

Treatment Monitoring of HIV-Infected Patients based on Mechanistic Models

Mélanie Prague,^{1,2,*} Daniel Commenges,^{1,2} Julia Drylewicz,³ and Rodolphe Thiébaud^{1,2}

¹University of Bordeaux, ISPED, Centre INSERM U897-Épidémiologie-Biostatistique, F-33000 Bordeaux, France

²INSERM, ISPED, Centre INSERM U897-Épidémiologie-Biostatistique, F-33000 Bordeaux, France

³Department of Immunology, University Medical Center Utrecht, The Netherlands

**email*: melanie.prague@isped.u-bordeaux2.fr

SUMMARY. For most patients, the HIV viral load can be made undetectable by highly active antiretroviral treatments highly active antiretroviral therapy: the virus, however, cannot be eradicated. Thus, the major problem is to try to reduce the side effects of the treatment that patients have to take during their life time. We tackle the problem of monitoring the treatment dose, with the aim of giving the minimum dose that yields an undetectable viral load. The approach is based on mechanistic models of the interaction between virus and the immune system. It is shown that the “activated cells model,” allows making good predictions of the effect of dose changes and, thus, could be a good basis for treatment monitoring. Then, we use the fact that in dynamical models, there is a nontrivial equilibrium point, that is with a virus load larger than zero, only if the reproductive number R_0 is larger than one. For reducing side effects, we may give a dose just above the critical dose corresponding to R_0 equal to 1. A prior distribution of the parameters of the model can be taken as the posterior arising from the analysis of previous clinical trials. Then the observations for a given patient can be used to dynamically tune the dose so that there is a high probability that the reproductive number is below one. The advantage of the approach is that it does not depend on a cost function, weighing side effects and efficiency of the drug. It is shown that it is possible to approach the critical dose if the model is correct. A sensitivity analysis assesses the robustness of the approach.

KEY WORDS: Bayes; Differential equations; Epidemiology; HIV; Metropolis–Hastings algorithm; Monitoring; Optimal control.

1. Introduction

In developed countries, most HIV-infected patients are treated with highly active antiretroviral therapies highly active antiretroviral therapy: which include a combination of several antiretroviral drugs. For many patients, such treatments succeed to control viral replication (i.e., the viral load becomes undetectable) and to restore the immune system sufficiently to dramatically decrease the risk of opportunistic diseases and AIDS (Egger et al., 2002). The virus is, however, still present in sanctuaries like resting memory cells (Chomont et al., 2009). Several side effects have been described in patients having long-term therapy (Carr and Cooper, 2000). To reduce these side effects, treatment interruptions have been attempted. However, these trials have been stopped at interim analysis showing an increased risk of opportunistic diseases (Ananworanich et al., 2006; Danel et al., 2006; Lundgren et al., 2008). Attempts have been made to decrease the number of drugs but did not succeed either (Girard et al., 2009). Reducing the doses of treatment may be a flexible way to alleviate side effects while maintaining efficiency (Sánchez-Conde et al., 2005; Milinkovic et al., 2007; Meynard et al., 2010). These studies, however, proposed an arm-specific reduction (same reduction for every patient from the same arm). A subject-specific dose adaptation (dose individualization) will probably lead to more significant results.

Recently, adaptive treatment strategies have been proposed in the statistical literature (Murphy, 2003; Murphy and McKay, 2004; Moodie, Richardson, and Stephens, 2007; Lavori and Dawson, 2008; Henderson, Ansell, and Alshibani, 2010). These approaches adapt some concepts of control theory, in particular dynamic programming, developed in engineering (Sage and White, 1977). In engineering, dynamical models on which optimal control is based are more complex than those generally used for adaptive treatment strategies and they are often expressed through a system of differential equations. Kirschner, Lenhart, and Serbin (1997) proposed to use control theory based on mechanistic models of the interaction of HIV and the immune system for adapting the dose of antiretroviral treatment. Indeed, since the pioneering work of Ho et al. (1995) and Wei et al. (1995), several mechanistic models have been developed. However, treatment monitoring raises different issues than control problems in engineering and the estimation of parameters in HIV dynamics models is challenging. Recent statistical approaches use random effect models with the aim to estimate the parameters of complex models using rather large samples from clinical trials: see Wu (2005) for a review and Putter et al. (2002); Huang, Liu, and Wu (2006); Guedj, Thiébaud, and Commenges (2007); Huang, Wu, and Acosta (2010) among others for original developments.

The aim of this article is to examine the possibility of using such sophisticated mechanistic models, with parameters that can be estimated on previous clinical trial data, for monitoring the treatment dose of an individual patient. A Bayesian approach seems here natural (Berger, 1985). Previous data give a prior for the parameters of the model. Observations for a new patient give additional information, especially on parameters varying between patients, and permit to update the priors. Then the dose can be adapted so as to minimize a risk function. One key idea in this article is that, for reasonable risk functions, the target dose is the critical dose which makes the reproductive number R_0 equal to (or just below) one. After presenting the method, we shall examine whether it can work in real life. This is a “proof of concept” article: real patients have not yet taken doses recommended by our method but we give arguments to convince that mechanistic models, and this particular method, can be used for treatment monitoring. We first examine, using real data, whether our model can detect an influence of the treatment dose; then we examine the predictive ability of the model, still using real data. Then, we apply the control method on simulated data.

The article is organized as follows. Section 2 recalls the mathematical and statistical models developed by Guedj et al. (2007) and Drylewicz et al. (2010). Section 3 tackles the control problem first treating the case with known parameters, then the case with unknown parameters. The optimal dose is defined as the dose which controls the probability that $R_0 < 1$ and this can be computed by use of a MCMC algorithm. In Section 4, a detailed analysis of the ALBI clinical trial is presented. In this trial, patients received different treatments or doses (patients in the third arm shifted treatment and some changed their dose) allowing us to study the predictive ability of the model. In Section 5, we present a simulation study of the proposed strategy. The posterior distribution of the parameters is used as a prior for a new patient whose data are generated from the model. We examine the sensitivity of the method to misspecification of the prior. A conclusion and perspectives are given in Section 6.

2. Models of Interaction Between HIV and Immune System

2.1 Mathematical Models

Perelson et al. (1996) proposed a three-dimensional Ordinary Differential Equation (ODE) system featuring uninfected $CD4^+$ T cells, infected $CD4^+$ T cells (T^*) and viruses (V) concentrations. We shall work with an extension of this model called the “activated cells model,” which distinguishes among uninfected cells those which are quiescent (Q) and those which are activated (T). This model has been shown to fit the ALBI trial data much better (Commenges et al., 2008). The definition of the parameters can be found in Table 1.

The model can be written as:

$$\begin{cases} \frac{dQ}{dt} = \lambda + \rho T - \alpha Q - \mu_Q Q, \\ \frac{dT}{dt} = \alpha Q - \gamma TV - \rho T - \mu_T T, \\ \frac{dT^*}{dt} = \gamma TV - \mu_{T^*} T^*, \\ \frac{dV}{dt} = \pi T^* - \mu_V V. \end{cases} \quad (1)$$

Table 1
Biological parameters for the “activated cells model”

Parameter	Meaning
α	Activation rate of Q cells (day^{-1}).
μ_{T^*}	Death rate of T^* cells (day^{-1}).
λ	Rate of Q cells production ($\mu_L^{-1} \text{day}^{-1}$).
μ_T	Death rate of T cells (day^{-1}).
π	Rate of virions per T^* cell (day^{-1}).
ρ	Rate of reversion to the Q state (day^{-1}).
γ	Infectivity: Infection rate of T cells per virion ($\text{day}^{-1} \mu_L$).
μ_Q	Death rate of Q cells (day^{-1}).
μ_V	Death rate of free virions (day^{-1}).

