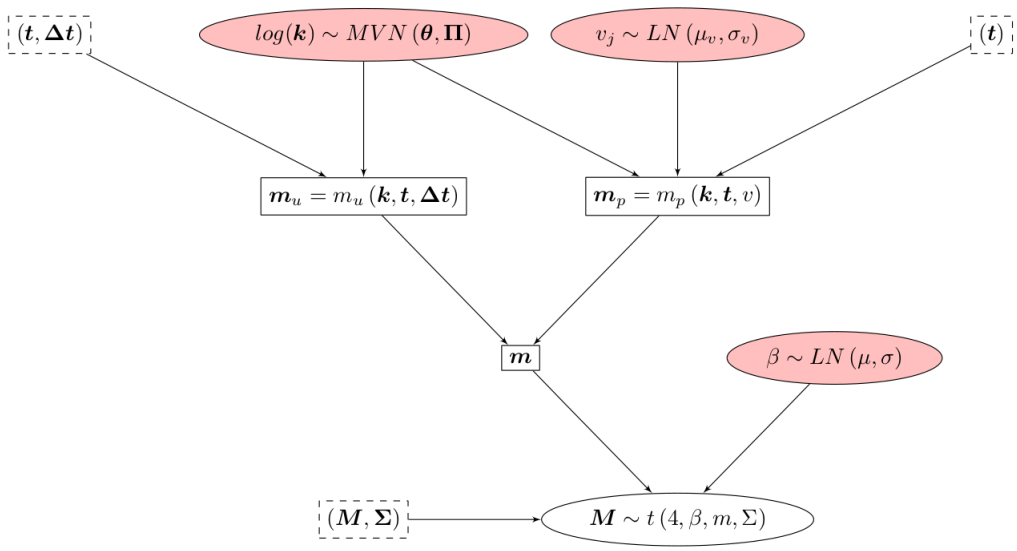

Figure 3.3 Bayesian model for an individual viewed from the perspective of the distributions of the priors (pink shaded nodes), the likelihood (unshaded node), the deterministic calculations (solid rectangle), and the fixed data (dashed rectangles).

point estimate of the intake and its associated uncertainty can be calculated (Hack, 2006, page 243).

In practice the integral given in (3.13) is intractable using conventional numerical and analytical methods and requires specialized methods like Markov Chain Monte Carlo (MCMC) to evaluate. We discuss this further in Section 3.4.5. The Bayes model for an individual has a simple hierarchical structure, which means that the values of some parameters depend on the values of higher-level parameters (see Figure 3.3). In the next section this hierarchical structure will be expanded to accommodate variability between subjects.

3.4.2 Population Level Model

The individual Bayes model can be extended to describe a population of individuals by adding hyperprior distributions, from which the parameters for the priors for an individual are drawn. This is a *population biokinetic model* (Bois, 2001; Wakefield, Aarons, and Racine-Poon, 1999), which allows us to model and account for hetero-

geneity between individuals in a population. The population model can be concisely expressed as

$$p\left(\boldsymbol{\psi}|\boldsymbol{M}\right) = \frac{p\left(\boldsymbol{M}|\boldsymbol{\psi}\right)p\left(\boldsymbol{\psi}\right)}{p\left(\boldsymbol{M}\right)}, \tag{3.14}$$

where the the parameter vector $\boldsymbol{\psi}$ is

$$\boldsymbol{\psi} = \{\boldsymbol{k}, \boldsymbol{\theta}, \boldsymbol{\Pi}, \boldsymbol{\Omega}, \tau, v\} \tag{3.15}$$

and the prior

$$p\left(\boldsymbol{\psi}\right) = p\left(\boldsymbol{k}|\boldsymbol{\theta}, \boldsymbol{\Pi}\right) \times p\left(\boldsymbol{\theta}\right) \times p\left(\boldsymbol{\Pi}|\boldsymbol{\Omega}, \tau\right) p\left(\boldsymbol{\Omega}\right) p\left(\tau\right) \times p\left(v\right) \tag{3.16}$$

The constant of proportionality is the marginal probability density of the data $\boldsymbol{M}$

$$p\left(\boldsymbol{M}\right) = \int p\left(\boldsymbol{M}|\boldsymbol{\psi}\right) p\left(\boldsymbol{\psi}\right) d\boldsymbol{\psi}. \tag{3.17}$$

The likelihood, prior, and hyperprior distributions for the population-level model shown in Figure 3.4 are discussed in the following sections.

3.4.3 Likelihood

The residuals in the regression of $\boldsymbol{M}$ on $\boldsymbol{m}$ in (3.7) often show more dispersion than expected from the normal distribution. This results in measurements frequently being viewed as outliers that skew the fit to the data. The regression errors can be modeled with a Student's t distribution with a low number of degrees of freedom ν, resulting in a regression that is more robust to outliers than a normal distribution (Kruschke, 2015, page 458). A Student's t distribution with $\nu = 4$ has been recommended as a default for robust regression for cases where ν is not estimated from the data (Lange, Little, and Taylor, 1989, pg 883). The sampling distribution given at the bottom of Figure 3.4 is a scaled and shifted Student's t distribution with $\nu = 4$ degrees of freedom

$$\boldsymbol{M}_j \sim t\left(4, \beta_j, \mathbf{m}_j, \boldsymbol{\Sigma}_j\right), \tag{3.18}$$

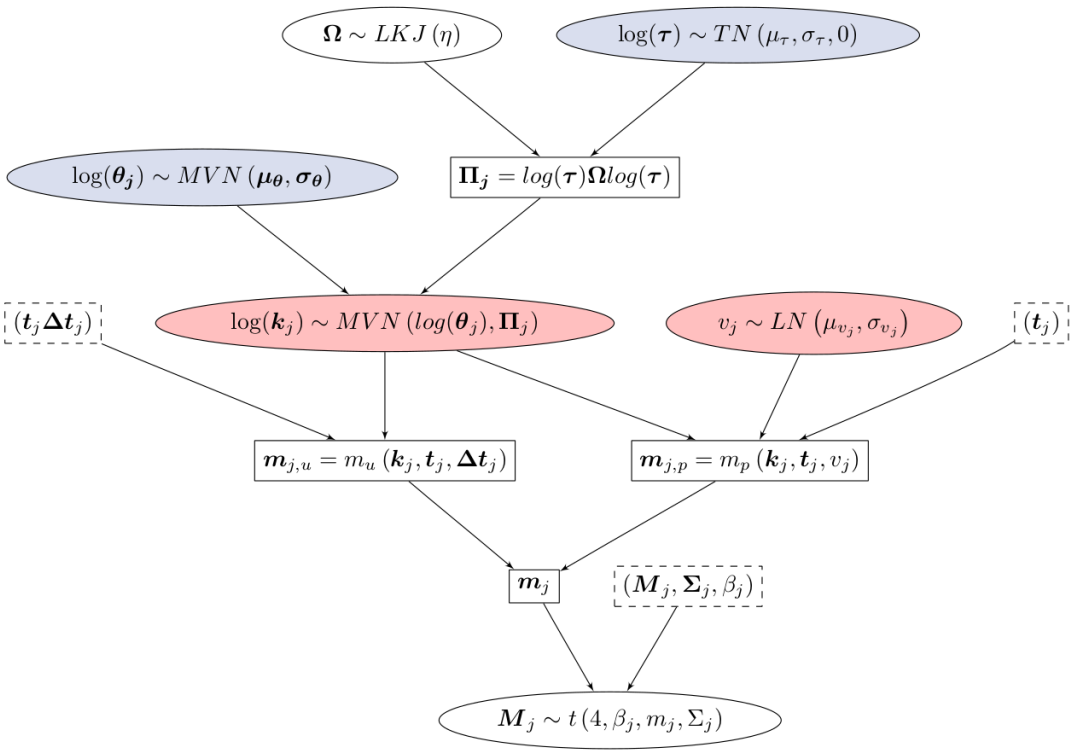

Figure 3.4 Bayesian population model for evaluating cases with the intent of estimating the posterior distributions of the hyperpriors in the blue shaded nodes. The nodes for parameter distributions are denoted by ellipses, computed parameters by rectangles, and data by dashed rectangles.. Note that $\beta = 1$ is a known constant for all subjects in the study and precise point estimates of their blood plasma volume are available.

where $\boldsymbol{\Sigma}_j$ is a $n_j \times n_j$ matrix of squared measurement uncertainties associated with the j^{th} individual's bioassay. For this error model only the main diagonal of $\boldsymbol{\Sigma}_j$ is non-zero, consisting of the n_j squared measurement uncertainties $\boldsymbol{u_j^2}$. The measurement uncertainty is only part of the overall uncertainty in bioassay measurements, and a covariance matrix consisting of only the measurement uncertainties is considered to have the smallest possible variance. One approach to inflating the variances to account for the additional uncertainty is to multiply the covariance matrix $\boldsymbol{\Sigma}$ times a proportionality constant $\sigma^2 \geq 1$ that is an unknown parameter to be estimated (Gelman, Carlin, et al., 2013, pg 372; Kutner, 2005, pg 424). This approach was not implemented here in order to have all the variation in excess of measurement

uncertainty to be propagated to the posterior distributions of the rate constants. The likelihood function for a robust regression with heteroscedastic Student's-t distributed errors is

$$p\left(\boldsymbol{M}|\beta,\boldsymbol{k}\right)=L\left(\beta,\boldsymbol{k}\right)=\prod_{i=1}^{n_j}f\left(4,M_i,\beta m_i,\Sigma_{i,i}\right) \tag{3.19}$$

3.4.4 Prior Specifications

Three types of priors discussed by Banner, Irvine, and Rodhouse (2020, page 883) are

- vague: A prior that is used to reflect that not much is known about the parameter of interest but is well justified.
- weakly informative: A prior that is used to reflect a diluted (or scaled back) amount of knowledge about the parameters.
- informative: A prior that is carefully designed to reflect the current knowledge (and uncertainty) of the parameter.

Vague priors let the data "speak for themselves", allowing one to gauge how much of an influence the prior is having on the posterior distribution. However, vague priors have limited use in this application because they give a posterior that is (more or less) equivalent to a maximum likelihood solution, which is problematic (see Section 2.4). Priors need to be informative enough to make the biokinetic model identifiable and the system of differential equations solvable.

Informative or weakly informative priors are the logical choice simply because we have a lot of prior knowledge about the biokinetics of zirconium. The prior knowledge concerning the biokinetics of zirconium summarized in ICRP 134 (Paquet, Leggett, et al., 2016, pg 269) is based on animal and human studies with zirconium, the biokinetics of its chemical analog niobium, and expert opinion. The compartmental structure

of the ICRP 134 zirconium biokinetic model and its associated rate constants are considered to embody the best available prior information on the biokinetics of zirconium. However, ICRP 134 only provides point estimates of the rate constants, so some generalizations concerning the prior distributions are required.

Priors in biokinetic models are generally assumed to follow normal or lognormal distributions (Davidian, 2006, pg 18). Here we primarily use lognormal priors because the true values of the parameters of interest (e.g., rate matrix, intake, blood plasma volume) are physically constrained to be non-negative. In addition, the multivariate normal (lognormal) distribution is often preferred for modeling multivariate priors because of its somewhat unique properties (e.g., the marginal distributions of a multivariate normal are all normally distributed). In this discussion lognormal distributions are often specified in terms of the geometric mean (gm) and geometric standard deviation (gsd). Specific values reported in the literature for rate constants include: gsd of 2 or 3 for lognormally distributed rate constants and coefficient of variation of 0.3 for normally distributed rate constants (Schmidl et al., 2012, Supplement, pg 8); mean gsd of ~ 1.9 (calculated from 99% confidence intervals for Zr rate constants given in Li, Greiter, et al. (2011a, Table 3)).

The time required to calculate the posterior distribution is a functions of the gsd used in the lognormal priors – the more diffuse the prior the more parameter space that must be searched. The default gsd used for rate constant priors is important because it is approximately the gsd that will be observed in the posterior for rate constants that are not identifiable. Therefore, a default $gsd = 1.5$ is used to limit computer run times (that are already measured in weeks) while not greatly overstating our degree of knowledge concerning these parameters.

Referring to Figure 3.4, the plasma volume v_j for the j^{th} individual has a lognormal prior

$$v_j \sim LN\left(\mu_v, \sigma_v\right), \tag{3.20}$$

with the scale parameter $\sigma_v = \log(1.1)$ is assumed to be constant for all individuals. The location parameter is $\mu_v = \log(V_j)$, where V_j was determined by non-linear regressions of the zirconium plasma concentrations for each individual (see Appendix B). The rate matrix $\mathbf{k}_j$ for the j^{th} individual has a multivariate lognormal prior for non-zero rate constants. That is, the log of the rate constant is modeled as a multivariate normal distribution

$$log(\boldsymbol{k}_j) \sim MVN(\log(\boldsymbol{\theta}_j), \boldsymbol{\Pi}_j), \tag{3.21}$$

where $\log(\boldsymbol{\theta}_j)$ is a $n_r \times n_r$ matrix of population level mean of the rates, and $\boldsymbol{\Pi}_j$ is a $n_r \times n_r$ covariance matrix that describes the dispersion of the rate parameters from the population level means. The population-level location parameters $\boldsymbol{\theta}_j$ have a MVLN hyperprior for non-zero rate constants

$$\log(\boldsymbol{\theta}_j) \sim MVN(\boldsymbol{\mu_\theta}, \boldsymbol{\sigma_\theta}), \tag{3.22}$$

where $\boldsymbol{\mu_\theta}$ is a n_r long vector of location parameters and $\boldsymbol{\sigma_\theta}$ is a $n_r \times n_r$ covariance matrix populated only on the main diagonal. The logs of the ICRP 134 zirconium model rate constants will populate $\boldsymbol{\mu_\theta}$ whereas $\boldsymbol{\sigma_\theta} = \log(1.5)$ as discussed above. The covariance matrix $\boldsymbol{\Pi}_j$ in the prior for $\boldsymbol{k}_j$ is decomposed (Barnard, McCulloch, and Meng, 2000) into the log of the geometric standard deviation τ and a $n_r \times n_r$ correlation matrix $\boldsymbol{\Omega}$

$$\boldsymbol{\Pi}_j = diag(\log(\tau)) \cdot \boldsymbol{\Omega} \cdot diag(\log(\tau)) \tag{3.23}$$

which is equivalent to

$$\boldsymbol{\Pi}_j = \left[\log(\tau)^2\right] \boldsymbol{\Omega} \tag{3.24}$$

since a single value of τ is used. The decomposition of $\boldsymbol{\Pi}$ into τ and $\boldsymbol{\Omega}$ allows us to focus our attention on the geometric standard deviations, for which we might have some intuition, and apply a simplified distribution to the correlations, for which we

probably do not. As a result, $\boldsymbol{\Omega}$ is modeled with a Lewandowski–Kurowicka–Joe (LKJ) distribution (Lewandowski, Kurowicka, and Joe, 2009),

$$\boldsymbol{\Omega} \sim LKJ\left(\eta\right), \tag{3.25}$$

where η is a scalar parameter that controls the dispersion of the off-diagonal elements of $\boldsymbol{\Omega}$. Population model runs were made with η being a parameter to estimate, but it was found that the posterior distribution of this parameter tends to be pretty much the same as the prior distribution and estimating this parameter significantly increases the time required to perform the calculations. As a result, a fixed value of $\eta = 10$ was used to allow for an appropriate distribution of the correlations given the sensitivity of the proposed model (see Appendix E). Finally, the log of the geometric standard deviation $\boldsymbol{\tau}$ is modeled with a normal distribution truncated at 0

$$\log\left(\tau\right) \sim TN\left(\mu_\tau, \sigma_\tau, 0\right), \tag{3.26}$$

which precludes the generation of a negative dispersion parameter. The log mean $\mu_\tau = \log\left(1.55\right)$ and the log standard deviation $\sigma_\tau = \log\left(1.1\right)$ were selected using subject matter knowledge. When estimating the zirconium biokinetic model parameters from the 16 subjects, $\beta = 1$ for each because we know the quantity of zirconium that was injected and the data were subsequently normalized to 1. When β is unknown it is a parameter to be estimated.

One way to assess the suitability of the priors is to consider "are observations drawn from the prior distributions plausible given our knowledge of the field?" Gelman, Simpson, and Betancourt (2017, page 9) explored this idea in detail, stating:

"The guiding principle for prior specification we have emphasized here can be encapsulated in the question: could this prior generate the type of data we expect to see? This is in accordance with the Jaynesian idea that a prior should reflect the known constraints on the system. Rather than looking for hard constraints that are

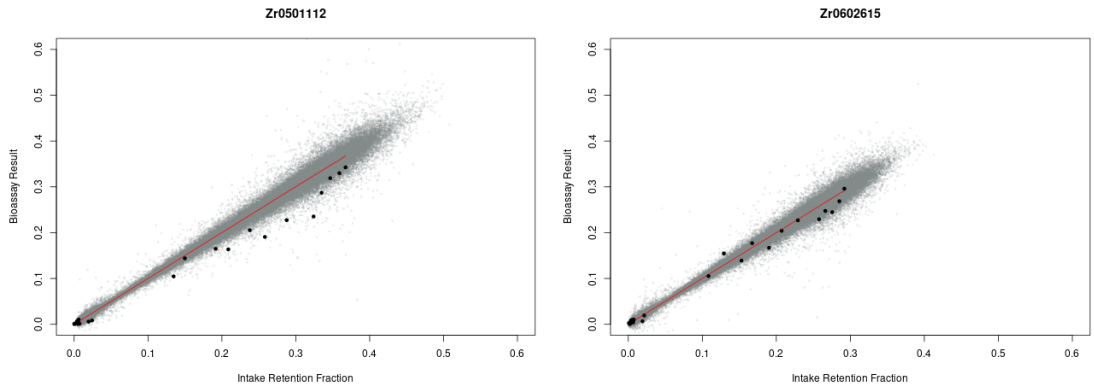

Figure 3.5 Simulations of bioassay data from the prior predictive distribution of the Bayes population model compared the the observed data for two subjects (Zr0501112 on left and Zr0602615 on right). The observed bioassay data are plotted against the mean of the respective reference bioassay function distribution.

difficult to elicit for complex models, however, we instead focus on ensuring that most of the prior mass is in parts of the parameter space that correspond to reasonable data generating processes."

If we have observed bioassay data $\boldsymbol{M}$, $p(\boldsymbol{M})$ in (3.17) is a normalization constant, a number that makes the posterior a proper probability distribution. If the data have yet to be observed, (3.13) is function of the data, the *prior predictive distribution*, that gives the probability of observing a given set of bioassay data. The prior predictive distribution is useful for judging the plausibility of the priors in complex models as discussed by Wesner and Pomeranz (2020), Gelman, Simpson, and Betancourt (2017), and McElreath (2018).

The prior predictive distributions in the Bayes population model for zirconium were compared to the observed bioassay results for the 16 subjects. The examples in Figure 3.5 show that bioassay data simulated from the prior distributions are plausible and not excessively diffuse.

3.4.5 Computational Details

The posterior distributions of the parameters given the data as described by (3.14) reflect all the information contained in the data given the prior distributions. All inferences are made with the posterior distributions of the parameters. In particular, the mean of each posterior distribution is commonly used as the point estimate of the parameter and the width of the posterior distribution is a measure of the uncertainty in the parameter estimate (Hack, 2006, pg 245). The statistical model for zirconium is too complex to obtain analytical solutions to the posterior distributions, so estimates are by drawing samples directly from the posterior distributions using Markov Chain Monte Carlo (MCMC) and analyzing these samples (Gelman, Carlin, et al., 2013).

The "No U-Turn Sampler" (NUTS) variant of Hamiltonian Monte Carlo (Lambert 2018, Ch 15; Monnahan, Thorson, and Branch, 2017; Kruschke 2015, Ch 14) as implemented in the the computer code Stan (Carpenter et al., 2017) is used to generate the samples from the posterior distributions. NUTS is essentially the Metropolis-Hastings sampler where the proposal step incorporates information about the local geometry of the log posterior distribution based on the motion of a particle on the surface of the log posterior. This information allows NUTS to make good proposals with a high acceptance rate, which results in efficient exploration of the complex parameter space. As far as MCMC samplers go, NUTS is somewhat unique in that it provides feedback on whether the sampler is working properly, which is discussed in the next section.

As a particular implementation of NUTS, Stan has the advantages that: (i) the performance of Stan/NUTS is often superior to the more traditional Metropolis or Gibbs MCMC samplers for high-dimensional models with correlated parameters like the zirconium population model (Hoffman and Gelman, 2014); (ii) an efficient matrix exponential ODE solver is part of the Stan language; (iii) the Stan language is compiled rather than interpreted, which makes it faster; (iv) the syntax of Stan

resembles C++ with a pinch of R, which makes it familiar.; (v) Stan is under active development by a team of expert statisticians and programmers.; (vi) Stan has an active support community, which includes direct access to developers.

The Stan MCMC calculations discussed in this dissertation were performed with cmdstan 27.1 interface using the codes and bash scripts discussed in Chapter 5. The computers used were an Intel Core i7-7820X CPU and an Intel Core i9-10900X, both with 128 GB of memory running the Linux Mint 20.2 operating system, and the Hyperion computer cluster at the University of South Carolina[3]. Calculations were performed with one chain per core. A total of 11 chains were run, with each chain started with 3000 iterations for warm up (to tune sampling parameters of NUTS) and 10000 iterations for sampling (total of $110,000$ sampling iterations). The calculations took approximately 20–25 days to complete.

3.4.6 Checking for Convergence of MCMC

Determining whether the Markov chain has converged and the posterior samples are from stationary chains is imperative. If so, we can proceed to the evaluation of the biokinetic model. Adequacy of the MCMC sampling of the posterior is a necessary prerequisite for assessing adequacy of the model.

The No-U-Turn Sampler is somewhat unique among MCMC samplers in that it provides diagnostics that give an indication of whether the process of sampling has been successful (Stan Development Team, 2021b, Ch 17). Perhaps the most useful of these diagnostics is the *divergent transition*, the occurrence of which is an indication that the sampler is not accurately following the contours of the log posterior distribution. No divergent transitions were reported with the Stan MCMC runs used in this analysis.

[3]I would like to acknowledge that the Research Computing program under the Division of Information Technology at the University of South Carolina contributed to the results in this research by providing High Performance Computing resources and expertise.

A properly constructed MCMC sampler is guaranteed to converge to the correct stationary distributions and generate useful samples from the posterior distributions only if it is run for an infinite length of time. There is no way to determine with certainty whether the samples from truncated chains (i.e., run for a finite length of time) are indeed from stationary chains and therefore useful. However, there are diagnostics that can be used to identify MCMC chains that are not stationary. Note that absence of proof is not proof of absence – just because no convergence problems were detected with the diagnostics does not prove that convergence to stationary chains was attained. Statistics used to judge convergence of the MCMC calculations are (Vehtari et al., 2021): potential scale reduction statistic $\hat{R}$, which gives an indication whether the sampling has adequately explored the parameter space of the posterior distribution; effective sample size ESS, which tells us whether we have enough samples of the posterior distributions to calculate stable estimates of the uncertainty in the parameter estimates; Monte Carlo standard error (mcse), which is the uncertainty in the mean of a given parameter estimate. Simply put, we are asking whether the MCMC sampling "worked". If it did, do we have enough information to calculate a stable estimate of the uncertainty in the parameters? If we do, what is the uncertainty and is it small enough to be useful?

Autocorrelation can reduce the amount of information given by Markov chain samples of the posterior distribution, making the number of MCMC samples equivalent to a smaller number of independent samples. The effective sample size is an indication of the degree to which the MCMC samples autocorrelated, giving an an estimate of the number of independent draws from the posterior distribution that give the same amount of information as the dependent draws. The ESS for the tails of a posterior distribution can be different than the center of the distribution, which motivated Vehtari et al. (2021) to propose a *bulk ESS* for the center of a distribution and a *tail ESS* for the tails (see lhs plot in Figure 3.6). They recommended that an

ESS greater than 400 is desirable to get a stable estimate of the posterior distribution. The potential scale reduction statistic $\hat{R}$ is an indication of the mixing of the chain, i.e., the degree to which the chain has explored parameter space and achieved stationarity. $\hat{R}$ compares within chain variance to between chain variance, and if all chains are at equilibrium these variances will be the same with $\hat{R} = 1$. Historically, an $\hat{R} > 1.1$ was taken to indicate poor mixing, but more recent research has suggested that a threshold of $\hat{R} > 1.01$ using a *rank-normalized and folded* version of $\hat{R}$ is more appropriate (Vehtari et al., 2021).

The mcse is the the standard error of the mean of the posterior draws divided by the mean. What constitutes an acceptable mcse is a domain specific decision, but a reasonable cutoff for this application is a relative mcse that is less than 10%, i.e., a relative mcse of ≤ 0.1. The plot on the lhs of Figure 3.7 shows the relative mcse for all parameters. The correlation matrix Ω and the covariance matrix Π can have mean values that are essentially equal to 0, which limits the usefulness of the relative mcse for those parameters. The plot on the rhs of Figure 3.7 omits the correlation matrix and the covariance matrix.

The trace plots can be useful for visualizing chains that are of interest for some other reason, e.g., low ESS. However, they are not of much use as a first-line diagnostic tool, especially if there are a large number of parameters as with the zirconium population model and a large number of iterations as with these Stan runs [4].

[4]Note that Markov chains having negative correlation on odd lags (antithetic chains) can have an ESS that is greater than the actual sample size. This result was observed for a number of parameters in this analysis.

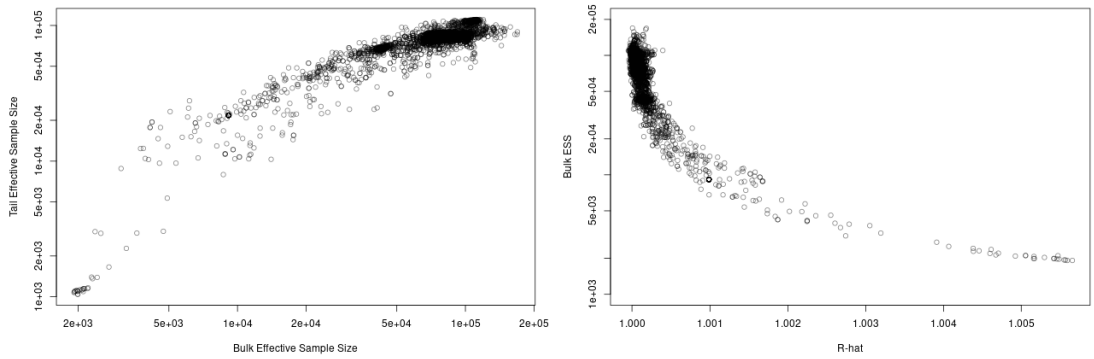

Figure 3.6 Tail ESS versus Bulk ESS (left) and Bulk ESS versus $\hat{R}$ (right).

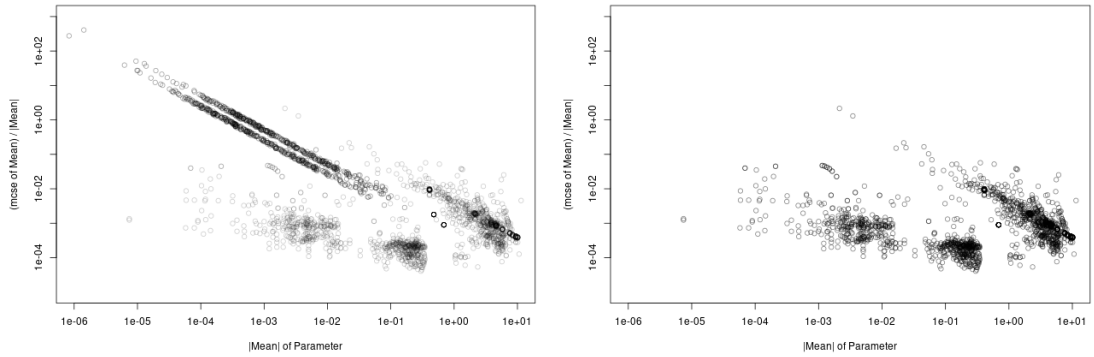

Figure 3.7 Relative Monte Carlo standard error of the mean versus the absolute value of the mean for all parameters (left), and the same plot without the correlation matrix $\mathbf{\Pi}$ and covariance matrix $\mathbf{\Omega}$ (right).

3.5 Results of Data Analysis

The plots of bulk ESS versus $\hat{R}$ and tail ESS versus $\hat{R}$ shown in Figure 3.6 show that $\hat{R}$ calculated in this Stan MCMC run meet the $\hat{R} > 1.01$ criterion and ESS meets the > 400 criterion for all parameters in the zirconium population model. As shown in Figure 3.7, the relative Monte Carlo standard error is < 0.1 for essentially all parameters once the parameters with extremely small means are excluded. Trace plots were examined if the ESS or $\hat{R}$ diagnostics indicated a problem. Finally, Stan did not report any divergences. Thus, there are no indications that the MCMC runs did not converge.

Next, we can explore whether the zirconium population model adequately describes the observed data. A good way to do this is to ask are observations generated using the posterior distributions of the parameters consistent with the observed data? The posterior predictive distribution $p\left(\tilde{\boldsymbol{M}}|\boldsymbol{M}\right)$ is used to answer this question

$$p\left(\tilde{\boldsymbol{M}}|\boldsymbol{M}\right) = \int p\left(\tilde{\boldsymbol{M}}|\boldsymbol{\psi}\right) p\left(\boldsymbol{\psi}|\boldsymbol{M}\right) d\psi, \tag{3.27}$$

(Gelman, Carlin, et al., 2013, pg 7) where $p\left(\tilde{M}|\boldsymbol{\psi}\right)$ is the sampling distribution of future data $\tilde{M}$ given the posteriors $\boldsymbol{\psi}$ of the model parameters, i.e., the likelihood, and $p\left(\boldsymbol{\psi}|\boldsymbol{M}\right)$ is the distribution of the parameters given the observed data $\boldsymbol{M}$, i.e., the posterior distribution. In practice, the observations from the posterior predictive distribution are generated during the the MCMC by using the values of the posterior distributions in a given iteration of the sampling distribution to generate a corresponding set of predicted observations. Examples of posterior predictive distributions obtained for 4 representative subjects are given in Figure 3.8. These plots clearly contrast the measurement uncertainty in the bioassay and the estimate of the total uncertainty that includes biological variability.

The posterior distributions of $\boldsymbol{\theta}$ and τ (the parameters in blue colored nodes in Figure 3.4) are of particular interest because the canned priors discussed in Chapter 4 will be derived from them. The posterior distributions of the blood plasma volumes for each individual are of interest because they are used to help define blood plasma volume priors for out-of-sample subjects.

Marginal density plots of the posterior density of the population level rate constants $\boldsymbol{\theta}$ appeared to be lognormally distributed. The medians and 95% credible intervals for the marginal posteriors of these parameters are given in Table 3.1 where they are compared to the default ICRP 134 rate constants. In Figure 3.9 a plot of the *gsd* for each parameter is presented along with a plot of the median of the parameter versus the default ICRP 134 value. The mean residence time for the urinary bladder, which is the mean length of time a zirconium atom resides there, is shown in Figure

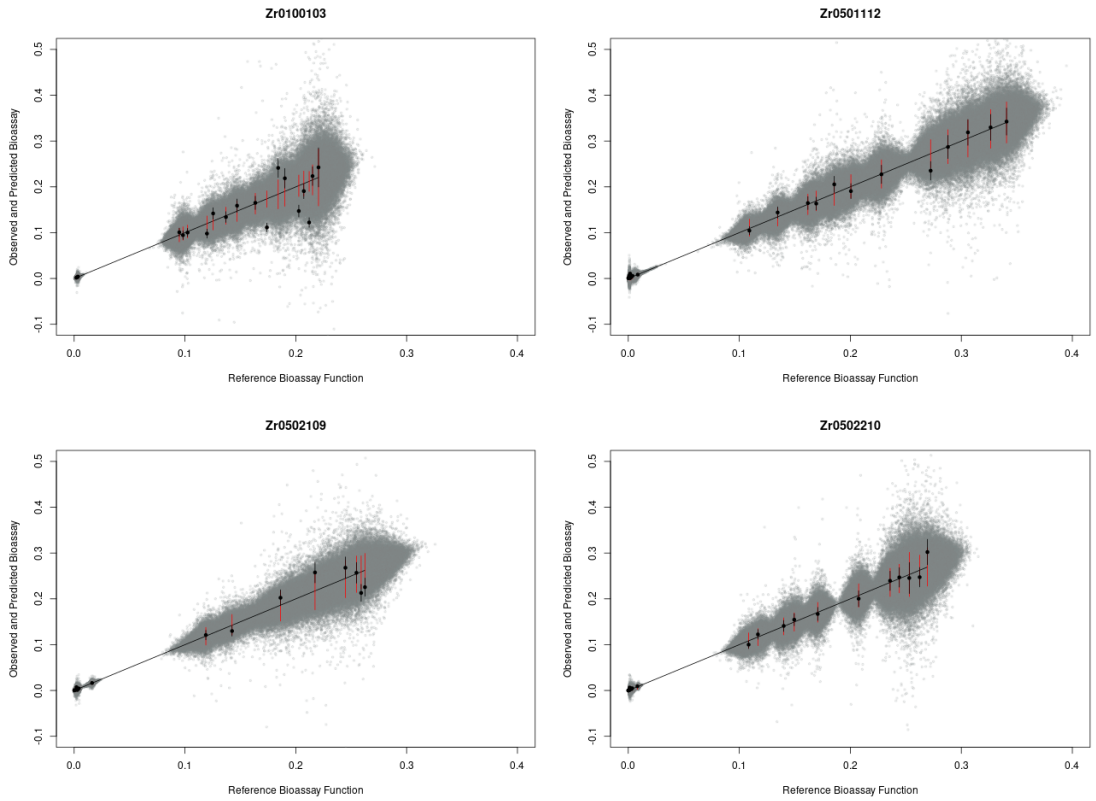

Figure 3.8 Posterior predictive distribution for 4 of the 16 subjects given as an example of posterior predictive distributions. The red error bars are the 95% credible interval of the posterior predictive distribution and the black error bars are the 95% coverage interval for the measurement uncertainty.

3.10. The residence times were calculated as discussed by Claudio Cobelli, Foster, and Toffolo (2007, pg 94) from the individual rate matrix $\boldsymbol{k}$ in each MCMC iteration. The correlation plot for $\boldsymbol{\theta}$ in Figure 3.11 shows that many of the rate constants involving the urine and blood plasma are correlated. Posterior distributions of log blood plasma volumes are summarized in a plot in Figure 3.12 of the *gm* versus the *gsd* of the 16 blood plasma posterior distributions for the study group.

