



Modeling Peak Expiratory Flow in Patients With Asthma and Quantifying Treatment Effects Using a Mixed-Effects Hidden Markov Model

Downloaded from: <https://research.chalmers.se>, 2026-06-24 19:29 UTC




Citation for the original published paper (version of record):

Jakobsson, L., Baaz, M., Leander, J. et al (2026). Modeling Peak Expiratory Flow in Patients With Asthma and Quantifying Treatment Effects Using a Mixed-Effects Hidden Markov Model. *CPT: Pharmacometrics and Systems Pharmacology*, 15(6). <http://dx.doi.org/10.1002/psp4.70281>

N.B. When citing this work, cite the original published paper.

ARTICLE OPEN ACCESS

Modeling Peak Expiratory Flow in Patients With Asthma and Quantifying Treatment Effects Using a Mixed-Effects Hidden Markov Model

 Ludvig Jakobsson^{1,2,3}  | Marcus Baaz²  | Jacob Leander³  | Philip Gerlee¹ | Mats Jirstrand^{2,4} 

¹Department of Mathematical Sciences, Chalmers University of Technology and University of Gothenburg, Gothenburg, Sweden | ²Fraunhofer-Chalmers Research Centre for Industrial Mathematics, Gothenburg, Sweden | ³Clinical Pharmacology and Quantitative Pharmacology, Clinical Pharmacology and Safety Sciences, R&D, AstraZeneca, Gothenburg, Sweden | ⁴Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden

Correspondence: Ludvig Jakobsson (ludvig.jakobsson@fcc.chalmers.se)

Received: 6 February 2026 | **Revised:** 20 April 2026 | **Accepted:** 22 May 2026

Keywords: dose response | mathematical modeling | mixed effects models | model-based drug development | respiratory

ABSTRACT

Clinical trials in asthma and chronic obstructive pulmonary disease often use exacerbation risk as the primary endpoint. However, exacerbations occur with low frequency, leading to long and costly clinical trials. Home-measured spirometry, which is becoming more common, provides an alternative and has previously been used to shorten the necessary trial duration. In this work, we develop a *mixed-effects hidden Markov model* (MEHMM) for analyzing home-measured *peak expiratory flow* (PEF), combining an observation model with a latent two-state disease process representing sustained periods of high and low PEF, respectively. An inference framework is implemented to estimate fixed and random effects together with measures of uncertainty. Data from a phase 2b dose-finding study of velsecorat in asthma are used to investigate dose–response relationships, complemented by an extensive simulation study. The results demonstrate reliable estimation of parameters and identify statistically significant treatment effects on multiple model components. These findings support the use of latent disease-state models for extracting meaningful information from home-measured spirometry.

1 | Introduction

In respiratory diseases, such as asthma and chronic obstructive pulmonary disease (COPD), lung function variability has been shown to [1, 2] be associated with the risk of acute worsening events, known as exacerbations. These events, during which patients experience increased symptoms and obstructed airflow, account for much of the mortality and treatment costs associated with respiratory diseases [3, 4], and are often chosen as the primary endpoint in late-stage clinical trials [5]. The incidence of exacerbations is low [6], leading to long and expensive clinical trials. This has prompted efforts to make fuller use of the available lung function data to shorten trial durations while preserving a strong link to exacerbation risk. Clinical advancements

have supported these efforts with home-measured spirometry, providing higher-resolution data. Detrended fluctuation analysis has previously been used to study the association between variability in home-measured peak expiratory flow (PEF) measurements and exacerbation risk [1, 7]. More recent work used a stochastic mixed effects model to quantify the effect of various sources of PEF variability on exacerbation risk and found that long-term variability, as opposed to day-to-day variability, was associated with both treatment effect and exacerbation risk [8].

Other recent advances in clinical drug development for respiratory diseases are the event endpoints CompEx Asthma [9] and COPDCompEx [10]. These events are defined by the occurrence of a diary event based on a change from baseline PEF and

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2026 The Author(s). *CPT: Pharmacometrics & Systems Pharmacology* published by Wiley Periodicals LLC on behalf of American Society for Clinical Pharmacology and Therapeutics.

Study Highlights

What Is the Current Knowledge on the Topic?

Lung function variability of patients with asthma fluctuates substantially over time, and increased variability has been linked to an elevated risk of worsening events. Home-measured spirometry can provide rich temporal data, but its irregular dynamics make it difficult to analyze using conventional modeling approaches.

What Question Did This Study Address?

Can a mixed-effects hidden Markov model be used to model home-measured peak expiratory flow (PEF) in asthmatic patients, and can any treatment effects be estimated using a modified stochastic approximation expectation–maximization algorithm (SAEM)?

What Does This Study Add to Our Knowledge?

A mixed-effects hidden Markov model was developed and used to analyze peak expiratory flow measurements in a clinical trial of patients with asthma. The results showed that parameter estimation and state-transition estimation were reliable and that several treatment effects were statistically significant. This serves as a proof of concept for modeling PEF using latent state-dependent variables, which are estimated using a modified SAEM algorithm.

How Might This Change Clinical Pharmacology or Translational Science?

The results add to the broader knowledge on variability in lung function in respiratory diseases, and on the utilization of high-resolution home-measured spirometry data. This has the potential to support clinical drug development through the refinement or development of respiratory endpoints.

reliever use, along with other symptom variables. The findings showed that CompEx Asthma resulted in 2.8 times more events than exacerbations and allowed for a 50% reduction in the number of patients needed, while allowing for a 3-month trial to evaluate the drug effect on exacerbation risk instead of the common study length of 12 months [9]. This has led to CompEx being used as a primary endpoint in several Phase 2 clinical trials in asthma and COPD (NCT06529419, NCT06376045, NCT05492877, NCT05251259) [11–14].

Our work investigates whether discrete, sustained worsening events can be identified from PEF time series at the individual level, rather than being defined by predetermined thresholds as in the PEF-component of the CompEx event. This is done by modeling PEF with a discrete-time mixed-effects hidden Markov model (MEHMM) [15]. In a paper by Delattre and Lavielle, an inference scheme was proposed for estimating population parameters, individual parameters, and the most likely path of the latent process for each individual in an MEHMM [16]. That approach adapts the standard nonlinear mixed-effects (NLME) framework by modifying the stochastic approximation expectation–maximization (SAEM) [17] algorithm so that observations from a hidden Markov model (HMM) may be considered. The implementation of the

proposed scheme used a combination of the Monolix software and MATLAB to perform the estimations.

We modeled PEF using an MEHMM with inference based on the proposed scheme. Our primary goal was to estimate treatment effects from disease-state dynamics, rather than modeling the full complexity of PEF trajectories. Accordingly, transitions were assumed to be instantaneous, providing a simplified representation at the observation level. A more extensive discussion on this assumption is given in Section 4. The model's robustness to different data scenarios was tested through a simulation study before data from a phase 2b clinical trial in asthma were analyzed.

The model and the corresponding inference algorithms were implemented in R, and the code is available in a public repository on GitHub (see [Supporting Information](#)). The implementation allows for common dose–response relationships to be included directly into the model, making inference about treatment effects practical and straightforward.

The article is organized as follows: First, the modeling framework and the estimation scheme are described in Sections 2.1 and 2.2. This is followed by a presentation of the PEF model and the simulation study framework in Sections 2.3 and 2.4, respectively. In Section 3.1, the results from the simulation study are presented, and in Section 3.2, the model is fitted to clinical trial data from a phase 2b study evaluating an asthma treatment. Finally, in Section 4, the results are discussed and summarized.

