

A Hierarchical Bayesian Approach for Parameter Estimation in HIV Models

H. T. Banks*, Sarah Grove*, Shuhua Hu*, and Yanyuan Ma†

*Center for Research in Scientific Computation
North Carolina State University
Raleigh, North Carolina 27695-8205

†Department of Statistics
Texas A&M University
College Station, Texas 77843-3143

April 20, 2005

Abstract

A hierarchical Bayesian approach is developed to estimate parameters at both the individual and the population level in a HIV model, with the implementation carried out by Markov Chain Monte Carlo (MCMC) techniques. Sample numerical simulations and statistical results are provided to demonstrate the feasibility of this approach.

Key Words: HIV dynamics, parameter estimation, hierarchical model, MCMC, Gibbs sampling, Metropolis-Hasting algorithm.

1 Introduction

Human immunodeficiency virus (HIV) is a retrovirus that targets the CD4+ T-cells in the immune system. Once the virus has taken control of a sufficiently large proportion of CD4+ T-cells, an individual is said to have AIDS. While there is still no cure or vaccine for AIDS, significant progress has been made in the treatment of HIV infected patients. Currently the most prevalent treatment strategy for acutely HIV infected patients is highly anti-retroviral therapy (HAART) which uses two or more drugs. Typically these drug cocktails consist of one or more reverse transcriptase inhibitors (RTI) as well as a protease inhibitor (PI). It was reported in [22] that this multi-drug regimen can lead to the rapid decay of the viral load in a relatively short term. However, it was also reported in [17] that this rapid decay can sometimes be followed by a resurgence of virus within a few months. The resurgence may be caused by two main factors in the long-term use of drugs. One is poor patient adherence which is promoted by high drug cost and/or grave pharmaceutical side effects. The other is increasing drug resistance due to virus mutations. Hence, instead of adopting a continuous

combination therapy strategy, a number of researchers have suggested that one needs to seek a strategy that can decrease the negative side effects and medication burden placed on the patients and perhaps slow mutation rates while maintaining some control over HIV replication through an individual's immune response.

A set schedule for doing this involving on-off treatment segments is called a *structured treatment interruption* (STI). STI therapy has received a great deal of attention since a patient known as the "Berlin patient" in [19] was eventually successful in removing himself from medication permanently and still maintain HIV replication control after starting and stopping treatment several times. Since then a number of research efforts have focused on how to stimulate the immune system while maintaining control over HIV replication. One such study (see [27]) reported that a sustained HIV-specific Cytotoxic T Lymphocyte (CTL) response was increased during one, two, and three periods of treatment interruption. It was observed that control of the viremia in all test patients occurred after the CTL response was established. These examples document the importance of including the immune response in any model of HIV infection when studying STI. We should also note that realistically the HIV-specific immune response may be either strong or almost nonexistent as one considers different patients.

There have been a large number of mathematical models proposed by researchers at both the cellular level and the population level in the last decade (for example, see [1, 7, 16, 20, 22, 23, 24, 28, 29]) that attempt to understand the pathogenesis of HIV infection or similar biological processes in antiviral activities. While a number of these models have not included the immune response system in an effector role, recent efforts related to STI (e.g., [7, 28]) often include such a compartment (see [1] for discussions of such models). More importantly to our efforts here, even simple simulation studies with most of these models require some type of parameter estimation procedure since almost all models contain some parameters which can not be measured directly. In this paper we investigate an inverse problem methodology to estimate not only parameter values but also to provide some measures of the uncertainty (i.e., reliability or lack thereof) associated with the parameter values and their estimation. This will be done in the context of patients undergoing STI therapies as well as patients who are untreated (i.e., HIV progression without HAART). The model we chose is sufficiently complex so as to demonstrate applicability of the methodology for a wide class of HIV models.

In order to capture a number of the desired features noted above that are important in

STI therapies, the authors in [1] proposed and investigated the following model:

$$\begin{aligned}
\text{Type 1 target:} \quad & \dot{T}_1 = \lambda_1 - d_1 T_1 - (1 - \epsilon) k_1 V T_1 \\
\text{Type 2 target:} \quad & \dot{T}_2 = \lambda_2 - d_2 T_2 - (1 - f\epsilon) k_2 V T_2 \\
\text{Type 1 infected:} \quad & \dot{T}_1^* = (1 - \epsilon) k_1 V T_1 - \delta T_1^* - m_1 E T_1^* \\
\text{Type 2 infected:} \quad & \dot{T}_2^* = (1 - f\epsilon) k_2 V T_2 - \delta T_2^* - m_2 E T_2^* \\
\text{Free virions:} \quad & \dot{V} = N_T \delta (T_1^* + T_2^*) - cV - [(1 - \epsilon) \rho_1 k_1 T_1 + (1 - f\epsilon) \rho_2 k_2 T_2] V \\
\text{Immune effectors:} \quad & \dot{E} = \lambda_E + \frac{b_E (T_1^* + T_2^*)}{(T_1^* + T_2^*) + K_b} E - \frac{d_E (T_1^* + T_2^*)}{(T_1^* + T_2^*) + K_d} E - \delta_E E,
\end{aligned} \tag{1.1}$$

with specified initial values for T_1 , T_2 , T_1^* , T_2^* , V and E at time $t = t_0$. Here two co-circulating populations of target cells, potentially representing $CD4^+$ T-lymphocytes (T_1) and macrophages (T_2), are considered in the model. The parameter λ_i denotes the rate at which new T_i cells are created from sources within the body, d_i denotes the death rate of T_i cells, k_i is the infection rate of T_i cells infected by the virus, m_i represents the immune-induced clearance rate for T_i and δ is the death rate of the infected T cells. The drug efficacy parameter ϵ models a reverse transcriptase inhibitor (RTI) that blocks new infections and is potentially more effective in the Type 1 populations than in the Type 2 populations, where the efficacy is $f\epsilon$ with $f \in [0, 1]$. The parameters ρ_1 and ρ_2 represent the average number of virions infecting a Type 1 cell and a Type 2 cell, respectively. Moreover, λ_E denotes the immune effector production rate, δ_E denotes the natural death rate for immune effectors, b_E is the maximum birth rate for immune effectors, and d_E is the maximum death rate for immune effectors. Finally, K_b represents the saturation constant for immune effector birth, and K_d represents the saturation constant for immune effector death. The reader may consult [1] for a more detailed discussion of this model and its qualitative and quantitative features.

For our study of the proposed estimation methodology, we adopt model (1.1) as a reasonable approximation to the dynamics for any specific individual, but hypothesis that the dynamics may vary across subjects. In fact, clinical data clearly exhibits a great deal of variability among patients and their response to therapy, which may be attributed to individual mechanistic variability in both viral transcription and production, and in immune response. As illustrated in [2], if the infectivity rates k_1 and k_2 in (1.1) are changed, stable equilibria associated with the parameters may be changed from a dangerously high viral load related equilibrium point to one corresponding to levels of low viral load. Hence, the dynamics (and thus the response to therapy) for different individuals may differ because of the difference of several crucial parameters. With this in mind, an approach employing hierarchical nonlinear mixed-effects (NLME) models [9] appears to be reasonable for modelling HIV dynamics. While hierarchical NLME models have enjoyed wide spread use for some time with physiologically based pharmacokinetic (PBPK) studies (see, e.g., [9, 21]

and the references therein), use of such models in HIV pathogenesis studies is more recent [14, 15, 16, 24, 29, 30]. Early uses [14, 15, 30] involved simplification of the models to obtain closed form solutions for viral loads. With the recent progress in MCMC, it is possible to entertain much more complex mathematical models. In particular, the authors of [24] consider a system of nonlinear differential equations with constant drug efficacy but use only short-term viral load data to estimate parameters. In [30], the authors consider a system of nonlinear differential equations with variability in drug resistance and adherence in the presence of antiretroviral therapy. But to our knowledge, none of these previous efforts have included the immune response effector in the HIV model.

In this paper, we report on our investigation of how MCMC based methods will perform on estimation in complex models for HIV progression in untreated patients as well as in patients undergoing STI therapy. Our studies were designed to assist in efforts with clinical data sets which pose several specific challenges. In particular, one needs to estimate parameters from *limited amounts of data differing from patient to patient both in duration and sampling frequency*. Moreover, the typical HIV data set involves *partial observations, measurements from combined compartments, and heavily censored viral load measurements*. We treat for, to our knowledge, the first time such data sets for complex models. We will use a Bayesian approach to estimate the parameters at both individual and population levels in a framework of a hierarchical NLME model along the lines described in [9]. Our goal is to develop methodology to assist in understanding the distribution of the individual-specific parameters across the population of patients with the eventual aim of using the knowledge to predict the future patterns of the response at each individual level. Ideally, this will allow one to select different individual treatment interruption therapy regimes and to develop “best” performance strategies for the individual, where the performance is measured by some specific cost criterion (see [2] for related work).

The paper is organized as follows. In section 2, a hierarchical Bayesian approach is adopted to estimate the dynamic system parameters for the model (1.1), with the implementation carried out using MCMC techniques. In sections 3 and 4, some simulation studies are presented to assess the performance of and illustrate use of the proposed methodologies in the cases that there is no treatment and that one adopts an STI strategy for the patient treatment, respectively. Finally, in section 5, some conclusions and remarks are made.

Before proceeding with our presentation, we note that in our considerations here the methodology investigated is parametric in nature (i.e., we assume functional parameterized forms for various prior distributions) as opposed to the nonparametric approaches developed and discussed in [2, 3, 4, 5, 6].

2 A Hierarchical Bayesian Approach for Parameter Estimation

We formulate our estimation problems in the context of typical HIV clinical data sets such as those described in [1]. These data sets include individual longitudinal measurements for a population of patients with the number of observations varying from individual to individual. As mentioned above, the viral load measurements (V) are censored (above and below) due to assay limitations. Moreover, typical data sets usually include $CD4^+T$ cell counts (essentially T_1 or $T_1 + T_1^*$ in the model (1.1)) and $CD8$ cell counts as a rough indication of the immune effector E levels (see [1]).

Suppose that we have M patients and the number of measurements on the i th subject is n_i . Let $\tilde{\theta}_i \in \mathbb{R}^p$ denote the individual-specific parameters for subject i . We will use $\theta_i = \ln(\tilde{\theta}_i)$ as parameters in all subsequent discussions and computations. As noted, we usually have measurements of the viral load (V), the uninfected Type 1 cell (T_1) (or actually the sum ($T_1 + T_1^*$) of the uninfected and infected Type 1 cell), and the immune response (E) and hence we only consider the observations on these three components in our inverse problem calculations. For notational convenience, let $F_1(t_{ij}, \theta_i) = \log(V(t_{ij}, \theta_i))$, $F_2(t_{ij}, \theta_i) = \log(T_1(t_{ij}, \theta_i))$ (or in some cases $= \log(T_1(t_{ij}, \theta_i) + T_1^*(t_{ij}, \theta_i))$), $F_3(t_{ij}, \theta_i) = \log(E(t_{ij}, \theta_i))$, and $F(t_{ij}, \theta_i) = (F_1(t_{ij}, \theta_i), F_2(t_{ij}, \theta_i), F_3(t_{ij}, \theta_i))$, where $V(t_{ij}, \theta_i)$, $T_1(t_{ij}, \theta_i)$, $T_1^*(t_{ij}, \theta_i)$ and $E(t_{ij}, \theta_i)$ denote the numerical solutions of $V(t)$, $T_1(t)$, $T_1^*(t)$ and $E(t)$, respectively, in (1.1) for the i th subject at time t_{ij} . Let \tilde{Y}_{ij1} , \tilde{Y}_{ij2} and \tilde{Y}_{ij3} denote the observations for V , T_1 (or $T_1 + T_1^*$), and E , respectively, at time points t_{ij} for subject i , and let $\tilde{Y}_{ij} = (\tilde{Y}_{ij1}, \tilde{Y}_{ij2}, \tilde{Y}_{ij3})$. A reasonable assumption for the observations is that $Y_{ij} = F(t_{ij}, \theta_i) + \varepsilon_{ij}$, where $Y_{ij} = (\log(\tilde{Y}_{ij1}), \log(\tilde{Y}_{ij2}), \log(\tilde{Y}_{ij3}))$ with the three components assumed to be independent, and $\varepsilon_{ij} = (\varepsilon_{ij1}, \varepsilon_{ij2}, \varepsilon_{ij3})$ is assumed to be normal with mean zero and an unknown diagonal covariance matrix Σ_e (i.e., the \tilde{Y}_{ijk} are lognormal). However, in our clinical data, the measurements of the viral load data are censored, i.e., they are confined to a certain range $[\Upsilon_L, \Upsilon_R]$, because the assay can accurately detect only down to the Υ_L limit, and the same for the upper limit Υ_R . Therefore, instead of assuming a normal error for Y_{ij1} , one should assume a truncated normal error model for Y_{ij1} with a censor window $[\Upsilon_L, \Upsilon_R]$. Hence, the intra-individual error model summarizing the first stage of hierarchical model is given by

$$\begin{aligned} Y_{ij1} &\sim \mathcal{N}_{[\Upsilon_L, \Upsilon_R]}(F_1(t_{ij}, \theta_i), \Sigma_{e11}) \\ Y_{ij2} &\sim \mathcal{N}(F_2(t_{ij}, \theta_i), \Sigma_{e22}) \\ Y_{ij3} &\sim \mathcal{N}(F_3(t_{ij}, \theta_i), \Sigma_{e33}) \end{aligned} \tag{2.1}$$

where $\mathcal{N}_{[\Upsilon_L, \Upsilon_R]}$ denotes a truncated normal distribution with left limits Υ_L and right limits Υ_R , and \mathcal{N} denotes a normal distribution. We will use this notation throughout the paper.

Having noted the censored aspects of viral load data which motivates special distributional assumptions, we proceed to formulate a hierarchical Bayesian estimation problem [9] with standard distributional assumptions except where appropriate modifications are necessary.

After one establishes the individual level model, one seeks a population level understanding of HIV progression with and without therapy. The NLME Bayesian approach allows one to do this while also enhancing the parameter estimation procedure on each individual by borrowing information from other individuals in the population. Thus, one views $\theta_1, \dots, \theta_M$ as M independent realizations of a population level random vector that follows a certain probability distribution. In our case, we assume a truncated normal distribution $\mathcal{N}_{[\gamma_L, \gamma_R]}$ with unknown mean γ and unknown variance Σ . This is based on an assumption of prior information (physical and/or biological) about the upper γ_R and lower γ_L bounds on components of the rate parameters in the p -dimensional vector θ . Hence, the inter-individual variation model summarizing the second stage of the hierarchical model is assumed to be given by

$$\theta_i \sim \mathcal{N}_{[\gamma_L, \gamma_R]}(\gamma, \Sigma), \quad i = 1, 2, \dots, M. \quad (2.2)$$

Thus, under the current model assumptions, the parameters we are interested in estimating are θ_i ($i = 1, 2, \dots, M$), γ , Σ and Σ_e . In order to do this, a third stage for the hyperprior distribution is required at which prior distributions are specified for the parameters defining the second stage distributional form and the intra-individual variance parameters. We therefore assume:

- $De_{ii} \sim \mathcal{G}(\alpha_i, \beta_i)$, $i = 1, 2, 3$, where De_{ii} is the diagonal element of De and the off-diagonals of De are 0, and the hyperparameters α_i, β_i are known.
- $\gamma \sim \mathcal{N}_{[\gamma_L, \gamma_R]}(\eta, \Lambda^{-1})$, where γ_L and γ_R are known constants, and the hyperparameters η, Λ are also known.
- $D \sim \mathcal{W}_p(\nu, \Omega^{-1})$, where the hyperparameters ν and Ω are known.

Here D and De denote the corresponding precision matrices of Σ and Σ_e , respectively, i.e. $D = \Sigma^{-1}$ and $De = \Sigma_e^{-1}$. The mutually independent Gamma (\mathcal{G}), truncated Normal ($\mathcal{N}_{[\gamma_L, \gamma_R]}$) and Wishart (\mathcal{W}) prior distributions are chosen in accordance with standard practice of using conjugate priors (see, for example the discussions in [9], p.71–; p.219–).

