



Published in final edited form as:

Ann Appl Stat. 2023 March ; 17(1): 333–356. doi:10.1214/22-aos1630.

BAYESIAN NON-HOMOGENEOUS HIDDEN MARKOV MODEL WITH VARIABLE SELECTION FOR INVESTIGATING DRIVERS OF SEIZURE RISK CYCLING

Emily T. Wang^{*}, Sharon Chiang[†], Zulfi Haneef[‡], Vikram R. Rao[†], Robert Moss[§], Marina Vannucci^{*}

^{*}Rice University

[†]University of California, San Francisco

[‡]Baylor College of Medicine

[§]SeizureTracker LLC

Abstract

A major issue in the clinical management of epilepsy is the unpredictability of seizures. Yet, traditional approaches to seizure forecasting and risk assessment in epilepsy rely heavily on raw seizure frequencies, which are a stochastic measurement of seizure risk. We consider a Bayesian non-homogeneous hidden Markov model for unsupervised clustering of zero-inflated seizure count data. The proposed model allows for a probabilistic estimate of the sequence of seizure risk states at the individual level. It also offers significant improvement over prior approaches by incorporating a variable selection prior for the identification of clinical covariates that drive seizure risk changes and accommodating highly granular data. For inference, we implement an efficient sampler that employs stochastic search and data augmentation techniques. We evaluate model performance on simulated seizure count data. We then demonstrate the clinical utility of the proposed model by analyzing daily seizure count data from 133 patients with Dravet syndrome collected through the *Seizure Tracker*TM system, a patient-reported electronic seizure diary. We report on the dynamics of seizure risk cycling, including validation of several known pharmacologic relationships. We also uncover novel findings characterizing the presence and

emilywang@rice.edu .

Department of Statistics, Rice University, Houston, TX, USA

Department of Neurology, Baylor College of Medicine, Houston, TX, USA

Department of Neurology, University of California, San Francisco, San Francisco, CA, USA

Seizure Tracker, LLC, Springfield, VA, USA

Supplementary Material. The software implementing our model, with setup instructions and codes replicating the simulation studies, can be found as online supplement. To request the decryption password from the authors please complete the form at <https://forms.gle/hc8UHsB3bSdisQX27>. The most updated version of the software can be downloaded from GitHub at <https://github.com/seizurerisk>.

Appendix A of the online supplement contains a detailed description of the MCMC algorithm implemented for our proposed method.

Appendix B contains example MCMC trace plots, auto-correlation plots, and results of Geweke's diagnostic tests to demonstrate convergence of the MCMC chain in the Seizure Tracker study, accompanied by a short discussion.

Appendix C contains a comparison of our proposed method, ZINB-NHMM-BVS, to the approach of Hubin (2019), ASA-EM, as well as to various submodels of ZINB-NHMM-BVS which employ simpler variable selection schemes for the transition probability regression coefficients on the Seizure Tracker dataset.

Appendix D contains supplemental Tables and Figures.

volatility of risk states in Dravet syndrome, which may directly inform counseling to reduce the unpredictability of seizures for patients with this devastating cause of epilepsy.

Keywords and phrases:

Bayesian inference; Count data; Dravet syndrome; Epilepsy; Hidden Markov Models; Markov chain Monte Carlo; Seizure risk; Zero-inflation

1. Introduction.

Epilepsy, a chronic neurological disorder, is characterized by frequent, unpredictable seizures arising from abnormal electrical disturbances in the brain. This disorder affects over 60 million people world-wide, which is approximately equivalent to a 1% disease prevalence (WHO, ILAE and IBE, 2019). The unpredictable nature of seizures not only makes treatment of epilepsy difficult but leads to increased morbidity and mortality and severely reduces patients' quality of life (Arthurs et al., 2010). Currently, clinical decision-making in epilepsy depends heavily on raw seizure counts and decisions about treatment are based primarily on whether the seizure frequency has increased or decreased after an intervention. It is recognized, however, factors other than treatment, such as cycles in underlying epileptiform discharges, may modulate the likelihood of seizures (Baud et al., 2020).

It is increasingly recognized that seizures are stochastic realizations of periods of heightened seizure risk (Chiang et al., 2018; Baud et al., 2020); given any underlying seizure risk level, the overt number of seizures that a patient has can vary due to natural probabilistic variation and the inherent unpredictability of seizures, making clinical prediction of seizures challenging (Goldenholz et al., 2018; Chiang et al., 2018). Consequently, raw seizure counts are only a surrogate measure and recognized not to be an accurate measure of a patient's true seizure risk. This concept was first formalized by Chiang et al. (2018), who introduced the notion of discrete unknown seizure risk "states" and showed that seizure risk can be estimated as a latent quantity based on the seizure counts. The validity of this conceptual approach to seizure risk has been validated against specialized epilepsy clinician experts (Chiang et al., 2020) and confirmed by empiric observations from chronic intracranial electrocorticography data, which have demonstrated that seizures tend to occur at specific phases of underlying fluctuations in interictal epileptiform activity (Baud et al., 2018; Karoly et al., 2018; Rao et al., 2020; Proix et al., 2020; Leguia et al., 2021).

Here, we build upon the literature described above by proposing a Bayesian non-homogeneous hidden Markov model for zero-inflated count data that yields substantially increased flexibility for accommodating highly granular data and that allows for simultaneous selection of high-dimensional covariates to identify drivers of seizure risk cycles. To handle high proportions of zeros as well as overdispersion, we model daily seizure counts using a zero-inflated negative binomial (ZINB) distribution. We incorporate the external clinical covariates into the estimation of both emission and transition probabilities through logistic regression frameworks that allow additional flexibility across different subjects and time points, and employ variable selection priors to simultaneously

identify the significant covariates. The proposed model offers significant advances as it allows for (1) high temporal granularity of the count data, (2) fine-tuned and personalized risk assessment via the modeling of the effect of subject-level clinical factors on both state transitions and the current expected number of seizures, and (3) variable selection to simultaneously identify covariates that drive seizure risk changes. These aspects are essential to developing a generalizable model capable of modeling seizure risk data in diverse datasets.

We optimize inference by implementing a Markov chain Monte Carlo method that uses stochastic search methods and data augmentation techniques for efficient sampling (Savitsky, Vannucci and Sha, 2011; Polson, Scott and Windle, 2013). We evaluate model performance on simulated seizure count data. We then demonstrate the clinical utility of our model by analysing daily clinical seizure counts recorded by patients with a severe genetic cause of epilepsy, Dravet syndrome (DS), a catastrophic developmental epileptic encephalopathy affecting one in 15,700 infants (Wu et al., 2015). Infants with DS often present with the first convulsive seizure within the first year of life. Shortly thereafter, multiple other seizure types emerge; eventually, patients develop refractory convulsive seizures, multiple seizure types, and intellectual disability. Mortality in Dravet syndrome is high and up to 17% of patients with Dravet syndrome die before adulthood, with 15–61% of deaths attributable to sudden unexpected death in epilepsy (SUDEP) (Cooper et al., 2016). The single greatest risk factor for SUDEP is the presence of generalized tonic-clonic seizures (GTCs) (Harden et al., 2017). Patients with more than three GTCs in the preceding year have more than an eight-fold increased risk of death from SUDEP (Walczak et al., 2001). Understanding patterns in GTC cycling as well as which anti-seizure medications (ASMs) and triggers reduce or increase the likelihood of transitioning to periods of lowered seizure risk for GTCs is of interest to guide understanding and prevention of SUDEP (Ayub et al., 2020). However, the stochastic nature of seizures makes this task difficult.

In our analysis, we consider data from 133 patients with Dravet syndrome from one of the world's largest seizure diary databases, Seizure Tracker LLC¹, and a total of 36 potential variables that may drive seizure risk. Our approach identifies the presence of three distinct seizure risk states in Dravet syndrome through which patients cycle, that roughly equate to a states reflecting low, moderate, and high propensity for GTCs. We report on the volatility and time dynamics of these states in order to contribute to knowledge on the natural history of seizure risk cycling in Dravet syndrome. Lastly, we show that our model accurately recovers several known pharmacologic relationships of ASMs in Dravet syndrome. This application improves understanding of seizure risk cycling in Dravet syndrome and illustrates the usefulness of our approach more broadly for individualized investigations of seizure risk cycling in epilepsy.

In Section 2 below we detail the proposed model and the inference strategy. In Section 3 we utilize simulated seizure count data to evaluate performances of our model. In Section 4 we report on the results from the analysis of daily seizure counts data from patients with Dravet syndrome and discuss our findings. We give concluding remarks in Section 5.

¹Seizure Tracker™ - Your comprehensive resource for tracking and sharing seizure information, <https://seizuretracker.com/>.

2. Methods.

A hidden Markov Model (HMM) is a statistical model which assumes that a system is in one of many latent (unobservable) states at any point in time (Rabiner, 1989). The system transitions from one state to another over time, and observations are generated from an emission distribution, conditional on the latent state sequence. Fig 1 depicts our proposed HMM, which we call ZINB-NHMM-BVS, for seizure risk assessment. The latent layer allows a probabilistic estimate of the sequence of seizure risk states at an individual level. The process in which observed seizure counts are emitted from a statistical distribution dependent on the risk state also allows for a degree of natural variance in counts to be accounted for. Our model formulation additionally incorporates clinical covariates and employs variable selection to perform simultaneous inference on variables significantly associated with changes in latent risk.

We now describe our model, choice of priors and MCMC algorithm for posterior inference in detail.

2.1. Bayesian Non-Homogeneous HMM for Zero-Inflated Count Data.

We assume to have data observed on N patients, with each patient $i \in \{1, \dots, N\}$ recording daily seizure counts over T_i days. Let $Y_{it} \in \{0, 1, \dots\}$ be the number of seizures of patient i at time t , for $i = 1, \dots, N$ and $t = 1, \dots, T_i$. The number of days recorded per patient and the initial starting day for each patient do not need to be the same. Let $\mathbf{W}_{it} = (W_{it1}, \dots, W_{itp})$ be the values of p time-varying covariates for patient i at time t , and $\mathbf{Z}_i = (Z_{i1}, \dots, Z_{iq})$ the values of q fixed covariates. Finally, let $\mathbf{X}_i = (\mathbf{W}_i, \mathbf{1}_{q \times 1} \mathbf{Z}_i)$ be the $((q + p) \times T_i)$ stacked matrix of all clinical covariates for patient i .

To accommodate the temporal structure in the seizure count process, we model the seizure counts as a K -state non-homogeneous first-order HMM, where $K \in \mathbb{N}$. Let ξ_{it} denote the latent seizure risk state of patient i at time t . By the Markov property, the probability of patient i transitioning to a specific state $k \in \{1, \dots, K\}$ at time t depends only upon the hidden state at the previous time point, $\xi_{i(t-1)}$. Furthermore, in a non-homogeneous HMM, the transition probabilities vary over time as a function of the covariates. Specifically, we use a multinomial logit regression to represent these transition probabilities as

$$\Pr(\xi_{it} = k \mid \xi_{i,t-1} = k', \mathbf{X}_{i,t-1}, \beta_k) = \frac{\exp(\mathbf{X}_{i,t-1}^T \beta_{k'k})}{1 + \sum_{l=1}^{K-1} \exp(\mathbf{X}_{i,t-1}^T \beta_{k'l})}. \quad (2.1)$$

The use of logistic regression to link covariate effects to transition probabilities was first proposed by Muenz and Rubinstein (1985), for the case of a two-state Markov chain. To maintain model identifiability, we establish a baseline state and set to zero the regression coefficients corresponding to transitions into this state. With this representation, the coefficient $\beta_{j,k'k}$ measures the effect of covariate j on the *relative* likelihood of patients transitioning from state k' to state k , compared to transitioning from state k' to the

predetermined baseline state. Positive (negative) values of $\beta_{i,k'k}$ indicate that covariate j increases (decreases) the chance of patients transitioning from state k' to k over transitioning from state k' to the baseline. Each individual transition from one state to another has its own set of regression coefficients. Thus, for p covariates there are a total of $K \times (K - 1) \times p$ coefficients governing the Markov transitions. An intercept may be modeled by appending a vector of ones to the data matrix X . We notice that using fixed regression coefficients for transitions from any state into a new state k' in Eq (2.1) would lead to a more parsimonious model, see for example Holsclaw et al. (2017). However, in our case study, separate sets of coefficients for each transition are essential as we are primarily interested in identifying covariates associated with escalations or de-escalations of seizure risk. In Section “Variable Selection Priors” below we employ specialized priors to achieve this objective.