The basic reproductive number R_0 (Gran et al., 2008) is given by:

$$R_0 = \frac{\gamma \pi \alpha \lambda}{\mu_{T^*} \mu_V (\rho \mu_Q + \alpha \mu_T + \mu_Q \mu_T)}. \quad (2)$$

If R_0 is lower than one, the only equilibrium point is the trivial one with $V = 0$, otherwise the trivial equilibrium is unstable and after introduction of the virus, the system stabilizes to a nontrivial equilibrium with $V > 0$ (see the formula in Web Appendix A).

2.2 Statistical Models

For the interindividual variability of the parameters, we use a statistical model with the same structure as proposed by Guedj et al. (2007). For $i = 1 \dots N$, we denote by ξ^i the vector of the nine individual biological parameters and by $\tilde{\xi}^i$ the vector of the log-transformed parameters:

$$\tilde{\xi}^i = (\tilde{\xi}_l^i, l = 1 \dots 9) = (\tilde{\alpha}^i, \tilde{\mu}_{T^*}^i, \tilde{\lambda}^i, \tilde{\mu}_T^i, \tilde{\pi}^i, \tilde{\rho}^i, \tilde{\mu}_Q^i, \tilde{\mu}_V^i, \tilde{\gamma}^i)^T.$$

The possible between-subjects variability and time variability of the parameters is modeled as:

$$\tilde{\xi}_l^i(t) = \phi_l + z_l^i(t) \beta_l + \omega_l^i u^i, \quad l = 1 \dots 9, \quad (3)$$

where ϕ_l is the intercept, z_l^i and ω_l^i are the vectors of (possibly time-dependent) explanatory variables associated to the fixed and random effects respectively of the l th biological parameter. The β_l are vectors of regression coefficients associated to the fixed effects; u^i is the individual vector of random effects. We assume $u^i \sim \mathcal{N}(0, \Sigma)$ with Σ a diagonal matrix.

Specifically, in the following, we shall use as explanatory variable only the treatments doses $d_j^i(t)$, where j is the treatment index. In the application, there are two treatments ($j = 1, 2$); for sake of simplicity, we present the theory for only one treatment. We assume that the treatment acts by diminishing the infectivity (this is indeed the case for the reverse transcriptase inhibitors used in the application). The model for the infectivity parameter is assumed to be:

$$\tilde{\gamma}^i = \tilde{\gamma}_0 + \beta \psi \{d^i(t)\}, \quad (4)$$

with $\beta < 0$ and where $\psi(\cdot)$ is a known increasing positive “pharmaco-dynamic” function. We assume in addition $\psi(0) = 0$ and $\psi(\cdot)$ unbounded.

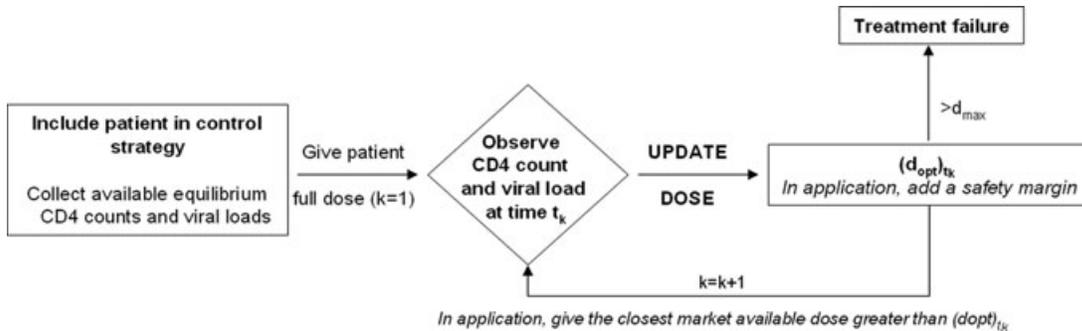


Figure 1. Flow chart for the control strategy procedure: individualized dose monitoring.

Let $X = (Q, T, T^*, V)$ in model (1). We used $g_1(X) = \log_{10}(V)$ and $g_2(X) = (Q + T + T^*)^{0.25}$ as transformations to achieve normality and homoscedasticity of measurement error distributions (Thiébaud et al., 2003) of the observed two compartments: the viral load (number of virions per μL) and the total CD4 count (number of cells per μL). Let Y_{jm}^i denote the j th measurements of the m th observable component for subject i at time t_{jm}^i ; we assume that:

$$\begin{aligned} Y_{j1}^i &= g_1\{X(t_{j1}^i, \tilde{\xi}^i)\} + \epsilon_{j1}^i & j = 1, \dots, n_1^i \\ Y_{j2}^i &= g_2\{X(t_{j2}^i, \tilde{\xi}^i)\} + \epsilon_{j2}^i & j = 1, \dots, n_2^i \end{aligned} \quad (5)$$

where ϵ_{j1}^i and ϵ_{j2}^i are independent Gaussian with zero mean and variances σ_{VL}^2 and σ_{CD4}^2 , respectively. The observation mechanism is generally complicated by a detection limit for the viral load, producing left-censoring for Y_{j1}^i (Jacqmin-Gadda et al., 2000).

3. Control Strategy

3.1 Targeting R_0 , Known Parameters

Our aim is to adapt the treatment dose at visits spaced by several weeks. Study of mathematical models and observed trajectories shows that the equilibrium state is essentially reached in a very short period of time. So, we can avoid the complexity of optimal control theory where the control is in continuous time and the cost function is an integral over time of a function of state and control. Instead, we define a cost function which depends only on the equilibrium state $\bar{X}(d; \xi^i)$ that is reached for a dose d : $C\{\bar{X}(d; \xi^i), d\}$, and more specifically $C\{\bar{V}(d; \xi^i), d\}$. The best dose for the patient minimizes $C\{\bar{V}(d; \xi^i), d\}$ on $[0, d_{max}]$. The function $C(\cdot, \cdot)$ is strictly increasing in both arguments. Then, we can use the fact that in our HIV dynamic model if $R_0(d, \xi^i) < 1$ there is only one equilibrium state which is the trivial one with $\bar{V} = 0$. The critical dose $d_{crit}(\xi^i)$ is defined as the solution of $R_0(d, \xi^i) = 1$. Because $\psi(\cdot)$ is an unbounded and increasing function in the dose, R_0 is a one-to-one function strictly decreasing with the dose. Thus, for any given ξ^i , a unique critical dose exists. We obtain the critical dose, $d_{crit}(\xi^i)$ from (2) and (4):

$$d_{crit}(\xi^i) = \psi^{-1} \left[\frac{1}{\beta} \ln \left\{ \frac{\mu_{T^*}^i \mu_V^i (\rho^i \mu_Q^i + \alpha^i \mu_T^i + \mu_Q^i \mu_T^i)}{\pi^i \alpha^i \lambda^i \gamma_0^i} \right\} \right] \quad (6)$$

Because $C(\cdot, \cdot)$ is increasing in its second argument, the optimal dose cannot be higher than $d_{crit}(\xi^i)$. If $C(\cdot, \cdot)$ rises

very sharply in its first argument, the optimal dose will be close to $d_{crit}(\xi^i)$. If $C(x, y)$ presents a discontinuity in $x = 0$, the optimal dose is $d_{crit}(\xi^i)$. This is a reasonable requirement from the medical point of view: $d_{crit}(\xi^i)$ is the minimum dose which achieves virus extinction (in practice, a very low viral load). This relieves us from the burden of specifying a cost function which would weigh the toxicity and the efficiency of the drug in an arbitrary way.

