Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

A computational framework for posterior probability vectors over discrete pharmacokinetic states

Jinming Zhang
Shanghai Chest Hospital, Shanghai Jiao Tong University School of Medicine
Zheng Jiao
zjiao@sjtu.edu.cn

Shanghai Chest Hospital, Shanghai Jiao Tong University School of Medicine

Research Article

Keywords: Bayesian inference, Discrete pharmacokinetic states, Posterior probabilities, Laplace approximation, Therapeutic drug monitoring, Medication adherence

Posted Date: November 24th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11757027/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

A computational framework for posterior probability vectors over discrete pharmacokinetic states

Jinming Zhang, Zheng Jiao

Shanghai Chest Hospital, Shanghai Jiao Tong University School of Medicine, Shanghai, China

Corresponding author: Zheng Jiao

Email: zjiao@sjtu.edu.cn

Telephone: 86 18930841095

Funding

Partial financial support was received from the National Natural Science Foundation of China (82073933).

Abstract

Aim In therapeutic drug monitoring, clinicians need to identify potential causes of unexpected concentrations, yet limited concentration observations may not distinguish competing explanations. Building on existing methods, this study aimed to formulate such assessments as calculating posterior probabilities of discrete pharmacokinetic states with explicit inputs and outputs, establish a reusable computational workflow, and evaluate it by distinguishing among dosing adherence histories.

Methods Using a Bayesian approach, the Python package uses Laplace approximation to calculate marginal likelihoods and combines them with prior probabilities to obtain posterior probabilities of candidate states. Workflow evaluation used published oxcarbazepine and tacrolimus models and compared posterior probabilities with the NONMEM Laplace implementation, external adaptive quadrature, and an existing virtual patient approach. Sensitivity analyses examined computational consistency, posterior probability allocation, and discrimination among candidate states across input settings.

Results Python and NONMEM Laplace implementations yielded highly consistent posterior probabilities. Compared with external adaptive quadrature, posterior distributions were close in most scenarios; the most probable state differed between methods in some tacrolimus scenarios where posterior probabilities were distributed nearly evenly across candidate states. In comparable oxcarbazepine scenarios, the Python implementation and virtual patient approach showed generally consistent main trends in posterior probabilities. Input variations could alter posterior probability allocation and the most probable state, while different dosing histories remained difficult to distinguish under some settings.

Conclusion This study provides a reusable framework and Python implementation for calculating posterior probabilities of candidate explanations in therapeutic drug monitoring. Its clinical utility requires validation in studies with independent state labels.

Keywords

Bayesian inference; Discrete pharmacokinetic states; Posterior probabilities; Laplace approximation; Therapeutic drug monitoring; Medication adherence

Introduction

Therapeutic drug monitoring (TDM) supports individualized dosing by measuring drug concentrations in patients and interpreting the results alongside clinical information [1,2]. When concentrations deviate from expectations under the current dosing regimen, clinicians need to identify potential causes of the deviation. For instance, a low concentration may be associated with either recently missed doses or the patient's physiological characteristics [3,4]. Limited concentration observations may not suffice to distinguish among these explanations, introducing uncertainty into the clinical assessment of unexpected drug concentrations.

Clinicians can use population pharmacokinetic models to link dosing histories and patient characteristics with changes in drug concentrations [5,6]. The challenge is to quantify how likely different causes are based on the model and observed data, and determine how these results can inform clinical judgment. A Bayesian approach offers a way to address this by calculating posterior probabilities across candidate explanations, thereby quantifying the support provided by available information for each explanation [7,8].

Specifically, the Bayesian approach multiplies the likelihood of observed data under each candidate explanation by its corresponding prior probability, and then normalizes the products across candidate explanations to obtain their posterior probabilities. Existing studies on dosing adherence have applied this approach to calculate the posterior probabilities of candidate dosing histories based on pharmacokinetic

models and observed concentrations [3,7,8]. Among these, the virtual patient approach by Liu et al. uses Monte Carlo simulation to estimate the concentration distribution under each candidate history, thereby obtaining the likelihood required for posterior calculation [8]. As a standard software package for population pharmacokinetic modeling, NONMEM provides multiple methods for numerically approximating integrals over random effects to compute the marginal likelihood of individual observations [9]. NONMEM also supports specifying mixture models via $MIX, using different submodels to describe distinct subpopulations within a population. For a given individual, the software computes the likelihood of the observed data under each submodel and combines it with the corresponding subpopulation prior probability to obtain the posterior probability of the individual belonging to each subpopulation [9,10].

Applying these methods to different pharmacokinetic problems requires adapting to variations in models, observation records, and candidate explanations. The existing implementation of the virtual patient approach primarily targets the included pharmacokinetic models, using a single concentration observation to distinguish combinations of taken and missed doses at scheduled dosing times [8]. Incorporating new pharmacokinetic models, combining multiple observations, or comparing candidate histories that differ in dosing times and doses requires corresponding extensions to the simulation and likelihood calculation. Although NONMEM allows flexible definition of submodels, incorporating different candidate explanations into the calculation still requires writing control streams, organizing data, and specifying the corresponding submodels and prior probabilities for each specific problem [9]. Therefore, it is necessary to formulate the comparison of such candidate explanations as a problem of calculating posterior probabilities across discrete pharmacokinetic states with well-defined inputs and outputs, and to establish a computational workflow capable of adapting to different models, observation records, and candidate definitions. At the same time, its numerical reliability and the impact of different input information on posterior probabilities need to be systematically evaluated.

This study aimed to establish a pharmacokinetic model-based Bayesian computational framework for calculating the posterior probabilities of individual discrete pharmacokinetic states, and to implement it as a Python package. Using scenarios for discriminating dosing adherence histories as a specific application, we evaluated this computational workflow under oxcarbazepine and tacrolimus models. To verify numerical reliability, we compared the posterior probabilities obtained from the Python package with those from a NONMEM implementation and an alternative numerical integration method, and, in oxcarbazepine scenarios with matched settings, with an existing virtual patient approach [8]. Subsequently, sensitivity analyses were performed to evaluate the applicability of the framework across varied input settings, examining how input variations affect posterior probability allocation and discrimination among candidate states, and testing consistency across the computational approaches in these scenarios.

Theoretical

Bayesian posterior inference

To compare alternative explanations for the concentration observations from the same individual, we define these explanations as discrete candidate states and calculate the posterior probability of each state by combining prior information with observational evidence.

For individual i, let $\mathbf{y}_i$ denote the vector of concentration observations to be explained, and let K_i candidate states form a finite set:

$$\mathcal{S}_i = \{1, \dots, K_i\}. \tag{1}$$

A discrete variable S_i takes a single value in this set, such that the candidate states represent mutually exclusive explanations in a given comparison. These states explain the same set of concentration

observations $\mathbf{y}_i$, whereas their corresponding input information may differ. Each state s corresponds to an input record z_{is} containing covariates, dosing history, sampling times, and state-specific model inputs. Letting $z_i = \{z_{is}: s \in \mathcal{S}_i\}$, the computational output is a posterior probability vector ordered by the candidate states:

$$\mathbf{p}_i = (p_{i1}, \ldots, p_{iK_i})^{\mathrm{T}}, \qquad p_{is} = P(S_i = s \mid \mathbf{y}_i, z_i). \tag{2}$$

The calculation requires first specifying prior probabilities for the candidate states, where the definition of candidate states and their prior probabilities should be grounded in literature, clinical evidence, or study design. For individual i, let $\pi_s = P(S_i = s \mid z_i)$ denote the probability assigned to state s prior to incorporating the concentration observations to be explained, satisfying:

$$\pi_s \geq 0, \qquad \sum_{s \in \mathcal{S}_i} \pi_s = 1. \tag{3}$$

Subsequently, the concentration observations are incorporated into the calculation through the marginal likelihood of the observed concentrations under each state. Let $m_s(\mathbf{y}_i \mid z_{is})$ denote the marginal likelihood for state s, also referred to as the model evidence. Multiplying this marginal likelihood by the corresponding prior probability and normalizing across the candidate states yields, by Bayes' theorem:

$$p_{is} = \frac{\pi_s \, m_s(\mathbf{y}_i \mid z_{is})}{\sum_{r \in \mathcal{S}_i} \pi_r \; m_r(\mathbf{y}_i \mid z_{ir})}. \tag{4}$$

The denominator in Eq. (4) must be finite and strictly positive; under this condition, nonnegative prior probabilities and model evidence yield nonnegative posterior probabilities that sum to 1.

These posterior probabilities quantify the degree of support for each candidate explanation given the candidate set, input records, prior probabilities, pharmacokinetic model, and variability settings. Altering the membership of the candidate set or the prior probabilities may change the allocation of posterior probabilities, and changing the number of candidate states correspondingly changes the dimension of the posterior probability vector. The following section describes how the marginal likelihood of each state is calculated from the pharmacokinetic model.

Marginal likelihood

To obtain the marginal likelihood in Eq. (4), we first formulate the likelihood of the observations conditional on the random effects, and then integrate this conditional likelihood over the distribution of the random effects. The population parameters and variability settings are treated as given in each calculation.

Let f denote the pharmacokinetic model, $\boldsymbol{\theta}$ the population parameters, and $\boldsymbol{\eta}_i$ the vector of random effects for individual i. For candidate state s, let t_{ijs} denote the sampling time of the jth observation in the input record z_{is}, for which the predicted concentration is:

$$\mu_{ijs}(\boldsymbol{\eta}_i) = f(t_{ijs}, z_{is}, \boldsymbol{\theta}, \boldsymbol{\eta}_i), \tag{5}$$

The predicted concentrations are linked to the observed concentrations through a residual error model. Letting y_{ij} denote the jth observed concentration and ϵ_{ij} the residual error between the observed and predicted concentrations, we write:

$$y_{ij} = \mu_{ijs}(\boldsymbol{\eta}_i) + \epsilon_{ij}. \tag{6}$$

Given a candidate state and the random effects, the residual error distribution defines the observation likelihood $p(\mathbf{y}_i \mid \boldsymbol{\eta}_i, s, z_{is})$ for the entire concentration record, with its normalizing term retained in the calculation. The residual variance may depend on the predicted concentrations according to the selected residual error model.

Multiple observations from the same individual are described by a single joint random-effect vector $\boldsymbol{\eta}_i$. If the model includes interoccasion variability, the random effects for each occasion are incorporated as distinct components of this vector and applied during prediction according to the predefined occasion structure.

Because this observation likelihood remains dependent on the random effects, obtaining the marginal likelihood for each state requires integrating it over the random-effect distribution. Assuming that $\boldsymbol{\eta}_i$ follows a q-dimensional Gaussian distribution with zero mean and a positive-definite covariance matrix $\boldsymbol{\Omega}$, we have:

$$m_s(\mathbf{y}_i \mid z_{is}) = \int_{\mathbb{R}^q} p\,(\mathbf{y}_i \mid \boldsymbol{\eta}_i, s, z_{is})\phi(\boldsymbol{\eta}_i; \mathbf{0}, \boldsymbol{\Omega})\,\mathrm{d}\boldsymbol{\eta}_i. \tag{7}$$

Here, ϕ denotes the normalized probability density of the Gaussian distribution specified above; in a given comparison, the same random-effect distribution is shared across all candidate states. The marginal likelihood obtained from Eq. (7) serves as the model evidence in Eq. (4). In the absence of random effects, it reduces to the observation likelihood evaluated at the fixed predictions for that state.