Next we describe the emission distribution of our HMM model. Count data is typically modeled using the Poisson, binomial, or negative binomial distributions. The Poisson distribution is constrained in that the expected value must be equal to the variance. However, daily seizure counts usually exhibit extreme zero-inflation and overdispersion in which the variance exceeds the mean (Tharayil et al., 2017). Therefore, we assume that seizure counts at each state k follow a zero-inflated negative binomial (ZINB) mixture distribution which assumes that, conditional upon being in state k , the observed count is either zero with probability p_k , or it comes from a negative binomial random variable, with probability $1 - p_k$. To facilitate computation, we introduce auxiliary variables Z_{it} that allow us to keep track of whether an observation belongs to the negative binomial mixture component or the “zero” component, that is $Z_{it} | p_k \sim \text{Bernoulli}(p_k)$ where $Z_{it} = 1$ if $Y_{it} = 0$ and $Z_{it} = 0$ if Y_{it} is negative binomial. We follow Pillow and Scott (2012) and parametrize the negative binomial distribution in terms of state-specific dispersion parameters r_k and subject- and state-dependent probabilities ψ_{ik} (equivalent to $1 - \text{success probability}$) as

$$p(Y_{it} | \psi_{ik}, r_k, \xi_{it} = k, Z_{it} = 0) = \frac{\Gamma(Y_{it} + r_k)}{\Gamma(r_k)Y_{it}!} (1 - \psi_{ik})^{r_k} \psi_{ik}^{Y_{it}}. \quad (2.2)$$

This parametrization, in particular, allows us to incorporate both subject-level covariates and time-varying effects into the emission distributions by letting the probabilities ψ_{ik} depend on the covariates as

$$\psi_{ik} = \frac{\exp(\mathbf{X}_{it}^T \boldsymbol{\rho}_k)}{1 + \exp(\mathbf{X}_{it}^T \boldsymbol{\rho}_k)}, \quad (2.3)$$

with $\boldsymbol{\rho}_k$ a state-dependent vector of regression coefficients. A separate set of covariates than those used for the transition probability regression can be considered, if desired. Parametrizing the negative binomial distribution in terms of ψ_{ik} instead of the mean yields a

closed-form full conditional posterior for the regression coefficients ρ . Mean parameters can be easily recovered as $\mu_{itk} = \frac{\psi_{itk} r_k}{1 - \psi_{itk}}$. We can now write our ZINB model as

$$[Y_{it} | \xi_{it} = k, \mathbf{r}, \boldsymbol{\psi}_{itk}, \mathbf{p}] \sim p_k 1_{\{Y_{it}=0\}} + (1 - p_k) \text{NB}(r_k, \psi_{itk}), \quad (2.4)$$

with $\text{NB}(r_k, \psi_{itk})$ given by Eq (2.2).

2.2. Variable Selection Priors.

We are interested in identifying clinical covariates which are associated with changes in seizure risk. We achieve this task by employing variable selection priors (George and McCulloch, 1997; Brown, Vannucci and Fearn, 1998). Let us first consider the transition probabilities (2.1). We introduce a $K \times (K - 1) \times p$ latent inclusion tensor γ such that $\gamma_{j,k'k} = 1$ if variable j has a non-zero effect on the likelihood of transitioning from state k' to state k . We then impose a *spike-and-slab* prior on the regression coefficient $\beta_{j,k'k}$, obtained as a mixture of a Gaussian distribution (the slab) and a point mass at zero (the spike) as

$$[\beta_{j,k'k} | \gamma_{j,k'k}, \mu_{\beta}, \sigma_{\beta}^2] \sim \gamma_{j,k'k} N(\mu_{\beta}, \sigma_{\beta}^2) + (1 - \gamma_{j,k'k}) \delta_0(\beta_{j,k'k}). \quad (2.5)$$

We further assume $\gamma_{j,k'k} \sim \text{Bernoulli}(\theta_{j,k'k})$, with a hyperprior on the inclusion probability $\theta_{j,k'k} \sim \text{Beta}(g_{\beta}, h_{\beta})$, and then integrate $\theta_{j,k'k}$ out. Multivariate *spike-and-slab* prior constructions of type (2.5) have been employed in multivariate regression models to identify covariates that are associated with individual responses (Stingo et al., 2010). In our framework, this construction allows us to select covariates that are associated with specific transitions, from one risk state to another one. We note that these inferences are conducted in relation to the baseline state. Selection of covariate j as important for transitions from state k' to state k implies that the value of covariate j has a significant effect on the likelihood of a patient transitioning from state k' to state k , over state k' to the baseline state.

We adopt a similar prior setup for the regression coefficients ρ_k of the state-dependent probabilities (2.3). Let δ_{jk} be a latent indicator of whether variable j has a non-zero effect on the probability parameter of state k , ψ_k . We impose a *spike-and-slab* prior on ρ_{jk} of the type

$$[\rho_{jk} | \delta_{jk}, \mu_{\rho}, \sigma_{\rho}^2] \sim \delta_{jk} N(\mu_{\rho}, \sigma_{\rho}^2) + (1 - \delta_{jk}) \delta_0(\rho_{jk}), \quad (2.6)$$

with $\delta_{jk} \sim \text{Bernoulli}(\theta_{jk})$ and a hyperprior $\theta_{jk} \sim \text{Beta}(g_{\rho}, h_{\rho})$. The two variable selection schemes outlined above offer separate interpretations. Variable selection for the emission distributions via ρ determines which covariates are associated with higher or lower seizure frequency, conditional on the underlying latent state. In contrast, variable selection for the transitions identifies covariates which are associated with worsening or improvement of seizure risk, from time t to $t + 1$.

Lastly, we specify a Beta prior for the zero-inflation parameters, $p_k \sim \text{Beta}(c, d)$ for $k = 1, \dots, K$, a Gamma prior on the dispersion parameters, $r_k \sim \text{Gamma}(e, f)$, and assume the distribution of the initial hidden state at time $t = 1$ to be $\xi_{i1} \sim \text{Multinomial}(1; \pi_1, \dots, \pi_K)$, with a conjugate Dirichlet prior on π .

2.3. Data Augmentation.

Given the model construction, the full likelihood is

$$\begin{aligned} \mathfrak{L}(\mathbf{Y} \mid \xi, \mathbf{r}, \boldsymbol{\psi}, \mathbf{p}, \mathbf{X}, \boldsymbol{\beta}) &\propto \prod_{i=1}^N p(\xi_{i1}) \times \prod_{i=1}^N \prod_{t=1}^{T_i} \frac{\exp(\mathbf{X}_{i,t-1}^T \boldsymbol{\beta}_{\xi_{i,t-1}, \xi_{it}})}{1 + \sum_{l=1}^{K-1} \exp(\mathbf{X}_{i,t-1}^T \boldsymbol{\beta}_{\xi_{i,t-1}, l})} \\ &\times \prod_{i=1}^N \prod_{t=1}^{T_i} \left[(1 - p_{\xi_{it}}) + p_{\xi_{it}} (1 - \psi_{i \xi_{it}})^{r_{\xi_{it}}} \right]^{1 - Z_{it}} \times \left[p_{\xi_{it}} \frac{\Gamma(Y_{it} + r_{\xi_{it}})}{\Gamma(r_{\xi_{it}}) Y_{it}!} (1 - \psi_{i \xi_{it}})^{r_{\xi_{it}}} \psi_{i \xi_{it}}^{Y_{it}} \right]^{Z_{it}}. \end{aligned}$$

In our sampling scheme for posterior inference, we make use of modern data augmentation techniques that allow us to optimize the inference.

2.3.1. Inference on the transition matrix elements.—In our model, the probability of a patient transitioning from risk state k on day t to state k' on the next day, $t + 1$, is dependent on the values of the patient’s clinical covariates on the original day t via the multinomial logistic regression of Eq (2.1). Sampling schemes for the multinomial logistic regression coefficients are challenging to implement, due to the intractable form of the likelihood and the lack of a conjugate prior for the coefficients. An efficient data augmentation approach for multinomial logistic regression models via Pólya-Gamma distributions was proposed by Polson, Scott and Windle (2013). This method essentially manipulates terms in the full likelihood such that they combine into a single Gaussian kernel in the posterior. The key result which makes Pólya-Gamma random variables useful in the logistic setting is that for $b > 0$,

$$\frac{(e^\psi)^a}{(1 + e^\psi)^b} = 2^{-b} e^{\kappa \psi} \int_0^\infty e^{-\omega \psi^2 / 2} p(\omega) d\omega, \tag{2.7}$$

where $\kappa = a - b/2$ and $\omega \sim \text{PG}(b, 0)$.

Let $Y_{itk} = I_{\xi_{it}}(k)$ be a binary representation of the latent states, where $I_{\xi_{it}}(k)$ is an indicator function equal to 1 if $\xi_{it} = k$ and 0 otherwise. Then using the multinomial logistic representation of transition probabilities in Eq (2.1) and following Holmes and Held (2006), the posterior can be rewritten as

$$p(\boldsymbol{\beta} \mid \mathbf{X}, \mathbf{Y}) = p(\boldsymbol{\beta}) \cdot \prod_{i=1}^N \prod_{t=2}^{T_i} \prod_{k=1}^K \left[\frac{\exp(\zeta_{itk} - C_{itk})}{1 + \exp(\zeta_{itk} - C_{itk})} \right]^{Y_{itk}},$$

where $\zeta_{itk} = \mathbf{X}_{i,t-1}^T \boldsymbol{\beta}_{k'k}$ and $C_{itk} = \log \sum_{j \neq k} \exp\{\zeta_{ij}\}$. Making use of the property of the Pólya-Gamma distribution in Eq (2.7), the full conditional posterior of $\boldsymbol{\beta}_{k'k}$, given all other parameters of the model, including $\boldsymbol{\beta}_{k'(-k)}$, is

$$\begin{aligned} p(\boldsymbol{\beta}_{k'k} | \boldsymbol{\beta}_{k'(-k)}, \mathbf{X}, \mathbf{Y}) &\propto p(\boldsymbol{\beta}_{k'k} | \boldsymbol{\beta}_{k'(-k)}) \prod_{i=1}^N \prod_{t=2}^{T_i} \left[\frac{e^{\zeta_{itk}} - C_{itk}}{1 + e^{\zeta_{itk}} - C_{itk}} \right]^{Y_{itk}} \left[\frac{1}{1 + e^{\zeta_{itk}} - C_{itk}} \right]^{1 - Y_{itk}} \\ &= p(\boldsymbol{\beta}_{k'k} | \boldsymbol{\beta}_{k'(-k)}) \prod_{i=1}^N \prod_{t=2}^{T_i} \left[e^{\kappa_{itk}(\zeta_{itk} - C_{itk})} \int_0^\infty e^{-\omega_{itk}(\zeta_{itk} - C_{itk})^2/2} p(\omega_{itk}) d\omega_{itk} \right], \end{aligned}$$

where $\kappa_{itk} = Y_{itk} - 1/2$ and $\omega_{itk} \sim \text{PG}(1,0)$, for $k = 1, \dots, K$. Using this method, we can assume conditionally conjugate priors for the regression coefficients, $\boldsymbol{\beta}_{k'k} \sim N(m_0, V_0)$. Then, conditioning on the Pólya-Gamma random variables ω_{itk} and with some straightforward algebra, the posterior collapses into a single Gaussian kernel, leading to a two-step sampling scheme which involves a Pólya-Gamma update for the latent variables ω_{itk} , followed by a joint Gaussian update for the regression coefficients $\boldsymbol{\beta}_{k'k}$. For additional details on this sampling step, see Appendix A of the Supplementary Material.

2.3.2. Inference on the negative binomial probability parameters.—Conditional on a latent state $\xi_{it} = k$, daily seizure counts are observed from a zero-inflated negative binomial (ZINB) emission distribution with dispersion parameter r_k , probability ψ_{itk} and zero-inflation parameter p_k . In this section, we consider inference on the probability parameter ψ_{itk} , which is dependent on the covariates \mathbf{X}_{it} through the logistic regression relationship of Eq (2.3). We focus on inference for the regression coefficients $\boldsymbol{\rho}_k$ for the probability parameter and consider only the negative binomial mixture component of the ZINB distribution.

Traditionally, Bayesian negative binomial regression is challenging due to the intractable form of the posterior and lack of a conjugate prior for the regression coefficients. Pillow and Scott (2012) describe an approach which circumvents these issues by taking advantage of Pólya-Gamma auxiliary variables to derive a two-step closed-form Gibbs update for the regression coefficients. Due to the similar logistic link connecting the covariates to the parameter of interest, the sampling procedure for the negative binomial regression coefficients is similar to that of the transition probability regression coefficients.