3.6 Discussion

In this chapter we took the ICRP 134 zirconium biokinetics model shown in Figure 3.2, defined priors and hyperpriors for the hierarchical Bayesian statistical model shown

Table 3.1 Transfer rate constants in units of day^{-1} for the ICRP 134 zirconium compartmental model (Paquet, Leggett, et al., 2016, page 274), where the rate constant $k_{13,1}$ was calculated with $f_A = 0.002$. The median and 95% credible interval of the marginal posterior distributions the parameters parameter are given in the last three columns. Complete names of the compartments are given in Appendix F.

						Posterior Rate Constants		
					ICRP 134			
Index	From	Comp	To	Comp	Rate Constants	2.5%	median	97.5%
1	bld1	1	bld2	2	2	2.283	3.283	4.633
2	bld1	1	liv0	3	0.075	0.03349	0.06964	0.1658
3	bld1	1	kid	5	0.0125	0.005859	0.01119	0.02139
4	bld1	1	ST0	6	2	1.892	2.573	3.472
5	bld1	1	ST1	7	0.0375	0.01675	0.03363	0.07349
6	bld1	1	ubc	12	0.1	0.05179	0.07543	0.1089
7	bld1	1	SI	13	0.025	0.01176	0.02223	0.04182
8	bld1	1	ts	8	0.375	0.1583	0.331	0.577
9	bld1	1	cs	10	0.375	0.13	0.2302	0.3821
10	bld2	2	bld1	1	0.462	0.5729	0.8335	1.216
11	liv0	3	SI	13	0.116	0.05678	0.1104	0.2158
12	liv0	3	bld1	1	0.116	0.06112	0.1207	0.2356
13	liv0	3	liv1	4	0.462	0.2347	0.4638	0.9136
14	liv1	4	bld1	1	0.01	0.005075	0.009721	0.01901
15	kid	5	bld1	1	0.01	0.004999	0.009736	0.01926
16	ST0	6	bld1	1	0.462	0.3429	0.5373	0.8349
17	ST1	7	bld1	1	0.02	0.01012	0.02029	0.04056
18	ts	8	bld1	1	0.000493	0.0003595	0.001015	0.001818
19	ts	8	tv	9	0.000247	0.0001229	0.0002368	0.0004588
20	tv	9	bld1	1	0.000493	0.0002536	0.0004932	0.0009616
21	cs	10	bld1	1	8.21E-05	4.106e-05	7.841e-05	0.0001494
22	cs	10	cv	11	4.11E-05	2.117e-05	4.119e-05	8.06e-05
23	cv	11	bld1	1	8.21E-05	4.213e-05	8.218e-05	0.0001602
24	ubc	12	urn	15	12	1.292	2.01	3.237
25	SI	13	col	14	1.5	0.7701	1.502	2.92
26	col	14	fec	16	1.5	0.777	1.506	2.957
27	SI	13	bld1	1	0.0030	0.001546	0.003004	0.005858

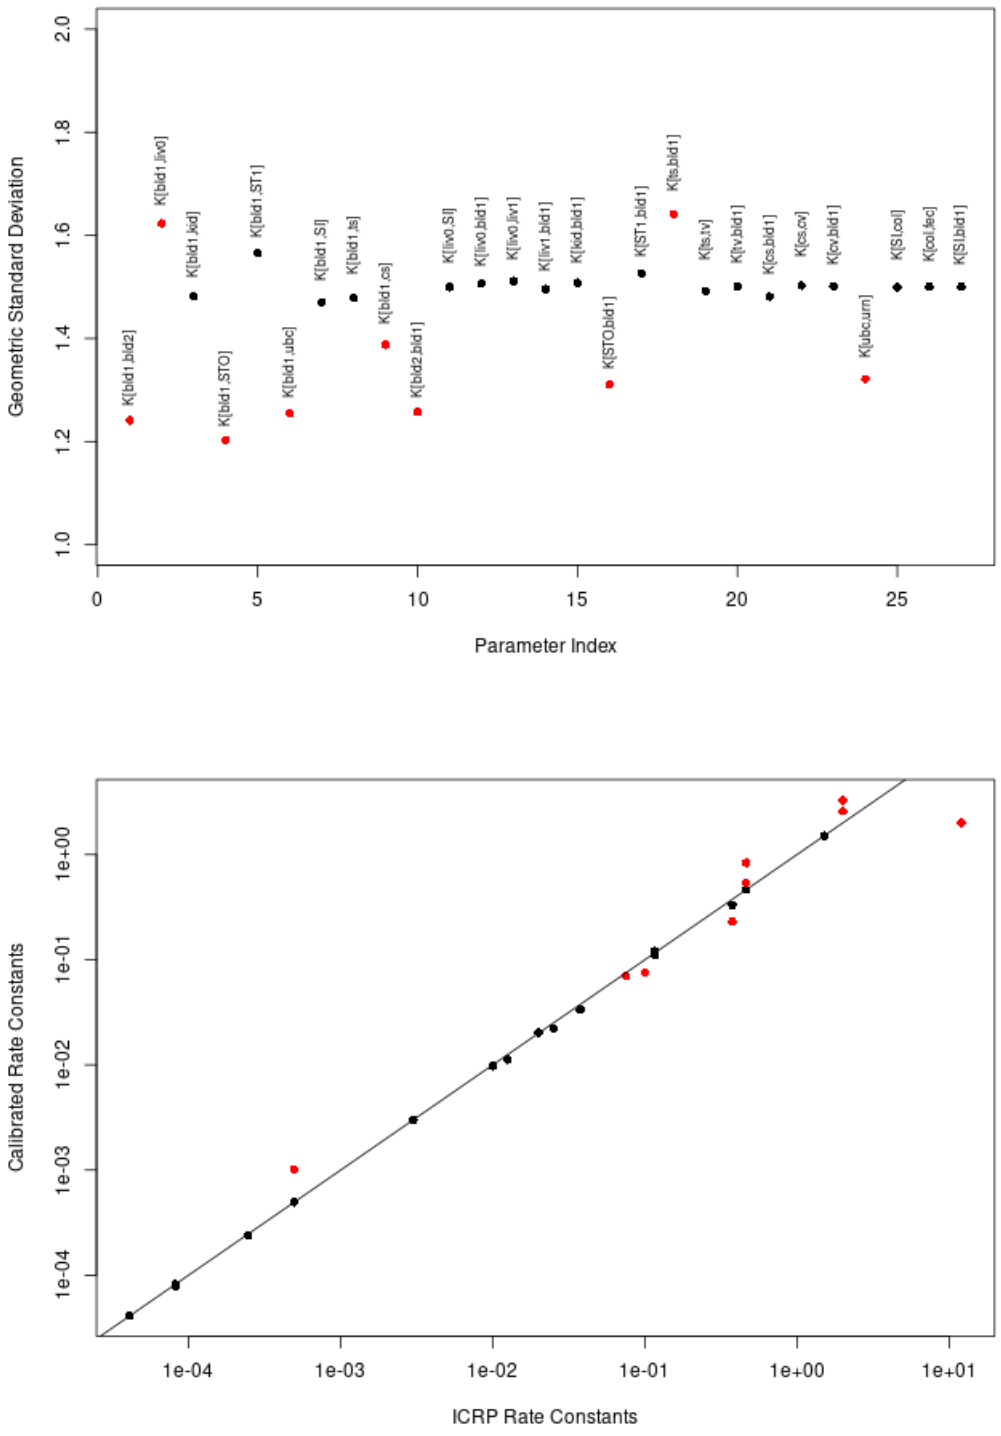

Figure 3.9 Plot of geometric standard deviations of posterior distributions of rate constants $\boldsymbol{\theta}$ (top). For reference, the priors for these parameters had a $gsd = 1.5$. Plot of median of posterior rate constants $\boldsymbol{\theta}$ versus the default ICRP 134 rate constants (bottom). The points in red on both plots have a $gsd < 1.4$ or a $gsd > 1.6$.

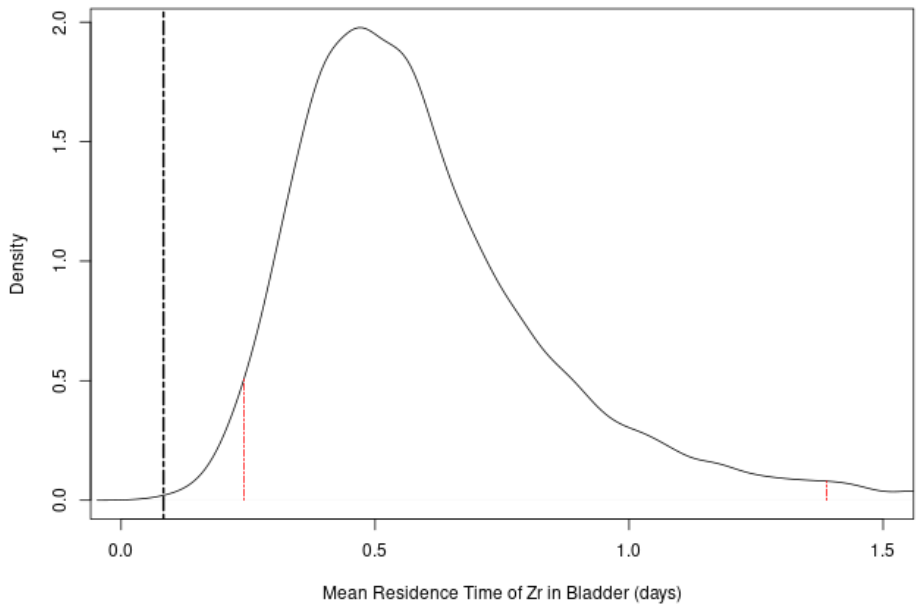

Figure 3.10 Mean residence time of zirconium atoms in the urinary bladder calculated with the standard ICRP 134 zirconium biokinetic model (black dashed line at 0.083 days) and the population model (black curve). The red dashed lines denote the 95% confidence interval of mean residence times calculated with the population model, which has a median of 0.547 days.

in Figure 3.4, and calibrated the biokinetic model with the urine and blood plasma bioassay data shown in Figure 3.1 from 16 subjects with known injection intakes of zirconium. MCMC run with Stan was used to generate samples from the posterior distributions of the parameters that can be used to estimate the mean, uncertainty and 95% credible intervals. These uncertainties incorporate measurement error as well as inherent heterogeneity between individuals in the population. The specification of physiologically realistic and informative priors is crucial because a large number of parameters are not identifiable. The plausibility of priors and hyperpriors defined for use in the population model was checked using prior predictive distribution plots like the ones in Figure 3.5, taking advantage of the generative property of the Bayesian model. The MCMC calculations took $\sim$ 4 weeks, and with such calculations non-convergence of the MCMC calculations is always a concern. The diagnostics used

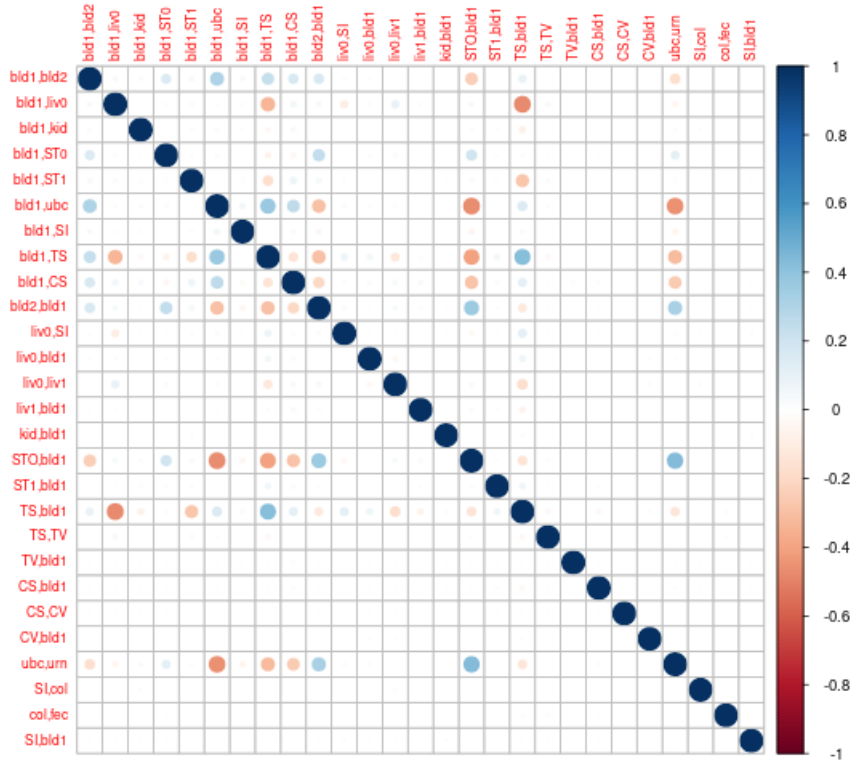

Figure 3.11 Correlation plot for samples from the posterior distribution of $\boldsymbol{\theta}$.

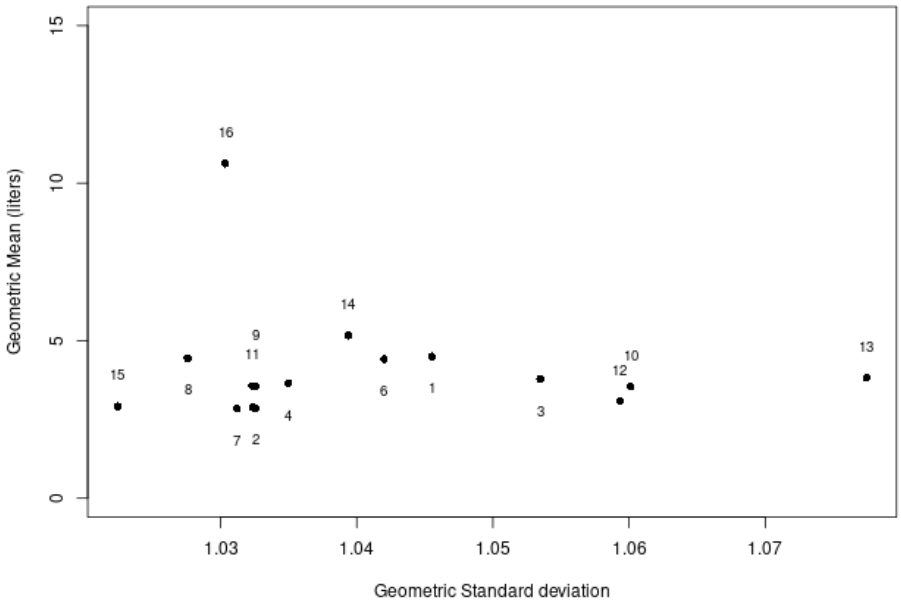

Figure 3.12 Plot of blood plasma geometric mean versus geometric standard deviation for 16 subjects.

to check for convergence problems included the ESS versus $\hat{R}$ plots in Figure 3.6 and relative Monte Carlo standard error plots in Figure 3.7. The information given in these figures along with the absence of any divergences (as reported by Stan) gives some assurance that the calculations converged and the number of samples from the posterior distributions of all parameters is adequate to define the posterior distributions with sufficient accuracy. Posterior predictive distribution plots like those in Figure 3.8 were used to check the plausibility of the posterior distributions by comparing simulated bioassay results generated from the population model with the observed bioassay results. This comparison gives us a good picture of the uncertainty in the predicted bioassay and is the only reality check available.

The parameters of primary interest are the population-level mean rate matrix $\boldsymbol{\theta}$ and the dispersion parameter τ. The posterior distribution of the dispersion parameter τ is well described by a lognormal distribution with a $gm \simeq 2$ and a $gsd \simeq 1.1$. The marginal medians and 95% credible intervals of $\boldsymbol{\theta}$ are tabulated in Table 3.1 where they are compared to the ICRP 134 values of the rate constants. The most significant differences between the posterior distribution of $\boldsymbol{\theta}$ and the ICRP 134 rate constants can be seen Figure 3.9 where the gsd of each θ is plotted versus its identity in the top plot and the median of each θ is plotted against the ICRP 134 rate constant in the bottom plot. The prior distribution for $\boldsymbol{\theta}$ had a $gsd = 1.5$ and the most significant changes in gsd were observed in 9 rate constants involving the blood plasma and urine. The same 9 rate constants had the greatest change in their median values compared to the ICRP 134 values. The posteriors of the other 18 rate constants did not have gm or gsd significantly different than the default values because they don't involve the urine or blood plasma. This is expected because these rate constants are most likely to be influenced by the bioassay data through the likelihood. To improve the estimates of these rate constants we would have to provide information about the

associated compartments in the form of more informative priors derived from previous studies or additional types of bioassay data that involved those compartments.

Of particular interest is the rate constant for the ubc -> urn pathway, which was reduced from the ICRP default of $12\,d^{-1}$ to marginal posterior median of $2\,d^{-1}$. The consequences of this change can be readily grasped by comparing the mean residence time of zirconium atoms in the urinary bladder calculated with the ICRP 134 zirconium biokinetic model with the residence times calculated with the zirconium population model as shown in Figure 3.10. The distribution of mean residence times for the population model does not agree with the residence time obtained from the standard ICRP 134 model. This can in part be attributed to the ICRP urinary bladder model not providing a distribution for the $k[ubc, urn]$ rate constant and the model not being intended for the evaluation of urine bioassay (see Appendix J). This result is nevertheless interesting because the informative prior on $k[ubc, urn]$ was overwhelmed by the data, suggesting that there is something non-trivial occurring. For example, the structure of the ICRP 134 zirconium biokinetic model shown in Figure 3.2 has a single excretion pathway from the blood to the urinary bladder, bypassing the kidney. This structure, combined with observed bioassay on either side of the urinary bladder, may be causing the large shift in $k[ubc, urn]$. This is will be a subject of future research.

Most correlations between posterior values of $\boldsymbol{\theta}$ are in range of approximately -0.4 to 0.4 (see Figure 3.11), which is consistent with the $LKJ\,(\eta = 10)$ prior distribution on the correlation matrix for $\boldsymbol{\theta}$ (see Appendix E). The physical interpretation of these correlations is complex. Consider there is a given amount of zirconium in the blood plasma and urine at a given time. If, for example, the $k[ts, bld1]$ rate constant increases, the $k[bld1, liv0]$ rate constant needs to decrease for the predicted content of these compartments to match the observed contents (i.e., they are negatively cor-

related). Expand this process to all the rate constants and one can see how it might be difficult to give a more detailed physical interpretation of the correlation matrix.

The plot of gm versus gsd of the posterior distributions of the blood plasma volumes for the 16 subjects shown in Figure 3.12 shows that there is not much inter-subject variability in the geometric mean volume except for Subject 16, who has a somewhat implausible blood plasma volume of excess of 10 liters. This volume is consistent with that estimated by other means and is not an anomaly in the Bayesian calculation. In contrast, the intrasubject variability in the estimate of the blood plasma volume as indicated by the geometric standard deviation of the posterior is quite small, ranging from 1.02 to 1.07. Schmidl et al. (2012) appear to have treated the blood plasma volume as a known constant in their Bayesian modeling rather than a parameter to estimate as was done here. The blood plasma volume has a significant influence on the reference bioassay function for blood plasma, so properly accounting for its uncertainty can provide more accurate estimates the uncertainties in the intake and radiation dose.

A population-level Bayesian analysis like the one discussed in this chapter has two primary practical applications in radiation protection. First, the calculation of radiation dose can easily be incorporated into the Stan code. The resulting posterior distribution of the radiation dose for members of the population can be used to establish population exposure limits for radioactive isotopes of zirconium that reflect sensitive members of the population (Bois et al., 1996, pg 1409; USEPA, 2006). Second, the knowledge gained from the population-level calibration can be incorporated into the individual-level evaluation of bioassay data from a new, out-of-study subject. This application is of primary interest in this dissertation. The idea of incorporating knowledge gained from previous studies into the current study is a strength of the Bayesian approach, and in the context of this dissertation there are two main ways of doing this: (i) add the data from the new subject to the data from the 16 study

subjects and repeat the population-level Bayesian analysis, and (ii) approximate the posterior distributions of $\boldsymbol{\theta}$ and τ from the controlled study in closed form, and use them as priors in the individual-level analysis of the new subject. These approaches will be compared in Chapter 4 using the population-level model developed here.

Ultimately, the population-level analysis in this chapter is a proof-of-principle exercise that shows how researchers who develop biokinetic models for the ICRP can go about can go about it in a way that provides realistic probability distributions that incorporate all prior knowledge. In an actual biokinetic model development, prior information from experimental studies and accidental exposures can be combined with expert elicitation (high quality SWAGs) to create a highly informative model. Thus, the goal of this chapter, i.e., to provide a example of a how the ICRP can create more informative an useful biokinetic models, is fundamentally different than that of the researchers at HMGU, LANL, and HPE. HMGU (Schmidl et al., 2012) was concerned with developing new and improved compartmental models for zirconium using data from controlled experimental studies. LANL (Poudel et al., 2018) and HPE (Puncher and Riddell, 2016) were concerned with incorporating Bayesian analyses into production environments, evaluating a large number of subjects with Bayesian models. Li was primarily concerned with sensitivity and uncertainty analysis of the zirconium biokinetic model (Li, Greiter, et al., 2011a) and the cesium biokinetic model (Li, Klein, et al., 2015). The uncertainty analysis of the cesium biokinetic model is notable because it gives an excellent example of how one might go about generating informative prior distributions for the parameters in the biokinetic model.

While the work in this chapter has many strengths, there are a number of weaknesses that warrant further work. First and foremost is the complexity of the calculations and the length of time required to perform them. This issue is ameliorated to some extent because the population level models are designed and evaluated by expert modelers, which means that the typical end user is not exposed to the full

complexity of the process. The time required to perform MCMC calculations can be reduced by optimizing the ODE solvers (e.g., use eigensystems) and to profile the codes to find and eliminate pinch points. In addition, Stan allows for the use of multiple cores for each chain, which would allow full exploitation of the modern multi-core processors and high performance computer clusters. The level of complexity end users are exposed to can be reduced by making MAP calculations available and creating user friendly wrappers for the Stan code a la those given in the *brms* R package by Burkner (2017). Second, to make this work of greater practical use methods need to be presented for coupling the systemic models like the one in Figure 3.2 to models for the respiratory tract and GI tract that feed the systemic compartments. This is a topic of ongoing research.

At the end of the day, all we know for certain about ICRP 134 zirconium biokinetic model is that it is only an approximation to the real world. For example, here we assume that the kinetics of zirconium in the body can be described by the compartmental analysis, the ICRP model has the correct compartmental structure, and that movement of material between compartments is governed by first order processes that are independent of time. In a system as complex as the human body these assumptions are unlikely to hold exactly.

Chapter 4

Development and Use of Canned Priors

4.1 Introduction

If a worker is inadvertently exposed to a radioactive material in the workplace it is often necessary to estimate the quantity of material taken into the body, the *intake*, from quantities of radioactive material measured in the body or excreta, the *bioassay*. Intake, denoted by β, can be estimated with the *standard method* which uses linear regression of the observed bioassay on the *reference bioassay functions*. The reference bioassay functions, denoted by $\boldsymbol{m}$, are deterministic functions designed to describe the dynamics of the assimilated material in the body (ICRP, 2015, pg 32) They are derived from biokinetic models, which are idealized mathematical representation of how a material is deposited in the body, subsequently translocated to various organs and tissues, and ultimately excreted (Boecker, 1998; Harrison, 2009). The International Commission on Radiological Protection (ICRP) publishes biokinetic models that are the international standard, being used to calculate intakes from bioassay data and to calculate permissible concentrations of radionuclides in air and water (Clarke and Valentin, 2009).

The standard method is a straightforward and quick analysis that gives point estimates and confidence intervals of the intake, from which corresponding values of radiation dose can be calculated. However, the standard method has two major flaws: (i) the ICRP biokinetic model used to calculate $\boldsymbol{m}$ may not adequately describe the biokinetics of the material in the person, which results in inaccurate estimates

of intake; (ii) since the standard method assumes deterministic reference bioassay functions, realistic estimates of the uncertainty in the intake can't be calculated. With the standard method, the reference bioassay functions do not vary by person and thereby miss the inherent heterogeneity in ability of people to naturally remove the radioactive material from their body. Further, to produce accurate estimates of uncertainty around the estimated intake we must account for model uncertainty, measurement error in the data, and uncertainty in the parameters. When using the standard method these aspects are commonly ignored, resulting in uncertainty that is underestimated.

Due to the above issues with the standard method, we seek to develop an approach that more effectively shares and uses the current state of knowledge about a biokinetic model. In Chapter 3, a Bayesian hierarchical model was used to calibrate the ICRP 134 zirconium biokinetic model shown in Figure 4.1 with the data from a controlled study where the intakes of zirconium administered to 16 subjects were known (Greiter et al., 2011). The end result of the Bayesian calibration was a version of the ICRP 134 zirconium biokinetic model that incorporated prior knowledge on the distribution of the model parameters, between person heterogeneity, and measurement error, resulting in estimated posterior distributions for the rate constants in the zirconium biokinetic model. In this chapter, we investigate computationally efficient methods to share and leverage the Bayesian population-level calibration results to provide point estimates and uncertainty intervals for individual level model parameters, i.e., intake and radiation dose. The use of previous information is similar in some respects to *transfer learning*, a subfield of machine learning, that allows a population prior to be directly represented in the sampling process for an individual (Christinaki et al., 2021).

The idea of incorporating data from previous biokinetic modeling studies into the priors for the current study is discussed by authors like Wakefield (1996, pg 64) and

Miller (2008, pg 394). In the context of this dissertation there are two ways of doing this: (i) combine the data from the out-of-study subject to the data from the 16 study subjects and repeat the population-level Bayesian analysis, or (ii) approximate the posterior distributions from the population-level study in closed form and use them as priors in the analysis of the out-of-study subject. The first approach is discussed by Gelman, Carlin, et al. (2013, Section 5.1) and is referred to here as an *add-one-in analysis.* This analysis requires access to all the original data used in the calibration and to repeat the calibration process with the addition of one out-of-study subject.

The second approach of using posteriors from the population-level calibration as priors for parameters in the out-of-sample individual-level Bayesian analysis is referred to as *posterior passing* (Brand et al., 2019). Passed posteriors expressed as custom probability density functions have been referred to in the literature as *bespoke priors* [1]. Here we are using familiar distributions like a multivariate normal, so the the passed posteriors will be referred to as *canned priors.* Canned priors incorporate between-person and within-person variability in the rate constants and encapsulate a considerable amount of prior knowledge concerning the biokinetics of a material in a concise form hat can be easily made available to in the open literature. As will be discussed in Section 4.4, canned priors also make it feasible to perform a relatively fast Bayesian evaluation of bioassay data and have the potential to make a Bayesian analysis accessible to individuals with minimal background in statistics.

The canned prior Bayesian model is evaluated using Markov Chain Monte Carlo (MCMC) as implemented in the computer code Stan (Carpenter et al., 2017), which provides full posterior distributions of all parameters. However, once this model is defined point estimates of parameters can be obtained from the maxima of the unnormalized posterior distribution. These quick and computationally inexpensive

[1]The term means "made for a particular customer or user", and in this context refers to a custom probability distribution created for a particular purpose.

maximum a posteriori (MAP) point estimates are superior to those obtained from the standard method but do not provide all the information as the full Bayesian analysis, i.e., no realistic estimates of uncertainty are provided,

A simple example of posterior passing used in an occupational internal dosimetry evaluation was performed by James, Birchall, and Puncher (2008) when they took the posterior distributions of 3 respiratory tract model parameters derived from the analysis of one individual and used them as the prior distributions for the analysis of data from another individual. On a larger scale, Puncher and Riddell (2016) calculated the radiation doses to over 11000 radiation workers using prior distributions for the respiratory tract model parameters derived from a previous study (Puncher, Birchall, and Bull, 2011). The details concerning how the prior distributions were constructed in that study are given in an internal report that is not available in the open literature.

In this chapter, the bioassay data from two out-of-study subjects will be evaluated. These subjects were part of the Greiter et al. (2011) study but were not included in their final analysis and were not used in the Bayesian calibration detailed in Chapter 3. Specifically, we will compare performing an individual-level analysis using using: (i) the standard method, (ii) the add-one-in Bayesian approach, and (iii) the canned prior Bayesian approach. In addition, MAP estimates based on the unnormalized posterior distribution of (iii) will be calculated. These methods are discussed along with the methods used to calculate the canned priors. Then the final canned priors are presented along with the intake estimates from the four methods. The chapter concludes with an assessment of the relative performance of each method and a discussion of their pros and cons.

4.2 ICRP 134 Zirconium Biokinetic Model

The biokinetic model for zirconium shown in Figure 4.1 is given in ICRP Publication 134 (Paquet, Leggett, et al., 2016, Page 273). This model consist of compartments between which the movement of material is described by a system of first-order ODE. The system of ODE can be solved to give the fraction of a unit intake of the zirconium that is present in each compartment, or any combination of compartments at any time after the intake.

The content of each of the $n_c = 16$ compartments in the ICRP zirconium model can be represented by the following system of n_c ODE

$$\begin{aligned}
\frac{dq_1(t)}{dt} = \dot{q_1}(t) =& -k_{11}q_1(t) + k_{21}q_2(t) + \cdots + k_{n_c1}q_{n_c}(t) && (4.1)\\
\frac{dq_2(t)}{dt} = \dot{q_2}(t) =& k_{12}q_1(t) - k_{22}q_2(t) + \cdots + k_{n_c2}q_{n_c}(t) && (4.2)\\
&\vdots\\
\frac{dq_{n_c}(t)}{dt} = \dot{q_{n_c}}(t) =& k_{1n_c}q_1(t) + k_{2n_c}q_2(t) + \cdots - k_{n_cn_c}q_{n_c}(t),
\end{aligned}$$

where $q_i(t)$ describes the content of compartment i at time t, $\frac{dq_i(t)}{dt}$ describes the rate of change in compartment i at time t, and k_{ij} describes the instantaneous rate at which zirconium is transferred from compartment i to compartment j. Note that rate constants for compartments with no connections are equal to 0. The rate constants for the ICRP 134 zirconium model are given in Table 3.1.

The rate of change of the i^{th} compartment can be expressed as

$$\dot{q_i}(t) = \sum_{j \neq i} k_{ji}q_j(t) - k_{ii}q_i(t), \tag{4.3}$$

where k_{ii} is the sum of all the rate constants that govern the transfer of material out of compartment i by all routes, and is given by

$$k_{ii} = \lambda + \sum_{j \neq i} k_{ij}, \tag{4.4}$$

where the constant λ is added to account for removal by radioactive decay (Turner, Downing, and Bogard, 2012, page 15; Claudio Cobelli, Foster, and Toffolo 2007, Section 2.3). For the stable isotopes of zirconium like those used in the Greiter study $\lambda = 0$.

The ODEs in (4.1) can be presented concisely in matrix notation as

$$\begin{pmatrix} \dot{q_1}(t) \\ \dot{q_2}(t) \\ \vdots \\ \dot{q_{n_c}}(t) \end{pmatrix} = \begin{pmatrix} -k_{11} & k_{21} & \cdots & k_{n_c 1} \\ k_{12} & -k_{22} & \cdots & k_{n_c 2} \\ \vdots & \vdots & \vdots & \vdots \\ k_{1 n_c} & k_{2 n_c} & \cdots & -k_{n_c n_c} \end{pmatrix} \begin{pmatrix} q_1(t) \\ q_2(t) \\ \cdots \\ q_{n_c}(t) \end{pmatrix} \tag{4.5}$$

or $\dot{\mathbf{q}} = \boldsymbol{kq}$, where $\boldsymbol{k}$ is referred to as the rate matrix. The goal is to solve this system of ODE for the quantity of zirconium $q_i(t)$ of each compartment i at time t. There are a number of different ways to accomplish this, but here we will use the *matrix exponential* which is a computationally simple and fast method for solving linear systems of ODE that is available in the MCMC computational software (Stan Development Team, 2021b, pg 181). The matrix exponential solution to the system of ODE is

$$\mathbf{q_t} = \exp(\mathbf{k}t)\,\mathbf{q_0}, \tag{4.6}$$

where $\mathbf{q_0}$ is the initial contents of the compartments at $t = 0$, i.e., the initial conditions needed to solve the ODE. The initial quantities $\boldsymbol{q_0}$ of zirconium in the compartments are 1 for blood plasma-1 and 0 for the other compartments because we are assuming an intravenous injection of a unit quantity of zirconium into blood plasma-1. Thus, for this model the initial content of the blood plasma-1 compartment is the intake.

The collection of compartments in the biokinetic model that corresponds to what was measured in a bioassay is referred to a reference bioassay function $\boldsymbol{m}_j = m(\boldsymbol{t}_j)$ at times $\boldsymbol{t}_j$ (Paquet, Leggett, et al., 2016). For blood plasma, the reference bioassay function will be the sum of blood-1 and blood-2 compartments. In practice, $m_{j,p}$ must

account for the volume of the blood plasma compartment, which is incorporated via

$$\boldsymbol{m}_{j,p} \equiv m_p\left(\boldsymbol{t}_{p_j}, v_j\right) = \frac{q_1\left(\boldsymbol{t}_{p_j}\right) + q_2\left(\boldsymbol{t}_{p_j}\right)}{v_j}, \tag{4.7}$$

where $\boldsymbol{t}_j$ is a vector of time points and v_j is the volume of blood plasma, respectively, for the jth subject (see Appendix B for more details on the calculation of blood volume). Note that $\boldsymbol{m}_{j,p}$ and $\boldsymbol{m}_{j,u}$ are also a function of the rate matrix $\boldsymbol{k}$ and the initial quantities $\boldsymbol{q}_0$ but these parameters are suppressed for the moment to make the notation more concise. The quantity of zirconium in an incremental urine sample has the reference bioassay function

$$\boldsymbol{m}_{j,u} \equiv m_u\left(\boldsymbol{t}_{u_j}, \Delta\boldsymbol{t}_j\right) = q_{12}\left(t_{u_j}\right) - q_{12}\left(t_{u_j} - \Delta t_{u_j}\right), \tag{4.8}$$

where $\boldsymbol{t}_{u_j}$ is the time at which the collection of urine ended and $\Delta\boldsymbol{t}_{u_j}$ is the time interval over which the urine was collected. In practice, an individual collects an incremental urine sample by completely voiding his bladder at time t_1. All subsequent urine voids are collected until a later time t_2, when the individual completely voids his bladder one last time, completing the incremental urine sample. By convention, in the incremental urine reference bioassay function the time of the sample $t_j = t_2$ and the collection interval $\Delta t_{u_j} = t_2 - t_1$. Collection intervals used in the Greiter zirconium study were either 12 hours or 24 hours. The quantity variable $\boldsymbol{m}_j = \left(\boldsymbol{m}_{p_j}, \boldsymbol{m}_{u_j}\right)$ is a $n_j \times 1$ vector that consists of the n_{p_j} plasma reference bioassay functions stacked on top of n_{u_j}urine reference bioassay functions.