Laplace approximation

To evaluate the random-effect integral in Eq. (7), we use the Laplace approximation, which approximates the marginal likelihood from the location of the conditional mode and the local curvature around it under each candidate state [11,12,13]. To determine the conditional mode and evaluate the local curvature, we define the negative logarithm of the integrand in Eq. (7) as:

$$\Phi_{is}(\boldsymbol{\eta}_i) = -\log\{p(\mathbf{y}_i \mid \boldsymbol{\eta}_i, s, z_{is})\phi(\boldsymbol{\eta}_i; \mathbf{0}, \boldsymbol{\Omega})\}. \tag{8}$$

Minimizing Φ_{is} yields the conditional mode $\hat{\boldsymbol{\eta}}_{is}$ of the random effects under this state, and evaluating the Hessian matrix $\mathbf{H}_{is}$ of Φ_{is} at this point characterizes its local curvature:

$$\begin{aligned} \hat{\boldsymbol{\eta}}_{is} &= \arg\min_{\boldsymbol{\eta}_i} \Phi_{is}(\boldsymbol{\eta}_i), \\ \mathbf{H}_{is} &= \left.\frac{\partial^2 \Phi_{is}(\boldsymbol{\eta}_i)}{\partial \boldsymbol{\eta}_i\, \partial \boldsymbol{\eta}_i^{\mathrm{T}}}\right|_{\boldsymbol{\eta}_i = \hat{\boldsymbol{\eta}}_{is}}. \end{aligned} \tag{9}$$

At an isolated conditional mode in the interior of the domain, assume that Φ_{is} is twice continuously differentiable in a neighborhood of this point and that $\mathbf{H}_{is}$ is positive definite. Because the first derivative of Φ_{is} vanishes at this point, its local quadratic expansion is:

$$\Phi_{is}(\boldsymbol{\eta}_i) \approx \Phi_{is}(\hat{\boldsymbol{\eta}}_{is}) + \frac{1}{2}(\boldsymbol{\eta}_i - \hat{\boldsymbol{\eta}}_{is})^{\mathrm{T}}\mathbf{H}_{is}(\boldsymbol{\eta}_i - \hat{\boldsymbol{\eta}}_{is}). \tag{10}$$

Using Eq. (10), the integrand in Eq. (7) can be approximated by a Gaussian kernel. Integrating this kernel yields the Laplace approximation to the marginal likelihood of each state:

$$\begin{aligned} m_s(\mathbf{y}_i \mid z_{is}) &\approx \exp[-\Phi_{is}(\hat{\boldsymbol{\eta}}_{is})] \\ &\quad \times \int_{\mathbb{R}^q} \exp\left[-\frac{1}{2}(\boldsymbol{\eta}_i - \hat{\boldsymbol{\eta}}_{is})^{\mathrm{T}}\mathbf{H}_{is}(\boldsymbol{\eta}_i - \hat{\boldsymbol{\eta}}_{is})\right] \mathrm{d}\boldsymbol{\eta}_i \\ &= \exp[-\Phi_{is}(\hat{\boldsymbol{\eta}}_{is})](2\pi)^{q/2}|\mathbf{H}_{is}|^{-1/2}. \end{aligned} \tag{11}$$

This approximation combines the height of the joint density at the mode with the local width characterized by the Hessian matrix. This single-mode approximation is most appropriate when the dominant contribution to the integral arises from an approximately Gaussian neighborhood around the selected mode.

When calculating the marginal likelihood of each state, all state-dependent normalizing terms in the likelihood and random-effects density must be retained; only factors common to all candidate states

cancel out during posterior normalization. These calculations can be performed on the logarithmic scale to mitigate numerical underflow or overflow. Finally, substituting the approximate marginal likelihood of each state into Eq. (4) yields the posterior probability vector under the Laplace approximation.

Alternative computational approaches

Under the same Bayesian updating principle, as alternatives to the primary Python-based Laplace approximation in this study, the posterior probabilities of candidate states can also be computed using different implementation platforms, integral approximation methods, or likelihood estimation pathways.

NONMEM implementation

The core difference between the NONMEM comparator and the primary approach lies in the implementation platform for the Laplace approximation. Through the `$MIX` mixture model mechanism, each candidate state can be defined as a distinct subpopulation within the mixture model and assigned a corresponding prior probability [9,10]. In this study, NONMEM likewise uses the Laplace approximation to compute the marginal likelihood of individual observations under each subpopulation, and combines it with the subpopulation prior probabilities to obtain the posterior mixture probabilities of the individual belonging to each candidate state.

Adaptive Gauss-Hermite quadrature

The core difference between the adaptive Gauss-Hermite quadrature (GHQ) comparator and the primary approach lies in the numerical integration method for the marginal likelihood. For the marginal likelihood integral in Eq. (7), adaptive GHQ places quadrature nodes based on the conditional mode of the random effects and its local curvature, approximating the marginal likelihood of each candidate state through a weighted sum at these nodes [13,14]. The resulting approximate marginal likelihoods of the states are subsequently combined with the state priors as before, while the corresponding target of posterior inference remains unchanged.

Virtual-patient simulation

The core difference between the virtual patient approach and the primary approach lies in the likelihood estimation pathway. Based on a given population pharmacokinetic model and variability settings, this approach uses Monte Carlo sampling to separately simulate the concentration distribution under each candidate dosing history, and estimates the likelihood of the observed concentrations under the corresponding history from the simulated distribution [8]. Once the likelihoods for the candidate histories are obtained, subsequent calculations likewise combine them with the corresponding prior probabilities according to Eq. (4) to yield the posterior probabilities of the states.

Methods

Software and implementations

The primary approach in this study implements Laplace approximation-based posterior probability calculations via an in-house Python package, with a NONMEM mixture model and adaptive GHQ from the external R package `aghq` serving as reference implementations. A supplementary comparison was made between the package results and the Dashboard outputs from the existing virtual patient approach [8]. The software versions, dependencies, computing environments, and detailed implementations used across the analyses are provided in Online Resources 1 and 4. Online Resource 4 also provides the code, input data, saved results, and instructions for reproducing the analyses. Re-running the NONMEM calculations requires a licensed NONMEM installation, and re-running the adaptive GHQ reference calculations requires the corresponding R environment.

Python Laplace implementation

We developed pmx-discrete-posterior (version 0.2.0) in Python to calculate the posterior probabilities of user-defined pharmacometric states. To allow different pharmacokinetic models to share the same workflow for posterior calculation, we divided the software into four components: input interface, model specification, numerical integration, and posterior output.

The input interface receives and validates the model information, individual information, observation history, candidate states, and their prior probabilities. The model specification component supports user-defined pharmacokinetic models and variability settings, and performs concentration predictions and conditional likelihood calculations under each candidate state for given random effects. The numerical integration component incorporates the random-effects distribution to approximate the marginal likelihood of each state via integration, and the posterior output component combines these likelihoods with the corresponding prior probabilities to return the complete posterior probability vector and numerical diagnostic information.

The package uses NumPy and SciPy [15] for numerical calculations. The numerical integration component uses the BFGS optimizer in SciPy to locate the conditional mode of the random effects and computes the Hessian matrix at that point via numerical differentiation to calculate the marginal likelihood according to the Laplace approximation described above.

NONMEM mixture model

In the NONMEM control stream, each candidate dosing history corresponds to a subpopulation in the mixture model. All subpopulations share the same pharmacokinetic model structure and fixed population and variability parameters, with the prior probabilities specified via $MIX. NONMEM uses the LAPLACIAN setting with interaction to approximate the marginal likelihood of each candidate history and combines these likelihoods with the prior probabilities to obtain the posterior mixture probabilities. We extracted the posterior probability of each candidate dosing history from the PMIX column of the .phm file.

R aghq

The adaptive GHQ reference was implemented using the external R package `aghq` [14]. This approach shares the same pharmacokinetic prediction and likelihood functions with the Python Laplace implementation. For each candidate state, Python constructs the log-integrand of the marginal likelihood integral and exposes it to R via a bridge. After completing the adaptive quadrature, `aghq` returns the log marginal likelihood of each candidate state to Python, which then combines these values with the corresponding state prior probabilities and normalizes them to obtain the posterior probabilities.

Models and scenarios

To evaluate the performance of the package in calculating posterior probabilities across different pharmacokinetic models, we selected two population pharmacokinetic models for oxcarbazepine (OXC) and tacrolimus (TAC). A reference scenario was defined for each model to serve as a common starting point for subsequent comparisons of computational approaches and sensitivity analyses.

In the reference scenarios for both models, candidate dosing histories for two prior scheduled doses at 0 h and 12 h were compared on the basis of a single concentration observation at 24 h. Each of the two scheduled doses could be either taken or missed, defining four candidate states: both doses missed (ω_{00}), only the 0 h dose taken (ω_{01}), only the 12 h dose taken (ω_{10}), and both doses taken (ω_{11}). The subscripts are ordered backward from the most recent dose, and each state was assigned a prior probability of 0.25.

One-compartment oxcarbazepine model

We adopted the pediatric population pharmacokinetic model of OXC developed by Lin et al. [16] to predict the concentrations of its active metabolite, the monohydroxy derivative (MHD). This model is a one-compartment model with first-order absorption and first-order elimination, in which apparent clearance varies with body weight.

The model incorporates interindividual variability (IIV) only in apparent clearance using exponential random effects, and uses an additive form for residual unexplained variability (RUV). Reference values for the model parameters and variability are presented in Table 1, along with the corresponding parameter definitions, functional relationships, and distributional assumptions in the table and its footnotes.

For the OXC reference scenario, we used a regimen of 300 mg every 12 h, following the example setting in the adherence study by Liu et al. [8]. The model parameters and variability settings used the reference values in Table 1, with a reference body weight of 25 kg, and the observed MHD concentration at 24 h was set to 12 mg/L. To determine the drug amount in the body at the start of the look-back window, all four candidate histories shared a preceding regular dosing history of 300 mg every 12 h, with the last shared dose at −12 h. Blood concentrations were assumed to have reached steady state under this regular dosing regimen before the start of the look-back window.

Two-compartment tacrolimus model

We adopted the final population pharmacokinetic model of TAC reported by Bergmann et al. [17] to predict TAC concentrations. This model is a two-compartment model with an absorption lag time, first-order absorption, and first-order elimination. Apparent clearance incorporates the effects of CYP3A5 genotype, hematocrit, body weight, and postoperative time, whereas the apparent central volume of distribution incorporates the effect of peak free prednisolone concentration [18].

The model incorporates IIV in apparent clearance, apparent central and peripheral volumes of distribution, and the absorption rate constant, as well as interoccasion variability (IOV) in apparent clearance and apparent central volume of distribution. Both IIV and IOV are modeled using exponential random effects, retaining the IIV correlations reported in the original model, and RUV is described using a proportional form. Reference values for the model parameters and variability are presented in Table 2, along with the corresponding parameter definitions, covariate relationships, and distributional assumptions in the table and its footnotes.

For the TAC reference scenario, we used the initial regimen of 6 mg twice daily proposed via simulation by Bergmann et al. for CYP3A5 nonexpressers [17], with a 12-h dosing interval in the present study. The model parameters and variability settings used the reference values in Table 2, with a reference body weight of 70 kg, and the observed TAC concentration at 24 h was set to 12 ng/mL. The reference individual was set as a CYP3A5 nonexpresser ($X = 0$), corresponding to the 6 mg regimen described above. Hematocrit (HEM), postoperative time (POD), and peak free prednisolone concentration were set to the centering reference values of the original model (0.33 L/L, 22.7 days, and 155.5 nmol/L, respectively), yielding covariate adjustment factors of 1. The effect coefficients for these covariates took the reference values in Table 2, and all covariates remained constant throughout the dosing and observation records.

To determine the drug amounts in the body at the start of the look-back window, all four candidate histories shared a preceding regular dosing history of 6 mg every 12 h, with the last shared dose at −12 h. Blood concentrations were assumed to have reached steady state under this regular dosing regimen before the start of the look-back window. This shared preceding history constituted occasion 1; given the individual random effects and the IOV effects for this occasion, all earlier doses used the corresponding parameters to determine the initial drug amounts under a conditional periodic steady state.

For TAC, subsequent occasions within the look-back window were divided according to the scheduled 12 h interval. Based on the parameters used for observations, 0 h and earlier were assigned to occasion 1, $0 < t \leq 12$ h to occasion 2, and $12 < t \leq 24$ h to occasion 3. If an observation occurred at the 0 h or 12 h boundary, it was evaluated first using the parameters of the preceding occasion, after which the parameters were switched and the dose at that time point was processed. For example, if a sample was collected at 0 h, that observation used the parameters for occasion 1, whereas the subsequent dose at the same time point belonged to occasion 2. In the reference scenario, the observation at 24 h belonged to occasion 3, and the two scheduled doses at 0 h and 12 h belonged to occasions 2 and 3, respectively. The IOV effects across different occasions were mutually independent; early, delayed, or missed doses did not alter this occasion assignment. The final occasion extended through the last observation, and no additional occasions were automatically created after 24 h.

Within the reference scenario for each model, the subsequent computational approaches used the same model settings, candidate histories, prior probabilities, and observation inputs to calculate the posterior probabilities.

Evaluation metrics

In the respective reference scenarios for OXC and TAC, we evaluated the three computational routes—Python Laplace, NONMEM, and the external R package aghq—across four aspects: output availability, repeatability, cross-method agreement, and posterior probability distribution. The subsequent sensitivity analyses used the same evaluation framework, and the definitions of all quantitative metrics are provided in Table 3.