Conditional on $Z_{it} = 0$ and latent state $\xi_{it} = k$, the distribution of seizure counts for patient i at time t is negative binomial with

$$p(Y_{it} | X_{it}, \mathbf{r}, \boldsymbol{\rho}, \xi_{it} = k, Z_{it} = 0) \propto (1 - \psi_{itk})^{r_k} \psi_{itk}^{Y_{it}}.$$

Then, leveraging the Pólya-Gamma property of Eq (2.7), we obtain the full conditional

$$\begin{aligned}
 p(\rho_k | Y_{it}, X_{it}, r, \xi_{it} = k, Z_{it} = 0) &= p(\rho_k) \prod_{i=1}^N \prod_{t=1}^{T_i} \frac{\{ \exp(X_{it}^T \rho_k) \}^{Y_{it}}}{\{ 1 + \exp(X_{it}^T \rho_k) \}^{r_k + Y_{it}}} \\
 &\propto p(\rho_k) \prod_{i=1}^N \prod_{t=1}^{T_i} e^{\kappa_{2it} \eta_{itk}} \int_0^\infty e^{-\omega_{2itk} \eta_{itk}^2} 2p(\omega_{2itk}) d\omega_{2itk},
 \end{aligned}$$

where $\kappa_{2itk} = \frac{(Y_{it} - r_k)}{2\omega_{2itk}}$, $\eta_{itk} = X_{it}^T \rho_k$ for brevity, and $\omega_{2itk} \sim \text{PG}(Y_{it} + r_k, 0)$. The final expression combines to form a Gaussian kernel which allows us to derive a two-step Gibbs sampler for the negative binomial regression coefficients. Details of this sampler are provided in Appendix A of the Supplementary Material.

2.3.3. Inference on the dispersion parameters.—For the update of the negative binomial dispersion parameters r_k , we adopt the data augmentation approach of Zhou and Carin (2015), which exploits the relationship between the negative binomial distribution and the compound Poisson distribution. Conditional on $Z_{it} = 0$, the negative binomial counts can be rewritten as $Y_{it} = \sum_{t=1}^{L_{itk}} u_{itkt}$, with $u_{itkt} \stackrel{iid}{\sim} \text{Logarithmic}(\psi_{itk})$ and $L_{itk} \sim \text{Poisson}(-r_k \ln(1 - \psi_{itk}))$, for $i = 1, \dots, N$ and $t = 1, \dots, T_i$ (Quenouille, 1949). Assuming the prior $r_k \sim \text{Gamma}(e, f)$, a two-step update for the dispersion parameter r_k at each state $k = 1, \dots, K$ iteratively draws

$$\begin{aligned}
 [L_{itk} | r_k] &\sim \text{CRT}(Y_{it}, I_{itk}(k), r_k) \\
 [r_k | L_{itk}, \psi_{itk}] &\sim \text{Gamma}\left(e + \sum_{i=1}^N \sum_{t=1}^{T_i} L_{itk}, f - \sum_{i,t} \log(1 - \psi_{itk})\right),
 \end{aligned}$$

where CRT is the Chinese restaurant table distribution.

2.4. Posterior Inference.

For posterior inference, we design a MCMC algorithm that iteratively samples from the joint posterior distribution of all parameters, $\Theta = \{\beta_k, \gamma, \pi, \xi_{it}, \rho_k, \delta, r_k, p_k, Z_{it}\}$. Our strategy combines stochastic search methods for variable selection and the data augmentation techniques previously discussed. Our complete sampler is a Metropolis-within-Gibbs and is briefly summarized in Algorithm 1.

Full details of the updates are given in Appendix A of the Supplementary Material. In our applications, starting values for all model parameters were initialized randomly. We discarded a fixed initial number of posterior samples, called the burn-in period, to control for the effect of initialization. Posterior estimates of latent states for each patient were obtained via the posterior mode. The other model parameters were estimated via posterior means and 95% credible intervals. Covariate effects were considered statistically significant if their marginal posterior probability of inclusion (MPPI) was greater than 0.5 (Barbieri and Berger, 2004).

Algorithm 1: Sampler for ZINB-NHMM-BVS

```

Input : Iterations  $S$ , states  $K$ , data  $X, Y$ 
Output: Posterior samples of model parameters  $\Theta$ 
1 for  $s = 1$  to  $S$  do
2   for  $k = 1$  to  $K$  do
3     Joint update of  $(\beta, \gamma)$  via the stochastic search MCMC of Savitsky,
       Vannucci and Sha (2011), combined with the Pólya-Gamma data
       augmentation scheme of Polson, Scott and Windle (2013). This sampler
       performs an Add/Delete step on one component of  $(\beta, \gamma)$  or a Swap step
       of two components, with equal probability.
4     Joint update of  $(\rho, \delta)$ , similarly to the update of  $(\beta, \gamma)$ .
5     Update  $r$  via the data augmentation method of Zhou and Carin (2015).
6     Update  $p$  from the full conditionals.
7     Update  $\pi$  by sampling from the full conditionals.
8   end
9   Update  $\xi$  via Scaled Forward-Backward algorithm (Scott, 2002).
10  for  $i = 1$  to  $N$  do
11    for  $t = 1$  to  $T_i$  do
12      Update  $Z$  from the full conditionals.
13    end
14  end
15 end

```

Finally, one common issue with MCMC methods for state-space models is label switching, which arises due to the invariance of the likelihood to permutations of the labels of the latent states (Scott, 2002). To address this, at each iteration of the MCMC algorithm we calculate the negative binomial means as $\mu_{ik} = \frac{\psi_{ik} r_k}{1 - \psi_{ik}}$ for $k = 1, \dots, K$, and enforce an ordering among the averaged state means as $\sum_{i=1}^N \sum_{t=1}^{T_i} \mu_{i1} < \sum_{i,t} \mu_{i2} < \dots < \sum_{i,t} \mu_{iK}$.

3. Simulation Study.

We performed an extensive simulation study on the performance, accuracy, and sensitivity of our proposed model, ZINB-NHMM-BVS, on simulated seizure count data. We also compare the performance of our method for latent state classification and variable selection to the adaptive simulated annealing expectation maximization (ASA-EM) algorithm of Hubin (2019), which performs simultaneous variable selection and inference on hidden states.

3.1. Data generation.

We generated data by simulating $N = 100$ patients such that each patient had between 100 and 110 time points of data. If patients record daily seizure frequencies, this corresponds to about three months of data, which is a typical time between outpatient epilepsy visits. We set the number of hidden states to $K = 3$ and set the corresponding dispersion parameters to $r = (3.8, 15)'$, the zero-inflation parameters to $p = (0.7, 0.05, 0.01)'$ and the initial state distribution to $\pi = (0.9, 0.08, 0.02)'$. We allowed covariates to affect the negative binomial parameters via construction (2.3). For each patient, $p = 7$ covariates were designed to include a mix of both discrete and continuous features. Negative binomial regression coefficients were set as $\rho_1 = (-0.7, -0.8, -0.8, 0, -0.8, -0.7, -0.7)$, $\rho_2 = (-0.4, 0, 0, -0.4, 0, -0.7, -0.6)$ and $\rho_3 = (0, -0.5, 0, -0.5, 0.5, 0.4, 0)$. This approximately corresponds to an average of 0.2, 2.8, and 12.6 seizures per time unit at states 1, 2, and 3, respectively. We also allow all covariates to affect transitions from state 1 to 1, covariates $\{X_1, X_2, X_3\}$ to affect transitions from state 1 to 2, covariates $\{X_2, X_3, X_7\}$ to affect transitions from state 2 to 1, covariates $\{X_3, X_7\}$ to affect transitions from state 2 to 2, and covariates $\{X_4, X_7\}$ to affect both transitions from state 3 to 1 and from state 3 to 2. Non-zero transition regression coefficients were set as

$\{\beta_{1,11}, \beta_{2,11}, \beta_{3,11}, \beta_{4,11}, \beta_{5,11}, \beta_{6,11}, \beta_{7,11}\} = 3.5$, $\{\beta_{1,12}, \beta_{2,12}, \beta_{3,12}\} = 2.9$, $\{\beta_{2,21}, \beta_{3,21}, \beta_{7,21}\} = 2.4$, $\{\beta_{3,22}, \beta_{7,22}\}$ and $= 3.0$, $\{\beta_{4,31}, \beta_{7,31}\} = -2.9$
 $\{\beta_{4,32}, \beta_{7,32}\} = -2.5$. Finally, we simulated the hidden states ξ and then generated the seizure counts Y from $[Y_{it} | \xi_{it} = k, \mathbf{r}, \mathbf{p}] \sim \text{ZINB}(r_k, \psi_{itk}, p_k)$.

3.2. Parameter settings.

When running MCMC chains, we set the following hyperparameters: We set non-informative Beta(1, 1) priors on the zero-inflation parameters, vague Gamma(.01, .01) priors on the NB dispersion parameters and a non-informative Dir(1, ..., 1) prior on the initial hidden state probabilities. We imposed Beta(1, 5) priors on the variable inclusion indicators for the transition probability regression and negative binomial regression coefficients, to allow for some sparsity in the selection model. We set mildly informative priors on the transition probability regression coefficients, $\beta_{k'k} \sim N(0_p, \text{diag}(1))$, and the negative binomial regression coefficients, $\rho_k \sim N(0_p, \text{diag}(1))$, to prevent the logistic terms (2.1) and (2.3) from taking on too extreme values (near 0 or 1). MCMC chains were run for 20,000 iterations, discarding the first 10,000 as burn-in, leaving 10,000 for inference.

3.3. Results.

We evaluated classification accuracy of the latent risk states using the following multi-class metrics:

$$\text{Accuracy} = \left(\sum_{k=1}^K \frac{\text{TP}_k + \text{TN}_k}{\text{TP}_k + \text{FP}_k + \text{TN}_k + \text{FN}_k} \right) / K$$

$$\text{Precision} = \left(\sum_{k=1}^K \frac{\text{TP}_k}{\text{TP}_k + \text{FP}_k} \right) / K$$

$$\text{Sensitivity} = \left(\sum_{k=1}^K \frac{\text{TP}_k}{\text{TP}_k + \text{FN}_k} \right) / K$$

$$\text{Specificity} = \frac{\left(\sum_{k=1}^K \frac{\text{TN}_k}{\text{TN}_k + \text{FP}_k} \right)}{K}$$

$$F_1 \text{ score} = 2 \cdot \frac{\text{Precision} \cdot \text{Sensitivity}}{\text{Precision} + \text{Sensitivity}}$$

that weigh all latent states equally, and where TP_k , FP_k , TN_k , and FN_k denote true positive, false positive, true negative, and false negative counts for state k . Using these macro-level

metrics allows us to evaluate model performance in the presence of possible class (i.e. latent state) imbalance.

Results on the inference of the latent states, averaged across 20 replicate datasets, are reported in the first row of Table 1, as the “default” model. Our method achieves good recovery of latent states with less than 5% false positive and false negative rates. It also estimates other model parameters, such as dispersion, zero-inflation, and initial hidden states, with high accuracy and low mean-squared errors, see Table S3. Furthermore, it exhibits excellent performance on the selection of the covariates and the estimation of the corresponding regression coefficients. Table 2 (“default” row) reports covariate selection results in terms of classification metrics, calculated based on a 0.5 threshold on the PPIs and averaged across 20 replicate datasets. Averaged estimates and root mean square errors (RMSE) on the corresponding regression coefficients are given in Tables S4 and S5.

Additionally, we briefly investigated scalability to larger p by augmenting the data matrix with 43 additional noise variables, for a total of $p = 50$ covariates. Convergence was reached at 30,000 MCMC iterations and the resulting chain took approximately 8.7 hours to run. The model achieved high (>0.90) F_1 scores for classification of latent states as well as covariate selection, indicating that the method can scale up somehow to larger values of p .

3.4. Sensitivity analysis.

Due to the hierarchical structure of our HMM, there are several hyperparameters that need to be set by the user before running the MCMC. In Section “Parameter settings”, we provide guidelines for specification of non-informative or vague priors on the zero-inflation parameters, the dispersion parameters and the initial hidden state probabilities. We performed sensitivity analysis of the beta priors on the variable inclusion indicators and the Gaussian slab components of the *spike-and-slab* priors (2.5) and (2.6) on the regression coefficients, β_{k_k} and ρ_k . Results in terms of the classification metrics are reported in Tables 1 and 2, for latent state estimation and covariate selection, respectively. As expected, we found that varying the prior probability of variable inclusion affects the number of covariates selected into the model, in both the transition and emission components. Based on F_1 score, a good balance between FPR and FNR can be achieved by setting the prior probability of variable inclusion to be low, centered around 5–30%. Overall, classification of latent risk states as assessed by the macro-level metrics was robust across the different prior specifications (Table 1).

