JIRSS (2026)
Vol. 25, No. 2, pp 157-188
DOI: 10.22034/jirss.2026.2067102.1127

Bayesian change-point detection in zero-inflated regression models: Analyzing COVID-19 mortality patterns

Masoud Majidizadeh[1]

[1]Department of Statistics, Faculty of Mathematical Sciences, Shahid Beheshti University, Tehran, Iran.

Received: 2025-07-26, Accepted: 2026-07-25, Published online: 2026-09-03

Abstract. Maintaining consistent experimental conditions over long periods is inherently challenging, and even when the underlying model structure remains stable, the assumption of constant model parameters may not hold. In this paper, we investigate the estimation of a single change-point within a flexible class of zero-inflated models, including well-known cases such as the Zero-Inflated Poisson (ZIP) model. Our goal is to characterize the change-point through shifts in model parameters by employing a fully Bayesian framework. To achieve this, we develop a Reversible Jump Markov Chain Monte Carlo (RJMCMC) algorithm for parameter estimation and model selection, specifically tailored to detect the presence and location of a potential change-point in zero-inflated regression models, with particular emphasis on the ZIP formulation. This Bayesian approach allows practitioners to incorporate expert knowledge directly into the analysis of zero-inflated count data. We assess the performance and properties of the proposed method through a series of simulation studies. In addition, we apply our approach to COVID-19 mortality data from two countries to identify change-points in death trends while accounting for relevant explanatory variables.

Keywords. Bayesian inference, Data augmentation, RJMCMC, Zero-inflated power series models, Change-point detection.

MSC: 62F15, 62M10, 65C40, 62J12.

Corresponding Author: M. Majidizadeh (m_majidizadeh@sbu.ac.ir).

1 Introduction

In epidemiological and public health research, the analysis of count data presents unique statistical challenges, particularly when such data exhibit a high prevalence of zero values. For instance, in studies of health services utilization, a significant portion of the population may report no service use, while in substance abuse research, many individuals might report zero days of consumption during a treatment period. This phenomenon, known as zero-inflation, renders standard count data models like the Poisson or negative binomial distributions inadequate, as they fail to account for the excess zeros (Lambert, 1992). The necessity for specialized statistical methods is not confined to public health but extends to diverse fields such as industrial manufacturing (Lambert, 1992) and biomedical sciences (Heilbron and Gibson, 1990; Hall, 2000). Addressing this issue with appropriate analytical techniques is crucial for uncovering accurate underlying patterns and informing effective policy-making.

The Zero-Inflated Poisson (ZIP) regression model, formally introduced by Lambert (1992), has become a cornerstone for analyzing such data. The model conceptualizes the data as arising from a two-part process: with a certain probability, an observation is "structurally" zero (an "always-zero" group), and with the complementary probability, it is generated from a standard Poisson distribution, which can itself produce zeros. The parameters governing both the zero-inflation probability and the Poisson mean can be modeled as functions of covariates. The utility of this approach has been widely recognized in biometrics and beyond, offering a robust alternative where standard models fail (Böhning, 1998). Further extensions, such as the Zero-Inflated Negative Binomial (ZINB) for handling over-dispersion in non-zero counts (Ridout et al., 2001; Perumean-Chaney et al., 2013) and the Zero-Inflated Binomial (ZIB) for upper-bounded counts (Hall, 2000), have broadened the applicability of this modeling framework. Hurdle models represent an alternative conceptualization, separating the analysis into a binary model for zero versus non-zero counts and a truncated count model for positive counts (Loeys et al., 2012; Cheung, 2002).

Estimation of these models can be approached from two primary statistical paradigms: frequentist and Bayesian. The frequentist approach, typically relying on Maximum Likelihood Estimation (MLE), is well-established. Score tests have been developed to formally compare standard Poisson models against their zero-inflated counterparts (Jansakul and Hinde, 2002; Van den Broek, 1995), and modern software packages like glmmTMB in R provide efficient and flexible tools for fitting a wide array of zero-inflated generalized linear (mixed) models (Brooks et al., 2017). However, the Bayesian framework offers distinct advantages, especially for complex hierarchical structures. By treating parameters as random variables, Bayesian methods provide a natural way to quantify uncertainty, incorporate prior knowledge, and handle complex dependencies, such as spatial correlations through random effects (Agarwal et al., 2002). Seminal works have developed Bayesian estimation procedures for various zero-inflated models using simulation-based techniques like Markov Chain Monte Carlo (MCMC), demonstrating their flexibility and robustness (Ghosh et al., 2006; Rodrigues, 2003; Angers and Biswas, 2003; Özmen and Demirhan, 2010). Recent advancements have further

enhanced computational efficiency through techniques like Pólya-Gamma data augmentation (Neelon, 2019) and the use of sophisticated shrinkage priors (Bhattacharyya et al., 2022; Klein et al., 2015).

The challenge of modeling zero-inflated data is compounded when analyzing longitudinal or time-series data, where the underlying data-generating process may undergo abrupt structural changes. Such change-points are of paramount interest in epidemiology, as they can signify the impact of public health interventions, the emergence of a new virus variant, or shifts in population behavior. The recent COVID-19 pandemic provides a compelling context, with numerous studies employing zero-inflated time-series models to analyze daily infection or mortality counts and identify trends (Khedhiri, 2021; Tawiah et al., 2021; Junnumtuam et al., 2023; Zhou et al., 2024). Several authors have explored modeling parameter changes in longitudinal count data, for instance, by developing CUSUM tests for zero-inflated autoregressive models (Lee et al., 2016) or by introducing random change-points to capture shifts in treatment efficacy over time (Wen et al., 2024; Grimm and Stegmann, 2019). Recent literature has also highlighted the utility of Bayesian change-point frameworks specifically for modeling COVID-19 infection and mortality dynamics (Majidizadeh and Taheriyoun, 2024; Majidizadeh, 2024).

However, a significant methodological gap remains. While standard frequentist methods for change-point detection, such as CUSUM tests, are well-established, they are typically designed for normally distributed data and are ill-suited for series characterized by the zero-inflation and over-dispersion inherent in epidemiological count data. Furthermore, methods based on MLE may provide a point estimate for a change-point's location but fail to offer a full probabilistic characterization of its uncertainty. To address this gap, we propose a novel and comprehensive solution: a Bayesian Reversible Jump Markov Chain Monte Carlo (RJMCMC) methodology for detecting a single change-point in Zero-Inflated Poisson regression models. The Bayesian framework is chosen to explicitly model the complex data structure and provide robust uncertainty quantification. The key innovation lies in the use of RJMCMC (Green, 1995), a powerful trans-dimensional MCMC technique. This approach allows the sampler to jump between competing models—specifically, a model with one segment (no change-point) and a model with two segments (one change-point). Consequently, our method not only estimates the location and magnitude of a potential change but also calculates the posterior probability of its very existence, providing a principled and integrated approach to model selection and parameter estimation. While Green (1995) pioneered this method, its application to discrete-time, zero-inflated regression models represents a significant and non-trivial extension.