Output availability

For the three computational routes—Python Laplace, NONMEM, and the external R package aghq—a single run was recorded as a valid output only if it returned posterior probabilities correctly corresponding to all candidate states, with each probability being a finite value within $[0,1]$ and the sum of probabilities equaling 1 within a specified numerical tolerance. Based on these criteria, we recorded whether each route obtained a valid output in each scenario; the detailed checking rules are provided in Online Resource 1.

Repeatability

Each route was independently run five times with identical inputs and numerical settings, and computational repeatability was evaluated using the sample standard deviation (SD) and range of the posterior probabilities for each candidate state across replicates. For routes obtaining five valid outputs, we normalized the probability vector from each output separately and then averaged them by candidate state to form the mean posterior probability vector, which was used for subsequent cross-method comparisons and descriptions of the posterior probability distribution.

Cross-method agreement

When both routes obtained five complete valid outputs, we conducted pairwise comparisons among the three computational routes based on their respective mean posterior probability vectors, with states matched across routes by the same candidate dosing history. We evaluated posterior probability differences using the maximum absolute probability difference (d_∞), total variation (TV), Euclidean distance (d_2), and Hellinger distance (d_H), and compared the most probable state and the full ranking of candidate states obtained from the different routes. The pairwise inclusion criteria and the comparison rules for the most probable state and the full ranking are provided in Online Resources 1 and 4.

In the OXC reference scenario, we obtained the posterior probabilities of the candidate dosing histories from the Dashboard of the existing virtual patient approach [8] under matched visible input settings for a

supplementary comparison with the package results. This comparison directly used the web-displayed probability values at their displayed precision without further normalization; the detailed settings and comparison rules are provided in Online Resources 1 and 4.

Posterior probability summaries

Based on the aforementioned mean posterior probability vectors, we used the maximum probability ($P_{\max}$), top-two probability gap (G), and normalized entropy (H_N) to describe the distribution of posterior probabilities across the candidate states for each computational route.

Input sensitivity analyses

To examine the effects of input settings on posterior probability outputs, as well as to evaluate the generalizability of the computational framework across a broader range of scenarios and the consistency across different computational routes, we conducted sensitivity analyses based on the reference scenarios for OXC and TAC. Most scenarios varied the value of a single parameter or a single attribute of one dosing event while keeping all other conditions unchanged. Tables 4 and 5 list the reference and varied settings for the two drugs, respectively.

Structural parameters

To examine the effects of changes in pharmacokinetic parameters on posterior probabilities, we varied the clearance coefficient, volume of distribution, and absorption parameter in each model, and evaluated the effect of changes in body weight on clearance. The TAC scenarios additionally included variations in intercompartmental clearance, peripheral volume of distribution, and absorption lag time. All adjustments retained the original model relationships while keeping the remaining parameters and covariates at their reference settings.

Variability

We then examined the effects of variability magnitude on posterior probabilities. IIV and IOV were varied on the coefficient of variation (CV) scale, adjusting one component at a time; RUV was varied separately by changing the additive residual SD for OXC and the SD of the proportional residual coefficient for TAC. The variation levels included different proportions of the reference values as well as near-zero settings.

For TAC, when an IOV component was varied, its CV was adjusted synchronously across all occasions, while the IOV effects across different occasions remained mutually independent. The IIV correlation coefficients were held fixed, and the corresponding covariances were recalculated from the altered log-scale SDs.

Actual dosing histories

To examine how observations generated by dosing histories outside the candidate set affect posterior probability allocation, we first varied the dosing history used to generate observations and then calculated the posterior probabilities under the fixed candidate set. For generating observations, the reference history consisted of taking the scheduled doses at both 0 h and 12 h, and each varied history adjusted the dose amount or timing of only one of these administrations. Both the reference and varied histories set the random effects to zero and added no residual noise, using the generated 24 h concentration directly as the observation. The posterior calculations retained the four candidate histories, equal prior probabilities, and variability settings of the reference scenario.

Observation records

To examine the effects of sampling time and the number of observations on posterior probabilities, we extracted observations from a fixed reference concentration profile for each drug. This profile was generated using the reference model parameters and the shared preceding dosing history, with a missed dose at 0 h, a dose taken at 12 h, and random effects set to zero; the entire profile was then scaled proportionally so that the 24 h concentration was 12 mg/L for MHD and 12 ng/mL for TAC. Each scenario read concentrations at the corresponding time points from this profile to form single- or multiple-observation records. Multiple observations were analyzed jointly as a single record to calculate posterior probabilities, sharing the individual random effects, whereas the IOV effects for TAC followed the established occasion assignment.

Prior probabilities

To examine the effect of prior information on posterior probability allocation, we fixed the four reference candidate histories and the observation record and varied only the prior probabilities. For both drugs, the prior probabilities for the candidate states were generated by varying the prior probability of each scheduled dose being taken; assuming that the two dosing events were mutually independent yielded the six sets of prior probability vectors listed in the tables.

Candidate states

In contrast to varying the dosing histories used to generate observations, this set of analyses fixed the reference observation record to examine how changes in the candidate histories during inference affected posterior probability allocation. First, we varied the dose amount or timing of a single candidate dosing event at either 0 h or 12 h, applying the adjustment to all candidate histories that included that dose.

Subsequently, we altered the membership of the candidate set by either removing one state at a time from the reference set or adding a new state to the full reference set, forming three-state or five-state sets, respectively. We also set the number of scheduled doses under evaluation within the look-back window to one, two, or three, and enumerated all combinations of taken and missed doses to yield 2, 4, or 8 candidate states, respectively. After changing the set membership or the number of scheduled doses assessed, equal prior probabilities were assigned to all candidates. A shared preceding dosing history was specified for each look-back window and shared among the candidate histories within that window; the specific dosing time points and initialization settings for each window are provided in Online Resource 4.

The sensitivity analyses described above followed the evaluation methods detailed in the Evaluation metrics section. Between-route differences in posterior probabilities were summarized using the median, 95th percentile, and maximum of each metric by model and route pair, with the same summaries also reported by sensitivity setting category. In addition, posterior probabilities for 43 of the 82 OXC sensitivity scenarios were obtained from the Dashboard of the existing virtual patient approach [8] for supplementary comparisons with the results from the package; details of the specific scenarios are provided in Online Resources 1 and 2.

Results

Computational evaluation

Output availability

In the respective reference scenarios for OXC and TAC, the Python package, NONMEM, and the external R package aghq all produced complete posterior probability vectors for the four candidate dosing histories.

Across the 82 OXC sensitivity scenarios, both the Python package and NONMEM obtained five valid posterior probability outputs in all 82 scenarios, whereas the external R package aghq obtained five valid outputs in 81 scenarios. Across the 134 TAC sensitivity scenarios, all three routes obtained five valid outputs in the same 133 scenarios.

In the sensitivity scenarios, when the additive residual SD for OXC was reduced to 1×10^{-6} mg/L, the external R package aghq did not yield posterior probability outputs. When the SD of the proportional residual coefficient for TAC was approximately 6.12×10^{-8}, none of the package Laplace calculation, the external R package aghq, or NONMEM yielded valid posterior probability outputs.

Repeatability

In the respective reference scenarios for OXC and TAC, with identical inputs and numerical settings, both the sample SD and the range of the posterior probabilities for each candidate state across the five replicate runs were zero for each route. In the sensitivity scenarios, all scenarios that successfully yielded outputs for Python, NONMEM, and the external R package aghq, respectively, also showed the same repeatability results.

Cross-method agreement

In the reference scenarios, the posterior probabilities obtained from the Python and NONMEM Laplace implementations showed high agreement, with maximum absolute probability differences (d_∞) of 1.04×10^{-8} for OXC and 1.93×10^{-8} for TAC. Between the package and the external R package aghq, the d_∞ values for the posterior probabilities were 3.47×10^{-4} and 1.25×10^{-2}, respectively. In both models, these differences did not alter the posterior probability ranking of the candidate states (Fig. 1a, b).

Across the 82 sensitivity scenarios for OXC, the median d_∞ between the Python and NONMEM Laplace implementations was 3.06×10^{-9}, with a maximum of 1.19×10^{-4}. Across the 81 scenarios with successful outputs from the external R package aghq, when each of the two Laplace implementations was compared with aghq, the median d_∞ was approximately 3.31×10^{-4} and the maximum was approximately 3.62×10^{-3} for both comparisons. The largest difference involving GHQ occurred in the scenario where the CV for clearance IIV increased from the reference value of 11.1% to 33.3% (Fig. 1c).

Across the 133 sensitivity scenarios for TAC with successful outputs from all three routes, the median d_∞ between the Python and NONMEM Laplace implementations was 1.51×10^{-8}, with a maximum of 4.50×10^{-7}. Across these 133 sensitivity scenarios, when each of the two Laplace implementations was compared with the external R package aghq, the median d_∞ was approximately 1.15×10^{-2} and the maximum was approximately 5.26×10^{-2} for both comparisons. The largest difference involving GHQ occurred in the scenario where the single sampling time was changed from the reference value of 24 h to 16 h (Fig. 1d).

Across the sensitivity scenarios included in the pairwise comparisons, the two Laplace implementations yielded identical rankings of the candidate states. When compared with the external R package aghq, the most probable state was consistent across all 81 OXC scenarios, whereas it differed in 18 of the 133 TAC scenarios (Fig. 1c, d).

Across the OXC sensitivity scenarios included in the comparison, the package and the existing virtual patient approach [8] yielded the same most probable state and showed similar trends in posterior probabilities across scenarios varying concentration, body weight, and single sampling time. Detailed concentration–posterior probability curves and scenario-by-scenario probability comparisons are provided in Online Resources 1 and 2.

Posterior probability summaries

In the respective reference scenarios, the posterior probability distribution was relatively concentrated for OXC, whereas the posterior probabilities were distributed relatively evenly across the four candidate states for TAC. Across the OXC sensitivity scenarios, the posterior probabilities ranged from relatively concentrated to relatively even distributions. In most TAC sensitivity scenarios, the posterior probabilities were distributed relatively evenly across the candidate states, with only a few scenarios showing relatively concentrated distributions. Across the 18 TAC scenarios in which the most probable state differed, the maximum posterior probability obtained by all three routes was below 0.293, and the normalized entropy was above 0.990. Detailed summary metrics of the posterior probability distributions for each scenario are provided in Online Resources 2 and 3.

Effects of input changes

Varying the input settings could alter both the most probable candidate dosing history and the posterior probability distribution. Below, the Laplace calculation results from the package are used to illustrate the effects of various input settings (Figs. 2–4).

Structural parameters

Changes in structural model parameters could alter the most probable candidate dosing history (Fig. 2a, d). In the OXC analysis, with the observation record and all other settings held unchanged, increasing only the clearance coefficient by 20% from the reference value of 1.68 L/h to 2.016 L/h shifted the most probable state from taking only the 12 h dose to taking both doses, with the posterior probability of taking both doses increasing from 0.351 to 0.566.

Variability

Increasing the magnitude of variability made the posterior probability distribution more uniform while leaving the most probable state unchanged. This pattern was observed in both the IIV and residual variability analyses for OXC (Fig. 2b, c). For example, when the residual SD was increased from the reference value of 2.992 mg/L to 11.968 mg/L, the highest posterior probability still corresponded to taking only the 12 h dose, although its probability decreased from 0.523 to 0.279. In the TAC analysis, when the IOV CV for central volume of distribution was increased from the reference value of 126.5% to 506%, the highest posterior probability likewise continued to correspond to taking only the 12 h dose, with its probability decreasing slightly from 0.273 to 0.271.

Actual dosing histories

Varying the dosing history used to generate observations altered the observed concentrations and thereby affected posterior probability allocation under the fixed candidate set. Starting from an OXC dosing history of 300 mg at both 0 h and 12 h, halving only the 12 h dose to 150 mg decreased the 24 h MHD concentration from 14.968 to 10.476 mg/L. When these two concentrations were used as observations while keeping the candidate dosing histories unchanged, the posterior probability of taking 300 mg at both times decreased from 0.589 to 0.234, and the most probable state shifted to taking 300 mg at 12 h only (Fig. 4a).

Observation records

With the same number of observations, different sampling times also altered the posterior probabilities. While retaining the 24 h observation, a single observation at either 0 h (pre-dose) or 18 h was added separately (both additional observations were taken from the same reference concentration profile within each model). The posterior probability of taking only the 12 h dose increased from 0.523 with a single 24

h observation to 0.546 and 0.699 for OXC, and from 0.273 to 0.298 and 0.344 for TAC, respectively (Fig. 3a, c).