Noting that Y_{ij1} may be left censored or right censored, we use an indicator parameter c_{ij} to denote whether or not a measurement Y_{ij1} is censored. Let $c_{ij} = 1$ denote a right censored measurement, $c_{ij} = -1$ denote a left censored measurement and $c_{ij} = 0$ denote an

uncensored measurement. Using these indicator parameters and recalling Bayes' Theorem, we assume the posterior distribution has the following form:

$$\begin{aligned}
& \pi(De, D, \gamma, \theta_1, \theta_2, \dots, \theta_M | \text{Data}) \\
& \propto \prod_{l=1}^3 g(De_{ll}; \alpha_l, \beta_l) \times w(D; \nu, \Omega^{-1}) \times \varphi_{[\gamma_L, \gamma_R]}(\gamma; \eta, \Lambda^{-1}) \\
& \quad \times \prod_{i=1}^M \varphi_{[\gamma_L, \gamma_R]}(\theta_i; \gamma, D^{-1}) \times \prod_{i=1}^M \prod_{j=1, c_{ij}=0}^{n_i} \varphi(Y_{ij}; F(t_{ij}, \zeta_i, \theta_i), De^{-1}) \\
& \quad \times \prod_{i=1}^M \prod_{j=1, c_{ij}=1}^{n_i} \int_{Y_{ij1} \geq \Upsilon_R} \varphi(Y_{ij}; F(t_{ij}, \zeta_i, \theta_i), De^{-1}) dY_{ij1} \\
& \quad \times \prod_{i=1}^M \prod_{j=1, c_{ij}=-1}^{n_i} \int_{Y_{ij1} \leq \Upsilon_L} \varphi(Y_{ij}; F(t_{ij}, \zeta_i, \theta_i), De^{-1}) dY_{ij1},
\end{aligned} \tag{2.3}$$

where ζ_i denotes the dosage mechanism for patient i , $g(De_{ll}; \alpha_l, \beta_l)$ represents the probability density function of a Gamma distribution with parameters α_l and β_l , $w(D; \nu, \Omega^{-1})$ represents the probability density function of a Wishart distribution with parameters ν and Ω^{-1} , $\varphi(Y_{ij}; F(t_{ij}, \zeta_i, \theta_i), De^{-1})$ represents the probability density function of a Normal distribution with mean $F(t_{ij}, \zeta_i, \theta_i)$ and variance matrix De^{-1} , $\varphi_{[\gamma_L, \gamma_R]}(Y_{ij}; F(t_{ij}, \zeta_i, \theta_i), De^{-1})$ represents the probability density function of a truncated Normal distribution with left limit γ_L and right limit γ_R and with mean $F(t_{ij}, \zeta_i, \theta_i)$ and variance matrix De^{-1} .

After we obtain the posterior distribution, we can make the prediction of the parameter θ_{M+1} for a new patient using the Bayesian posterior prediction in the following formula:

$$\begin{aligned}
\pi(\theta_{M+1} | \text{Data}) &= \int \pi(\theta_{M+1} | \gamma, D) \pi(\gamma, D | \text{Data}) d\gamma dD \\
&= \frac{1}{K} \sum_{i=1}^K \varphi_{[\gamma_L, \gamma_R]}(\theta_{M+1}; \gamma^i, D^i),
\end{aligned}$$

where $\{(\gamma^i, D^i) | i = 1, 2, \dots, K\}$ are K samples from the posterior distribution (2.3). We then use the estimated θ_{M+1} to predict a time course of the state variables V, T_1, \dots under estimation following any possible dosage mechanism.

When new data becomes available, we can update our estimate for the population level and individual level parameters by using the current posterior distribution as the prior and using the new data to form the likelihood function. Let $\{Y_{i, n_{i+1}}, \dots, Y_{i, \tilde{n}_i}\}$ denote the new

measurement for individual i , then the posterior distribution can be written as follows:

$$\begin{aligned}
& \pi(De, D, \gamma, \theta_1, \theta_2, \dots, \theta_M | \text{NewData}) \\
& \propto \pi(De, D, \gamma, \theta_1, \theta_2, \dots, \theta_M | \text{Data}) \times \prod_{i=1}^M \prod_{j=n_i+1, c_{ij}=0}^{\tilde{n}_i} \varphi(Y_{ij}; F(t_{ij}, \zeta_i, \theta_i), De^{-1}) \\
& \quad \times \prod_{i=1}^M \prod_{j=n_i+1, c_{ij}=1}^{\tilde{n}_i} \int_{Y_{ij} \geq \Upsilon_R} \varphi(Y_{ij}; F(t_{ij}, \zeta_i, \theta_i), De^{-1}) dY_{ij1} \\
& \quad \times \prod_{i=1}^M \prod_{j=n_i+1, c_{ij}=-1}^{\tilde{n}_i} \int_{Y_{ij1} \leq \Upsilon_L} \varphi(Y_{ij1}; F(t_{ij}, \zeta_i, \theta_i), De^{-1}) dY_{ij1} \\
& \propto \prod_{l=1}^3 g(De_{ll}; \alpha_l, \beta_l) \times w(D; \nu, \Omega^{-1}) \times \varphi_{[\Upsilon_L, \Upsilon_R]}(\gamma; \eta, \Lambda^{-1}) \\
& \quad \times \prod_{i=1}^M \varphi_{[\Upsilon_L, \Upsilon_R]}(\theta_i; \gamma, D^{-1}) \times \prod_{i=1}^M \prod_{j=1, c_{ij}=0}^{\tilde{n}_i} \varphi(Y_{ij}; F(t_{ij}, \zeta_i, \theta_i), De^{-1}) \\
& \quad \times \prod_{i=1}^M \prod_{j=1, c_{ij}=1}^{\tilde{n}_i} \int_{Y_{ij1} \geq \Upsilon_R} \varphi(Y_{ij}; F(t_{ij}, \zeta_i, \theta_i), De^{-1}) dY_{ij1} \\
& \quad \times \prod_{i=1}^M \prod_{j=1, c_{ij}=-1}^{\tilde{n}_i} \int_{Y_{ij1} \leq \Upsilon_L} \varphi(Y_{ij}; F(t_{ij}, \zeta_i, \theta_i), De^{-1}) dY_{ij1}.
\end{aligned}$$

We note that there is no closed form solution for the model (1.1), and this prevents an easy implementation of standard statistics software such as SAS, Splus or even the popular Bayesian software BUGS. Although the MCMC software MCSim can treat such situations, we are still unable to readily use it since we have a heavily censored data involving the measurements on the viral load. Current versions of MCSim do not incorporate censored data in its formulation. Therefore, we coded the algorithm independently. Since we do not have a closed form posterior distribution, we adopt the MCMC approach to approximate it. To do this, we first need to rewrite the terms in (2.3) contributed by the censored data as a truncated normal distribution using characteristic or indicator functions I to obtain

$$\int_{Y_{ij} \geq \Upsilon_R} \varphi(Y_{ij}; F(t_{ij}, \zeta_i, \theta_i), De^{-1}) dY_{ij1} = \int_{-\infty}^{\infty} \varphi(Y_{ij}; F(t_{ij}, \zeta_i, \theta_i), De^{-1}) I(Y_{ij1} \geq \Upsilon_R) dY_{ij1},$$

and

$$\int_{Y_{ij} \leq \Upsilon_L} \varphi(Y_{ij}; F(t_{ij}, \zeta_i, \theta_i), De^{-1}) dY_{ij1} = \int_{-\infty}^{\infty} \varphi(Y_{ij}; F(t_{ij}, \zeta_i, \theta_i), De^{-1}) I(Y_{ij1} \leq \Upsilon_L) dY_{ij1}.$$

We ignore the integration (hence treat the corresponding measurement Y_{ij1} as a parameter

of interest as well), and then the posterior distribution (2.3) can be rewritten as

$$\begin{aligned}
& \pi(De, D, \gamma, \theta_1, \theta_2, \dots, \theta_M, Y_{ij1} | \text{Data}) \\
& \propto \prod_{l=1}^3 g(De_l; \alpha_l, \beta_l) \times w(D; \nu, \Omega^{-1}) \times \varphi_{[\gamma_L, \gamma_R]}(\gamma; \eta, \Lambda^{-1}) \\
& \quad \times \prod_{i=1}^M \varphi_{[\gamma_L, \gamma_R]}(\theta_i; \gamma, D^{-1}) \times \prod_{i=1}^M \prod_{j=1, c_{ij}=0}^{n_i} \varphi(Y_{ij}; F(t_{ij}, \zeta_i, \theta_i), De^{-1}) \\
& \quad \times \prod_{i=1}^M \prod_{j=1, c_{ij}=1}^{n_i} \varphi(Y_{ij}; F(t_{ij}, \zeta_i, \theta_i), De^{-1}) I(Y_{ij1} \geq \Upsilon_R) \\
& \quad \times \prod_{i=1}^M \prod_{j=1, c_{ij}=-1}^{n_i} \varphi(Y_{ij}; F(t_{ij}, \zeta_i, \theta_i), De^{-1}) I(Y_{ij1} \leq \Upsilon_L),
\end{aligned} \tag{2.4}$$

where only the Y_{ij1} 's whose corresponding $c_{ij} \neq 0$ are treated as parameters and hence appear on the left side of (2.4). We also can argue that the conditional posterior distribution for the censored Y_{ij1} still follows the truncated normal distribution. By collecting all the terms with De_l ($l = 1, 2, 3$) and making some routine calculations, we find that the conditional posterior distribution of De_l is a Gamma distribution:

$$De_l \sim \mathcal{G} \left(\frac{\sum_{i=1}^M n_i}{2} + \alpha_l, \left(\frac{\sum_{i=1}^M \sum_{j=1}^{n_i} (Y_{ijl} - F_l(t_{ij}, \zeta_i, \theta_i))^2}{2} + \frac{1}{\beta_l} \right)^{-1} \right).$$

Similarly, we obtain that the full conditional posterior distribution of γ is a truncated Normal distribution:

$$\gamma \sim \mathcal{N}_{[\gamma_L, \gamma_R]} \left((\Lambda + MD)^{-1} (\Lambda \eta + D \sum_{i=1}^M \theta_i), (\Lambda + MD)^{-1} \right),$$

and that the full conditional posterior distribution of D is a Wishart distribution:

$$D \sim \mathcal{W}_p \left(M + \nu, \left(\Omega + \sum_{i=1}^M (\theta_i - \gamma)(\theta_i - \gamma)^T \right)^{-1} \right).$$

The full conditional distribution of θ_i , given the remaining parameters and the data in (2.4), can not be written explicitly. To generate a sample for each θ_i , we again follow standard practice and choose a Metropolis-Hastings algorithm with the proposal distribution from a random walk chain.

The resulting MCMC algorithm is outlined as follows:

Step 1: Select the values for hyperparameters α_l, β_l for $l = 1, 2, 3$, η, Λ, ν and Ω . Pick the candidate variance Δ for θ_i . Select the starting values for $\theta_1^0, \theta_2^0, \dots, \theta_M^0, \gamma^0, D^0$ (Σ^0), De^0 (Σ_e^0), and Y_{ij}^0 if $c_{ij} \neq 0$. Set $k = 1$.

Step 2: Calculate $F(t_{ij}, \zeta_i, \theta_i^{k-1})$, $j = 1, 2, \dots, n_i$, and $i = 1, 2, \dots, M$.

Step 3: Generate a sample De_{ll}^k ($l = 1, 2, 3$) from the following distribution:

$$\mathcal{G} \left(\frac{\sum_{i=1}^M n_i}{2} + \alpha_l, \left(\frac{\sum_{i=1}^M \sum_{j=1}^{n_i} (Y_{ijl}^{k-1} - F_l(t_{ij}, \zeta_i, \theta_i^{k-1}))^2}{2} + \frac{1}{\beta_l} \right)^{-1} \right).$$

Step 4: Generate a sample γ^k from the following distribution by Gibbs sampling(see [9], p.222–) method:

$$\mathcal{N}_{[\gamma_L, \gamma_R]} \left((\Lambda + MD^{k-1})^{-1} (\Lambda \eta + D^{k-1} \sum_{i=1}^M \theta_i^{k-1}), \quad (\Lambda + MD^{k-1})^{-1} \right).$$

Step 5: Generate a sample D^k from $\mathcal{W}_p \left(M + \nu, \left(\Omega + \sum_{i=1}^M (\theta_i^{k-1} - \gamma^k)(\theta_i^{k-1} - \gamma^k)^T \right)^{-1} \right)$.

Step 6: Generate a sample Y_{ij1}^k if $c_{ij} \neq 0$.

Generate it from the first component of $\mathcal{N}(F(t_{ij}, \zeta_i, \theta_i^{k-1}), \Sigma_e^{k-1})I(Y_{ij1} \leq \Upsilon_L)$ if $c_{ij} = -1$.

Generate it from the first component of $\mathcal{N}(F(t_{ij}, \zeta_i, \theta_i^{k-1}), \Sigma_e^{k-1})I(Y_{ij1} \geq \Upsilon_R)$ if $c_{ij} = 1$.

Step 7: Generate a sample $\hat{\theta}_i$ from $\mathcal{N}(\theta_i^{k-1}, \Delta)$.

$$\varrho = \frac{\varphi(\hat{\theta}_i; \gamma^k, \Sigma^k) \prod_{j=1}^{n_i} \varphi_{[\gamma_L, \gamma_R]}(Y_{ij}; F(t_{ij}, \zeta_i, \hat{\theta}_i), \Sigma_e^k)}{\varphi(\theta_i^{k-1}; \gamma^k, \Sigma^k) \prod_{j=1}^{n_i} \varphi_{[\gamma_L, \gamma_R]}(Y_{ij}; F(t_{ij}, \zeta_i, \theta_i^{k-1}), \Sigma_e^k)}.$$

Generate a sample u from a uniform distribution $\mathcal{U}(0, 1)$. Set

$$\theta_i^k = \begin{cases} \theta_i^{k-1} & \text{if } \varrho < u \\ \hat{\theta}_i & \text{if } \varrho \geq u. \end{cases}$$

Step 8: Set $k = k + 1$ and return to Step 2 until convergence is reached.

To test the algorithm presented above, it is necessary to draw samples from several different probability distributions. The random number samples required in the algorithm drawn from the Gamma distribution and Wishart distribution can be generated by the random generators *gamrnd* and *wishrnd* in MATLAB, respectively. However, there is no random number generator in MATLAB that can generate the random number samples from the truncated multivariate normal distributions. Instead, there are two common methods to draw samples from a truncated norm distribution (see [13]). One is the rejection method,

which is straightforward and easy to implement but tends to be wasteful and very slow (see [12, 13]). The other method is Gibbs sampling which exploits the fact that if a random variable follows a truncated multivariate normal distribution, then the distribution of each element of this random variable conditional on all of its other elements is a truncated univariate normal. It was shown in [11] that the rate of convergence is geometric in the L^1 norm. Hence, we adopt the Gibbs sampling in Step 4 to generate a sample γ , and the random number sample from the truncated univariate normal distribution is generated by the inverse cumulative distribution function method introduced in [10].

Note that, in our implementation, the proposal distribution for $\hat{\theta}_i$ is chosen to be a multivariate normal distribution centered at the current value of θ_i , as it can be easily sampled and is symmetric. A critical issue here is how to choose the value for the proposal variance. As illustrated in [8, 26], if the proposal variance is too large, the moves are large but are often unaccepted. This may lead to high autocorrelation and poor mixing, and may require a much longer chain. If the proposal variance is too small, the moves are generally accepted but will move around the parameter space slowly which again leads to high autocorrelation and poor mixing. Hence, we need to adjust the value of the proposal variance to stabilize the posterior distribution faster.

In our implementation, we use trace plots to monitor the convergence because it is a simple and effective tool. Note that the initial simulations usually do not come from the stationary distribution targeted by the Markov chain. We need to discard a number of initial burn-in simulations. So we prefer to use the trace plot of a single long-run sequence in all of our simulations below. In order to test the mixing behavior of a Markov chain, we use autocorrelation plots.

We tested the proposed algorithm with simulated data for “patients” undergoing no therapy (Section 3) and “patients” undergoing STI therapy (Section 4). Simulated data was prepared in both cases in a similar manner. For all the simulations used in this paper, the initial values of T_1 , T_2 , T_1^* , T_2^* , V and E are chosen as 10^6 , 3198, 10^{-4} , 10^{-4} , 1, 10, respectively. In order to solve (1.1), we also need to specify all the parameter values. The default values of these parameters are given as (see [1, 2] for motivation for some of the values):

$$\lambda_1 = 10000, \lambda_2 = 31.98, f = 0.34, \delta = 0.7, N_T = 100, b_E = 0.3, K_b = 100, d_E = 0.25, \\ K_d = 500, \delta_E = 0.1, \lambda_E = 1, d_1 = d_2 = 0.01, m_1 = m_2 = 10^{-5}, \rho_1 = \rho_2 = 1, k_2 = 10^{-4}, \\ k_1 = 8 \times 10^{-7}, c = 13.$$

When some of the parameters are to be estimated, we treat them as random effects, i.e., the parameter values vary across the population, and we will specify how to generate these values in its corresponding simulation. All the other parameters are treated as fixed effects,

i.e., all the patients take the same values as the default values listed above.