The remainder of this paper is organized as follows. Section 2 details the proposed change-point model for ZIP regression. Section 3 describes the Bayesian estimation procedure, with a focus on the data augmentation strategy and the RJMCMC algorithm. In Section 4, we conduct a simulation study to rigorously evaluate the performance of our proposed method. Section 5 demonstrates the practical utility of our approach through an application to real-world COVID-19 mortality data from Cyprus and Qatar, illustrating its ability to uncover meaningful structural changes. Finally, Section 6 dis-

cusses the implications of our findings, limitations of the current study, and directions for future research.

2 Model and Notations

In this section, we specify the change-point model for zero-inflated count data. We assume that the time series of length n is potentially divided into two segments by a single change-point at an unknown location $m \in \{1, \ldots, n-1\}$. If no change-point exists, we can formally set $m = n$.

The model parameters are assumed to be constant within each segment but may differ between segments. To formalize this, we define an indicator variable s_t such that $s_t = 1$ for observations $t = 1, \ldots, m$ (before the change-point) and $s_t = 2$ for observations $t = m+1, \ldots, n$ (after the change-point). Consequently, the parameters of the zero-inflated distribution for observation Y_t depend on the segment it belongs to. We denote these segment-specific parameters as (p_{s_t}, θ_{s_t}). Our primary objective is to estimate the change-point location m and the sets of parameters corresponding to each segment.

A zero-inflated random variable Y_t can be conceptualized as $Y_t = V_t(1 - B_t)$, where B_t is a Bernoulli variable with the segment-specific probability p_{s_t}, and V_t follows a discrete distribution (e.g., Poisson(θ_{s_t}), Negative Binomial(θ_{s_t}, r_{s_t}), or a Power Series distribution with parameter θ_{s_t}). The mean and variance of Y_t are then given by:

$$\begin{aligned} E(Y_t) &= (1 - p_{s_t})E(V_t), \\ \mathrm{Var}(Y_t) &= \frac{p_{s_t}}{1 - p_{s_t}}[E(Y_t)]^2 + \delta_{s_t}E(Y_t). \end{aligned}$$

Here, $\delta_{s_t} = \mathrm{Var}(V_t)/E(V_t)$ is the variance-to-mean ratio (or index of dispersion) of the underlying count distribution V_t. For instance, if V_t follows a Poisson distribution, $\delta_{s_t} = 1$, whereas for a Negative Binomial distribution, $\delta_{s_t} > 1$, capturing overdispersion in the count component of the model.

We focus on the Power Series distribution for its theoretical insights. The probability mass function of the ZIPS distribution, denoted as ZIPS(p_t, θ_t), is:

$$\begin{aligned} P(Y_t = 0) &= p_{s_t} + (1 - p_{s_t})\frac{b(0)}{c(\theta_{s_t})}, \\ P(Y_t = k) &= (1 - p_{s_t})\frac{b(k){\theta_{s_t}}^k}{c(\theta_{s_t})}, \quad k = 1, 2, \ldots, \end{aligned}$$

where $c(\theta_t) = \sum_{k=0}^{\infty} b(k){\theta_t}^k$, $0 \le p_t < 1$, and $\theta_t > 0$. When $p_t = 0$, the ZIPS distribution reduces to a regular PS(θ_t) distribution. The Power Series (PS) class of distributions has a probability mass function (PMF) of the form

$$P(Y = y) = \frac{b(y)\,\theta^y}{c(\theta)}, \qquad y \in \{0, 1, 2, \ldots\}.$$

The form of these functions defines the specific distribution. For instance:

- **Poisson distribution:**

$$\theta = \lambda > 0, \qquad b(y) = \frac{1}{y!}, \qquad c(\theta) = e^{\theta}.$$

- **Negative Binomial distribution** (with known size parameter r):

$$\theta = p \in (0,1), \qquad b(y) = \binom{y+r-1}{y}, \qquad c(\theta) = (1-\theta)^{-r}.$$

Our framework is general, but in this paper, we focus on the ZIP model.

In Zero-Inflated Power Series (ZIPS) regression models, covariates are linked to the model parameters, the zero-inflation parameter $\boldsymbol{p}$ and the underlying count distribution parameter $\boldsymbol{\theta}$, via link functions. Specifically, for an observation at time t, these relationships are defined as follows:

$$\log(\mu_t) = \boldsymbol{z}_t^T \boldsymbol{\beta}_{s_t} \quad \text{and} \quad \text{logit}(p_t) = \boldsymbol{w}_t^T \boldsymbol{\eta}_{s_t}.$$

Here, $\boldsymbol{z}_t$ and $\boldsymbol{w}_t$ are vectors of covariates, and $\boldsymbol{\beta}_{s_t}$ and $\boldsymbol{\eta}_{s_t}$ are the corresponding regression coefficient vectors for the segment to which observation t belongs (i.e., $s_t \in \{1,2\}$). The parameter $\mu_t = E(V_t)$ is the mean of the underlying count distribution V_t, which is a function of θ_t. For the Zero-Inflated Poisson special case, which is a primary focus of this paper, this relationship simplifies to $\mu_t = \theta_t$. The link functions are chosen to match the parameter space of the model. The log link for μ_t ensures that the mean of the count distribution remains positive, while the logit link for p_t transforms the linear predictor from the real line $(-\infty, \infty)$ to the probability space $(0,1)$.

Given a random sample $\boldsymbol{Y} = (Y_1, Y_2, \ldots, Y_n)$ from the ZIPS count process, the likelihood function is given by:

$$\mathcal{L}(\boldsymbol{\beta}_1, \boldsymbol{\beta}_2, \boldsymbol{\eta}_1, \boldsymbol{\eta}_2, m \mid \boldsymbol{Y}) \propto \prod_{t=1}^{n} \left\{ \left[p_t + (1-p_t)\frac{b(0)}{c(\theta_t)} \right]^{\mathbb{I}(Y_t=0)} \times \left[(1-p_t)\frac{b(Y_t)\theta_t^{Y_t}}{c(\theta_t)} \right]^{\mathbb{I}(Y_t>0)} \right\},$$

where $\mathbb{I}(\cdot)$ is the indicator function. The structure of this likelihood function is inherently complex. Its mixture nature, arising from the two potential sources of zero counts (from the Bernoulli component with probability p_t or the base count component with probability $(1-p_t)P(V_t = 0)$), combined with the non-linear link functions, makes direct inference from the posterior distribution of the parameters intractable.