4.3 Data

Greiter et al. (2011) performed a controlled study where known quantities of non-radioactive zirconium were injected in volunteer subjects. The zirconium in the blood plasma and urine of the subjects was measured at various times afterwards. Subjects Zr0101102 and Zr0102004 were participants in the Greiter study but their data were

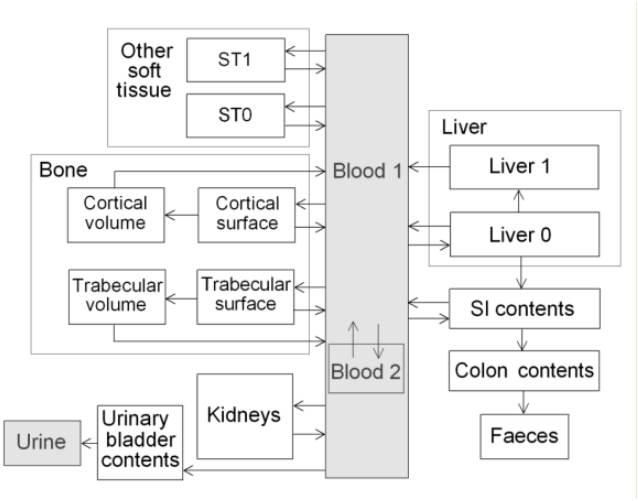

Figure 4.1 ICRP 134 biokinetic model for zirconium.

not used in the development of the zirconium biokinetic model discussed in Greiter et al. (2011). The quantity of injected zirconium is normalized to be $\beta = 1$ for both subjects and their blood plasma volumes are known, so evaluation of the bioassay data from these out-of-study subjects allows us to compare the predictive power and ability to accurately characterize uncertainty of the various approaches for these key parameters. The blood plasma and urine bioassay measurements for these subjects are shown in Figure 4.2, where the error bars are the 95% coverage intervals calculated according to guidance given in the *ISO Guide to the Expression of Uncertainty in Measurement (GUM)* (ISO, 2008).

4.4 Methods

Four different ways of calculating intake from bioassay data are discussed: standard method, add-one-in Bayesian analysis, canned prior Bayesian analysis, and MAP analysis. With reference to the canned prior Bayesian analysis, four different distributions are examined for fitting samples from the posterior distribution of $\boldsymbol{\theta}$ in the population model to create canned priors: multivariate Gaussian mixture (MGM),

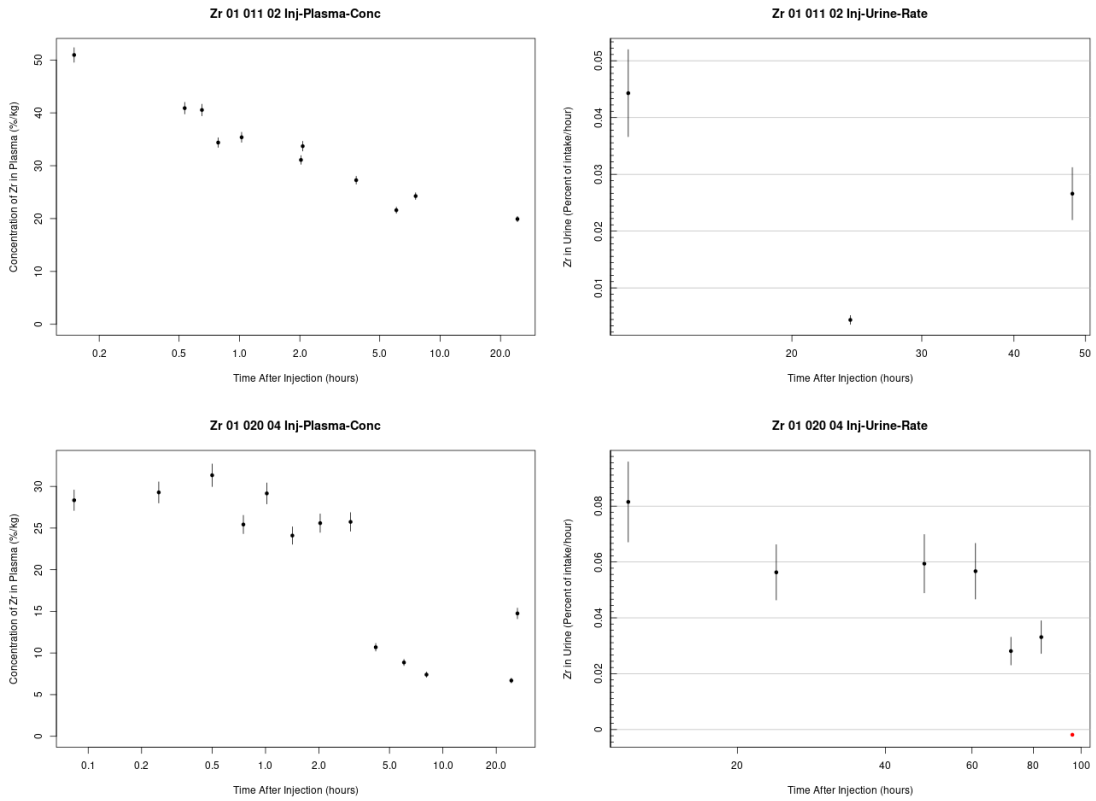

Figure 4.2 Concentration of zirconium in blood plasma (left column) in units of percent of injected dose per kg of blood plasma and urinary excretion rate (right column) of zirconium in urine with units of percent of injected dose per hour for subject Zr0101102 (top row) and Zr0102004 (bottom row). Error bars are the 95% GUM coverage intervals.

multivariate normal (MVN), and two multivariate Student's t distributions – one with 20 degrees of freedom (MVT20) and one with 4 (MVT4).

4.4.1 Standard Method

For n bioassay measurements from an individual we assume the relationship between $\boldsymbol{M}$ and $\boldsymbol{m}$ is represented as

$$M(t_i) = \beta \cdot m(t_i) + \varepsilon_i, \quad i = 1 \ldots, n \tag{4.9}$$

where $\boldsymbol{\varepsilon}$ is the $n \times 1$ vector of heteroscedastic regression error with $E(\boldsymbol{\varepsilon}) = 0$ and $var(\varepsilon_i) = \sigma_i^2$. The values of the reference bioassay functions for zirconium in urine

and blood plasma were calculated by solving a system of ODE that yields predictions for the amount of material in various biological systems at time t as described in Section 4.2. The system of ODE are functions of rate constants and the content of the compartments at $t = 0$, i.e., the initial conditions; the ICRP models use fixed values for the rate constants (this is relaxed in the proposed methods).

The ratio $E\left[M\left(t\right)\right]/m\left(t\right) = \beta$ represents the quantity of zirconium injected into the individual (the intake). The estimate of the intake β can be made using a weighted linear regression of $\boldsymbol{M}$ on $\boldsymbol{m}$, which gives a straight line with a slope equal to the intake. The regression minimizes the weighted residual sum of squares WRSS

$$WRSS = \sum_{j=1}^{n} w_i \left[M\left(t_j\right) - \beta m\left(t_j\right)\right]^2, \tag{4.10}$$

where w_i are the weighting factors that are required because bioassay measurement errors are heteroscedastic (Claudio Cobelli, Foster, and Toffolo 2007, Chapter 10; Skrable, Chabot, et al., 1994). Commonly used weighting factors are analytical measurement uncertainties u or functions of the observed or predicted bioassay measurements (Skrable, French, et al., 2002; Wakefield, 1996; Wakefield, Aarons, and Racine-Poon, 1999). Here, the weight w of each datum as its taken to be its squared relative uncertainty

$$w_i = \left(\frac{u_i}{M\left(t_i\right)}\right)^2, \tag{4.11}$$

because the large difference in the magnitude of the urine and blood plasma measurements. The estimated intake is

$$\hat{\beta} = \frac{\sum\limits_{i=1}^{n}\left[w_i M\left(t_i\right) m\left(t_i\right)\right]}{\sum\limits_{i=1}^{n} w_i \left[m\left(t_i\right)\right]^2}, \tag{4.12}$$

with an estimated standard error

$$\hat{\sigma}_\beta = \sqrt{\frac{\sum\limits_{i=1}^{n} w_i \left(M\left(t_i\right) - \hat{\beta} m\left(t_i\right)\right)^2}{(n-1)\sum\limits_{i=1}^{n} w_i \left[m\left(t_i\right)\right]^2}}. \tag{4.13}$$

The 95% confidence interval of the intake is $\hat{\beta} \pm \left(t_{1-\alpha/2,n-1}\right) \hat{\sigma}_b$, where $t_{1-\alpha/2,n-1}$ is the Student's t quantile with $n-1$ degrees of freedom and a significance level of $\alpha = 0.05$.

4.4.2 Prior for Blood Plasma Volume

The blood plasma volume used to calculate the reference bioassay functions has a significant impact on the estimate intake. In the standard method the blood plasma volume is assumed to be 3 liters for all individuals (see Appendix B). However, prior distributions for blood plasma volume are needed for the add-one-in, canned prior, and MAP analyses. When evaluating the data from an out-of-study individual whose intake and blood plasma volume are not known, the prior for the blood plasma volume is

$$v \sim LN\left(\mu_v, \sigma_v\right) \tag{4.14}$$

where $\mu_v = \log\left(v_r\right)$, with v_r calculated using published regression fits of blood plasma volume to height, weight, and sex as described in Appendix B. The log standard deviation is $\sigma_v = \log\left(1.2\right)$, where geometric standard deviation of 1.2 was chosen to be reasonably informative based on what was observed in the study population (see Figure 3.12) while allowing for relatively large deviations of the true blood plasma volume expected when it is predicted using the published regressions. Subject Zr0101102 is a 59 kg female with a measured blood plasma volume of 1.75 liters and an estimated blood plasma volume of

$$v_r = (24.0 \cdot 59 + 872)/1000 = 2.29\ liters. \tag{4.15}$$

Subjects Zr0102004 is a 78 kg female with a measured blood plasma volume of 3.13 liters and an estimated volume of

$$v_r = (24.0 \cdot 78 + 872)/1000 = 2.74\ liters. \tag{4.16}$$

For each subject the posterior distribution of blood plasma can be compared to the measured value to gauge the predictive power of the model for this important parameter.

4.4.3 Add-One-In Bayesian Analysis

The calibration performed in Chapter 3 provided probability distributions for all the parameters in the ICRP 134 zirconium biokinetic model. The most direct way to perform a Bayesian analysis of an out-of-study subject is to add the subject's data to that of the 16 subjects used in the calibration and repeat the calibration, i.e., perform an add-one-in analysis.

The add-one-in Bayesian model can be concisely expressed as

$$p\left(\boldsymbol{\psi}|\boldsymbol{M}\right) = \frac{p\left(\boldsymbol{M}|\boldsymbol{\psi}\right)p\left(\boldsymbol{\psi}\right)}{p\left(\boldsymbol{M}\right)}, \tag{4.17}$$

where the the parameter vector $\boldsymbol{\psi}$ is

$$\boldsymbol{\psi} = \{\boldsymbol{k}, \boldsymbol{\theta}, \boldsymbol{\Pi}, \boldsymbol{\Omega}, \tau, \beta, \upsilon\} \tag{4.18}$$

and the prior

$$p\left(\boldsymbol{\psi}\right) = p\left(\boldsymbol{k}|\boldsymbol{\theta}, \boldsymbol{\Pi}\right) \times p\left(\boldsymbol{\theta}\right) \times p\left(\boldsymbol{\Pi}|\boldsymbol{\Omega}, \tau\right) p\left(\boldsymbol{\Omega}\right) p\left(\tau\right) \times p\left(\beta\right) \times p\left(\upsilon\right). \tag{4.19}$$

The constant of proportionality is the marginal probability density of the data $\boldsymbol{M}$

$$p\left(\boldsymbol{M}\right) = \int p\left(\boldsymbol{M}|\boldsymbol{\psi}\right) p\left(\boldsymbol{\psi}\right) d\boldsymbol{\psi}. \tag{4.20}$$

A directed graph of this model (which describes both the add-one-in and canned prior analyses) is shown in Figure 4.3. The bioassay $\boldsymbol{M} = \{M_1, \ldots, M_{16}, M_{17}\}$ have the data from the 16 study subjects used to calibrate the ICRP 134 zirconium biokinetic model with the data from the out-of-study subject 17 appended on the end of the vector. The only significant differences between the calibration and the add-one-in calculation are: (i) the data from out-of-study subject are added to that of the

16 study subjects, giving 17 subjects in the analysis; (ii) the out-of-study subject's intake is an unknown parameter to be estimated and is assumed to have a lognormal prior with a $gm = 1$ and a $gsd = 1.2$ whereas the 16 study subjects have known intakes $\beta = 1$, i.e., have a degenerate prior for β; (iii) the geometric mean of the lognormal prior distribution of the blood plasma volume for the out-of-study subject is determined as discussed in Section 4.4.2 whereas the blood plasma volume of the calibration subjects were determined experimentally at the time of the study.

Priors and hyperpriors in biokinetic models are generally assumed to follow normal or lognormal distributions (Davidian, 2006, pg 18), and lognormal distributions are primarily used in (4.19) because the true values of the parameters of interest are physically constrained to be non-negative and tend to be right skewed (see Appendix D). The hyperprior $\boldsymbol{\theta}_j$ describes the population level mean of the rates while the covariance matrix $\mathbf{\Pi}_j$ describes the dispersion of the individual level rate parameters $\mathbf{k}_j$ from the population level means. The covariance matrix $\mathbf{\Pi}_j$ is decomposed (Barnard, McCulloch, and Meng, 2000) into the log of the geometric standard deviation τ and a correlation matrix $\mathbf{\Omega}$, which is modeled with the same $LKJ\,(\eta = 10)$ distribution (Lewandowski, Kurowicka, and Joe, 2009) used for the population-level analysis. The decomposition allows us to focus our attention on estimating the geometric standard deviations, for which we might have some intuition, and then apply a simplified distribution to the correlations, for which we probably do not. The log of the geometric standard deviation $\boldsymbol{\tau}$ is modeled with a normal distribution truncated at 0 and the plasma volume v_j is modeled with a lognormal distribution. The reader is referred to Chapter 3 for a more detailed description of the prior and hyperprior distributions and Section 4.5.2 for a discussion of the how the MCMC code Stan was used used to generate the posterior samples. The reference bioassay functions for urine and blood plasma, $\boldsymbol{m}_u$ and $\boldsymbol{m}_p$ respectively, are calculated as described in Section 4.2. The reference bioassay functions combined with the bioassay data $\mathbf{M}$ and their associated

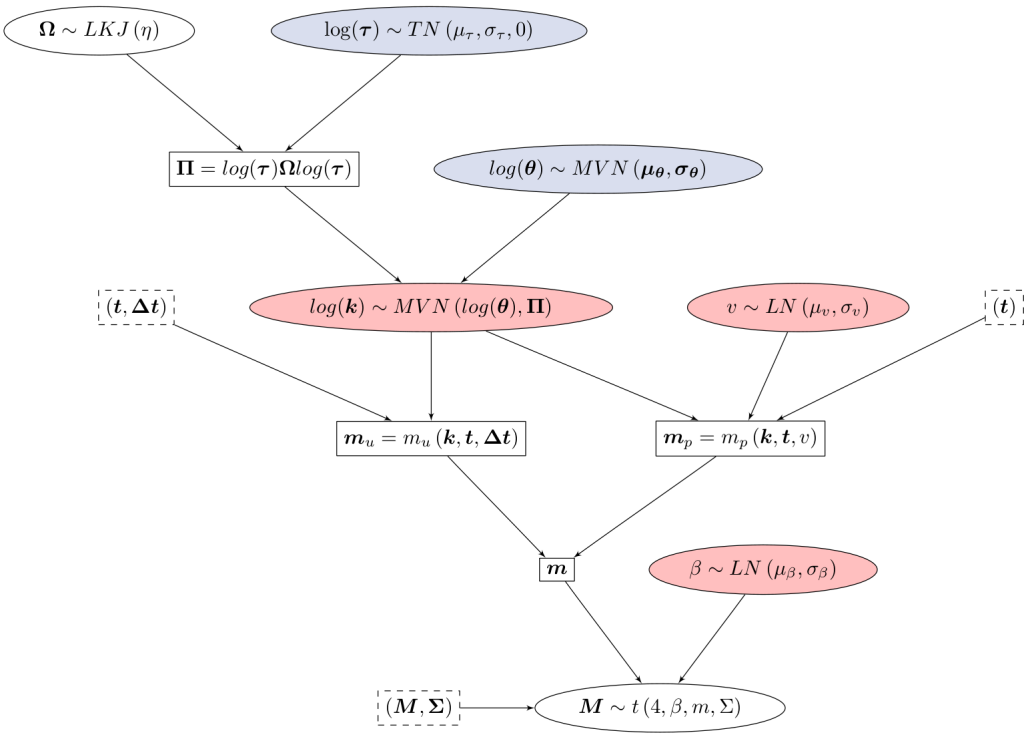

Figure 4.3 Bayesian model for evaluating bioassay data with the intent of estimating the intake, the individual level rate matrix, and blood plasma volume (nodes in pink). The nodes for parameter distributions are denoted by ellipses, computed parameters by rectangles, and data by dashed rectangles. For canned prior Bayesian analysis the nodes in blue are the canned priors.

uncertainties $\boldsymbol{\Sigma}$ comprise the likelihood, which is a Student's t distribution with 4 degrees of freedom.

4.4.4 Canned Prior Bayesian Analysis

In this Section we assume that we have fitted distributions for for the population level parameters $\boldsymbol{\theta}$ and τ that can be easily sampled from, i.e., canned priors. We will discuss how these distributions can be used to get posterior samples of the intake, blood plasma volume, and rate constants for a single subject not used in the calibration of the model parameters.

The canned prior Bayesian model is the same as the add-one-in model defined in equations (4.17) through (4.20) and Figure 4.3 except the parameter vector $\boldsymbol{\psi}$ has

the canned priors $p(\boldsymbol{\theta_c})$ and $p(\tau_c)$

$$p(\boldsymbol{\psi}) = p(\boldsymbol{k}|\boldsymbol{\theta}, \boldsymbol{\Pi}) \times p(\boldsymbol{\theta_c}) \times p(\boldsymbol{\Pi}|\boldsymbol{\Omega}, \tau_c)\, p(\boldsymbol{\Omega})\, p(\tau_c) \times p(\beta) \times p(v)\,, \tag{4.21}$$

and $\boldsymbol{M} = M_{17}$, i.e., the data vector is for the out-of-study subject only. As with the add-one-in analysis, the intake is assumed to have a lognormal prior with a $gm = 1$ and a $gsd = 1.2$. A discussion of the MCMC calculations used to generate the posterior samples is given in Section 4.5.2.

The bioassay data for the two out-of-study subjects were analyzed using the truncated normal canned prior for $\log(\tau_c)$ along with the MGM and MVN canned priors for $\log(\boldsymbol{\theta_c})$. Intake evaluations were not performed with the MVT4 or MVT20 canned prior for $\log(\boldsymbol{\theta})$ because of their relatively poor fit compared to the MGM and MVN distributions. However, if longer tails are desired in the sampling distribution of $\boldsymbol{\theta}$ in a particular analysis, a useful tactic is to plug the parameters for the MVN canned prior into a MVT sampling distribution with the desired degrees of freedom in Stan. This is referred to as an MVT/MVN parameterisation of the MVN canned prior, which was used in an intake evaluation to explore its relative performance. MAP estimates of the parameters in the MVN canned prior model were calculated with the Stan *optimize* function, which utilizes the LBFGS algorithm (Stan Development Team, 2021a, Ch 10).

4.5 Computational Details

In this Section we discuss the estimation of the canned prior, which is basically an exercise in density estimation. The results of this analysis are given in Section 4.6. In addition, details are given on how the add-one-in and canned prior Bayesian MCMC calculations were performed.

4.5.1 Estimation of Canned Priors

In this Section we discuss the estimation of the canned priors using samples from the posterior distributions $\boldsymbol{\theta}$ and τ (the blue shaded nodes in Figure 4.3) from the posterior samples obtained from the population-level analysis in Chapter 3. We want to approximate the distributions of these parameters by fitting the samples from the posterior with probability distributions that can be easily sampled from. These distributions are the canned priors; in the following Section we discuss how they can be utilized to estimate the posterior distributions of the intake β, the blood plasma volume v, and the rate constant matrix $\boldsymbol{k}$ for an individual. Given the distributions for these parameters the distributions for all other quantities of interest like the radiation dose can be calculated.

We estimate the canned priors for $\boldsymbol{\theta}$ and τ separately, as there was no indication of dependence between the posterior samples of these quantities. For τ, we found that posterior samples of $\log(\tau)$ reasonably resembled a normal distribution. This model was fit using the *fitdistr* function from the MASS R package (Venables and Ripley, 2002). For $\boldsymbol{\theta}$, we considered four different multivariate distributions for $\log(\boldsymbol{\theta})$ in an effort to identify the distribution that best fit the samples. First, since the posterior samples of $\log(\boldsymbol{\theta})$ for all parameters were reasonably bell shaped, we considered a 27 dimension multivariate normal distribution (MVN)

$$f(\log(\boldsymbol{\theta})) = \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}), \tag{4.22}$$

where $\boldsymbol{\mu}$ is the vector of location parameters and $\boldsymbol{\Sigma}$ is the scale matrix. This model can be fitted using the *mvn* function from the mclust package (Scrucca et al., 2016). To allow for some potential departures from normality we tested the two component multivariate Gaussian mixture model (MGM)

$$f(\log(\boldsymbol{\theta})) = \pi\mathcal{N}(\boldsymbol{\mu_1}, \boldsymbol{\Sigma_2}) + (1-\pi)\mathcal{N}(\boldsymbol{\mu_2}, \boldsymbol{\Sigma_2}), \tag{4.23}$$

where π is the mixing fraction, and $\{\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1\}$and $\{\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2\}$ are the parameters of the first and second multivariate normal component of the mixture, respectively. MGM have been shown to be able to well represent a large variety of distributional forms (Everitt, 1996; Green, 2018). This model can be fitted using the *mvnormalmixEM* function from the mixtools package (Benaglia et al., 2009). Finally, we considered the scaled, shifted and skewed multivariate Student's t distribution

$$f\left(\log\left(\boldsymbol{\theta}\right)\right) = t_{\nu}\left(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\lambda}\right), \tag{4.24}$$

where ν is the degrees of freedom, $\boldsymbol{\mu}$ is the vector of location parameters, $\boldsymbol{\Sigma}$ is the scale matrix, and $\boldsymbol{\lambda}$ the vector of skew parameters. Two different MVT distributions were fitted, one with $\nu = 4$ degrees of freedom and one with $\nu = 20$ degrees of freedom (MVT4 and MVT20, respectively), using the *mstFit* function from the fMultivar package (Wuertz, Setz, and Chalabi, 2020).

4.5.2 Details of MCMC Calculations

The posterior distributions of the parameters given the data as described by (4.17) reflect all the information contained in the data given the prior distributions. All inferences are made with the posterior distributions of the parameters. In particular, the mean of each posterior distribution is commonly used as the point estimate of the parameter and the width of the posterior distribution is a measure of the uncertainty in the parameter estimate (Hack, 2006, pg 245). The statistical model for zirconium is too complex to obtain analytical solutions to the posterior distributions, so estimates are by drawing samples directly from the posterior distributions using MCMC and analyzing these samples (Gelman, Carlin, et al., 2013).

The "No U-Turn Sampler" (NUTS) variant of Hamiltonian Monte Carlo (Lambert, 2018, Ch 15; Monnahan, Thorson, and Branch, 2017; Kruschke, 2015, Ch 14) as implemented in the the computer code Stan (Carpenter et al., 2017) is used to generate

the samples from the posterior distributions. NUTS is essentially the Metropolis-Hastings sampler where the proposal step incorporates information about the local geometry of the log posterior distribution based on the motion of a particle on the surface of the log posterior. This information allows NUTS to make good proposals with a high acceptance rate, which results in efficient exploration of the complex parameter space. NUTS is somewhat unique in that it provides feedback on whether the sampler is working properly, as discussed below. The Stan MCMC calculations discussed in this dissertation were performed with cmdstan 27.1 interface (Stan Development Team, 2021a) using the codes and bash scripts discussed in Chapter 5. The two computers used had Intel Core i7-7820X and Intel Core i9-10900X CPUs, both with 128 GB of memory running the Linux Mint 20.2 operating system. Calculations were performed with one chain per core. Data preparation and analysis were performed with R 4.1.1 (R Core Team, 2021).

For the add-one-in MCMC analysis, 1 chain consisting of 3000 iterations for warm up (to tune sampling parameters of NUTS) and 10000 iterations was calculated for each subject. For the analysis using the MVN canned priors, 3 chains were calculated for each subject, with 4000 iterations for warm up and 15000 iterations for sampling in each. For the analysis using the MGM canned priors, 2 chains were calculated for each subject, with 4000 iterations for warm up and 15000 iterations for sampling in each. Finally, for the MVN/MVT parameterization of the MVN canned priors, 1 chain was calculated for each subject, with 4000 iterations for warm up and 15000 iterations for sampling in each. The number of chains and samples obtained for a particular canned prior reflects the level of interest in that canned prior, with the MVN being of most interest.

Table 4.1 Comparison of fits to the samples from the posterior distribution of $\boldsymbol{\theta}$, where loglik is the log likelihood, k is the number of parameters in the model, and BIC is the Bayesian Information Criterion.

Model	loglik	k	BIC
MGM	-1.025952×10^6	758	2.060702×10^6
MVN	-1.044802×10^6	378	2.093992×10^6
MVT20	-1.062339×10^6	405	2.12938×10^6
MVT4	-1.113732×10^6	405	2.232166×10^6

4.6 Results

The results of the fits of the 4 different canned prior distributions to the population-level posterior samples of $\boldsymbol{\theta}$ are presented. Next, results of the evaluations of the two out-of-study subjects using the 4 different methods described in Section 4.4 are presented: a summary of the intakes and 95% intervals given in Table 4.2; the M versus m scatter plots in Figure 4.4; density plots of the blood plasma posterior distributions in Figure 4.6.

4.6.1 Canned Posterior Distributions

The canned prior τ_c is the truncated normal distribution

$$\log(\tau_c) \sim TN(\log(1.991), \log(1.061), 0). \tag{4.25}$$

Samples from the posterior distribution of $\boldsymbol{\theta}$ where fit with four different distributions. The results are summarized in Table 4.1 with details of the calculations given Appendix H. The models, in order of lowest to highest BIC – best to worst are: $MGM < MVN < MVT20 < MVT4$.

The skew of the Student's t distributions were essentially 0, i.e., the fitted distributions were symmetric scaled and shifted t distributions. The comparison of BIC for the 4 models indicates that the data are fit better with distributions having normal tails than with distributions having longer tails (i.e., the Student's t distribution with 4 degrees of freedom). Both the MVN and MGM canned priors were used in the eval-

uation of the data from the two out-of-study subjects and it was observed that: the calculations take approximately twice as long with the MGM canned prior versus the MVN canned prior; the MCMC using the MGM prior is less efficient than with the MVN prior, i.e., the effective sample sizes are 10 to 100 times smaller for the MGM prior for a given number of iterations; the distributions of intakes are essentially the same for both priors (see Table 4.2).

4.6.2 MCMC Diagnostics

Stan did not report divergences for any of the add-one-in or canned prior Bayesian MCMC calculations, which is an indication that there are no overt signs of non-convergence. MCMC diagnostics were performed with the R package *posterior* (Vehtari et al., 2021). The effective sample size (ESS), which an indication of the degree to which the MCMC samples are autocorrelated, was > 400 in all cases. $\hat{R}$, which is an indication of non-convergence of the MCMC chains, was < 1.01 for all parameters. The relative Monte Carlo standard error of the mean versus the absolute value of the mean is below 10% for all parameters. Additional details on these diagnostics are presented in Appendix I.

4.6.3 Results of Evaluations Using Different Methods

The results for the standard method analysis of the bioassay from the two out-of-study subjects are shown in the scatter plots of $\boldsymbol{M}$ versus $\boldsymbol{m}$ in Figure 4.4. The estimated intake $\hat{\beta}$ and standard error of the intake $\hat{\sigma}$ are shown on each plot, with the $\hat{\beta}$ and their 95% confidence intervals given in Table 4.2.

The add-one-in analysis took ~ 30 days for each subject whereas the canned prior Bayesian analysis with the MVN prior took ~ 7 hours. For these two analyses: the scatter plots of $\boldsymbol{M}$ versus $\boldsymbol{m}$ for both subjects are shown in Figure 4.4; in Figure 4.5 posterior intake distributions from the add-one-in analyses, the standard method,

Table 4.2 Summary of intake estimates for subjects Zr0101102 and Zr0102004 calculated using: the standard method (Standard); MGM, MVN, and MVT/MVN canned Bayesian analysis; add-on-in Bayesian analysis (Bayes AOI), and maximum a posteriori (MAP) analysis. Uncertainties for MAP are not reported by Stan. The true value of the intake is $\beta = 1$ for both subjects.

		Quantiles		
Subject	Method	0.025	0.5	0.975
Zr0101102	MGM	0.7826	1.008	1.3
Zr0101102	MVN	0.778	1.007	1.304
Zr0101102	MVT/MVN	0.7823	1.009	1.31
Zr0101102	Standard	1.148	1.263	1.378
Zr0101102	Bayes AOI	0.8504	1.05	1.309
Zr0101102	MAP		1.031	
Zr0102004	MGM	0.7538	0.9692	1.247
Zr0102004	MVN	0.7546	0.9706	1.254
Zr0102004	MVT/MVN	0.7527	0.9658	1.246
Zr0102004	Standard	0.7686	0.8549	0.9412
Zr0102004	Bayes AOI	0.7498	0.9129	1.122
Zr0102004	MAP		0.9249	

and the canned prior analysis are compared; the intake estimate and its 95% credible interval are given in Table 4.2; in Figure 4.6 the posterior distributions of the blood plasma volumes for both subjects are compared to the true values. The scatter plots of $\boldsymbol{M}$ versus $\boldsymbol{m}$ for both subjects using the MAP analysis of the MVN canned prior posterior distribution are given in Figure 4.4 and the intake estimate in Table 4.2.

The canned prior Bayesian analysis performed with the MGM prior took $\sim$ 15 hours and the intake evaluation using MVT/MVN parameterization of the MVN prior took $\sim$ 7 hours. The intakes and credible intervals obtained from evaluations of bioassay data preformed with the MGM and MVT/MVN evaluations are shown in Table 4.2 and the scatter plots of $\boldsymbol{M}$ versus $\boldsymbol{m}$ for both subjects are shown in Figure 4.4.

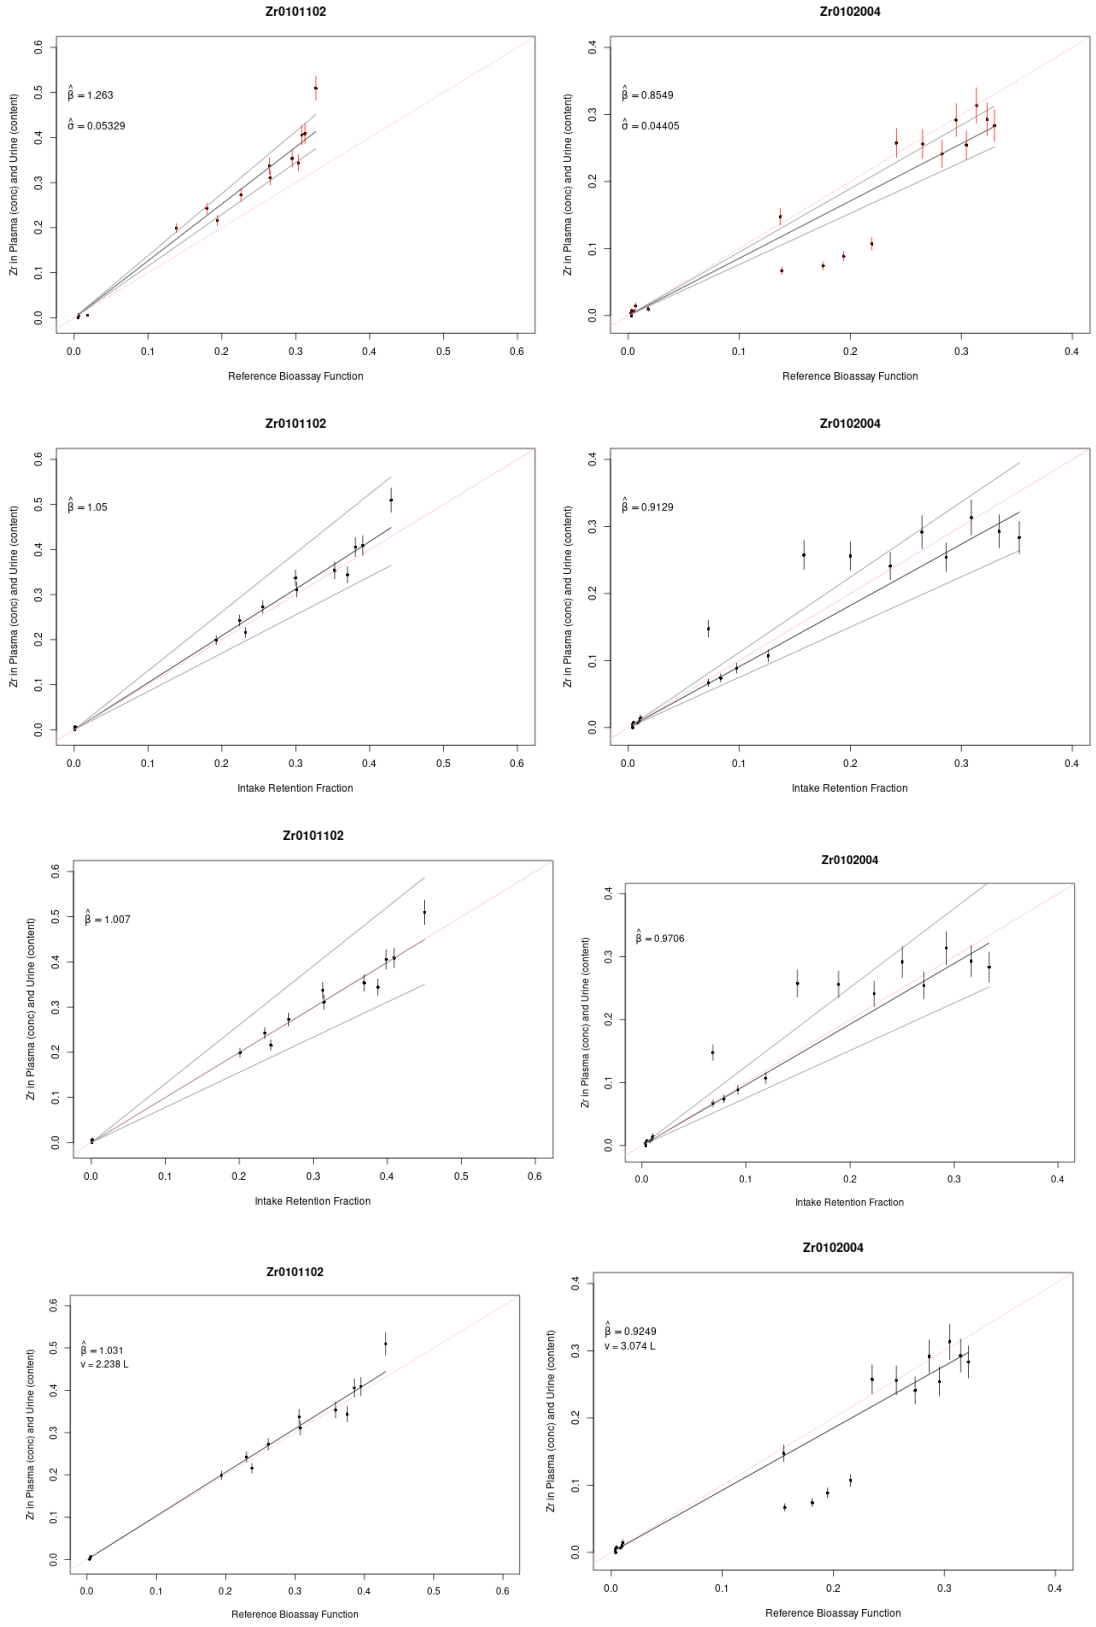

Figure 4.4 Comparison of results for Subject Zr0101102 (left column) and Subject Zr0102004 (right column) using, from top to bottom: standard method, add-one-in, MVN canned prior, MAP.