3.2 Unknown Parameters: Controlling the Probability that $R_0 < 1$

In practice, the values of the parameters of a particular subject are unknown. We have, however, a certain amount of knowledge about ξ^i . We make observations of both viral load and CD4 count at times before and after the initiation of the control strategy: $t_0, t_{-1}, t_{-2}, \dots$, and t_1, t_2, \dots , respectively; this gives essentially information about the values of the random effects u^i for this subject. Given $\mathcal{F}_{t_k}^i$, the information at time t_k , $R_0(d, \xi^i)$ has a posterior distribution. At time t_k , we propose to choose $d_{opt}^{t_k}$, the readjusted optimum dose, as the minimum dose which gives a high posterior probability (ω) that R_0 is below 1:

$$\mathbb{P}^{\mathcal{F}_{t_k}^i} [R_0\{d_{opt}^{t_k}(\omega), \xi^i\} < 1] = \omega. \quad (7)$$

We may take $\omega = 90\%$; in the following, we write $d_{opt}^{t_k}$ for the optimal dose at time t_k . R_0 is a decreasing one-to-one function, thus, $d_{opt}^{t_k}$ is unique. We denote by ξ^{i*} the true parameters value for patient i and $d_{crit}(\xi^{i*})$ his critical dose. If we assume that the information increases in time so that Doob's consistency theorem can be applied (Van der Vaart, 2000), $d_{opt}^{t_k}$ will tend to the critical dose (see Web Appendix B for the proof):

$$d_{opt}^{t_k} \xrightarrow[t_k \rightarrow \infty]{P} d_{crit}(\xi^{i*}).$$

We have that $\mathbb{P}^{\mathcal{F}_{t_k}^i} \{R_0(d, \xi^i) < 1\} = \mathbb{P}^{\mathcal{F}_{t_k}^i} \{d_{crit}(\xi^i) < d\}$ (see Web Appendix B). Thus $d_{opt}^{t_k}$ is the ω -quantile of the posterior distribution of $d_{crit}(\xi^i)$ which can be computed by MCMC using formula (6).

Figure 1 displays a flow chart for the control strategy. Time after time, observations of CD4 counts and viral loads are collected and an optimal dose given the reaction to the previous doses is provided. Optimal dose greater than d_{max} make the patient leave the control strategy because of treatment failure.

4. Dose Effect and Predictive Ability in the ALBI Trial

4.1 The ALBI Trial

We use the data of the ALBI ANRS 070 trial (Molina et al., 1999). In this three arms controlled trial, 151 antiretroviral-naïve patients (with viral loads between 10,000 and 100,000 copies/ml and CD4 counts greater than 200 cells/mm³) received 24 weeks of treatment. The 51 patients from arm 1, received the d4T + ddI treatment, that is stavudine (d4T)/didanosine (ddI) whose dosage depended on their weight; d4T 250 mg plus ddI 60 mg for patients less than 60 kg, otherwise d4T 400 mg plus ddI 80 mg. For arm 2, 51 patients received the AZT + 3TC treatment, that is 500 mg of zidovudine (AZT) and 300 mg of lamivudine (3TC). The 49 patients from arm 3, the switch arm, received the d4T + ddI treatment for 12 weeks followed by AZT + 3TC, with the same doses as in arm 1 and 2. Blood samples were collected every 4 weeks until 24 weeks. CD4 counts and viral loads (with a lower quantification limit of 50 copies/ml) were measured. Two patients dropped out before the first blood sample collection time, so 149 patients were available for the analysis. For each patient, changes in dose were either self reported or declared by the clinician. Twenty-seven patients (18%) changed their doses during the trial; some of them even interrupted the treatment during a period of time. Molina et al. (1999) found that the d4T + ddI treatment was significantly better than the AZT + 3TC treatment in an intent-to-treat (ITT) analysis.

4.2 Model For the ALBI Trial

We performed a thorough analysis of the ALBI trial using the “activated cells model,” taking into account all the complexities of the design and the doses actually taken by the patients. Doses were included in the analysis, by considering treatment as a time-varying covariate. We took into account the doses of the two ALBI treatments, denoted $\{d_{11}^i(t), d_{12}^i(t)\}$ and $\{d_{21}^i(t), d_{22}^i(t)\}$. To overcome unities problem, we took as reference dose the median dose for each drug, denoted (d_{11}^*, d_{12}^*) and (d_{21}^*, d_{22}^*) . We constructed an indicator of the doses for treatment $j = 1$ (d4T + ddI) or 2 (AZT + 3TC) as:

$$d_j^i(t) = \frac{1}{2} \left\{ \frac{d_{j1}^i(t)}{d_{j1}^*} + \frac{d_{j2}^i(t)}{d_{j2}^*} \right\}.$$

Moreover, we smoothed the dose indicators by taking the mean of the three doses taken in previous days, as this provided a better fit. Therefore, the dose indicator for treatment j of subject i at time t can be written:

$$\bar{d}_j^i(t) = \{d_j^i(t-2) + d_j^i(t-1) + d_j^i(t)\}/3.$$

As for the choice of ψ in (4), we first tried a linear function, but this achieved a poor fit of the data. We obtained a much better fit with power functions, with a power around 0.1. However, such a simple power function does not have an acceptable shape in the sense that the effect is still high for very low doses (for which we had no data). We tried sigmoid functions as in Shen et al. (2008) but this yielded shapes very similar to the power functions. To overcome this problem, we constructed a function which was a power function in the range where information was available $[0.6; 1]$, connected to

zero by a linear function. Finally, the “pharmaco-dynamical” function for patient i , treatment j at time t was:

$$\psi \{d_j^i(t)\} = \begin{cases} d_j^i(t)^{\kappa_d} & \text{if } d_j^i(t) > 0.6, \\ (0.6)^{\kappa_d-1} d_j^i(t) & \text{otherwise.} \end{cases} \quad (8)$$

We also tried to introduce the weights of the subjects as explanatory variable but did not find any effect of it; so these results are not shown.

4.3 Elicitation of a Prior, Algorithm, and Posterior

We denote by $\theta = (\phi, \beta, \sigma_\alpha, \sigma_\lambda, \sigma_{\mu_T^*}, \sigma_{VL}, \sigma_{CD4})$, the vector of the parameters. Here, $\phi = (\phi_l, l = 1 \dots 9)$, appearing in (3), is the vector of the nine baseline biological parameters and β the two-dimensional vector of regressors for treatment doses. There are also three standard deviations of random effects (on α , λ and μ_{T^*} as suggested in Guedj et al. (2007)) and the two standard deviations of errors measurements.

Normal independent priors were assumed for the components of ϕ with expectation and variance elicited in accordance with the literature. Essentially, the rule was to define the lowest and the largest values proposed in the literature ϕ_l^{lower} and ϕ_l^{upper} and take the mean and standard deviation as $(\phi_l^{\text{upper}} + \phi_l^{\text{lower}})/2$ and $(\phi_l^{\text{upper}} - \phi_l^{\text{lower}})/4$, respectively. The literature was not precise enough to elicit these bounds for the infectivity parameter $\tilde{\gamma}_0$; we took a very wide range: $\phi_9^{\text{lower}} = \log(10^{-6}) = -13.8$ and $\phi_9^{\text{upper}} = \log(10) = 2.3$. For the effects of the treatment doses, we also assumed normal priors with expectation corresponding to dividing the infectivity by 3 for the standard dose, and standard deviation equals to about the third of the absolute value of this expectation: this expresses the fact that it is unlikely that treatments increase the infectivity while yielding rather weakly informative priors. The priors did not value one treatment over the other. Half-Cauchy priors were taken for the variance of the random effects, as recommended by Gelman (2006). For the variances of the measurement errors, conventional Jeffreys’s improper noninformative priors were taken. Table 2 summarizes the priors for the different parameters.