To overcome this computational challenge, particularly within a Bayesian framework, we employ a standard and powerful approach known as data augmentation (Ghosh et al., 2006; Tanner and Wong, 1987). The core idea is to introduce latent variables that resolve the mixture structure of the model. By conditioning on these latent variables, the likelihood function decomposes into simpler, well-known forms, which facilitates efficient sampling from the full conditional distributions of the parameters using a Gibbs sampling algorithm. The specifics of this method are detailed in the subsequent section.

2.1 Bayesian Analysis

For $i = 1, 2$, let $\boldsymbol{\theta}_i = (\boldsymbol{\beta}_i, \boldsymbol{\eta}_i)$ and $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, m)$. We are interested in the posterior distribution $\pi(\boldsymbol{\theta} \mid \boldsymbol{Y})$, where $\boldsymbol{Y} = (Y_1, \ldots, Y_n)$ denotes n observations from the count process $\{Y_t\}$.

We assume independent prior distributions for the parameters:

$$p_t \sim \text{Beta}(b_1, b_2) \quad \text{and} \quad \theta_t \sim p(\theta),$$

with $p(\theta) \propto \frac{[\theta]^{a_1}}{[c(\theta)]^{a_2}}$ as a conjugate prior for the PS family. The hyper-parameters a_1, a_2, b_1, b_2 are assumed known, typically with $b_1 = b_2 = 1$ for a uniform prior on $(0, 1)$. In our model fitting (Section 4), we set the hyperparameters to $b_1 = b_2 = 0.5$ and $a_1 = a_2 = 0.001$. The choice of a Beta$(0.5, 0.5)$ prior for p_t corresponds to the Jeffreys prior for a binomial proportion, while setting a_1 and a_2 to small positive values (e.g., 0.001) is a common strategy to create a diffuse prior for the parameter of the Power Series distribution. These choices are intended to be weakly informative, allowing the data to dominate the posterior distribution (Gelman et al., 1995).

2.1.1 The Data Augmentation Method

To extend Gibbs sampling to regression problems, we augment the data Y_t with latent variables (V_t, B_t) using $Y_t = V_t(1 - B_t)$. Instead of sampling directly from the posterior of (p_t, θ_t) given Y_t, we sample from the posterior of $(p_t, \theta_t, V_t, B_t)$. The conditional distributions are:

$$P(V_t = v, B_t = 0 | Y_t = y) = \begin{cases} \frac{(1-p_t)P(V_t=0)}{p_t+(1-p_t)P(V_t=0)} & \text{if } v = y = 0; \\ 1 & \text{if } v = y > 0; \\ 0 & \text{otherwise} \end{cases}$$

$$P(V_t = v, B_t = 1 | Y_t = y) = \begin{cases} \frac{pP(V_t=v)}{p_t+(1-p_t)P(V_t=0)} & \text{if } y = 0; \\ 0 & \text{otherwise} \end{cases}$$

For each $t = 1, \ldots, n$, if $Y_t = 0$, sample a Bernoulli variable with probability p_t. If heads, set $B_t = 1$ and draw V_t from the distribution with parameter θ_t; if tails, set $B_t = 0$ and $V_t = 0$. If $Y_t > 0$, set $B_t = 0$ and $V_t = Y_t$. Then, sample p_t from Beta$\left(\sum_{t=1}^n B_t + b_1, \sum_{t=1}^n (1 - B_t) + b_2\right)$ using the completed data (V_t, B_t). Sample θ_{s_t} from a distribution in the exponential family, e.g., if $V_t \sim \text{Poisson}(\theta_t)$, then $\theta_t \sim \text{Gamma}(a_1 + \sum_{t=1}^n V_t, a_2 + n)$. This process iterates until convergence (Ghosh et al., 2006).

2.1.2 Fitting ZIPS Regression

In this article, we assume fixed covariates affecting $\boldsymbol{p}$ and $\boldsymbol{\mu}$. The parameters $\boldsymbol{\beta}_{s_t}$ and $\boldsymbol{\eta}_{s_t}$ are a priori independent, with normal priors:

$$\boldsymbol{\beta}_{s_t} \sim N_q(\boldsymbol{\beta}_0, \sigma_\beta^2 \boldsymbol{I}_q) \quad \text{and} \quad \boldsymbol{\eta}_{s_t} \sim N_r(\boldsymbol{\beta}_0, \sigma_\eta^2 \boldsymbol{I}_r), \quad s_t = 1, 2,$$

where $\boldsymbol{\beta}_0, \boldsymbol{\eta}_0, \sigma_\beta^2$, and σ_η^2 are known constants. We use $\boldsymbol{\beta}_0 = 0, \boldsymbol{\eta}_0 = \mathbf{0}, \sigma_\beta^2 = \sigma_\eta^2 = 10^3$ to express prior ignorance. The components of $\boldsymbol{\beta}_{s_t}$ and $\boldsymbol{\eta}_{s_t}$ are not independent a posteriori. When informative priors are available, the prior variance-covariance matrix can be adjusted accordingly.

These data augmentation methods are standard in modern computing, and we provide implementation details for the regression model when the latent variable follows a Poisson distribution. Modifications for other distributions in the PS family are straightforward.

2.1.3 Data Augmentation for ZIP Regression

In ZIP regression models, we have $c(\theta) = e^\theta$ and $b(k) = \frac{1}{k!}$, leading to the mean parameter $\mu = \mu(\theta) = \theta$. The regression model can be inverted to express the parameters as follows:

$$\theta_t = \exp({z_t}^T \boldsymbol{\beta}_{s_t}) \quad \text{and} \quad p_t = \left[1 + \exp(-{w_t}^T \boldsymbol{\eta}_{s_t})\right]^{-1}.$$

Similar to the methods proposed in Section 2.1.1, a data augmentation technique can be employed to generate samples from the joint posterior of $\boldsymbol{\beta}_{s_t}$ and $\boldsymbol{\eta}_{s_t}$. The data augmentation step is modified such that the probability of obtaining a HEAD is given by $1/\{1 + \exp\left(-\sum_{t=1}^n ({w_t}^T \boldsymbol{\eta}_{s_t})\right)\}$ and the mean parameter for the Poisson distribution becomes $\exp\left(\sum_{t=1}^n {z_t}^T \boldsymbol{\beta}_{s_t}\right)$.