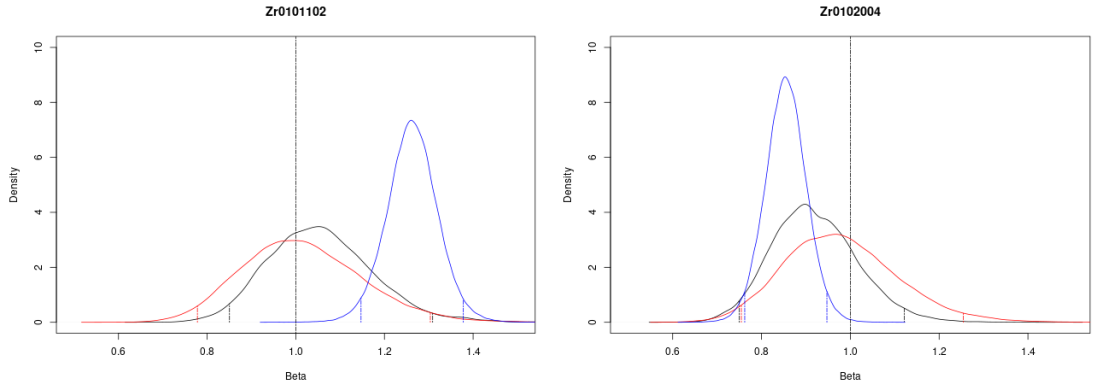

Figure 4.5 The posterior distribution of $\hat{\beta}$ calculated for Zr0101102 (left) and Zr0102004 (right) with standard method (blue), add-one-in Bayesian analysis (black), and MVN canned prior Bayesian analysis (red). The true value of the intake is $\beta = 1$ for both subjects. The vertical lines on the distribution of β in denote the 95% confidence or credible interval.

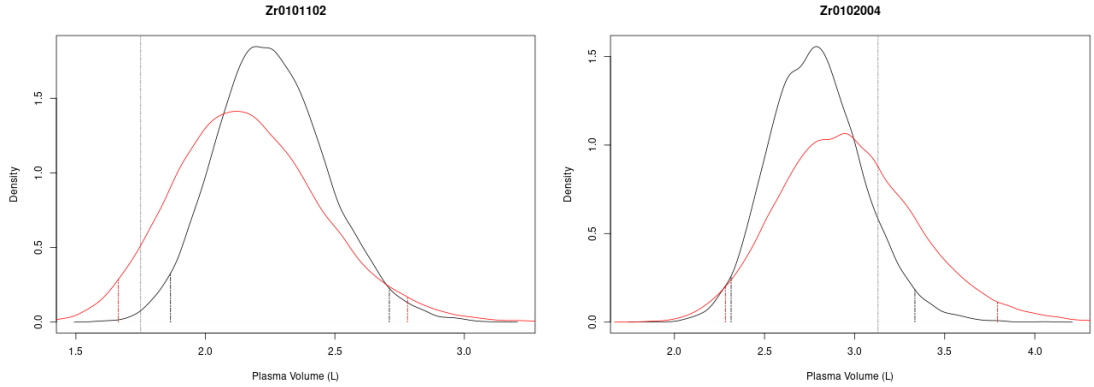

Figure 4.6 Posterior distribution of blood plasma volume for subjects Zr010102 (left) and Zr0102004 (right) from MVN canned prior analysis (red) and add-one-in analysis (black) . The vertical dashed lines on the distributions of denote the 95% credible interval and the vertical dashed line is the true blood plasma volume.

4.7 Discussion

In this chapter we addressed the problem of evaluating bioassay data from a specific individual who has been exposed to radioactive materials. The bioassay data from out-of-study subjects Zr0101102 and Zr0102004 shown in Figure 4.2 were analyzed using four methods: (i) the standard method that consists of a linear regression of $\boldsymbol{M}$ on $\boldsymbol{m}$, where $\boldsymbol{m}$ is calculated from the unmodified ICRP 134 zirconium biokinetic model; (ii) the add-one-in Bayesian analysis where data from the out-of-study subject is combined with the data from the 16 study subjects and the calibration described in Chapter 3 repeated; (iii) the canned prior Bayesian analysis where familiar distributions are fit to posterior samples from the calibration are used as priors in a Bayesian analysis of a single subject; (iv) MAP analysis of the posterior distribution from the canned prior Bayesian analysis. In the development of the canned priors, 4 different multivariate distributions were fit to the log of the posterior samples from $\boldsymbol{\theta}$ in an effort to identify the distribution that best fit the samples.

The canned Bayesian analysis and comparison of it with the standard method is of primary interest, and most of this discussion will focus on that comparison. The add-one-in Bayesian analysis is included primarily to highlight the advantages of the canned prior Bayesian analysis over the add-one-in. The MAP analysis illustrates an inexpensive and potentially useful way to calculate intakes if estimates of uncertainty are not needed.

The relative performance of each method is judged by its ability to recover the true values of the intake and blood plasma volume. In other words, are the true values of these parameters in the 95% intervals? The analyses are also compared on the agreement between the predicted and observed bioassay, i.e., is there lack-of-fit? For the standard method only the intake and lack-of-fit checks are performed, and with MAP analysis only the lack-of-fit check is performed. Finally, the time required to

perform each analysis is an important metric in practice that will be used to compare the performance of the methods.

4.7.1 Analysis with the Standard Method

The standard method of regressing $\boldsymbol{M}$ on $\boldsymbol{m}$ is a simple and quick analysis used in occupational internal dosimetry to give estimates of the intake (IAEA, 2004; Skrable, Chabot, et al., 1994). This method was used on the bioassay data for the two out-of-study subjects, the results of which can be seen in the $\boldsymbol{M}$ versus $\boldsymbol{m}$ scatter plots in the first row of Figure 4.4. Outliers are seen in the bioassay data for both out-of-study subjects. Subject Zr0101102 has 3 urine samples and Zr1012004 has 4 blood plasma results that do not fall near the regression line. This is likely the result of a combination of measurement error in those bioassay results that is not accounted for in their stated uncertainties and the biokinetics of zirconium in the individuals being different than the biokinetics specified by the ICRP 134 zirconium model, i.e., model misspecification.

In the linear regression the sum of the weighted squared difference between the observed and predicted bioassay results are minimized. As clearly shown by (4.13), the confidence intervals on the slopes shown in the plots are a function of this difference and the the measurement uncertainty. The uncertainty in the intake does not consider the uncertainty in the reference bioassay functions, which are treated as known constants. This means that the 95% confidence intervals on the intake are too narrow because they do not account for the uncertainty in the reference bioassay functions. Nevertheless, since the 95% confidence intervals on the intakes for both subjects do not contain the true value of the intakes from these plots we conclude that the intake estimates are inaccurate because of misspecification of the biokinetic model.

In most practical applications, the shortcomings of the standard method can be safely ignored because the radiation doses are usually much lower than the applicable regulatory limits and don't warrant further refinements. However there are cases where radiation doses exceed regulatory limits and steps should be taken to reduce model misspecification and to estimate realistic uncertainties (NCRP, 2010, pg 2), which is the motivation behind the Bayesian analyses.

4.7.2 Add-One-In Bayesian Analysis

If a Bayesian analysis has been performed on a given dataset and then additional data are obtained, one can update the analysis by adding the new data to the existing dataset and repeating the analysis (Gelman, Carlin, et al., 2013, Section 5.1). Thus, the Bayesian analysis of an out-of-study subject can be performed by adding the that subject to the 16 calibration study subjects and performing the population level Bayesian analysis described in Chapter 3, which is referred to here as an add-one-in analysis. Overall, the results of the add-one-in analysis are similar to those obtained from the MVN canned Bayesian analysis discussed in the next section, but take much longer than the canned prior Bayes analysis to run (weeks versus hours).

The intakes and blood plasma volumes of the two out-of-study subjects are known because they were part of the study by Greiter et al. (2011), but they were not used by Greiter and were not part of the calibration in Chapter 3. To test the predictive power of the add-in-in analysis the intakes of these two out-of-study subjects are treated as unknown parameters to be estimated and their blood plasma volumes are estimated using published regressions. Treating the intake of one subject as an unknown parameter and the others as known constants required a minor change in the Stan code used for the calibration. The plots of $\boldsymbol{M}$ versus $\boldsymbol{m}$ in the second row of Figure 4.4 show that the true value of the slope, i.e., $\beta = 1$, is in the 95% credible interval for the estimated slope for both subjects. Comparisons of the intake distri-

butions calculated with the standard methods are compared to intake distributions calculated using add-one-in Bayesian analysis in the first row of Figure 4.5. Estimated intakes and associated 95% credible intervals are given in Table 4.2. The true value of the intakes are in the 95% credible intervals of the intakes calculated with the add-on-in Bayesian method, which is not the case for the intakes calculated with the standard method.

The distribution of blood plasma volumes calculated in the add-one-in analysis were compared to the true blood blood plasma volumes determined experimentally. The true value of the blood plasma volume is in the 95% credible interval of blood plasma volume calculated for subject Zr0102004 but not for subject Zr0101102 (see Figure 4.6). The blood plasma volume has a direct influence on the calculation of the reference bioassay function for blood plasma. Thus uncertainty in the blood plasma volume can be a significant contributor to the uncertainty in many parameter estimates.

4.7.3 Canned Priors

The main problems with the add-one-in Bayesian analysis are: in practice the internal dosimetrist will not have access to the data used to calibrate the ICRP biokinetic model, and it is expensive in terms of time and resources to repeat the calibration procedure in order to get results for one subject. Rather than repeating the calibration we can summarize the posterior distributions of key parameters probability distributions that are easily sampled and use them as priors for the out-of-subject analysis. The procedures used to create these canned priors for the rate matrix $\boldsymbol{\theta}$ and dispersion parameter τ (the purple colored nodes in Figure 4.3) are described in Section 4.6.1.

The distribution of samples from the posterior of $\log(\tau)$ are fit well with a normal distribution given in (4.25). The fits of four different multivariate distributions to the

posterior samples of $\log(\boldsymbol{\theta})$ were considered: multivariate Gaussian mixture (MGM) with 2 components, multivariate normal distribution (MVN), skewed t distribution with 20 degrees of freedom (MVT20), and a skewed t distribution with 4 degrees of freedom (MVT4). The distributions with normal tails (MGM and MVN) fit the data better than the distributions with long tails (MVT4 and MVT20) using the Bayesian Information Criterion to measure the goodness of fit. The canned prior constructed with MVN has essentially the same means and standard deviations as the posterior samples from the population-level model for all parameters, indicating a good fit was achieved. The MVN and MGM canned priors produced estimates of intake and predicted bioassay distributions (see Table 4.2) that were essentially the same, but models using the MVN canned priors were sampled much more efficiently by Stan than models using the MGM canned priors and ran twice as fast. Thus, most of the discussion in this chapter focuses on the MVN canned priors because of their relative simplicity and resulting ease of presentation. However, MGM canned priors will be considered in the future for use with biokinetic models for materials other than zirconium (e.g., plutonium) where the addition flexibility of the MGM distribution might be advantageous.

4.7.4 Canned Prior Bayesian Analysis

The individual-level canned Bayesian analyses were considerably less complex than the add-on-in Bayesian analysis performed. Stan efficiently generated samples from the posterior distributions in a matter of hours that resulted in precise parameter estimates. The results of the canned Bayesian intake analysis can be seen in the $\boldsymbol{M}$ versus $\boldsymbol{m}$ scatter plots in the third row of Figure 4.4 show that the true value of the slope, i.e., $\beta = 1$, is in the 95% credible interval for the estimated slope for both subjects. This is made even more obvious in Figure 4.5, where the probability distributions of the intakes calculated by the canned Bayesian analysis are compared to the

intake distribution of intakes calculated with the standard method. The 95% credible intervals of the canned prior Bayesian analysis are wider than the 95% confidence intervals produced by the standard method because they incorporate the uncertainty in the biokinetic model and the data model and are thus more realistic representations of the true uncertainty in the intake estimates.

The main cause of the inaccurate estimate of intake for Zr0102004 calculated by the standard method is the 4 blood plasma outliers. This can be seen by comparing the plots in the first row of of Figure 4.4 (standard method) with those in the third row (canned Bayesian method). The 4 blood plasma outliers in the standard method plot are shifted to the left in the canned Bayesian analysis plot, resulting in a better overall fit to the data. This horizontal shifting is the result of the rate constants in the zirconium biokinetic model being modified during the Bayesian analysis in order to achieve a better fit to the data. This disciplined modification is possible because the rate constants in the biokinetic model have a probability distribution rather than being constants and informative priors are used. Outliers like these are commonly observed in bioassay data and can be caused by some combination of measurement errors and systematic differences between the biokinetics of the individual and the biokinetics of the ICRP model. The choice of the data model can also be related to whether data are considered outliers, so a Student's t distribution was chosen for the data model (likelihood) rather than a normal distribution in order to better accommodate outliers.

The inaccurate intake estimated by the standard method for Zr0101102 is caused primarily by her small blood plasma volume compared to that used by the ICRP 134 biokinetic model. This can be seen in the second row of Figure 4.6 where the posterior distributions of blood plasma volume for both subjects are compared to their true blood plasma volume. The posterior distribution of blood plasma volume for subject Zr0102004 estimated in the canned Bayesian analysis has a median of

$\sim$ 3.1 liters, which is quite close to her true blood plasma volume of 3.13 liters and the 3 liters assumed in the ICRP 134 biokinetic model. On the other hand, the posterior distribution of blood plasma volume of subject Zr0101102 has a median of $\sim$ 2.2 liters and a true value of 1.75 liters, which are significantly different than the 3 liters assumed in the ICRP 134 zirconium biokinetic model model.

Overall, the canned Bayesian analysis and add-one-in Bayesian analysis give comparable results that are equivalent for all practical intents and purposes. The canned Bayesian analysis and add-one-in Bayesian analysis produce similar distributions of intakes as can be seen in the third row of Figure 4.5, with the true value of the intake falling within the 95% credible intervals for both methods. The plots in Figure 4.6 show that the true blood plasma volume was recovered in the canned Bayesian analysis, i.e., the true value of the blood plasma volume is in the 95% credible interval. The plots for the add-one-in analysis in Figure 4.6 show that the blood plasma volume for Zr0101102 was not recovered, whereas the blood plasma volume for both subjects was recovered in the canned Bayesian analysis.

4.7.5 MAP Estimates

One drawback of the Bayesian analyses discussed here is that they take much longer to complete than the standard method, with the canned prior Bayesian evaluation of each subject taking $\sim$ 7 hours while the standard method is essentially instantaneous. A maximum a posteriori (MAP) analysis consists of finding the maxima of the unnormalized posterior distribution formed by the product of the canned priors and the likelihood. As can be seen in Figure 4.4, MAP evaluations give some of the benefits of a full Bayesian analysis with the speed of the standard method while requiring no additional coding in Stan. In particular, MAP estimates incorporate customization of the blood plasma volume to the individual much like the full Bayesian analysis, which results in intake estimates that more closely align with the true values.

4.8 Summary

In Chapter 3, a proof-of-principle example was given that showed how one can define informative priors for the distribution of rate constants in the ICRP 134 zirconium biokinetic model and perform a population-level Bayesian calibration of the model using bioassay data from a controlled study. In principle, a similar calibration can be performed with any biokinetic model using different types of bioassay data including that from controlled experiments, occupational exposures, and autopsy (Avtandilashvili, Brey, and Birchall, 2013). The end products of this calibration are the MCMC samples of the posterior distributions of key population-level parameters. In this chapter, fits of familiar probability distributions to these posterior samples compactly encapsulate all available information about the parameters into canned priors, which can then be used as priors in the Bayesian analysis of bioassay data from a new subject.

The canned prior Bayesian analysis was illustrated using two out-of-study subjects, contrasting its results with those obtained with 3 other methods: standard method, add-in-in Bayesian analysis, and MAP analysis of the canned prior Bayesian model. A summary of the intakes calculated by all 4 methods for both subjects is given in Table 4.2. The intake evaluations for these 2 subjects highlight the main advantages of the Bayesian method using canned priors over the standard method, specifically: meaningful uncertainties can be estimated for quantities of interest, and parameters of the biokinetic model can be customized to an individual in a disciplined fashion that avoids physiologically implausible or impossible parameters. This adjustment of parameters in the biokinetic model can lead to better agreement between the observed and predicted bioassay results, giving increased accuracy in intake estimates, and more realistic estimates of uncertainties. The add-one-in and canned Bayesian approaches yield similar results. However, the canned Bayesian approach is much faster and offers the advantage using canned priors that can be disseminated

along with the biokinetic models when they are published. The MAP estimates of intake are essentially as fast as the standard method and come at no cost once the canned prior Bayesian model is specified, but are not true Bayesian estimates and therefore do not provide meaningful estimates of uncertainty. The canned Bayesian analysis would not be practical for evaluating a large number of cases as described by Puncher and Riddell (2016) and Poudel et al. (2018), but would be ideal for evaluating bioassay from a small number of individuals who had significant exposures to radioactive materials like those described by Carbaugh and LaBone (2003).

Chapter 5

The Details

5.1 Introduction

All of the data preparation and calculations in this dissertation were performed with R scripts and cmdstan code run from bash scripts. While completely reproducible, this body of work is complex and difficult to communicate to others concisely. Therefore, to provide additional details on the calculations presented in Chapters 3 and 4 here multiple sets of simulated bioassay data for 16 hypothetical subjects are generated. The simulated data from some subset of the 16 subjects (e.g., 5 subjects to speed up the calculations) will then be used as the input data for the population model calculations discussed in Chapter 3. The canned priors calculated from the simulated population data will then be used to evaluate the simulated bioassay data from a couple of individuals not used in development of the canned priors, which is akin to the calculations performed in Chapter 4. Finally, a description of the codes and procedures are put on GitHub along with key Stan and R codes. The goal is to allow interested individuals to download and run the codes, which will facilitate them reproducing and expanding on what was done in this dissertation.

5.2 Simulated Data

Simulated bioassay data were calculated starting with the sex, weight, $\boldsymbol{t}$, $\Delta\boldsymbol{t}$, $\boldsymbol{u}$, $\boldsymbol{n}$, $\boldsymbol{n}_u$, and $\boldsymbol{n}_p$ from the 16 subjects as defined in Chapter 3. For each subject 100 prior predictive distributions were generated using the canned priors for $\boldsymbol{\theta}$ and τ from

Table 5.1 Simulated bioassay data for Subject Zr_01 from run 50.

run	subject	sex	weight	vol	t	Δt	M	u
50	Zr_01	m	81	4.956	0.0034722	0	0.1981	0.018209
50	Zr_01	m	81	4.956	0.013889	0	0.18535	0.0092612
50	Zr_01	m	81	4.956	0.020139	0	0.19931	0.0085196
50	Zr_01	m	81	4.956	0.030556	0	0.19551	0.0083008
50	Zr_01	m	81	4.956	0.040972	0	0.17867	0.0082824
50	Zr_01	m	81	4.956	0.073611	0	0.20219	0.0085674
50	Zr_01	m	81	4.956	0.091667	0	0.17114	0.0072409
50	Zr_01	m	81	4.956	0.12639	0	0.15155	0.0067768
50	Zr_01	m	81	4.956	0.16944	0	0.1416	0.0064528
50	Zr_01	m	81	4.956	0.25972	0	0.14022	0.0059817
50	Zr_01	m	81	4.956	0.33681	0	0.1152	0.0052956
50	Zr_01	m	81	4.956	0.4625	0	0.10957	0.0048248
50	Zr_01	m	81	4.956	0.54653	0	0.10568	0.0052279
50	Zr_01	m	81	4.956	1.0035	0	0.088244	0.0039411
50	Zr_01	m	81	4.956	1.1771	0	0.087258	0.0041543
50	Zr_01	m	81	4.956	1.3375	0	0.082417	0.0036531
50	Zr_01	m	81	4.956	0.5	0.5	0.0043563	0.00079519
50	Zr_01	m	81	4.956	1	0.5	0.0049678	0.00098987
50	Zr_01	m	81	4.956	1.375	0.375	0.003174	0.00071988

Chapter 4 and the blood plasma volume posterior distributions from Chapter 3. After the simulated blood plasma data were generated the blood plasma volume for each subject was calculated using the methods described in Appendix B. For example, the bioassay data for Subject Zr_01 from run 50 of the simulation are shown in Table 5.1 and spaghetti plots of the data from all 16 subjects for run 50 are shown in Figure 5.1. Data with $\Delta t = 0$ are blood plasma samples and those with $\Delta t \neq 0$ are urine samples The units of the blood plasma bioassay are the fraction of a unit quantity of injected zirconium per liter of blood plasma and the units of the incremental urine bioassay are fraction of a unit quantity of injected zirconium present in the urine sample collected from $t - \Delta t$ to t. Weight has units of kg, vol liters, t days, and Δt days. These data are in the file *ZrSimData-612.csv*.

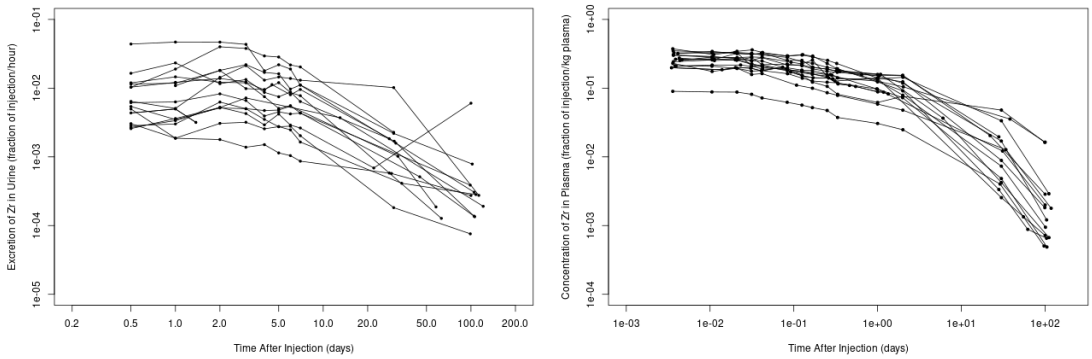

Figure 5.1 Simulated zirconium urine and blood plasma data for 16 subjects from run 50 of the simulation.

The R script *ZrDataSim.R* reads in the data from a specified run of the simulated data, the 50*th* run in this example, and creates a data file *Zr612-50.data.R* and an init file *Zr612-50.init.R* that are readable by the cmdstan code *zircon-all.stan*.

5.3 Computing Environment

The calculations in this chapter were performed on two computers running Linux Mint 20.2 having 128Gb of RAM, one of which had an i7-7820x CPU and the other an i9-10900X CPU. The statistical computing software R[1] was run in the integrated development environment RStudio[2] that was used to prepare datasets for and analyze the output from the program cmdstan[3] that did the MCMC calculations. Bash scripts were used to execute the cmdstan code as a batch file running in the background of the Linux operating system. I tried disabling hyper threading to speed up the calculations but found that the calculations ran faster with hyper threading on.

Although Stan was used in this dissertation, the MCMC software package NIMBLE (Valpine et al., 2017) was also considered. NIMBLE is very promising but it

[1]Version 4.1 available from https://cran.r-project.org/

[2]Version 1.4.1717 available from https://www.rstudio.com/products/rstudio/

[3]Version V2.27 available from https://mc-stan.org/users/interfaces/cmdstan

did not offer methods for solving the system of differential equations, i.e., the matrix exponential function was not available. NIMBLE did offer eigenvalue decomposition, which can also be used to solve the systems of differential equations, but it tended to return imaginary numbers for rate matrices (which is not physically possible) and was therefore impractical to use. NIMBLE is designed to be extensible, and my future plans include implementing the matrix exponential and eigenvalue decomposition method so that its performance can be compared to Stan's. Stan did not offer an eigenvalue decomposition method for non-symmetric matrices like those encountered with biokinetic models, and I also plan on extending Stan to handle these calculations.

5.4 Overview

The following steps give instructions on how to assemble simulated bioassay datasets for populations and individuals, analyze them with Stan using the cmdstan interface, and evaluate the posterior distributions with R code.

1. Run *ZrDataSim.R* to generate cmdstan input data files for the 16 individuals from run 50 of the simulated dataset and cmdstan input data files for the first 5 of these individuals who are designated as the study population in this example.

2. Run cmdstan program *zircon-all.stan* from the bash script *Zr-all-50.sh*.

3. Run *Zranal-Sim.R* to analyze Stan output and create canned prior dataset.

4. Run the first part of *Zrind-Sim.R* to prepare the cmdstan data file for subject 10 from run 50 of the simulated data and to perform a standard regression of the reference bioassay functions on the simulated bioassay data for this subject.

5. Run cmdstan program *zircon-ind.stan* from the bash script *Zr_10.sh*.

6. Run the second half of *Zrind-Sim.R* to analyze the Stan output.

The three R scripts are in the main level of the GitHub directory and are not discussed any further because they are only used to prepare and analyze cmdstan output data. The sub directories of GitHub are

- cmdstan-home: bash scripts (and when you run it, cmdstan)
  - Zr: cmdstan code
- data: input data
- diag: MCMC diagnostic plots
- plots: output plots
- subjects: datasets for individuals from run 50 of the simulated data.

The cmdstan code and bash scripts are discussed in more detail in the following sections.

5.5 Bash script for Population Model

The bash script *Zr-all-50.sh* shown below runs the cmdstan code with 3000 samples for the warm up followed by 10000 samples from the posterior distribution. The biokinetic model parameters, prior distributions, and bioassay data are read from *Zr612-50.data.R* and the starting values for the parameters from *Zr612-50.init.R*. Note that if a starting value is not defined for a given parameter Stan uses a randomly drawn value. The posterior samples for each parameter are written to *Zr612-50.csv*. This script assumes that it and the folder Zr reside in the default cmdstan-home folder as described in the documentation for cmdstan (Stan Development Team, 2021a).

```
#!/bin/bash
Zr/zircon-all \
   sample num_warmup=3000 num_samples=10000 \
```

```
data file=Zr/Zr612-50.data.R \
init=Zr/Zr612-50.init.R \
output file=Zr/Zr612-50.csv
```

The command used to submit the bash script as a background job is

```
nohup ./Zr-all-50.sh > Zr-50.out &
```

where the warnings, errors, and progress information are written to *Zr-50.out*.

5.6 Stan Code for Population Model

Cmdstan code resembles a mix of C++ and R, and consists of the following sections:

- Functions - user defined functions.
- Data - anything that has a fixed value and is not subject to stochastic variation.
- Parameters - anything that is assigned a probability distribution in the Model section.
- Transformed Parameters - anything that takes a parameter and performs a calculation with it before drawing a sample from the posterior distribution. Transformed parameter affect the sampling of the posterior.
- Model - the declaration of the probability distributions of the parameters.
- Generated Quantities - anything that takes a parameter and performs a calculation with it after drawing a sample from the posterior distribution. Generated quantities do not affect the sampling of the posterior.

The cmdstan code *zircon-all.stan* for the zirconium population model discussed in this section must be compiled as described in the cmdstan documentation before use. Note that all data structures used in cmdstan code must be declared before use.

5.6.1 Functions

The code in the Function section is broken down into two parts, the first giving functions used to solve the system of ordinary differential equations and the second to track and assign data and parameters to each of the 16 individuals. The key function in the first part is *q_c,* which calculates reference bioassay functions given the compartment number *comp*, times x, rate matrix kt, adjacency matrix H, number of compartments nc, and, number of non-zero rate constants nt that are read from *Zr612-50.data.R.* The adjacency matrix concisely represents how the nodes in a digraph like the one in Figure 3.4 are connected (Birchall and James, 1989).

```
functions {

//Calculate total removal rate constants in matrix K
matrix trrc( matrix k, real lambda) {
  real Ksum;
  matrix[rows(k),rows(k)] K;
  K = k;
  for(i in 1:rows(K)) {
    Ksum = 0;
    for(j in 1:rows(K)){
      Ksum = Ksum + k[i,j];
    }
    K[i,i] = -Ksum - lambda;
  }
  return K;
}

//calculate compartment content using the matrix exponential method
vector q_me(int c,vector x, vector q0, matrix k) {
  matrix[rows(k),rows(k)] a;
  vector[rows(x)] q;
  q = rep_vector(0,rows(x));
  for (i in 1:rows(x)) {
    a = matrix_exp(k*x[i]);
    for (j in 1:rows(k)) {
      q[i] = q[i] + a[j,c]*q0[j];
    }
```

```
  }
  return q;
}

//calculate content at times x and rate constants kt
vector q_c( int comp, vector x, vector kt, int[,] H, int nc, int nt ) {
  vector[nc] q0;
  matrix[nc,nc] k;
  vector[rows(x)] content;
  q0 = rep_vector(0,nc);
  q0[1] = 1.0;
  k = rep_matrix(0,nc,nc);
  for(i in 1:nt) {
    k[H[i,1],H[i,2]] = kt[i];
  }
  k = trrc(k,0);
  content = q_me(comp,x,q0,k);
  return content ;
}
```

The bioassay results for all subjects are arranged in one data vector and there is another vector that contains a number from 1 to 16 that indicates the subject to whom each bioassay result belongs. The following functions parse the data vector into blood plasma and urine data for a given subject. The data were arranged this way because each subject can have a different number of bioassay results and Stan does not support ragged matrices.

```
//gives the number of bioassay results for a given subject
int count_sub(int[] sub,int index) {
  int count;
  count = 0;
  for(i in 1:num_elements(sub)) {
    if(sub[i] == index) {
        count = count + 1;
    }
  }
  return(count);
}
```

```
//gives the number of blood plasma results for a given subject
int count_bld(int[] sub,vector dT,int index) {
  int count;
  count = 0;
  for(i in 1:num_elements(sub)) {
    if(sub[i] == index) {
      if(dT[i] == 0.0) {
        count = count + 1;
      }
    }
  }
  return(count);
}

//gives the number of urine results for a given subject
int count_urn(int[] sub,vector dT,int index) {
  int count;
  count = 0;
  for(i in 1:num_elements(sub)) {
    if(sub[i] == index) {
      if(dT[i] > 0.0) {
        count = count + 1;
      }
    }
  }
  return(count);
}

//returns the indices for all results for a given subject
int[] inds(int[] sub,int index) {
  int res[count_sub(sub,index)];
  int ci;
  ci = 1;
  for(i in 1:num_elements(sub)) {
    if(sub[i] == index) {
      res[ci] = i;
      ci = ci + 1;
    }
  }
  return(res);
}
```

```
//returns the indices for urine results for a given subject
int[] indu(int[] sub,vector dT,int index) {
  int res[count_urn(sub,dT,index)];
  int ci;
  ci = 1;
  for(i in 1:num_elements(sub)) {
    if(sub[i]==index) {
      if(dT[i] > 0.0) {
        res[ci] = i;
        ci = ci + 1;
      }
    }
  }
  return(res);
}

//returns the indices for urine results for a given subject
int[] indb(int[] sub,vector dT,int index) {
  int res[count_bld(sub,dT,index)];
  int ci;
  ci = 1;
  for(i in 1:num_elements(sub)) {
    if(sub[i]==index) {
      if(dT[i] == 0.0) {
        res[ci] = i;
        ci = ci + 1;
      }
    }
  }
  return(res);
}}
```

5.6.2 Data

The biokinetic model and bioassay data are read into the cmdstan code from *Zr612-50.data.R*.

```
data {
  int ntot;                  //total number of bioassay results
```

```
  int N;                      //number of individuals
  int n[N];                   //number of results for each individual
  int sub[ntot];              //the subject number for each bioassay result
  vector[ntot] T;             //the time of each bioassay result
  vector[ntot] dT;            //the delta t for each bioassay result
  vector[ntot] M;             //the bioassay results
  vector[ntot] u;             //the measurement uncertainties
  int nt;                     //the number of transfer rate constants
  int nc;                     //the number of compratments
  int H[nt,2];                //the adjacency matrix
  vector[N] mu_v;             //the log mean of the plasma volume
  vector[N] sigma_v;          //the log sd of the plasma volume
  vector[nt] mu_theta;        //log means of theta
  matrix[nt,nt] sigma_theta; //log covariance matrix of theta
  real mu_tau;                //log mean of tau
  real sigma_tau;             //log sd of tau
  real Eta;                   //Eta for LKJ distribution
  real df;                    //degrees of freedom for Student t
}
```

5.6.3 Parameters

Parameters have probability density models that are defined in the Model section.

```
parameters {
  vector[N] V;
  matrix[nt,N] Kt;
  corr_matrix[nt] Omega;
  real tau;
  vector[nt] theta;
}
```

5.6.4 Transformed Parameters

In the Transformed Parameter section values of the parameters drawn in the Model section are manipulated to give variables that are a function of one or more parameters. The data functions are first used in this section. For example, `T[indb(sub,dT,i)]` retrieves the times of all blood plasma measurements for the ith subject.

```
transformed parameters{
  cov_matrix[nt] Pi;
  vector[nt] Tau;
  vector[nt] kt;
  real v;
  vector[ntot] m;
  Tau = rep_vector(tau,nt);
  Pi = quad_form_diag(Omega,Tau);
  for(i in 1:N) {
    kt = exp(Kt[1:nt,i]);
    v = exp(V[i]);
    m[indb(sub,dT,i)] = (
      q_c(1, T[indb(sub,dT,i)], kt, H, nc, nt) +
      q_c(2, T[indb(sub,dT,i)], kt, H, nc, nt)) / v;
    m[indu(sub,dT,i)] = q_c(15, T[ indu(sub,dT,i)], kt, H, nc, nt) -
      q_c(15, T[indu(sub,dT,i)] - dT[ indu(sub,dT,i)], kt, H, nc, nt);
    }
}
```

5.6.5 Models

The Model section is where probability distributions are assigned.

```
model {
  Omega ~ lkj_corr(Eta);
  theta ~ multi_normal(mu_theta,sigma_theta);
  tau ~ normal(mu_tau,sigma_tau);
  for(i in 1:N) {
    Kt[1:nt,i] ~ multi_normal(theta,Pi);
    V[i] ~  normal(mu_v[i],sigma_v[i]);
    M[inds(sub,i)] ~ student_t(df,m[inds(sub,i)],u[inds(sub,i)]);
  }
}
```

5.6.6 Generated Quantities.

In the Generated Quantities section the current sample of the posterior distributions of the parameters are used to generate the predicted bioassay results for the posterior predictive plot and the log likelihood in case one wants to compare different models.

```
generated quantities{
  real M_hat[ntot];
  real log_lik[ntot];
  for(i in 1:N) {
    M_hat[inds(sub,i)] =  student_t_rng(df,m[inds(sub,i)],u[inds(sub,i)]);
    for(j in inds(sub,i)) {
      log_lik[j] = student_t_lpdf(M[j] | df,m[j],u[j]);
    }
  }
}
```

5.7 Bash script for Individual Model

The bash script *Zr_10.sh* shown below contains the *sample* command that is used to calculate the posterior distributions of the parameters in the individual model (analogous to the bash script for the population model) and the *optimize* command that is used to generate MAP estimates of the parameters. One nice thing about Stan is that the same code is used for both calculations, i.e., we don't have to write any new code to get the MAP estimates.

```
Zr/zircon-ind \
   sample num_warmup=4000 num_samples=15000 \
   data file=Zr/Zr_10.data.R \
   init=Zr/Zr_10.init.R \
   output file=Zr/Zr_10.csv

Zr/zircon-ind \
  optimize \
  data file=Zr/Zr_10.data.R \
  init=Zr/Zr_10.init.R \
  output file=Zr/Zr_10opt.csv
```

5.8 Stan Code for Individual Model

The Stan code for the individual model is somewhat simpler than the code for the population model because it deals with a single individual, which means that the 6

functions in Section 5.6.1 used to parse the bioassay data for multiple subjects are not needed and there are no loops in the Model and Transformed Parameters sections of the code.