In the case of simulated data for patients undergoing therapy, the dosage mechanism ζ_i is generated computationally as follows: we divide the patients into 5 groups, and we assume that patients in the same group have the same treatment schedules. The number of treatment cycles for group 1, group 2, group 3, group 4 and group 5 are 1, 2, 3, 4 and 5, respectively. In order to clearly state how we generate the dosage schedule, we explain group 2 as an example. We first generate 4 numbers uniformly from day 0 to day 1460 (4 years range). For example, [7, 428, 558, 1147] are the 4 numbers we generated, then we add 2 to each number to get another 4 numbers. Thus we have 8 numbers [7, 9, 428, 430, 558, 560, 1147, 1149], and then the time schedule is chosen as follows:

$$\begin{bmatrix} 0 & 7 & 9 & 428 & 430 & 558 & 560 & 1147 & 1149 & \text{last day} \\ 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 \end{bmatrix},$$

where 0 denotes treatment off and 1 denotes treatment on with drug efficacy 0.8, and we always assume that the treatment in day 0 and last day is 0. If the day is between the above two indicated days, then we use linear interpolation to calculate the drug efficacy. For example, in day 8, the drug efficacy is $\varepsilon = 0.8 \left(\frac{1-0}{9-7}(8-7) + 0 \right) = 0.4$. Hence, for group 2, there is no treatment in the days between 0 and 7, there is a treatment with drug efficacy 0.4 in day 8, there is a treatment with drug efficacy 0.8 in the days between 9 and 428, there is a treatment with drug efficacy 0.4 in day 429, there is no treatment in the days between 430 and 558, and then we go to next treatment cycle with the same process.

The data that we used is generated computationally as follows:

- Use the ODE solver to obtain the values for $F(t_{ij}, \zeta_i, \theta_i)$. In the results reported here, we used a custom-compiled Fortran 77 code for a Backward Difference Formula (BDF) method (Gear type method) to solve the stiff system (1.1). This was done by employing the Livermore Solver for Ordinary Differential Equations (LSODE) [25] in double precision. In this routine, the Jacobian matrix is treated either as a dense or a banded matrix and is supplied by the user. The error tolerances for the data generation ODE solver in Section 3 and Section 4 are $\{1e-9, 1e-24\}$ and $\{1e-3, 1e-3\}$, respectively.
- Generate the data $Y_{ij} = F(t_{ij}, \zeta_i, \theta_i) + \varepsilon_{ij}$, where ε_{ij} is a random sample generated from a multivariate norm random number generator *mvnrnd* in MATLAB with mean zero and some diagonal covariance matrix Σ_e .
- If $Y_{ij1} < \Upsilon_L$, then we set $c_{ij} = -1$. In this case, we need to update its value. The starting value for Y_{ij1} is chosen as Υ_L , and its subsequent value is generated as the same process as in MCMC algorithm Step 6.

If $Y_{ij1} > \Upsilon_R$, then we set $c_{ij} = 1$. In this case, we also need to update its value. The starting value for Y_{ij1} is chosen as Υ_R , and its subsequent value is generated with the same process as in MCMC algorithm Step 6.

3 Simulation Studies I: Subjects Without Therapy

In this section, we apply the MCMC method that we developed in section 2 to estimate parameters of interest in the case that without treatment, i.e. $\epsilon = 0$. We assume that Y_{ij1} is confined to be in the range $[\log(50), \log(750000)]$, and the measurements are taken every five days from day 0 to day 200. The data that we use in this section was generated computationally with variance $\Sigma_e = 0.04I_3$, where I_3 is the identity matrix with dimension 3.

Trace plots, histograms and autocorrelation plots are presented for the estimated parameters γ , an arbitrarily chosen patient θ_i , D and De in each simulation. We also present in tables their corresponding statistical inference: mean, mode, 95% Highest Density Region (HDR).

3.1 Simulation: k_1

In this subsection, we assume that all the parameters in (1.1) are fixed and known except for k_1 . We want to test the following three effects on the parameter estimation procedure:

1. The number of patients;
2. The type of the data $((V, T_1, E)$ vs. $(V, T_1 + T_1^*, E)$);
3. The differences in tolerances for the ODE solver when solving the inverse problem and generating the data.

The first comparison will yield information about the quality of the population level estimation procedure as a function of the number of patients observed. The second will reveal the effects of combined compartments data on the estimates. The third will shed light on possible degradation of the process when the data is generated with approximations of the model different than those used in the inverse problem calculations.

We present results for the following five situations:

- $M = 150$, and the observed three components are V , T_1 and E with the same tolerances;
- $M = 100$, and the observed three components are V , T_1 and E with the same tolerances;

- $M = 150$, and the observed three components are V , $T_1 + T_1^*$ and E with the same tolerances;
- $M = 100$, and the observed three components are V , $T_1 + T_1^*$ and E with the same tolerances;
- $M = 150$, and the observed three components are V , T_1 and E with different tolerances.

The values for γ_L and γ_R are chosen to be $\ln(10^{-16})$ and $\ln(0.1)$, respectively. The true parameter values of θ_i are produced from a truncated normal distribution with mean $\ln(8 \times 10^{-7})$ and standard deviation 0.1. The proposed variance is $\Delta = 5e - 5$. The value of the hyperparameters are chosen as follows:

$$\alpha_l = 2.5, \beta_l = 10 \text{ for } l = 1, 2, 3, \quad \eta = \ln(8 \times 10^{-7}), \Lambda = 0.1, \quad \nu = 1, \Omega = 0.01.$$

We ran one long chain with length 10000 for each simulation in this subsection.

3.1.1 $M = 150$, and the observed three components are V , T_1 and E with the same tolerances

After discarding the first 10 samples, we retain every third sample from the remaining samples. The trace plots and histograms for γ , θ , and D are presented in Figure 1 and those for De are presented in Figure 2. From these figures, we see that all chains are stabilized.

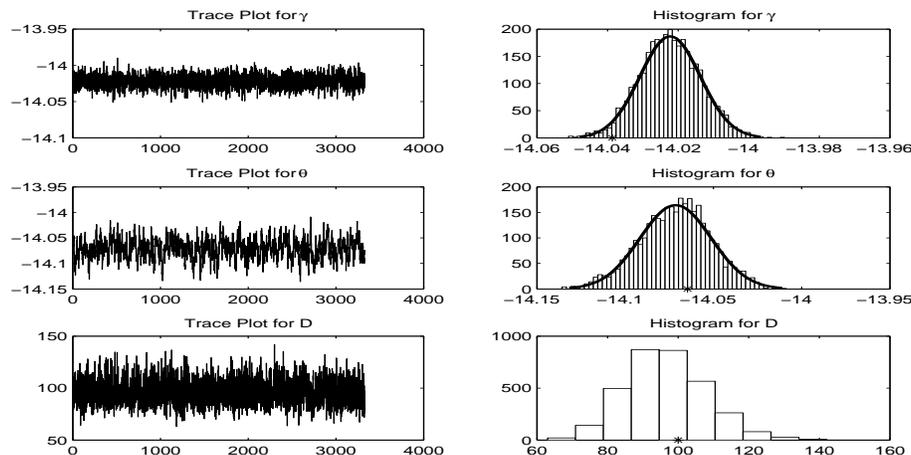


Figure 1: Trace plots and histograms for γ , θ and D . The point $*$ in the histogram plots denotes the true value for the corresponding parameter.

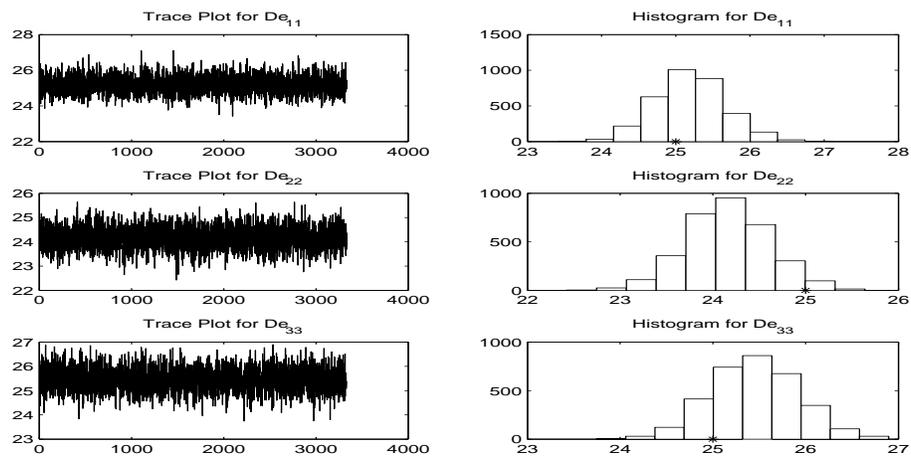


Figure 2: Trace plots and histograms of De .

In order to test the mixing behavior of the chain, autocorrelation plots for γ , θ and D are presented in Figure 3, and those for De are presented in Figure 4. From these graphs, we see that each chain is well-mixed.

The statistical results are presented in Table 1. From the results of this table, we see that we obtain very good estimates for all the parameters of interest.

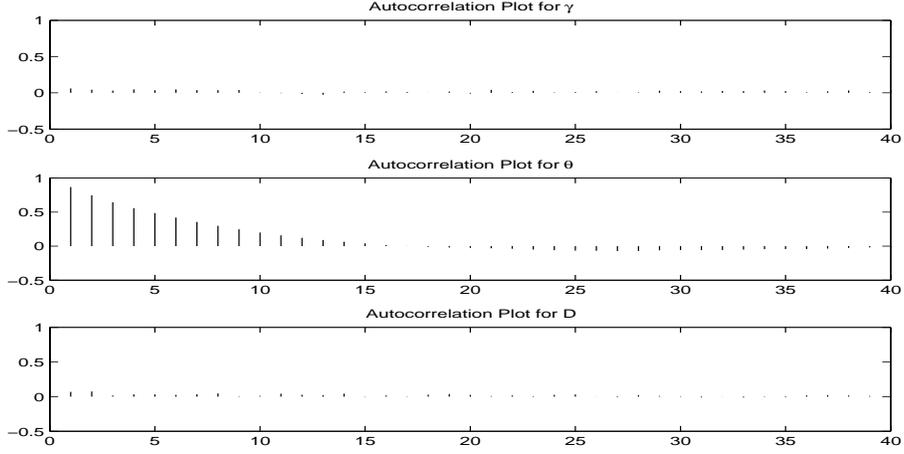


Figure 3: Autocorrelation plots for γ , θ and D , respectively.

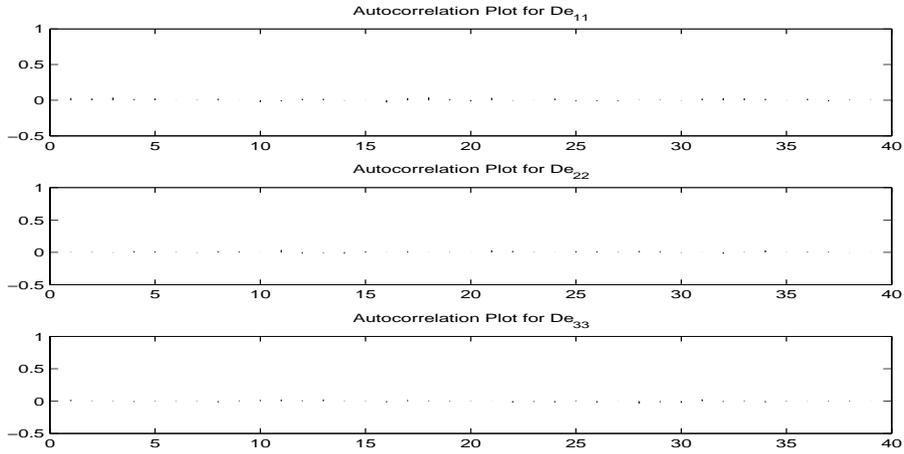


Figure 4: Autocorrelation plots for elements of De .

	true	mean	mode	95% HDR
$\ln(k_1)$	-14.0387	-14.0223	-14.0235	$[-14.0387, -14.0058]$
θ_i	-14.0647	-14.0714	-14.0660	$[-14.1136, -14.0328]$
D	100.0000	96.4980	90.6896	$[76.0024, 120.9687]$
De_{11}	25.0000	25.1942	25.0824	$[24.3024, 26.1647]$
De_{22}	25.0000	24.1635	24.1926	$[23.2934, 25.0461]$
De_{33}	25.0000	25.4498	25.4783	$[24.5703, 26.3806]$

Table 1: Statistical results for the estimated parameters when there are 150 patients and the observed three components are V , T_1 and E .

3.1.2 $M = 100$, and the observed three components are V , T_1 and E with the same tolerances

After discarding the first 10 samples, we again retain every third sample from the remaining samples. The trace plots and histograms γ , θ , and D are presented in Figure 5 and those for De are presented in Figure 6. From these figures, we see that all chains are stabilized.

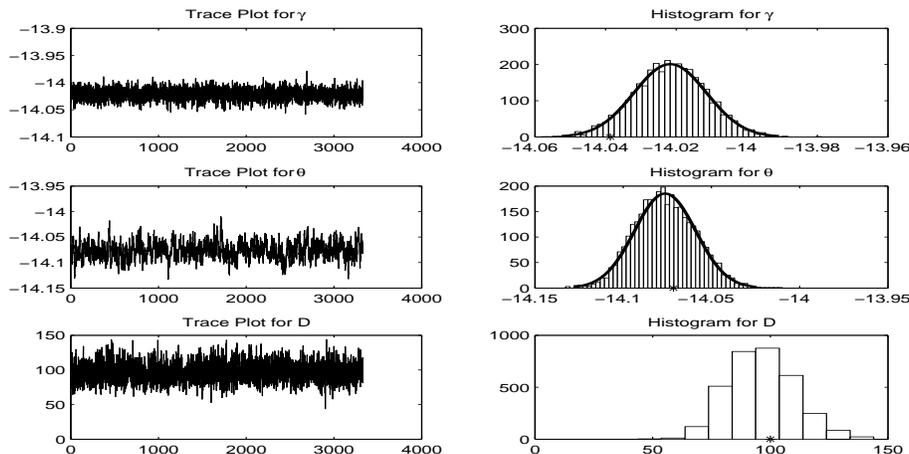


Figure 5: Trace plots and histograms for γ , θ and D .

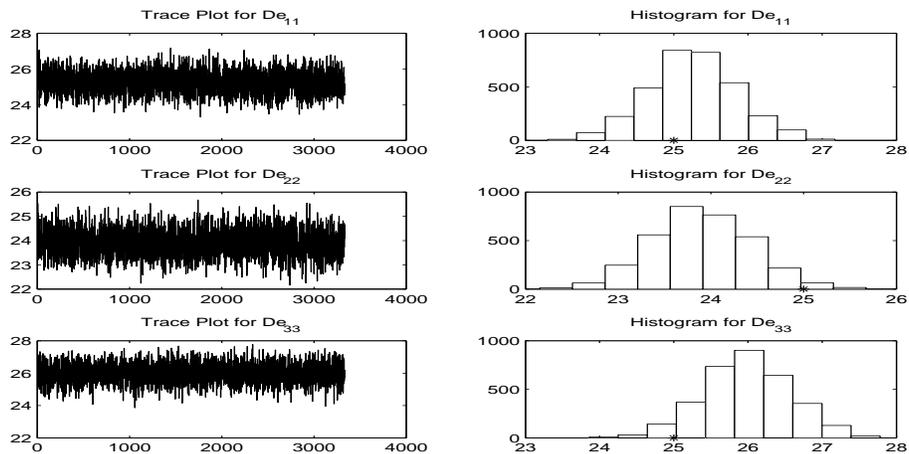


Figure 6: Trace plots and histograms of De .

In order to test the mixing behavior of the chain, autocorrelation plots for γ , θ and D are presented in Figure 7, and those for De are presented in Figure 8. From these graphs, we see that again each chain is well-mixed.

The statistical results for the estimated parameters are presented in Table 2. Comparing the results of this table with those in Table 1, we see that a larger number of patients can lead to perceptibly better estimates, which is expected since a larger number of patients should provide more information than a smaller number of patients.

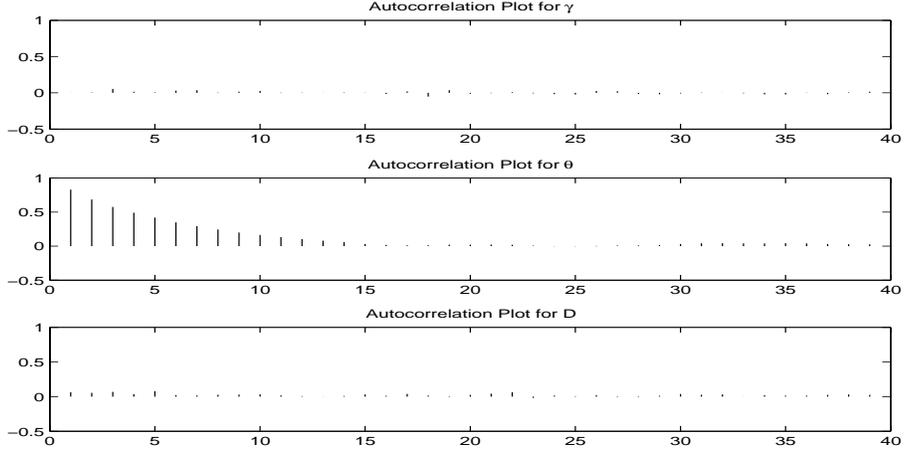


Figure 7: Autocorrelation plots for γ , θ and D , respectively.