Furthermore, we found that varying the slab variances Σ_ρ of the negative binomial regression parameters ρ had minimal effect on both the covariates selected as well as the classification of the latent risk states. This was not the case for the slab variances Σ_β of the transition probability regression parameters β_k . Setting the diagonal elements of the covariance matrix Σ_β to extreme values increases the false positive rate for variable selection. However, classification of latent states remained robust even when the false positive rate for variable selection was high.

In the above simulations, variable selection performance was assessed based on the median probability model, i.e. covariates were selected if their MPPI > 0.5 . Furthermore, we investigated alternative ways to variable selection. For example, the most probable model chooses the model with the largest probability mass. In our implementation, each individual transition and state is treated separately and the most probable model is chosen based on the model that is visited the most often. In contrast, model selection based on the information criteria AIC or BIC chooses the model based on the MCMC iteration that minimizes the information criterion. We provide a comparison of these various approaches in Table 3. Based on these simulations, the median probability model has very similar performance compared to the most probable model, while the models chosen by AIC or BIC lag behind both the median and most probable models.

Finally, we investigated the impact of various effect sizes on model performance. Specifically, we multiplied the effect sizes specified in Section 3.1 by a scale factor $F \in \{0.2, 0.4, 0.6, 0.8, 1\}$ and reassessed variable selection and latent state classification performances. Results are presented in Table S6 and indicate a measure of robustness even with the smallest effect sizes.

3.5. Comparison to ASA-EM.

We compare the performance of our method to the non-homogeneous HMM modeling approach of Hubin (2019) that employs the adaptive simulated annealing expectation maximization (ASA-EM) algorithm for estimation. Similar to our approach, the ASA-EM approach allows covariates to affect both emissions and transitions. However, ASA-EM operates under the restriction that regardless of latent state, firstly, the same covariates affect the transitions and, secondly, the same covariates affect the emission distributions. The expectation maximization (EM) algorithm is used to infer unknown parameters, and the simulated annealing procedure is used for model exploration. The adaptive component of the method comes from the automatic learning of the model's tuning parameters over time. For additional details see Hubin (2019).

The ASA-EM approach requires emission distributions to come from exponential families. Therefore, we re-simulated the count data using a Poisson emission distribution generated through a simple generalized linear model as $Y_{it} | \xi_{it} = k, \mu_{itk} \sim \text{Poisson}(\mu_{itk})$, with $\log(\mu_{itk}) = \mathbf{X}_{it}^T \boldsymbol{\zeta}_k$. Again, we simulated data for $N = 100$ patients such that each patient had between 100 and 110 time points of data. We set the number of hidden states to $K = 2$ and the initial state distribution to $\boldsymbol{\pi} = (0.9, 0.1)'$. For each patient, $p = 15$ covariates were simulated and included a mix of both discrete and continuous features. The Poisson regression coefficients were fixed for each state at $\boldsymbol{\zeta}_1 = (-0.7, -0.7, 0, 0, -4, 0, -0.7, 0, 0, 0, 0, 0, 0, 0, 0)$ and $\boldsymbol{\zeta}_2 = (0.5, -0.4, 0, 0, 0.7, 0, 0.5, 0, 0, 0, 0, 0, 0, 0, 0)$. Thus, regardless of state, covariates $\{X_1, X_2, X_5, X_7\}$ affected the emission distributions. This construction approximately corresponds to an average of 0.1 and 3.3 daily seizures at states 1 and 2, respectively. Finally, regardless of state we allowed covariates $\{X_2, X_3, X_5, X_7\}$ to affect the transitions.

Results reported here were obtained using the same hyperparameter settings described in the previous simulations. For the proposed ZINB-NHMM-BVS model, we considered three

different choices for the variable inclusion hyperparameters in conjunction with the median probability model, to explore the effect of the beta prior. We also considered two alternative approaches to model selection, AIC and BIC. ASA-EM was run on 24 parallel threads for 3 epochs each and with parameters suggested in Hubin (2019). AIC and BIC were used as the model selection criteria for ASA-EM, with a separate set of results reported for each. The choice of model selection criterion had a non-trivial impact on performance, and better performance gains may be realized if an information criterion more suitable for complex hierarchical models is used instead. Final results were averaged across 30 simulated datasets.

Variable selection performance metrics, including number of covariates selected, selection false negative rates, false positive rates and other metrics are reported in Table 4. Latent state classification performance metrics are reported in Table 5. Overall, our proposed model achieves superior performance over ASA-EM in selecting covariates affecting the emission distributions, in selecting covariates that influence transitions, and in terms of overall classification of the latent states. The difference in performances could be attributed to several design choices which differ from ASA-EM. For instance, ZINB-NHMM-BVS implements a stochastic search algorithm based on an add-delete-swap sampler and modern data augmentation tools such as Pólya-Gamma augmentation, which proved to be very effective in recovering regression coefficients affecting both the Markov transitions as well as emissions.

4. Dynamics of risk cycling in Dravet syndrome.

4.1. Study subjects.

The Seizure TrackerTM system contains patient-reported data on seizure timestamps, seizure characteristics, triggers, and medications recorded by more than 30,000 people with epilepsy across the world. Deidentified and unlinked data on all patients with self-identified Dravet syndrome in the SeizureTracker database between 2007 – 2020 were exported from [SeizureTracker.com](https://www.seizuretracker.com) on February 26, 2020. At the time of export, a total of 133 people with Dravet syndrome, with ages ranging between 2 months and 47 years, recorded 34, 431 GTCs between December 1, 2007 and February 26, 2020, spanning 141, 499 patient-days of data with an average of 1, 064 days and 259 GTCs recorded per patient. Demographic information is reported in Table 6.

4.2. Data pre-processing.

A total of 618 patients with Dravet syndrome were initially identified from 2007 – 2020 from the Seizure Tracker database. Patients were excluded if their [SeizureTracker.com](https://www.seizuretracker.com) account was used for less than 30 days, if fewer than 20 GTCs were recorded, or if patients averaged less than one seizure per year. This allowed for exclusion of patients who did not consistently record seizures and to increase the precision of estimates. Patients with missing or invalid dates of birth or sex were also excluded. Invalid seizure durations were imputed as the mean duration of that seizure type for the patient or, if unavailable, the mean duration of that seizure type for all patients. Continuous covariates were normalized to the 0–1 range.

Individual medication records were excluded if the medication name or recorded start date were missing or invalid. For medications with missing end dates, the end date was imputed as three weeks (21 days) after the medication start date. A three-week imputation was considered reasonable as ASMs are often uptitrated over 2–4 weeks to attain a therapeutic dosage (Panayiotopoulos, 2005). Rescue medications were not included in analysis. Generic and brand name medications were analyzed in the same class (Privitera et al., 2016; Vossler, Anderson and Bainbridge, 2016).

The influence of $p = 36$ clinical covariates on changes in seizure risk were evaluated, including 20 classes of daily dosed ASMs, three adjuvant medications, nine common seizure triggers, vagus nerve stimulation (VNS) therapy usage, and three clinical characteristics (age, sex, and menstrual cycling). A complete list of medication classes, seizure triggers, and clinical characteristics is reported in Table S7.

4.3. Seizure risk states in Dravet syndrome.

We applied our model to GTCs in Dravet syndrome in order to evaluate the characteristics and dynamics of seizure risk cycles in patients with this condition, while simultaneously identifying ASMs and triggers associated with changes in GTC risk. In addition to the $p = 36$ clinical covariates from Seizure Tracker, we also estimated the effect of the intercept in both the transition and emission components of the model. Results were obtained by running MCMC chains for 20,000 iterations with 10,000 sweeps as burn-in. Hyperparameters were set to be weakly informative as described in the simulation studies. Convergence of the MCMC chains was assessed via visual inspection of trace plots and auto-correlation plots, and by the Geweke's diagnostic test (Geweke, 1992). Convergence results are shown for a few selected parameters that are sampled via Pólya-Gamma data augmentation in Appendix B of the Supplementary Material. The optimal number of states, K , was chosen based on minimization of the deviance information criterion (DIC, defined as in Gelman et al., 2013) over a grid of possible values $K \in \{2, 3, 4, 5, 6, 7\}$. DIC values were obtained using the BEST approach of selecting covariates based on the median probability model (as opposed to the FULL model which includes all covariates or the NULL model which includes no covariates). More specifically, let $\hat{\theta}_{\text{Bayes}} = E(\theta | y)$ be the posterior mean and θ^s be posterior draws at MCMC iteration $s = 1, \dots, S$. Then, the DIC is obtained as

$$\text{DIC} = -2\log p(y | \hat{\theta}_{\text{Bayes}}) + 2p_{\text{DIC}},$$

where

$$p_{\text{DIC}} = 2 \left(\log p(y | \hat{\theta}_{\text{Bayes}}) - \frac{1}{S} \sum_{s=1}^S \log p(y | \theta^s) \right).$$

We found evidence for $K = 3$ distinct seizure risk states in Dravet syndrome based on minimization of the DIC. These three states can be interpreted with respect to their pro-ictal tendency as states in which the patient is at low, moderate, or high risk for GTCs. The empirical distribution of seizures within each of the identified states provides information

on the distribution of seizures that a patient might expect on days in each state (Fig 2). State 1 (“low” risk) was associated with the lowest seizure frequency. Patients in this state could expect to be seizure-free from GTCs on 93% of days in this state, with a mean GTC seizure frequency of 1.31 (SD, 0.82) on days when at least one GTC occurred. Patients with Dravet syndrome spent the majority of time (63.8% of all days) in this low risk state. In state 2 (“moderate” risk), patients also had reasonable expectation for freedom from GTCs, and could expect to be seizure-free from GTCs on 75% of days in this state, with a mean GTC seizure frequency of 1.42 (SD, 0.88) on days when at least one GTC occurred. On average, patients were in this moderate risk state about 34.3% of all days. In state 3 (“high” risk), patients had a dramatically higher risk for GTCs. In this state, patients could expect to have at least one GTC with more than 85% likelihood, and at least two GTCs per day on two-thirds of days spent this state. These estimates may be informative for counseling patients on seizure-related injury and death. For example, given the association between GTCs and SUDEP risk, patients in risk state 3 would be predicted to be at extremely high risk of SUDEP on days in this state. Patients spent the least amount of time (1.9% of all days) in this high risk state.

4.4. Mean sojourn time and volatility of risk states in Dravet syndrome.

Our inference yields information on the mean sojourn times (MST) for each seizure risk state, which provide a measure of the expected duration of time that a patient with Dravet syndrome is likely to spend in each seizure risk state before transitioning out of the state. This quantity is empirically estimated based on the posterior mode of the latent risk states (see Figure S3). The MST for the low risk state was 117.6 days (about 3.9 months) with an interquartile range (IQR) of 12 to 123 days, indicating large variance in the sojourn time for the low risk state. We found that the sojourn time of the low risk state was <12 months in 91.4% of patients. The MST for the moderate risk state was 59.7 days with an IQR of 3.0 to 53.5 days. In contrast, the MST for the high risk state was a shorter 5.8 days (IQR, 2 to 6 days).

We also investigated the volatility of the dynamics of transitions between seizure risk states, based on the distribution of the estimated transition probabilities (2.1), which provide estimates of the likelihood a patient will be in a certain state on the next day. We found that seizure risk states in Dravet syndrome have a strong tendency toward self-transitioning (i.e., from State X to State X) on a daily timescale. In other words, once a patient enters a given risk state, s/he is more likely to remain in that state the following day, than to transition to a different state. The self-transitioning property was strongest for state 1, less for state 2, and least for state 3; for example, see the average transition matrix in Fig 3. Figure S4 also shows the distribution of transition probabilities for each individual transition.

4.5. Incorporation of external clinical covariates improves accuracy of seizure risk estimation.

Next, we show that accounting for external variables on modulatory factors improves accuracy of the estimates for seizure risk states. Fig 4a shows the expected number of GTCs based on our proposed model (ZINB-NHMM-BVS) for the sequence of daily seizure counts from a randomly selected patient with Dravet syndrome. The expected number of

GTCs based on a *homogeneous* HMM (ZINB-HHMM), in which patient-level covariates are not incorporated into estimation of the mean seizure process μ_{it} , is shown for comparison in Fig 4b. We found that the use of a non-homogeneous HMM corresponds to a 8.6% improvement in mean absolute error, defined as the average absolute difference between the observed response Y_{it} and the mean value estimated by the model $\hat{\mu}_{it}$, over a homogeneous HMM that does not incorporate covariates into the seizure risk estimate. In terms of deviance information criterion, the DIC of the proposed method (ZINB-NHMM-BVS) was 119,790, compared to the DIC of the homogeneous model (ZINB-HHMM) which was a higher 138,854. These metrics suggest that accounting for clinical variables into statistical algorithms for estimating seizure cycles may improve the accuracy of seizure risk estimation, particularly compared to methods which do not consider external variables into their estimates of seizure risk, and emphasizes the utility of incorporating multimodal data streams into seizure risk algorithms.