In view of the complexity of the problem, MCMC methods would be too time consuming and the INLA method (Approximate Bayesian Inference for Latent Gaussian Models; Rue, Martino, and Chopin, 2009) can not be applied to this problem. We turned to a normal approximation of the posterior which is justified by the Bernstein–Von Mises Theorem (Van der Vaart, 2000). Numerically, this amounts to compute the maximum a posteriori estimator (MAP); this is identical to penalized likelihood maximization. Specifically, the function that must be maximized is $pl = L - J(\theta)$, where the penalty term $J(\theta)$ is:

$$J(\theta) = \sum_{j=1}^9 \frac{[\phi_j - \mathbf{E}^0(\phi_j)]^2}{2\text{var}^0(\phi_j)} + \sum_{j=1}^2 \frac{[\beta_j - \mathbf{E}^0(\beta_j)]^2}{2\text{var}^0(\beta_j)} \\ + \log(\sigma_\alpha^2 + s_\alpha^2) + \log(\sigma_{\mu_{T^*}^*}^2 + s_{\mu_{T^*}^*}^2) \\ + \log(\sigma_{VL}) + \log(\sigma_{CD4}),$$

where \mathbf{E}^0 and var^0 stand for the means and variances of the priors. Here, L is the log likelihood which is described in Guedj et al. (2007). The maximization (or rather the

Table 2
Priors and posteriors for the “activated cells model” parameters estimated from ALBI trial data

		Priors	Posteriors	
Biological parameter (Normal priors):				
Parameter	Mean (SD)	References	Mean	SD
$\tilde{\alpha}$	-4.00 (2.00)	Ribeiro et al. (2002)	-3.19	0.14
$\tilde{\mu}_{T^*}$	-0.05 (0.68)	Althaus, De Vos, and De Boer (2009)	-0.52	0.12
		Brandin et al. (2006)		
$\tilde{\lambda}$	2.55 (1.90)	Mohri et al. (1998)	2.52	0.10
$\tilde{\mu}_T$	-2.59 (0.34)	Ribeiro et al. (2002)	-2.57	0.10
$\tilde{\pi}$	4.04 (2.66)	Chen et al. (2007)	2.49	0.54
$\tilde{\rho}$	-4.34 (1.38)	Ribeiro et al. (2002)	-5.13	0.54
$\tilde{\gamma}_0$	-5.76 (4.02)	See Section 4.3	-5.38	0.03
$\tilde{\mu}_Q$	-9.00 (1.00)	Vriskoop et al. (2008)	-11.2	0.99
$\tilde{\mu}_V$	2.90 (0.68)	Ramratnam et al. (1999)	1.70	0.59
Regressors for treatments doses (Normal priors):				
Parameter	Mean (SD)	References	Mean	SD
β_1	-1.10 (0.37)	See Section 4.3	-0.97	0.09
β_2	-1.10 (0.37)	See Section 4.3	-1.03	0.09
Standard deviation for random effects (Half-Cauchy priors):				
Parameter	Median	References	Mean	SD
σ_α	0.53	Guedj et al. (2007)	0.38	0.03
$\sigma_{\mu_{T^*}}$	0.37	Guedj et al. (2007)	0.03	0.01
σ_λ	0.10	Guedj et al. (2007)	0.03	0.01
Standard deviation for error measurement (Jeffrey’s priors):				
Parameter			Mean	SD
σ_{CV}	-	-	0.45	0.01
σ_{CD4}	-	-	0.20	0.01

minimization of $-pl$) can be done using the algorithm described in Guedj et al. (2007), modified to take into account the penalty brought by the prior. This algorithm uses the so-called RVS algorithm which approximates the Hessian using first derivatives. For penalized likelihood, the way the Hessian is approximated must be modified as described in section 4.5 of Commenges et al. (2006). There is, however, a sign error in the original version. The correct formula for the Hessian of $-pl$ in θ_k (at iteration k) is:

$$G(\theta_k) = \sum_{i=1}^N U_i(\theta_k)U_i^T(\theta_k) - n^{-1}U(\theta_k)U^T(\theta_k) + \frac{\partial^2 J(\theta)}{\partial \theta^2}, \tag{9}$$

where $U(\theta_k) = -\frac{\partial pl}{\partial \theta}|_{\theta_k}$ and the $U_i(\theta_k)$ ’s are the individual gradients ($U(\theta_k) = \sum_{i=1}^N U_i(\theta_k)$). The addition of $\frac{\partial^2 J(\theta)}{\partial \theta^2}$ (here a diagonal matrix with positive diagonal) generally improves the condition number of the approximate Hessian. $G(\theta_k)$ is close to $-\frac{\partial^2 pl}{\partial \theta^2}(\theta_k)$ near the maximum. It is generally easy to compute $\frac{\partial^2 J(\theta)}{\partial \theta^2}$.

We evaluated κ_d in (8) by profile likelihood, leading to $\kappa_d = 0.13$. For other parameters, the posterior was approximated by a normal distribution with expectation given by the MAP and variance given by the inverse of the approximation of the Hessian of $-pl$ (9). We present their means and standard deviations of posteriors in Table 2.

Finally, we checked reproducibility by starting from 10 different initial values drawn at random at one standard

deviation from the mean of the prior. The algorithm converged toward the same region obtaining the same two significant digits for the standard deviations of measurement errors and main random effect, one significant digit for most other parameters, except $\tilde{\pi}$, $\tilde{\mu}_Q$, and $\tilde{\mu}_P$, for which we only had the order of magnitude. This is probably due to lack of practical identifiability: actually, predictions are not impacted by this relative lack of precision. In Web Appendix C, we show that the fits for three stopping points with rather different values of $\tilde{\pi}$, $\tilde{\mu}_Q$, and $\tilde{\mu}_P$ are nearly indistinguishable. Because we use the model essentially for prediction, the problem is not too severe.

4.4 Results: Predictive Ability for Treatment Change

The ALBI design offers a very good opportunity to assess the predictive ability of our model, especially thanks to the switch arm. We can use the first period to assess the random effects for these patients and then predict how they react to the next treatment. We selected patients from the switch arm in an objective manner: we took them at the quartile values of the distribution of the viral load at the end of the study. We computed the MAP estimates removing these patients. Then, we estimated the parameters with random effects thanks to a parametric empirical Bayes (PEB) procedure (Kass and Steffey, 1989) based on the first 12 weeks; viral loads and CD4 counts were computed for the value of the parameters using MAP estimates for fixed effects and PEB for parameters with random effects, taking into account the

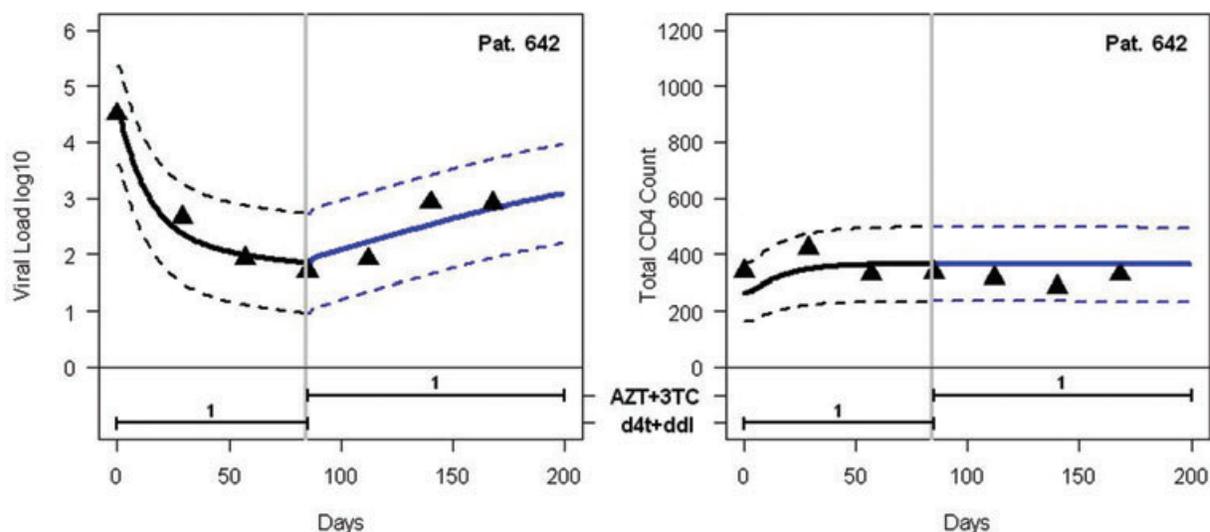


Figure 2. Viral load (log₁₀ copies/ml) et CD4 count (cells/L) predictions for patient 642 from the switch arm who took full dose of d4T + ddI then switched after 85 days to full dose AZT + 3TC. ▲ are observations. Dashed lines represent 95% “measurement error predictive interval.” Vertical line materializes when the treatment switch occurred. Left side: fit; right side: predictions. Treatment adherence is presented on the lower part of the graph.