2 | Methods

2.1 | Mixed-Effects Hidden Markov Models

Hidden Markov models (HMMs) provide a flexible framework for modeling time series data with underlying discrete latent structures. In this setting, the observed data are assumed to arise from an unobserved, discrete-time Markov process representing latent states, where each state governs the outcome of observations through a state-dependent observation distribution. Mixed-effects hidden Markov models (MEHMM) [15] were defined as an extension of HMMs to include random effects for population approaches, allowing for between-subject variability in both transition dynamics and observation distributions. This extension is particularly important when individual time series provide limited information for subject-specific HMM inference, e.g., when no state transitions are observed. Covariates can be incorporated into MEHMMs to alter transition probabilities or observation distributions, enabling the models to account for both individual heterogeneity and treatment effects.

For a standard HMM, we let $Y_i = (Y_{i,1}, \dots, Y_{i,T})$ be the sequence of observed data for individual i , taking on discrete or continuous values, and let $Z_i = (Z_{i,1}, \dots, Z_{i,T})$ be the corresponding latent state sequence with $Z_{i,t}$ evolving with subject-specific transition probabilities

$$p_{i,j,l} = P(Z_{i,t} = l | Z_{i,t-1} = j)$$

where $j, l = 0, \dots, L$ represent the possible latent states. The observations $\{Y_{i,t}\}_{t=1}^T$ are then assumed conditionally independent

given the latent state, with an observation distribution $P(Y_{i,t} | Z_{i,t})$ that may depend on other parameters as well. A schematic view of an HMM is shown at the top of Figure 1.

In the extension to MEHMMs, individual heterogeneity is introduced by allowing the parameters of the HMM to vary across individuals via random effects. Let ψ_i denote the vector of individual-specific parameters for individual i , including transition probabilities and any parameters of the observation distribution. Assume that we have a transformation g so that $\phi_i = g(\psi_i)$ is Gaussian in each component. The transformed parameters ϕ_i are then defined as

$$\phi_i = \mu + \beta(x_i) + \eta_i$$

where $\eta_i \sim \mathcal{N}(0, \Omega)$ is the vector of random effects with covariate matrix Ω , μ is the vector of population reference values, and x_i is a vector of covariates giving effects $\beta(x_i)$. The population parameters of the model are denoted $\theta = (\mu, \beta, \Omega)$, with β containing all covariate parameters.

2.2 | Estimation Scheme

When estimating the parameters of a standard HMM, a version of the expectation–maximization (EM) algorithm is most

used. This version is often referred to as the Baum–Welch algorithm [18] and utilizes the forward-backward algorithm [18] in the expectation step to solve the smoothing problem, which is to compute $P(Z_t | Y_{1:T})$, for all $t = 1, \dots, T$.

Building on the MEHMM framework and the Baum–Welch algorithm, an estimation scheme [16] was proposed using the stochastic approximation expectation–maximization algorithm (SAEM) [17]. The method enables the estimation of both fixed and random effects, including covariate effects, as well as individual-specific empirical Bayes estimates (EBE) along with latent process decoding via the Viterbi [18] algorithm and estimation of standard errors. The following notes on the method are included for the convenience of the reader.

The standard SAEM algorithm for NLME models requires a slight modification to work with MEHMMs. In the simulation step, the goal is to sample from the posterior distribution $P(\phi_i | Y_i, \theta^{(k)})$, where $k = 1, \dots, K$ is the current iteration of the algorithm. When sampling from the posterior distribution is intractable, samples are often drawn via a Markov chain Monte Carlo (MCMC) method, which requires evaluation of the full likelihood

$$P(Y_i, \phi_i^{(m)} | \theta^{(k)}) = P(Y_i | \phi_i^{(m)}, \theta^{(k)}) \times P(\phi_i^{(m)} | \theta^{(k)})$$

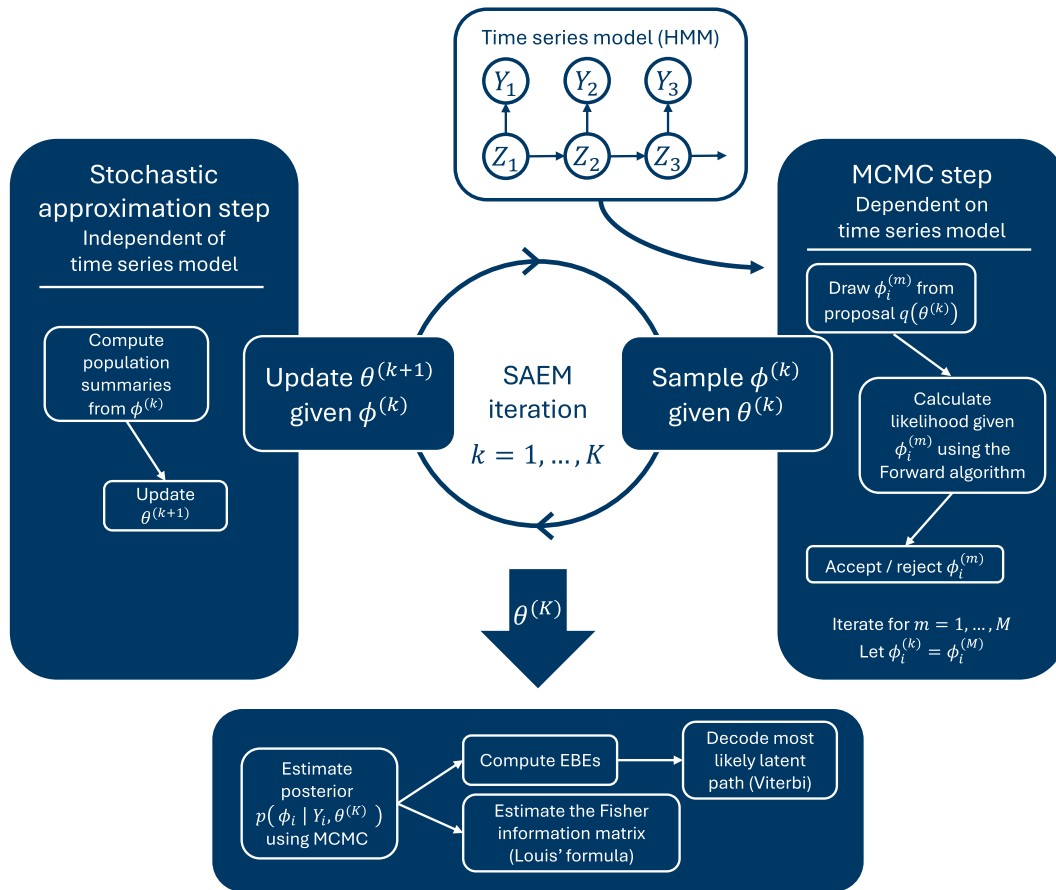


FIGURE 1 | A schematic view of the estimation framework used in this work. The algorithm alternates between a Markov chain Monte Carlo (MCMC) step where patient-specific parameters $\phi_i^{(k)}$ are sampled conditioned on the current population parameter estimates $\theta^{(k)}$ and the patient’s time series data, and a stochastic approximation step where population summaries are computed from the sampled $\phi_i^{(k)}$ and used to compute $\theta^{(k+1)}$. After convergence, the final $\hat{\theta} = \theta^{(K)}$ are used to compute empirical Bayes estimates, estimate the Fisher matrix, and decode the most likely latent path.

for each set of proposed parameters $\phi_i^{(m)} \sim q(\theta^{(k)})$, $m = 1, \dots, M$, given a proposal density q . In classical NLME models, this can be evaluated analytically. However, for MEHMMs, evaluation of the first factor above requires the use of the forward recursions of the forward-backward algorithm since the latent process is unknown. This modification is straightforward when implementing the SAEM algorithm manually, but requires workarounds when using existing software.