We also illustrate the methodological advances of the current model compared to our original model in Chiang et al. (2018), which we refer to as ZIP-NHMM, and demonstrate different situations in which the two models may be used. First, as shown in Fig 4c, using ZIP-NHMM, the state-specific estimates of the expected number of seizures at each risk state (μ_k) are fixed and do not vary based on patient-level covariates. In contrast, ZINB-NHMM-BVS produces subject- and state-specific estimates of the expected number of seizures in each risk state (μ_{itk} , Eq 2.3). This will provide improved estimation of the expected number of seizures at each time point in situations where the distribution of seizure counts occurring during a specific risk state is more variable. The difference between the two models can be seen by contrasting Fig 4a and c. During the first 300 days, the estimated mean process is better for ZIP-NHMM, as the patient has a less variable range of seizure counts occurring in these risk states (here, 0 or 1 seizures per day) (Fig 4c). In the following 200 days, the patient has a more variable range of seizure counts that may occur, with better estimation of the mean process by ZINB-NHMM-BVS (Fig 4a). Second, the current model improves on the ability to distinguish between distinct risk states in datasets when the number of events is small. As shown in Fig 4a and c, ZINB-NHMM-BVS distinguishes between low, moderate, and high risk states, whereas ZIP-NHMM groups the moderate and high risk states into a single high risk state. This becomes relevant in data with increased temporal granularity, such as intracranial electrocorticography data or real-time wearable data (Chiang et al., 2021), for which counts in smaller time bins are often close to zero.

We conduct additional experiments to compare the performance of ZINB-NHMM-BVS to the method of Hubin (2019), ASA-EM, as well as to various submodels of ZINB-NHMM-BVS that utilize simpler variable selection schemes for the transition probability regression coefficients, on the Dravet syndrome Seizure Tracker dataset. These results are presented in Appendix C of the Supplementary Material.

4.6. Drivers of risk cycles in Dravet syndrome.

Next, we demonstrate how the methodological addition of variable selection to our model results in a method that can be used to study drivers of seizure risk cycles. By allowing for covariate selection in our model, in particular, we can distinguish between clinical variables

that may have differential effects on seizure risk, which may either have an acute effect on the number of seizures the patient is likely to have at time t , or a more subacute effect on how high risk the patient is likely to become at time $t + 1$. The rationale for this model development is to capture the fact that some factors, such as electrolyte abnormalities or missing a medication dose, may exert a more acute influence on the expected number of seizures at the current time point; other factors, such as hormonal cycles, may exert a more long-term influence on the probability of being high risk at a future time point.

Marginal posterior probabilities of inclusion (MPPI) for transition probability regression coefficients β_k , by covariate, and for negative binomial regression coefficients ρ_k , by state and covariate, are shown in Fig 5. This figure illustrates the significance of covariate effects on either (a) transitions from one risk state to another or (b) the number of seizures given the current risk state. The reference state was set to $k = 3$ (the high seizure risk state) such that the regression coefficients corresponding to transitions into this state were zero, to maintain model identifiability. Selection of a different reference state resulted in a similar model fit based on DIC and variable inclusion agreement for the majority of regression coefficients. Selected covariates (i.e., $> 50\%$ marginal posterior probability) for state transitions are reported in Table 7, together with posterior means and credible intervals of the corresponding regression coefficients. For patients currently in low or moderate seizure risk states (1 or 2), the use of cannabidiol (CBD) was associated with a greater likelihood of remaining in those states, as opposed to worsening to the higher risk state 3. For patients currently in risk state 2 (moderate risk for GTCs), older patient age and treatment with zonisamide were associated with an increased likelihood of remaining at the same risk the following day, as opposed to worsening to the higher risk state 3.

Various triggers and ASMs were found to have an acute effect on the number of GTCs a patient could expect on a given day, given the current risk state (see Table S8 and Figure S5). The influence of patient-level covariates on modulating the number of seizures on a given day was greatest for patients currently in a “moderate” risk state. For patients in this risk state, bad mood, sudden changes in daily medication regimen (e.g. missed/late medications or changing medications), illness, and tiredness were associated with a greater number of expected seizures that day. This is consistent with research finding that sleep and anxiety are relevant in patient self-prediction of seizures (Haut et al., 2007). Among ASMs, we found that perampanel, triple or potassium bromide, and verapamil were associated with reducing the expected number of GTCs in state 2. Preliminary evidence of efficacy of these medications as adjunctive therapies has been identified in other studies as well (Yoshitomi et al., 2019).

5. Conclusion.

In this paper, we have developed a rigorous Bayesian non-homogeneous hidden Markov modeling approach for estimating seizure risk based on observed seizure counts and demonstrated its usage for investigating drivers of risk cycling in epilepsy. The approach accounts for the overdispersion exhibited in daily seizure count data due to seizure unpredictability and natural variability. It also incorporates exogenous clinical covariates into the estimation of model parameters, allowing for a more precise assessment of seizure

risk. Unlike existing approaches, we relax the assumption that clinical covariates that drive risk are known *a priori*, but rather simultaneously infer important drivers through variable selection priors. Our approach allows improved granularity of seizure risk assessment as well as an integrated framework for identifying drivers of risk cycles.

The utility of this model is demonstrated through our application analyzing self-recorded seizures by people with Dravet syndrome through [SeizureTracker.com](https://www.seizuretracker.com), a widely used seizure diary. We demonstrate the presence of three distinct states in Dravet syndrome at different risk for GTCs. We have characterized the volatility of these states, showing that patients with Dravet syndrome tend to self-transition (i.e., remain in the same state once in the state). We have shown that ASMs and triggers may drive fluctuations in seizure risk through differential effects on either (1) future state transitions, or (2) expected seizure frequency given the current risk state. This suggests that drivers of risk cycles can either exhibit a short-term effect, by reducing the number of expected seizures in the current time-frame, or a longer term effect, by modulating which risk state the patient is likely to transition to in the future. We have shown that incorporating the influence of external modulatory variables into statistical inference on seizure risk, such as medications and seizure triggers, improves the accuracy of risk estimation, compared to estimation based on seizure counting data alone.

The usefulness of our model for characterizing risk cycles in Dravet syndrome suggests potential for broader utility in the field of seizure risk modeling in epilepsy. Estimation of the MST of the high GTC risk state may potentially be relevant for guiding preventative counseling in risk for SUDEP and seizure-related injury. In particular, mean estimates and confidence intervals for the MST of high risk states are useful for counseling patients on how much time the average patient can expect to remain at high risk for GTCs once they enter a high risk state. For patients in high GTC risk states (and increased risk for SUDEP and seizure-related injury), an estimated MST of 5.8 days may suggest a possible approach of counseling patients in this state to take heightened seizure precautions for at least one week.

There are several limitations to this study. The medication data analyzed here is self-reported, and therefore results about the efficacy of pharmacological treatments on reducing seizure risk in Dravet syndrome should be considered with this caveat. This study primarily serves to validate our model's usage for investigating pharmacological relationships and extension to randomized clinical trial data is needed for rigorous studies of medication efficacy. Generally speaking, clinical applications of our method can directly improve patient care through a better understanding of the natural history of risk cycles in epilepsy, allowing clinicians to act proactively with drug interventions based on the understanding of expected fluctuations in risk cycles. Furthermore, the model can be used to understand the pharmacological effectiveness of therapies in modulating risk cycles in epilepsy. Although we have demonstrated our model's utility on patient-reported clinical seizure count data, our model is applicable to seizure counting data from any source, including intracranial, subdural, or scalp electroencephalography.

In our model parametrization, the number of parameters can grow quickly as the number of predictors or latent states increases, potentially resulting in overfitting. This is especially true for the transition model, as a separate set of regression coefficients is specified for each individual transition. In our approach, this issue is mitigated by the use of sparse priors for covariate selection.

Recent studies have suggested that individual gene expression may affect the resistance of seizures to ASMs (Naimo et al., 2019). In this context, our model is well-equipped to evaluate the effect and significance of varying genomic profiles on drug-refractory epilepsy. Furthermore, in the field of seizure forecasting and prediction, in addition to the predicted next risk state, the predicted absolute sequence of states may also be of clinical interest. Evaluation of how well the model performs in prospective on-line predictions of seizure risk will be of interest in future work, to determine contribution of this model to other seizure forecasting systems.

Our approach could be extended to infer the optimal number of states K via a sampling scheme that can accommodate posterior distributions of varied dimensions, for example reversible jump MCMC, though at the expense of an increase in computational cost. Our proposed model offers deep insights into the state-dependence of risk factors contributing to seizure propensity at both the Markov transitions and count emissions. Such inferences may provide great clinical value, for example through helping clinicians make more informed decisions regarding treatment plans for patients by leveraging knowledge of which medications or treatments are most effective conditional on the current seizure risk state of the patient. On the other hand, in situations where the clinical interpretation is not essential, a simpler model such as one of the submodels suggested in Appendix C should be preferred. As another example, a nonparametric approach to estimation of the count-valued emission distribution or transition probabilities would simultaneously improve flexibility and reduce the complexity of the model, though potentially at a cost of computation time and interpretability. For example, the general case of the method of Canale and Dunson (2011) models the count observations via a kernel mixture of counts but, unlike our model, does not assess the effect of additional covariates on the response. A multivariate extension allows for joint modeling of count observations with continuous and categorical predictors via a multivariate rounded mixture of Gaussians. However, this introduces additional complexities into the model and further complicates interpretation.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments.

We thank the people living with Dravet syndrome using [SeizureTracker.com](https://seizuretracker.com), who allowed their de-identified data to be used in this analysis. Data utilized in the seizure count case study was provided by Seizure Tracker™ - <https://seizuretracker.com/>, and should be directly requested from Seizure Tracker, LLC. ETW is supported by a fellowship from the Gulf Coast Consortia, on the NLM Training Program in Biomedical Informatics and Data Science T15LM007093. SC is supported by the National Institute of Neurological Disorders and Stroke, National Institutes of Health (5R25NS070680-12). The contents are solely the responsibility of the authors and do not necessarily represent the views of the NIH.

References.