The posterior distribution of $(\boldsymbol{\beta}_{s_t}, \boldsymbol{\eta}_{s_t})$ is obtained by multiplying the likelihood function (based on "complete data") and the prior, expressed as:

$$\begin{aligned}
&\exp\left\{\left(\textstyle\sum_{t=1}^n V_t z_t\right)^T \boldsymbol{\beta}_{s_t}\right\} \exp\left\{\left(\textstyle\sum_{t=1}^n B_t w_t\right)^T \boldsymbol{\eta}_{s_t}\right\} \\
&\times \textstyle\prod_{t=1}^n \exp\left(-\exp(z_t^T \boldsymbol{\beta}_{s_t})\right) (1 + \exp(w_t^T \boldsymbol{\eta}_{s_t}))^{-1} \\
\times (\sigma_{\boldsymbol{\beta}_{s_t}})^{-q} \exp&\left(-\tfrac{1}{2\sigma_\beta^2} (\boldsymbol{\beta}_{s_t} - \boldsymbol{\beta}_0)^T (\boldsymbol{\beta}_{s_t} - \boldsymbol{\beta}_0)\right) \times (\sigma_\eta^2)^{-r} \exp\left(-\tfrac{1}{2\sigma_\eta^2} (\boldsymbol{\eta}_{s_t} - \boldsymbol{\eta}_0)^T (\boldsymbol{\eta}_{s_t} - \boldsymbol{\eta}_0)\right).
\end{aligned}$$

To implement the modified Gibbs sampling algorithm, the full conditional densities are required. The full conditional distributions are given by:

$$\begin{aligned}
p[\boldsymbol{\beta}_{s_t} | \boldsymbol{\eta}_{s_t}, m, \boldsymbol{Y}] \quad &\propto \quad (\sigma_\beta^2)^{-q} \exp\left(-\frac{1}{2\sigma_\beta^2} (\boldsymbol{\beta}_{s_t} - \boldsymbol{\beta}_0)^T (\boldsymbol{\beta}_{s_t} - \boldsymbol{\beta}_0)\right) \\
&\times \quad \exp\left\{\left(\sum_{t=1}^n V_t z_t\right)^T \boldsymbol{\beta}_{s_t}\right\} Y_t \times \prod_{t=1}^n \exp\left(-\exp(z_t^T \boldsymbol{\beta}_{s_t})\right), \qquad (2.1)
\end{aligned}$$

and

$$p[\boldsymbol{\eta}_{s_t}|\boldsymbol{\beta}_{s_t}, m, \boldsymbol{Y}] \propto (\sigma_\eta^2)^{-r} \exp\left(-\frac{1}{2\sigma_\eta^2}(\boldsymbol{\eta}_{s_t} - \boldsymbol{\eta}_0)^T(\boldsymbol{\eta}_{s_t} - \boldsymbol{\eta}_0)\right) \times \exp\left\{\left(\sum_{t=1}^{n} B_t \boldsymbol{w}_t\right)^T \boldsymbol{\eta}_{s_t}\right\} Y_t \times \prod_{t=1}^{n}(1 + \exp(\boldsymbol{w}_t^T \boldsymbol{\eta}_{s_t}))^{-1}, \quad (2.2)$$

and

$$p[m|\boldsymbol{\beta}_{s_t}, \boldsymbol{\eta}_{s_t}, \boldsymbol{Y}] = \frac{p(\boldsymbol{\theta}_{-m}, m, \boldsymbol{Y}, \boldsymbol{z}, \boldsymbol{w})}{\sum_{j=1}^{n-1} p(\boldsymbol{\theta}_{-j}, j, \boldsymbol{Y}, \boldsymbol{z}, \boldsymbol{w})},$$

where

$$p(\boldsymbol{\theta}_{-m}, m, \boldsymbol{Y}, \boldsymbol{z}, \boldsymbol{w}) \propto \prod_{s_t=1}^{2} \left\{ \exp\left(\left(\sum_{t=1}^{n} V_t \boldsymbol{z}_t\right)^T \boldsymbol{\beta}_{s_t}\right) \exp\left(\left(\sum_{t=1}^{n} B_t \boldsymbol{w}_t\right)^T \boldsymbol{\eta}_{s_t}\right) \times \prod_{t=1}^{n} \exp\left(-\exp(\boldsymbol{z}_t^T \boldsymbol{\beta}_{s_t})\right)(1 + \exp(\boldsymbol{w}_t^T \boldsymbol{\eta}_{s_t}))^{-1} \right\}.$$

These densities are derived by retaining only the relevant terms from the posterior expression.

3 The Reversible Jump MCMC Algorithm

In developing the single change-point model, our objective goes beyond simply inferring the model parameters; we aim to incorporate a procedure within our Bayesian framework to test for the existence of a change-point in the data. In this section, we detail the RJMCMC algorithm designed for testing the presence of a single change-point. This approach can be extended to multiple change-points, as discussed in Appendix. The algorithm operates through two main updating steps: within-model updates and model-switching steps.

3.1 Within-Model Movement

This updating scheme involves sampling the parameters of the current model using standard Metropolis-Hastings updates. First, we perform a data augmentation step to sample (V_t, B_t) based on the current values of p_t, μ_t (or equivalently θ_t) and the data $\boldsymbol{Y}$. For $i = 1, 2$, the conditional distribution of $\boldsymbol{\beta}_i$ does not have a straightforward form. We can use different methods to update these parameters. One option is to update the parameters using a random walk Metropolis step, or to sample $\boldsymbol{\beta}_i$ using Adaptive Rejection Sampling (ARS) (see Gilks and Wild, 1992) given the sampled (V_i, B_i) and $\boldsymbol{\eta}_i$. However, we employ a different method to update $\boldsymbol{\beta}_i$. The pair of vectors $\boldsymbol{\beta}_1$ and $\boldsymbol{\beta}_2$ is

drawn from normal approximations to their posterior conditional distributions in (2.1). For example, $\boldsymbol{\beta}_i$ is drawn from $N(\boldsymbol{\beta}_{\max}, \Sigma_{\max})$ where

$$\boldsymbol{\beta}_{max} = arg_{\boldsymbol{\beta}_{s_t}} \max p[\boldsymbol{\beta}_{s_t}|\boldsymbol{\eta}_{s_t}, m, \boldsymbol{Y}],$$

and

$$\Sigma_{max} = \left\{ -\frac{\partial^2 \log p[\boldsymbol{\beta}_{s_t}|\boldsymbol{\eta}_{s_t}, m, \boldsymbol{Y}]}{\partial \boldsymbol{\beta}_{s_t} \partial \boldsymbol{\beta}'_{s_t}} |_{\boldsymbol{\beta}_{s_t}=\boldsymbol{\beta}^{max}} \right\}^{-1}.$$