Prior probabilities

With the observation record and candidate dosing histories held fixed, adjusting the prior probabilities also altered the most probable state (Fig. 3b, d). In the OXC analysis, increasing the prior probability of each scheduled dose being taken from 0.5 to 0.9 shifted the most probable state from taking only the 12 h dose to taking both doses. The posterior probability of taking both doses rose from 0.351 to 0.831, whereas that of taking only the 12 h dose decreased from 0.523 to 0.137.

Candidate states

With the observation record held fixed, varying the dose amounts in the candidate dosing histories also altered the posterior probability allocation. In the OXC analysis, keeping the combinations of taken and missed doses, the remaining dosing settings, and the prior probabilities unchanged while reducing only the 12 h dose in the candidate histories from 300 mg to 150 mg made taking both doses (300 mg at 0 h and 150 mg at 12 h) the most probable state, with its posterior probability rising from 0.351 before the adjustment to 0.611 (Fig. 4b).

Altering the membership of the candidate set also redistributed the posterior probabilities. Adding a candidate dosing history with a different dosing time and dose amount to the four-state set for OXC and assigning equal prior probabilities across the new set yielded a posterior probability of 0.158 for the new candidate. The probabilities of the four original histories decreased proportionally while maintaining their relative ranking; taking only the 12 h dose remained the most probable state, with its probability decreasing from 0.523 to 0.440. Furthermore, with the 24 h reference observation held fixed, increasing the number of doses evaluated within the look-back window from one to three—expanding the candidate dosing histories from 2 to 8—also redistributed the posterior probabilities across the new set of candidate states.

Discussion

Flexible posterior computation

This study addresses a practical need in TDM: comparing the plausibility of different explanations for the same concentration record while preserving the uncertainty among these explanations. We formulated such problems as the calculation of posterior probabilities across discrete states and established a pharmacokinetic model-based Bayesian computational framework. Users can define candidate states according to the specific problem and calculate the posterior probability of each state by combining the provided model information, observation history, and prior probabilities. Using the discrimination of dosing adherence as a specific application, we set different dosing histories as candidate states to perform numerical verification.

In TDM, Bayesian methods can combine population pharmacokinetic models with individual concentration observations to estimate individual pharmacokinetic parameters [5,6,19]. This inferential approach has also been applied to adherence assessment by calculating the posterior probabilities of different candidate dosing histories from concentration observations and prior information [3,7,8]. Among these methods, the virtual patient approach obtains concentration distributions under each candidate history through simulation and then performs inference by combining the observed concentrations with prior probabilities [8]. In this study, the Python package uses the Laplace approximation to integrate over the random effects and calculate the marginal likelihood and posterior probability of each state, without the need to generate a virtual patient population. Users can set the number of observations and sampling times, and define candidate histories by specifying the time and dose of each administration, rather than

being restricted to binary combinations of "taken or missed" doses at fixed dosing intervals. We also utilized the submodel probability calculation mechanism provided by NONMEM mixture models and IMA-related methods [9,10,20], representing candidate dosing histories as subpopulations in a mixture model to provide a corresponding cross-platform implementation for the same posterior calculation objective.

The Python package connects to user-specified pharmacokinetic models through an interface, enabling different models to reuse the same posterior probability calculation procedures. Once the corresponding model is connected, users can directly call the package to analyze different observation records and candidate histories; when switching models or adjusting input settings, there is no need to rewrite the code for random effects integration and posterior probability calculation.

Numerical reliability

Across the reference scenarios and multiple types of sensitivity scenarios for the two pharmacokinetic models, the Python and NONMEM implementations based on the Laplace approximation [9,12] yielded highly consistent posterior probabilities. The close results obtained across the two software platforms for the same computational objective support the reliable implementation of this algorithm by the package.

Comparison with the external R package aghq [14] further examined the impact of the integration method on the posterior probabilities. In most scenarios, the probability distributions obtained from the Laplace approximation and adaptive GHQ were relatively close, with more pronounced differences occurring mainly under the TAC settings. Across the 18 TAC sensitivity scenarios in which the most probable state changed, the posterior support across candidate histories was relatively close, with the maximum posterior probability obtained by all three routes remaining below 0.293 and the normalized entropy exceeding 0.990. Because no single candidate history held a clear probability advantage, probability differences arising from the integral approximation could alter the most probable state.

Across the OXC scenarios varying concentration, body weight, and single sampling time that were supported by the existing virtual patient approach [8] and had matchable settings, the main trends in posterior probabilities obtained from the package and this approach were generally consistent. This external comparison with a published approach provides further support for the reliability of the package’s computational results.

Posterior sensitivity and uncertainty

The sensitivity analyses showed that changes in model information, predefined candidate states and their priors, and observation information all affected the posterior probabilities of dosing histories. With all other settings held unchanged, increasing the OXC clearance coefficient from 1.68 to 2.016 L/h shifted the most probable history from taking only the 12 h dose to taking both doses (Fig. 2a), whereas increasing the residual variability made the posterior probabilities of the candidate histories closer to one another (Fig. 2c).

With the observations held fixed, adjusting the dose amounts in the OXC candidate histories altered the most probable combination of taken and missed doses (Fig. 4b). When both the observations and the candidate histories were kept unchanged, increasing the prior probability of each scheduled dose being taken from 0.5 to 0.9 raised the posterior probability of taking both doses from 0.351 to 0.831 (Fig. 3b). The candidate definition determines which explanations are compared, whereas the priors specify their relative weights before incorporating observations; both jointly contribute to the posterior allocation.

For OXC, when the 24 h observation was retained and an observation at either 0 h (pre-dose) or 18 h was added separately, the posterior probability of taking only the 12 h dose rose from 0.523 to 0.546 and 0.699, respectively (Fig. 3a). Although both records contained two observations, they yielded different

posterior probabilities, indicating that sampling times and their corresponding concentrations still affect the posterior allocation across candidate histories.

Under the high-variability setting of the TAC reference model, with equal prior probabilities assigned across the four candidate histories, a single 24 h observation yielded a maximum posterior probability of only 0.273 (Fig. 3c), indicating that this combination of information remained insufficient to clearly distinguish among different dosing histories.

Limitations and future directions

The validation in this study was based on simulated datasets, and the clinical utility of the framework still requires evaluation in data with independent dosing-history labels. With these labels, one can examine whether the reported posterior probabilities match the actual frequencies of the corresponding dosing histories. Future studies using real clinical data should also compare this framework with patient self-reports [21], clinicians' judgment, and a Bayesian approach based on Monte Carlo virtual patient simulations [8] to evaluate the accuracy of identifying dosing histories. Further work is also needed to investigate how subsequent dosing regimens should be adjusted on the basis of these posterior probabilities to address patients' prior nonadherence.

Given the population parameters, this study compared candidate dosing histories that shared a pharmacokinetic model structure, random-effects distribution, and residual model. Each candidate used the same observation history, with both the dosing history and the state prior probabilities prespecified, and had the same types and dimensions of random effects to be integrated out. Future work can further incorporate the uncertainty in population parameter estimates and extend the comparison to candidate states with different pharmacokinetic model structures or with different types or numbers of unknown parameters. For example, when comparing on-time dosing with delayed dosing in which the dosing time is unknown within a certain interval, the latter introduces an additional continuous unknown quantity, requiring a probability distribution to be specified for this time and incorporated into the marginal likelihood calculation. In addition, under the near-zero proportional residual setting for TAC, which was far below the reference residual level, none of the three routes yielded valid outputs. This failure may be related to difficulties in optimization and curvature calculation caused by the high sensitivity of the likelihood to prediction errors under extremely small residuals.

Given the model information, candidate states, and their prior probabilities, it is necessary to examine whether the observational information is sufficient to support discrimination among the candidate states. In the TAC reference scenario, the posterior probabilities of the four candidate dosing histories obtained from a single 24 h observation remained close to one another, suggesting that this observation provides limited discriminatory information under the current setting. Building on existing TDM records, future work can combine information theory with optimal sampling design methods to evaluate the information provided by different observation schedules for distinguishing among candidate states, and thereby optimize sampling times, numbers of samples, and their combinations [22].

While this study focused on discriminating among dosing adherence histories as a specific application, the framework can be extended in future work to calculate posterior probabilities for a broader range of discrete states. One class of applications involves defining subpopulations or treatment states with different typical pharmacokinetic parameters or parameter distributions as candidates [9], where such differences may be related to genotype, drug-drug interactions, or immunogenicity [23,24,25]. For example, in the TAC model adopted in this study, CYP3A5 expressers and nonexpressers correspond to different typical apparent clearances [17]. Future studies could accordingly define these two groups as candidate subpopulations to calculate the posterior probability of individual membership. Another class of applications involves comparing different explanations for the same observation record in clinical studies, such as whether a recorded sampling time is accurate or biased, using the corresponding posterior probabilities to aid in assessing the credibility of the data. Each application requires specifying the

candidate states and their prior probabilities according to the specific problem so that the candidate set covers plausible explanations, as well as validating the approach using data with independent state labels.

Conclusions

Starting from the interpretation of unexpected drug concentrations in TDM, this study formulated such problems as the calculation of posterior probabilities across discrete pharmacokinetic states and developed a corresponding Python package. Using the discrimination of dosing adherence histories for two drugs as an application scenario, we performed numerical verification across multiple scenarios, platforms, and algorithms.

Declarations

Funding

Partial financial support was received from the National Natural Science Foundation of China (82073933).

Author contributions

Jinming Zhang: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing.

Zheng Jiao: Conceptualization, Methodology, Project administration, Resources, Supervision, Validation, Writing – original draft, Writing – review & editing.

Both authors approved the manuscript for submission.

Competing interests

The authors declare no competing interests.

Ethics approval and consent to participate

This study used simulation data and published pharmacokinetic models. Ethics approval and consent to participate were not applicable.

Consent for publication

Not applicable.

Data availability

The scenario inputs, raw computational outputs, and result tables supporting the findings of this study are provided in Online Resource 4, and the result summaries are provided in Online Resources 2 and 3.

Code availability

The Python package and calculation scripts, NONMEM control streams and data files, reference scripts using the R package aghq, and reproduction instructions are provided in Online Resource 4. Software versions and dependencies are detailed in Online Resource 1. Re-running the NONMEM calculations requires a licensed NONMEM installation.

Use of AI-assisted tools

AI-assisted tools were used for code development and debugging, summarizing and organizing experimental results, writing plotting code, and organizing supporting files and supplementary materials. They were also used for manuscript translation and language editing, as well as internal manuscript review. The translated text was reviewed and revised paragraph by paragraph by the authors. The authors take full responsibility for the accuracy and integrity of the manuscript and supplementary materials, including the code, analyses, and conclusions.

References

1.Patsalos PN, Berry DJ, Bourgeois BFD, Cloyd JC, Glauser TA, Johannessen SI, Leppik IE, Tomson T, Perucca E (2008). Antiepileptic drugs–best practice guidelines for therapeutic drug monitoring: a position paper by the subcommission on therapeutic drug monitoring, ILAE Commission on

Therapeutic Strategies. Epilepsia. 49(7):1239-1276. https://doi.org/10.1111/j.1528-1167.2008.01561.x.

2.Patsalos PN, Spencer EP, Berry DJ (2018). Therapeutic Drug Monitoring of Antiepileptic Drugs in Epilepsy: A 2018 Update. Therapeutic Drug Monitoring. 40(5):526-548. https://doi.org/10.1097/FTD.0000000000000546.

3.Barriere O, Li J, Nekka F (2013). Compliance spectrum as a drug fingerprint of drug intake and drug disposition. Journal of Pharmacokinetics and Pharmacodynamics. 40(1):41-52. https://doi.org/10.1007/s10928-012-9285-x.

4.Li J, Nekka F (2007). A Pharmacokinetic Formalism Explicitly Integrating the Patient Drug Compliance. Journal of Pharmacokinetics and Pharmacodynamics. 34(1):115-139. https://doi.org/10.1007/s10928-006-9036-y.

5.Keizer RJ, ter Heine R, Frymoyer A, Lesko LJ, Mangat R, Goswami S (2018). Model-Informed Precision Dosing at the Bedside: Scientific Challenges and Opportunities. CPT: Pharmacometrics & Systems Pharmacology. 7(12):785-787. https://doi.org/10.1002/psp4.12353.

6.Darwich AS, Polasek TM, Aronson JK, Ogungbenro K, Wright DFB, Achour B, Reny JL, Daali Y, Eiermann B, Cook J, Lesko L, McLachlan AJ, Rostami-Hodjegan A (2021). Model-Informed Precision Dosing: Background, Requirements, Validation, Implementation, and Forward Trajectory of Individualizing Drug Therapy. Annual Review of Pharmacology and Toxicology. 61:225-245. https://doi.org/10.1146/annurev-pharmtox-033020-113257.