- Arthurs S, Zaveri HP, Frei MG and Osorio I (2010). Patient and caregiver perspectives on seizure prediction. *Epilepsy & Behavior* 19 474–477. [PubMed: 20851054]
- Ayub N, Chiang S, Moss R and Goldenholz D (2020). Natural history of generalized motor seizures: a retrospective analysis. *Seizure* 80 109–112. [PubMed: 32563169]
- Barbieri MM and Berger JO (2004). Optimal predictive model selection. *Ann. Statist* 32 870–897.
- Baud MO, Kleen JK, Mirro EA, Andrechak JC, King-Stephens D, Chang EF and Rao VR (2018). Multi-day rhythms modulate seizure risk in epilepsy. *Nature Communications* 988.
- Baud MO, Proix T, Rao VR and Schindler K (2020). Chance and risk in epilepsy. *Current Opinion in Neurology* 33 163–172. [PubMed: 32049738]
- Brown PJ, Vannucci M and Fearn T (1998). Multivariate Bayesian variable selection and prediction. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 60 627–641.
- Canale A and Dunson DB (2011). Bayesian Kernel Mixtures for Counts. *Journal of the American Statistical Association* 106 1528–1539. [PubMed: 22523437]
- Chiang S, Vannucci M, Goldenholz DM, Moss R and Stern JM (2018). Epilepsy as a dynamic disease: A Bayesian model for differentiating seizure risk from natural variability. *Epilepsia Open* 3 236–246. [PubMed: 29881802]
- Chiang S, Goldenholz DM, Moss R, Rao VR, Haneef Z, Theodore WH, Kleen JK, Gavvala J, Vannucci M and Stern JM (2020). Prospective validation study of an epilepsy seizure risk system for outpatient evaluation. *Epilepsia* 61 29–38. [PubMed: 31792970]
- Chiang S, Khambhati AN, Wang ET, Vannucci M, Chang EF and RAO VR (2021). Evidence of state-dependence in the effectiveness of responsive neurostimulation for seizure modulation. *Brain Stimulation* 14 366–375. [PubMed: 33556620]
- Cooper MS, Mcintosh A, Crompton DE, McMahon JM, Schneider A, Farrell K, Ganesan V, Gill D, Kivity S, Lerman-Sagie T, McLellan A, Pelekanos J, Ramesh V, Sadleir L, Wirrell E and Scheffer IE (2016). Mortality in Dravet syndrome. *Epilepsy Research* 12843–47.
- Gelman A, Carlin JB, Stern HS, Dunson DB, Vehtari A and Rubin DB (2013). *Bayesian Data Analysis*, 3rd ed. Chapman & Hall/CRC Texts in Statistical Science. Taylor & Francis.
- George EI and McCulloch RE (1997). APPROACHES FOR BAYESIAN VARIABLE SELECTION. *Statistica Sinica* 7 339–373.
- Geweke J (1992). Evaluating the Accuracy of Sampling-Based Approaches to the Calculation of Posterior Moments. In *IN BAYESIAN STATISTICS* 169–193. University Press.
- Goldenholz DM, Goldenholz SR, Moss R, French J, Lowenstein D, Kuzniecky R, Haut S, Cristofaro S, Detyniecki K, Hixson J, Karoly P, Cook M, Strashny A and Theodore WH (2018). Is seizure frequency variance a predictable quantity? *Annals of Clinical and Translational Neurology* 5 201–207. [PubMed: 29468180]
- Harden C, Tomson T, Gloss D, Buchhalter J, Cross JH, Donner E, French JA, Gil-Nagel A, Hesdorffer DC, Smithson WH et al. (2017). Practice guideline summary: sudden unexpected death in epilepsy incidence rates and risk factors: Report of the Guideline Development, Dissemination, and Implementation Subcommittee of the American Academy of Neurology and the American Epilepsy Society. *Epilepsy Currents* 17 180–187. [PubMed: 28684957]
- Haut SR, Hall CB, Masur J and Lipton RB (2007). Seizure occurrence: precipitants and prediction. *Neurology* 69 1905–1910. [PubMed: 17998482]
- Holmes CC and Held L (2006). Bayesian auxiliary variable models for binary and multinomial regression. *Bayesian Anal* 1 145–168.
- Holsclaw T, Greene AM, Robertson AW and Smyth P (2017). Bayesian non-homogeneous Markov models via Pólya-Gamma data augmentation with applications to rainfall modeling. *Ann. Appl. Stat* 11 393–426.
- Hubin A (2019). An adaptive simulated annealing EM algorithm for inference on non-homogeneous hidden Markov models. *Proceedings of the International Conference on Artificial Intelligence, Information Processing and Cloud Computing - AIIPCC '19*.

- Karoly PJ, Goldenholz DM, Freestone DR, Moss RE, Grayden DB, Theodore WH and Cook MJ (2018). Circadian and circaseptan rhythms in human epilepsy: a retrospective cohort study. *Lancet Neurology* 17 977–985. [PubMed: 30219655]
- Leguia MG, Andrzejak RG, Rummel C, Fan JM, Mirro EA, Tcheng TK, Rao VR and Baud MO (2021). Seizure cycles in focal epilepsy. *JAMA Neurology*.
- Muenz LR and Rubinstein LV (1985). Markov models for covariate dependence of binary sequences. *Biometrics* 41 91–101. [PubMed: 4005390]
- Naimo GD, Guarnaccia M, Sprovieri T, Ungaro C, Conforti FL, Andò S and Cavallaro S (2019). A Systems Biology Approach for Personalized Medicine in Refractory Epilepsy. *International Journal of Molecular Sciences* 203717.
- Panayiotopoulos C (2005). Chapter 4, Principles of Therapy in Epilepsies. Panayiotopoulos CP. *The Epilepsies: Seizures, Syndromes and Management*. Oxford-shire, UK: Bladon Medical Publishing.
- Pillow JW and Scott J (2012). Fully Bayesian inference for neural models with negative-binomial spiking. In *Advances in Neural Information Processing Systems 25* (Pereira F, Burges CJC, Bottou L and Weinberger KQ, eds.) 1898–1906. Curran Associates, Inc.
- Polson NG, Scott JG and Windle J (2013). Bayesian Inference for Logistic Models Using Pólya-Gamma Latent Variables. *Journal of the American Statistical Association* 108 1339–1349.
- Privitera MD, Welty TE, Gidal BE, Diaz FJ, Krebill R, Szaflarski JP, Dworetzky BA, Pollard JR, Elder EJ Jr, Jiang W et al. (2016). Generic-to-generic lamotrigine switches in people with epilepsy: the randomised controlled EQUIGEN trial. *The Lancet Neurology* 15 365–372. [PubMed: 26875743]
- Proix T, Truccolo W, Leguia MG, Tcheng TK, King-Stephens D, Rao VR and Baud MO (2020). Forecasting seizure risk in adults with focal epilepsy: a development and validation study. *The Lancet Neurology*.
- Quenouille MH (1949). A Relation between the Logarithmic, Poisson, and Negative Binomial Series. *Biometrics* 5 162–164. [PubMed: 18151958]
- Rabiner LR (1989). A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition. *Readings in Speech Recognition* 77 257–286.
- Rao VR, G Leguia M, Tcheng TK and Baud MO (2020). Cues for seizure timing. *Epilepsia*.
- Savitsky T, Vannucci M and Sha N (2011). Variable Selection for Nonparametric Gaussian Process Priors: Models and Computational Strategies. *Statist. Sci* 26 130–149.
- Scott SL (2002). Bayesian Methods for Hidden Markov Models: Recursive Computing in the 21st Century. *Journal of the American Statistical Association* 97 337–351.
- Stingo FC, Chen YA, Vannucci M, Barrier M and Mirkes PE (2010). A Bayesian graphical modeling approach to microRNA regulatory network inference. *The annals of applied statistics* 4 2024–2048. [PubMed: 23946863]
- Tharayil JJ, Chiang S, Moss R, Stern JM, Theodore WH and Goldenholz DM (2017). A big data approach to the development of mixed-effects models for seizure count data. *Epilepsia* 58 835–844. [PubMed: 28369781]
- Vossler DG, Anderson GD and Bainbridge J (2016). AES position statement on generic substitution of antiepileptic drugs. *Epilepsy Currents* 16 209–211. [PubMed: 27330454]
- Walczak TS, Leppik IE, D’amelio M, Rarick J, So E, Ahman P, Ruggles K, Cascino GD, Annegers J and Hauser W (2001). Incidence and risk factors in sudden unexpected death in epilepsy: A prospective cohort study. *Neurology* 56 519–525. [PubMed: 11222798]
- WHO, ILAE and IBE (2019). Epilepsy: A public health imperative.
- Wu YW, Sullivan J, McDaniel SS, Meisler MH, Walsh EM, Li SX and Kuzniewicz MW (2015). Incidence of Dravet syndrome in a US population. *Pediatrics* 136 e1310–e1315. [PubMed: 26438699]
- Yoshitomi S, Takahashi Y, Yamaguchi T, Imai K, Ishit A, Hirose S and Inoue Y (2019). Efficacy and tolerability of perampanel in pediatric patients with Dravet syndrome. *Epilepsy Research* 154 34–38. [PubMed: 31035242]
- Zhou M and Carin L (2015). Negative Binomial Process Count and Mixture Modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 37 307–320. [PubMed: 26353243]

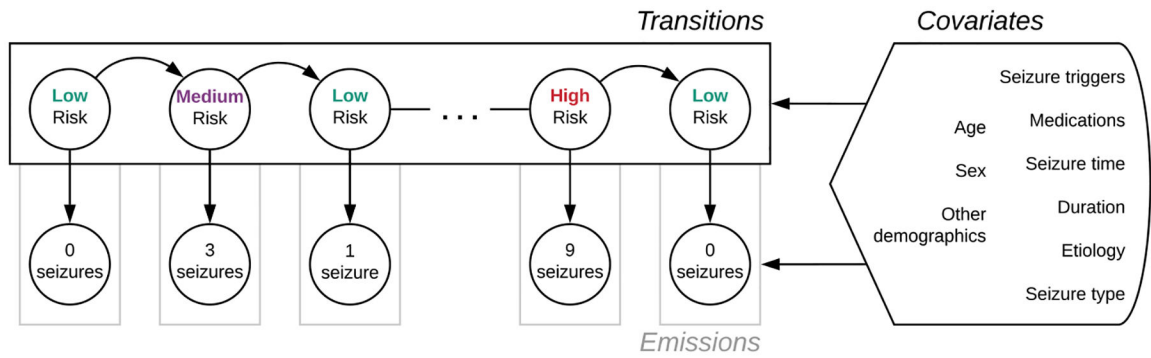


Fig 1. Proposed hidden Markov model for seizure risk estimation (ZINB-NHMM-BVS). Depending on exogenous covariates, patients transition between different seizure risk states from day to day. Then, conditional upon the states, seizure counts are emitted from zero-inflated negative binomial distributions, which are also dependent on covariates.

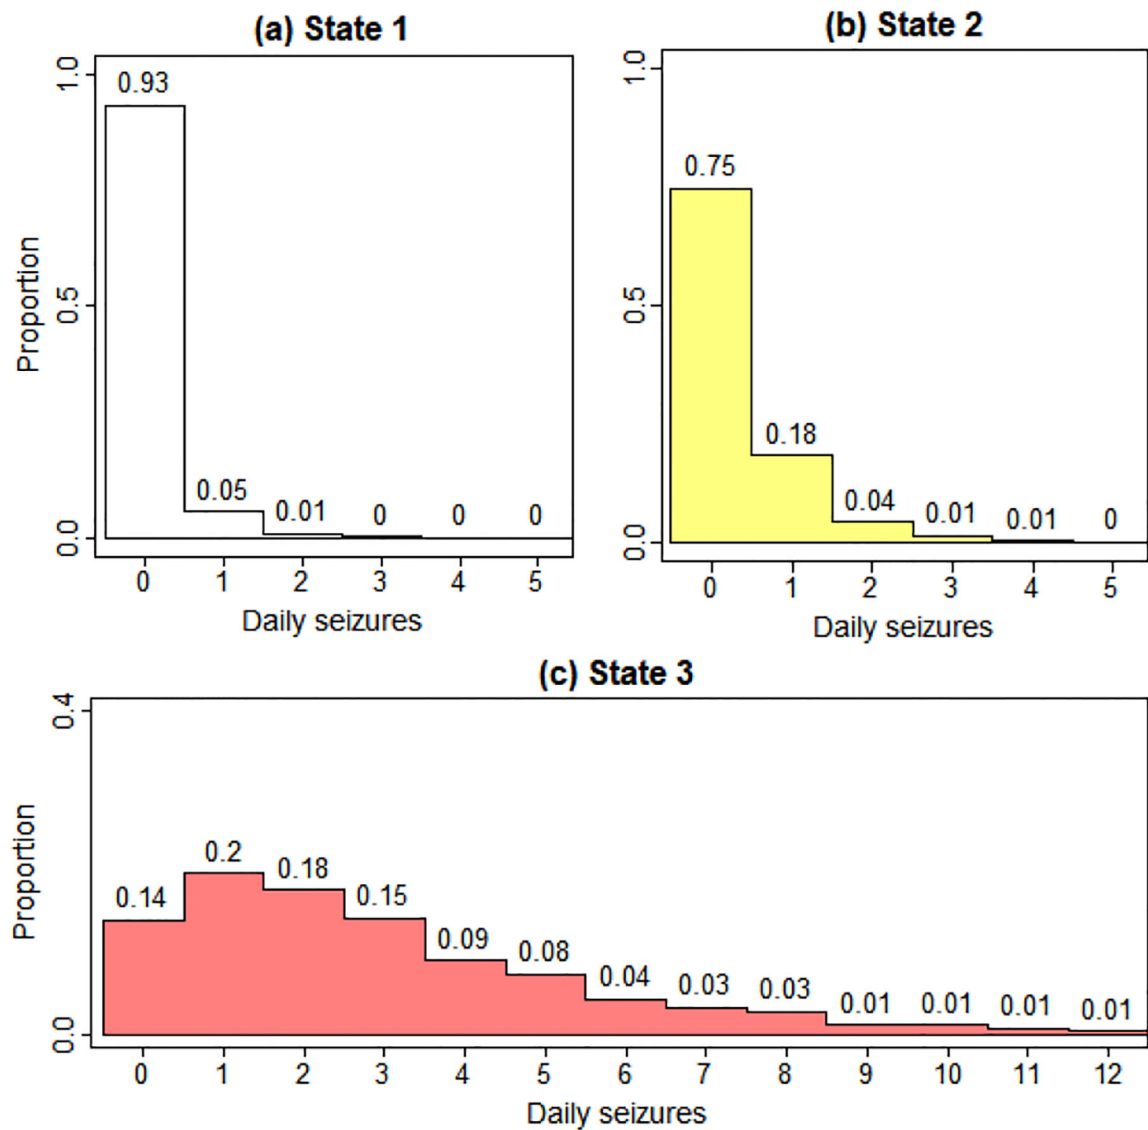


Fig 2. Dravet syndrome Seizure Tracker data analysis: Distributions of daily GTC counts in each latent risk state.