5.8.1 Functions

The functions to solve the system of ordinary differential equations are the same as those in the first part of Section 5.6.1.

```
functions {

//Calculate total removal rate constants in matrix K
matrix trrc( matrix k, real lambda) {
  real Ksum;
  matrix[rows(k),rows(k)] K;
  K = k;
  for(i in 1:rows(K)) {
    Ksum = 0;
    for(j in 1:rows(K)){
      Ksum = Ksum + k[i,j];
    }
    K[i,i] = -Ksum - lambda;
  }
  return K;
}

//calculate compartment content using the matrix exponential method
vector q_me(int c,vector x, vector q0, matrix k) {
  matrix[rows(k),rows(k)] a;
  vector[rows(x)] q;
  q = rep_vector(0,rows(x));
  for (i in 1:rows(x)) {
    a = matrix_exp(k*x[i]);
    for (j in 1:rows(k)) {
      q[i] = q[i] + a[j,c]*q0[j];
    }
  }
  return q;
}
```

```
//calculate content at times x and rate constants kt
vector q_c( int comp, vector x, vector kt, int[,] H, int nc, int nt ) {
  vector[nc] q0;
  matrix[nc,nc] k;
  vector[rows(x)] content;
  q0 = rep_vector(0,nc);
  q0[1] = 1.0;
  k = rep_matrix(0,nc,nc);
  for(i in 1:nt) {
    k[H[i,1],H[i,2]] = kt[i];
  }
  k = trrc(k,0);
  content = q_me(comp,x,q0,k);
  return content ;
}}
```

5.8.2 Data

The data are for one subject, but otherwise are essentially the same as in the Data section of the population code.

```
data {
  int<lower=1> np;
  int<lower=1> nu;
  int<lower=1> n;
  vector[n] M;
  vector[n] u;
  vector[n] ru;
  vector[n] T;
  vector[n] dT;
  int<lower=1> nt;
  int<lower=1> nc;
  int H[nt,2];
  vector[nt] mu_theta;
  matrix[nt,nt] sigma_theta;
  real mu_tau;
  real sigma_tau;
  real mu_beta;
  real sigma_beta;
  real mu_v;
```

```
  real sigma_v;
  real Eta;
  real df;
}
```

5.8.3 Parameters

The intake β for an individual subject is a parameter to estimate unlike in the population model where it is a known constant.

```
parameters {
  real beta;
  corr_matrix[nt] Omega;
  vector[nt] theta;
  real<lower=0> tau;
  real v;
  vector[nt] Kt;
}
```

5.8.4 Transformed Parameters

```
transformed parameters{
  vector[np] m_p;
  vector[nu] m_u;
  vector[n] m;
  vector[nt] kt;
  real Beta;
  real V;
  cov_matrix[nt] Pi;
  vector[nt] Tau;
  vector[n] mu;
  vector[n] sigma;
  Tau = rep_vector(tau,nt);
  Pi = quad_form_diag(Omega,Tau);
  kt = exp(Kt);
  Beta = exp(beta);
  V = exp(v);
  m_p = ( q_c(1,T[1:np],kt,H,nc,nt) + q_c(2,T[1:np],kt,H,nc,nt) ) / V;
  m_u = q_c(15,T[(np+1):n],kt,H,nc,nt) - q_c(15,T[(np+1):n]-dT[(np+1):n],kt,H,nc,nt);
```

```
  m = append_row(m_p,m_u);
  sigma = u;
  mu = Beta * m;
}
```

5.8.5 Model

The all the subjects in the population model had known quantities of zirconium injected that were normalized to $\beta = 1$ in order to facilitate comparisons of bioassay data from subject to subject. When the data from an individual are modeled the quantity injected β is an unknown parameter to be estimated. Having the intake and blood plasma volume as parameters to estimate is a significant source of uncertainty in the final parameter estimates for a out-of-sample subject.

```
model {
  Omega ~ lkj_corr(Eta);
  theta ~ multi_normal(mu_theta,sigma_theta);
  tau ~ normal(mu_tau,sigma_tau);
  v ~  normal(mu_v,sigma_v);
  Kt ~ multi_normal(theta,Pi);
  beta ~ normal(mu_beta,sigma_beta);
  M ~ student_t(df,mu,sigma);
}
```

5.8.6 Generated Quantities

The main difference between the Generated Quantities sections for the population versus the individual model is that σ for the t distribution is estimated from the product of the predicted bioassay measurement `Beta*m[i]` times the relative uncertainty in the observed bioassay measurement `ru[i] = u[i]/M[i]`.

```
generated quantities{
  vector[n] M_hat;
  for(i in 1:n) {
    M_hat[i] = student_t_rng(df,Beta*m[i],Beta*m[i]*ru[i]);
  }
```

}

5.9 Procedure

To run the population model:

1. Compile *zircon-all.stan*, noting that it is assumed to be in the cmdstan-home/Zr directory.

2. Place *Zr612-50.data.R* and *Zr612-50.init.R.* in the Zr directory.

3. Place *Zr-all-50.sh* in the cmdstan-home directory, set its permissions, and submit it as a background job using

   a) chmod 700 Zr-all-50.sh

   b) nohup ./Zr-all-50.sh > Zr-all-50.out &

4. Note that the output goes to the text file *Zr-all-50.out* and it takes about 4 days for the example to run using the first 5 subjects as the population. In comparison, analyzing all 16 subjects takes about 25 days.

5. Run the R code *Zranal-Sim.R* to perform diagnostics and generate canned priors.

To run the individual model:

1. Compile *zircon-ind.stan*, noting that it is assumed to be in the cmdstan-home/Zr directory.

2. Run the first part of the R code Zrind-Sim.R, which reads in the canned priors from *Zrcan-50.R* and the bioassay data for the 10th subject from *Zr_10.R.*

3. Place *Zr_10.data.R* and *Zr_10.init.R.* in the Zr directory.

4. Place *Zr-ind-10.sh* in the cmdstan-home directory, set its permissions, and submit it as a background job using

    a) chmod 700 Zr-ind-10.sh

    b) nohup ./Zr-ind-10.sh > Zr-ind-10.out &

5. Note that the output goes to the text file *Zr-ind-10.out* and it takes about 3 hours for the example to run.

6. Run the second part of the R code *Zrind-Sim.R* to perform diagnostics and generate output.

5.10 GitHub

The code and instructions discussed here are available in the GitHub repository *ZrStan* in the *tomlabone* account.

Chapter 6

Summary

6.1 Essential Background

Radioactive materials spontaneously emit ionizing radiation, which has been shown to increase the risk of an exposed individual developing cancer. In order to limit the risk of detrimental health effects in individuals, regulatory limits are placed on the *radiation dose* an individual can receive from radioactive materials taken into the body. The quantity of material taken into the body is referred to as an *intake*, which can occur through various pathways such as inhalation, ingestion, and absorption through intact skin. It is not possible in practice to measure an intake directly, so the intake is inferred from measurements of the quantity of the radioactive material present in various regions of the body or in excreta at various times after the intake. Such measurements are referred to *bioassay*. Once an intake is estimated the radiation dose to the organ of interest is can be calculated. Radiation dose is the amount of ionizing radiation absorbed in tissues and organs and is used as an index of harm (e.g., the risk of developing cancer).

The process of using bioassay to calculate an intake requires a *biokinetic model* that provides a mathematical relationship between the intake and the bioassay. Thus, a biokinetic model is an idealized mathematical representation of how a material is deposited in the body, subsequently translocated to various organs and tissues, and ultimately excreted. The International Commission on Radiological Protection (ICRP) is an international non-governmental organization that publishes biokinetic models

that relate measurements $\boldsymbol{M}$ of radioactive material in the body and excreta to the amount β of the material taken into the body. These biokinetic models are used world-wide as the common basis for internal dosimetry calculations. As discussed in Chapter 2, the ICRP approximates the biokinetics of radioactive materials in the body with compartmental models (like the one for the element zirconium shown in Figure6.2) expressed mathematically as a system of ordinary differential equations (ODE), for which they provide point estimates for the rate constants. The ODE are solved to give the proportion $\boldsymbol{m}$ of an intake that is in a group of biokinetic model compartments that corresponds to what the bioassay is measuring. This proportion is called a *reference bioassay function*. If the biokinetics of the individual perfectly match the biokinetics of the ICRP biokinetic model and the bioassay has no measurement error, the intake is given by the ratio of the bioassay to its corresponding reference bioassay function.

6.2 Concise Statement of the Problem

In practice the biokinetics of an individual will diverge to some (possibly large) extent from the biokinetics of the ICRP model and the bioassay will have measurement errors, resulting in inaccurate estimates of intake and radiation dose. In addition, the ICRP models are deterministic and don't allow for the estimation of realistic uncertainties in the intakes and radiation doses, which may be needed for cases where the regulatory limits are exceeded. To address these limitations one must (i) define physiologically realistic probability distributions for the rate constants in the ICRP model, and (ii) provide a way to use the enhanced ICRP biokinetic model in the evaluation of bioassay from an individual.

6.3 Proposed Solution

In Chapter 3 (see Figure 6.1) a proof-of-principle example is given showing how the ICRP Publication 134 zirconium biokinetic model in Figure 6.2 can be calibrated with bioassay data from 16 subjects who participated in the Greiter et al. (2011) experimental study (Figure 6.3) using Bayesian hierarchical methods and informative priors. Samples from the posterior distributions of the hierarchical model shown in Figure 6.4 were obtained using Markov Chain Monte Carlo (MCMC) implemented in the computer code Stan, with the calculations requiring weeks of run time on capable computers. The posterior distributions from which these samples were drawn incorporate all the prior knowledge from the ICRP model and the variation in biokinetics observed in the study population.

In Chapter 4, probability distributions like a multivariate lognormal and a multivariate lognormal mixture were fit to the samples from the posterior distributions in the calibrated biokinetic model. These fitted posterior distributions are referred to as *canned priors* and are used for posterior passing, i.e., using the posterior distributions from one study as the priors in another. In the end, a multivariate lognormal distribution was chosen for the canned prior for $\boldsymbol{\theta}$ and a truncated lognormal distribution for τ (the blue shaded nodes in Figure 6.4). These canned priors can be used in the Bayesian analysis of an individual, which permits the disciplined modification of the rate constants in the ICRP 134 biokinetic model to better match the biokinetics of the individual, increasing the accuracy of the estimated intake and radiation doses. In addition, expressing the rate constants in the ICRP biokinetic model as distributions facilitates the estimation of realistic uncertainties in the intake and radiation doses.

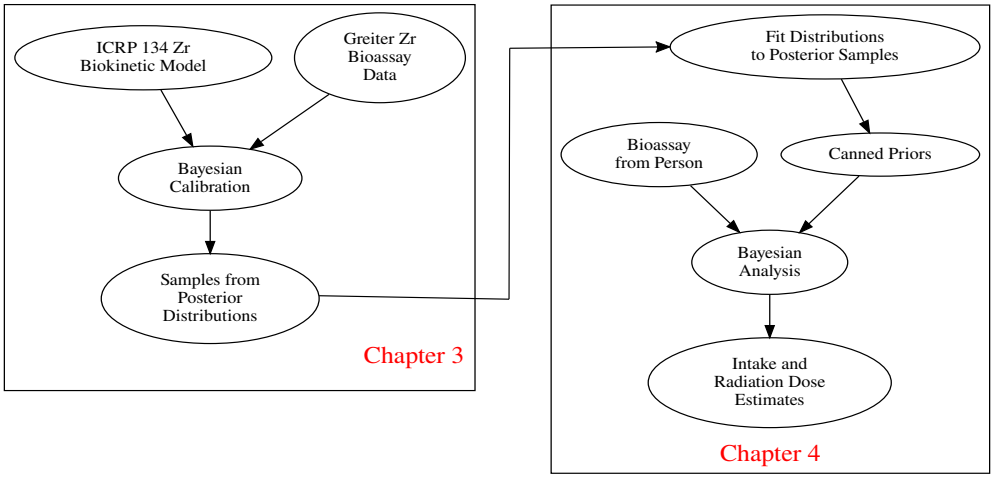

Figure 6.1 Outline of Chapter 3 and Chapter 4.

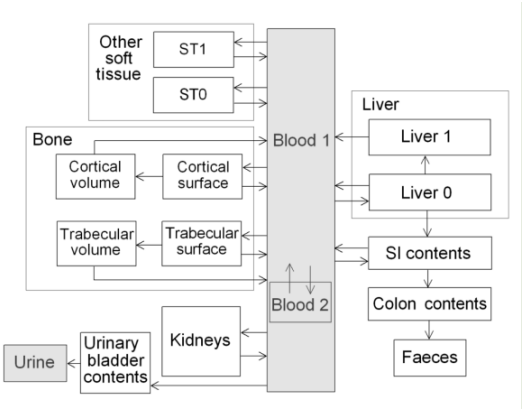

Figure 6.2 ICRP 134 biokinetic model for zirconium.

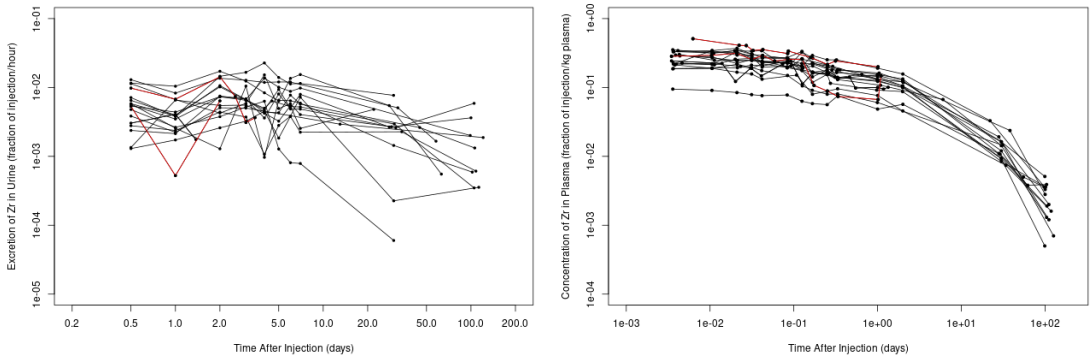

Figure 6.3 Urine and blood plasma data for 16 subjects from the Greiter zirconium study used in the population model and 2 subjects (in red) that were not used in the model. These two out-of-sample individuals are discussed in Chapter 4.

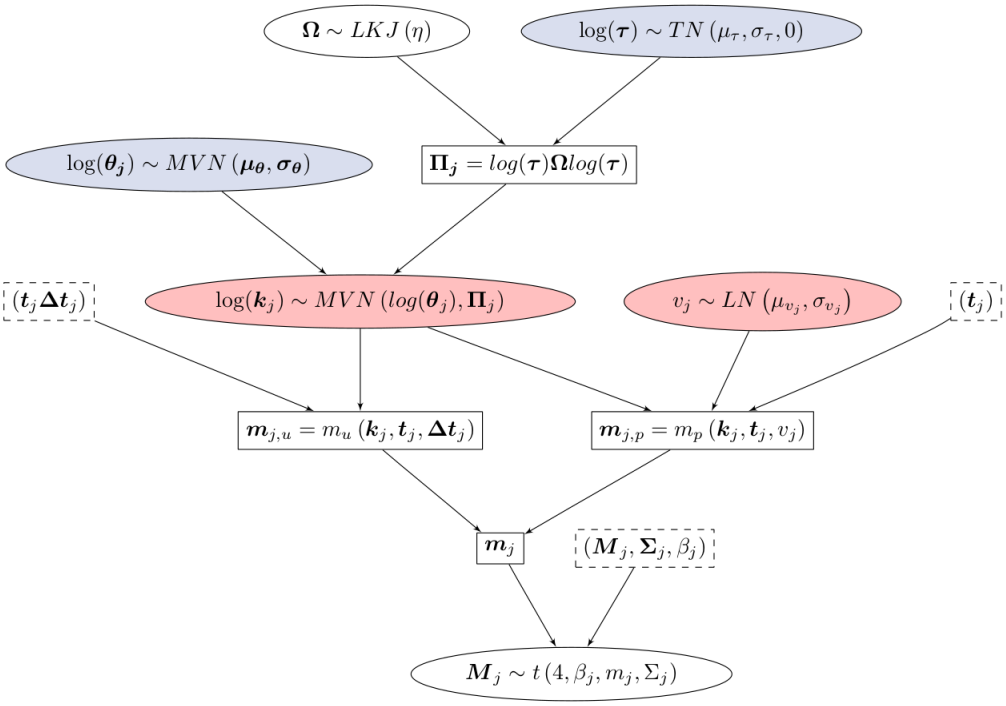

Figure 6.4 Bayesian population model for evaluating cases with the intent of estimating the posterior distributions of the hyper priors in the blue shaded nodes. The nodes for parameter distributions are denoted by ellipses, computed parameters by rectangles, and data by dashed rectangles..

6.4 Key Results

The end result of Chapter 3 is somewhat abstract — 110000 samples from the posterior distribution of the Bayesian hierarchical model. Things becomes a little more concrete in Chapter 4 when probability distributions are fit to these posterior samples to create canned priors. Then, at the end of Chapter 4 the purpose behind the process becomes clear with examples of the canned prior Bayesian analyses of bioassay performed on two out-of-study subjects (designated Zr0101102 and Zr0102004) from the Greiter study. The results of these canned prior Bayesian analyses are compared to the known values of intake and the results obtained by 3 other methods: (i) regression of $\boldsymbol{M}$ on $\boldsymbol{m}$, where $\boldsymbol{m}$ is calculated with the standard ICRP 134 zirconium model (i.e., the standard method); (ii) Bayesian analysis where the out-of-study subject is

added to the 16 study subjects and the Bayesian calibration repeated (i.e., an add-one-in analysis); (iii) the MAP analysis using the unnormalized posterior distribution of the canned prior Bayesian model.

Of primary interest are the results of the standard method versus the results of the canned prior Bayesian analysis. In Figures 6.5 and 6.6 the plots of M versus m are given for each subject. In these plots the slope of the regression line is the intake β of zirconium and the uncertainty in the slope is the uncertainty in the intake. For both subjects the standard method gives inaccurate estimates of the intake, which is known to be $\beta = 1$, whereas the canned Bayesian analysis gives more accurate estimates of the intake. The uncertainty in the intake estimate for the standard method is smaller than the uncertainty estimated with the canned Bayesian method because the standard method does not account for the uncertainty in the biokinetic model whereas the canned Bayesian analysis does. The intake distributions calculated with the two methods are directly compared in Figure 6.7 where the inaccuracy of the intake and underestimated uncertainty in the standard method are clearly seen.

The add-one-in Bayesian analysis was performed to see how how it compared to the canned prior Bayesian analysis. The add-one-in analysis took weeks to calculate and required the data from the 16 subjects used to calibrate the biokinetic model. Contrast this with the canned prior Bayesian analysis, which gave comparable results in hours without needing the calibration dataset. The advantages of the canned prior approach are obvious. The MAP analysis gave some of the benefits of the Bayesian analysis with the speed of the standard method, and could be useful as an alternative to the canned Bayesian analysis if uncertainties in estimated quantities are not needed.

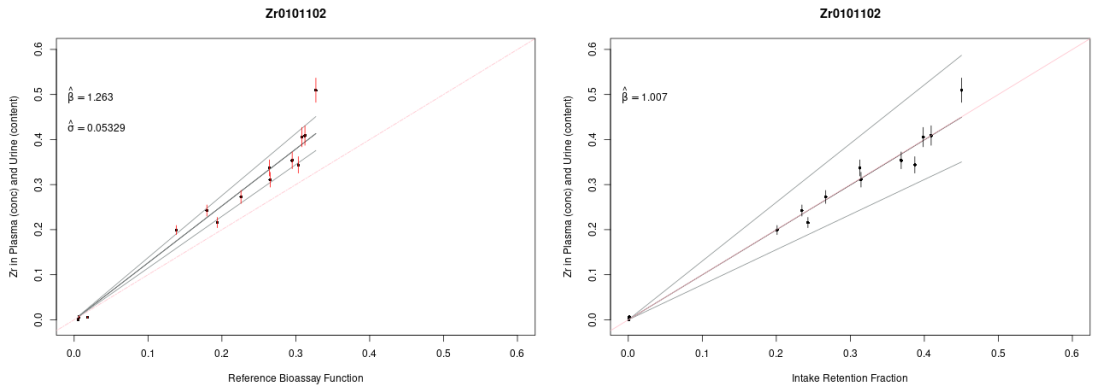

Figure 6.5 Comparison of standard method (right) to canned prior Bayesian analysis (right) for Subject Zr0101102. The grey lines denote the 95% confidence interval and 95% credible interval, respectively, on the regression lines, which correspond to the 95% intervals on the intakes. The error bars denote the 95% coverage interval on the measurements.

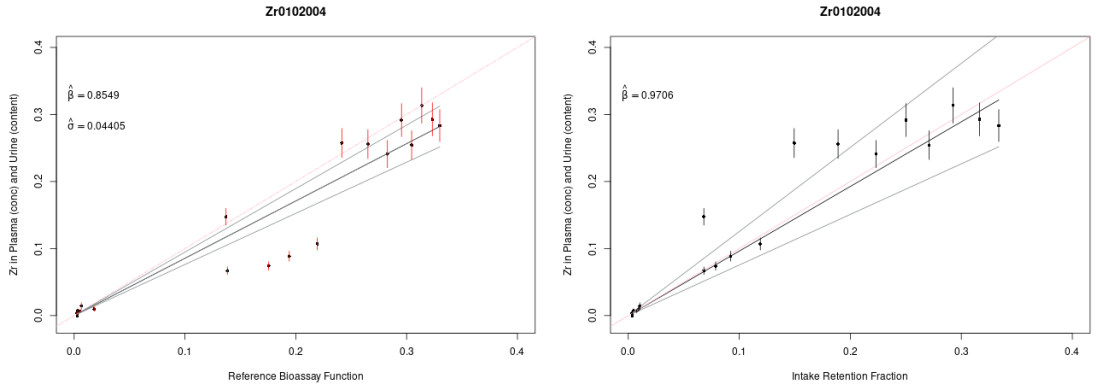

Figure 6.6 Comparison of standard method (right) to canned prior Bayesian analysis (right) for Subject Zr0102004. The grey lines denote the 95% confidence interval and 95% credible interval, respectively, on the regression lines,which correspond to the 95% intervals on the intakes. The error bars denote the 95% coverage interval on the measurements.

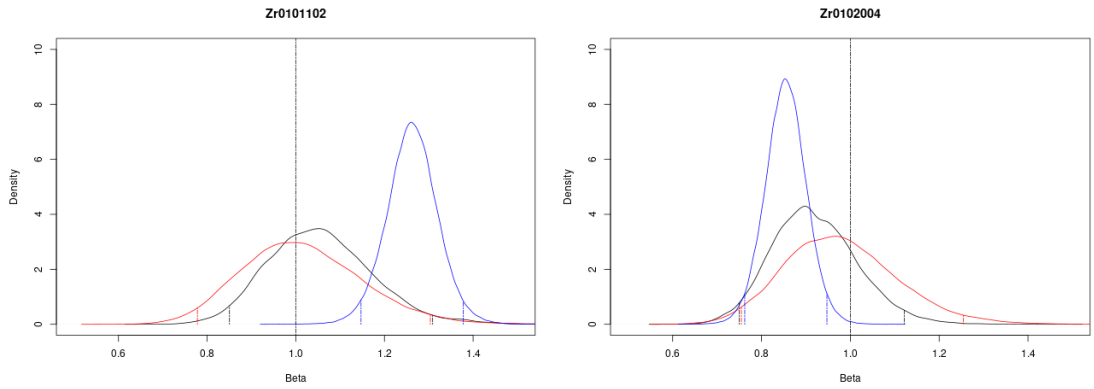

Figure 6.7 The posterior distribution of $\hat{\beta}$ calculated for Zr0101102 (left) and Zr0102004 (right) with standard method (blue), add-one-in Bayesian analysis (black), and MVN canned prior Bayesian analysis (red). The true value of the intake is $\beta = 1$ for both subjects. The vertical lines on the distribution of β in denote the 95% confidence or credible interval.

6.5 Significance of Research

This dissertation provides a proof-of-principle exercise that shows how hierarchical Bayesian models can be used to incorporate experimental data into an ICRP biokinetic model and generate realistic probability distributions for all parameters of the model. I am proposing is a fundamental change in the level of information provided by the ICRP in their published models, which is a prerequisite for increasing the accuracy in estimated quantities like intake and making meaningful estimates of uncertainty in these quantities. Further, I have given examples of how this additional information can be summarized in canned priors that can be published by the ICRP and used in subsequent Bayesian evaluations of bioassay data from an individual. Finally, I have shown that the canned prior Bayesian analysis can provide intake and radiation dose estimates that have higher accuracy and more realistic uncertainties than those obtained from the standard method. The presentation of this information in one coherent package is novel and in my opinion can be of great practical significance to the community of scientists who develop and apply biokinetic models in the field of internal dosimetry.

6.6 Future Research

My future research will focus on making the MCMC calculations faster by: exploring the use of samplers other than HMC (e.g, use NIMBLE rather than Stan); using eigenvalues/eigenvectors to solve the system of differential equations rather than the matrix exponential; running each MCMC chain on multiple computer cores (currently each chain is run on a single core); trying different parameterisations of the priors and hyperpriors (the nodes in Figure 6.4), which can have a significant impact on the speed of the calculation. Further, I want to collaborate with expert biokinetic modelers to see what we can do with a more complicated biokinetic model with more data, like the one for plutonium. Finally, I want to work on making a user interface for the canned Bayesian analysis that makes the calculation more accessible to individuals with a modest background in statistics and computer programming.

Bibliography

Allen, Bruce C., C. Eric Hack, and Harvey J. Clewell (2007). "Use of Markov Chain Monte Carlo Analysis with a Physiologically-Based Pharmacokinetic Model of Methylmercury to Estimate Exposures in U.S. Women of Childbearing Age". In: *Risk Analysis* 27.4, pp. 947–959. ISSN: 1539-6924.

Avtandilashvili, M., R. Brey, and A. Birchall (2013). "Application of Bayesian Inference to the Bioassay Data from Long-term Follow-up of Two Refractory PuO2 Inhalation Cases". In: *Health Physics* 104.4, pp. 394–404. ISSN: 0017-9078. DOI: 10.1097/HP.0b013e31827fd5cf.

Banner, Katharine M., Kathryn M. Irvine, and Thomas J. Rodhouse (2020). "The use of Bayesian priors in Ecology: The good, the bad and the not great". In: *Methods in Ecology and Evolution* 11.8, pp. 882–889. ISSN: 2041-210X.

Barnard, John, Robert McCulloch, and Xiao-Li Meng (2000). "Modeling Covariance Matrices in Terms of Standard Deviations and Correlations, with Application to Shrinkage". In: *Statistica Sinica* 10.4. Publisher: Institute of Statistical Science, Academia Sinica, pp. 1281–1311. ISSN: 1017-0405.

Benaglia, Tatiana et al. (2009). "mixtools: An R Package for Analyzing Finite Mixture Models". In: *Journal of Statistical Software* 32.6, pp. 1–29. URL: http://www.jstatsoft.org/v32/i06/.

Bernillon, P and F. Y. Bois (2000). "Statistical Issues in Toxicokinetic Modeling: A Bayesian Perspective". In: *Environmental Health Perspectives* 108, pp. 883–893. ISSN: 0091-6765. DOI: 10.2307/3454322.

BIPM, IEC et al. (2012). "The international vocabulary of metrology - basic and general concepts and associated terms (VIM), JCGM 200:2012". In: *JCGM (Joint Committee for Guides in Metrology)*.

Birchall, A. and A. C. James (1989). "A microcomputer algorithm for solving first-order compartmental models involving recycling." In: *Health physics* 56.6, pp. 857–868. ISSN: 0017-9078. DOI: 10.1097/00004032-198906000-00003.

Boecker, B.B. (1998). "Development and Use of Biokinetic Models for Incorporated Radionuclides". In: *Radiation Protection Dosimetry* 79.1, pp. 223–228. ISSN: 0144-8420. DOI: 10.1093/oxfordjournals.rpd.a032397.

Bois, F. Y. (2001). "Applications of population approaches in toxicology". In: *Toxicology Letters* 120.1, pp. 385–394. ISSN: 0378-4274. DOI: 10.1016/S0378-4274(01)00270-3.

Bois, F. Y. et al. (1996). "Population toxicokinetics of tetrachloroethylene". In: *Archives of Toxicology* 70.6, pp. 347–355. ISSN: 0340-5761, 1432-0738. DOI: 10.1007/s002040050284.

Bolker, Ben (2020). *bbmle: Tools for General Maximum Likelihood Estimation*. Version 1.0.23.1. URL: https://CRAN.R-project.org/package=bbmle.

Bonate, Peter L. (2011). *Pharmacokinetic-Pharmacodynamic Modeling and Simulation*. Google-Books-ID: kfS5oLko95kC. Springer Science & Business Media. 634 pp. ISBN: 978-1-4419-9485-1.

Box, George EP (1979). "Robustness in the strategy of scientific model building". In: *Robustness in statistics*. Elsevier, pp. 201–236.

Brand, Charlotte Olivia et al. (2019). "Cumulative Science via Bayesian Posterior Passing:" in: *Meta-Psychology* 3. ISSN: 2003-2714. DOI: 10.15626/MP.2017.840.

Burkner, Paul-Christian (2017). "brms: An R Package for Bayesian Multilevel Models Using Stan". In: *Journal of Statistical Software* 80.1. Number: 1, pp. 1–28. ISSN: 1548-7660. DOI: 10.18637/jss.v080.i01.

Carbaugh, E. H. and T. R. LaBone (2003). "Two case studies of highly insoluble plutonium inhalation with implications for bioassay". In: *Radiation Protection Dosimetry* 105.1, pp. 133–138. ISSN: 0144-8420. DOI: 10.1093/oxfordjournals.rpd.a006208. (Visited on 09/23/2019).

Carpenter, Bob et al. (2017). "Stan: A Probabilistic Programming Language". In: *Journal of Statistical Software* 76.1. Number: 1, pp. 1–32. ISSN: 1548-7660. DOI: 10.18637/jss.v076.i01.

Carson, Ewart R., L. Finkelstein, and C. Cobelli (1983). "Mathematical modeling of metabolic and endocrine systems: Model formulation, identification, and validation." In: *John Wiley & Sons, New York*.

Christinaki, Eirini et al. (2021). "Well-being Forecasting using a Parametric Transfer-Learning method based on the Fisher Divergence and Hamiltonian Monte Carlo". In: *EAI Endorsed Transactions on Bioengineering and Bioinformatics* 1.1. Num-

ber: 1 Publisher: European Alliance for Innovation, e6. ISSN: 2709-4111. DOI: 10.4108/eai.16-10-2020.166661.

Clarke, R. H. and J. Valentin (2009). "The History of ICRP and the Evolution of its Policies". In: *Annals of the ICRP* 39.1.

Cobelli, Claudio, David Foster, and Gianna Toffolo (2007). *Tracer kinetics in biomedical research: from data to model*. Springer Science & Business Media.

Davidian, Marie (2006). "Introduction to statistical population modeling and analysis for pharmacokinetic data". In: *Invited white paper for the International Workshop on Uncertainty and Variability in Physiologically Based Pharmacokinetic (PBPK) Models*. Vol. 89. Citeseer.

DiStefano, J. J. and E. M. Landaw (1984). "Multiexponential, multicompartmental, and noncompartmental modeling. I. Methodological limitations and physiological interpretations". In: *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology* 246.5. Publisher: American Physiological Society, R651–R664. ISSN: 0363-6119. DOI: 10.1152/ajpregu.1984.246.5.R651.

Everitt, BS (1996). "An introduction to finite mixture distributions". In: *Statistical Methods in Medical Research* 5.2. Publisher: SAGE Publications Ltd STM, pp. 107–127. ISSN: 0962-2802. DOI: 10.1177/096228029600500202.

Fairlie, Ian (2005). "Uncertainties in Doses and Risks from Internal Radiation". In: *Medicine, Conflict and Survival* 21.2, pp. 111–126. ISSN: 1362-3699. DOI: 10.1080/13623690500073414.

Gelman, Andrew, John B. Carlin, et al. (2013). *Bayesian Data Analysis, Third Edition*. Google-Books-ID: ZXL6AQAAQBAJ. CRC Press. 677 pp. ISBN: 978-1-4398-4095-5.