7.Barriere O, Li J, Nekka F (2011). A Bayesian approach for the estimation of patient compliance based on the last sampling information. Journal of Pharmacokinetics and Pharmacodynamics. 38(3):333-351. https://doi.org/10.1007/s10928-011-9196-2.

8.Liu XQ, Li ZR, Lin WW, Wang J, Gu FQ, Ding JJ, Jiao Z (2026). Bayesian-Based Pharmacokinetic Framework Integrated with Therapeutic Drug Monitoring for Assessing Adherence to Antiseizure Medications: A Clinical Trial Simulation Study. Journal of Medical Internet Research. 28:e77917. https://doi.org/10.2196/77917.

9.Bauer RJ (2019). NONMEM Tutorial Part II: Estimation Methods and Advanced Examples. CPT: Pharmacometrics & Systems Pharmacology. 8(8):538-556. https://doi.org/10.1002/psp4.12422.

10.Carlsson KC, Savic RM, Hooker AC, Karlsson MO (2009). Modeling Subpopulations with the $MIXTURE Subroutine in NONMEM: Finding the Individual Probability of Belonging to a Subpopulation for the Use in Model Analysis and Improved Decision Making. The AAPS Journal. 11(1):148-154. https://doi.org/10.1208/s12248-009-9093-4.

11.Wolfinger R (1993). Laplace's approximation for nonlinear mixed models. Biometrika. 80(4):791-795. https://doi.org/10.1093/biomet/80.4.791.

12.Kristensen K, Nielsen A, Berg CW, Skaug H, Bell BM (2016). TMB: Automatic Differentiation and Laplace Approximation. Journal of Statistical Software. 70(5):1-21. https://doi.org/10.18637/jss.v070.i05.

13.Pinheiro JC, Bates DM (1995). Approximations to the Log-Likelihood Function in the Nonlinear Mixed-Effects Model. Journal of Computational and Graphical Statistics. 4(1):12-35. https://doi.org/10.1080/10618600.1995.10474663.

14.Stringer A (2021). Implementing Approximate Bayesian Inference using Adaptive Quadrature: the aghq Package. arXiv:2101.04468. https://doi.org/10.48550/arXiv.2101.04468.

15.Virtanen P, Gommers R, Oliphant TE, et al (2020). SciPy 1.0: fundamental algorithms for scientific computing in Python. Nature Methods. 17(3):261-272. https://doi.org/10.1038/s41592-019-0686-2.

16.Lin WW, Li XW, Jiao Z, Zhang J, Rao X, Zeng DY, Lin XH, Wang CL (2019). Population pharmacokinetics of oxcarbazepine active metabolite in Chinese paediatric epilepsy patients and its application in individualised dosage regimens. European Journal of Clinical Pharmacology. 75(3):381-392. https://doi.org/10.1007/s00228-018-2600-8.

17.Bergmann TK, Hennig S, Barraclough KA, Isbel NM, Staatz CE (2014). Population pharmacokinetics of tacrolimus in adult kidney transplant patients: impact of CYP3A5 genotype on starting dose. Therapeutic Drug Monitoring. 36(1):62-70. https://doi.org/10.1097/FTD.0b013e31829f1ab8.

18.Staatz CE, Bergmann TK, Barraclough KA, Isbel NM, Hennig S (2013). Population pharmacokinetics of tacrolimus in adult kidney transplant recipients. Poster presented at the Population Approach Group Europe (PAGE) meeting. https://www.page-meeting.org/wp-content/uploads/pdf_assets/7204-PAGE2013Staatz.pdf. Accessed 21 September 2026.

19.Leven C, Coste A, Mane C (2022). Free and Open-Source Posologyr Software for Bayesian Dose Individualization: An Extensive Validation on Simulated Data. Pharmaceutics. 14(2):442. https://doi.org/10.3390/pharmaceutics14020442.

20.Chasseloup E, Tessier A, Karlsson MO (2021). Assessing Treatment Effects with Pharmacometric Models: A New Method that Addresses Problems with Standard Assessments. The AAPS Journal. 23(3):63. https://doi.org/10.1208/s12248-021-00596-8.

21.Osterberg L, Blaschke T (2005). Adherence to Medication. New England Journal of Medicine. 353(5):487-497. https://doi.org/10.1056/NEJMra050100.

22.Busetto AG, Hauser A, Krummenacher G, Sunnaker M, Dimopoulos S, Ong CS, Stelling J, Buhmann JM (2013). Near-optimal experimental design for model selection in systems biology. Bioinformatics. 29(20):2625-2632. https://doi.org/10.1093/bioinformatics/btt436.

23.Birdwell KA, Decker B, Barbarino JM, Peterson JF, Stein CM, Sadee W, Wang D, Vinks AA, He Y, Swen JJ, Leeder JS, van Schaik RHN, Thummel KE, Klein TE, Caudle KE, MacPhee IAM (2015). Clinical Pharmacogenetics Implementation Consortium (CPIC) Guidelines for CYP3A5 Genotype and Tacrolimus Dosing. Clinical Pharmacology & Therapeutics. 98(1):19-24. https://doi.org/10.1002/cpt.113.

24.Mueck W, Kubitza D, Becka M (2013). Co-administration of rivaroxaban with drugs that share its elimination pathways: pharmacokinetic effects in healthy subjects. British Journal of Clinical Pharmacology. 76(3):455-466. https://doi.org/10.1111/bcp.12075.

25.Berends SE, van Steeg TJ, Ahsman MJ, Singh S, Brandse JF, D'Haens GRAM, Mathot RAA (2019). Tumor necrosis factor-mediated disposition of infliximab in ulcerative colitis patients. Journal of Pharmacokinetics and Pharmacodynamics. 46(6):543-551. https://doi.org/10.1007/s10928-019-09652-5.

Table 1. Oxcarbazepine model parameters and reference values [16]

Parameter	Reference	Definition
Structural parameters[a]		
CL_p (L/h)	1.68	Typical apparent clearance at a body weight of 70 kg
V_p (L)	14.7	Typical apparent volume of distribution
k_a (h^{-1})	0.83	First-order absorption rate constant
Random-effect variability[b]		
$\omega_{CL/F}$	11.1%	Exponential interindividual variability
Residual variability[c]		
σ (mg/L)	2.992	Additive residual variability

Typical parameter relationships

$$CL/F = CL_p\left(\frac{BW}{70}\right)^{0.624-\frac{0.233\,BW^{219}}{8.97^{219}+BW^{219}}},$$

$$V/F = V_p.$$

a. Structural model. Reference values follow the final model reported by Lin et al. The one-compartment model has first-order absorption and elimination and describes concentrations of the monohydroxy derivative (MHD), the active metabolite of oxcarbazepine. The equations give typical parameter relationships before random effects. Here, CL is clearance, V is volume of distribution, k_a is the first-order absorption rate constant, F is oral bioavailability, and BW is body weight in kg.

b. Random-effect variability. Individual CL/F equals its typical value multiplied by the exponential factor $\exp(\eta_{CL})$; the reference coefficient of variation (CV) for interindividual variability is 11.1%.

c. Residual variability. Observed MHD concentrations equal model-predicted concentrations plus residual error, which is normally distributed with mean zero and a standard deviation of 2.992 mg/L.

Table 2. Tacrolimus model parameters and reference values [17,18]

Parameter	Reference	Definition
Structural parameters[a]		
CL/F, $\theta_{CL/F}$ (L/h)	25.5	Typical apparent clearance at reference covariates
$V1/F$, $\theta_{V1/F}$ (L)	113.0	Typical apparent central volume of distribution at reference covariates
Q/F (L/h)	67.9	Apparent intercompartmental clearance
$V2/F$ (L)	1060	Apparent peripheral volume of distribution
ka (h^{-1})	0.35	First-order absorption rate constant
Lag time (h)	0.44	Delay before absorption begins
Covariates and coefficients[b]		
X	0	CYP3A5 genotype indicator (nonexpresser)
θ_{CYP3A5}	1.60	
HEM (L/L)	0.33	Hematocrit
θ_{HEM}	-1.01	
POD (days)	22.7	Postoperative time
θ_{POD}	-0.21%	
$\mathrm{Pred}C_{\mathrm{max,free}}$ (nmol/L)	155.5	Peak free prednisolone concentration
θ_{Pred}	-0.28%	
Random-effect variability[c]		
$\mathrm{IIV}_{CL/F}$	29.5%	Exponential interindividual variability
$\mathrm{IIV}_{V1/F}$	46.8%	Exponential interindividual variability
$\mathrm{IIV}_{V2/F}$	89.4%	Exponential interindividual variability
IIV_{ka}	47.6%	Exponential interindividual variability
$\mathrm{IOV}_{CL/F}$	29.9%	Exponential interoccasion variability
$\mathrm{IOV}_{V1/F}$	126.5%	Exponential interoccasion variability
Correlation: $V1/F, ka$	67.7%	Correlation between log-scale interindividual random effects
Correlation: $V1/F, V2/F$	-4.9%	Correlation between log-scale interindividual random effects
Correlation: $ka, V2/F$	-1.3%	Correlation between log-scale interindividual random effects
Residual variability[d]		
σ	18.3%	Proportional residual variability

Typical parameter relationships

$$(CL/F)_i = \theta_{CL/F} \times \theta_{CYP3A5}^{X}[1 + \theta_{\mathrm{HEM}}(\mathrm{HEM} - 0.33)] \times \left(\frac{WT}{70}\right)^{0.75} [1 + \theta_{\mathrm{POD}}(\mathrm{POD} - 22.7)],$$

$$(V1/F)_i = \theta_{V1/F}[1 + \theta_{\mathrm{Pred}}(\mathrm{Pred}C_{\mathrm{max,free}} - 155.5)].$$

Table 2 (continued)

a. Structural model. Model parameter reference values are the final-model estimates reported by Bergmann et al. The two-compartment model has an absorption lag, first-order absorption and first-order elimination and describes tacrolimus (TAC) concentrations. The equations give covariate-dependent typical parameter values before random effects. Here, CL is clearance, $V1$ is central volume of distribution, Q is intercompartmental clearance, $V2$ is peripheral volume of distribution, ka is the first-order absorption rate constant, F is oral bioavailability.

b. Covariates and their coefficients. The covariate reference values specify the individual used in this evaluation. Covariates were fixed at the values shown so that the model's typical parameter values corresponded to those of CYP3A5 nonexpressers.

c. Random-effect variability. Interindividual and interoccasion variability are incorporated as exponential multipliers on the corresponding parameters. The reported correlations are between the corresponding interindividual random effects.

d. Residual variability. The residual standard deviation is 18.3% of the predicted TAC concentration.

Table 3. Evaluation metrics

Measure	Definition	Interpretation
Repeatability[a]		
Sample standard deviation (SD) and range	$s_r = \sqrt{\frac{1}{4}\sum_{j=1}^{5}(r_j - \bar{r})^2}$ $R_r = \max_j r_j - \min_j r_j$	Run-to-run variability in state probabilities
Cross-method agreement[b]		
Maximum absolute probability difference, d_∞	$d_\infty = \max_s \lvert p_s - q_s \rvert$	Largest absolute probability difference for any state
Total variation, TV	$\mathrm{TV} = \frac{1}{2}\sum_{s=1}^{K} \lvert p_s - q_s \rvert$	Minimum probability reallocated to match the two distributions
Euclidean distance, d_2	$d_2 = \sqrt{\sum_{s=1}^{K}(p_s - q_s)^2}$	Overall difference in state probabilities
Hellinger distance, d_H	$d_H = \sqrt{\frac{1}{2}\sum_{s=1}^{K}(\sqrt{p_s} - \sqrt{q_s})^2}$	Overall difference on the square-root probability scale
Posterior probability summaries[c]		
Maximum probability, $P_{\max}$	$P_{\max} = p_{(1)}$	Posterior support for the most probable state
Top-two probability gap, G	$G = p_{(1)} - p_{(2)}$	Separation between the two most probable states
Normalized entropy, H_N	$H_N = -\frac{\sum_{s=1}^{K} p_s \log p_s}{\log K}$	Dispersion across states: 0, all probability on one state; 1, equal probabilities

Note. p and q are posterior probability vectors over K candidate states, with p_s and q_s denoting the posterior probabilities of state s ($s = 1, \ldots, K$). For the Python Laplace, NONMEM, and external R aghq computational routes, in scenarios yielding five valid outputs, the probability vector from each output was normalized before averaging by state; the existing virtual patient approach used the web-displayed probabilities directly without further normalization.

a. r_j is the posterior probability of a given candidate state in replicate run j, and $\bar{r}$ is the mean across five runs. Posterior probabilities before and after normalization were summarized separately, reporting the maximum sample standard deviation (SD) and maximum range across all candidate states.

b. Pairwise comparisons among the three computational routes matched scenarios, input data, and prior settings, whereas supplementary comparisons with the virtual patient approach matched only visible input settings. In each formula, p_s and q_s correspond to the same candidate dosing history; smaller distances indicate closer posterior probability distributions between the two methods.

c. $K \geq 2$; $p_{(1)}$ and $p_{(2)}$ are the largest and second-largest posterior probabilities in vector p. When $p_s = 0$, $p_s \log(p_s)$ is defined as 0 and the state remains counted in K.