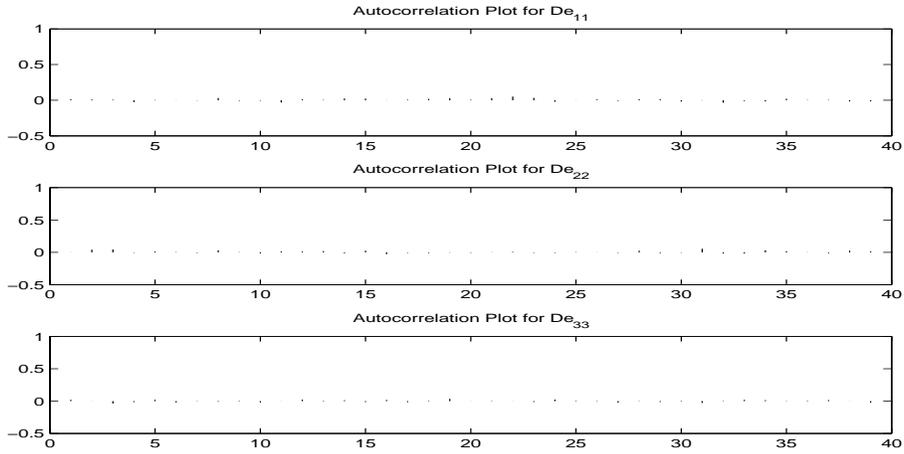


Figure 8: Autocorrelation plots for elements of De .

	true	mean	mode	95% HDR
$\ln(k_1)$	-14.0387	-14.0217	-14.0225	$[-14.0423 \quad -14.0018]$
θ	-14.0716	-14.0763	-14.0774	$[-14.1081 \quad -14.0424]$
D	100.0000	96.3566	98.8505	$[71.1895, 125.3455]$
De_{11}	25.0000	25.2566	25.0440	$[24.0865, 26.4564]$
De_{22}	25.0000	23.9016	23.7391	$[22.8963, 24.9544]$
De_{33}	25.0000	25.9859	26.0166	$[24.8209, 27.1394]$

Table 2: Statistical results for the estimated parameters when there are 100 patients and the observed three components are V , T_1 and E .

3.1.3 $M = 150$, and the observed three components are V , $T_1 + T_1^*$ and E with the same tolerances

After discarding the first 10 samples, we retain every third sample from the remaining samples. The trace plots and histograms γ , θ , and D are presented in Figure 9 and those for De are presented in Figures 10. From these figures, we see that all chains are stabilized.

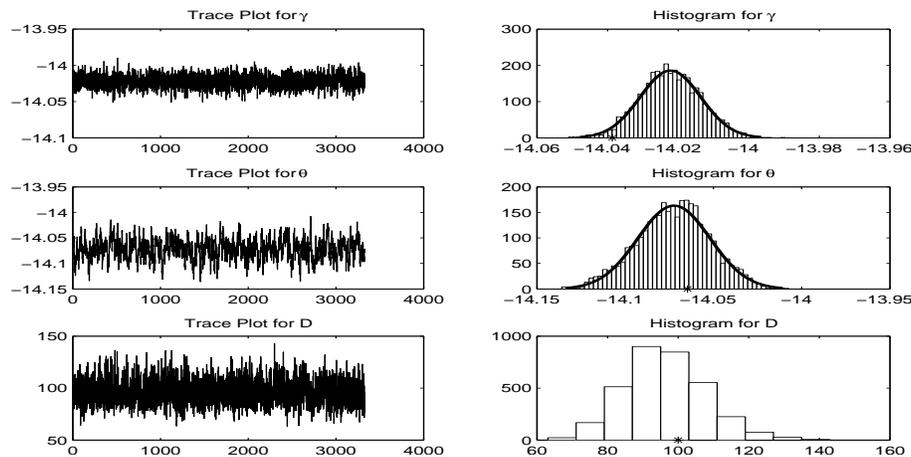


Figure 9: Trace plots and histograms for γ , θ and D .

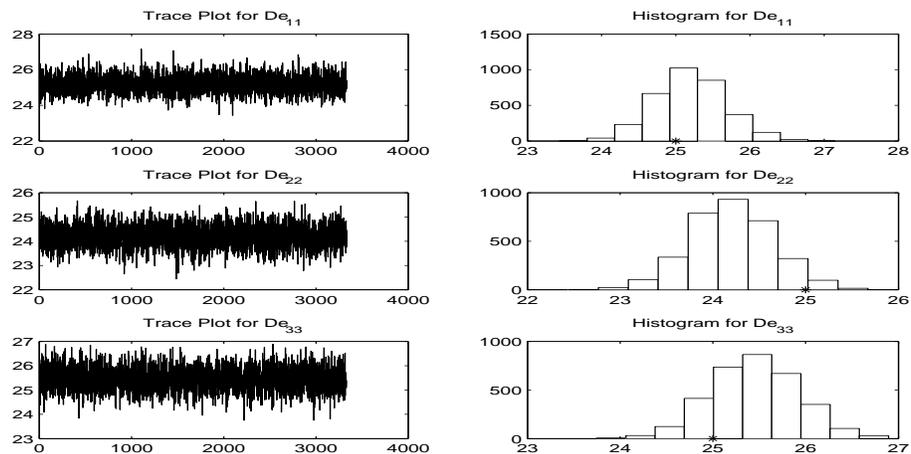


Figure 10: Trace plots and histograms of De .

In order to test the mixing behavior of the chain, autocorrelation plots for γ , θ and D are presented in Figure 11 and those for De are presented in Figure 12. From these graphs, we see that each chain is well-mixed.

The statistical results for the estimated parameters in Section 3.1.3 are presented in Table 3. Comparing the results here with those in Table 1, we do not see much difference.

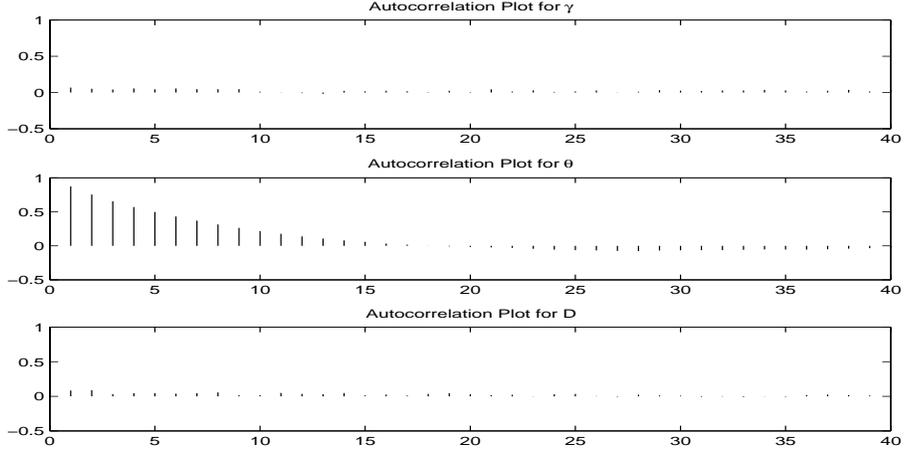


Figure 11: Autocorrelation plots for γ , θ and D , respectively.

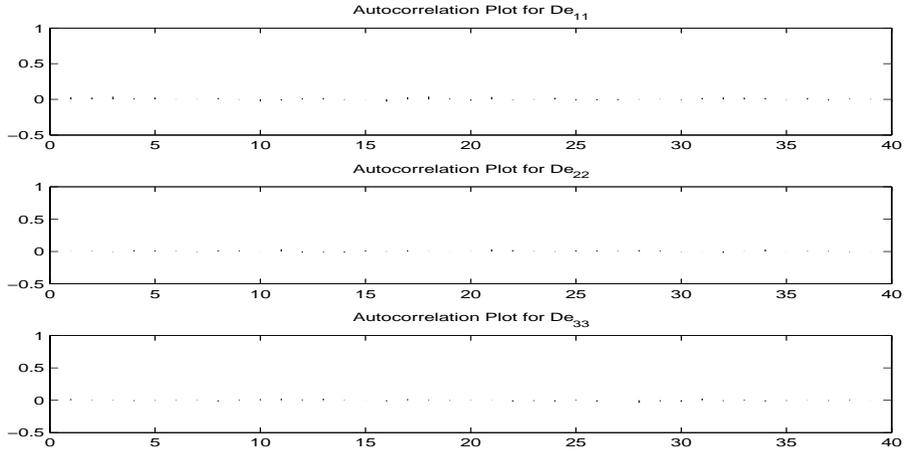


Figure 12: Autocorrelation plots for elements of De .

	true	mean	mode	95% HDR
$\ln(k_1)$	-14.0387	-14.0222	-14.0234	$[-14.0386, -14.0054]$
θ	-14.0647	-14.0724	-14.0652	$[-14.1139, -14.0330]$
D	100.0000	96.3046	91.0545	$[75.7526, 120.6870]$
De_{11}	25.0000	25.1940	25.1074	$[24.3063, 26.1580]$
De_{22}	25.0000	24.2023	24.2183	$[23.3256, 25.0855]$
De_{33}	25.0000	25.4489	25.4783	$[24.5670, 26.3778]$

Table 3: Statistical results for the estimated parameters when there are 150 patients and the observed three components are V , $T_1 + T_1^*$ and E .

3.1.4 $M = 100$, and the observed three components are V , $T_1 + T_1^*$ and E with the same tolerances

After discarding the first 10 samples, we retain every third sample from the remaining samples. The trace plots and histograms γ , θ , and D are presented in Figure 13 and those for De are presented in Figures 14. From these figures, we can see that all chains are stabilized.

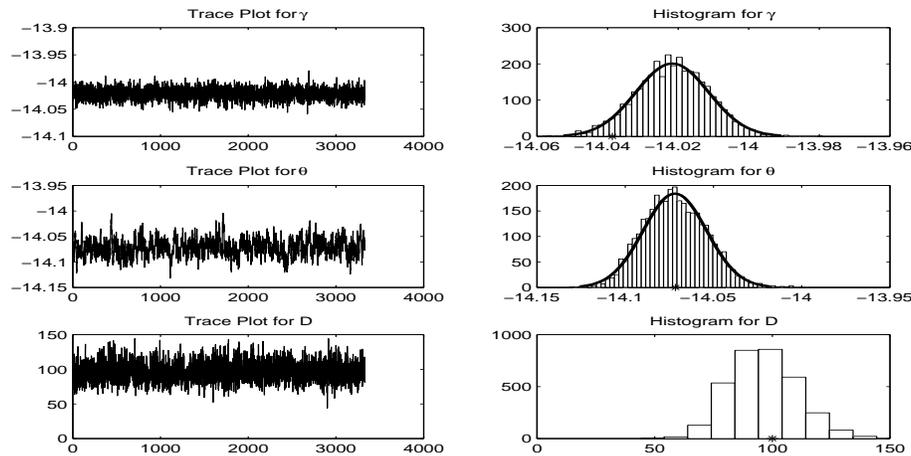


Figure 13: Trace plots and histograms for γ , θ and D .

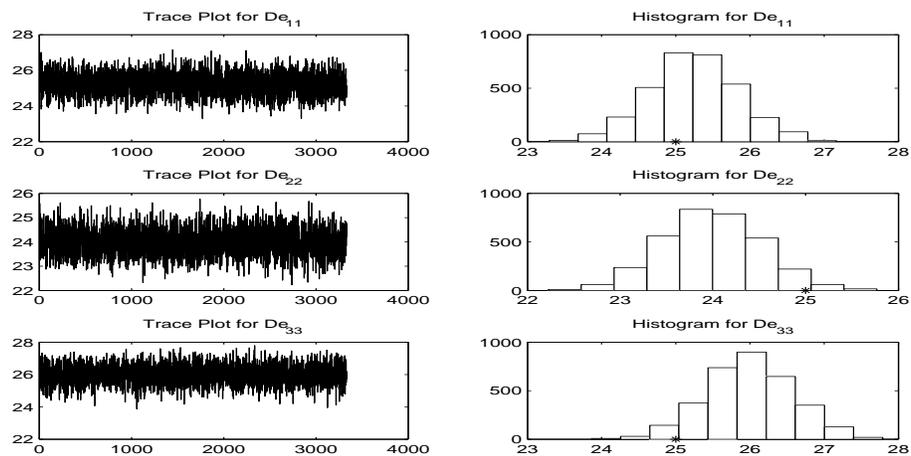


Figure 14: Trace plots and histograms of De .

In order to test the mixing behavior of the chain, autocorrelation plots for γ , θ and D are presented in Figure 15 and those for De are presented in Figure 16. From these graphs, we again see that each chain is well-mixed.

The statistical results for the estimated parameters in section 3.1.4 are presented in Table 4. Comparing the results of this table with those in Table 3, we see that this further illustrates the fact that a larger number of patients can lead to substantially better estimates.

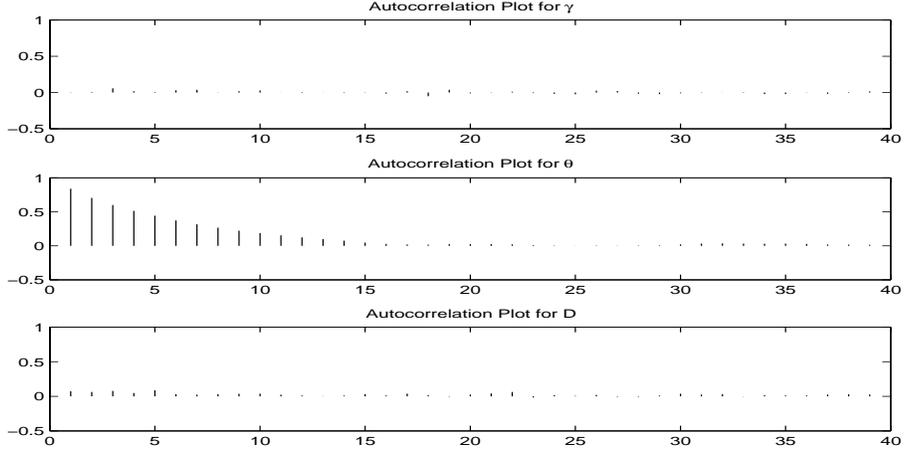


Figure 15: Autocorrelation plots for γ , θ and D , respectively.

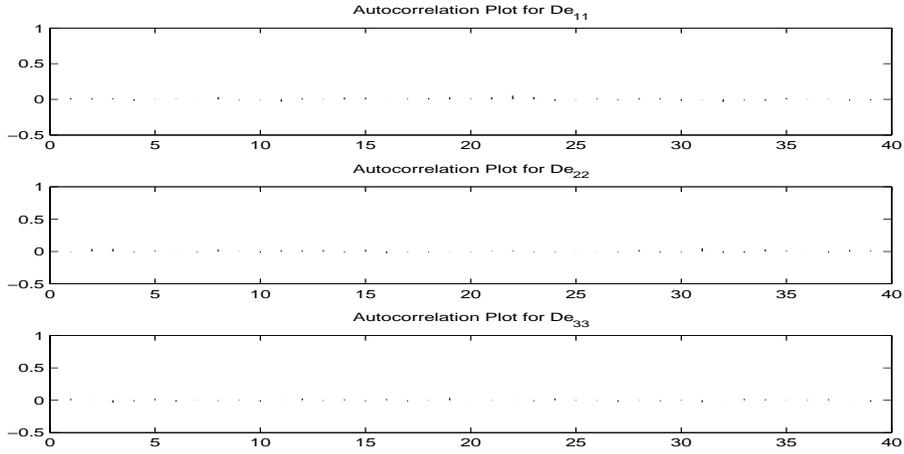


Figure 16: Autocorrelation plots for elements of De .

	true	mean	mode	95% HDR
$\ln(k_1)$	-14.0387	-14.0217	-14.0228	$[-14.0422, -14.0019]$
θ	-14.0716	-14.0718	-14.0745	$[-14.1042, -14.0361]$
D	100.0000	96.3266	99.1819	$[71.1410, 125.9427]$
De_{11}	25.0000	25.2379	25.0383	$[24.0615, 26.4394]$
De_{22}	25.0000	23.9891	23.8198	$[22.9683, 25.0332]$
De_{33}	25.0000	25.9892	26.0233	$[24.8265, 27.1449]$

Table 4: Statistical results for the estimated parameters when there are 100 patients and the observed three components are V , $T_1 + T_1^*$ and E .