Gelman, Andrew, Daniel Simpson, and Michael Betancourt (2017). "The Prior Can Often Only Be Understood in the Context of the Likelihood". In: *Entropy* 19.10, p. 555. ISSN: 1099-4300. DOI: 10.3390/e19100555.

Gill, Jeff and Gary King (2004). "What to Do When Your Hessian is Not Invertible: Alternatives to Model Respecification in Nonlinear Estimation". In: *Sociological Methods & Research* 33.1. Publisher: SAGE Publications Inc, pp. 54–87. ISSN: 0049-1241. DOI: 10.1177/0049124103262681.

Green, Peter J. (2018). "Introduction to finite mixtures". In: *arXiv:1705.01505 [stat]*. arXiv: 1705.01505.

Greiter, M. B. (2008). "Study of the biokinetics of zirconium isotopes in humans and its relevance to internal dosimetry". PhD thesis. Technische Universitat Munchen.

Greiter, M. B. et al. (2011). "Human biokinetic data and a new compartmental model of zirconium - A tracer study with enriched stable isotopes". In: *Science of The Total Environment* 409.19, pp. 3701–3710. ISSN: 0048-9697. DOI: 10.1016/j.scitotenv.2011.06.031.

Hack, C. Eric (2006). "Bayesian analysis of physiologically based toxicokinetic and toxicodynamic models". In: *Toxicology*. A Research Strategy to Improve Risk Estimates for Bromate in Drinking Water 221.2, pp. 241–248. ISSN: 0300-483X. DOI: 10.1016/j.tox.2005.12.017.

Harrison, John (2009). "Biokinetic and dosimetric modelling in the estimation of radiation risks from internal emitters". In: *Journal of Radiological Protection* 29.2, A81–A105. ISSN: 0952-4746. DOI: 10.1088/0952-4746/29/2A/S06.

Hoffman, Matthew D. and Andrew Gelman (2014). "The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo." In: *J. Mach. Learn. Res.* 15.1, pp. 1593–1623.

Hubal, Elaine A. Cohen et al. (2019). "Advancing internal exposure and physiologically-based toxicokinetic modeling for 21st-century risk assessments". In: *Journal of Exposure Science & Environmental Epidemiology* 29.1, p. 11. ISSN: 1559-064X. DOI: 10.1038/s41370-018-0046-9.

IAEA (2004). *IAEA Report 37, Methods for Assessing Occupational Radiation Doses Due to Intakes of Radionuclides*. International Atomic Energy Agency.

ICRP (1975). *Report of the Task Group on Reference Man. ICRP Publication 23*. Pergamon Press.

— (1979). *ICRP Publication 30 (Part 1): Limits for Intakes of Radionuclides by Workers*. Annals of the ICRP 2 (3-4).

— (1990). *ICRP Publication 56: Age-dependent Doses to Members of the Public from Intake of Radionuclides*. Publication 56. Ann. ICRP 20 (2).

— (1993). *ICRP Publication 67: Age-dependent Doses to Members of the Public from Intake of Radionuclides - Part 2 Ingestion Dose Coefficients*. Annals of the ICRP 23 (3-4).

— (1995). *ICRP Publication 68: Dose Coefficients for Intakes of Radionuclides by Workers*. Annals of the ICRP 24 (4).

ICRP (2015). *ICRP Publication 130: Occupational Intakes of Radionuclides: Part 1*. Annals of the ICRP 44(2).

— (2017). *ICRP Publication 137 Occupational Intakes of Radionuclides: Part 3*. Ann. ICRP 46(3/4).

ISO (2008). *ISO/IEC Guide 98-3:2008 - Uncertainty of measurement – Part 3: Guide to the expression of uncertainty in measurement (GUM:1995)*.

Jacquez, John Alfred (1985). *Compartmental Analysis in Biology and Medicine*. Google-Books-ID: Uc1qAAAAMAAJ. University of Michigan Press. 584 pp. ISBN: 978-0-472-10063-7.

James, A. C., A. Birchall, and M. Puncher (2008). "Uncertainty in internal doses: using Bayes to transfer information". In: *Proceedings of the 12th International Congress of the International Radiation Protection Association,(Buenos Aires, Argentina)*. IRPA 12. Buenos Aires, Argentina.

JCGM (2008). *Evaluation of measurement data - Guide to the expression of uncertainty in measurement*. 100:2008.

Kacker, Raghu N. (2015). "Probability distributions and coverage probability in GUM, JCGM documents, and statistical inference". In: *Measurement* 65, pp. 61–70. ISSN: 0263-2241. DOI: 10.1016/j.measurement.2014.12.056.

Kennedy, M. C. and A. O'Hagan (2002). "Bayesian calibration of computer models". In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 63.3, pp. 425–464. ISSN: 1369-7412. DOI: 10.1111/1467-9868.00294.

Kruschke, John (2015). *Doing Bayesian data analysis: A tutorial with R, JAGS, and Stan*. 2nd ed. Publisher: Academic Press. Elsevier. ISBN: 978-0-12-405888-0.

Kutner, Michael H. (2005). *Applied Linear Statistical Models*. Google-Books-ID: 0xqCAAAACAAJ. McGraw-Hill Irwin. 1396 pp. ISBN: 978-0-07-112221-4.

LaBone, T. R. (2010). *Use of Weighted Regression in Internal Dosimetry*. Master's Essay. University of South Carolina.

Lambert, Ben (2018). *A Student's Guide to Bayesian Statistics*. Sage.

Lange, Kenneth L., Roderick J. A. Little, and Jeremy M. G. Taylor (1989). "Robust Statistical Modeling Using the t Distribution". In: *Journal of the American Statistical Association* 84.408. Publisher: [American Statistical Association, Taylor & Francis, Ltd.], pp. 881–896. ISSN: 0162-1459. DOI: 10.2307/2290063.

Leggett, R. W. and L. R. Williams (1995). "A proposed blood circulation model for Reference Man." In: *Health physics* 69.2, pp. 187–201.

Lewandowski, Daniel, Dorota Kurowicka, and Harry Joe (2009). "Generating random correlation matrices based on vines and extended onion method". In: *Journal of Multivariate Analysis* 100.9, pp. 1989–2001. ISSN: 0047-259X. DOI: 10.1016/j.jmva.2009.04.008.

Li, W. B. (2018). "Internal Dosimetry". In: *Japanese Journal of Health Physics* 53.2, pp. 72–99. DOI: 10.5453/jhps.53.72.

Li, W. B., M. B. Greiter, et al. (2011a). "Reliability of a new Biokinetic Model of Zirconium in Internal Dosimetry: Part I, Parameter Uncertainty Analysis". In: *Health Physics* 101.6, p. 660. ISSN: 0017-9078. DOI: 10.1097/HP.0b013e3181fbfba9.

— (2011b). "Reliability of a new Biokinetic Model of Zirconium in Internal Dosimetry: Part II, Parameter Sensitivity Analysis". In: *Health Physics* 101.6, pp. 677–692. ISSN: 0017-9078. DOI: 10.1097/HP.0b013e318226edc0.

Li, W. B., W. Klein, et al. (2015). "Parameter uncertainty analysis of a biokinetic model of caesium". In: *Radiation Protection Dosimetry* 163.1, pp. 37–57. ISSN: 0144-8420. DOI: 10.1093/rpd/ncu055. (Visited on 02/10/2018).

McElreath, Richard (2018). *Statistical Rethinking: A Bayesian Course with Examples in R and Stan*. Google-Books-ID: 1yhFDwAAQBAJ. CRC Press. 489 pp. ISBN: 978-1-4822-5348-1.

Miller, Guthrie (2008). "Variability and uncertainty of biokinetic model parameters: the discrete empirical Bayes approximation". In: *Radiation Protection Dosimetry* 131.3. ISSN: 0144-8420. DOI: 10.1093/rpd/ncn180.

— (2017). *Probabilistic Interpretation of Data*. Glasstree Academic Publishing.

Moler, C. and C. Van Loan (2003). "Nineteen Dubious Ways to Compute the Exponential of a Matrix, Twenty-Five Years Later". In: *SIAM Review* 45.1, pp. 3–49. ISSN: 0036-1445. DOI: 10.1137/S00361445024180.

Monnahan, C. C., J. T. Thorson, and T. A. Branch (2017). "Faster estimation of Bayesian models in ecology using Hamiltonian Monte Carlo". In: *Methods in Ecology and Evolution* 8.3, pp. 339–348. ISSN: 2041-210X. DOI: 10.1111/2041-210X.12681.

National Research Council (1988). *Health risks of radon and other internally deposited alpha-emitters: BEIR IV*. Publisher: National Academies Press.

National Research Council (2009). *Science and Decisions: Advancing Risk Assessment*. Washington, DC: The National Academies Press. 422 pp. ISBN: 978-0-309-12046-3. DOI: 10.17226/12209.

NCRP (2010). *NCRP Report No. 164, Uncertainties in Internal Radiation Dose Assessment*.

Norton, J. P. (1982). "An investigation of the sources of nonuniqueness in deterministic identifiability". In: *Mathematical Biosciences* 60.1, pp. 89–108. ISSN: 0025-5564. DOI: 10.1016/0025-5564(82)90033-5. (Visited on 03/13/2021).

Paquet, F., M.R. Bailey, et al. (2016). "Assessment and interpretation of internal doses: uncertainty and variability". In: *Annals of the ICRP* 45.1. Publisher: SAGE Publications Ltd, pp. 202–214. ISSN: 0146-6453. DOI: 10.1177/0146645316633595.

Paquet, F., R. W. Leggett, et al. (2016). "ICRP Publication 134: Occupational Intakes of Radionuclides: Part 2." In: *Annals of the ICRP* 45.3, pp. 7–349.

Pawitan, Yudi (2001). *In all likelihood: statistical modelling and inference using likelihood*. Oxford University Press.

Polig, E. (2001). "Modeling the Distribution and Dosimetry of Internal Emitters: A Review of Mathematical Procedures Using Matrix Metbods". In: *Health Physics* 81.5, p. 492. ISSN: 0017-9078.

Poudel, Deepesh et al. (2018). "Bayesian Analysis of Plutonium Bioassay Data at Los Alamos National Laboratory". In: *Health Physics* 115.6, p. 712. ISSN: 0017-9078. DOI: 10.1097/HP.0000000000000933.

Puncher, M., A. Birchall, and R. K. Bull (2011). "Uncertainties on Lung Doses from Inhaled Plutonium". In: *Radiation Research* 176.4, pp. 494–507. ISSN: 0033-7587. DOI: 10.1667/RR2410.1.

— (2012). "A method for calculating bayesian uncertainties on internal doses resulting from complex occupational exposures". In: *Radiation Protection Dosimetry* 151.2, pp. 224–236. ISSN: 0144-8420. DOI: 10.1093/rpd/ncr475.

— (Dec. 1, 2014). "An intake prior for the Bayesian analysis of plutonium and uranium exposures in an epidemiology study". In: *Radiation Protection Dosimetry* 162.3, pp. 306–315. ISSN: 0144-8420. DOI: 10.1093/rpd/nct268. (Visited on 2021).

Puncher, M. and A. E. Riddell (2016). "A Bayesian analysis of plutonium exposures in Sellafield workers". In: *Journal of Radiological Protection* 36.1, p. 1. ISSN: 0952-4746. DOI: 10.1088/0952-4746/36/1/1.

R Core Team (2021). *R: Language and Environment for Statistical Computing*. Version 4.1. URL: https://www.R-project.org/.

Retzlaff, John A. et al. (1969). "Erythrocyte Volume, Plasma Volume, and Lean Body Mass in Adult Men and Women". In: *Blood* 33.5, pp. 649–667. ISSN: 0006-4971. DOI: 10.1182/blood.V33.5.649.649.

Riggs, Douglas S. (1952). "Quantitative Aspects of Iodine Metabolism in Man". In: *Pharmacological Reviews* 4.3, pp. 284–370. ISSN: 0031-6997, 1521-0081.

Robertson, J. S. (1983). *Compartmental distribution of radiotracers*. CRC Press.

Roosa, Kimberlyn and Gerardo Chowell (2019). "Assessing parameter identifiability in compartmental dynamic models using a computational approach: application to infectious disease transmission models". In: *Theoretical Biology and Medical Modelling* 16.1, p. 1. ISSN: 1742-4682. DOI: 10.1186/s12976-018-0097-6.

Schmidl, D. et al. (2012). "Bayesian model selection validates a biokinetic model for zirconium processing in humans". In: *BMC Systems Biology* 6.1, p. 95. ISSN: 1752-0509. DOI: 10.1186/1752-0509-6-95.

Scrucca, Luca et al. (2016). "mclust 5: Clustering, Classification and Density Estimation Using Gaussian Finite Mixture Models". In: *The R Journal* 8.1, pp. 289–317. ISSN: 2073-4859.

Skrable, K. W., G. E. Chabot, et al. (1988). "Intake retention functions and their applications to bioassay and the estimation of internal radiation doses". In: *Health Physics* 55.6, pp. 933–950. ISSN: 0017-9078.

— (1994). "Estimation of intakes from repetitive bioassay measurements". In: *Internal Radiation Dosimetry (Raabe, OG, Ed.), Medical Physics Publishing, Madison, WI*, pp. 431–460.

Skrable, K. W., C. S. French, et al. (2002). *Variance Models for Estimating Intakes from Repetitive Bioassay Measurements IN: Practical Applications of Internal Dosimetry WE Bolch eds*. Medical Physics Publishing, Madison Wi.

Spiegel, Murray R. (1981). *Applied differential equations*. Englewood Cliffs: Prentice-Hall.

Stan Development Team (2021a). *Cmdstan User Guide*. Version 2.27. URL: https://mc-stan.org.

— (2021b). *Stan User Guide*. Version 2.27. URL: https://mc-stan.org.

Thorne, M. C. (1992). *ICRP publication 60: 1990 recommendations of the international commission on radiological protection: Annals of the ICRP, 21 (1-3), 1991.*

Turner, James E., Darryl J. Downing, and James S. Bogard (2012). *Statistical Methods in Radiation Physics*. Google-Books-ID: rhLIjvjgjkMC. John Wiley & Sons. 1 p. ISBN: 978-3-527-64656-2.

USEPA (2006). *Approaches for the Application of Physiologically Based Pharmacokinetic (PBPK) Models and Supporting Data in Risk Assessment*. Reports & Assessments.

Valentin, J. (2002). "Basic anatomical and physiological data for use in radiological protection: reference values: ICRP Publication 89". In: *Annals of the ICRP* 32.3, pp. 1–277. ISSN: 0146-6453. DOI: 10.1016/S0146-6453(03)00002-2.

Valpine, Perry de et al. (2017). "Programming With Models: Writing Statistical Algorithms for General Model Structures With NIMBLE". In: *Journal of Computational and Graphical Statistics* 26.2, pp. 403–413. ISSN: 1061-8600. DOI: 10.1080/10618600.2016.1172487.

Vehtari, Aki et al. (2021). "Rank-Normalization, Folding, and Localization: An Improved R-hat for Assessing Convergence of MCMC". In: *Bayesian Analysis* -1.-1. Publisher: International Society for Bayesian Analysis, pp. 1–28. ISSN: 1936-0975, 1931-6690. DOI: 10.1214/20-BA1221.

Venables, W. N. and B. D. Ripley (2002). *Modern Applied Statistics with S*. Fourth. ISBN 0-387-95457-0. New York: Springer.

Wackerly, Dennis, William Mendenhall, and Richard L. Scheaffer (2014). *Mathematical statistics with applications*. Cengage Learning.

Wakefield, Jon (1996). "The Bayesian Analysis of Population Pharmacokinetic Models". In: *Journal of the American Statistical Association* 91.433, pp. 62–75. ISSN: 0162-1459. DOI: 10.2307/2291383.

Wakefield, Jon, Leon Aarons, and Amy Racine-Poon (1999). "The Bayesian approach to Population pharmacokinetic/pharmacodynamic modeling". In: *Case Studies in Bayesian Statistics*. Ed. by Constantine Gatsonis et al. Lecture Notes in Statistics. Springer New York, pp. 205–265. ISBN: 978-1-4612-1502-8.

Wesner, Jeff S. and Justin P. F. Pomeranz (2020). "Choosing priors in Bayesian ecological models by simulating from the prior predictive distribution". In: *bioRxiv*. Publisher: Cold Spring Harbor Laboratory Section: Confirmatory Results, p. 2020.12.10.419713. DOI: 10.1101/2020.12.10.419713.

Wuertz, Diethelm, Tobias Setz, and Yohan Chalabi (2020). *fMultivar: Rmetrics - Analysing and Modeling Multivariate Financial Return Distributions*. R package version 3042.80.1. URL: https://CRAN.R-project.org/package=fMultivar.

Appendix A

The Inverse Problem and Identifiability of the Biokinetic Model

In theory, the maxima of the likelihood for the Riggs iodine models give the maximum likelihood estimates of the intake and rate constants, but as discussed in Chapter 2 there are significant problems encountered if we try to do this. In the Inverse Problem a given biokinetic model is said to be *identifiable* if all the unknown rate constants in the model can be uniquely determined from the available bioassay measurements (Bonate, 2011, page 29; Jacquez, 1985, Section 14.3; Claudio Cobelli, Foster, and Toffolo, 2007, Ch 5; Carson, Finkelstein, and C. Cobelli, 1983, Ch 7&8). A biokinetic model that is not identifiable because one or more parameters in the biokinetic model can't be uniquely determined is referred to as an *unidentifiable* model. Biokinetic models that are unidentifiable given an infinite amount of error-free bioassay measurements are referred to as being *structurally unidentifiable*. For example, the Riggs model is unidentifiable given an infinite amount of error-free thyroid bioassay measurements (Norton, 1982, Model 6). To make the Riggs model structurally identifiable either bioassay measurements of the blood compartment or the Other compartment must be provided in addition to the thyroid bioassay. Looking in more detail, given only the thyroid bioassay measurements the rate constant is $k_{34} = k_{oth,fec}$ is unidentifiable, which causes the whole biokinetic model to be unidentifiable. Note that it is quite difficult to determine analytically whether a biokinetic model with

more than 2 or 3 compartments is structurally identifiable (Claudio Cobelli, Foster, and Toffolo, 2007, pg 110).

Those who are unaware of the structural unidentifiability in the Riggs model may proceed to feed the likelihood function for model and the thyroid bioassay measurements into an optimizer program (e.g., mle2 from the bbmle package in R) to obtain estimates of all the rate constants in the model. We are lucky if the the optimizer fails (i.e., does not converge to a solution). More problematic is the case when estimates of the rate constants including $k_{oth,fec}$ are generated. In this case the error in the estimate of $k_{oth,fec}$ will be noticeably much larger than the error in the estimates of the other rate constants. In practice this can go unnoticed because

- the Hessian matrix used to calculate the errors in not invertable and errors are not reported, or
- the errors may be calculated using a generalized inverse of an uninvertable Hessian (Gill and King, 2004), which can produce error estimates that are difficult to interpret[1].

Structural identifiability is a necessary but not sufficient property for a biokinetic model to have because even if a biokinetic model is structurally identifiable – it may not be *practically identifiable* because of insufficient quantity or inadequate quality of data. One way to check for unidentifiability in a biokinetic model is a parametric bootstrap analysis similar to the one discussed by (Roosa and Chowell, 2019), which goes something like this:

1. Generate a large number of simulated error-free data using the biokinetic model.
2. Add $N(0, \varepsilon M)$ distributed error to each datum, where ε is the relative uncertainty in each datum M.

[1]Note that mle2 does this.

3. Determine the rate constants a number of times to estimate the variation in their estimates.

4. Repeat steps 2-3, increasing the value of ε.

This parametric bootstrap analysis was performed on the Riggs iodine model with $\varepsilon = \{0.001, 0.002, 0.005, 0.007, 0.01\}$ and 30 iterations at each relative uncertainty. The results in Figure A.1 show that the error in the estimates of $k_{oth,fec}$ is over 100 times the measurement error in the content of the thyroid compartment – essentially imperceptible measurement errors result in large errors in the parameter estimate. The variation in $k_{oth,bld}$ is somewhat smaller, but the variation in the estimates of the other three rate constants is much smaller and fairly constant. This analysis shows that the Riggs iodine model is for all practical intents and purposes unidentifiable. This conclusion will be extended to all *realistic* ICRP biokinetic models (like zirconium) without proof, i.e., all ICRP biokinetic models will be assumed to be unidentifiable.

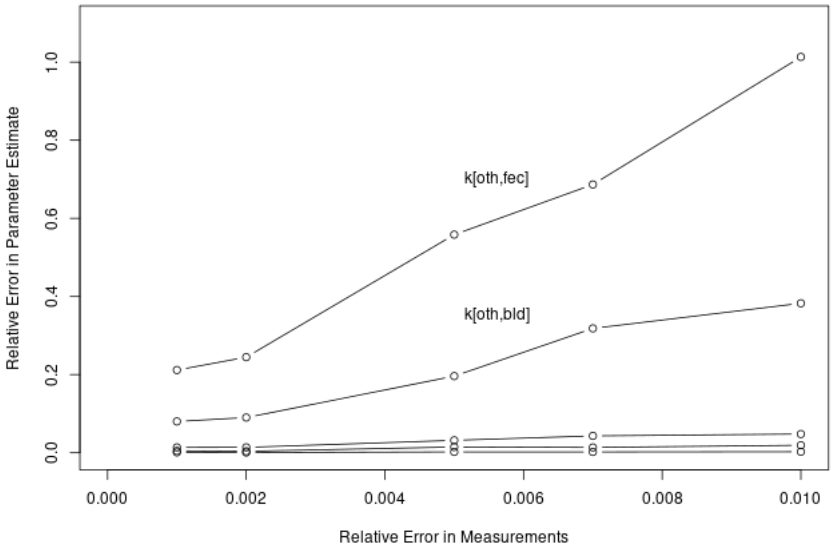

Figure A.1 Bootstrap estimate of parameter uncertainty for a maximum likelihood (ML) estimate of the rate constants in the Riggs iodine model as a function of measurement error in the thyroid bioassay.

Appendix B

Volume of Blood Plasma

B.1 Blood Volume Calculated by Isotopic Dilution

In the study by Greiter et al. (2011) measurements of zirconium in the blood are made by drawing 10 mL of blood at specified times after the iv injection of zirconium. The measured content $M_{bld}(t)$ of the blood plasma compartment is calculated by multiplying the measured concentration $C_{bld}(t)$ by the volume V of the blood plasma compartment. The volume of the blood plasma compartment was calculated by fitting a two exponential function to the zirconium concentration as a function of time (Figure B.1)

$$C_{bld}(t) = A_1 e^{-B_1 t} + A_2 e^{-B_2 t} \tag{B.1}$$

For example, using the non-linear regression function nls in R, the fitted function for subject Zr0602615 is

$$C_{bld}(t) = 11.515580 e^{-0.491357 t} + 17.126996 e^{-0.009952 t} \tag{B.2}$$

which at $t = 0$ is

$$C_{bld}(0) = 11.515580 + 17.126996 = 28.64258 \tag{B.3}$$

both having units of percent of the injected intake per kg of plasma. Thus, 0.2864258 of the injected zirconium is present in a kg of plasma at time $t = 0$. Assuming the plasma has a density of 1.027 kg/L, the volume of the is plasma compartment in units of L is

$$V = \frac{1}{(1.027)(0.2864258)} = 3.399519 \tag{B.4}$$

This is same volume calculated in Greiter (2008, Section 6.3.2).

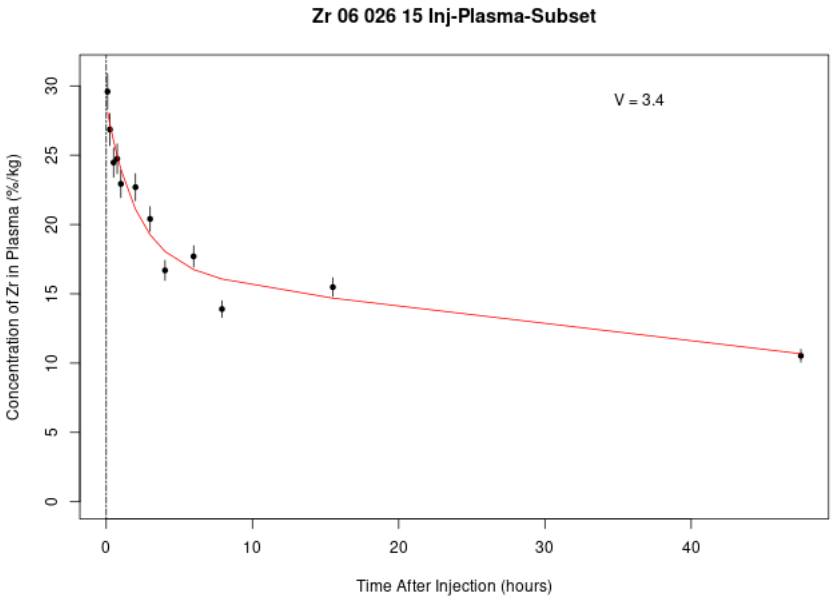

Figure B.1 Blood plasma volume in liters of subject Zr0602615 determined by calculating the concentration of zirconium in the blood at $t = 0$ days after the injection intake.

B.2 Blood Plasma Volume

The dilution method used by Greiter to calculate the volume of the blood plasma compartment is applicable only if the intake is known, which will not be the case in practice. When the intake is not known the volume of the blood plasma compartment for an individual can be estimated using one of the following equations (Retzlaff et al., 1969) based sex, height H in meters, and weight W in kg:

Male blood plasma volume in mL:

- $V_m = 14.5W + 2035$ (when only weight is known).
- $V_m = 30.2H + 2119$ (when only height is known).
- $V_m = 23.7H + 9W - 1709$ (when height and weight are known).

Female blood plasma volume in mL:

- $V_f = 24.0W + 872$ (when only weight is known).
- $V_f = 48.1H + 5550$ (when only height is known).
- $V_f = 40.5H + 8.4W - 4811$ (when height and weight are known).

The geometric mean of a lognormal blood plasma volume prior distribution for an individual is set equal to the appropriate value calculated above. The geometric standard deviation of the blood volume prior is set equal to 1.2, resulting in 95% of the prior values being within a factor of ~ 1.3 of the geometric mean. Although ICRP Publication 89 (Valentin, 2002, pg 138) gives a blood plasma volume of 3.0 L for males and 2.4 L for females, the volume of 3.0 L should be used when evaluating bioassay data with the unmodified ICRP 134 zirconium biokinetic model because the model is based primarily on data from males[1].

[1]Personal communication with Dr. Richard Leggett via email on 7/8/2021.

Appendix C

Variability, Uncertainty, and Error

Some key terms used in the dissertation are (BIPM et al., 2012;National Research Council, 2009, pg 91):

- The *measurand* in metrology (the science of measurement) is the thing we are measuring in an effort to estimate its true value. We will seldom know the true value of the measurand, i.e., for all applications of interest it is unknown and unknowable.

- *Variability* refers to the inherent heterogeneity of the true values of the measurand in a sample or population. Variability is a property of the sample or population that cannot be reduced or eliminated but whose characterization can be improved.

- An *error* is the distance between the value of the measurement and the true value of the measurand. The distance is usually expressed as an absolute or relative difference between the measurement and the true value. Note that because the true value is not knowable the error is also not knowable.

- A *residual* is the distance between the measurement and an estimate of the true value of the measurand, usually expressed as an absolute or relative difference. The residual is not the error, but nevertheless is knowable and gives us information about the error.

- *Uncertainty* refers to our lack of knowledge concerning the true value of the measurand. Sources of uncertainty are measurement errors, sampling errors, and misspecified statistical models. Steps can be taken to improve our understanding and reduce uncertainty, e.g., collect additional data or use a better model.

In conversation, error and uncertainty are often used interchangeably, but strictly speaking they are very different. For example, we often talk of measurement *error*, which is an unknowable quantity that is of little practical use. On the other hand, measurement *uncertainty* is a knowable quantity that provides our current state of knowledge about the measurand and should be reported with every measurement. In metrology, uncertainty is formally defined (JCGM, 2008) as a "parameter, associated with the result of a measurement, that characterizes the dispersion of the values that could reasonably be attributed to the measurand." They are referring specifically to the *combined standard uncertainty* u in a measurement, which can be thought of as the variation in repeated observations as described by the standard deviation plus all of the variation that we know is being generated by other sources. Thus, u incorporates the effects of all known sources of uncertainty. Note that when we refer to the uncertainty of a measurement in this report we are referring to the combined standard uncertainty.

The uncertainty in a measurement can be reported as a point estimate $\pm u$. For example, the mass of a cube of tungsten might be reported as 10.0 ± 0.1 grams, where $u = 0.1$ grams. The uncertainty can also be multiplied by an *expansion factor* like 2 to expand the uncertainty to a given probability level, e.g., we are 95% confident that the mass of the cube is 10.0 ± 0.2 grams. This uncertainty can also be expressed as the 95% *coverage* interval $(9.8, 10.2)$. These expressions of uncertainty are the metrology equivalents how we report a mean and its standard error in statistics. Statements of GUM measurement uncertainty intervals appear to be interpreted like

Bayesian credible intervals (see Kacker, 2015) and the definitions of *coverage interval* and *coverage probability* in BIPM et al. (2012) even if the uncertainty intervals were calculated with frequentist methods. This means that the true value of the measurand is said to have a given probability of being in the uncertainty interval even if it is a frequentist confidence interval.

Appendix D

Lognormal Distribution

The lognormal distribution is used extensively as the prior distribution for many parameters in biokinetic models, so it is appropriate to review some of the basic properties of this important distribution.

D.1 Probability Density Functions

The probability density function for the normal distribution is

$$f_X(x) = \frac{1}{\sigma_x\sqrt{2\pi}} e^{-\frac{1}{2}(x-\mu_x)^2/\sigma_x^2}, \tag{D.1}$$

where $-\infty < \mu_x < \infty$ is the mean, $\sigma_x > 0$ the standard deviation, and $-\infty < x < \infty$ is the support. Note that while μ_x and σ_x are the parameters of the normal pdf they are also the mean and standard deviation of X. The random variable X is normally distributed, but we can also entertain the situation where it is the log of another random variable Y, i.e., $X = \log(Y)$. The random variable Y has a *lognormal distribution* , which has the pdf

$$f_Y(y) = \frac{1}{y\sigma_x\sqrt{2\pi}} e^{-\frac{1}{2}(\log(y)-\mu_x)^2/\sigma_x^2}, \tag{D.2}$$

where $0 < y < \infty$ is the support. In contrast to the normal distribution, note that μ_x and σ_x are the parameters of the lognormal pdf but are not the mean and standard deviation of Y, which will be derived in the next section.

To show how we transform X to Y, let's first define the transform function $g(x)$ and inverse transform function $g^{-1}(y)$:

$$g(x) = y = e^x, -\infty < x < \infty, \tag{D.3}$$

$$g^{-1}(y) = x = \log(y), 0 < y. \tag{D.4}$$

These functions are depicted in Figure D.1[1], where the lines for $x = \log(7) = 1.94591$ and $y = 7$ are shown for reference. The mean μ_x of the normal distribution X is related to the *geometric mean* μ_g of the lognormal distribution Y by the same transform functions:

$$\mu_x = \log(\mu_g) \tag{D.5}$$

and

$$\mu_g = \exp(\mu_x). \tag{D.6}$$

The same applies to the standard deviation σ_x of the normal distribution and the *geometric standard deviation* σ_g of the lognormal distribution:

$$\sigma_x = \log(\sigma_g) \tag{D.7}$$

and

$$\sigma_g = \exp(\sigma_x). \tag{D.8}$$

Note that $0 < \mu_g < \infty$ and $1 < \sigma_g < \infty$. Substituting the transform and inverse transform functions in the following probability statements

$$P(Y \leq y) = P\left(g^{-1}(Y) \leq g^{-1}(y)\right) = P\left(X \leq g^{-1}(y)\right) \tag{D.9}$$

gives us

$$P(Y \leq y) = P\left(X \leq g^{-1}(y)\right), \tag{D.10}$$

which can be expressed as cumulative distribution functions (cdf):

$$F_Y(y) = F_X\left(g^{-1}(y)\right) = \int_{-\infty}^{g^{-1}(y)} \frac{1}{\sigma_x\sqrt{2\pi}} e^{-\frac{1}{2}(x-\mu_x)^2/\sigma_x^2} dx \tag{D.11}$$

[1]These three plots make what I call a DESY plot, which clearly illustrates the relationship among $g(x)$, $f_X(x)$, and $f_Y(y)$. DESY is a particle physics research center in Germany. Gerhard Bohm and Günter Zech from DESY wrote the book *Introduction to Statistics and Data Analysis for Physicists*, where I first saw this type of plot. The book can be downloaded at http://www-library.desy.de/preparch/books/vstatmp_engl.pdf.

The pdf of the lognormal distribution is derived by differentiating the cdf:

$$f_Y(y) = \frac{d}{dy}F_Y(y) = \frac{d}{dy}F_X\left(g^{-1}(y)\right) = \left[f_X\left(g^{-1}(y)\right)\right]\left[\frac{d}{dy}g^{-1}(y)\right], \tag{D.12}$$

$$f_Y(y) = \left[f_X(\log(y))\right]\left[\frac{d}{dy}\log(y)\right], \tag{D.13}$$

$$f_Y(y) = \left[\frac{1}{\sigma_x\sqrt{2\pi}}e^{-\frac{1}{2}(\log(y)-\mu_x)^2/\sigma_x^2}\right]\left[\frac{1}{y}\right] = \frac{1}{\sigma_x y\sqrt{2\pi}}e^{-\frac{1}{2}(\log(y)-\mu_x)^2/\sigma_x^2}, 0 < y. \tag{D.14}$$

Let's assign values to the parameters and look at some plots. The pdf in bottom right of Figure D.1 is a $N(\mu = \log(5), \sigma = \log(2.5))$. The area A under the curve at $F_X(\log(7) = 1.94591)$ is shaded. The pdf in the top left of Figure D.1 is a $LN(\mu = \log(5), \sigma = \log(2.5))$. The area B under the curve at $F_Y(\exp(1.94591) = 7)$ is shaded. The significance of (D.9) is that area A and area B are equal.

The method that we just used to find the distribution of a lognormal random variable is referred to as the *Distribution Function Method* (Wackerly, Mendenhall, and Scheaffer, 2014). Another way of doing this is to say that $f_Y(y)$ is the height of the density curve and dy is a tiny width – so tiny that $f_Y(y)$ is constant over dy. Under these conditions $f_Y(y)\,dy$ can be loosely thought of as probability. This sets up the following equality between two probabilities:

$$\left[f_Y(y)\right]\left[dy\right] = \left[f_X\left(g^{-1}(y)\right)\right]\left[dg^{-1}(y)\right]. \tag{D.15}$$

Bringing the differential dy over to the rhs gives us the relationship we ended up with previously in (D.12):

$$f_Y(y) = \left[f_X\left(g^{-1}(y)\right)\right]\left[\frac{d}{dy}g^{-1}(y)\right] = f_X\left(g^{-1}(y)\right)\left|\frac{d}{dy}g^{-1}(y)\right|. \tag{D.16}$$

Although it is not relevant for this transform, in general the absolute value of the derivative is used because the pdf must be non-negative. At this point the transform is completed as shown above. This is referred to as the *Method of Transformations* (Wackerly, Mendenhall, and Scheaffer, 2014).