The likelihood calculation is simplified by noting that Y_i is conditionally independent of $\theta^{(k)}$ given $\phi_i^{(m)}$, that is

$$P(Y_i | \phi_i^{(m)}, \theta^{(k)}) = P(Y_i | \phi_i^{(m)})$$

This further simplifies the maximization step of the SAEM algorithm and the approximation of the standard errors since only the second factor of the likelihood now depends on the population parameters θ , and $P(\phi_i^{(m)} | \theta^{(k)})$ is a known distribution.

Standard errors can be calculated from the obtained population parameter estimates $\hat{\theta}$ using Louis' formula to estimate the Fisher information matrix [19, 20]. The expected values present in Louis' formula are estimated by running an MCMC chain to approximate the posterior distribution of each parameter given the final population parameter estimates $\hat{\theta}$. The entire algorithm is described schematically in Figure 1.

2.3 | Modeling Peak Expiratory Flow

Here, we will present an MEHMM for modeling PEF measurements in respiratory diseases. Building on the framework presented above, we assume that each patient's observed PEF depends on a latent process with two states, *high PEF* and *low PEF*, representing the patient being in a normal health condition and in a worsened health condition, respectively, as well as several individual-specific model parameters following some population distribution. The patients are assumed to be given either a placebo drug or an active drug at some dose x_i . The treatment effect is included in the model as a dose-dependent additive effect on each transformed parameter, as in Section 2.1. Dose was selected as the exposure variable since, for inhaled therapies, the pharmacologically relevant exposure is commonly the local drug concentration in the lung and not plasma concentration [21]. However, the model could be modified to consider fuller pharmacokinetic-pharmacodynamic relationships, including, e.g., a delayed response and other routes of administration.

The PEF measurement $Y_{i,t}$ and the state transition probabilities for subject $i = 1, \dots, N$ at time $t = 1, \dots, T$ are modeled as

$$Y_{i,t} = m_{i,Z_{i,t}} + \varepsilon_{i,t},$$

$$P(Z_{i,t} = l | Z_{i,t-1} = j) = p_{ij,l},$$

where $\varepsilon_{i,t} \sim N(0, \sigma_i^2)$ and $j, l \in \{0, 1\}$, with individual-specific model parameters

$$\psi_i = (m_{i,0}, d_i, \sigma_i^2, p_{i,0,1}, p_{i,1,0}).$$

The parameter $d_i \in [0, 1]$ defines the relative drop from high PEF $m_{i,0}$ to low PEF $m_{i,1} = d_i \times m_{i,0}$. Note that the pair $(p_{i,0,1}, p_{i,1,0})$ fully defines the transition matrix P_i for individual i since each row sums to one. The individual parameters are in turn transformed using a transformation $\phi_i = g(\psi_i)$ that applies logarithmic or logit functions depending on the parameter and are then modeled as:

$$\log(m_{i,0}) = \mu_1 + \beta_1(x_i) + \eta_{1,i}$$

$$\text{logit}(d_i) = \mu_2 + \beta_2(x_i) + \eta_{2,i}$$

$$\log(\sigma_i^2) = \mu_3 + \beta_3(x_i) + \eta_{3,i}$$

$$\text{logit}(2p_{i,0,1}) = \mu_4 + \beta_4(x_i) + \eta_{4,i}$$

$$\text{logit}(2p_{i,1,0}) = \mu_5 + \beta_5(x_i) + \eta_{5,i}$$

where $\eta_{c,i} \sim N(0, \omega_c^2)$ for each $i = 1, \dots, N$ and $c = 1, \dots, 5$. The factor 2 in the definitions of $p_{i,0,1}$ and $p_{i,1,0}$ restricts these parameters to the interval [0,0.5], reflecting the assumption of sustained rather than rapid transitions, which is consistent with the aim of this work. The random effects are assumed to be independent across the components of ϕ_i , meaning that the covariate matrix Ω is diagonal. This simplifies the MCMC algorithm since the samples can be drawn independently across parameters, but a full covariance matrix could also be considered.

Four types of dose-response relationships were implemented in the code:

$$\text{Categorical } \beta(x_i) = \sum \beta_p(x_{i,p})$$

$$\text{Constant } \beta(x_i) = \beta \mathbf{1}_{\{x_i > 0\}}$$

$$\text{Linear } \beta(x_i) = \beta x_i$$

$$\text{Emax } \beta(x_i) = \frac{E_{\max} x_i}{ED_{50} + x_i}$$

where x_i is a continuous dose in the constant, linear, and Emax models, while $x_i = (x_{i,1}, \dots, x_{i,p}) \in \{0, 1\}^p$ with $\sum_p x_{i,p} = 1$ is a categorical dose in the categorical model. To explore treatment effects on each parameter, we first fit a flexible model using categorical treatment indicators, henceforth denoted the exploratory model. This is analogous to an ANOVA-type analysis. After running the estimation scheme, non-significant treatment effects were removed to obtain a more parsimonious model. Finally, based on the pattern observed in the exploratory step, we evaluated parametric dose-response models.

2.4 | Simulation Study

Datasets were simulated in R using the PEF model above with population parameter values chosen to reproduce the mean and variance of twice-daily home-measured PEF observed in clinical trials. A similar number of patients per treatment group and similar time series lengths were also used. One active treatment

TABLE 1 | Summarized results from estimating parameters using simulated data sets of length $T = 200$ with one active treatment group and a placebo group. The treatment effects β are categorical and additive on the transformed individual-specific parameters for the active group.

Parameter	True value	Mean estimate	Standard deviation of estimate	Mean standard error	Coverage (%)
High PEF m_0					
μ_1	5.521	5.526	0.030	0.032	97
β_1	0.250	0.246	0.043	0.045	96
ω_1^2	0.100	0.101	0.010	0.010	95
Drop d					
μ_2	1.000	0.994	0.054	0.050	91
β_2	0.000	0.001	0.076	0.072	93
ω_2^2	0.250	0.250	0.026	0.026	96
Inter-state variability σ^2					
μ_3	5.704	5.702	0.024	0.025	96
β_3	0.000	0.004	0.034	0.035	96
ω_3^2	0.050	0.051	0.007	0.006	94
Transition probability $p_{0,1}$					
μ_4	-2.197	-2.181	0.086	0.080	93
β_4	-0.500	-0.527	0.122	0.115	94
ω_4^2	0.400	0.388	0.071	0.068	92
Transition probability $p_{1,0}$					
μ_5	-2.197	-2.203	0.081	0.077	94
β_5	0.000	-0.003	0.124	0.111	92
ω_5^2	0.300	0.312	0.072	0.062	92

arm was simulated alongside a placebo, and a categorical treatment effect was used for all parameters. The algorithms described in Section 2.2 were then used for the estimation of parameters from the simulated datasets.

Given the point estimates $\hat{\theta}$ of fixed and random effects in each run, 200 samples from the posterior distribution $P(\phi_i | y_i, \hat{\theta})$ were drawn for each subject i using the MCMC algorithm. These samples were then used to calculate EBEs as the means of the posterior samples and to estimate the Fisher information matrix. The EBEs were, in turn, used to calculate the most likely latent path via the Viterbi algorithm.