These are estimated in our model by first estimating the latent states ξ via the posterior mode and then plotting the empirical distribution of the observed daily counts in each state.

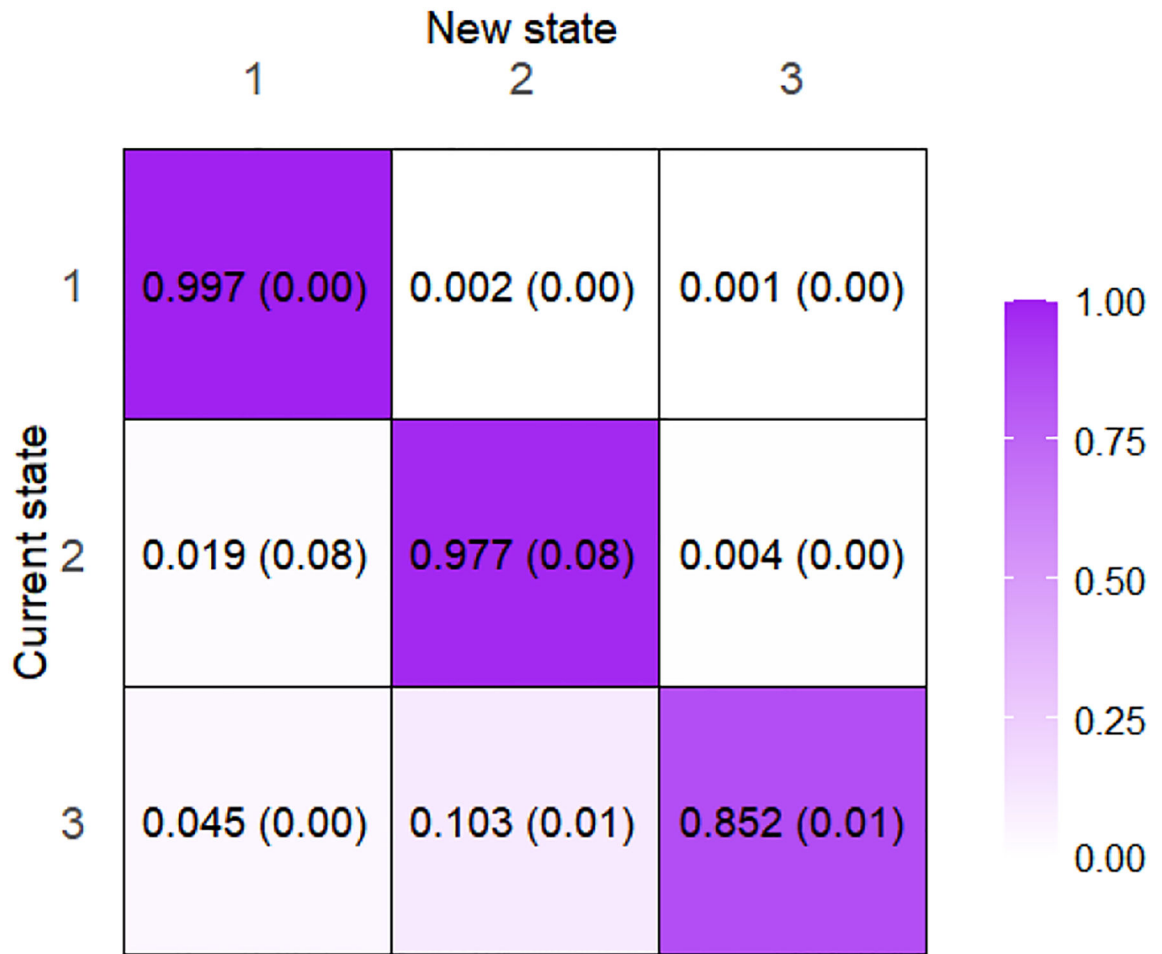


Fig 3. Dravet syndrome Seizure Tracker data analysis: Mean transition matrix.

These are estimated in our model by first calculating $Pr(\xi_{it} = k \mid \xi_{i,t-1} = k', \cdot)$ via the multinomial logistic equation (2.1) with coefficients $\beta_{k'}$ estimated by the posterior mean, then averaging across all patients and days for each transition. Standard deviations (SD) are shown in parentheses.

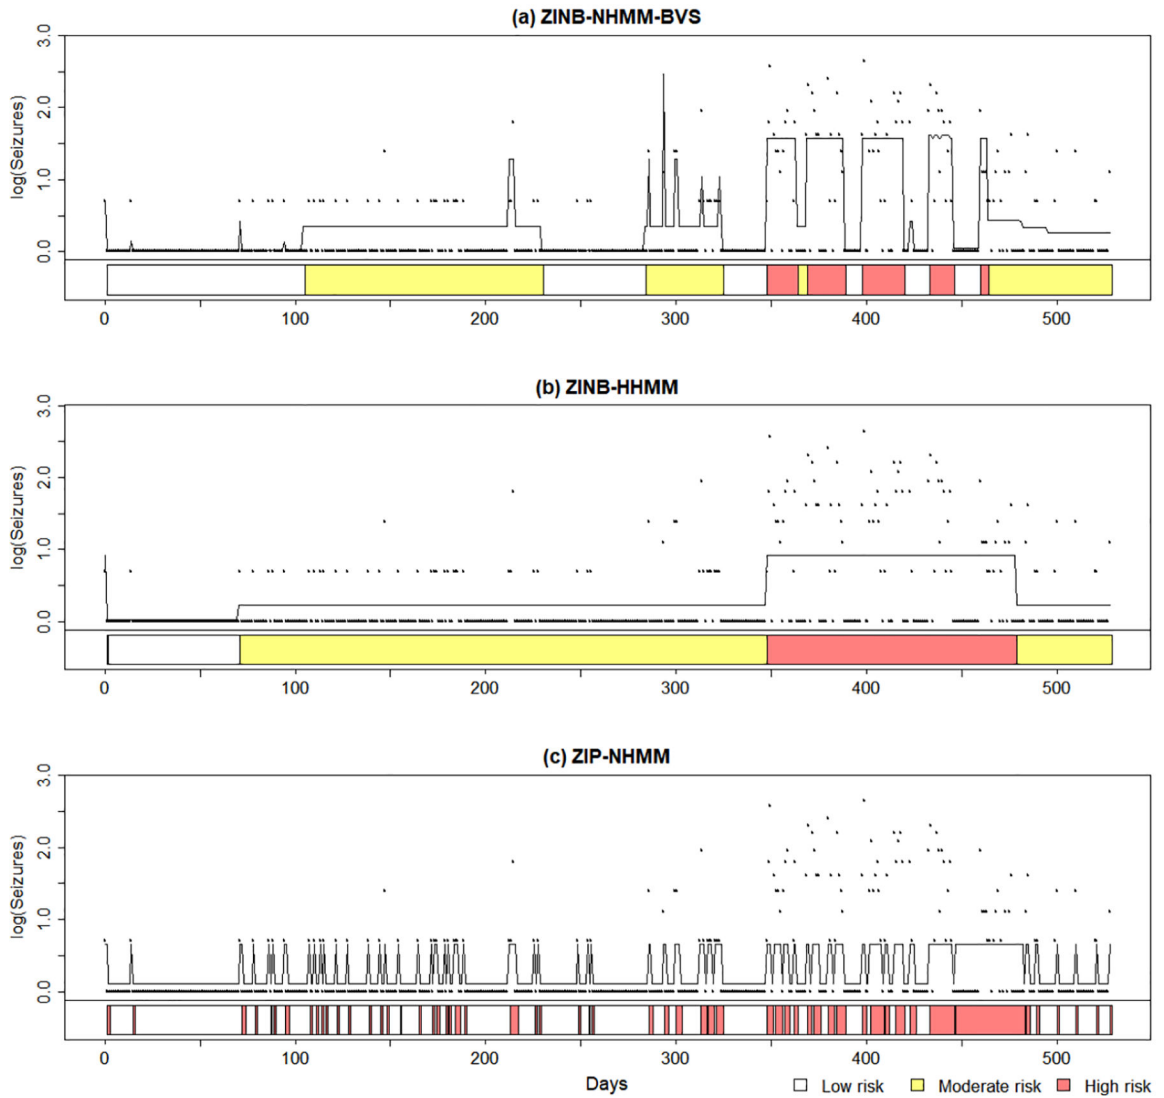


Fig 4. Dravet syndrome Seizure Tracker data analysis: Expected number of seizures for the sequence of daily seizure counts from one patient. These are estimated with (a) the proposed non-homogeneous HMM (ZINB-NHMM-BVS), (b) a baseline homogeneous HMM (ZINB-HHMM), and (c) the method of Chiang et al. (2018) (ZIP-NHMM). Values are shown on the log scale. Unlike ZIP-NHMM and ZINB-HHMM, which produce state-specific estimates of risk, ZINB-NHMM-BVS results in subject- and state-specific estimates, obtained as the estimated NB mean parameters, $\mu_{itk} = \frac{\psi_{itk} r_k}{1 - \psi_{itk}}$, with overdispersion parameters r_k estimated via posterior mean and probabilities ψ_{itk} calculated via (2.3), with latent risk states ξ_{it} estimated by posterior mode and regression coefficients ρ_k via posterior mean. The heatmap at the bottom of each panel shows the most likely sequence of risk states for that method.

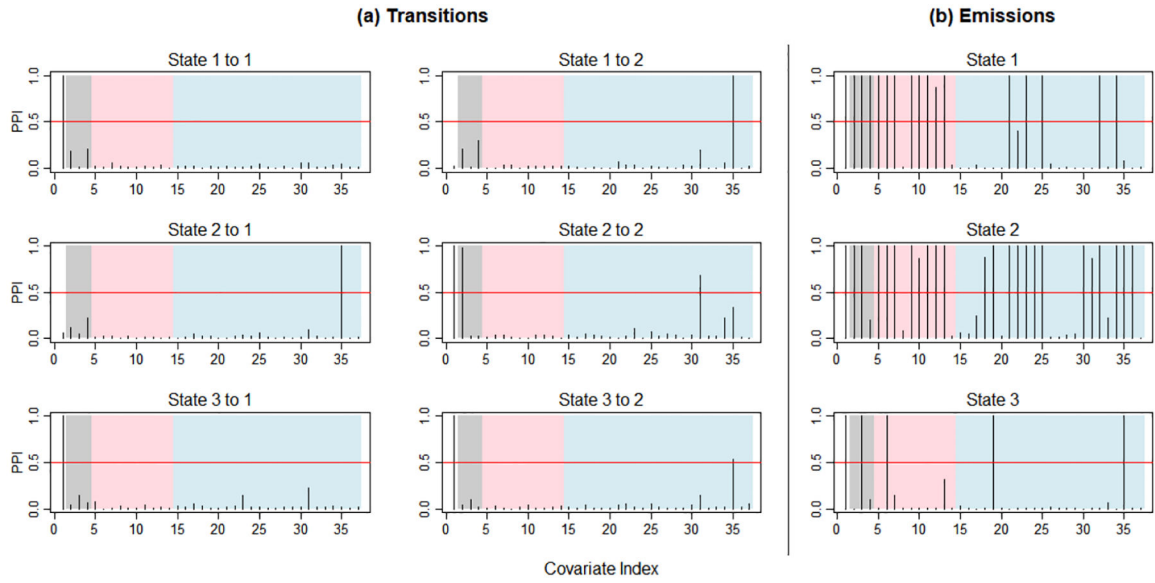


Fig 5. Dravet syndrome Seizure Tracker data analysis: Marginal posterior probabilities of inclusion (MPPI). *MPPI are shown for (a) transition probability regression coefficients β_k by covariate, and for (b) negative binomial regression coefficients ρ_k by state and covariate. Variable inclusion criterion is set at $MPPI > 0.5$ and is indicated by the red line. Covariates are split into three broad categories: patient demographics (gray area), seizure triggers (pink area), and medication classes (blue area). Left-most covariate (position 1) is the bias term. Baseline state is $k = 3$. A key indicating the variable name corresponding to each covariate index can be found in Table S7.*

Table 1
Simulation study (negative binomial): Sensitivity analysis. Classification metrics for latent state estimation.