3.1.5 $M = 150$, and the observed three components are V , T_1 and E with different tolerances

The tolerances for the ODE solver when solving the inverse problem are $\{1e-4, 1e-6\}$. After discarding the first 10 samples, we retain every third sample from the remaining samples. The trace plots and histograms for γ , θ , and D are presented in Figure 17 and those for De are presented in Figures 18. From these figures, we see that all chains are stabilized.

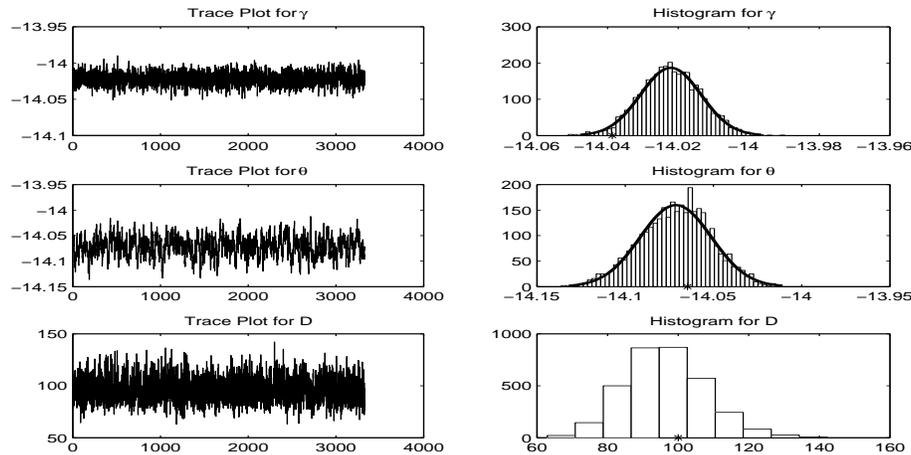


Figure 17: Trace plots and histograms for γ , θ and D .

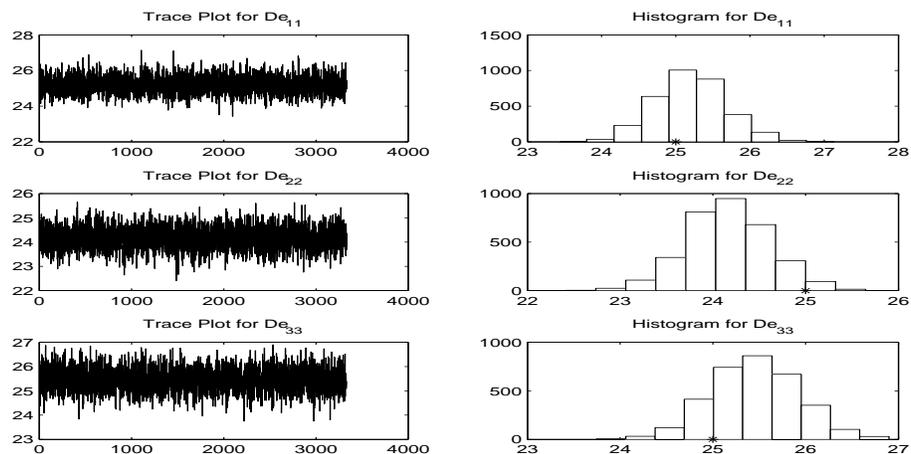


Figure 18: Trace plots and histograms of De .

In order to test the mixing behavior of the chain, autocorrelation plots for γ , θ and D are presented in Figure 19, and those for De are presented in Figure 20. From these graphs, we see that each chain is well-mixed.

The statistical results are presented in Table 5. Comparing the results of this table with those in Table 1, we see that allowing different tolerances can lead to slightly worse estimates, but the estimates are still acceptable.

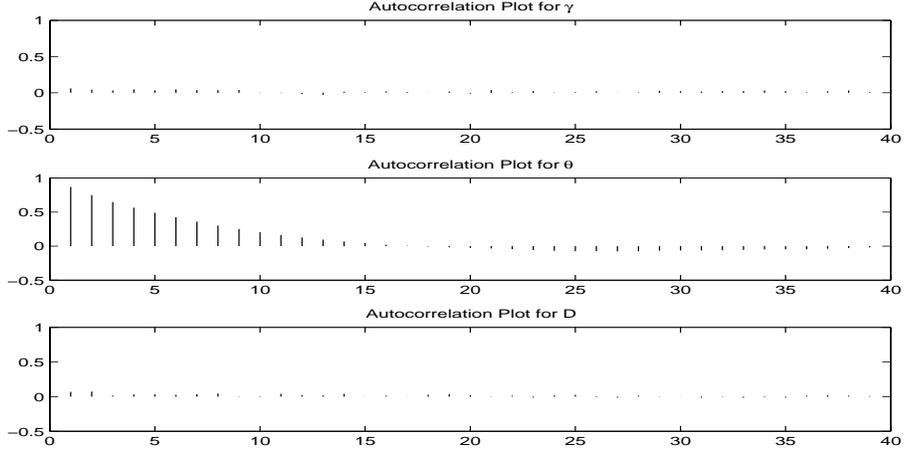


Figure 19: Autocorrelation plots for γ , θ and D , respectively.

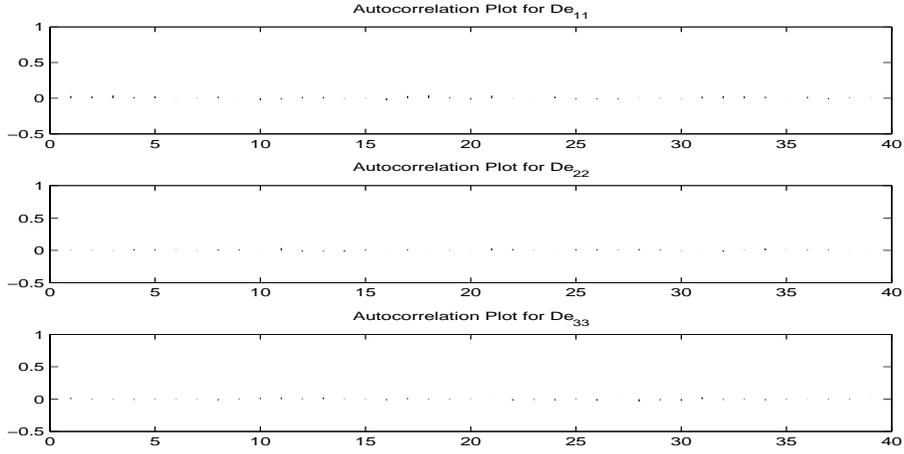


Figure 20: Autocorrelation plots for elements of De .

	true	mean	mode	95% HDR
$\ln(k_1)$	-14.0387	-14.0221	-14.0235	$[-14.0382, -14.0056]$
θ	-14.0647	-14.0713	-14.0681	$[-14.1128, -14.0331]$
D	100.0000	96.3787	98.6070	$[75.8971, 120.7532]$
De_{11}	25.0000	25.1982	25.0950	$[24.2996, 26.1663]$
De_{22}	25.0000	24.1631	24.1920	$[23.2893, 25.0420]$
De_{33}	25.0000	25.4498	25.4800	$[24.5696, 26.3790]$

Table 5: Statistical results for the estimated parameters when there are 150 patients and the observed three components are V , T_1 and E with different tolerances.

3.2 Simulation: k_1 and c with different tolerances

In this subsection, we assume that $M = 150$, the measurements are taken on V , T_1 and E . The tolerances for the ODE solver when solving the inverse problem are $\{1e-4, 1e-6\}$. The left limits γ_L of γ are chosen to be $(\ln(10^{-16}), \ln(0.1))$ and the right limits γ_R of γ are chosen to be $(\ln(0.1), \ln(10^4))$. The true parameter values of θ_i are produced from a truncated normal distribution with mean $(\ln(8 \times 10^{-7}), \ln(13))$ and variance $\Sigma = 0.01I_2$. The proposed variance $\Delta = 10^{-3}I_2$. The values of the hyperparameters are chosen as follows:

$$\alpha_l = 2.5, \beta_l = 10 \text{ for } l = 1, 2, 3, \quad \eta = [\ln(8 \times 10^{-7}), \ln(13)], \quad \Lambda = I_2, \\ \nu = 2, \quad \Omega = 0.01I_2.$$

We ran one long chain with length 80000. After discarding the first 15000 samples, we take every tenth from the remaining samples. The trace plots and histograms for γ , θ , D and De are presented in Figures 21, 22, 23 and 24, respectively. From these figures, we see that all chains are stabilized.

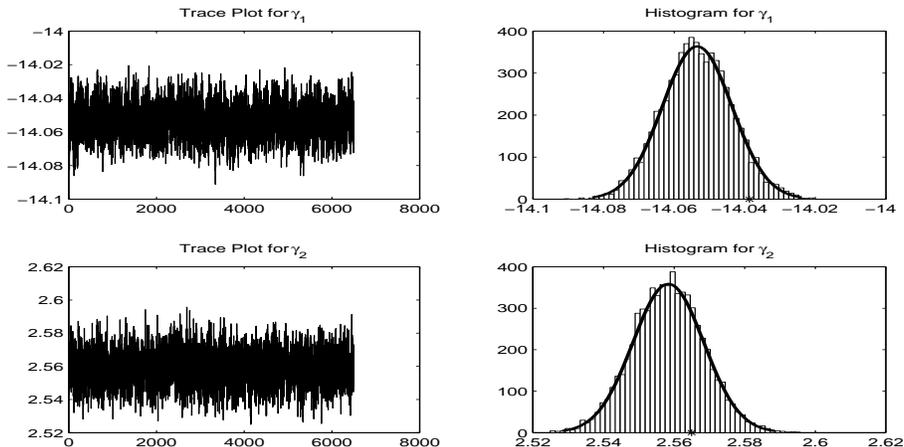


Figure 21: Trace plots and histograms for γ .

In order to test the mixing behavior, the autocorrelation plots of γ , θ , D and De are presented in Figures 25, 26, 27 and 28, respectively. From these graphs, we see that all the chains are well-mixed.

The statistical results are presented in Table 6. Although the results are somewhat worse than the results presented for the case with only one parameter unknown ($p = 1$), they are still reasonably good, which provides positive affirmation on the efficiency of our method when dealing with the case with multiple unknown parameters ($p > 1$).

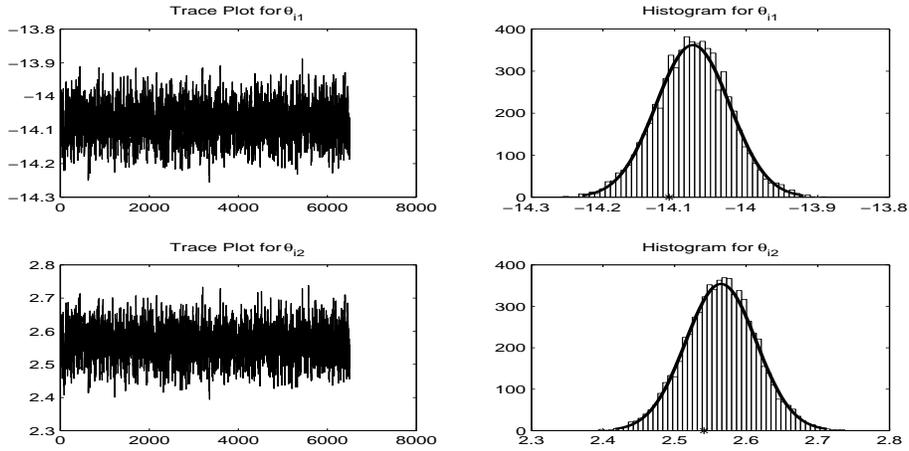


Figure 22: Trace plots and histograms for θ .

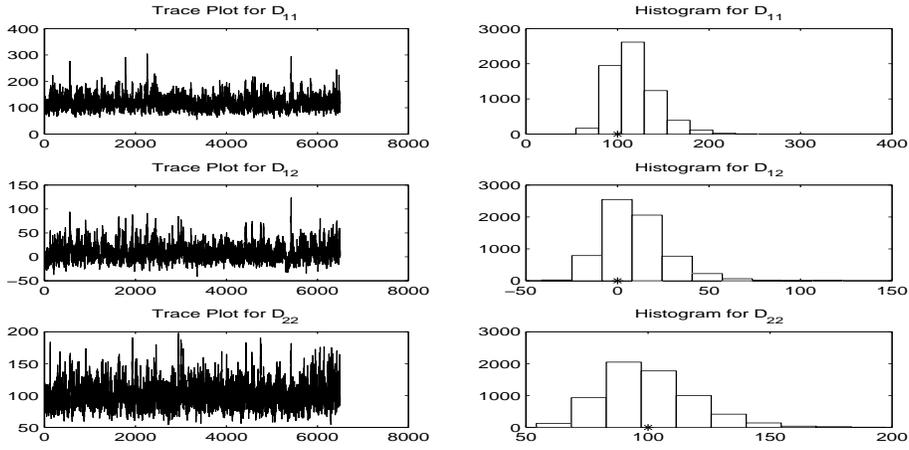


Figure 23: Trace plots and histograms for D .

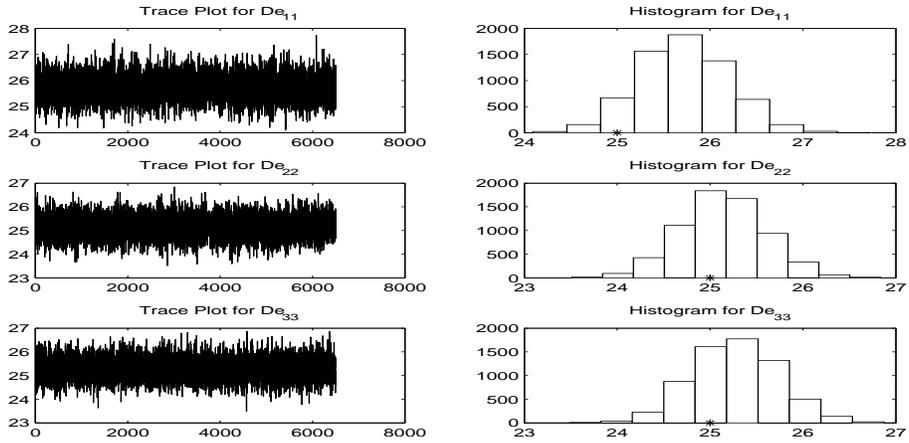


Figure 24: Trace plots and histograms of De .

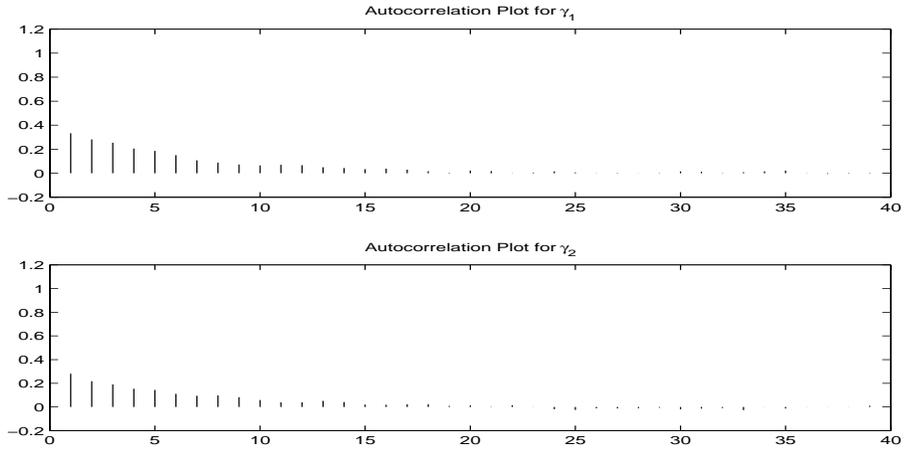


Figure 25: Autocorrelation plots for γ .

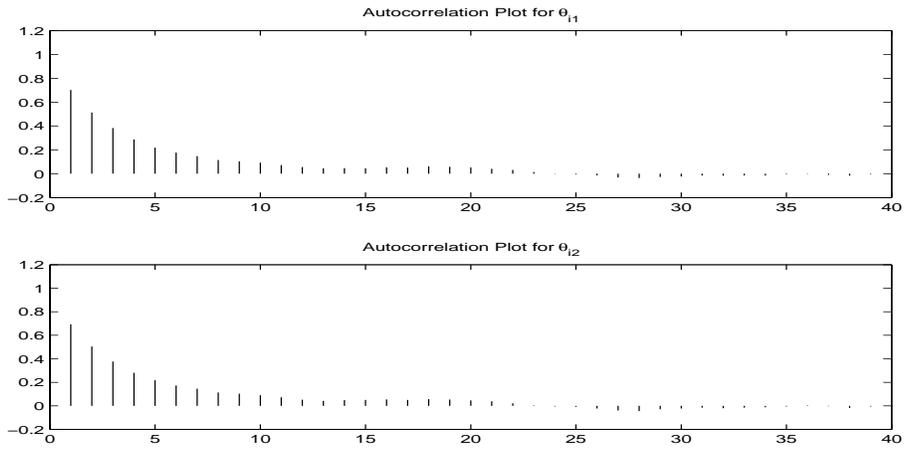


Figure 26: Autocorrelation plots for θ_i .