Table 4. Reference and varied settings for the oxcarbazepine sensitivity analyses

Setting	Reference	Variations
1. Structural model parameters and body weight[a]		
Clearance coefficient, CL_p (L/h)	1.68	0.84, 1.344, 1.68, 2.016, 3.36
Apparent volume, V_p (L)	14.7	7.35, 11.76, 14.7, 17.64, 29.4
Absorption rate constant, k_a (h^{-1})	0.83	0.415, 0.664, 0.83, 0.996, 1.66
Body weight, BW (kg)	25	12.5, 20, 25, 30, 50
2. Random-effect and residual variability[b]		
$IIV_{CL/F}$	11.1%	0.01%, 2.78%, 5.55% 8.33%, 11.10%, 16.65% 22.20%, 33.30%, 44.40%
RUV_{add}	2.992 mg/L	10^{-6}, 0.748, 1.496 2.244, 2.992, 4.488 5.984, 8.976, 11.968 mg/L
3. Dosing histories used to generate observations[c]		
Dose amounts (mg)	(300,300)	(150,300), (600,300) (300,150), (300,600)
Dose times (h)	(0,12)	(−4,12), (−2,12) (2,12), (4,12) (0,8), (0,10) (0,14), (0,16)
4. Observation records[d]		
Single observation, (t, C) (h, mg/L)	(24,12)	$t \in \{16,18,20,22,24,$ $26,28,30,32\}$ h $C = C_{\text{ref}}(t)$
Observation-time vectors (h)	(24)	$n = 1$: (24) $n = 2$: (0,24), (18,24), (24,30) $n = 3$: (0,24,30), (18,24,30)

Table 4 (continued)

Setting	Reference	Variations
5. Prior probabilities[e]		
Prior probabilities, ordered by state $(\omega_{00}, \omega_{01}, \omega_{10}, \omega_{11})$	$(0.25,0.25,0.25,0.25)$	$(0.81,0.09,0.09,0.01)$ $(0.49,0.21,0.21,0.09)$ $(0.25,0.25,0.25,0.25)$ $(0.09,0.21,0.21,0.49)$ $(0.04,0.16,0.16,0.64)$ $(0.01,0.09,0.09,0.81)$
6. Candidate states and sets[f]		
Candidate dose amounts (mg)	$(300,300)$	$(150,300),\ (600,300)$ $(300,150),\ (300,600)$
Candidate dose times (h)	$(0,12)$	$(-4,12),\ (-2,12)$ $(2,12),\ (4,12)$ $(0,8),\ (0,10)$ $(0,14),\ (0,16)$
Candidate-state set	$\{\omega_{00}, \omega_{01}, \omega_{10}, \omega_{11}\}$	$\{\omega_{01}, \omega_{10}, \omega_{11}\}$ $\{\omega_{00}, \omega_{10}, \omega_{11}\}$ $\{\omega_{00}, \omega_{01}, \omega_{11}\}$ $\{\omega_{00}, \omega_{01}, \omega_{10}\}$ $\{\omega_{00}, \omega_{01}, \omega_{10}, \omega_{11}, \omega_{\text{add}}\}$
Candidate window: dose times and state count	$(0,12)$ h; $K = 4$	(12) h; $K = 2$ $(-12,0,12)$ h; $K = 8$

Table 4 (continued)

Reference settings. Reference settings define the baseline inputs for the sensitivity analyses. All other inputs remain unchanged when the listed variations are applied.

a. Structural model parameters and body weight. Each parameter or body weight is varied separately at 0.5, 0.8, 1, 1.2 and 2 times its reference value. Body weight affects clearance through the unchanged model relationship.

b. Random-effect and residual variability. One component is varied at a time: IIV on the CV scale and additive RUV on the residual SD scale, using 0.25, 0.5, 0.75, 1, 1.5, 2, 3 and 4 times the reference value. For exponential random effects, the log-scale variance is $\log(1 + \mathrm{CV}^2)$, with CV expressed as a decimal. The additional near-zero IIV level has a log-scale variance of 10^{-8} (approximately 0.01% CV); the near-zero RUV level is 10^{-6} mg/L. IIV CVs are expressed as percentages and shown to two decimal places.

c. Dosing histories used to generate observations. The amount or time of one dose is changed while the candidate histories used for inference remain fixed. For both the reference and altered histories, the 24 h concentration is generated with random effects set to zero and no residual noise added. This concentration is then used as the observation to calculate posterior probabilities for the reference candidate states.

d. Observation records. The typical profile for missing the 0 h dose and taking the 12 h dose is scaled to 12 mg/L at 24 h. The resulting profile, $C_{\mathrm{ref}}(t)$, provides the concentrations at the observation times listed in the table. Each listed set of observation times defines one complete observation record. A 0 h observation precedes dosing.

e. Prior probabilities. Only the prior probabilities change; candidate histories and observations remain fixed. The state order is specified in the table.

f. Candidate states and sets. The 24 h reference concentration is used as the fixed observation. A change to the amount or time of a candidate dose is applied to every candidate history containing that dose. The added candidate ω_{add} comprises 200 mg at 2 h and 400 mg at 12 h.

For retrospective windows covering one, two or three scheduled doses, all combinations of taken and missed doses are enumerated, giving 2, 4 or 8 candidate states, respectively. Equal priors are assigned after changes to set membership or the number of doses assessed.

Symbols. OXC, oxcarbazepine; MHD, monohydroxy derivative, its active metabolite. CL, clearance; V, volume of distribution; F, oral bioavailability; CL/F and V/F, apparent clearance and volume of distribution; CL_p, typical apparent clearance coefficient at 70 kg; $V_p = V/F$, typical apparent volume of distribution; k_a, first-order absorption rate constant; BW, body weight.

IIV, interindividual variability (the subscript identifies the affected parameter); RUV, residual unexplained variability; $\mathrm{RUV}_{\mathrm{add}}$, additive residual SD; CV, coefficient of variation; SD, standard deviation. t, observation time; C, MHD concentration; n, observations per record; K, candidate-state count.

ω denotes a candidate dosing state. Digits run from the most recent dose backward: 1 means taken and 0 means missed. Thus ω_{00}, ω_{01}, ω_{10} and ω_{11} denote neither dose taken, only the 0 h dose taken, only the 12 h dose taken and both doses taken, respectively.

Table 5. Reference and varied settings for the tacrolimus sensitivity analyses

Setting	Reference	Variations
1. Structural model parameters and body weight[a]		
Clearance coefficient, $\theta_{CL/F}$ (L/h)	25.5	12.75, 20.4, 25.5, 30.6, 51
Central-volume coefficient, $\theta_{V1/F}$ (L)	113.0	56.5, 90.4, 113.0, 135.6, 226
Q/F (L/h)	67.9	33.95, 54.32, 67.9, 81.48, 135.8
$V2/F$ (L)	1060	530, 848, 1060, 1272, 2120
Absorption rate constant, k_a (h^{-1})	0.35	0.175, 0.28, 0.35, 0.42, 0.7
Lag time (h)	0.44	0.22, 0.352, 0.44, 0.528, 0.88
Body weight, WT (kg)	70	35, 56, 70, 84, 140
2. Random-effect and residual variability[b]		
$\mathrm{IIV}_{CL/F}$	29.5%	0.01%, 7.38%, 14.75% 22.13%, 29.50%, 44.25% 59.00%, 88.50%, 118.00%
$\mathrm{IIV}_{V1/F}$	46.8%	0.01%, 11.70%, 23.40% 35.10%, 46.80%, 70.20% 93.60%, 140.40%, 187.20%
$\mathrm{IIV}_{V2/F}$	89.4%	0.01%, 22.35%, 44.70% 67.05%, 89.40%, 134.10% 178.80%, 268.20%, 357.60%
IIV_{k_a}	47.6%	0.01%, 11.90%, 23.80% 35.70%, 47.60%, 71.40% 95.20%, 142.80%, 190.40%
$\mathrm{IOV}_{CL/F}$	29.9%	0.01%, 7.48%, 14.95% 22.43%, 29.90%, 44.85% 59.80%, 89.70%, 119.60%
$\mathrm{IOV}_{V1/F}$	126.5%	0.01%, 31.63%, 63.25% 94.88%, 126.50%, 189.75% 253.00%, 379.50%, 506.00%
$\mathrm{RUV}_{\mathrm{prop}}$	18.3%	6.1163×10^{-6}%, 4.575%, 9.15% 13.725%, 18.3%, 27.45% 36.6%, 54.9%, 73.2%

Table 5 (continued)

Setting	Reference	Variations
3. Dosing histories used to generate observations[c]		
Dose amounts (mg)	(6,6)	(3,6), (12,6) (6,3), (6,12)
Dose times (h)	(0,12)	(−4,12), (−2,12) (2,12), (4,12) (0,8), (0,10) (0,14), (0,16)
4. Observation records[d]		
Single observation, (t, C) (h, ng/mL)	(24,12)	$t \in \{16,18,20,22,24,$ $26,28,30,32\}$ h $C = C_{\mathrm{ref}}(t)$
Observation-time vectors (h)	(24)	$n = 1$: (24) $n = 2$: (0,24), (18,24), (24,30) $n = 3$: (0,24,30), (18,24,30)
5. Prior probabilities[e]		
Prior probabilities, ordered by state $(\omega_{00}, \omega_{01}, \omega_{10}, \omega_{11})$	(0.25,0.25,0.25,0.25)	(0.81,0.09,0.09,0.01) (0.49,0.21,0.21,0.09) (0.25,0.25,0.25,0.25) (0.09,0.21,0.21,0.49) (0.04,0.16,0.16,0.64) (0.01,0.09,0.09,0.81)
6. Candidate states and sets[f]		
Candidate dose amounts (mg)	(6,6)	(3,6), (12,6) (6,3), (6,12)
Candidate dose times (h)	(0,12)	(−4,12), (−2,12) (2,12), (4,12) (0,8), (0,10) (0,14), (0,16)
Candidate-state set	$\{\omega_{00}, \omega_{01}, \omega_{10}, \omega_{11}\}$	$\{\omega_{01}, \omega_{10}, \omega_{11}\}$ $\{\omega_{00}, \omega_{10}, \omega_{11}\}$ $\{\omega_{00}, \omega_{01}, \omega_{11}\}$ $\{\omega_{00}, \omega_{01}, \omega_{10}\}$ $\{\omega_{00}, \omega_{01}, \omega_{10}, \omega_{11}, \omega_{\mathrm{add}}\}$
Candidate window: dose times and state count	(0,12) h; $K = 4$	(12) h; $K = 2$ (−12,0,12) h; $K = 8$

Table 5 (continued)

Reference settings. Reference settings define the baseline inputs for the sensitivity analyses. All other inputs remain unchanged when the listed variations are applied.

a. Structural model parameters and body weight. Each parameter or body weight is varied separately at 0.5, 0.8, 1, 1.2 and 2 times its reference value. Body weight affects clearance through the unchanged model relationship. Other covariates and their coefficients remain fixed.

b. Random-effect and residual variability. One component is varied at a time: IIV and IOV on the CV scale and proportional RUV through the SD of the residual coefficient, using 0.25, 0.5, 0.75, 1, 1.5, 2, 3 and 4 times the reference value. For exponential random effects, the log-scale variance is $\log(1 + \mathrm{CV}^2)$, with CV expressed as a decimal. The additional near-zero IIV and IOV levels have a log-scale variance of 10^{-8} (approximately 0.01% CV). CVs are expressed as percentages and shown to two decimal places. The near-zero RUV level is approximately 6.1163×10^{-6}%, using the same relative reduction as the OXC near-zero level ($10^{-6}/2.992$).