The point estimates $\hat{\theta}$ were assessed by bias, calculated as the average difference between the estimates and the true parameters used to simulate the data. The EBEs were evaluated by their predictive accuracy with respect to the true simulated individual parameters and by the η -shrinkage [22]. The accuracy of the paths generated by the Viterbi algorithm was measured by confusion matrices for each simulated dataset. Finally, the standard error estimates were assessed per-parameter by their average difference from the standard deviation of the parameter estimates $\hat{\theta}$ and by the coverage of 95%-confidence intervals for $\hat{\theta}$.

2.5 | Clinical Trial Data

Patient-level data from a phase 2b dose-finding study of velsecorat in asthmatic patients (NCT03622112) [23] was used in model development and in the analysis. The clinical trial featured 5 active treatment arms and a placebo group with twice-daily home-measured PEF collected during a 12-week treatment period. Before the analysis, all informed consent forms were reviewed for data reuse in accordance with AstraZeneca data-sharing rules.

3 | Results

3.1 | Simulation Study

The results from 200 simulated datasets are summarized in Table 1 which shows the true parameter values, the mean and standard deviation of the population parameter estimates, the mean standard error estimates, and the coverage of the 95% confidence intervals calculated from the per-dataset estimates. The results are shown for models with a categorical treatment effect on all parameters and datasets with $N = 100$ subjects per treatment group and timeseries lengths $T = 200$. The analysis was also carried out for datasets with timeseries lengths $T = 50$

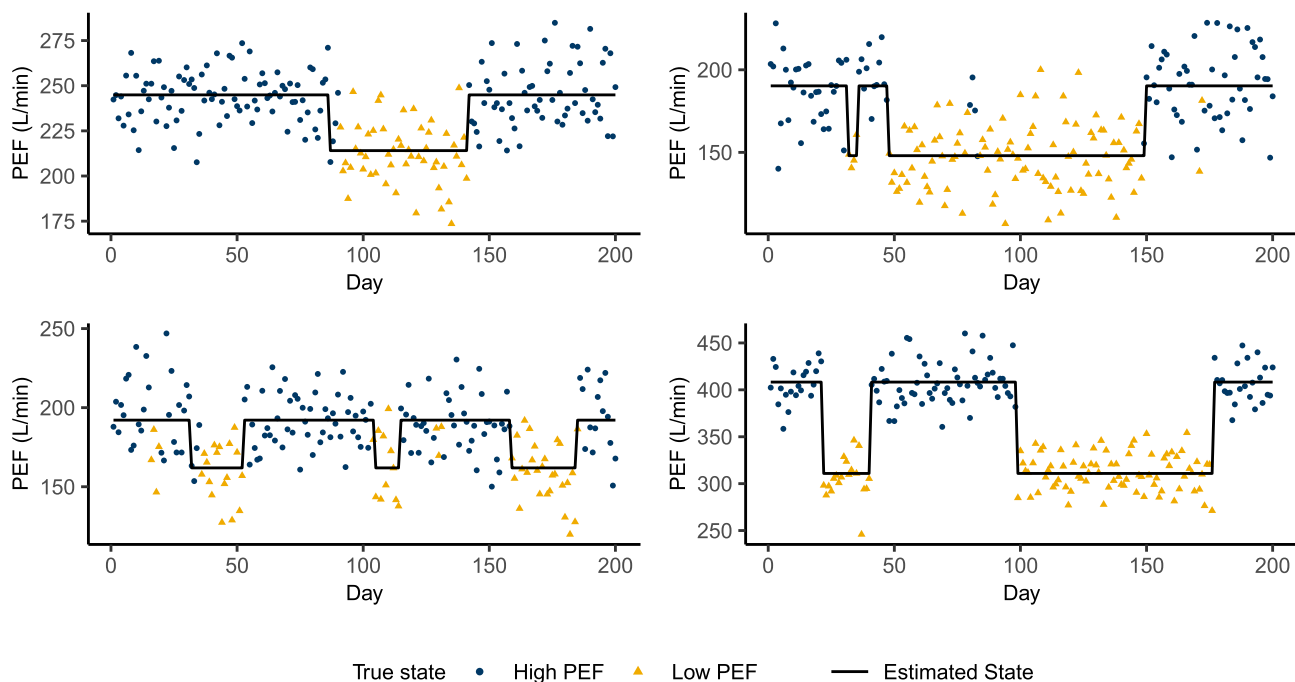


FIGURE 2 | Four simulated individual time series from a mixed-effects hidden Markov model with the population parameters from Table 1. The true latent state for each observation is indicated by the point color, and the estimated most likely latent path from the fitted model is shown as an overlaid line. The y-position of the line corresponds to the estimated mean of each state.

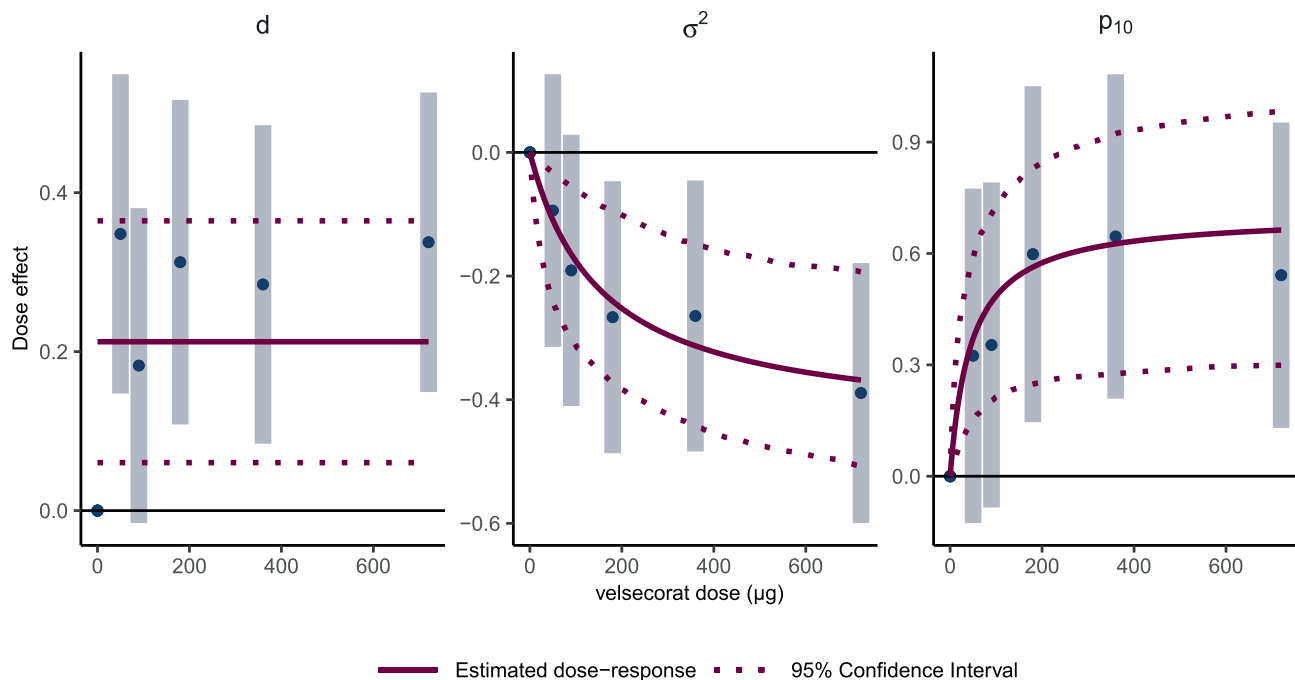


FIGURE 3 | Estimated dose effect for the drop parameter d , the inter-state variability σ^2 , and the transition probability $p_{1,0}$. The fitted dose-response curve (constant effect for d , and an Emax model for σ^2 , and $p_{1,0}$) is shown as a solid line with 95%-confidence interval as dotted lines. The points and corresponding bars indicate the estimated dose effect in each treatment group from the exploratory model.

and $T = 1000$, respectively, as reported in Tables S1 and S2 in the Supporting Information. For the time series of length $T = 200$, all estimates converged, and the coverage ranged between 91%–97%, with a mean of 94%.