Different slab variance parameters, Σ_ρ and Σ_β , and different Beta priors on the inclusion indicators are considered. The default set is $\Sigma_\beta = \text{diag}(1)$, $g_\beta = 1$, $h_\beta = 5$, $\Sigma_\rho = \text{diag}(1)$, and $g_\rho = 1$, $h_\rho = 5$, as described in Section “Parameter settings”. Results are averaged across 20 replicated datasets.

Model	Acc.	Prec.	Sens.	Spec.	F ₁
Default	0.99	0.98	0.96	0.97	0.97
$\Sigma_\beta = \text{diag}(0.1)$	0.99	0.97	0.95	0.97	0.96
$\Sigma_\beta = \text{diag}(10)$	0.99	0.97	0.96	0.97	0.97
$\Sigma_\beta = \text{diag}(100)$	0.99	0.97	0.96	0.97	0.96
$g_\beta = 1, h_\beta = 1$	0.99	0.98	0.96	0.97	0.97
$g_\beta = 1, h_\beta = 2$	0.99	0.98	0.96	0.97	0.97
$g_\beta = 1, h_\beta = 20$	0.99	0.97	0.96	0.97	0.97
$\Sigma_\rho = \text{diag}(0.1)$	0.99	0.98	0.96	0.97	0.97
$\Sigma_\rho = \text{diag}(10)$	0.99	0.98	0.96	0.97	0.97
$\Sigma_\rho = \text{diag}(100)$	0.99	0.98	0.96	0.97	0.97
$g_\rho = 1, h_\rho = 1$	0.99	0.98	0.96	0.97	0.97
$g_\rho = 1, h_\rho = 2$	0.99	0.98	0.96	0.97	0.97
$g_\rho = 1, h_\rho = 20$	0.99	0.98	0.96	0.97	0.97

Table 2
Simulation study (negative binomial): Sensitivity analysis. Classification metrics for covariate selection.

Different slab variance parameters of the regression coefficients, Σ_ρ and Σ_β , and different Beta priors on the variable inclusion indicators are considered. The default set is $\Sigma_\rho = \text{diag}(1)$, $g_\beta = 1$, $h_\beta = 5$, $\Sigma_\rho = \text{diag}(1)$, and $g_\rho = 1$, $h_\rho = 5$, as described in Section “Parameter settings”. “% included” refers to the prior expected number of included variables, and “# Selected” refers to the number of variables selected by the median probability model. Results are averaged over 20 simulated datasets. The true data-generating model has 19 true non-zero covariates affecting Markov transitions, and 14 non-zero covariates affecting emissions.

<i>Transitions</i>								
	% included	# Selected	FNR	FPR	Prec.	Sens.	Spec.	F ₁
Default	17%	20.60	0.02	0.09	0.91	0.98	0.91	0.94
$\Sigma_\rho = \text{diag}(0.1)$	17%	26.50	0.05	0.37	0.69	0.95	0.63	0.80
$\Sigma_\rho = \text{diag}(10)$	17%	20.65	0.04	0.10	0.89	0.96	0.90	0.92
$\Sigma_\rho = \text{diag}(100)$	17%	24.55	0.02	0.26	0.79	0.98	0.74	0.86
$g_\beta = 1, h_\beta = 1$	50%	41.90	0.00	1.00	0.45	1.00	0.00	0.62
$g_\beta = 1, h_\beta = 2$	33%	26.60	0.01	0.34	0.73	0.99	0.66	0.83
$g_\beta = 1, h_\beta = 20$	5%	19.10	0.05	0.04	0.95	0.95	0.96	0.95
<i>Emissions</i>								
	% included	# Selected	FNR	FPR	Prec.	Sens.	Spec.	F ₁
Default	17%	13.95	0.02	0.04	0.98	0.98	0.96	0.98
$\Sigma_\rho = \text{diag}(0.1)$	17%	14.00	0.02	0.04	0.98	0.98	0.96	0.98
$\Sigma_\rho = \text{diag}(10)$	17%	13.95	0.02	0.04	0.98	0.98	0.96	0.98
$\Sigma_\rho = \text{diag}(100)$	17%	14.05	0.02	0.05	0.98	0.98	0.95	0.98
$g_\rho = 1, h_\rho = 1$	50%	16.60	0.00	0.38	0.85	1.00	0.62	0.91
$g_\rho = 1, h_\rho = 2$	33%	15.05	0.01	0.16	0.93	0.99	0.84	0.96
$g_\rho = 1, h_\rho = 20$	5%	13.50	0.04	0.00	1.00	0.96	1.00	0.98

Table 3
Simulation study (negative binomial): Sensitivity analysis. Classification metrics for covariate selection, using different model selection approaches.

Results are averaged across 20 replicates each. “# Selected” refers to the total number of variables selected by each model selection approach.

<i>Transitions</i>							
Model selection method	# Selected	FNR	FPR	Prec.	Sens.	Spec.	F_1
Median probability	20.90	0.03	0.10	0.89	0.97	0.90	0.93
Most probable	20.55	0.03	0.10	0.90	0.97	0.90	0.93
AIC	21.10	0.08	0.16	0.84	0.92	0.84	0.87
BIC	19.50	0.09	0.10	0.89	0.91	0.90	0.90
<i>Emissions</i>							
Model selection method	# Selected	FNR	FPR	Prec.	Sens.	Spec.	F_1
Median probability	13.95	0.02	0.04	0.98	0.98	0.96	0.98
Most probable	13.95	0.02	0.03	0.99	0.98	0.97	0.98
AIC	14.45	0.02	0.11	0.95	0.98	0.89	0.96
BIC	13.90	0.03	0.04	0.98	0.97	0.96	0.97

Table 4
Simulation study (Poisson): Classification metrics for covariate selection.

The proposed model and the ASA-EM method are compared in terms of false positive and false negative rates, precision (Prec.), sensitivity (Sens.), specificity (Spec.), and F_1 score (F_1), for both the transition and emission components of the HMM. For the proposed ZINB model, different prior inclusion hyperparameters with the median model and alternative model selection approaches (AIC, BIC) were considered. Additionally, “% included” refers to the prior expected number of included variables, and “# Selected” refers to the number of variables selected by the models. Results are averaged across 30 replicates.

<i>Transitions</i>								
Model	% included	# Selected	FNR	FPR	Prec.	Sens.	Spec.	F_1
$g_\beta = 1, h_\beta = 3$	25%	8.87	0.00	0.04	0.91	1.00	0.96	0.95
$g_\beta = 1, h_\beta = 5$	17%	8.20	0.00	0.01	0.97	1.00	0.99	0.98
$g_\beta = 1, h_\beta = 10$	9%	8.30	0.00	0.01	0.97	1.00	0.99	0.98
ZINB, AIC	17%	10.33	0.00	0.11	0.79	1.00	0.89	0.87
ZINB, BIC	17%	9.47	0.00	0.07	0.86	1.00	0.93	0.92
ASA-EM, AIC	-	15.13	0.29	0.43	0.38	0.71	0.57	0.52
ASA-EM, BIC	-	8.80	0.52	0.22	0.48	0.48	0.78	0.57
<i>Emissions</i>								
Model	% included	# Selected	FNR	FPR	Prec.	Sens.	Spec.	F_1
$g_\rho = 1, h_\rho = 3$	25%	11.00	0.00	0.14	0.75	1.00	0.86	0.85
$g_\rho = 1, h_\rho = 5$	17%	9.83	0.00	0.08	0.83	1.00	0.92	0.90
$g_\rho = 1, h_\rho = 10$	9%	9.57	0.00	0.07	0.85	1.00	0.93	0.92
ZINB, AIC	17%	11.37	0.00	0.15	0.72	1.00	0.85	0.83
ZINB, BIC	17%	10.57	0.00	0.12	0.78	1.00	0.88	0.87
ASA-EM, AIC	-	15.13	0.07	0.35	0.54	0.93	0.65	0.67
ASA-EM, BIC	-	14.00	0.05	0.29	0.62	0.95	0.71	0.73

Table 5
Simulation study (Poisson): Classification metrics for latent states.

We compare (1) our proposed ZINB model with varying prior inclusion hyperparameters used in conjunction with the median model, and two alternative model selection approaches (AIC, BIC) and (2) ASA-EM, averaged across 30 simulated datasets. Metrics included are accuracy (Acc.), precision (Prec.), sensitivity (Sens.), specificity (Spec.) and F_1 score (F_1), with the first state set as the baseline state.

Model	Acc.	Prec.	Sens.	Spec.	F_1
$g_\beta = 1, h_\beta = 3$	0.97	0.98	0.98	0.95	0.98
$g_\beta = 1, h_\beta = 5$	0.97	0.98	0.98	0.95	0.98
$g_\beta = 1, h_\beta = 10$	0.97	0.98	0.98	0.95	0.98
$g_\rho = 1, h_\rho = 3$	0.97	0.98	0.98	0.95	0.98
$g_\rho = 1, h_\rho = 5$	0.97	0.98	0.98	0.95	0.98
$g_\rho = 1, h_\rho = 10$	0.97	0.98	0.98	0.95	0.98
ZINB, AIC	0.97	0.98	0.97	0.97	0.98
ZINB, BIC	0.97	0.99	0.97	0.97	0.98
ASA-EM, AIC	0.90	0.88	0.98	0.72	0.93
ASA-EM, BIC	0.89	0.87	0.98	0.69	0.92

Table 6Characteristics of Dravet syndrome study sample ($n = 133$).

Characteristic	Value
Demographic characteristics	
Age in years, mean (SD)	8.25 (6.55)
Male sex [‡]	51.1%
Seizure triggers [‡]	
Change in medications	7.1%
Tiredness	13.3%
Overheating	3.7%
Illness	5.8%
Bad mood	2.9%
Seizure burden [*]	
Duration in minutes, mean (SD)	1.94 (5.42)
Status epilepticus [§]	4.8%
Use of VNS therapy [‡]	21.1%
Seizure frequency [*]	
Seizures per day, mean (SD)	0.24 (0.92)
Mean proportion of seizure-free days	85.4%
Mean proportion of days with one seizure	10.2%
Mean proportion of days with 2+ seizures	4.4%
Anti-seizure medication usage [‡]	
Benzodiazepines	43.6%
Valproic acid	41.4%
Stiripentol	24.8%
Levetiracetam or brivaracetam	21.8%
Topiramate	20.3%
Cannabidiol	18.0%
Tetrahydrocannabinols	7.5%
Zonisamide	6.8%
Phenobarbital or phenobarbital-containing compounds	6.0%
Felbamate	4.5%
Ethosuximide or methsuximide	3.8%
Lacosamide	3.0%
Carbonic anhydrase inhibitors	2.3%
Lamotrigine	2.3%
Prednisolone	1.5%
Triple bromide or potassium bromide	1.5%
Carbamazepine, eslicarbazepine, oxcarbazepine	1.5%
Phenytoin	0.8%

Characteristic	Value
Primidone	0.8%
Perampanel	0.8%
Other medications [‡]	
Verapamil	3.0%
Risperidone	2.3%
Vitamin B6	2.3%

For continuous variables, mean (SD) are shown.

For categorical variables, proportion of seizures ([†]) or patients ([‡]) is shown.

* Only generalized tonic-clonic seizures were considered.

[§] Status epilepticus defined as seizure with duration >5 minutes.

Table 7
Dravet syndrome Seizure Tracker data analysis: Posterior

means and SD for transition probability regression coefficients, $\beta_{k|k}$, of selected covariates. These are reported with 95% credible intervals (defined as the interval $[a, b]$ such that $Pr(a \leq \beta_{j,k|k} \leq b|data) = 0.95$) and marginal posterior probabilities of inclusion (MPPI). Only covariates with MPPI > 0.5 are shown. Baseline state is $k = 3$.

Transition	Covariate	Post. mean (SD)	MPPI	95% CI
1→1	Intercept	6.65 (0.16)	1.00	(6.38,6.95)
2→2	Intercept	4.98 (0.19)	1.00	(4.64,5.36)
3→1	Intercept	-2.85 (0.33)	1.00	(-3.54,-2.23)
3→2	Intercept	-2.04 (0.33)	1.00	(-2.89,-1.58)
2→2	Age	2.99 (0.59)	0.98	(1.81,4.16)
2→2	Zonisamide	5.47 (0.94)	0.68	(3.29,7.06)
1→2	Cannabidiol	3.15 (0.49)	1.00	(2.33,4.12)
2→1	Cannabidiol	5.50 (0.85)	1.00	(3.54,6.62)
3→2	Cannabidiol	-2.36 (0.70)	0.53	(-3.90,-1.13)