IIV correlations for $V1/F$ with k_a, $V1/F$ with $V2/F$, and k_a with $V2/F$ remain 67.7%, -4.9% and -1.3%, respectively; covariances are recalculated from log-scale SDs. When an IOV component is varied, the same CV is used for that component in all occasions.

Occasions follow the planned dosing intervals; shifted or missed doses do not move the boundaries. At each boundary, observations are evaluated using the preceding occasion's parameters. The parameters are then switched before any dose at that time is processed. The final occasion extends to the last observation. IOV effects are independent across occasions.

c. Dosing histories used to generate observations. The amount or time of one dose is changed while the candidate histories used for inference remain fixed. For both the reference and altered histories, the 24 h concentration is generated with random effects set to zero and no residual noise added. This concentration is then used as the observation to calculate posterior probabilities for the reference candidate states.

d. Observation records. The typical profile for missing the 0 h dose and taking the 12 h dose is scaled to 12 ng/mL at 24 h. The resulting profile, $C_{\mathrm{ref}}(t)$, provides the concentrations at the observation times listed in the table. Each listed set of observation times defines one complete observation record. A 0 h observation precedes dosing.

e. Prior probabilities. Only the prior probabilities change; candidate histories and observations remain fixed. The state order is specified in the table.

f. Candidate states and sets. The 24 h reference concentration is used as the fixed observation. A change to the amount or time of a candidate dose is applied to every candidate history containing that dose. The added candidate ω_{add} comprises 4 mg at 2 h and 8 mg at 12 h.

For retrospective windows covering one, two or three scheduled doses, all combinations of taken and missed doses are enumerated, giving 2, 4 or 8 candidate states, respectively. Equal priors are assigned after changes to set membership or the number of doses assessed.

Symbols. TAC, tacrolimus; OXC, oxcarbazepine. CL, clearance; $V1$ and $V2$, central and peripheral volumes of distribution; Q, intercompartmental clearance; F, oral bioavailability. Division by F denotes the corresponding apparent parameter. $\theta_{CL/F}$ and $\theta_{V1/F}$, coefficients for typical apparent clearance and central volume at the reference covariate values; k_a, first-order absorption rate constant; lag time, absorption delay; WT, body weight.

IIV, interindividual variability; IOV, interoccasion variability (subscripts identify affected parameters); an occasion is a dosing period defining IOV effects. RUV, residual unexplained variability; $\mathrm{RUV}_{\mathrm{prop}}$, proportional residual-coefficient SD, expressed as a percentage; CV, coefficient of variation; SD, standard deviation. t, observation time; C, TAC concentration; n, observations per record; K, candidate-state count.

ω denotes a candidate dosing state. Digits run from the most recent dose backward: 1 means taken and 0 means missed. Thus ω_{00}, ω_{01}, ω_{10} and ω_{11} denote neither dose taken, only the 0 h dose taken, only the 12 h dose taken and both doses taken, respectively.

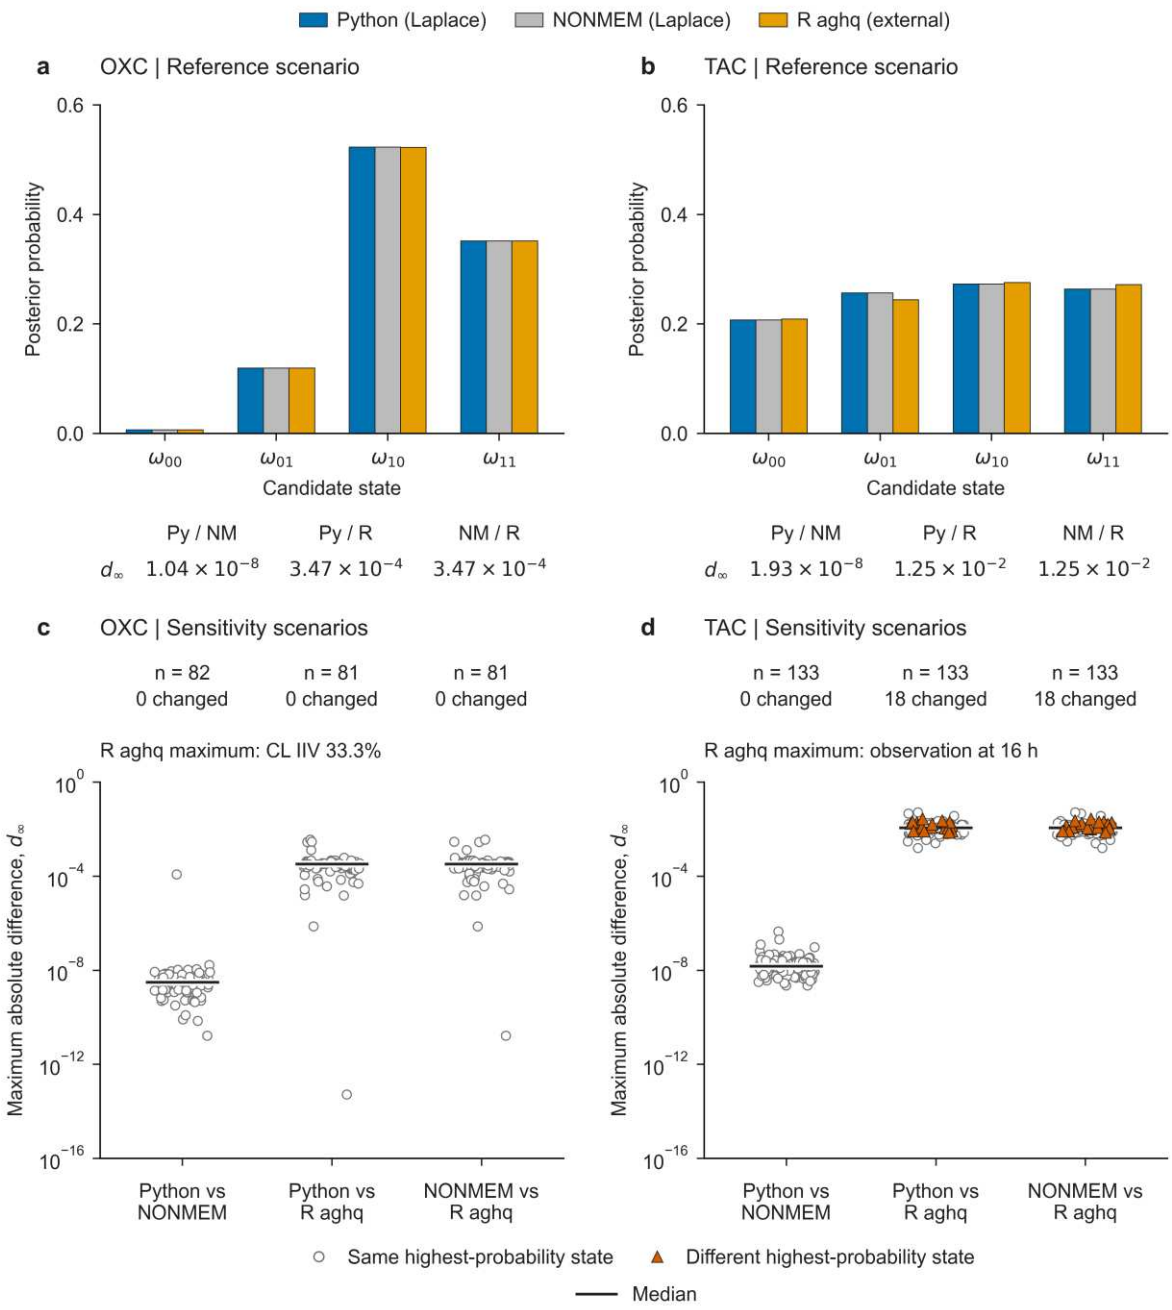

Fig. 1 Agreement between posterior probabilities obtained by different computational routes

a, b Posterior probabilities for the oxcarbazepine (OXC) and tacrolimus (TAC) reference scenarios, respectively. Values below show the maximum absolute probability difference across candidate states between each pair of routes (d_∞). c, d Corresponding differences across sensitivity scenarios for the two models. Text above the panels identifies the input settings with the largest differences in comparisons involving R aghq.

Python (Py) and NONMEM (NM) use the Laplace approximation; R denotes adaptive Gauss-Hermite quadrature implemented using the external R package aghq. ω_{00}, ω_{01}, ω_{10} and ω_{11} represent neither dose taken, only the 0 h dose taken, only the 12 h dose taken and both doses taken, respectively. CL IIV denotes interindividual variability in clearance, expressed as a coefficient of variation (%).

In c and d, each point represents one sensitivity scenario. Circles and triangles indicate the same and different highest-probability states between routes, respectively; horizontal lines indicate medians. n is the number of scenarios in each pairwise comparison; changed is the number with different highest-probability states. Comparisons use mean posterior probabilities after normalizing each of five runs; sensitivity comparisons exclude the reference scenarios

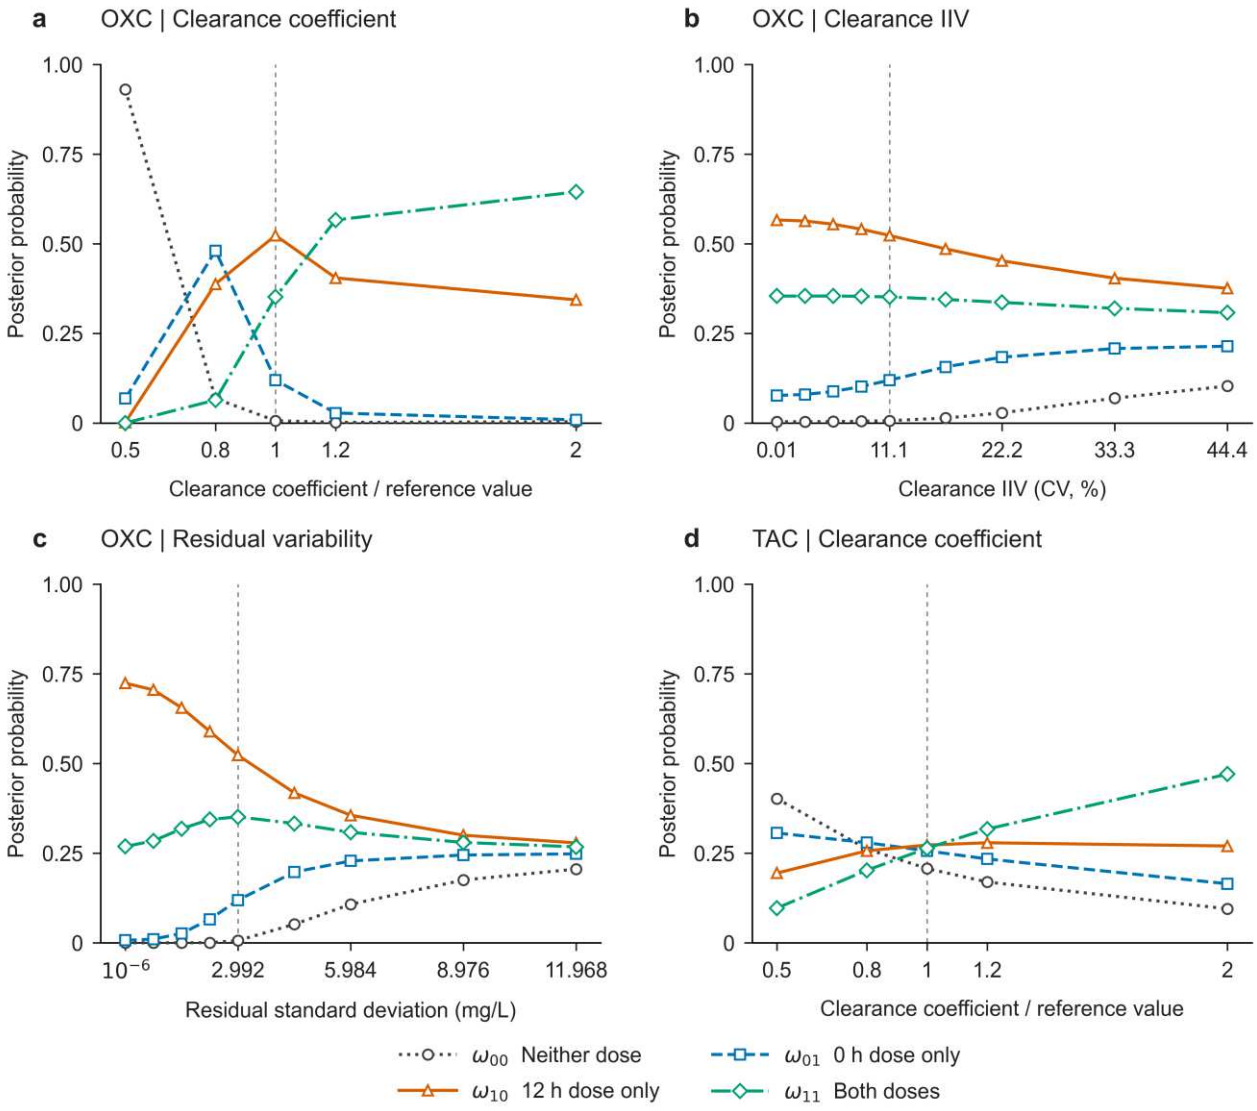

Fig. 2 Effects of model parameters and variability settings on posterior probability distributions

Starting from each model's reference scenario, panels a, b and c show the posterior probabilities of the four candidate states as the oxcarbazepine (OXC) clearance coefficient, interindividual variability (IIV) in clearance and residual variability were varied, respectively. Panel d shows the corresponding results as the tacrolimus (TAC) clearance coefficient was varied.