adherence for the remaining 12 weeks. Using notations in (5), we also computed a 95% “measurement error predictive interval” as $[\hat{Y}_{j1}^i \pm 1.96\sigma_{VL}]$ and $[\hat{Y}_{j2}^i \pm 1.96\sigma_{CD4}]$.

For space concern, we only present the median patient, but predictions for Q_1 and Q_3 patients at quartiles are available in Web Appendix D, others on request. Figure 2 presents the viral load and the total CD4 count with a fit in the first 12 weeks and predictions for the last 12 weeks after the treatment switch. One can see that we are able to predict quite well the viral load rebound after the treatment change (because the AZT + 3TC treatment is less efficient).

4.5 Results: Predictive Ability for Dose Change

We selected analyzed patients in an objective manner among the 27 patients with dose changes: we took those at each quartile values in function of the distribution of the viral load at the end of the study. We estimated again all the parameters excluding this new subset of patients and did step-by-step prediction. Each time we had an additional observation, we updated the knowledge about random effect and computed the predicted value taking into account the reported adherence.

For sake of illustration, we present a patient who particularly well filled his adherence information and who had a large range of behaviors (drug dose reduction, stop, switch...). Figure 3 presents predicted of viral loads and total CD4 counts for each time taking into account the increase of information in time. The 95% “measurement error predictive interval” is represented. Predictions taking into account drugs doses are much better than the ITT fits. This can be quantified by computing the mean square error (MSE) over all the patients with dose changes. The MSE for the viral load is 0.23 for the prediction versus 0.61 for the ITT fit. For the CD4 counts, the MSE are 7768 versus 7962. The predictions together with the MSE for selected patients are available in Web Appendix E.

5. Simulation Study For Dose Optimization

5.1 Principle of the Data Simulation

In this section, we aim at analyzing the efficiency of our drug dose reduction algorithm in term of precision and safety (recommended doses should be higher than critical doses). Data were simulated using the MAP estimates from the analysis of the ALBI trial (Table 2). We simulated samples of 100 patients (n). Parameters with random effects, α , μ_{T^*} and λ were drawn from the multinormal Gaussian posterior for every patient. Fixed parameters were drawn with the constraint $(\theta - \hat{\theta})^T G(\hat{\theta})(\theta - \hat{\theta}) = 1$ and were the same for all the patients of a sample. We only selected patients who had a baseline R_0 (that is without treatment) higher than 1 and a critical dose lower than the standard dose; thus, we excluded long-term nonprogressors and only kept patients needing a dose reduction. We assumed a plausible observation schedule: observation times were at 0, 7, and 14 days whereas dose readjustments were performed every 15 days. The initial dose was taken as the reference, $d_0 = 1$. We simulated the viral loads and CD4 counts with measurement errors variances equal to $(0.45)^2$ and $(0.2)^2$; Table 2). Finally, as ALBI treatments are under-efficient compared to those given in 2011, we took a treatment effect equal to twice the MAP estimate of treatment effect for d4T + ddI (β_2), that is $\beta = -2.06$.

5.2 Illustration of the Monitoring Method

We ran the control strategy algorithm described in Section 3.2 with a burn-in phase consisting in a 100,000 updates, a sampling phase of 50,000 updates, $d_{\max} = 2$ and $\omega = 90\%$. All the nine biological parameters were updated even if significant updates were only noticeable for parameters with random effects. It never recommended a dose greater than d_{\max} .

We first present the results for a particular patient drawn at random as visual illustration (10 updates were performed). His critical dose was 0.306. We can see on Figure 4 (left) that at the first readjustment time we could advise a reduction by

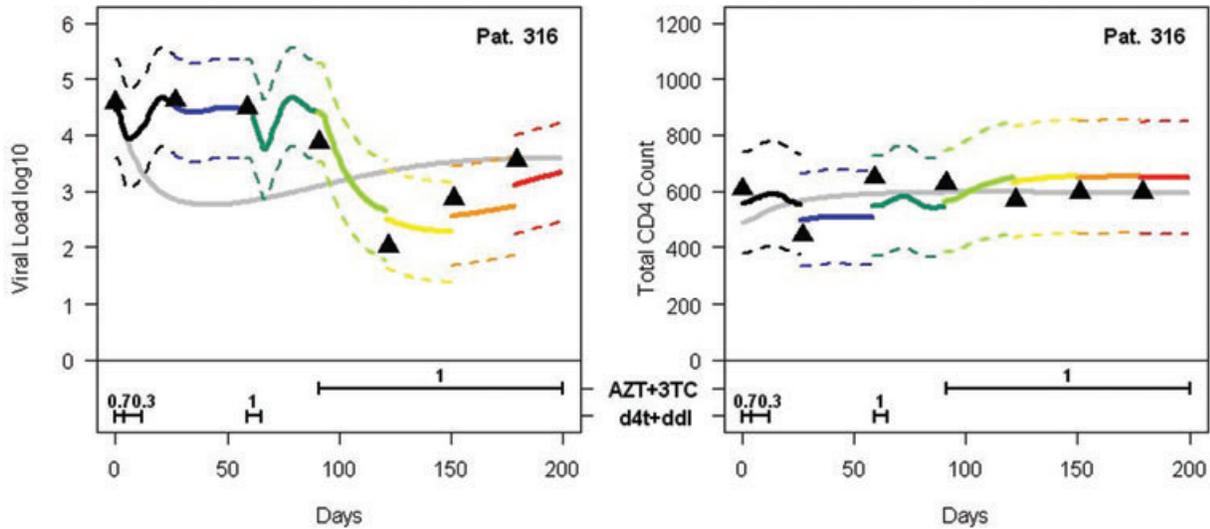


Figure 3. Viral load (log10 copies/ml) et CD4 count (cells/L) predictions for patient 316 who took 0.7% of d4T + ddI during 4 days, then decreased by about two during 8 days and then stopped until day 59. He then took again his treatment until day 65 but with reference posology (probably because he gained weight during the first treatment part and then passed the threshold of 60 kg). He stopped again the treatment and then from day 91 to the end of the study, he switched to AZT + 3TC treatment. ▲ are observations. Each color represents the predictions knowing the information up to the previous observation time. Dashed lines represent 95% “measurement error predictive interval.” Treatments are presented on the lower part of the graph. The gray plain line presents the fit without taking into account adherence.

more than 55% of the dose given to the patient. It is noticeable that the final dose was close to the critical dose. Moreover, this result was obtained while remaining above the critical dose during the tuning period. Furthermore, we can check on Figure 4 (right) that neither the viral load nor the CD4 count are impacted by drug reduction.