The average per-parameter relative difference between the true individual parameters and the EBEs across all simulated datasets

were small for the parameters m_0 , d , and σ^2 , and slightly larger for the transition probabilities (–10% and –12%, respectively). The EBEs are plotted against the true individual parameters, together with a coefficient of determination (R^2) corresponding to a linear regression, in Figure S1 in the Supporting Information. Figure S2 shows the shrinkage in each model parameter by comparing the estimated population variance to the sample variance of the EBEs.

Figure 2 shows four typical simulated time series colored by true latent states and estimated latent-state paths. Finally, given each subject's time series and latent-state path, the average accuracy over all data sets, calculated as the sum of the diagonal of the average confusion matrix, was 98.9%.

3.2 | Application to Clinical Trial Data

The parameter estimates from the exploratory model with categorical treatment effects are shown in Table S3 in the [Supporting Information](#) along with the estimated standard errors. This model allowed each treatment group to have its own effects and made no assumptions about dose–response relationships. Several treatment groups showed statistically significant effects vs. placebo for the drop parameter d , the inter-state variability σ^2 , and the transition probability $p_{1,0}$. The treatment effects for the remaining parameters m_0 and $p_{0,1}$ were subsequently removed before dose–response models were considered.

Based on visual inspection, an Emax model was selected and fitted for the parameters σ^2 and $p_{1,0}$, whereas a constant effect was

TABLE 2 | Mean model parameters in the placebo group and the highest-dose treatment group, with 95%-confidence intervals.

Parameter	Placebo	720 μg
m_0	317 (309, 326)	317 (309, 326)
d	0.879 (0.863, 0.893)	0.900 (0.895, 0.905)
σ	30.5 (28.4, 32.8)	25.4 (24.5, 27.4)
$p_{0,1}$	0.031 (0.028, 0.034)	0.031 (0.028, 0.034)
$p_{1,0}$	0.023 (0.017, 0.031)	0.043 (0.037, 0.050)

assumed for the parameter d , giving equal effect for all active treatment groups and no effect for the placebo group. The resulting fitted dose–response curves are overlaid with the categorical treatment effects obtained in the exploratory model and are shown in Figure 3. For the Emax dose–response curves, confidence intervals were obtained by sampling parameters using the estimated covariance matrix and computing the 2.5% and 97.5% quantiles for the response. For the constant dose–response curve and the categorical treatment effects, confidence intervals were calculated from the corresponding estimated standard errors. All population parameter estimates and standard errors for this model are given in Table S4 in the [Supporting Information](#). Given these estimates, the final mean parameter values, which have been transformed using $g^{-1}(\phi)$, for the placebo group and the highest dose group are given in Table 2, along with 95% confidence intervals.

Akaike's information criteria (AIC) for this model (denoted as 'mixed') was compared to the AIC of the exploratory model ('categorical') and a model without any dose effects ('none'). This is shown in Table S5 in the [Supporting Information](#) along with the R^2 of a linear regression between the PEF measurements and the predictions obtained for each model. Furthermore, Figure S3 shows the distribution of the estimated random effects obtained from the mixed dose–response model.

Given the mixed dose–response model, PEF time series with their estimated most likely latent paths, calculated using the EBEs, are shown in Figure 4 for selected individuals. Some shrinkage was observed in the EBEs, mainly for the transition probabilities (η -shrinkage 66% and 57%, respectively), which is illustrated in Figure 5 by the estimated population distributions along with the empirical distributions obtained from the EBEs. Figure 5 also shows predicted PEF plotted against the PEF measurements together with the R^2 of a linear regression between the two.

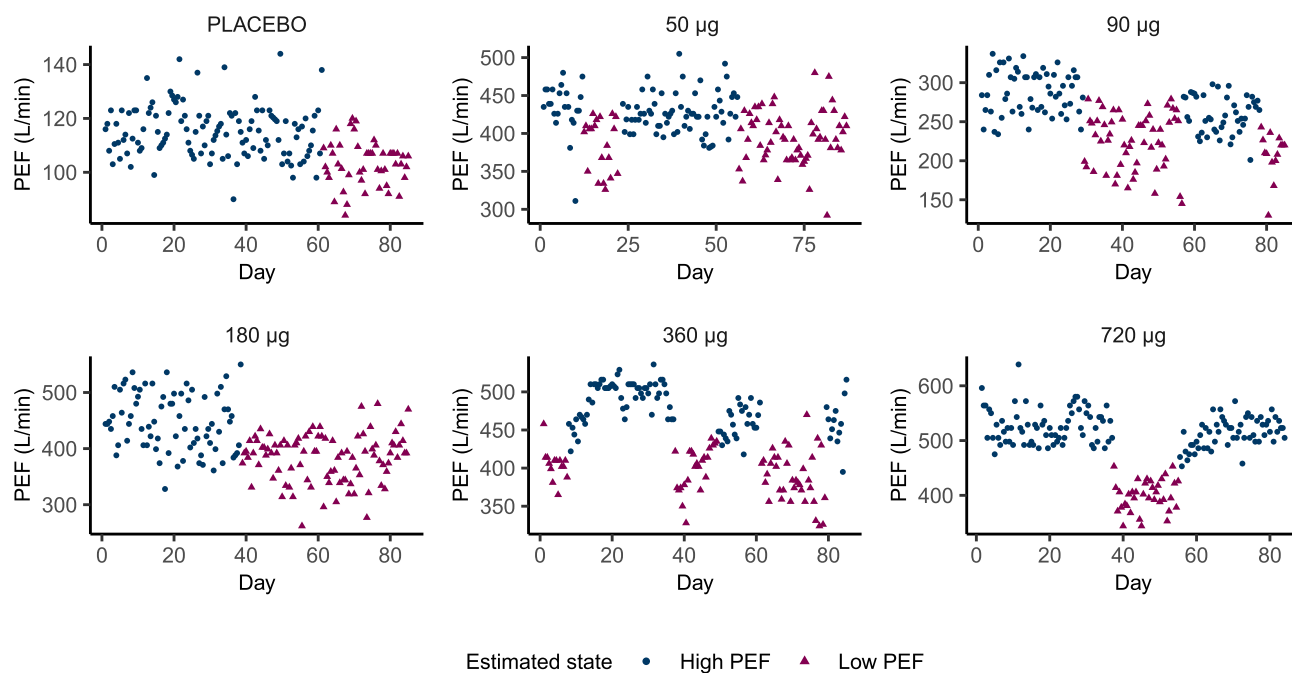


FIGURE 4 | Six individual PEF time series from the clinical trial, one from each treatment group. The most likely latent path, estimated using the Viterbi algorithm and the individual empirical Bayes estimates, is indicated by the point color and shape.