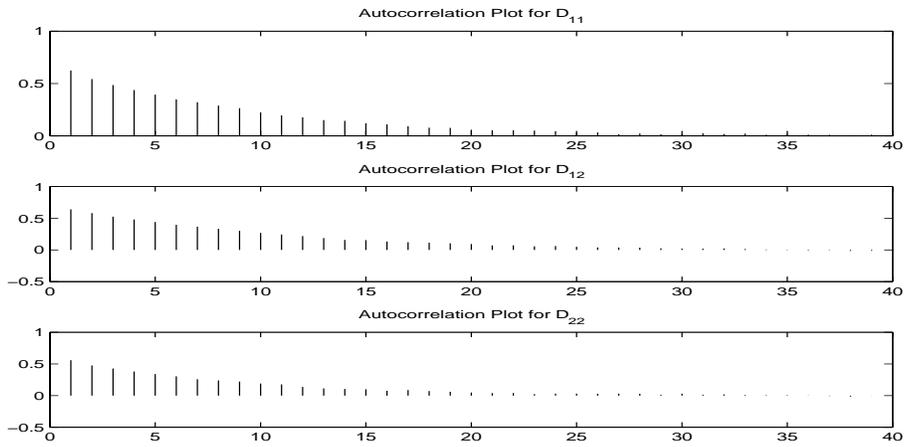


Figure 27: Autocorrelation plots for D .

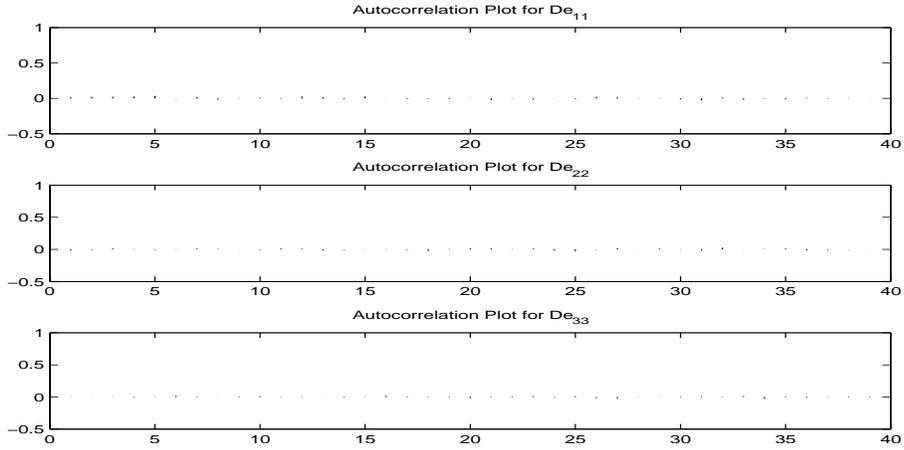


Figure 28: Autocorrelation plots for elements of De .

	true	mean	mode	95% HDR
$\ln(k_1)$	-14.0387	-14.0535	-14.0523	$[-14.0726, -14.0335]$
$\ln(c)$	2.5649	2.5584	2.5567	$[2.5389, 2.5782]$
θ_{i1}	-14.1076	-14.0774	-14.0903	$[-14.1782, -13.9721]$
θ_{i2}	2.5411	2.5649	2.5833	$[2.4658, 2.6635]$
D_{11}	100.0000	117.7546	116.8856	$[79.2181, 176.8763]$
D_{12}	0.0000	9.3624	-0.2182	$[-17.5255, 49.7824]$
D_{22}	100.0000	100.6734	90.0745	$[70.5071, 144.6248]$
De_{11}	25.0000	25.7205	25.7325	$[24.8019, 26.6694]$
De_{22}	25.0000	25.1390	25.0072	$[24.2604, 26.0487]$
De_{33}	25.0000	25.3470	25.3473	$[24.4032, 26.1895]$

Table 6: Statistical results for γ , an arbitrarily chosen patient θ_i , D and De .

4 Simulation Studies II: Subjects Undergoing Therapy

In this section, we apply the proposed Bayesian approach to estimate the parameters of interest in the case that an STI strategy is adopted for the patient treatment. We assume that each patient has at least one treatment interruption, and the duration of interruption is randomly produced, from several weeks to several months or even years. We assumed $M = 150$, Y_{ij1} is confined to be in the range $[\log(50), \log(750000)]$, and the measurements are taken on V , T_1 and E . The number of the measurements of each individual n_i is uniformly distributed between 2 and 50, and these measurements are uniformly distributed in a 4 year time range. The data set we used in this section is generated computationally with the variance $\Sigma_e = 9I$ and we used the same tolerances for the ODE solver in generating the data and in solving the inverse problem. Note that there is a large variability in our data and we also have a very sparse data set which is in some way similar to the clinical data. We carried out this simulation study since we think it can further assist in testing and understanding the capability of our estimation technique. We also think it is a necessary first step before use of this method with clinical data containing some generally unknown sources of variability and some missing observations.

We consider three situations in this section: (1) k_1 is unknown, (2) k_2 is unknown, (3) k_1 and c are both unknown. In each situation, trace plots, histograms and autocorrelation plots are presented for the estimated parameters γ , an arbitrarily chosen patient θ_i , D and De . We also present their corresponding statistical inference results: mean, mode, 95% Highest Density Region.

4.1 Simulation: k_1

In this subsection, we assume that all the parameters are known except for k_1 . The values for γ_L and γ_R are chosen to be $\ln(10^{-16})$ and $\ln(0.1)$, respectively. The true parameter values of θ_i are produced from a truncated normal distribution with mean $\ln(8 \times 10^{-7})$ and standard deviation 1. Note that the variability between the true parameter values of θ_i is also bigger than that in section 3. The proposed variance is $\Delta = 1$. The value of the hyperparameters are chosen as follows:

$$\alpha_l = 1, \beta_l = 1 \text{ for } l = 1, 2, 3, \quad \eta = \ln(8 \times 10^{-7}), \Lambda = 0.1, \quad \nu = 1, \Omega = 1.$$

We ran one long chain with length 10000, and we discarded the first 10 samples. Trace plots and histogram plots for γ , θ , and D are presented in Figure 29 and those for De are presented in Figures 30. From these figures, we see that all the chains are stabilized.

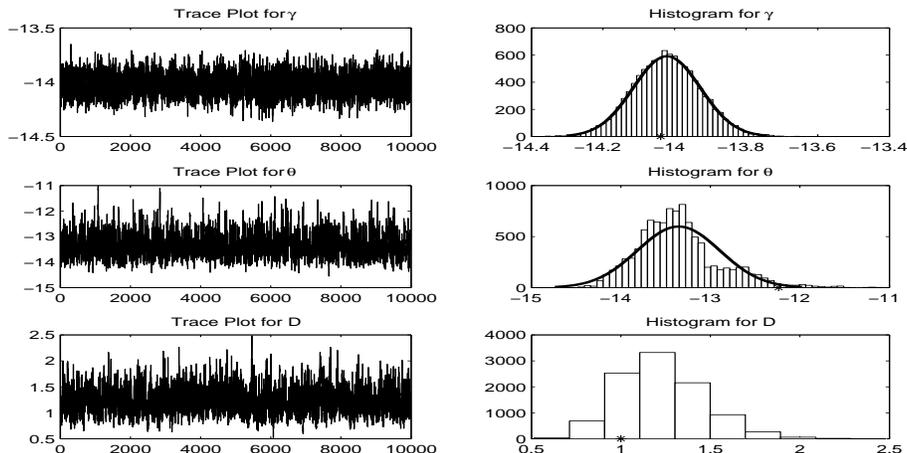


Figure 29: Trace plots and histograms for γ , θ and D . The point $*$ in the histogram plots denotes the true value for the corresponding parameter.

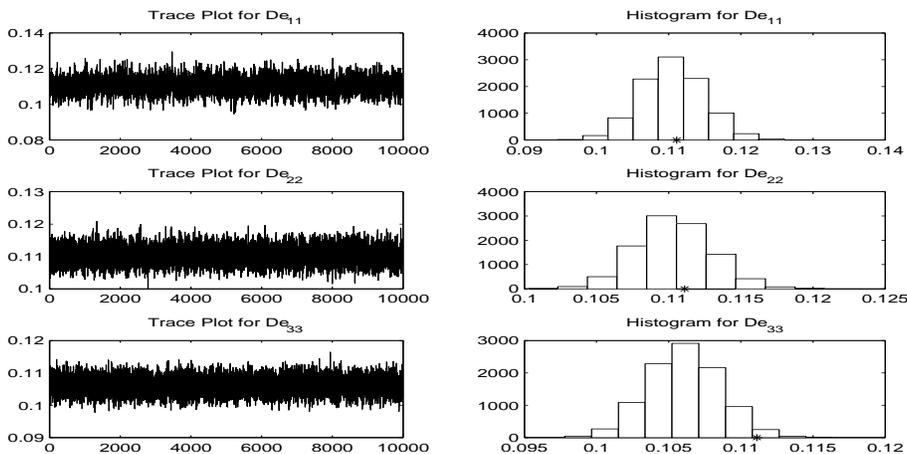


Figure 30: Trace plots and histograms of De .

In order to test the mixing behavior of the chain, autocorrelation plots for γ , θ and D are presented in Figure 31, and those for De are presented in Figure 32. From each figure, we see that the autocorrelation coefficient is geometrically decreasing to 0 as the lag increases, which means that all the chains are well-mixed.

In order to examine the dependence of the parameter estimates on the starting values, we run three chains with length 10000 starting from different initial values. The trace plots for γ , θ and D are presented in Figure 33, and those for De are presented in Figure 34. From each of the graphs, we can see that the three chains are well-mixed. Hence, the parameter estimates are not sensitive to the starting values.

The statistical results are presented in Table 7. From this table, we see that we still obtain reasonably good estimates on all the parameters of interest even in the case that

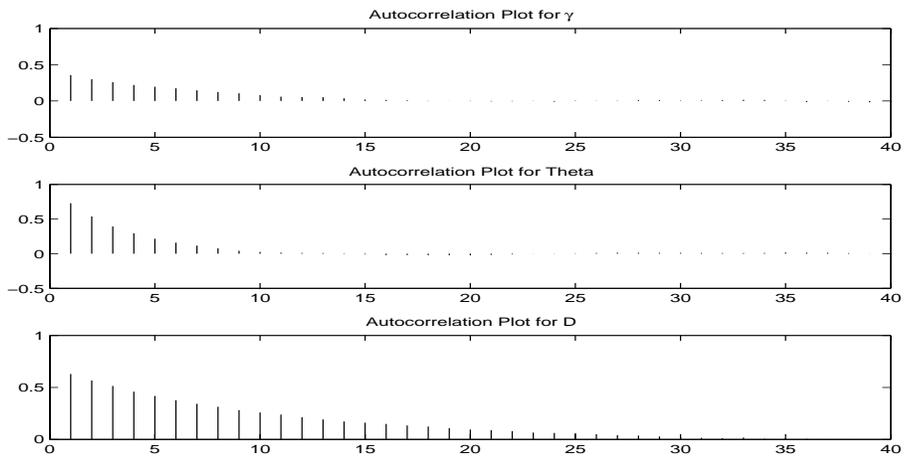


Figure 31: Autocorrelation plots for γ , θ and D , respectively.

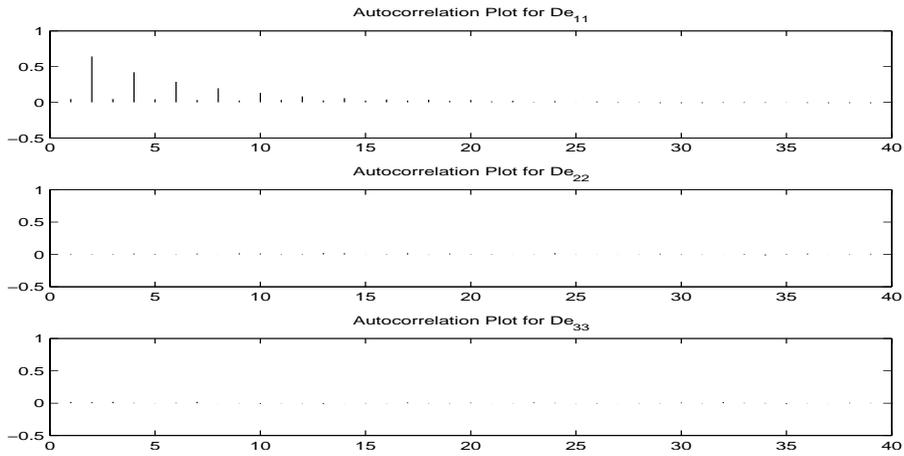


Figure 32: Autocorrelation plots for each element of De .

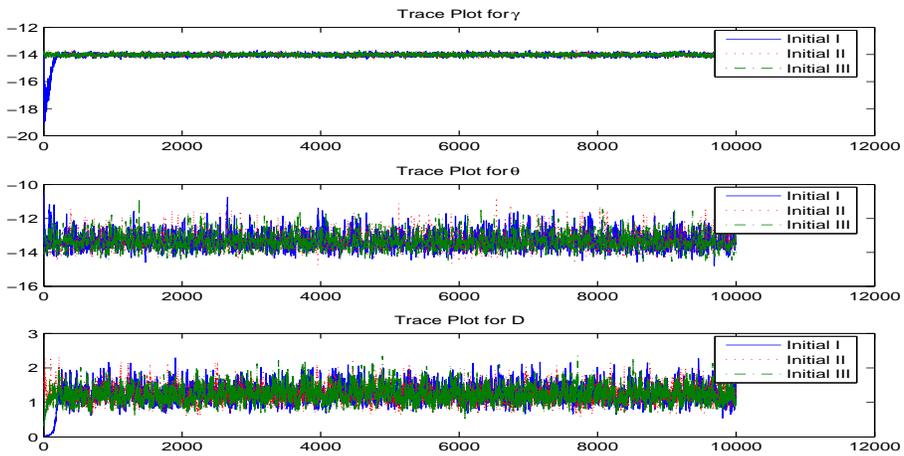


Figure 33: Autocorrelation plots for γ , θ and D , respectively.

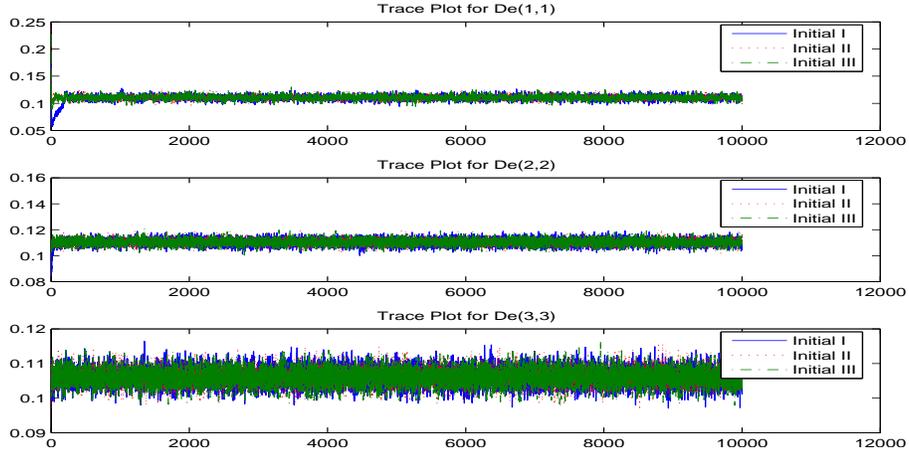


Figure 34: Autocorrelation plots for each element of De .

	true	mean	mode	95% HDR
$\ln(k_1)$	-14.0387	-14.0215	-14.0420	$[-14.2038, -13.8353]$
θ_i	-12.2372	-13.3548	-13.3126	$[-14.1213, -12.3100]$
D	1.0000	1.2270	1.2042	$[0.8256, 1.7364]$
De_{11}	0.1111	0.1106	0.1065	$[0.1021, 0.1194]$
De_{22}	0.1111	0.1104	0.1095	$[0.1054, 0.1156]$
De_{33}	0.1111	0.1061	0.1062	$[0.1013, 0.1110]$

Table 7: Statistical results for γ , θ_i , D and De .

the data is sparse and has a large variability. This provides some confidence regarding the capability of our method in future dealing with clinical data.