Table 3 presents the MSEs between the optimal dose and the critical dose for each readjustment time computed as:

$$MSE(t_k) = \frac{1}{n} \sum_{i=1}^n \{d_{opt}^{t_k}(i) - d_{crit}(\xi^{i*})\}^2,$$

together with the percentage of doses lower than the critical dose for the patient and the percentage of doses greater than d_{max} . In a well-specified model, the MSE decreases while information about the patient increases and we effectively control the probability to give a dose below the critical dose.

5.3 Sensitivity Analysis and A Safer Monitoring Method

We performed a sensitivity analysis to evaluate the robustness of our control strategy in misspecified models. First, we investigated the effect of priors misspecification. In the same fashion as in part 5.1, we simulated samples of 100 patients with parameters such that $(\theta - \hat{\theta})^T G(\hat{\theta})(\theta - \hat{\theta}) = k$, $k = 1.5, 2, 2.5, 3$, and 5. For each k , 10 different samples were taken to explore different profiles of prior misspecification. This extensive simulation study was made possible by using parallel computing. Results are presented in Table 3. MSE were good up to $k = 2.5$ but the percentage of $d_{opt}^{t_k} < d_{crit}$ reached 31% in average at the fifth readjustment. To fix this problem, we propose to add a safety margin to the readjusted dose. We focused on $k = 2$ because priors should not be too badly specified if patients come

from the same population as in ALBI trial. The value of the safety margin was determined such that our key proposition $\forall k \in 0..5, \mathbb{P}^{\mathcal{F}^i}_{t_k} [R_0\{d_{opt}^{t_k}(\omega), \xi^i\} < 1] < 10\%$ hold in the worst simulated case for $k = 2$. We ran again the control procedure, applying the safety margin, over a new validation sample of 1000 patients with misspecified priors. Table 3 shows that MSE are increased but are still reasonable while, as expected, the percentage of $d_{opt}^{t_k} < d_{crit}$ is smaller than 10%.

To finish, we tried a misspecified model. Data were generated from a “productive cell model” (Drylewicz et al., 2010) which has five compartments instead of four. To do so, we took a “productive cell model” with the same equilibrium state as in the “activated cell model.” Productive cell production rate was fixed at -0.12 , the estimates found by Drylewicz et al. (2010). For all the readjustment times, the percentage of optimal doses lower than the critical dose was always lower than 10%. MSE ranged from 1.21 at time t_1 to 0.12 at time t_5 .

6. Discussion

In this article, we have proposed a method of dose individualization of HAART for HIV infected patients. We have attempted to demonstrate its feasibility. First, a pharmacodynamic model was developed for the ALBI trial and we showed that it could be applied to real data. The model fitted the viral loads and CD4 counts data quite well, providing an in vivo estimation of the treatment efficacy. More importantly, the model had good prediction abilities. Then, the simulation study showed that when information is rich enough, the critical dose can be found in a rather small number of readjustments, while staying most of the time above the critical dose.

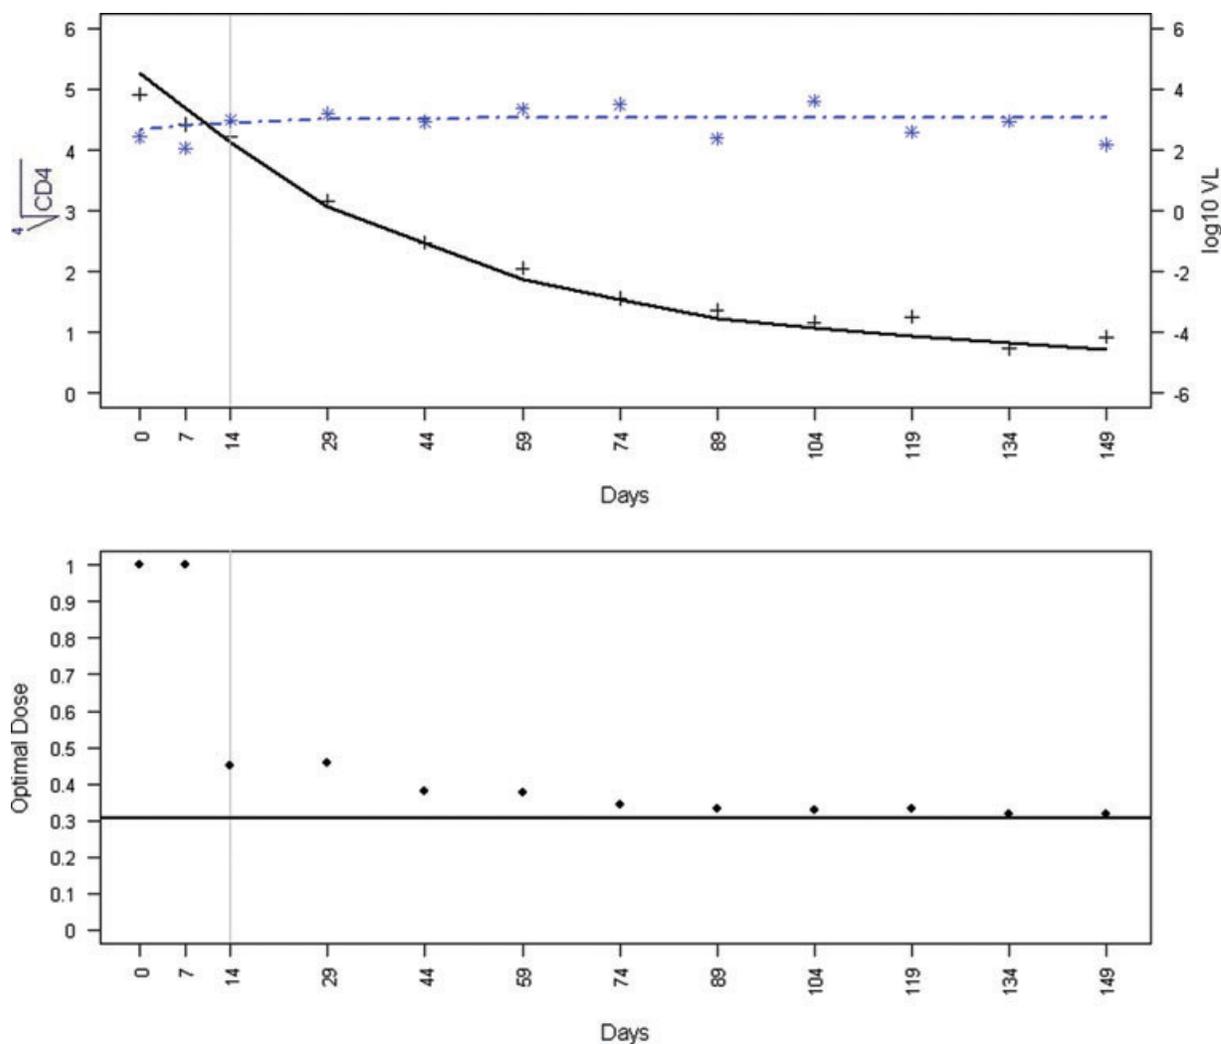


Figure 4. Simulation of 10 doses readjustments spaced by 15 days after three observation times: (up) Viral load (+) and CD4 count (*) simulated observations (down) Dose readjustment simulation, $d_{\text{opt}}^{t,k}$ become closer by above to $d_{\text{crit}} = 0.306$ (horizontal line).

An issue is the possible lack of information due to left-censoring when the treatment is very effective. Then, Doob's consistency theorem invoked in Section 3.2, would not apply. In that case, more intensive schedules and/or the measurements of more compartments would be needed. Other concerns could arise from model misspecification. A misspecification could lead to biased results and wrong recommendations. We have made a sensitivity analysis to test the robustness of the procedure to misspecification of the sampling distribution and to misspecification of the prior (which happens if the new patients are different from those used for defining the prior). In both cases, the convergence toward the critical dose is slower but the procedure, which can be improved by using a safety margin, generally works.