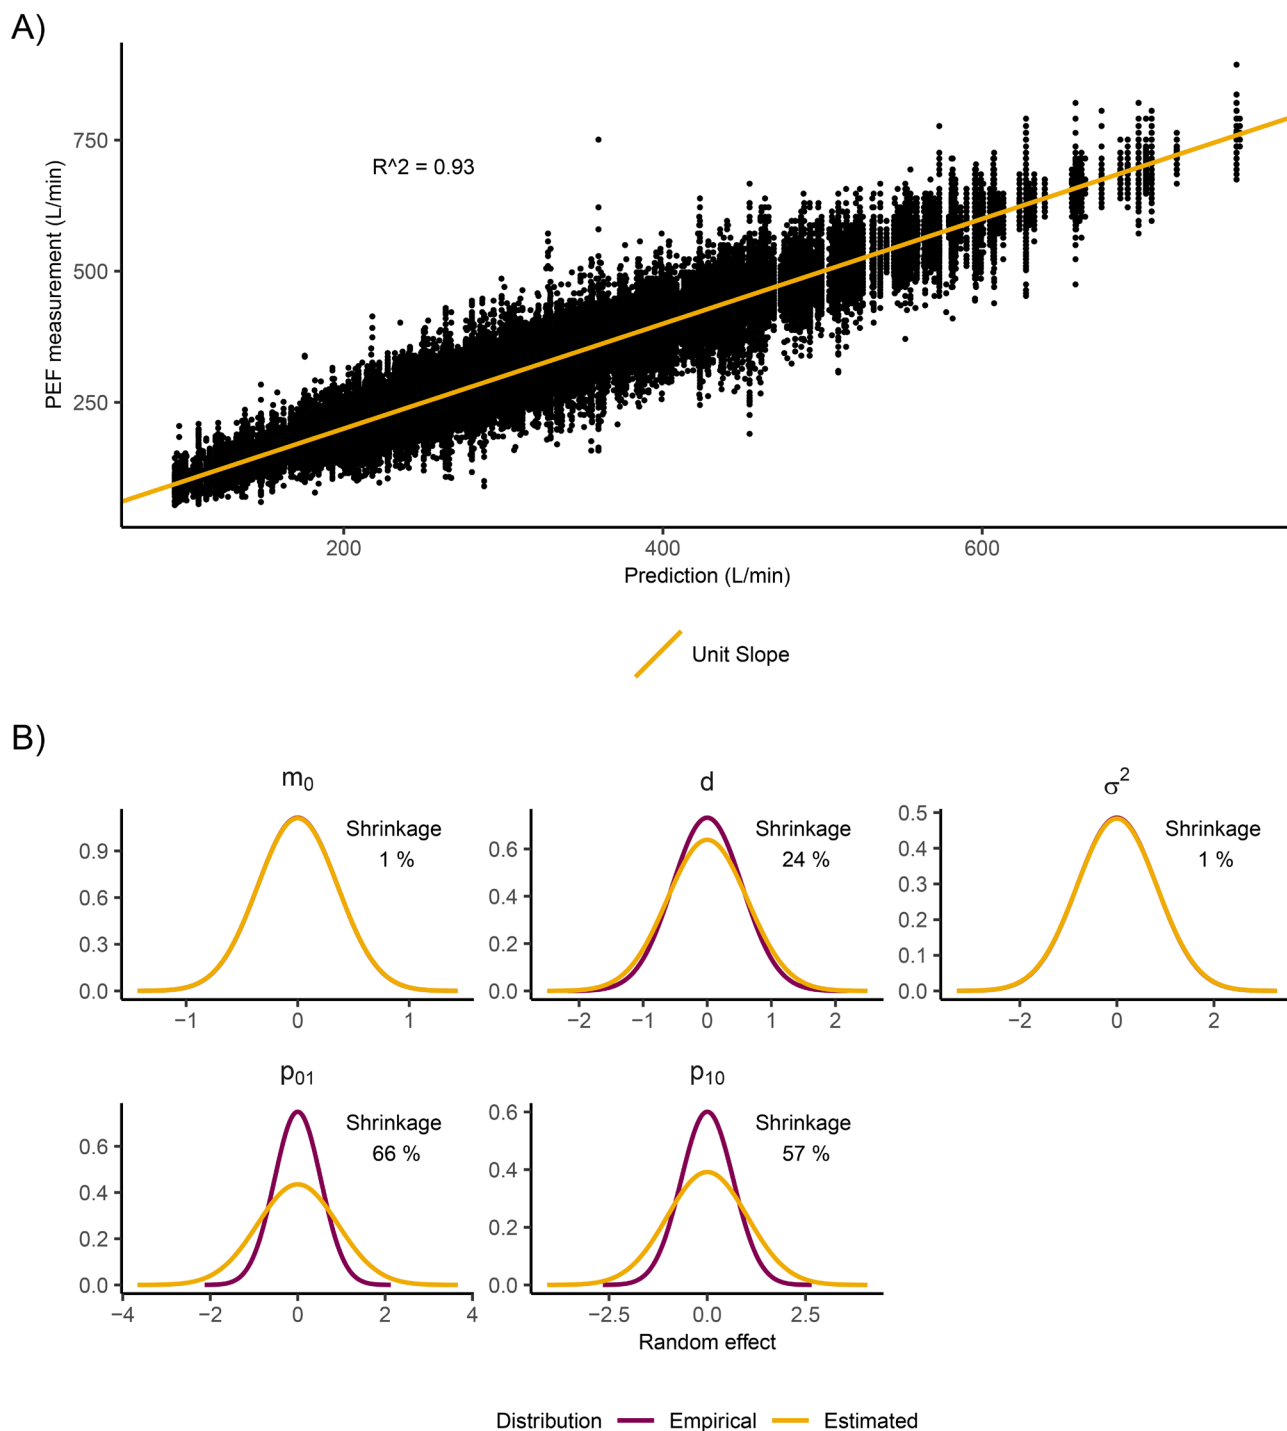


FIGURE 5 | (A) Individual PEF predictions plotted against PEF measurements with a unit slope line and the coefficient of determination (R^2) of a linear regression. (B) Distribution plots showing η -shrinkage for each model parameter in the clinical trial. The estimated distributions have variance equal to the estimated population variance ω_c^2 for $c = 1, \dots, 5$ and the empirical distributions have variance equal to the sample variance of the individual empirical Bayes estimates.

4 | Discussion

This work considered an application of the SAEM algorithm adapted to handle observations described by an MEHMM, as proposed by Delattre and Lavielle [16]. A two-state Gaussian MEHMM was developed to model home-measured PEF with potential treatment effects on all model parameters. The implementation of the estimation scheme was verified in a simulation

study and subsequently applied to home-measured spirometry data from a clinical trial of patients with asthma to evaluate its ability to distinguish treatment effects. The results of the simulation study showed that the estimation scheme produces estimates of fixed and random effects with low bias and standard errors that reflect the variability in parameter estimates obtained from repeated simulations of datasets. Estimation of individual-specific EBEs was found to reflect the true simulated

parameters to some degree, but shrinkage was observed. This indicates limited patient-level identifiability for time series of length $T = 200$ and motivates the use of random effects to get reliable population parameter estimates. The most likely latent path, estimated using the Viterbi algorithm, was compared to the true latent path by a mean confusion matrix over all simulated datasets and by visual inspection of latent paths plotted together with simulated data. Figure 2 shows how some short state transitions are difficult to estimate, but overall, the results show a reliable estimation of state transitions given EBEs, despite the observed shrinkage.