4.2 Simulation: k_2

In this subsection, we assume that all the parameters are known except for k_2 . The values for γ_L and γ_R are chosen to be $\ln(10^{-16})$ and $\ln(0.1)$, respectively. The true parameter values of θ_i are produced from a truncated normal distribution with mean $\ln(10^{-4})$ and standard deviation 1. The proposed variance is $\Delta = 1$. The value of the hyperparameters are chosen as follows:

$$\alpha_l = 1, \beta_l = 1 \text{ for } l = 1, 2, 3, \quad \eta = \ln(10^{-4}), \Lambda = 0.1, \quad \nu = 1, \Omega = 1.$$

We ran one long chain with length 10000. After discarding the first 10 samples, we retain every second from the remaining samples. Trace plots and histogram plots for γ , θ , and D are presented in Figure 35 and those for De are presented in Figures 36. From these figures, we see that all the chains are stabilized.

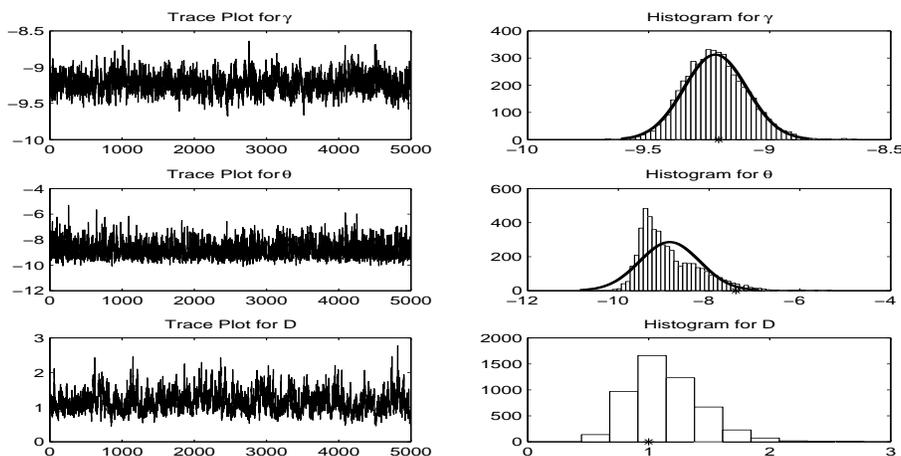


Figure 35: The trace plots and histogram plots for γ , θ and D .

In order to test the mixing behavior of the chain, autocorrelation plots for γ , θ and D are presented in Figure 37, and those for De are presented in Figure 38. From these graphs, we see that the chain is well-mixed.

In order to examine the dependence of the parameter estimates on the starting values, we run three chains with length 10000 starting from different initial values. The trace plots for γ , θ and D are presented in Figure 39 and those for De are presented in Figure 40. From each of the graphs, we see that the three chains are well-mixed. Hence, it appears that the estimated parameters are not sensitive to the starting values.

The statistical results are presented in Table 8. Comparing the results in Table 7 and the results here, we see that the estimates in Table 8 are somewhat worse than the estimates in Table 7. However, the estimates are still acceptable.

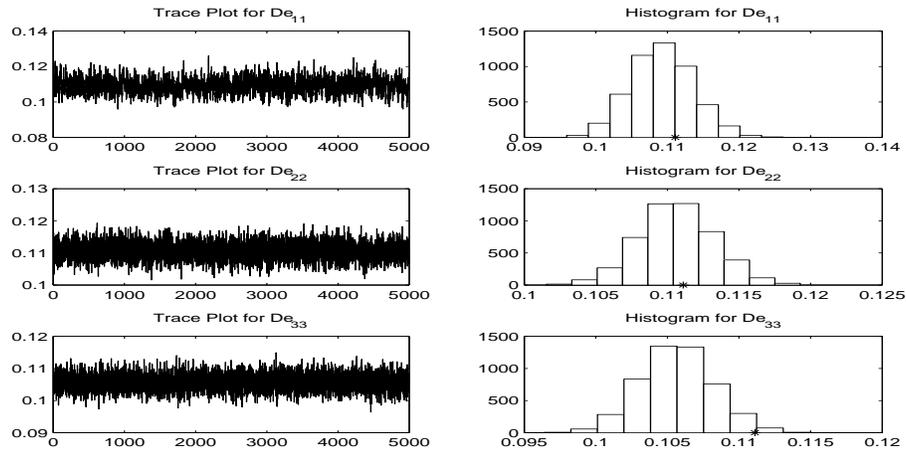


Figure 36: The trace plots and histograms for each element of De .

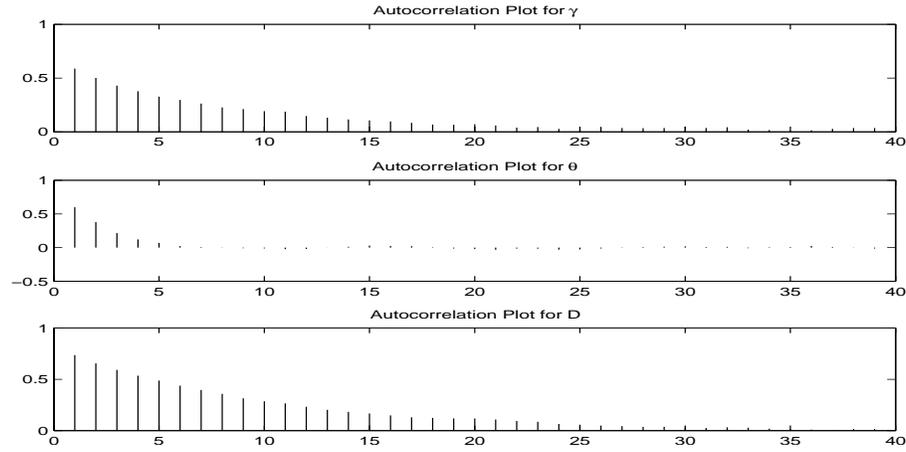


Figure 37: Autocorrelation plots for γ , θ_i and D , respectively.

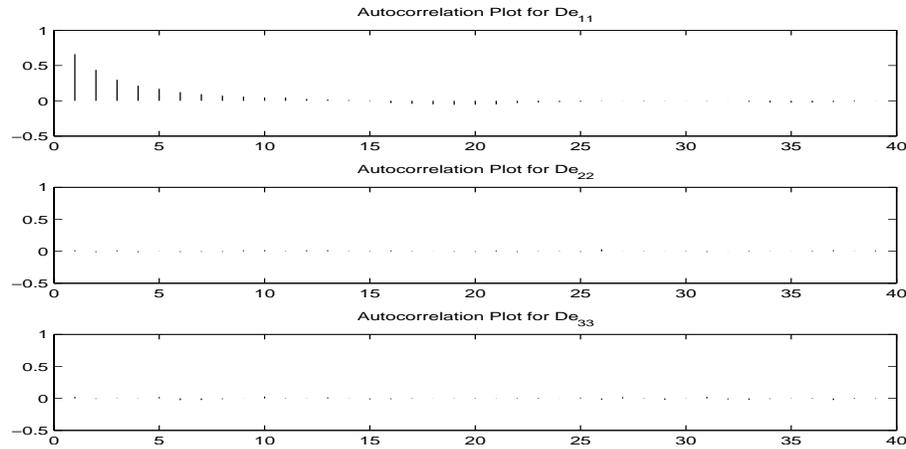


Figure 38: Autocorrelation plots for elements of De .

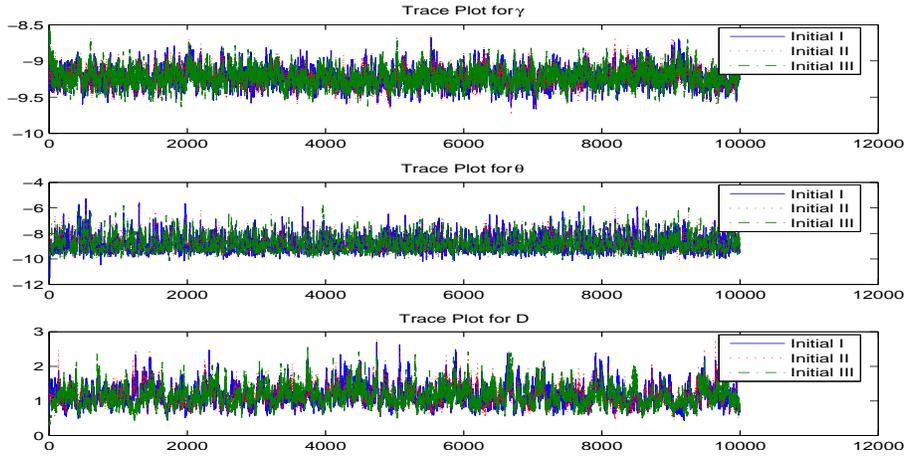


Figure 39: Autocorrelation plots for γ , θ and D , respectively.

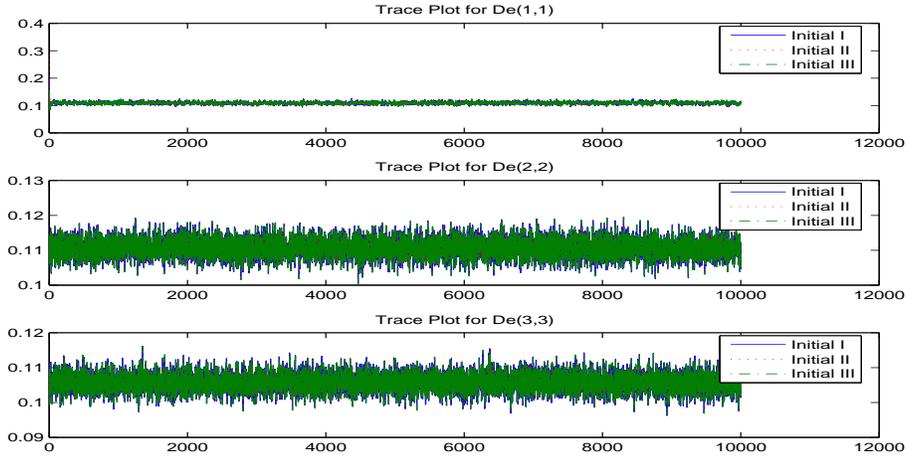


Figure 40: Autocorrelation plots for each element of De .

	true	mean	mode	95% HDR
$\ln(k_2)$	-9.2103	-9.2239	-9.2112	$[-9.4609, -8.9578]$
θ	-7.4089	-8.8682	-9.4019	$[-9.7192, -7.2858]$
D	1.0000	1.1431	1.0295	$[0.6728, 1.8013]$
De_{11}	0.1111	0.1092	0.1095	$[0.1010, 0.1181]$
De_{22}	0.1111	0.1106	0.1113	$[0.1055, 0.1159]$
De_{33}	0.1111	0.1057	0.1047	$[0.1009, 0.1107]$

Table 8: Statistical results for γ , an arbitrarily chosen patient θ_i , D and De .

4.3 Simulation: k_1 and c

In this subsection, we assume that all the parameters are known except for k_1 and c . The left limits γ_L of γ are chosen to be $(\ln(10^{-16}), \ln(0.1))$ and the right limits γ_R of γ are chosen to be $(\ln(0.1), \ln(100))$. The true parameter values of θ_i are produced from a truncated normal distribution with mean $(\ln(8 \times 10^{-7}), \ln(13))$ and variance $\Sigma = I_2$. The proposed variance is chosen as $\Delta = I_2$. The values of the hyperparameters are chosen as follows:

$$\alpha_l = 0.1De, \beta_l = 10 \text{ for } l = 1, 2, 3, \quad \eta = (\ln(8 \times 10^{-7}), \ln(13)), \nu = 2,$$

$$\Lambda = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \Omega = \begin{pmatrix} 2 & 0 \\ 0 & 2 \end{pmatrix}.$$

We ran one long chain with length 80000. After discarding the first 40000 samples, we retain every tenth from the remaining samples. An informal graphical check of convergence for γ , θ , D and De are presented in Figures 41-44, respectively. From all these figures, we see that all the chains are stabilized.

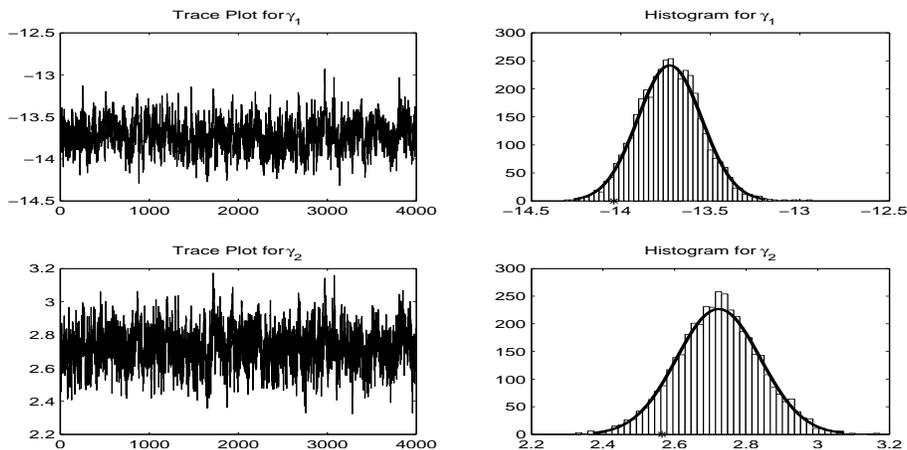


Figure 41: The trace plots and histogram plots of γ .

In order to test the mixing behavior of the chain, autocorrelation plots for γ , θ_i , D and De are presented in Figure 45-48. From these graphs, we see that all the chains are well-mixed.

The statistical results are presented in Table 9. Considering the large variability of the data, the results presented here are still quite reasonable.

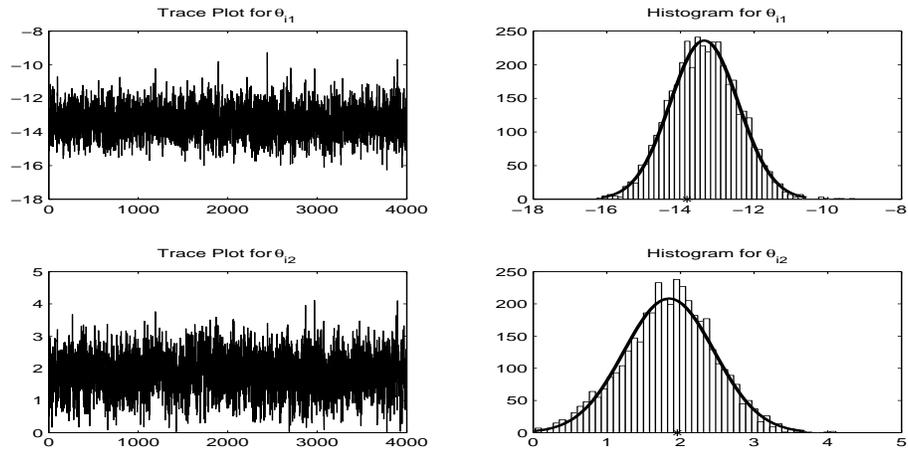


Figure 42: The trace plots and histograms of θ_i for a randomly chosen patient.

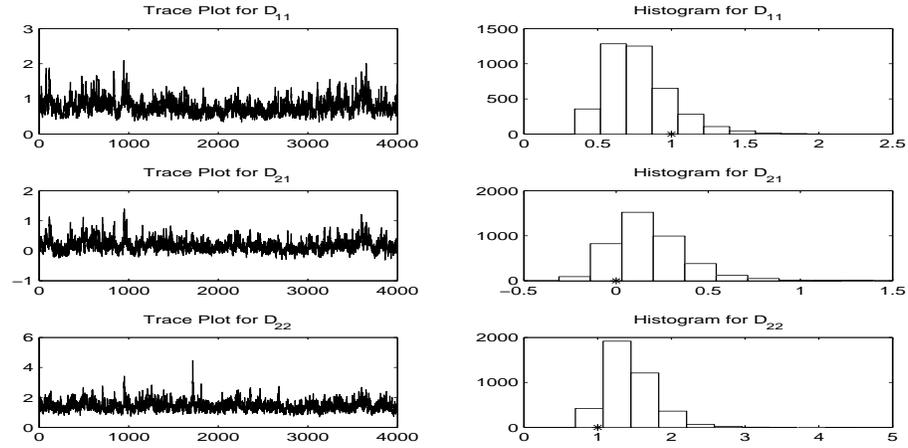


Figure 43: The trace plots and histogram plots of D .

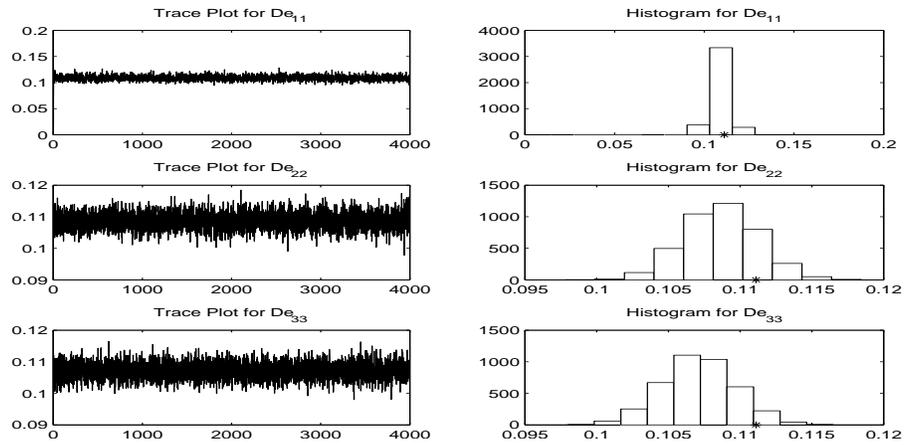


Figure 44: The trace plots and histogram plots for each element of De .