A major issue is development of drug resistance. If this happens during the monitoring, then viral load will raise leading to higher recommended dose until we reach the maximum admissible dose. This is a failure of the treatment and another treatment has to be prescribed. The search of the optimal

dose is relevant in a time period where no major resistance mutation develops.

We conclude that the use of such an adaptive scheme is to be further tested on more informative data (especially with well-documented adherence); however, this work shows that dose individualization is possible. The further step would be to validate the procedure in a clinical trial before using it in clinical practice. Such a clinical trial would typically have two arms, one in which the proposed procedure of dose monitoring would be applied, the other in which standard clinical practice would be applied; the endpoints would be the doses given at the end of the trial, measures of adverse effects and the proportion of virological failures.

7. Supplementary Materials

Web Appendices referenced in Sections 2.1, 3.2, 4.3, 4.4, and 4.5 are available with this article at the Biometrics website on Wiley Online Library.

Table 3
Analysis of sensitivity to prior misspecification: evaluation of the convergence of $(d_{opt}^k)_{k=0...5}$ toward d_{crit}

	t_0	t_1	t_2	t_3	t_4	t_5
	MSE(d_{opt}^k ; % $d_{opt}^k < d_{crit} - \% d_{opt}^k > d_{max}$)					
Well specification	0.62 (0%-0%)	0.04 (0%-0%)	0.03 (0%-0%)	0.01 (1%-0%)	0.005 (4%-0%)	0.001 (9%-0%)
Prior Misspecification						
$k = 1.5$, Act. T cell	0.59 (0%-0%)	0.04 (9%-0%)	0.03 (9%-0%)	0.01 (14%-0%)	0.01 (20%-0%)	0.002 (25%-0%)
$k = 2$, Act. T cell	0.66 (0%-0%)	0.09 (15%-0.5%)	0.04 (14%-0%)	0.02 (19%-0%)	0.02 (27%-0%)	0.01 (31%-0%)
$k = 2.5$, Act. T cell	0.69 (0%-0%)	0.13 (22%-1.5%)	0.06 (20%-0.1%)	0.03 (28%-0%)	0.02 (37%-0%)	0.02 (40%-0%)
$k = 3$, Act. T cell	0.67 (0%-0%)	0.24 (39%-4%)	0.18 (37%-1.3%)	0.11 (42%-0%)	0.08 (45%-0%)	0.08 (47%-0%)
$k = 5$, Act. T cell	0.74 (0%-0%)	0.71 (41%-16%)	0.74 (40%-15%)	0.61 (40%-12%)	0.56 (39%-11%)	0.55 (37%-10%)
A safer monitoring method : "safety margin" = 0.07						
$k = 2$, Act. T cell	0.71 (0%-0%)	0.18 (3%-0%)	0.09 (2%-0.9%)	0.07 (2%-0%)	0.05 (1%-0%)	0.04 (0.2%-0%)

ACKNOWLEDGEMENTS

The authors would like to thank the investigators of the ALBI ANRS 070 trial, particularly J. M. Molina (principal investigator), G. Chêne (methodologist) and V. Journot for data management help. Parallel computing was used thanks to the computing facilities MCIA (Mésocentre de Calcul Intensif Aquitain) of the Université de Bordeaux and of the Université de Pau et des Pays de l'Adour.

REFERENCES

Althaus, C., De Vos, A., and De Boer, R. (2009). Reassessing the human immunodeficiency virus type 1 life cycle through age-structured modeling: Life span of infected cells, viral generation time, and basic reproductive number, R0. *Journal of Virology* **83**, 7659–7667.

Ananworanich, J., Gayet-Ageron, A., Le Braz, M., Prasithsirikul, W., Chetchotisakd, P., Kiertiburanakul, S., Munsakul, W., Raksakulkarn, P., Tansuphasawasdikul, S., Sirivichayakul, S., Cavassini, M., Karrer, U., Genné, D., Nuesch, R., and Vernazza, P. (2006). CD4-guided scheduled treatment interruptions compared with continuous therapy for patients infected with HIV-1: Results of the Staccato randomised trial. *The Lancet* **368**, 459–465.

Berger, J. (1985). *Statistical Decision Theory and Bayesian Analysis*. Springer-Verlag New York.

Brandin, E., Thorstenson, R., Bonhoeffer, S., and Albert, J. (2006). Rapid viral decay in simian immunodeficiency virus-infected macaques receiving quadruple antiretroviral therapy. *Journal of Virology* **80**, 9861–9864.

Carr, A. and Cooper, D. (2000). Adverse effects of antiretroviral therapy. *The Lancet* **356**, 1423–1430.

Chen, H., Di Mascio, M., Perelson, A., Ho, D., and Zhang, L. (2007). Determination of virus burst size in vivo using a single-cycle SIV in rhesus macaques. *Proceedings of the National Academy of Sciences* **104**, 19079–19084.

Chomont, N., El-Far, M., Ancuta, P., Trautmann, L., Procopio, F., Yassine-Diab, B., Boucher, G., Boulassel, M., Ghattas, G., Brechley, J., Schacker, T., Hill, B., Douek, D., Routy, J., Haddad, E., and Sékaly, R. (2009). Hiv reservoir size and persistence are driven by t cell survival and homeostatic proliferation. *Nature Medicine* **15**, 893–900.

Commenges, D., Jacqmin-Gadda, H., Proust, C., and Guedj, J. (2006). A newton-like algorithm for likelihood maximization: The robust-variance scoring algorithm. *Arxiv math/0610402*.

Commenges, D., Sayyareh, A., Letenneur, L., Guedj, J., and Bar-Hen, A. (2008). Estimating a difference of Kullback–Leibler risks using a normalized difference of aic. *The Annals of Applied Statistics* **2**, 1123–1142.

Danel, C., Moh, R., Minga, A., Anzian, A., Ba-Gomis, O., Kanga, C., Nzunetu, G., Gabillard, D., Rouet, F., Sorho, S., Chaix, M., Eholié, S., Menan, H., Sauvageot, D., Bissagnene, E., Salamon, R., and Anglaret, X. (2006). Cd4-guided structured antiretroviral treatment interruption strategy in HIV-infected adults in West Africa (trivacan anrs 1269 trial): A randomised trial. *The Lancet* **367**, 1981–1989.

Drylewicz, J., Guedj, J., Commenges, D., and Thiébaud, R. (2010). Modeling the dynamics of biomarkers during primary HIV infection taking into account the uncertainty of infection date. *The Annals of Applied Statistics* **4**, 1847–1870.