For $i = 1, 2$, we apply the same procedure for updating $\boldsymbol{\eta}_i$. When updating m, we found that a hybrid proposal scheme performed well. We propose the following hybrid proposal scheme for updating the change-point m:

1. Proposal Type Determination:

$$r \sim \text{Uniform}(0, 1)$$

If $r < p$ (where $p = 0.9$), we perform a local move; otherwise, we perform a global move.

2. Local Move Proposal: For a local move, we propose to shift the change-point to:

$$m' = m + V \cdot Q$$

where:

- $V \sim \text{Categorical}(\{-1, 0, 1\}; (0.33, 0.34, 0.33))$
- $Q \sim \text{Binomial}(n - 1, q)$ with $q = 0.5$

Ensure that m' remains within the bounds $1 \leq m' \leq n - 1$. If m' is out of bounds, adjust or reject the proposal.

3. Global Move Proposal: For a global move, we propose:

$$m' \sim \text{Uniform}[1, n - 1].$$

In this context, if we denote the proposed values with a prime symbol ($'$), the corresponding acceptance probability for this draw is given by min $\min\{1, A\}$, where

$$A = \frac{\mathcal{L}(\boldsymbol{Y}, \boldsymbol{w}, \boldsymbol{z}|\boldsymbol{\theta}'_1, \mathcal{M}', m')}{\mathcal{L}(\boldsymbol{Y}, \boldsymbol{w}, \boldsymbol{z}|\boldsymbol{\theta}_1, \mathcal{M}, m)} \times \frac{p(\mathcal{M}'|\boldsymbol{\theta}'_1)}{p(\mathcal{M}|\boldsymbol{\theta}_1)} \times \frac{q(\boldsymbol{\theta}'_1|\boldsymbol{\theta}_1)}{q(\boldsymbol{\theta}_1|\boldsymbol{\theta}'_1)},$$

where $\frac{\mathcal{L}(\boldsymbol{Y}, \boldsymbol{w}, \boldsymbol{z}|\boldsymbol{\theta}'_1, \mathcal{M}', m')}{\mathcal{L}(\boldsymbol{Y}, \boldsymbol{w}, \boldsymbol{z}|\boldsymbol{\theta}_1, \mathcal{M}, m)}$ denotes the likelihood ratio for $\mathcal{M}'$ and $\mathcal{M}$, $\frac{p(\mathcal{M}'|\boldsymbol{\theta}'_1)}{p(\mathcal{M}|\boldsymbol{\theta}_1)}$ represents the prior ratio, and $\frac{q(\boldsymbol{\theta}'_1|\boldsymbol{\theta}_1)}{q(\boldsymbol{\theta}_1|\boldsymbol{\theta}'_1)}$ is the ratio of the proposal distributions. We can also express the acceptance probability for this move as follows:

$$\begin{aligned} A &= \frac{\mathcal{L}(\boldsymbol{Y}, \boldsymbol{w}, \boldsymbol{z}|\boldsymbol{\beta}'_1, \boldsymbol{\beta}'_2, \boldsymbol{\eta}'_1, \boldsymbol{\eta}'_2, \mathcal{M}', m')}{\mathcal{L}(\boldsymbol{Y}, \boldsymbol{w}, \boldsymbol{z}|\boldsymbol{\beta}_1, \boldsymbol{\beta}_2, \boldsymbol{\eta}_1, \boldsymbol{\eta}_2, \mathcal{M}, m)} \times \frac{p(\boldsymbol{\beta}'_1, \boldsymbol{\beta}'_2 \mid m')p(\boldsymbol{\eta}'_1, \boldsymbol{\eta}'_2 \mid m')}{p(\boldsymbol{\beta}_1, \boldsymbol{\beta}_2 \mid m)p(\boldsymbol{\eta}_1, \boldsymbol{\eta}_2 \mid m)} \\ &\times \frac{q(\boldsymbol{\beta}_1 \mid \boldsymbol{\mathcal{Y}}_1, m)q(\boldsymbol{\beta}_2 \mid \boldsymbol{\mathcal{Y}}_2, m)}{q(\boldsymbol{\beta}'_1 \mid \boldsymbol{\mathcal{Y}}'_1, m', \boldsymbol{\beta}_1)q(\boldsymbol{\beta}'_2 \mid \boldsymbol{\mathcal{Y}}'_2, m', \boldsymbol{\beta}_2)} \times \frac{q(\boldsymbol{\eta}_1 \mid \boldsymbol{\mathcal{Y}}_1, m)q(\boldsymbol{\eta}_2 \mid \boldsymbol{\mathcal{Y}}_2, m)}{q(\boldsymbol{\eta}'_1 \mid \boldsymbol{\mathcal{Y}}'_1, m', \boldsymbol{\eta}_1)q(\boldsymbol{\eta}'_2 \mid \boldsymbol{\mathcal{Y}}'_2, m', \boldsymbol{\eta}_2)}. \end{aligned}$$

In summary, the following algorithm can be employed to generate samples for this movement, specifically for identifying the location of the change-point and the corresponding parameters in each segment.

- Start with dispersed initial values of m, $\boldsymbol{\beta}_{s_t}$, and $\boldsymbol{\eta}_{s_t}$.
- Data augmentation step: Sample (V_t, B_t) based on the current values of p_t, μ_t (or equivalently θ_t), and the data $\boldsymbol{Y}$.
- Update m using the proposed hybrid proposal scheme.
- Sample $\boldsymbol{\beta}_{s_t}$ using normal approximations to their posterior conditional distributions in (2.1), given the sampled (V_i, B_i) and $\boldsymbol{\eta}_{s_t}$.
- Sample $\boldsymbol{\eta}_{s_t}$ using normal approximations to their posterior conditional distributions in (2.2), given the sampled (V_t, B_t) and $\boldsymbol{\beta}_{s_t}$.
- Move to the between-model step (3.2) with the current sampled values of parameters $(m, \boldsymbol{\beta}_{s_t}, \boldsymbol{\eta}_{s_t})$ if the move is accepted. Otherwise, return to the data augmentation step to update (V_t, B_t) if the move is rejected, and repeat this process until acceptance.