D.2 Expectation and Variance of a Lognormal RV

Start with the definition of expectation

$$E(Y) = \int_0^\infty \left(\frac{y}{\sigma_x y \sqrt{2\pi}} \right) e^{-\frac{1}{2}\left(\frac{\log(y)-\mu_x}{\sigma_x}\right)^2} dy. \tag{D.17}$$

Make the following substitutions:

$$z = \frac{\log(y) - \mu_x}{\sigma_x}, \tag{D.18}$$

$$y = e^{\sigma_x z + \mu_x}, \tag{D.19}$$

$$dy = \sigma_x e^{\sigma_x z + \mu_x} dz, \tag{D.20}$$

$$E(Y) = e^{\mu_x} \int_{-\infty}^\infty \left(\frac{1}{\sqrt{2\pi}} \right) e^{-\frac{1}{2}z^2 + \sigma_x z} dz. \tag{D.21}$$

The lower limit on the integral with respect to y is 0, so the lower limit with respect to z is $-\infty$. Finally, complete the square in the exponent

$$E(Y) = e^{\mu_x} \int_{-\infty}^\infty \left(\frac{1}{\sqrt{2\pi}} \right) e^{-\frac{1}{2}(z-\sigma_x)^2 + \frac{1}{2}\sigma_x^2} dz, \tag{D.22}$$

and evaluate the integral, whose integrand is a normal pdf, to give the desired result

$$E(Y) = e^{\mu_x + \frac{1}{2}\sigma_x^2} \int_{-\infty}^\infty \left(\frac{1}{\sqrt{2\pi}} \right) e^{-\frac{1}{2}(z-\sigma_x)^2} dz = e^{\mu_x + \frac{1}{2}\sigma_x^2}. \tag{D.23}$$

Do the same trick for the second moment. Take Equation (21) and put $y = e^{\sigma_x z + \mu_x}$ in the integrand:

$$E\left(Y^2\right) = e^{\mu_x} \int_{-\infty}^\infty \left(\frac{1}{\sqrt{2\pi}} \right) e^{\sigma_x z + \mu_x} e^{-\frac{1}{2}z^2 + \sigma_x z} dz, \tag{D.24}$$

$$E\left(Y^2\right) = e^{2\mu_x} \int_{-\infty}^\infty \left(\frac{1}{\sqrt{2\pi}} \right) e^{-\frac{1}{2}z^2 + 2\sigma_x z} dz. \tag{D.25}$$

Complete the square in the exponent

$$-\frac{1}{2}\left(z^2 - 4\sigma_x z + 4\sigma_x - 4\sigma_x\right) = -\frac{1}{2}\left(z - 2\sigma_x\right)^2 + 2\sigma_x, \tag{D.26}$$

which, with a bit of coaxing, leads us to the second moment:

$$E\left(Y^2\right) = e^{2\mu_x + 2\sigma_x^2} \int_{-\infty}^{\infty} \left(\frac{1}{\sqrt{2\pi}}\right) e^{-\frac{1}{2}(z-2\sigma_x)^2} dz = e^{2\mu_x + 2\sigma_x^2}. \tag{D.27}$$

Recall that the variance of Y is given by

$$var\left(Y\right) = E\left(Y^2\right) - E\left(Y\right)^2, \tag{D.28}$$

so

$$var\left(Y\right) = e^{2\mu_x + 2\sigma_x^2} - \left(e^{\mu + \frac{1}{2}\sigma^2}\right)^2, \tag{D.29}$$

$$var\left(Y\right) = e^{2\mu_x + 2\sigma_x^2} - e^{2\mu_x + \sigma_x^2}, \tag{D.30}$$

and finally

$$var\left(Y\right) = \left(e^{2\mu_x + \sigma_x}\right)\left(e^{\sigma_x^2} + 1\right). \tag{D.31}$$

D.3 Say it Again Simply

If Z is a standard normal random variable then X is a normally distributed random variable

$$X = \mu_x + \sigma_x Z, \tag{D.32}$$

where μ_x is the mean of X (the mean log parameter) and σ_x the standard deviation of X (the standard deviation log parameter). If Y is given by

$$Y = e^X = e^{\mu_x + \sigma_x Z}, \tag{D.33}$$

then Y is a lognormally distributed random variable. The geometric mean (the median) of Y occurs when $Z = 0$, which is

$$\mu_g = e^{\mu_x}. \tag{D.34}$$

The standard deviation of $X = \mu_x + \sigma_x(1) - \mu_x = \sigma_x$, so the geometric standard deviation of Y is

$$\sigma_g = e^{\mu_x + \sigma_x(1) - \mu_x} = e^{\sigma_x}. \tag{D.35}$$

Note that

$$\sigma_g = e^{\sqrt{\sigma_x^2}} \tag{D.36}$$

where σ_x^2 is the variance of X. The arithmetic mean of Y is

$$\mu_y = e^{\mu_x + \frac{1}{2}\sigma_x^2} \tag{D.37}$$

and the arithmetic standard deviation

$$\sigma_y = \left(e^{2\mu_x + \sigma_x}\right)\left(e^{\sigma_x^2} + 1\right). \tag{D.38}$$

Note that the geometric parameters are always smaller than the arithmetic parameters

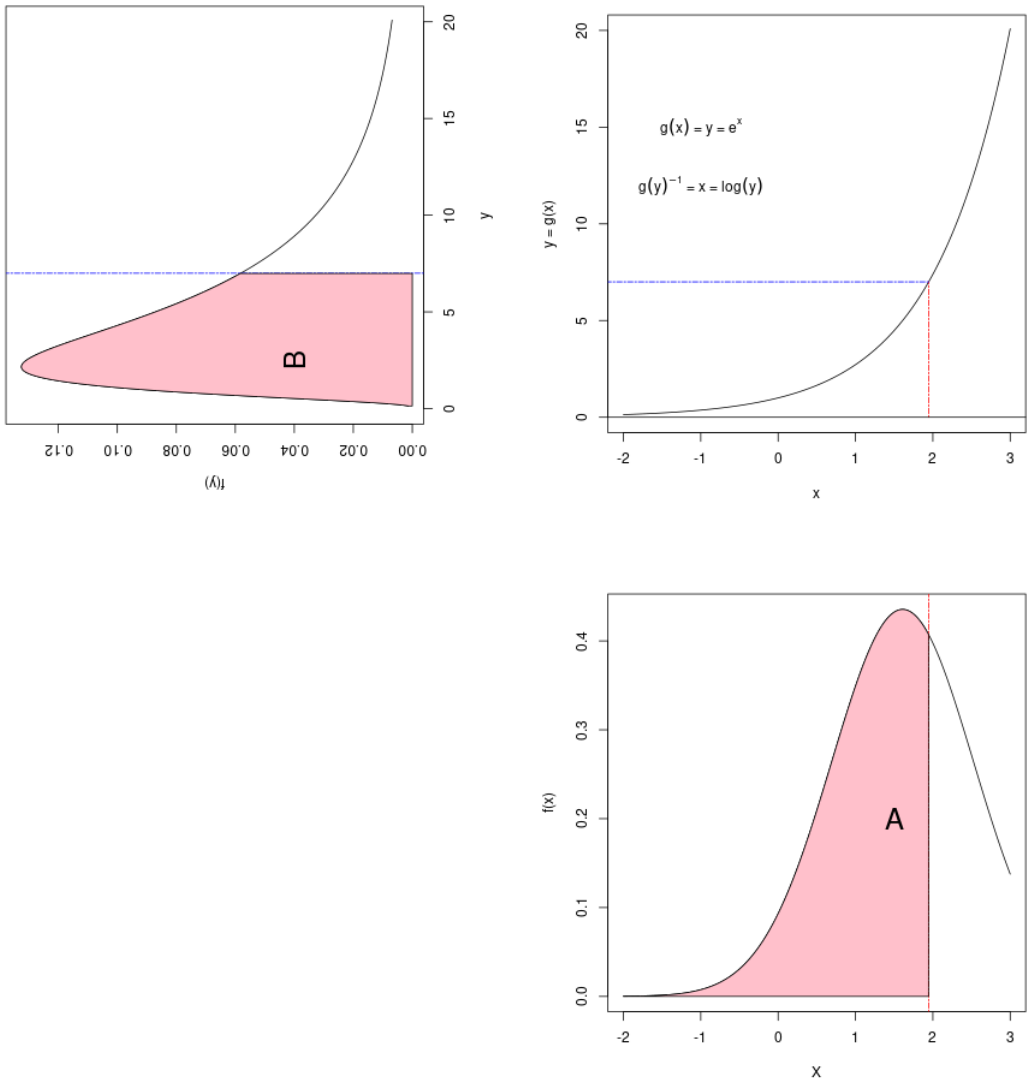

Figure D.1 DESY plot. Starting in the lower right hand corner and going counter-clockwise: pdf of the normally distributed random variable X, the transform function g(x), and the pdf of the lognormally distributed random variable Y.

Appendix E

LKJ Distribution

The approach recommended by Barnard, McCulloch, and Meng (2000) and Stan Development Team (2021b, pg 34) for defining a prior on a covariance matrix $\mathbf{\Pi}$ is to decompose it into a correlation matrix Ω and a diagonal matrix of the standard deviations σ

$$\mathbf{\Pi} = \boldsymbol{\sigma} \cdot \mathbf{\Omega} \cdot \boldsymbol{\sigma}. \tag{E.1}$$

This decomposition allows the modeler to think in terms of the standard deviation of the parameters and a correlation matrix, which is usually easier than thinking about a covariance matrix. The Lewandowski–Kurowicka–Joe (LKJ) distribution is used to model the distribution of $\mathbf{\Omega}$. The LKJ distribution has one argument η, and it places a shifted beta distribution with a support of -1 to 1 on each of the off-diagonal elements of the correlation matrix.

To illustrate, let's generate $m = 10^4$ $n \times n$ random correlation matrices using the LKJ distribution as implemented in the *rethinking* R package with $\eta = 1$ and $n = 3$. These random correlation matrices are stored in the 3-dimenional array *Omega*.

```
Omega <- rlkjcorr(1e4,K=3,eta=1)


dim(Omega)
## [1] 10000     3     3
```

The first index of Omega gives the ith 3×3 correlation matrix where $i = 1, 2, \ldots, m$. For example, setting the first index to 1 gives the first of 10^4 random correlation matrices,

```
Omega[1,,]
##              [,1]      [,2]         [,3]
## [1,]  1.000000000 0.3022829 -0.007865985
## [2,]  0.302282874 1.0000000  0.801085116
## [3,] -0.007865985 0.8010851  1.000000000
```

and setting it equal to 2 gives the second of 10^4.

```
Omega[2,,]
##             [,1]       [,2]       [,3]
## [1,]  1.0000000 -0.4419927 -0.5949270
## [2,] -0.4419927  1.0000000 -0.2047289
## [3,] -0.5949270 -0.2047289  1.0000000
```

The second index gives the jth row of the ith correlation matrix where $j = 1, 2, \ldots, n$. For example, the first row of the first correlation matrix is

```
Omega[1,1,]
## [1]  1.000000000  0.302282874 -0.007865985
```

Finally, the third index gives the element in the kth column of the ith correlation matrix where $k = 1, 2, \ldots, n$. For example the second element in the first row of the first correlation matrix is

```
Omega[1,1,2]
## [1] 0.3022829
```

This allows us to look at the density plots of the three off-diagonal elements of the 10^4 correlation matrices with $\eta = 1$.

```
plot(density(Omega[,1,2]),
   main="Eta = 1",xlab="Correlation Coefficient",
   xlim=c(-1,1)
```

```
)
lines(density(Omega[,1,3]),col="red")
lines(density(Omega[,2,3]),col="blue")
```

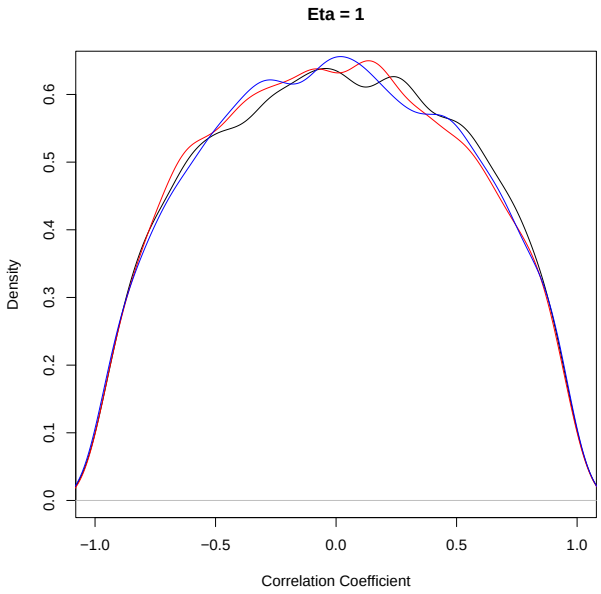

Setting $\eta = 0.5$ gives distributions that look more or less uniform.

```
Omega <- rlkjcorr(1e4,K=3,eta=0.5)
```

```
plot(density(Omega[,1,2]),
   main="Eta = 0.5",xlab="Correlation Coefficient",
   xlim=c(-1,1)
)
lines(density(Omega[,1,3]),col="red")
lines(density(Omega[,2,3]),col="blue")
```

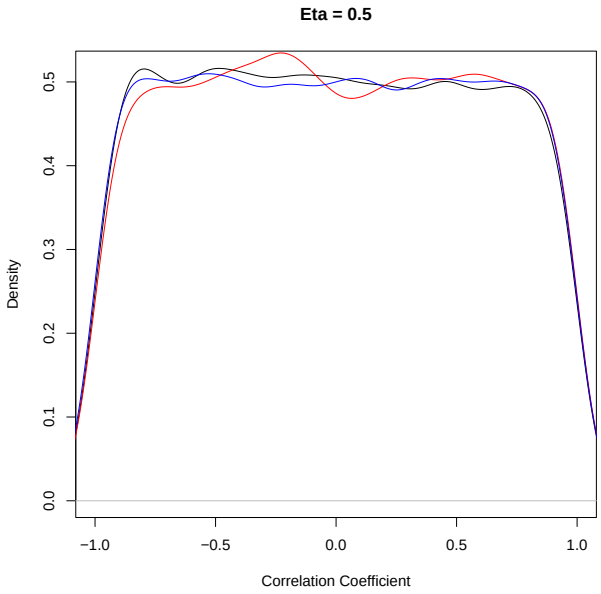

Setting $\eta = 4$ concentrates the density more to the middle of the distributions, giving something that looks Gaussian.

```
Omega <- rlkjcorr(1e4,K=3,eta=4)
```

```
plot(density(Omega[,1,2]),
   main="Eta = 4",xlab="Correlation Coefficient",
   xlim=c(-1,1)
)
lines(density(Omega[,1,3]),col="red")
lines(density(Omega[,2,3]),col="blue")
```

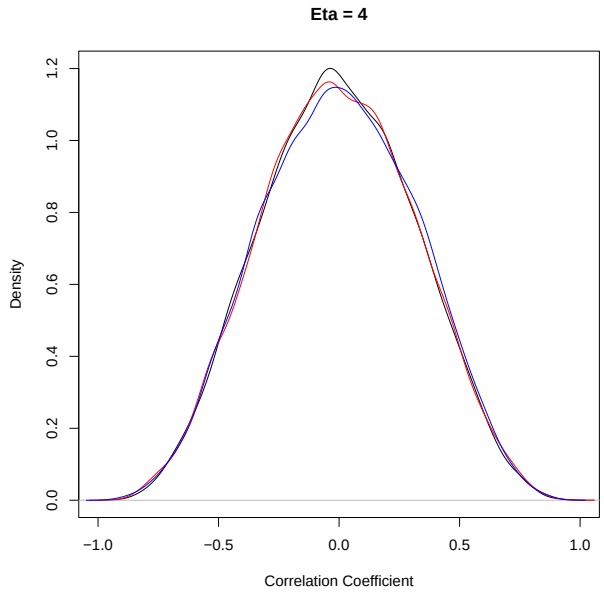

Setting $\eta = 10$ results in a tighter distribution ...

```
Omega <- rlkjcorr(1e4,K=3,eta=10)
```

```
plot(density(Omega[,1,2]),
   main="Eta = 10",xlab="Correlation Coefficient",
   xlim=c(-1,1)
)
lines(density(Omega[,1,3]),col="red")
lines(density(Omega[,2,3]),col="blue")
```

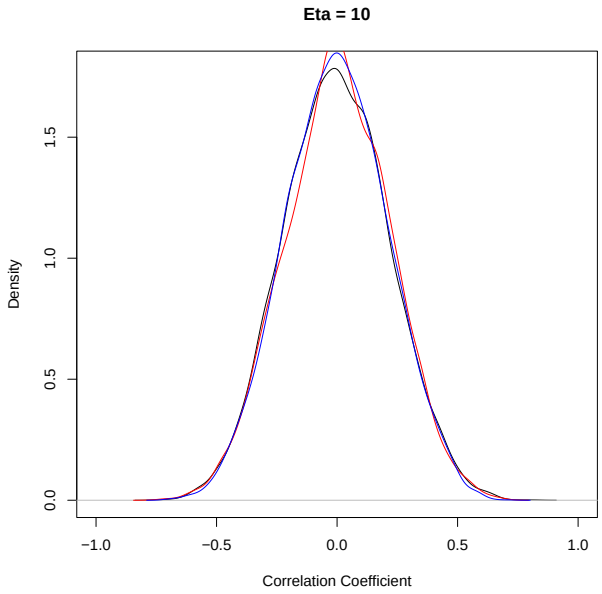

... and setting $\eta = 100$ gives essentially no correlation.

```
Omega <- rlkjcorr(1e4,K=3,eta=100)
```

```
plot(density(Omega[,1,2]),
   main="Eta = 100",xlab="Correlation Coefficient",
   xlim=c(-1,1)
)
lines(density(Omega[,1,3]),col="red")
lines(density(Omega[,2,3]),col="blue")
```

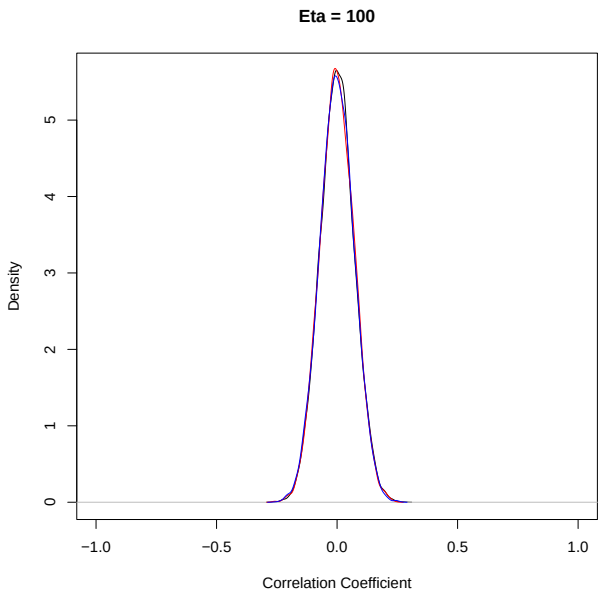

If all the elements of $\boldsymbol{\sigma}$ are equal, σ is a scalar and equation (E.1) can be simplified as

$$\boldsymbol{\Pi} = \sigma^2 \cdot \boldsymbol{\Omega} \, . \tag{E.2}$$

The standard deviation is referred to as $\tau = \sigma$ elsewhere in this dissertation.

Appendix F

ICRP 134 Zirconium Model

Table F.1 Names of the compartments in the ICRP 134 zirconium model (Paquet, Leggett, et al., 2016, page 274). Note that ST0 and ST1 are soft tissue compartments, SI is the small intestines, and trabecular and cortical refer to parts of skeletal bone.

Index	From	Symbol	Comp	To	Symbol	Comp
1	Blood 1	bld1	1	Blood 2	bld2	2
2	Blood 1	bld1	1	Liver 0	liv0	3
3	Blood 1	bld1	1	Kidneys	kid	5
4	Blood 1	bld1	1	ST0	ST0	6
5	Blood 1	bld1	1	ST1	ST1	7
6	Blood 1	bld1	1	Urinary bladder cont	ubc	12
7	Blood 1	bld1	1	SI contents	SI	13
8	Blood 1	bld1	1	Trabecular surface	ts	8
9	Blood 1	bld1	1	Cortical surface	cs	10
10	Blood 2	bld2	2	Blood 1	bld1	1
11	Liver 0	liv0	3	SI contents	SI	13
12	Liver 0	liv0	3	Blood 1	bld1	1
13	Liver 0	liv0	3	Liver 1	liv1	4
14	Liver 1	liv1	4	Blood 1	bld1	1
15	Kidneys	kid	5	Blood 1	bld1	1
16	ST0	ST0	6	Blood 1	bld1	1
17	ST1	ST1	7	Blood 1	bld1	1
18	Trabecular surface	ts	8	Blood 1	bld1	1
19	Trabecular surface	ts	8	Trabecular volume	tv	9
20	Trabecular volume	tv	9	Blood 1	bld1	1
21	Cortical surface	cs	10	Blood 1	bld1	1
22	Cortical surface	cs	10	Cortical volume	cv	11
23	Cortical volume	cv	11	Blood 1	bld1	1
24	Urinary bladder cont	ubc	12	Urine	urn	15
25	SI contents	SI	13	Colon contents	col	14
26	Colon contents	col	14	Feces	fec	16
27	SI contents	SI	13	Blood 1	bld1	1

Appendix G

Likelihood and Bayes Rule

To clarify the notation used in Bayes Rule, let's look at an example where the probability of observing k successes in n trials for a binomial data generating process having a probability θ of success:

$$P\left(K=k|\theta\right)=\left(\begin{array}{c} n \\ k \end{array}\right)\theta^{k}\left(1-\theta\right)^{n-k},\ k=\left\{0,1,2,\ldots,n\right\},\ 0\leq\theta\leq1. \tag{G.1}$$

Note that the data K are random and the probability of success θ is constant. This is a proper probability mass function since

$$\sum_{i=0}^{i=n}P\left(i|\theta\right)=1. \tag{G.2}$$

If we want to calculate the probability that a random probability of success Θ equals θ for a given value of k we need to use Bayes Rule

$$P\left(\Theta=\theta|k\right)=\frac{P\left(k|\theta\right)P\left(\theta\right)}{P\left(k\right)}, \tag{G.3}$$

where

$$P\left(k\right)=\int_{\theta}P\left(k|\theta\right)P\left(\theta\right)d\theta \tag{G.4}$$

Note that in Equation (G.3) we have

$$P\left(k|\theta\right)=\left(\begin{array}{c} n \\ k \end{array}\right)\theta^{k}\left(1-\theta\right)^{n-k}, \tag{G.5}$$

the rhs of which looks just like the rhs of Equation (G.1). However, in Equation (G.5) the data k are fixed and the probability of success θ is random. Equation (G.5) is

not a proper probability density function since

$$\int_{\theta=0}^{\theta=1} P\left(k|\theta\right) d\theta \neq 1. \tag{G.6}$$

To make the difference between Equation (G.1) and Equation (G.5) clear, Equation (G.5) is referred to as the *likelihood* of θ given the data k

$$L\left(\theta|k\right) = L\left(\theta\right) = P\left(k|\theta\right) = \left(\begin{array}{c} n \\ k \end{array} \right) \theta^k \left(1-\theta\right)^{n-k}. \tag{G.7}$$

So, sometimes we see Bayes Rule expressed like

$$P\left(\Theta = \theta|k\right) = \frac{L\left(\theta\right) P\left(\theta\right)}{P\left(k\right)}, \tag{G.8}$$

which is a bit confusing in my opinion.

When we are working with continuous data X rather than discrete data, the probability of observing $X = x$ is 0. If the data are measured with *good precision* the appropriate probability density function is used as the likelihood (Pawitan, 2001, page 23). For example, if the data come from a normal data generating process and have good precision, the normal pdf is used as the likelihood in Equation (3). Examples of data that do not have good precision are censored data and binned data.

Appendix H

Fits to Samples from the Posterior Distribution of Theta

In the population-level Bayesian calibration in Chapter 3, 110000 posterior samples are generated from the rate matrix hyper prior $\boldsymbol{\theta}$. A familiar distribution must be fit to these samples to make canned priors that are used for an out-of-sample evaluation in Chapter 4. Here, four different distributions are fit to the posterior samples: multivariate normal (MVN), multivariate skewed t with 4 degrees of freedom (MVT4), multivariate skewed t with 20 degrees of freedom (MVT20), and a multivariate Gaussian mixture (MGM). AIC and BIC are used to compare the models. To begin, let's read in the posterior samples.

```
theta <- read.csv("data/theta.csv")
dim(theta)
## [1] 110000     27
n <- nrow(theta)
```

Fit the samples with a multivariate normal distribution. The number of parameters estimated is 27 for the means plus $(27^2 - 27)/2$ for the covariance matrix.

```
library(mclust)


fit.mvn <- mvn("Ellipsoidal", data = theta)
Mu <- as.numeric(fit.mvn$parameters$mean)
Sigma <- fit.mvn$parameters$variance$Sigma
```

```
loglik.mvn <- fit.mvn$loglik
k.mvn <- 27 + (27^2 - 27)/2
AIC.mvn <- 2*k.mvn - 2*loglik.mvn
BIC.mvn <- k.mvn*log(n) - 2*loglik.mvn
```

```
loglik.mvn
## [1] -1044802
k.mvn
## [1] 378
AIC.mvn
## [1] 2090361
BIC.mvn
## [1] 2093992
```

Now try a multivariate skew t with $df = 4$. This is a shifted and scaled t distribution that can be non-symmetric. Alpha is the skew parameter, which can be in the range of -1 to 1. Here Alpha is essentially 0, which means the t distributions are symmetric, i.e., they are standard t distributions.

```
library(fMultivar)
library(sn)

fit.mvt <- mstFit(theta, fixed.nu=4)
Beta <- fit.mvt@fit$dp$beta
Omega <- fit.mvt@fit$dp$Omega
Alpha <- fit.mvt@fit$dp$alpha
Alpha
##       theta.1       theta.2       theta.3       theta.4       theta.5
## -1.932956e-06  3.550768e-05  3.172413e-07 -6.052230e-07  1.721517e-05
##       theta.6       theta.7       theta.8       theta.9      theta.10
## -7.904359e-07 -2.205950e-06 -3.415253e-05 -1.144856e-05  4.809289e-08
##      theta.11      theta.12      theta.13      theta.14      theta.15
```

```
##  2.573121e-06 -6.984653e-07 -2.047055e-06  3.795500e-06  3.505281e-06
##      theta.16      theta.17      theta.18      theta.19      theta.20
## -2.253202e-06  7.319625e-07 -1.287202e-04  1.509607e-06  2.789305e-07
##      theta.21      theta.22      theta.23      theta.24      theta.25
## -1.908804e-06  1.944516e-06 -2.546741e-07  4.334784e-06 -1.359099e-06
##      theta.26      theta.27
##  2.454062e-06  2.093020e-07
```

```
tmp <- apply(theta,1,dmst,xi=Beta,Omega=Omega,alpha=Alpha, nu=4,log=TRUE)
loglik.mvt4 <- sum(tmp)
k.mvt4 <- 27 + (27^2 - 27)/2 + 27
AIC.mvt4 <- 2*k.mvt4 - 2*loglik.mvt4
BIC.mvt4 <- k.mvt4*log(n) - 2*loglik.mvt4
```

```
loglik.mvt4
## [1] -1113547
k.mvt4
## [1] 405
AIC.mvt4
## [1] 2227903
BIC.mvt4
## [1] 2231795
```

Now let's look at a multivariate skew t with $df = 20$.

```
fit.mvt <- mstFit(theta, fixed.nu=20)
Beta <- fit.mvt@fit$dp$beta
Omega <- fit.mvt@fit$dp$Omega
Alpha <- fit.mvt@fit$dp$alpha
Alpha
##       theta.1       theta.2       theta.3       theta.4       theta.5
## -4.175385e-06  7.269093e-05  1.021783e-06 -1.274062e-06  3.526286e-05
```

```
##       theta.6       theta.7       theta.8       theta.9      theta.10
## -1.695614e-06 -4.252899e-06 -7.042894e-05 -2.302558e-05  1.306991e-07
##      theta.11      theta.12      theta.13      theta.14      theta.15
##  4.491104e-06 -1.750862e-06 -3.220527e-06  7.588377e-06  6.883004e-06
##      theta.16      theta.17      theta.18      theta.19      theta.20
## -3.946342e-06  9.130683e-07 -2.623345e-04  3.448467e-06  3.995365e-07
##      theta.21      theta.22      theta.23      theta.24      theta.25
## -3.631377e-06  3.807114e-06 -5.545758e-07  9.237500e-06 -2.663658e-06
##      theta.26      theta.27
##  5.006424e-06  4.630544e-07
```

```
tmp <- apply(theta,1,dmst,xi=Beta,Omega=Omega,alpha=Alpha, nu=20,log=TRUE)
loglik.mvt20 <- sum(tmp)
k.mvt20 <- 27 + (27^2 - 27)/2 + 27
AIC.mvt20 <- 2*k.mvt20 - 2*loglik.mvt20
BIC.mvt20 <- k.mvt20*log(n) - 2*loglik.mvt20
```

```
loglik.mvt20
## [1] -1062657
k.mvt20
## [1] 405
AIC.mvt20
## [1] 2126123
BIC.mvt20
## [1] 2130015
```

Finally, let's try fitting a 2 component multivariate Gaussian mixture model. The number of parameters estimated is 2 times the number of parameters for the multivariate normal plus 2 for the mixing fractions. Note that this takes a couple of hours to run with 110000 samples.

```
library(mixtools)


fit.mgm <- mvnormalmixEM(theta, lambda = c(0.5,0.5))
```

```
loglik.mgm <- fit.mgm$loglik
k.mgm <- 2*(27 + (27^2 - 27)/2) + 2
AIC.mgm <- 2*k.mgm - 2*loglik.mgm
BIC.mgm <- k.mgm*log(n) - 2*loglik.mgm
```

```
loglik.mgm
## [1] -1025952
k.mgm
## [1] 758
AIC.mgm
## [1] 2053419
BIC.mgm
## [1] 2060702
```

Here are the mixing fractions ...

```
fit.mgm$lambda
## [1] 0.381987 0.618013
```

... and the means.

```
fit.mgm$mu
## [[1]]
##  [1]   1.1587667  -2.3449260  -4.4660132   0.9553475  -3.2217875  -2.6374344
##  [7]  -3.8139564  -1.4153139  -1.4761301  -0.1329207  -2.2608834  -2.1553302
## [13]  -0.6771897  -4.6067532  -4.6176883  -0.5562580  -3.9300067  -7.4190567
## [19]  -8.3163004  -7.6139695  -9.4441286 -10.0954131  -9.4090838   0.7609981
## [25]   0.4118240   0.4188153  -5.8091308
##
```

```
## [[2]]
##  [1]   1.2011352  -2.8199514  -4.5079234   0.9367942  -3.4748857  -2.5542416
##  [7]  -3.8044409  -0.9671406  -1.4834241  -0.2122716  -2.1665035  -2.0904156
## [13]  -0.8256816  -4.6447395  -4.6362342  -0.6643702  -3.8780702  -6.7139352
## [19]  -8.3655028  -7.6149308  -9.4615370 -10.0960571  -9.4056106   0.6691862
## [25]   0.4023371   0.4071487  -5.8065572
```

The models, in order of both AIC and BIC (lowest to highest – best to worst) are $MGM < MVN < MVT20 < MVT4$.

Table H.1 Comparison of fits to the samples from the posterior distribution of $\boldsymbol{\theta}$.

Model	loglik	k	AIC	BIC
MGM	-1.0259516×10^6	758	2.0534192×10^6	2.0607023×10^6
MVN	-1.0448023×10^6	378	2.0903605×10^6	2.0939924×10^6
MVT20	-1.0626567×10^6	405	2.1261235×10^6	2.1300148×10^6
MVT4	-1.1135466×10^6	405	2.2279032×10^6	2.2317946×10^6

Appendix I

Chapter 4 Addendum

Additional plots of interest from Chapter 4 are presented in this appendix.

I.1 Canned Priors

The multivariate normal distribution fit to the log of the posterior distribution of the rate matrix $\boldsymbol{\theta}$ gives the canned prior $\boldsymbol{\theta}_c$. The gm and gsd of the marginal distributions of the rate constants in θ_c are given in Table I.1 along with the ICRP 134 rate constants for comparison. Comparing the medians of the posterior distribution of $\boldsymbol{\theta}$ in Table 3.1 with the geometric means of the MVN canned prior $\boldsymbol{\theta}_c$ in Table I.1 we can see that the MVN canned prior $\boldsymbol{\theta}_c$ captures the essence of the posterior distribution of $\boldsymbol{\theta}$ very well. The marginal distributions of the canned prior $\boldsymbol{\theta}_c$ are shown in Figure I.1. The distributions of most of the individual rate constants in $\boldsymbol{\theta}_c$ strongly resemble their priors (those whose peak density is ~ 1 in the plot), which means that those parameters are more or less non-identifiable given the bioassay data available in the study. However, the distributions of the 7 rate constants associated with the blood plasma and urine (distributions in red in the plots in Figure I.1 are clearly influenced by the data. The reduction in the gsd of rate constants influenced by the data can be seen in Figure I.3. The correlation plot of $\boldsymbol{\theta}_c$ in Figure I.2 shows that there are strong correlations between some rate constants in the zirconium biokinetic model that are accounted for in the canned prior. Such correlations are important when making modifications to the zirconium biokinetic model to better fit the data from an individual.