The application of the model and estimation scheme to home-measured PEF time series data shows promising results. Overall, the obtained parameter estimates were in a reasonable range, and the relative standard errors were in most cases small. Figure 4 shows that the inferred latent paths broadly follow the expected transition dynamics with sustained periods in both states. Furthermore, statistically significant treatment effects were found for several treatment groups for the drop parameter d , the inter-state variability σ^2 , and the transition probability $p_{1,0}$ without any assumptions on the dose–response relationship. These treatment effects were consistent with the expected direction of effect since a higher value of the parameter d indicates less severe worsened states, a lower value of the parameter σ^2 indicates more stable measurements within each state, and a higher value of the parameter $p_{1,0}$ indicates faster recovery from the low PEF state. No discernible effect was observed on the parameter m_0 , despite prior expectation that the treatment would increase average PEF. A possible explanation for this is the substantial dropout observed in the placebo group of the clinical trial, which may have inflated the estimated mean value of m_0 , thus underestimating the dose effects. These cases of informative censoring may be addressed in future work by including a risk model of early discontinuation in the PEF model presented here and estimating all parameters jointly, similar to the approach of Król et al. [24]. Another explanation for the lack of treatment effect is the assumption of an instantaneous effect following randomization. In reality, the effect is likely to increase gradually over the course of the study until it reaches its maximum. An extension of our work could consider a turnover model for one or several treatment effects to account for this delay, see, e.g., Leander et al. [8].

Two constraints were imposed on the individual parameters: An upper limit of 0.5 for the transition probabilities and a diagonal covariance matrix Ω . The upper limit was imposed to avoid unrealistically rapid transitions driven by outlier measurements. Without this constraint, transition probabilities were at times estimated close to one, corresponding to near-instant switches between the states. The independence between parameters was chosen for ease of implementation. However, moderate correlations were observed between the EBEs of the parameters m_0 , d , and σ^2 , suggesting that the covariance structure may be somewhat restrictive and that allowing for correlations in selected parameters could provide a more realistic representation of the inter-individual variability.

The parameters σ^2 and $p_{1,0}$ showed clear signs of saturation in the exploratory model, motivating the use of an Emax dose–response model. The resulting fitted curves reproduced the

categorical estimates closely and with reduced uncertainty. However, while the Emax parameter estimates were precise, the standard errors of the ED50 parameter estimates were substantial. This is expected when the available dose levels do not adequately characterize the steep part of the dose–response curve, but it nonetheless reduces confidence that Emax is the most appropriate representation of the underlying dose effect. Despite this, the model selection criteria AIC favored this model over the model without treatment effects and the model with categorical treatment effects. Furthermore, the estimated random effects from this model, stratified on treatment group, were centered at zero for all model parameters.

Future work could extend the MEHMM model to include a temporal dependence in the observation model, e.g., by modeling the observations as an autoregressive or stochastic process [8]. This would allow for smoother state transitions on the observation level, which are present in the data in Figure 4, instead of the assumed instant transitions considered in this work, and may offer more insight into the role of discrete states in lung function. Such an extension would require another method for computing the likelihood on the patient level, but would otherwise be straightforward to include within the provided implementation of the SAEM algorithm. A final natural extension is to broaden the model beyond only PEF time series. Other variables, such as daily patient-reported symptom scores and rescue medication use, are known to be affected by treatment and are informative of patient health [25, 26]. Future work could therefore extend the model to include these variables alongside PEF.

Overall, the results show potential for modeling home-measured PEF using a latent time-dependent discrete disease-state variable with the possibility to include treatment effects in the model parameters. If the proposed model is to support clinical decision making, additional work needs to determine whether the estimated treatment effects in the parameters correlate with other clinical outcomes in asthma, for instance, the risk of experiencing an exacerbation or a CompEx event.

Author Contributions

L.J., M.B., J.L., P.G., and M.J. wrote the manuscript; L.J., M.B., J.L., P.G., and M.J. designed the research; L.J. and J.L. performed the research; L.J. analyzed the data.

Acknowledgments

The authors acknowledge AstraZeneca for providing the data, computing facilities, and funding. L.J. participated in a PhD program supported by AstraZeneca. The authors also acknowledge that ChatGPT (OpenAI, San Francisco, CA, USA) was used in the drafting of this manuscript to improve language and readability.

Funding

The work was funded by AstraZeneca.

Conflicts of Interest

J. Leander is an employee of AstraZeneca and may own stock/stock options. All other authors declared no competing conflicts of interest for this work.

Data Availability Statement

R code sufficient for performing the simulation and estimation in this work is publicly available at <https://github.com/fcc-ludvig-jakobsson/pef-mehmm.git>.