3.2 Between-Model Movement

This updating step of the RJMCMC sampler involves proposing transitions between competing models. Parameter estimation for model $\mathcal{M}_0$ can be performed using the procedures described above, excluding the steps for updating $\boldsymbol{\theta}_2$ and m. Since model $\mathcal{M}_{00}$ has more parameters than model $\mathcal{M}_0$, it is necessary to introduce auxiliary parameters to transition from model $\mathcal{M}_0$ to model $\mathcal{M}_{00}$ [see Green, 1995]. The auxiliary parameters are $\tilde{b}_i$ (for $i = 1, \ldots, q$) and $\tilde{d}_i$ (for $i = 1, \ldots, r$), drawn from $N(0, \sigma_b^2)$ and $N(0, \sigma_d^2)$, respectively. The variances σ_b^2 and σ_d^2 are typically pilot-tuned for efficient moves. For m, we choose a position for the proposed change-point uniformly at random from the interval, $m \sim \text{DU}[1, n-1]$. The $\tilde{b}_i$ and $\tilde{d}_i$ compensate for the absence of the second set of regression coefficients in $\mathcal{M}_0$. We then propose to move to model $\mathcal{M}_{00}$ with model parameters $\boldsymbol{\theta}$ and a change-point at $m = k$:

$$
\begin{aligned}
(\boldsymbol{\beta}_1', \boldsymbol{\eta}_1') &= (\boldsymbol{\beta}_1, \boldsymbol{\eta}_1) + \left(1 - \frac{k}{n}\right)(\tilde{\boldsymbol{b}}, \tilde{\boldsymbol{d}}), \\
(\boldsymbol{\beta}_2', \boldsymbol{\eta}_2') &= (\boldsymbol{\beta}_1, \boldsymbol{\eta}_1) - \left(\frac{k}{n}\right)(\tilde{\boldsymbol{b}}, \tilde{\boldsymbol{d}}).
\end{aligned} \tag{3.1}
$$

For the reverse move, no auxiliary parameters are needed. Consequently, we propose the parameters of model $\mathcal{M}_0$ as the weighted means of the corresponding parameters of model $\mathcal{M}_{00}$:

$$
\boldsymbol{\beta}_1' = \frac{m}{n}\boldsymbol{\beta}_1 + \left(1 - \frac{m}{n}\right)\boldsymbol{\beta}_2, \quad \boldsymbol{\eta}_1' = \frac{m}{n}\boldsymbol{\eta}_1 + \left(1 - \frac{m}{n}\right)\boldsymbol{\eta}_2. \tag{3.2}
$$

This proposal scheme is based on the concept of moment matching [see Green, 1995; Brooks et al., 2003] and has proven effective in our context.

The acceptance probability of the algorithm for this step is (Green, 1995)

$$\min\Big\{1, (\text{ likelihood ratio}) \times (\text{prior ratio}) \times (\text{proposal ratio}) \times (\text{Jacobian})\Big\}.$$

At each iteration of the algorithm, when proposing to switch models, the acceptance probability for moving from $\mathcal{M}_0$ to $\mathcal{M}_{00}$ is given by:

$$\alpha_{0\to 00} = \min\left(1, A_{0\to 00}(\theta_1, \theta)\right).$$

The acceptance ratio $A_{0\to 00}(\boldsymbol{\theta}_1, \boldsymbol{\theta})$ is defined as:

$$A_{0\to 00}(\boldsymbol{\theta}_1, \boldsymbol{\theta}) = \frac{\mathcal{L}(\boldsymbol{Y}, \boldsymbol{w}, \boldsymbol{z}|\boldsymbol{\theta}, \mathcal{M}_{00})}{\mathcal{L}(\boldsymbol{Y}, \boldsymbol{w}, \boldsymbol{z}|\boldsymbol{\theta}_1, \mathcal{M}_0)} \times \frac{p(\mathcal{M}_{00}|\boldsymbol{\theta})}{p(\mathcal{M}_0|\boldsymbol{\theta}_1)} \times \frac{q(\boldsymbol{\theta}|\boldsymbol{\theta}_1)}{q(\boldsymbol{\theta}_1|\boldsymbol{\theta})} \times |J_{0\to 00}|, \tag{3.3}$$

where $\frac{\mathcal{L}(\boldsymbol{Y}, \boldsymbol{w}, \boldsymbol{z}|\boldsymbol{\theta}, \mathcal{M}_{00})}{\mathcal{L}(\boldsymbol{Y}, \boldsymbol{w}, \boldsymbol{z}|\boldsymbol{\theta}_1, \mathcal{M}_0)}$ denotes the likelihood ratio for $\mathcal{M}_{00}$ and $\mathcal{M}_0$, $\frac{p(\mathcal{M}_{00}|\boldsymbol{\theta})}{p(\mathcal{M}_0|\boldsymbol{\theta}_1)}$ represents the prior ratio, and $\frac{q(\boldsymbol{\theta}|\boldsymbol{\theta}_1)}{q(\boldsymbol{\theta}_1|\boldsymbol{\theta})}$ is the ratio of the proposal distributions for the transition between the models. Note that $q(\boldsymbol{\theta}_1|\boldsymbol{\theta}) = 1$ since the move from $\mathcal{M}_{00}$ to $\mathcal{M}_0$ is deterministic. Finally, $|J_{0\to 00}|$ is the Jacobian for the transformation between the parameters of the two models. Due to the block diagonal nature of $J_{0\to 00}$, it can be shown that $|J_{0\to 00}| = 1$. Note that, although we are calculating the current or proposed values in the prior ratio in (3.3), we have

$$p(\mathcal{M}, m, \boldsymbol{\theta}) = p(\mathcal{M}, m, \boldsymbol{\beta}, \boldsymbol{\eta}) = p(m \mid \mathcal{M}) p(\boldsymbol{\beta}, \boldsymbol{\eta} \mid \mathcal{M}, m) p(\boldsymbol{\eta} \mid \mathcal{M}, m, \boldsymbol{\beta}).$$