In a and d, the horizontal axis shows the clearance coefficient as a multiple of its reference value. In b, IIV is expressed as a coefficient of variation (CV); in c, additive residual variability is expressed as a standard deviation. Vertical dashed lines indicate the reference settings.

The four candidate states represent neither dose taken (ω_{00}), only the 0 h dose taken (ω_{01}), only the 12 h dose taken (ω_{10}), and both doses taken (ω_{11}). Each point represents the posterior probability at the corresponding setting, and lines connect adjacent evaluated points.

Values are mean posterior probabilities after normalizing each of five runs of the Python package's Laplace calculation

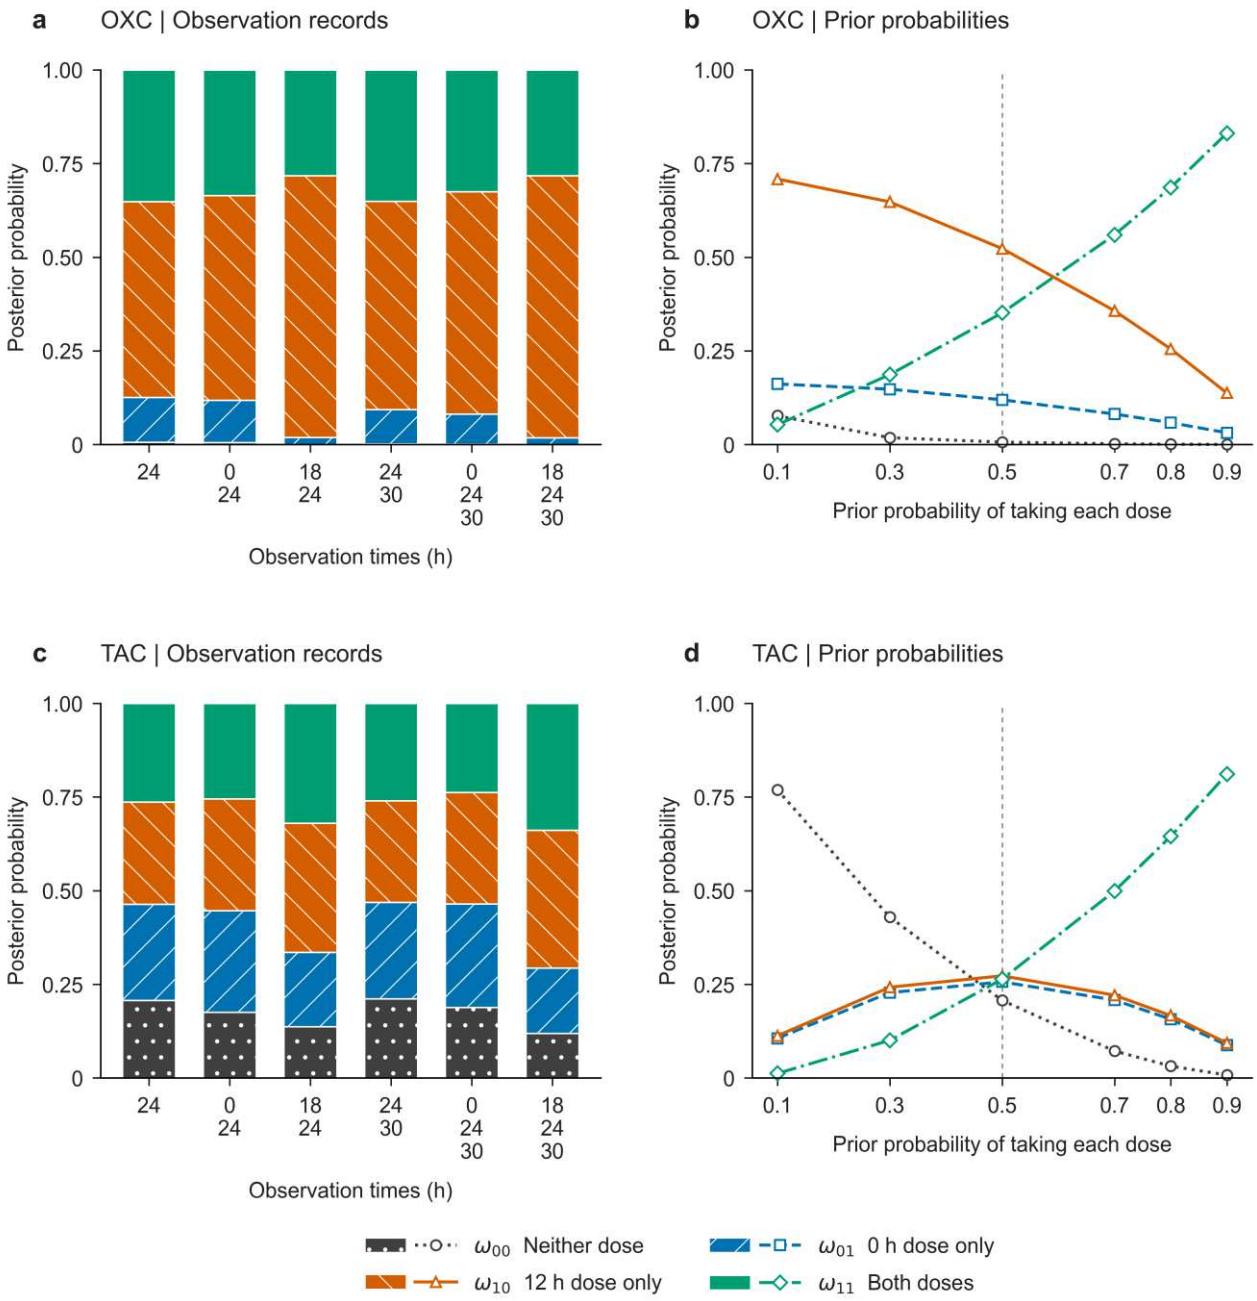

Fig. 3 Effects of observation records and prior probabilities on posterior probability distributions

Starting from each model's reference scenario, panels a and c show the posterior probabilities of the four candidate states after adding observations at different times in the oxcarbazepine (OXC) and tacrolimus (TAC) models, respectively. Panels b and d show the corresponding results as the prior probability of taking a dose was varied.

In a and c, the horizontal axis lists the observation times in each record. Additional observations were taken from the same reference concentration curve within each model; the 0 h observation preceded dosing. In b and d, the horizontal axis shows the prior probability of taking each dose, and vertical dashed lines indicate the reference setting.

Colored segments and curves represent neither dose taken (ω_{00}), only the 0 h dose taken (ω_{01}), only the 12 h dose taken (ω_{10}), and both doses taken (ω_{11}). Points on the curves show the posterior probabilities at the corresponding settings, and lines connect adjacent evaluated points.

Values are mean posterior probabilities after normalizing each of five runs of the Python package's Laplace calculation

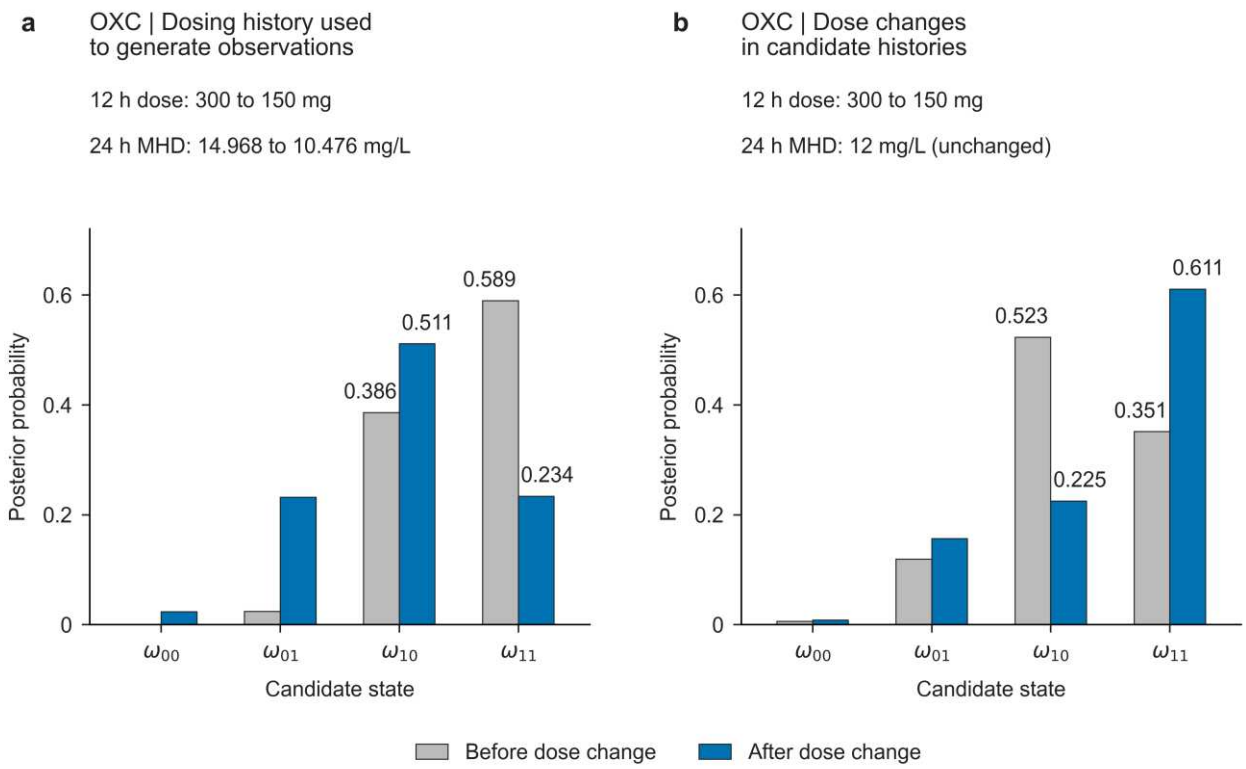

Fig. 4 Effects of changes in dosing histories on posterior probability distributions

Starting from the oxcarbazepine (OXC) reference scenario, panel a shows posterior probabilities after changing the dosing history used to generate observations. The 12 h dose was reduced from 300 to 150 mg, and the 24 h concentration of the monohydroxy derivative (MHD), the active metabolite of OXC, decreased from 14.968 to 10.476 mg/L. The corresponding concentrations were used as observations, with the reference candidate histories held fixed.

Panel b shows posterior probabilities after changing the doses in candidate histories. With the reference observation held fixed, the 12 h dose was reduced from 300 to 150 mg in all candidate histories containing that dose.

Gray and blue bars show posterior probabilities before and after the dose change, respectively. The four candidate states represent neither dose taken (ω_{00}), only the 0 h dose taken (ω_{01}), only the 12 h dose taken (ω_{10}), and both doses taken (ω_{11}). Numbers above bars give the corresponding posterior probabilities.

Values are mean posterior probabilities after normalizing each of five runs of the Python package's Laplace calculation

Supplementary Files

Online Resource 1. Software information, validation scenarios, reference calculations, evaluation metrics, and comparison rules.

Online Resource 2. Computational results for the OXC reference and sensitivity scenarios, including comparison results for the virtual patient approach.

Online Resource 3. Computational results for the TAC reference and sensitivity scenarios.

Online Resource 4. Code, input files, raw computational outputs, virtual patient comparison records, and reproduction materials.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- ESM1.pdf
- ESM3.pdf
- ESM2.pdf
- ESM4.zip