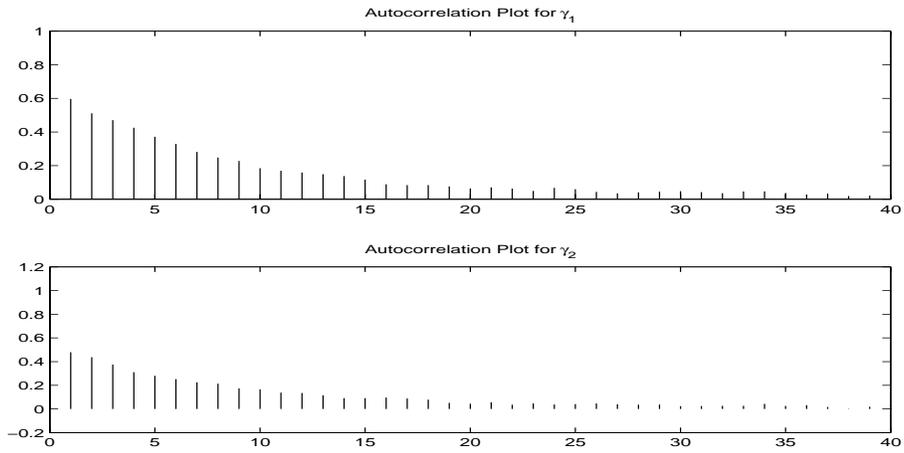


Figure 45: Autocorrelation plots for γ .

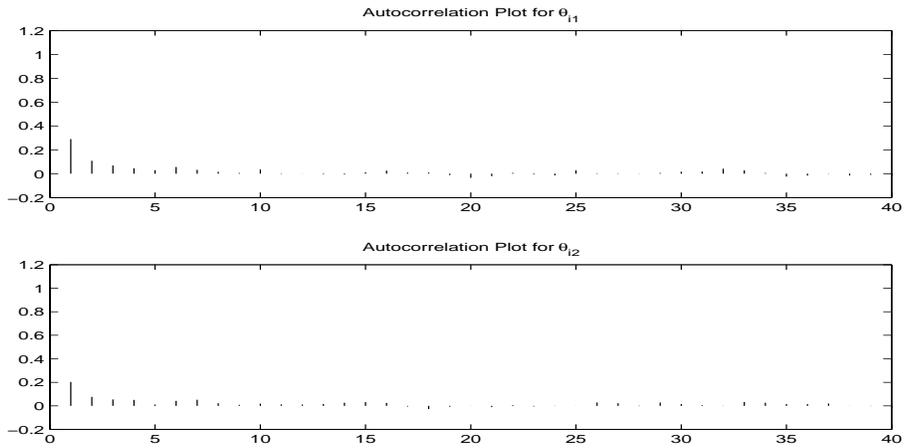


Figure 46: Autocorrelation plots for elements of θ_i .

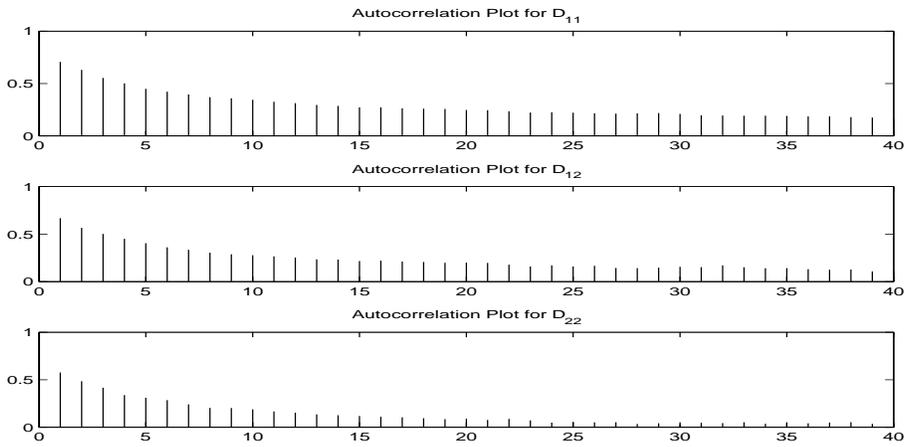


Figure 47: Autocorrelation plots for D .

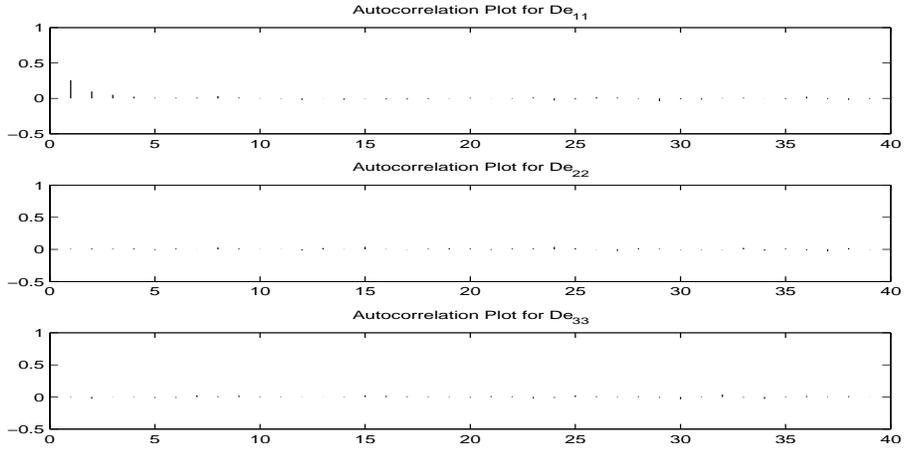


Figure 48: Autocorrelation plots for elements of De .

	true	mean	mode	95% HDR
$\ln(k_1)$	-14.0387	-13.7267	-13.6935	$[-14.0670, -13.3738]$
$\ln(c)$	2.5649	2.7228	2.7058	$[2.4909, 2.9543]$
θ_{i1}	-13.8228	-13.3551	-13.1125	$[-15.1261, -11.5572]$
θ_{i2}	1.9575	1.8464	1.8660	$[0.5791, 3.0058]$
D_{11}	1.0000	0.7748	0.6065	$[0.4529, 1.3211]$
D_{12}	0.0000	0.1764	0.1165	$[-0.1352, 0.6557]$
D_{22}	1.0000	1.4333	1.2627	$[0.9391, 2.1882]$
De_{11}	0.1111	0.1090	0.1094	$[0.1002, 0.1180]$
De_{22}	0.1111	0.1087	0.1092	$[0.1038, 0.1136]$
De_{33}	0.1111	0.1071	0.1063	$[0.1022, 0.1121]$

Table 9: Statistical results for γ , θ_i , D and De .

5 Conclusion and Remarks

In this paper we presented a hierarchical Bayesian NLME parameter estimation algorithm which incorporates and accounts for uncertainty and population level variability that arises in HIV pathogenesis. Modifications to allow for censored data as well as biologically/physically-based parameter constraints were incorporated. We investigated use of this algorithm with simulated data for subjects without and with ongoing STI therapy. Several conclusions can be drawn from our studies:

1. As expected, use of data sets with larger numbers of patients led to improved performance in the estimation procedures (compare, for example, the results in Table 1 with those in Table 2, and the results in Table 3 with those in Table 4).
2. Any effects of the type of the data used ((V, T_1, E) vs. $(V, T_1 + T_1^*, E)$) on the reliability of the estimation procedure is not obvious (compare the results in Table 1 with those in Table 3, and the results in Table 2 with those in Table 4). There is no clear degradation of results when using observations from combined compartments (a practical necessity in many cases).
3. The test results do not appear to be very sensitive to the ODE solver tolerances used in the MCMC simulations (compare the results in Table 1 with those in Table 5).
4. We obtain reasonable estimation results even in the case that the available data set is sparse and has large variability. This suggests that one might have confidence when using this approach to deal with the clinical data in our planned future efforts. Our experimental data sets (see [1]) do exhibit significant variability.

In closing, we need to remark on one possible shortcoming due to the parametric nature of the proposed method. Note that in using a Bayesian parametric method to estimate the parameters, we must impose a fixed form for the parameter distributions (the priors). If the assumed parameter distributional forms are correct, this method can provide a very efficient method for estimation since it can fully use the knowledge of the distributional structures. This is readily demonstrated by our numerical and statistical results in Sections 3 and 4. However, if the assumed distributional forms are incorrect, parametric methods can often lead to serious difficulties. For example, they could either produce a model that fits the data poorly or a reasonable model fit to the data with an incorrect distribution (see the examples in [6]). In this case, a nonparametric estimation method such as one based on the Prohorov Metric Framework developed and discussed in [3, 4, 6] could be more appropriate since it does not require one to provide a distributional structure for the estimated parameters. One

such method has been successfully used in PBPK modeling efforts (see [5, 6]). It is currently being investigated by our group with the model (1.1).

Acknowledgements: This research was supported in part by the Joint DMS/NIGMS Initiative to Support Research in the Area of Mathematics Biology under grant 1R01GM67299-01, and was facilitated through visits of the authors to (and previous postdoctoral support of Y. M. by) the Statistical and Applied Mathematical Science Institute (SAMSI), which is funded by NSF under grant DMS-0112069. We would like to thank Brian M. Adams for providing the ODE software (a Linux version of the Gear stiff solver) for solving the model (1.1). We are also grateful to Dr. Marie Davidian for numerous helpful conversations on statistical methods and to Dr. Eric Rosenberg for assistance in understanding the nature and limitations of the data collected in his clinical trials.

References

- [1] B. M. Adams, H. T. Banks, M. Davidian, Hee-Dae Kwon, H. T. Tran, S. N. Wynne, and E. S. Rosenberg, HIV dynamics: Modeling, data analysis, and optimal treatment protocols, Tech. Rep. CRSC-TR04-05, NCSU, Raleigh, February, 2004; *J. Comp. and Appl. Math.*, to appear.
- [2] B. M. Adams, H. T. Banks, H. D. Kwon, and H. T. Tran, Dynamic multidrug therapies for HIV: Optimal and STI control approaches, Tech. Rep. CRSC-TR04-18, NCSU, Raleigh, April, 2004; *Math. Biosci. and Engr.*, **1** (2004), 223–241.
- [3] H. T. Banks and K. L. Bihari, Modeling and estimating uncertainty in parameter estimation, *Inverse Problems*, **17** (2001), 95–111.
- [4] H. T. Banks, D. Bortz, G. A. Pinter and L. K. Potter, Modeling and imaging techniques with potential for application in bioterrorism, Chapter 6 in *Bioterrorism: Mathematical Modeling Applications in Homeland Security*, (H.T. Banks and C. Castillo-Chavez, eds.), Frontiers in Applied Mathematics, SIAM, Philadelphia, 2003, 129–154.
- [5] H. T. Banks and L. K. Potter, Probabilistic methods for addressing uncertainty and variability in biological models: Application to a toxicokinetic model, Tech. Rep. CRSC-TR02-27, NCSU, Raleigh, September, 2002; *Math. Biosci.*, **192** (2004), 193–225.
- [6] H. T. Banks, Y. Ma, and L. K. Potter, A simulation-based comparison between parametric and nonparametric estimation methods in PBPK models, Tech. Rep. CRSC-

- TR04-25, NCSU, Raleigh, June, 2004; *J. Inverse and Ill-Posed Problems*, **13** (2005), 1–26.
- [7] S. Bonhoeffer, M. Rembiszewski, G.M. Ortiz, and D.F. Nixon, Risks and benefits of structured antiretroviral drug therapy interruptions in HIV-1 infection, *AIDS*, **14** (2000), 2313–2322.
- [8] B. P. Carlin and T. A. Louis, *Bayes and Empirical Bayes Methods for Data Analysis*, Chapman and Hall, London, 1996.
- [9] M. Davidian and D. M. Giltinan, *Nonlinear Models for Repeated Measurement Data*, Chapman and Hall, London, 1995.
- [10] L. Devroye, *Non-Uniform Random Variate Generation*, Springer-Verlag, New York, 1986.
- [11] A. E. Gelfand, and A. F. M. Smith, Sampling based approaches to calculating marginal densities, *Journal of the American Statistical Association*, **85** (1990), 398–409.
- [12] A. E. Gelfand, S. E. Hills, A. Racine-Poon, and A. F. M. Smith, Illustration of Bayesian inference in normal data models using Gibbs sampling, *Journal of the American Statistical Association*, **85** (1990), 972–985.
- [13] J. Geweke, Efficient Simulation from the Multivariate Normal and Student-t Distributions Subject to Linear Constraints and the Evaluation of Constraint Probabilities, presented at the meeting Computing Science and Statistics: the Twenty-Third Symposium on the Interface, Seattle, April 22–24, 1991.
- [14] C. Han, K. Chaloner, and A. S. Perelson, Bayesian analysis of a population HIV dynamic model, *Case Studies in Bayesian Statistics*, **6** (2002), 223–237, Springer-Verlag, New York.
- [15] C. Han and K. Chaloner, Bayesian experimental design for nonlinear mixed-effects models with applications to HIV dynamics, *Biometrics*, **60** (2004), 25–33.
- [16] Y. Huang, D. Liu, and H. Wu, Hierarchical Bayesian methods for estimation of parameters in a longitudinal HIV dynamic system, Dept. Biostat and Comp. Bio Tech Rpt 04-06, Univ. Rochester, July, 2004; *Biometrics*, submitted.
- [17] G. R. Kaufmann, et al., Patterns of Viral Dynamics During Primary Human Immunodeficiency Virus Type 1 Infection, *Journal of Infectious Disease*, **178** (1998), 1812–1815.

- [18] N. Lange, B. Carlin, and A. Gelfand, Hierarchical Bayes Models for the Progression of HIV Infection Using Longitudinal CD4 T-cell number, *Journal of Americal Statistical Association*, **87** (1992), 615–626.
- [19] J. Lisziewicz, E. Rosenberg, and J. Liebermann, Control of HIV despite the discontinuation of anti-retroviral therapy, *New England J. Med.*, **340** (1999), 1683–1684.
- [20] J. E. Mittler, B. Sulzer, A. U. Neumann and A. S. Perelson, Influence of Delayed Viral Production on Viral Dynamics in HIV-1 Infected Patients, *Math. Biosci.*, **152** (1998), 143–163.
- [21] P. Muller and G. L. Rosner, A Bayesian population model with hierarchical mixture priors applied to blood count data, *Journal of Americal Statistical Association*, **92** (1997), 1279–1292.
- [22] A. S. Perelson, P. Essunger, et al., Decay Characteristics of HIV-1-Infected Compartments During Combination Therapy, *Nature*, **387** (1997), 188–191.
- [23] A. S. Perelson and P. W. Nelson, Mathematical Analysis of HIV-1 Dynamics in vivo, *SIAM Review*, **41** (1999), 3–44.
- [24] H. Putter, S. H. Heisterkamp, J. M. A. Lange, and F. DeWolf, A Bayesian Approach to parameter estimation in HIV dynamic Models, *Statistics in Medicine*, **21** (2002), 2199–2214.
- [25] K. Radhakrishnan and A. C. Hindmarsh, Description and use of LSODE, the Livermore Solver for Ordinary Differential Equations, LLNL Report UCRL-ID-113855, Livermore, CA, December, 1993.
- [26] G. O. Roberts, Markov Chain concepts related to sampling algorithms, in *Markov Chain Monte Carlo in Practice*, eds. W. R. Gilks, S. Richardson and D. J. Spiegelhalter, Chapman and Hall, London, 1996, 45–57.
- [27] E.S. Rosenberg, M. Altfeld, and S.H. Poon, et al, Immune control of HIV-1 after early treatment of acute infection, *Nature*, **407** (2000), 523–526.
- [28] D. Wodarz and M. A. Nowak, Mathematical models of HIV pathogenesis and treatment, *Bioessays*, **24** (2002), 1178–1187.
- [29] H. Wu, A. A. Ding, and V. Gruttola, Estimation of HIV dynamic parameters, *Statistics in Medicine*, **17** (1998), 2463–2485.

- [30] H. Wu and A. A. Ding, Population HIV-1 dynamics *in vivo*: Applicable models and inferential tools for virological data from AIDS clinical trials, *Biometrics*, **55** (1999), 410–418.