The denominator in the proposal ratio refers to the conditional density of the proposed number of segments and parameters given the current state as follows:

$$\begin{aligned} q(\mathcal{M}_{00}, \boldsymbol{\theta} \mid \mathcal{M}_0, \boldsymbol{\theta}_1) &= q(\mathcal{M}_{00} \mid \mathcal{M}_0) q(\boldsymbol{\theta} \mid \mathcal{M}_0, \mathcal{M}_{00}, \boldsymbol{\theta}_1) \\ &= q(\mathcal{M}_{00} \mid \mathcal{M}_0) q(m', \boldsymbol{\beta}, \boldsymbol{\eta} \mid \mathcal{M}_0, \mathcal{M}_{00}, \boldsymbol{\theta}_1) \\ &= q(\mathcal{M}_{00} \mid \mathcal{M}_0) q(m' \mid \mathcal{M}_0, \mathcal{M}_{00}, \boldsymbol{\theta}_1) \\ &\times \; q(\boldsymbol{\beta}, \boldsymbol{\eta} \mid m', \mathcal{M}_0, \mathcal{M}_{00}, \boldsymbol{\theta}_1) q(\boldsymbol{\beta} \mid \boldsymbol{\eta}, m', \mathcal{M}_0, \mathcal{M}_{00}, \boldsymbol{\theta}_1). \end{aligned}$$

The acceptance ratio for the reverse move (i.e., $\mathcal{M}_{00} \to \mathcal{M}_0$) is simply obtained using the inverse of all the terms in equation (3.3).

In summary, the following algorithm can be used to generate samples for the between-model movement, specifically for selecting the model (whether or not there is a change-point), as well as for determining the location of the change-point and the corresponding parameters in each segment

- Start with dispersed initial values of m, $\boldsymbol{\beta}_{s_t}$, and $\boldsymbol{\eta}_{s_t}$.
- Data augmentation step: Sample (V_t, B_t) based on the current values of p_t, μ_t (or equivalently θ_t), and the data $\boldsymbol{Y}$.

- Decide between $\mathcal{M}_0$ and $\mathcal{M}_{00}$.
- Sample m using $m \sim \text{DU}[1, n-1]$.
- If $\mathcal{M}_{00}$ is selected, propose $\boldsymbol{\beta}_1$, $\boldsymbol{\beta}_2$, $\boldsymbol{\eta}_1$, and $\boldsymbol{\eta}_2$ using (3.1).
- If $\mathcal{M}_0$ is selected, propose $\boldsymbol{\beta}_1$ and $\boldsymbol{\eta}_1$ using (3.2).
- Move to the within-model step (3.1) with the current sampled values of parameters $(m, \boldsymbol{\beta}, \boldsymbol{\eta})$ if the move is accepted. Otherwise, return to the data augmentation step to update (V_t, B_t) if the move is rejected, and continue this process until acceptance.

3.3 The Complete Computational Algorithm

To provide a holistic view of the inferential procedure, we now synthesize the within-model updates (Section 3.1) and between-model jumps (Section 3.2) into a single, comprehensive RJMCMC algorithm. At each iteration, the sampler decides whether to explore the parameter space of the current model or to attempt a jump to a model with a different number of segments (i.e., from one segment to two, or vice versa). Let $\boldsymbol{\Theta}_0 = (\boldsymbol{\beta}_1, \boldsymbol{\eta}_1)$ be the parameter set for the one-segment model ($\mathcal{M}_0$) and $\boldsymbol{\Theta}_{00} = (\boldsymbol{\beta}_1, \boldsymbol{\beta}_2, \boldsymbol{\eta}_1, \boldsymbol{\eta}_2, m)$ be the parameter set for the two-segment model ($\mathcal{M}_{00}$).

The complete algorithm for a single iteration of the RJMCMC sampler is detailed in Algorithm 1.

After running the sampler for a sufficient number of iterations, the initial burn-in period is discarded. The posterior distribution of the change-point location m can be estimated from the iterations where the two-segment model, $\mathcal{M}_{00}$, was active. The posterior probability of a change-point existing is simply the proportion of post-burn-in iterations where the sampler was in state $\mathcal{M}_{00}$. Similarly, posterior summaries for the segment-specific parameters $(\boldsymbol{\beta}_1, \boldsymbol{\eta}_1)$ and $(\boldsymbol{\beta}_2, \boldsymbol{\eta}_2)$ are computed using only the samples from the $\mathcal{M}_{00}$ state.

3.4 Proof of the Jacobian Determinant

In this section, we provide the detailed derivation of the Jacobian for the dimension-matching transformation used in the reversible jump MCMC algorithm, specifically for the move from model $\mathcal{M}_0$ to $\mathcal{M}_{00}$.

Let $\boldsymbol{\theta}$ denote the vector of parameters in model $\mathcal{M}_0$ (e.g., regression coefficients $\boldsymbol{\beta}$ or zero-inflation parameters $\boldsymbol{\eta}$). To split this vector into two new vectors $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$ for model $\mathcal{M}_{00}$ at a change-point k, we introduce an auxiliary vector $\mathbf{u}$ of the same dimension as $\boldsymbol{\theta}$.

Let $w = k/n$ be the proportion of data points before the change-point. The transformation defined in (3.1) can be written element-wise as:

$$\begin{aligned} \theta_1' &= \theta + (1-w)u \\ \theta_2' &= \theta - wu \end{aligned}$$

The Jacobian matrix, J, for this linear transformation with respect to (θ, u) is a block-diagonal matrix. For a single parameter component, the corresponding block takes the form:

$$B = \begin{pmatrix} \frac{\partial \theta_1'}{\partial \theta} & \frac{\partial \theta_1'}{\partial u} \\ \frac{\partial \theta_2'}{\partial \theta} & \frac{\partial \theta_2'}{\partial u} \end{pmatrix} = \begin{pmatrix} 1 & 1-w \\ 1 & -w \end{pmatrix} = \begin{pmatrix} 1 & 1-\frac{k}{n} \\ 1 & -\frac{k}{n} \end{pmatrix}.$$

We compute the determinant of this generic block B:

$$\begin{aligned} \det(B) &= (1)\left(-\frac{k}{n}\right) - (1)\left(1-\frac{k}{n}\right) \\ &= -\frac{k}{n} - 1 + \frac{k}{n} \\ &= -1. \end{aligned}$$

Since the transformation is applied element-wise to the parameter vectors, the full Jacobian determinant is the product of the determinants of these blocks. If the total dimension of the auxiliary parameters is d, then $\det(J_{0\to 00}) = (-1)^d$. Consequently, the absolute value of the Jacobian determinant is:

$$|J_{0\to 00}| = |-1|^d = 1.$$

For the reverse move (death move from $\mathcal{M}_{00}$ to $\mathcal{M}_0$), which involves a weighted mean, the Jacobian is simply the inverse:

$$|J_{00\to 0}| = |J_{0\to 00}|^{-1} = 1.$$

4 Simulation Study

To assess the performance of the proposed change-point model and the accompanying RJMCMC algorithm, we conducted a simulation study aimed at evaluating the model's ability to accurately detect and localize a change-point.