References

1. U. Frey, T. Brodbeck, A. Majumdar, et al., "Risk of Severe Asthma Episodes Predicted From Fluctuation Analysis of Airway Function," *Nature* 438, no. 7068 (2005): 667–670, <https://doi.org/10.1038/nature04176>.
2. G. C. Donaldson, T. A. R. Seemungal, J. R. Hurst, and J. A. Wedzicha, "Detrended Fluctuation Analysis of Peak Expiratory Flow and Exacerbation Frequency in COPD," *European Respiratory Journal* 40, no. 5 (2012): 1123–1129, <https://doi.org/10.1183/09031936.00180811>.
3. T. A. Seemungal, J. R. Hurst, and J. A. Wedzicha, "Exacerbation Rate, Health Status and Mortality in COPD - a Review of Potential Interventions," *COPD* 4 (2009): 203–223, <https://doi.org/10.2147/COPD.S3385>.
4. M. Engelkes, M. A. de Ridder, E. Svensson, et al., "Multinational Cohort Study of Mortality in Patients With Asthma and Severe Asthma," *Respiratory Medicine* 165 (2020): 105919, <https://doi.org/10.1016/j.rmed.2020.105919>.
5. N. Marco-Ariño, A. Jauhiainen, J. Betts, and C. A. Da Silva, "Complexity of Design in Early Phase Respiratory Clinical Trials; Evaluating Impact of Primary Endpoints and Sponsor Size," *Contemporary Clinical Trials Communications* 22 (2021): 100793, <https://doi.org/10.1016/j.conctc.2021.100793>.
6. E. Wang, M. E. Wechsler, T. N. Tran, et al., "Characterization of Severe Asthma Worldwide: Data From the International Severe Asthma Registry," *Chest* 157, no. 4 (2020): 790–804, <https://doi.org/10.1016/j.chest.2019.10.053>.
7. C. Thamrin, R. Nydegger, G. Stern, et al., "Associations Between Fluctuations in Lung Function and Asthma Control in Two Populations With Differing Asthma Severity," *Thorax* 66, no. 12 (2011): 1036–1042, <https://doi.org/10.1136/thx.2010.156489>.
8. J. Leander, M. Jirstrand, U. G. Eriksson, and R. Palmér, "A Stochastic Mixed Effects Model to Assess Treatment Effects and Fluctuations in Home-Measured Peak Expiratory Flow and the Association With Exacerbation Risk in Asthma," *CPT: Pharmacometrics & Systems Pharmacology* 11, no. 2 (2022): 212–224, <https://doi.org/10.1002/psp4.12748>.
9. A. L. Fuhlbrigge, T. Bengtsson, S. Peterson, et al., "A Novel Endpoint for Exacerbations in Asthma to Accelerate Clinical Development: A Post-Hoc Analysis of Randomised Controlled Trials," *Lancet Respiratory Medicine* 5, no. 7 (2017): 577–590, [https://doi.org/10.1016/S2213-2600\(17\)30218-7](https://doi.org/10.1016/S2213-2600(17)30218-7).
10. C. F. Vogelmeier, A. Fuhlbrigge, A. Jauhiainen, et al., "COPDCompEx: A Novel Composite Endpoint for COPD Exacerbations to Enable Faster Clinical Development," *Respiratory Medicine* 173 (2020): 106175, <https://doi.org/10.1016/j.rmed.2020.106175>.
11. AstraZeneca, A Phase 2b, Randomised, Double-Blind, Placebo-Controlled Dose Range-Finding Study to Assess Efficacy and Safety of Multiple Dose Levels of Inhaled AZD8630 Given Once Daily for 12 Weeks in Adults With Uncontrolled Asthma at Risk of Exacerbations (LEVANTE) accessed at, March 23, 2026, <https://clinicaltrials.gov/study/NCT06529419>.
12. Amgen, A Phase 2, Randomized, Double-Blind, Placebo-Controlled, Dose Ranging Study to Assess the Efficacy and Safety of Rocatinlimab in Adult Subjects With Moderate-to-Severe Asthma accessed at, March 23, 2026, <https://clinicaltrials.gov/study/NCT06376045>.
13. AstraZeneca, A Phase IIa Randomised, Double Blind, Placebo Controlled, Parallel Arm, Multi-Centre Study to Evaluate the Efficacy and Safety of Mitiperstat (AZD4831), for 12–24 Weeks, in Patients With Moderate to Severe Chronic Obstructive Pulmonary Disease (COPD) accessed at, April 10, 2026, <https://clinicaltrials.gov/study/NCT05492877>.
14. AstraZeneca, A Phase 2a, Randomised, Double-Blind, Placebo-Controlled Study to Assess Efficacy and Safety of Atuliflapon Given Orally Once Daily for Twelve Weeks in Adults With Moderate to Severe Uncontrolled Asthma accessed at, March 23, 2026, <https://clinicaltrials.gov/study/NCT05251259>.
15. R. M. Altman, "Mixed Hidden Markov Models: An Extension of the Hidden Markov Model to the Longitudinal Data Setting," *Journal of the American Statistical Association* 102, no. 477 (2007): 201–210, <https://doi.org/10.1198/01621450600001086>.
16. M. Delattre and M. Lavielle, "Maximum Likelihood Estimation in Discrete Mixed Hidden Markov Models Using the SAEM Algorithm," *Computational Statistics & Data Analysis* 56, no. 6 (2012): 2073–2085, <https://doi.org/10.1016/j.csda.2011.12.017>.
17. B. Delyon, M. Lavielle, and E. Moulines, "Convergence of a Stochastic Approximation Version of the EM Algorithm," *Annals of Statistics* 27, no. 1 (1999): 94–128, <https://doi.org/10.1214/aos/1018031103>.
18. L. R. Rabiner, "A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition," *Proceedings of the IEEE* 77, no. 2 (1989): 257–286, <https://doi.org/10.1109/5.18626>.
19. E. Kuhn and M. Lavielle, "Maximum Likelihood Estimation in Non-linear Mixed Effects Models," *Computational Statistics & Data Analysis* 49, no. 4 (2005): 1020–1038, <https://doi.org/10.1016/j.csda.2004.07.002>.
20. T. A. Louis, "Finding the Observed Information Matrix When Using the EM Algorithm," *Journal of the Royal Statistical Society. Series B, Statistical Methodology* 44, no. 2 (1982): 226–233.
21. S. Prothon, M. Aurivillius, U. Tehler, U. G. Eriksson, A. Aggarwal, and Y. Chen, "Safety, Pharmacokinetics and Pharmacodynamics of the Selective Glucocorticoid Receptor Modulator Velsecorat (AZD7594) Following Inhalation in Healthy Volunteers," *Drug Design, Development and Therapy* 16 (2022): 485–497, <https://doi.org/10.2147/DDDT.S334960>.
22. R. M. Savic and M. O. Karlsson, "Importance of Shrinkage in Empirical Bayes Estimates for Diagnostics: Problems and Solutions," *AAPS Journal* 11, no. 3 (2009): 558–569, <https://doi.org/10.1208/s12244-8-009-9133-0>.
23. AstraZeneca, A Phase 2b Randomised, Double Blind, Placebo-Controlled, Parallel Arm, Multi-Centre Study to Assess Efficacy and Safety of Multiple Dose Levels of AZD7594 DPI Given Once Daily for Twelve Weeks, Compared to Placebo, in Asthmatics Symptomatic on Low Dose ICS. clinicaltrials.gov, 2020 accessed at, March 23, 2026, <https://clinicaltrials.gov/study/NCT03622112>.
24. A. Król, R. Palmér, V. Rondeau, S. Rennard, U. G. Eriksson, and A. Jauhiainen, "Improving the Evaluation of COPD Exacerbation Treatment Effects by Accounting for Early Treatment Discontinuations: A Post-Hoc Analysis of Randomized Clinical Trials," *Respiratory Research* 21, no. 1 (2020): 158, <https://doi.org/10.1186/s12931-020-01419-8>.
25. A. Król, R. Palmér, J. Leander, C. Proust-Lima, and A. Jauhiainen, "Location-Scale Latent Process Model for Repeated Ordinal Patient-Reported Outcomes," *Statistics in Medicine* 45, no. 6–7 (2026): e70482, <https://doi.org/10.1002/sim.70482>.
26. B. Cokorudy, J. Harrison, and A. H. Y. Chan, "Digital Markers of Asthma Exacerbations: A Systematic Review," *ERJ Open Research* 10, no. 6 (2024): 00014-2024, <https://doi.org/10.1183/23120541.00014-2024>.

Supporting Information

Additional supporting information can be found online in the Supporting Information section. **Figure S1:** Individual empirical Bayes estimates (EBE) plotted against true simulated individual parameters with unit slope lines and per-parameter coefficients of determination (R^2). The points show all pairs of individual parameter values and EBEs

from the first 25 simulated data sets. **Figure S2:** Distribution plots showing η -shrinkage for each model parameter across all 200 simulated datasets. The estimated distributions have variance equal to the estimated population variance $\omega c^2 \omega c^2$ for $c = 1, \dots, 5$, and the empirical distributions have variance equal to the sample variance of the individual empirical Bayes estimates. **Table S1:** Summarized results from estimating parameters using simulated data sets of length $T = 50$. Runs that resulted in failed standard error estimations were excluded from the coverage calculation. **Table S2:** Summarized results from estimating parameters using simulated data sets of length $T = 1000$. **Table S3:** Population parameter estimates from the exploratory model with categorical treatment groups fitted to the clinical trial data. **Table S4:** Population parameter estimates from the mixed dose-response model fitted to the clinical trial data. **Table S5:** Comparison of three models incorporating different dose-response relationships using Akaike's information criteria (AIC) and the coefficient of determination (R^2) of a linear regression between observed and predicted PEF values. The model denoted none includes no treatment effects; the categorical model includes a categorical effect for every parameter and treatment group; and the mixed model includes a constant treatment effect for the parameter d and E_{\max} dose-response relationships for σ^2 and $p_{1,0}$. The AIC is computed as $2p - \ln(\hat{L})$, where p is the number of parameters in the model, and $\ln(\hat{L})$ is the estimated log-likelihood computed from 25,000 posterior samples per individual. Thus, a lower AIC is considered better. **Figure S3:** Boxplots showing the distribution of estimated random effects from the model denoted "mixed", per model parameter, and stratified on treatment group. The horizontal line is drawn at $y = 0$.