Egger, M., May, M., Chêne, G., Phillips, A., Ledergerber, B., Dabis, F., Costagliola, D., Monforte, A., de Wolf, F., Reiss, P., Lundgren, J., Justice, A., Schlomo, S., Lepout, C., Hogg, R., Sabin, C., Gill, J., Salzberger, B., and Sterne, J. (2002). Prognosis of HIV-1-infected patients starting highly active antiretroviral therapy:

- A collaborative analysis of prospective studies*. *The Lancet* **360**, 119–129.
- Gelman, A. (2006). Prior distributions for variance parameters in hierarchical models. *Bayesian Analysis* **1**, 515–533.
- Girard, P., Cabié, A., Michelet, C., Verdon, R., Katlama, C., Mercié, P., Morand-Joubert, L., Pétour, P., Monchecourt, F., Chêne, G., and Trylesinski, A. (2009). A randomized trial of two-drug versus three-drug tenofovir-containing maintenance regimens in virologically controlled HIV-1 patients. *Journal of Antimicrobial Chemotherapy* **64**, 126–134.
- Gran, J., Wasmuth, L., Amundsen, E., Lindqvist, B., and Aalen, O. (2008). Growth rates in epidemic models: Application to a model for HIV/aids progression. *Statistics in Medicine* **27**, 4817–4834.
- Guedj, J., Thiébaud, R., and Commenges, D. (2007). Maximum likelihood estimation in dynamical models of HIV. *Biometrics* **63**, 1198–1206.
- Henderson, R., Ansell, P., and Alshibani, D. (2010). Regret-regression for optimal dynamic treatment regimes. *Biometrics* **66**, 1192–1201.
- Ho, D., Neumann, A., Perelson, A., Chen, W., Leonard, J., Markowitz, M., (1995). Rapid turnover of plasma virions and cd4 lymphocytes in HIV-1 infection. *Nature* **373**, 123–126.
- Huang, Y., Liu, D., and Wu, H. (2006). Hierarchical Bayesian methods for estimation of parameters in a longitudinal HIV dynamic system. *Biometrics* **62**, 413–423.
- Huang, Y., Wu, H., and Acosta, E. (2010). Hierarchical Bayesian inference for HIV dynamic differential equation models incorporating multiple treatment factors. *Biometrical Journal* **52**, 470–486.
- Jacqmin-Gadda, H., Thiebaud, R., Chene, G., and Commenges, D. (2000). Analysis of left-censored longitudinal data with application to viral load in HIV infection. *Biostatistics* **1**, 355–368.
- Kass, R. and Steffey, D. (1989). Approximate Bayesian inference in conditionally independent hierarchical models (parametric empirical bayes models). *Journal of the American Statistical Association* **87**, 717–726.
- Kirschner, D., Lenhart, S., and Serbin, S. (1997). Optimal control of the chemotherapy of HIV. *Journal of Mathematical Biology* **35**, 775–792.
- Lavori, P. and Dawson, R. (2008). Adaptive treatment strategies in chronic disease. *Annual Review of Medicine* **59**, 443–453.
- Lundgren, J., Babiker, A., El-Sadr, W., Emery, S., Grund, B., Neaton, J., Neuhaus, J., and Phillips, A. (2008). Inferior clinical outcome of the cd4 (+) cell count guided antiretroviral treatment interruption strategy in the smart study: Role of cd4 (+) cell counts and hiv rna levels during follow-up. *Journal of Infectious Diseases* **197**, 1145–1155.
- Meynard, J., Morand-Joubert, L., Lacombe, K., Poirier, J., Slama, L., Valantin, M., and Girard, P. (2010). Impact of reduced dosing of lopinavir/ritonavir in virologically controlled hiv-infected patients: The kaleidose trial. *Journal of Antimicrobial Chemotherapy* **65**, 125–128.
- Milinkovic, A., Martinez, E., López, S., de Lazzari, E., Miró, O., Vidal, S., Blanco, J., Garrabou, G., Laguno, M., Arnaiz, J., Leon, A., Larrousse, M., Lonca, M., Mallolas, J., and Gatell, J. (2007). The impact of reducing stavudine dose versus switching to tenofovir on plasma lipids, body composition and mitochondrial function in HIV-infected patients. *Antiviral Therapy* **12**, 407–415.
- Mohri, H., Bonhoeffer, S., Monard, S., Perelson, A., and Ho, D. (1998). Rapid turnover of T lymphocytes in SIV-infected rhesus macaques. *Science* **279**, 1223–1227.
- Molina, J., Chêne, G., Ferchal, F., Journot, V., Pellegrin, I., Sombardier, M., Rancinan, C., Cotte, L., Madeline, I., Debord, T., and Decazes, J. (1999). The ALBI trial: A randomized controlled trial comparing stavudine plus didanosine with zidovudine plus lamivudine and a regimen alternating both combinations in previously untreated patients infected with human immunodeficiency virus. *The Journal of Infectious Diseases* **180**, 351–358.
- Moodie, E., Richardson, T., and Stephens, D. (2007). Demystifying optimal dynamic treatment regimes. *Biometrics* **63**, 447–455.
- Murphy, S. (2003). Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **65**, 331–355.
- Murphy, S. and McKay, J. (2004). Adaptive treatment strategies: An emerging approach for improving treatment effectiveness. *Clinical Science* **12**, 7–13.
- Perelson, A., Neumann, A., Markowitz, M., Leonard, J., and Ho, D. (1996). HIV-1 dynamics in vivo: virion clearance rate, infected cell life-span, and viral generation time. *Science* **271**, 1582–1586.
- Putter, H., Heisterkamp, S., Lange, J., and De Wolf, F. (2002). A Bayesian approach to parameter estimation in HIV dynamical models. *Statistics in Medicine* **21**, 2199–2214.
- Ramratnam, B., Bonhoeffer, S., Binley, J., Hurley, A., Zhang, L., Mittler, J., Markowitz, M., Moore, J., Perelson, A., and Ho, D. (1999). Rapid production and clearance of HIV-1 and hepatitis C virus assessed by large volume plasma apheresis. *The Lancet* **354**, 1782–1785.
- Ribeiro, R., Mohri, H., Ho, D., and Perelson, A. (2002). In vivo dynamics of t cell activation, proliferation, and death in hiv-1 infection: Why are cd4+ but not cd8+ t cells depleted? *Proceedings of the National Academy of Sciences of the United States of America* **99**, 15572–15577.
- Rue, H., Martino, S., and Chopin, N. (2009). Approximate Bayesian inference for latent gaussian models by using integrated nested laplace approximations. *Journal of the royal statistical society: Series B (Statistical Methodology)* **71**, 319–392.
- Sage, A. and White, C. (1977). *Optimum Systems Control*. NJ: Prentice-Hall Englewood Cliffs.
- Sánchez-Conde, M., de Mendoza, C., Jiménez-Nacher, I., Barreiro, P., Gonzalez-Lahoz, J., and Soriano, V. (2005). Reductions in stavudine dose might ameliorate mitochondrial-associated complications without compromising antiviral activity. *HIV Clinical Trials* **6**, 197–202.
- Shen, L., Peterson, S., Sedaghat, A., McMahon, M., Callender, M., Zhang, H., Zhou, Y., Pitt, E., Anderson, K., Acosta, E., and Siliciano, R. (2008). Dose-response curve slope sets class-specific limits on inhibitory potential of anti-hiv drugs. *Nature Medicine* **14**, 762–766.
- Thiébaud, R., Jacqmin-Gadda, H., Lepout, C., Katlama, C., Costagliola, D., Moing, V., Morlat, P., Chêne, G., and Group, A. (2003). Bivariate longitudinal model for the analysis of the evolution of HIV RNA and CD4 cell count in HIV infection taking into account left censoring of HIV RNA measures. *Journal of Biopharmaceutical Statistics* **13**, 271–282.
- Van der Vaart, A. (2000). *Asymptotic Statistics*. Cambridge University Press, UK.
- Vrisekoop, N., Den Braber, I., De Boer, A., Ruiter, A., Ackermans, M., Van Der Crabben, S., Schrijver, E., Spierenburg, G., Sauerwein, H., Hazenberg, M., De Boer, R., Miedema, F., Borghans, J., and Tesselaar, K. (2008). Sparse production but preferential incorporation of recently produced naive T cells in the human peripheral pool. *Proceedings of the National Academy of Sciences* **105**, 6115–6120.
- Wei, X., Ghosh, S., Taylor, M., Johnson, V., Emini, E., Deutsch, P., Lifson, J., Bonhoeffer, S., Nowak, M., and Hahn, B. (1995). Viral dynamics in human immunodeficiency virus type 1 infection. *Nature* **373**, 117–122.
- Wu, H. (2005). Statistical methods for HIV dynamic studies in AIDS clinical trials. *Statistical Methods in Medical Research* **14**, 171–192.