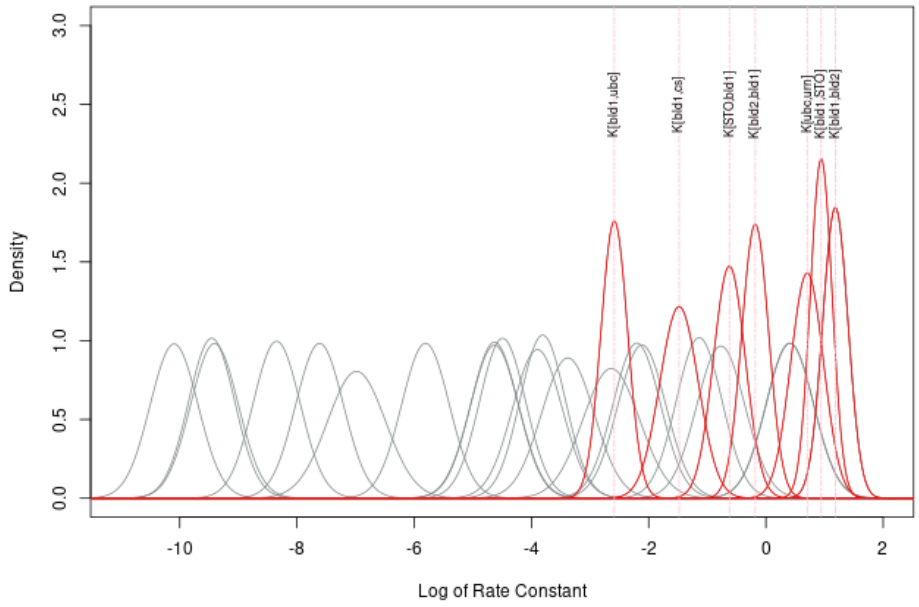

Figure I.1 Plot of the marginal distributions of the canned prior for the rate constants $\boldsymbol{\theta_c}$.

I.2 Evaluation of Out-of-Study Subjects

Stan MCMC diagnostic plots for the canned Bayesian analysis of bioassay from the subjects are shown in Figure I.4. An $\hat{R} < 1.01$ for all parameters indicates that the MCMC chains converged while the large number of effective samples ESS allow for precise parameter estimates as indicated by the relative Monte Carlo standard errors being less than 10%. No divergences were reported by Stan and there are no overt signs of non-convergence.

The plot of the geometric standard deviations for each rate constant in the rate matrix $\boldsymbol{k}$ are given in Figure I.7 for Zr0101102 and Figure I.10 for subject Zr0102004. The correlation plots for $\boldsymbol{k}$ are given in Figure I.6 and Figure I.8 show that the correlations in the rate matrices $\boldsymbol{k}$ for the two subjects are dissimilar as are the medians of the rate constants whose posterior is influenced by the data (Table 3.1, Figures I.7 and I.10). This indicates that the ICRP 134 zirconium biokinetic model

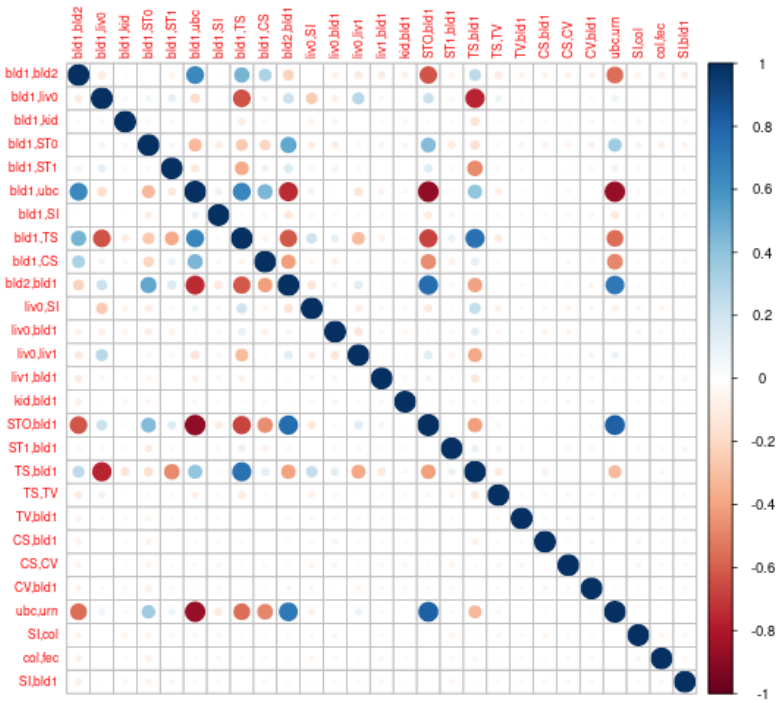

Figure I.2 Correlation plot for canned prior of $\boldsymbol{\theta_c}$.

was customized in the Bayesian analysis in order to accommodate peculiarities in each subject's the bioassay data (which are discussed below). Of particular interest is subject Zr0101102's rate constant k[ubc,urn] for the urinary bladder to urine pathway. All the posterior distributions of her rate constants (and those of Zr0102004 in Figure I.9) were more or less lognormally distributed except for this one, which can be seen in Figure I.5. The k[ubc,urn] rate constant for Zr0101102 is highly left skewed, almost looking bimodal, which results in a relatively large empirical geometric standard deviation as shown in Figure I.7. This behavior is attributed to the urine bioassay data for Zr0101102 having outliers, which results in significant changes to this rate constant in the Bayesian analysis.

Scatter plots of $\boldsymbol{M}$ versus $\boldsymbol{m}$ on log-log scales in Figures I.11 more clearly show the influence of the urine data on the fits. Posterior predictive distribution plots give an idea of how future bioassay datasets might compare to the bioassay actually observed and hence how well our biokinetic model agrees with the observed data.

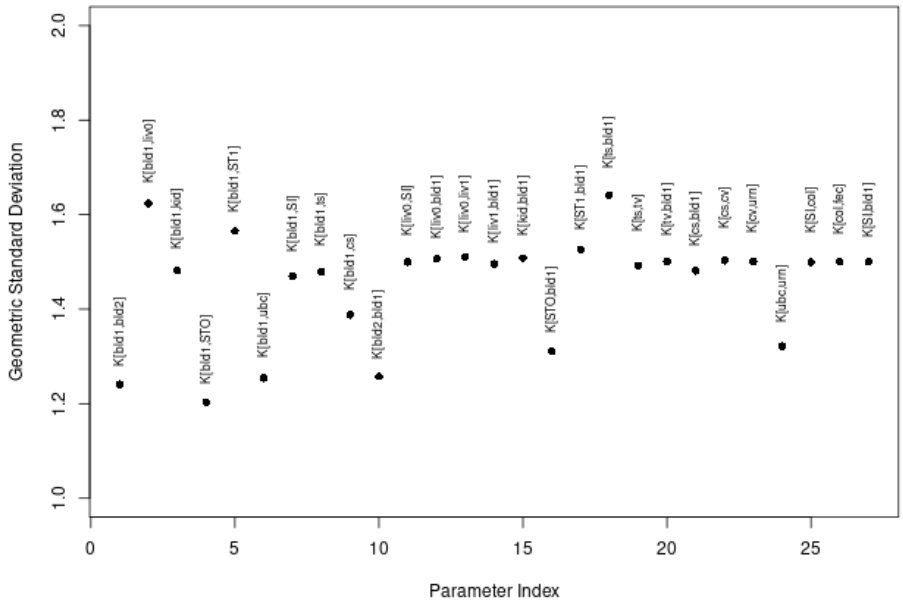

Figure I.3 Plot of $\boldsymbol{\theta_c}$ geometric standard deviations for canned prior distributions of rate constants. Rate constants associated with each parameter index are given in Table I.1.

In the plots in Figure I.12 the x-axis value for each of the points is a reference bioassay function m calculated from a draw of biokinetic model parameters from their posterior distributions. These reference bioassay functions and draws from the posterior distribution of the intake are inserted in the data sampling distribution (the likelihood) to generate predicted bioassay result $\hat{\boldsymbol{M}}$ that is the y-axis value. The line through the origin in each plot is formed by plotting the means of the values of each m versus the means of corresponding values of $\hat{\boldsymbol{M}}$, and the slope of this line is the estimated quantity $\hat{\beta}$ of zirconium injected into the subject. The error bars are the 95% credible interval of the model predictions. The amount of scatter in the plots is impressive, and gives us an idea of the true level of uncertainty in the biokinetic models. The greater degree of scatter observed with the canned prior analysis compared to the add-one-in analysis is attributed to the greater number of MCMC iterations in the canned prior analysis.

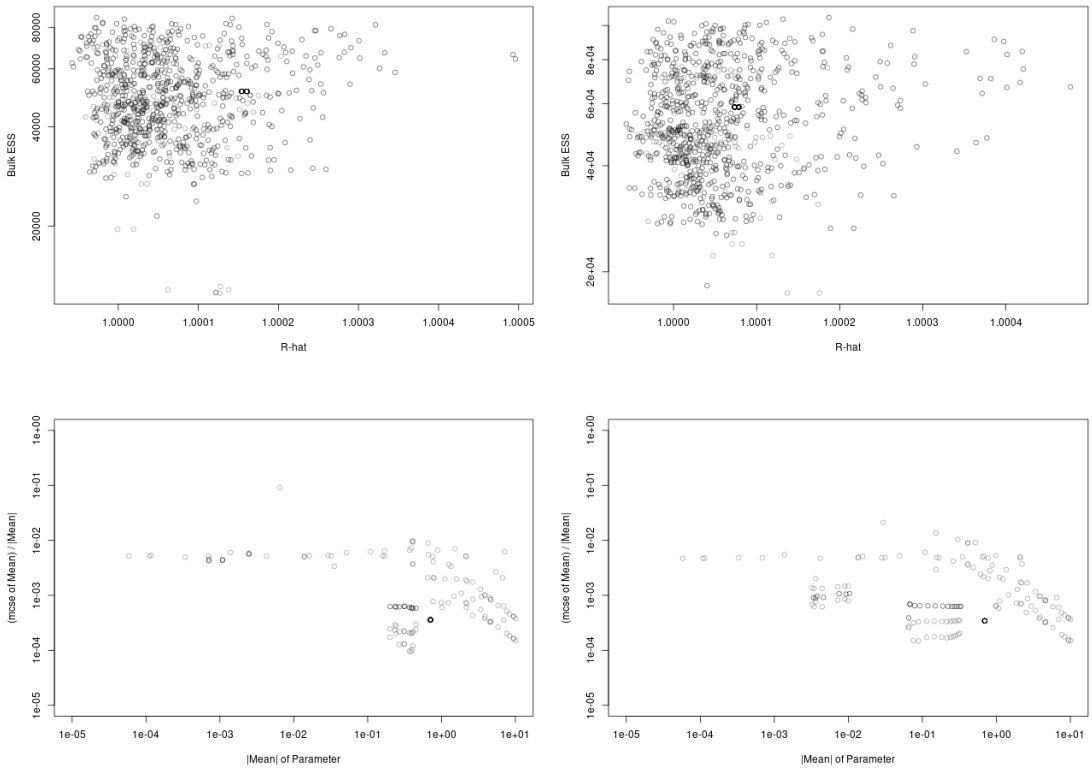

Figure I.4 Bulk ESS versus $\hat{R}$ for the posterior distributions of all parameters for Zr0101102 (top left) and Zr0102004 (top right). Relative Monte Carlo standard error of the mean versus the absolute value of the mean for Zr0101102 (bottom left) and Zr0102004 (bottom right). The standard error plots do not include the correlation matrix $\mathbf{\Pi}$ and covariance matrix $\mathbf{\Omega}$

Table I.1 Transfer rate constants in units of day^{-1} for the ICRP 134 zirconium compartmental model (Paquet, Leggett, et al., 2016), the MVN canned priors, and the posterior distributions for the out-of-study subjects. The geometric standard deviation for each calculated rate constant is given.

		ICRP 134	Canned	Canned	Zr0101102	Zr0101102	Zr0102004	Zr0102004
From	To	rate const	gm	gsd	gm	gsd	gm	gsd
bld1	bld2	2	3.276	1.04758	6.343	1.41264	1.163	1.38985
bld1	liv0	0.075	0.07181	1.26506	0.07203	2.4639	0.07411	2.41873
bld1	kid	0.0125	0.01123	1.16798	0.01116	2.31675	0.01133	2.25106
bld1	ST0	2	2.571	1.03511	4.982	1.53694	6.497	1.14656
bld1	ST1	0.0375	0.03388	1.22335	0.03468	2.40423	0.03438	2.33268
bld1	ubc	0.1	0.07532	1.05236	0.03132	1.65294	0.1461	1.37253
bld1	SI	0.025	0.02221	1.16146	0.02238	2.29202	0.02245	2.25601
bld1	ts	0.375	0.3217	1.1638	0.4069	2.60368	0.3707	2.26753
bld1	cs	0.375	0.2263	1.11338	0.2621	2.38575	0.2529	2.22777
bld2	bld1	0.462	0.8345	1.05357	0.6688	1.74955	0.4174	1.77419
liv0	SI	0.116	0.1107	1.17522	0.1112	2.29696	0.1108	2.23692
liv0	bld1	0.116	0.121	1.18291	0.1201	2.31078	0.1208	2.24973
liv0	liv1	0.462	0.4626	1.18772	0.4586	2.33802	0.4648	2.2652
liv1	bld1	0.01	0.009729	1.17553	0.009782	2.30123	0.009722	2.28468
kid	bld1	0.01	0.009726	1.18391	0.009842	2.29646	0.009755	2.26171
ST0	bld1	0.462	0.5356	1.07516	0.2879	2.01261	0.7429	1.57119
ST1	bld1	0.02	0.02034	1.19453	0.0204	2.32891	0.02003	2.27056
ts	bld1	0.000493	0.0009293	1.27893	0.0009473	2.4306	0.0009385	2.37207
ts	tv	0.000247	0.0002375	1.17311	0.0002349	2.30528	0.0002359	2.25652
tv	bld1	0.000493	0.0004923	1.18091	0.0004933	2.30262	0.0004947	2.25616
cs	bld1	8.21E-05	7.753e-05	1.16484	7.868e-05	2.28167	7.81e-05	2.24287
cs	cv	4.11E-05	4.127e-05	1.17966	4.168e-05	2.32836	4.158e-05	2.26758
cv	bld1	8.21E-05	8.154e-05	1.17827	8.112e-05	2.31484	8.28e-05	2.28499
ubc	urn	12	2.027	1.07966	4.555	3.06541	1.978	1.61825
SI	col	1.5	1.502	1.17878	1.502	2.33164	1.515	2.26836
col	fec	1.5	1.515	1.18028	1.501	2.31636	1.513	2.25847
SI	bld1	0.0030	0.002999	1.18063	0.00296	2.30146	0.003013	2.25914

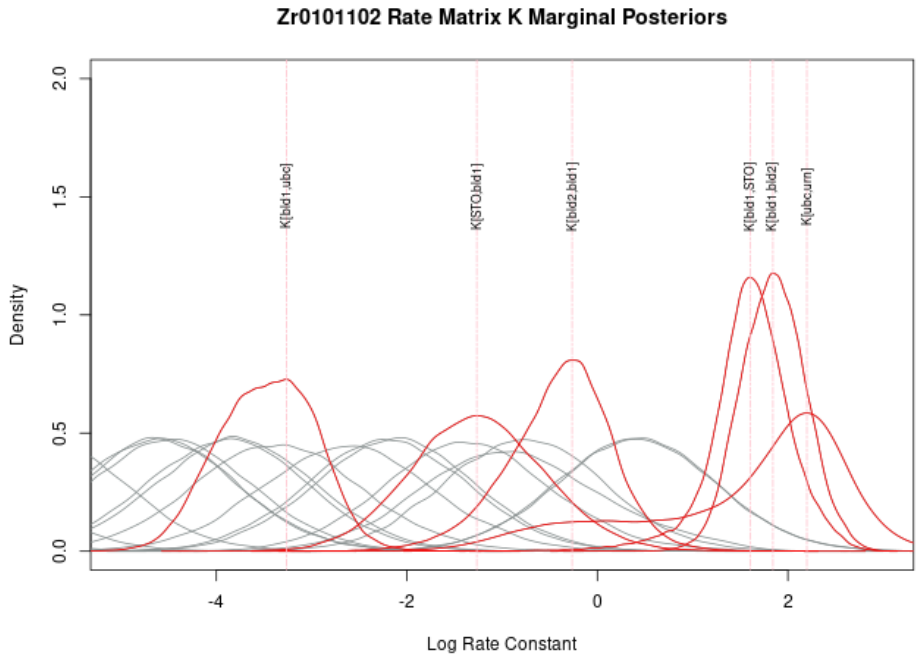

Figure I.5 Marginal posterior distributions for individual-specific rate matrix $\boldsymbol{k}$ for subject Zr0101102. Density curves for rate constants with distributions having a log standard deviation less than the one prescribed in the prior are in red and are identified by name.

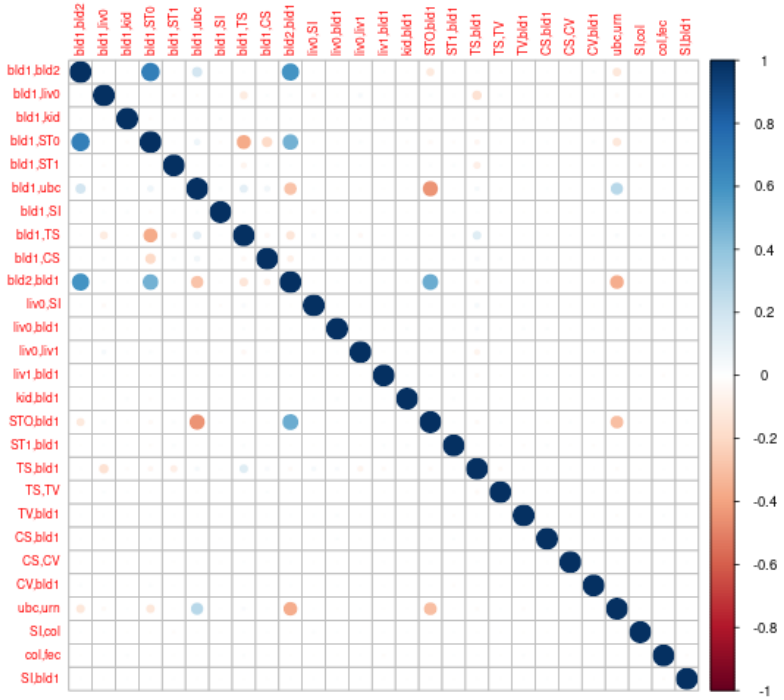

Figure I.6 Correlation plot for posterior of rate matrix $\boldsymbol{k}$ for subject Zr0101102.

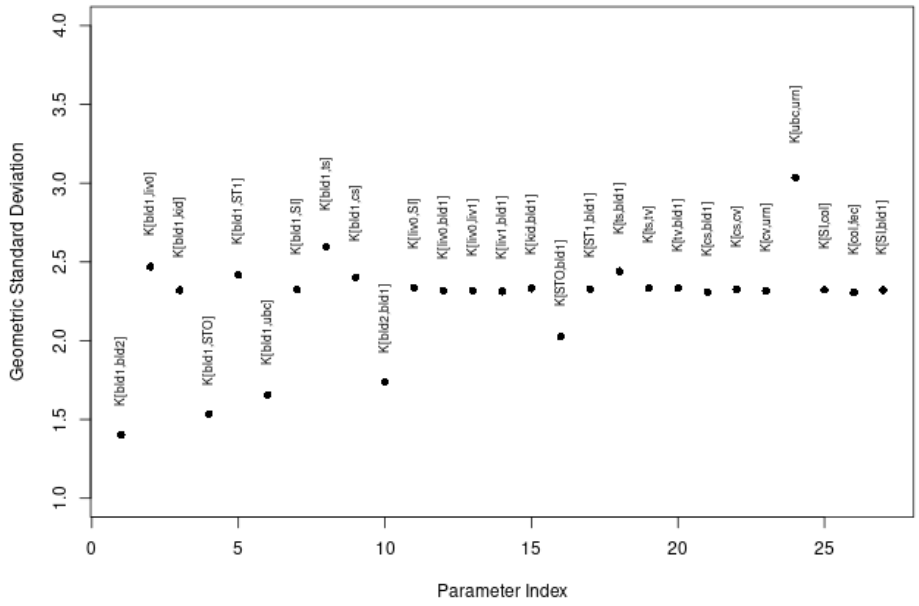

Figure I.7 Plot of geometric standard deviations for posterior distributions of rate constants $\boldsymbol{k}$ for subject Zr0101102.

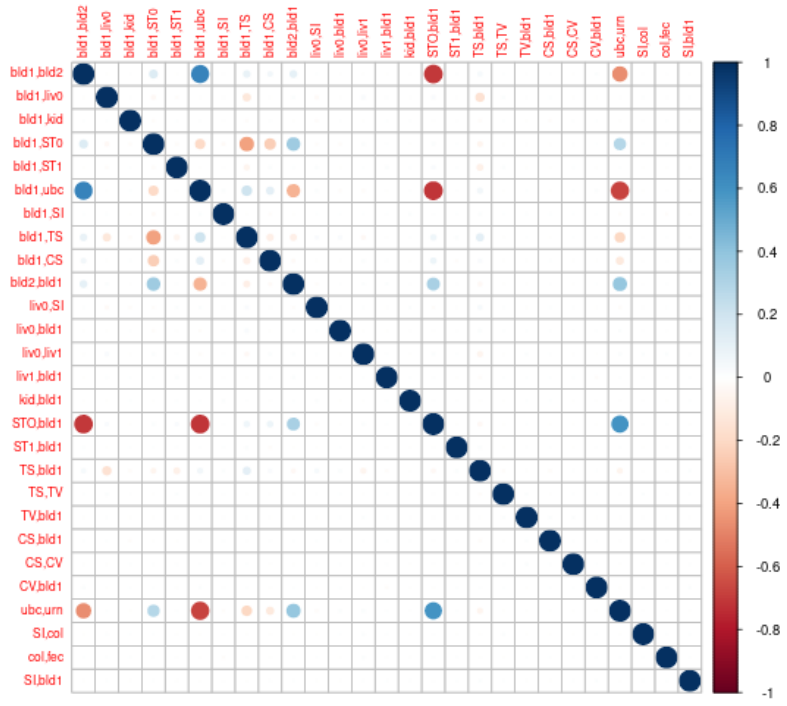

Figure I.8 Correlation plot for posterior of rate matrix $\boldsymbol{k}$ for subject Zr0102004.

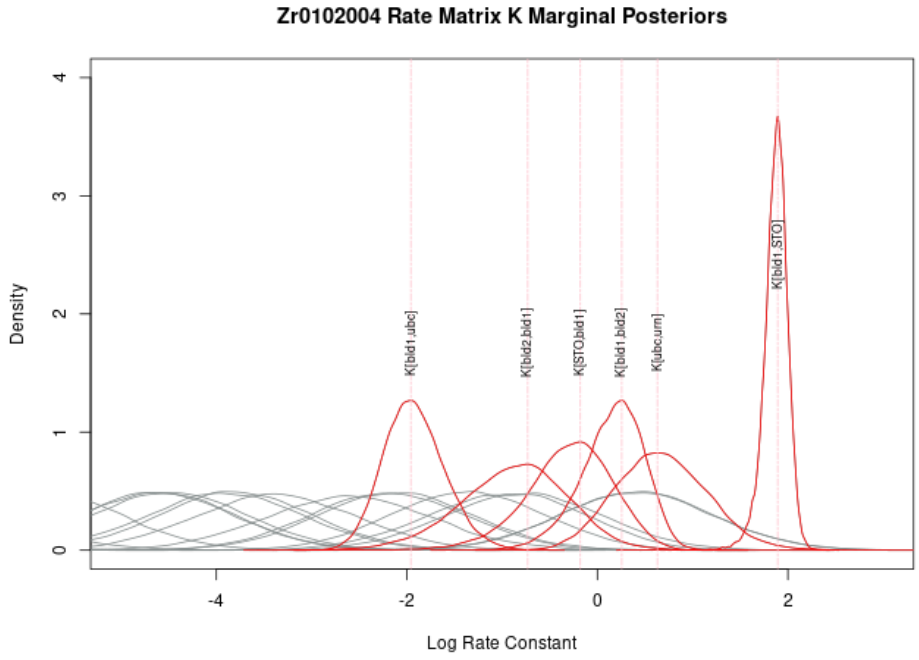

Figure I.9 Marginal posterior distributions for individual-specific rate matrix $\boldsymbol{k}$ for subject Zr0102004. Density curves for rate constants with distributions having a log standard deviation less than the one prescribed in the prior are in red and are identified by name.

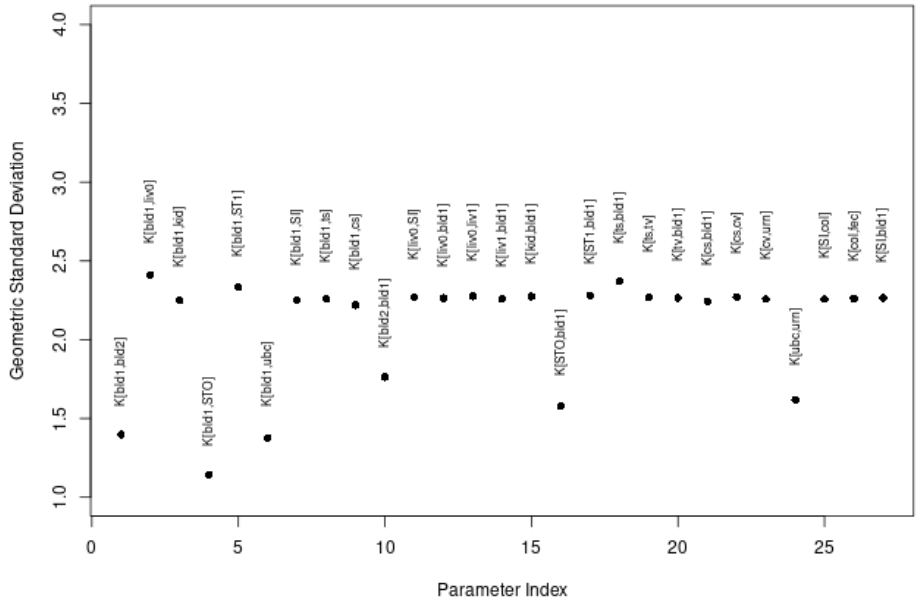

Figure I.10 Plot of geometric standard deviations for posterior distributions of rate constants $\boldsymbol{k}$ for subject Zr0102004.

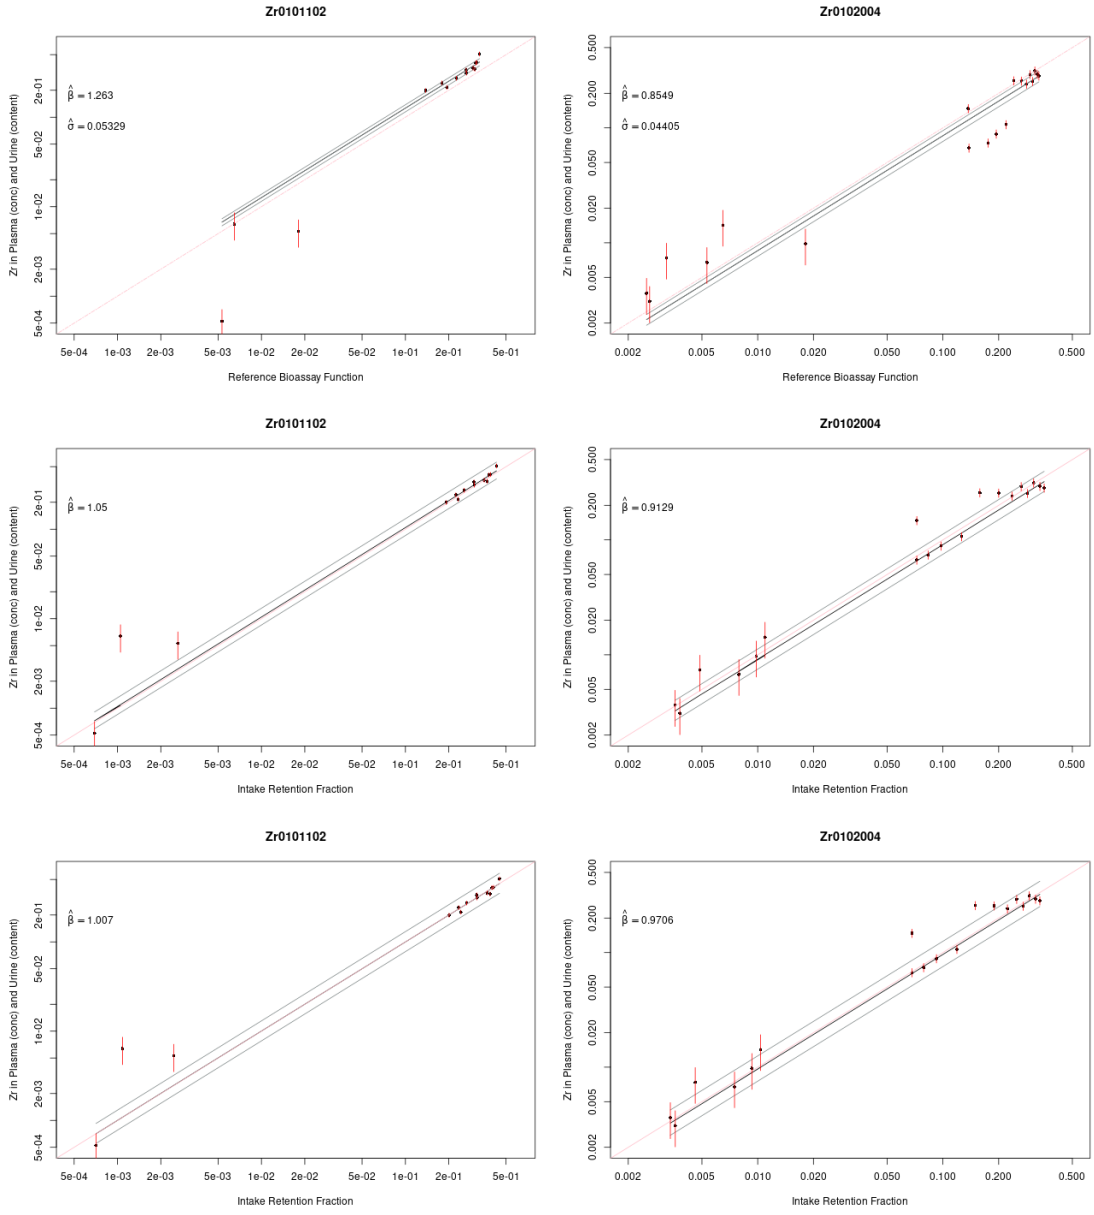

Figure I.11 Fits on a log-log scale to more clearly show the urine bioassay results: standard method top row, add-one-in Bayesian analysis middle row, and canned prior Bayesian analysis bottom row. The grey lines denote the 95% intervals for the intake. The error bars denote the 95% coverage interval on the measurements.

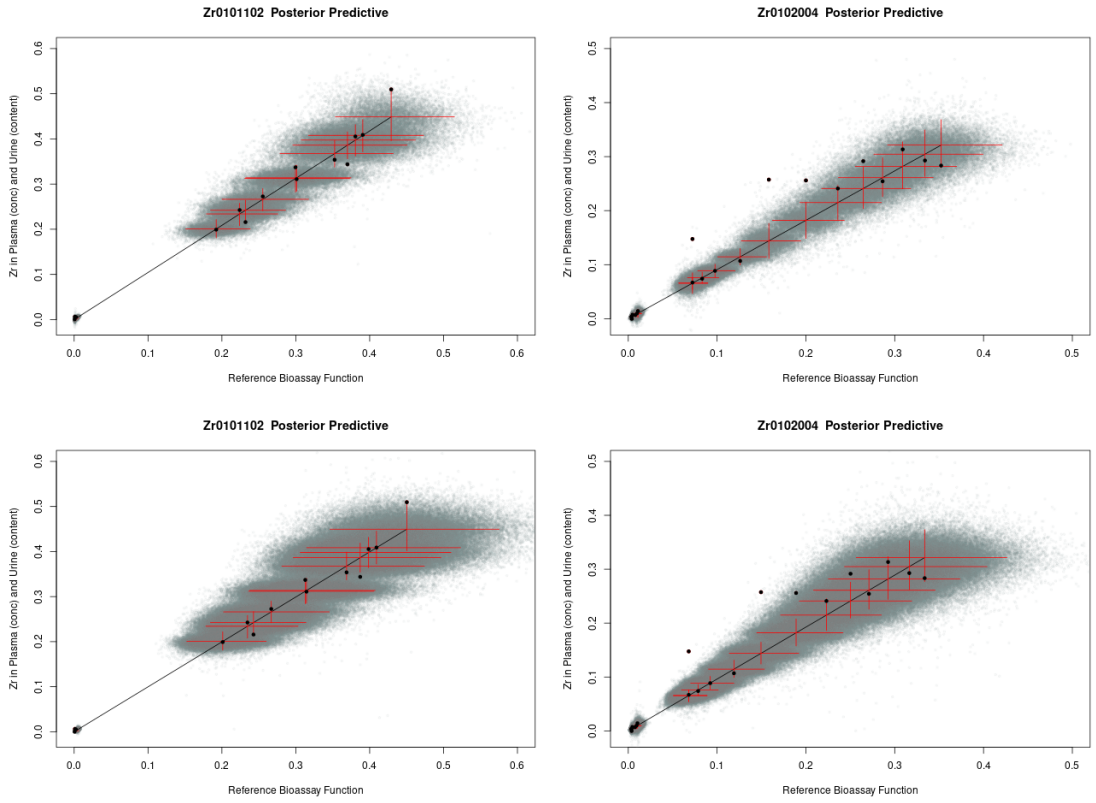

Figure I.12 Posterior predictive distributions from add-one-in analysis for Subject Zr0101102 (left top) and Zr0102004 (right top) and from MVN canned prior analysis for Subject Zr0101102 (left bottom) and Zr0102004 (right bottom). The error bars denote the 95% credible intervals on the posterior predictive distribution.

Appendix J

Urinary Bladder Model

The urinary bladder was assigned an explicit tissue tissue weighted factor in ICRP 1990 (Thorne, 1992), which prompted the ICRP to define a biokinetic model for the urinary bladder of an adult worker in ICRP Publication 67 (ICRP, 1993, pg 3). In this model the bladder has a fixed volume of $115\,mL$, which represents the average volume between voids. The number of voids per day is taken to be 6 for a worker, with $2 \times 115\,mL = 230\,mL$ per void. This gives $6 \times 230\,mL = 1380\,mL$ excreted per day, which matches the $1400\,mL$ for Reference Man in ICRP Publication 23 (ICRP, 1975). The first order rate constant for the bladder to urine pathway is therefore defined as

$$k_{ubc,urn} = \frac{1380\,mL/d}{115\,mL} = 12\,d^{-1}$$

as given in ICRP Publication 67, ICRP Publication 68 (ICRP, 1995, pg 13), and ICRP Publication 130 (ICRP, 2015, pg 85). In Chapter 3, a $LN\left(\log\left(12\right), \log\left(1.5\right)\right)$ prior was used for $k_{ubc,urn}$ whereas Schmidl et al. (2012) used a $T\left(6, 8, 24\right)$ and Li, Greiter, et al. (2011a) a $T\left(6, 12, 24\right)$. A lognormal was used in the dissertation to facilitate making a multivariate normal canned prior and because the authors of Stan advise us to avoid prior distributions like the triangular that have "kinks".