The simulation study was conducted as follows. We constructed a covariate vector:

$$\boldsymbol{w}^T = \boldsymbol{z}^T = (1, N(0,1), U(0,1)),$$

and generated two datasets, each containing $n = 200$ observations. For scenarios involving a change-point, its location was fixed at $m = 100$. For the two datasets, the parameter vector was defined as follows:

$$\begin{aligned} \boldsymbol{\theta}_1 &= \left(\boldsymbol{\beta}_1^T \in \mathbb{R}^{1\times 3}, \boldsymbol{\eta}_1^T \in \mathbb{R}^{1\times 3}\right) \\ &= (0.50, 0.80, -1.0, -10.0, 0.10, 0.04), \end{aligned} \tag{4.1}$$

and for the first dataset, we define $\boldsymbol{\theta}_2 = \boldsymbol{\theta}_1$. For the second dataset, the parameters of the second segment were defined as:

$$\begin{aligned} \boldsymbol{\theta}_2 &= \left(\boldsymbol{\beta}_2^T \in \mathbb{R}^{1\times 3}, \boldsymbol{\eta}_2^T \in \mathbb{R}^{1\times 3}\right) \\ &= (-0.50, -0.80, 2.0, -2.0, 3.0, 6.0). \end{aligned} \tag{4.2}$$

Thus, in Dataset (4.1), there is effectively no change-point, whereas in Dataset (4.2), a pronounced change-point is present in the parameters at $m = 100$. Using 200 observations, this simulation setting was replicated 50 times. The focus of this study is on estimating the potential change-point and its location using summaries of the posterior samples. The algorithm generates 50,000 pseudo-random numbers, with the first 10,000 values excluded from all parameter estimates as a burn-in period. Visual examination of the trace plots for the parameters indicated that convergence was achieved before the end of the burn-in period, with most parameters being well estimated in nearly all cases.

Figure 1 depicts a realization of the model without any change-point, showcasing a dataset consisting of 200 observations, which serves as a baseline for understanding how the model behaves in the absence of structural changes in the data. In contrast, Figure 2 illustrates a realization of the model with a single change-point, where the first segment comprises 100 observations followed by another 100 observations in the second segment; the distinct colors used in the plot effectively highlight the differences between the two segments, emphasizing the impact of the change-point on the data structure. Figure 3 presents histograms of the count data for both segments of the model with a single change-point, with Panel (a) depicting the histogram for the first segment and Panel (b) showcasing the histogram for the second segment; these histograms reveal notable differences in the distribution of the data, particularly in the prevalence of zeros and non-zero values across the two segments, highlighting the effect of the change-point on the underlying data distribution. Figure 4 illustrates the posterior probabilities for the model without a change-point across 50 replications. It should be noted that we employ maximum a posteriori (MAP) estimator to determine the number of segments - either one or two segments- and use the mean of the empirical posterior density to estimate the location of the change-point. The MAP estimator indicates that the model consistently identifies the absence of a change-point, demonstrating its robustness in scenarios where no structural changes are present. This suggests that the method exhibits no false positive performance in models without a change-point. In Figure 5, the posterior probabilities for the model with a single change-point are presented. Out of 50 replications, the method successfully detects the single change-point in 47 instances based on MAP estimators, showcasing the model's effectiveness in identifying structural changes in the data. Finally, Figure 6 presents the empirical posterior density for the location of the change-point in the model with a single change-point, where the dashed line indicates the estimated change-point based on the mean of the posterior density, calculated to be 97, while the true change-point is located at 100; notably, the mode of the posterior density suggests a biased estimation, indicating a mode at 94, which highlights the challenges in accurately pinpointing the location of change-points in practice. The estimated parameters in each segment are presented in Table 1. We utilize the modified Bayesian Information Criterion (mBIC) defined by Zhang and Siegmund (2007) for model comparison. The average mBIC calculated from the estimates obtained based on the model with a single change-point is -313.783. When the possibility of a change-point is not considered in this model, the average mBIC increases to -241.588. This increase further substantiates the accuracy of the

model that incorporates a change-point, despite having more than twice the number of parameters compared to the model without a change-point. For the model without a change-point, the average mBIC for the case of no change-point and for the case with a single change-point are -293.031 and -245.25, respectively. This again demonstrates the accuracy of the model based on the mBIC.

To further demonstrate the necessity and superior performance of our proposed ZIP-RJMCMC framework, we conducted a small empirical comparison against a standard Poisson change-point regression model. The standard model natively ignores zero-inflation (i.e., assuming $p_t = 0$ for all t) but uses the same RJMCMC algorithm to detect structural breaks. We applied this standard model to the 50 replications of Dataset (4.2), which contains both a true change-point and structural zeros. The results clearly highlight the limitations of the standard frequentist or basic Bayesian approaches when dealing with complex count data.

When zero-inflation is ignored, the standard Poisson model struggles significantly to accurately detect and localize the change-point. Out of the 50 replications, the standard model successfully identified a single change-point in only 29 instances (a 58% success rate, compared to 94% for our ZIP model). Furthermore, in the instances where a change-point was detected, the location estimation was highly erratic and biased. The mean estimated location was shifted to $\hat{m} \approx 81$ (true $m = 100$), accompanied by a substantially wider posterior variance. This occurs because the standard model attempts to inappropriately absorb the structural zeros into the base Poisson mean, thereby confusing the true structural break with random fluctuations in zero counts.

In terms of model fit, the average mBIC for the standard Poisson change-point model on Dataset (4.2) was -178.450. This is significantly worse than the average mBIC of -313.783 achieved by our proposed ZIP change-point model. This empirical comparison provides compelling evidence that failing to account for zero-inflation in count time series not only degrades the overall goodness-of-fit but also severely compromises the reliability and accuracy of structural break detection.

Visual examination of the trace plots indicated that convergence was achieved well before the end of the burn-in period. For instance, Figure 7 shows the trace and posterior density for the β_0 parameter. The trace plot is stable, and the posterior distribution is well-defined around the true value.

In summary, the simulation study demonstrates the model's capabilities in identifying change-point and provides valuable insights into the underlying data structure, illustrating both the effectiveness and limitations of the proposed approach.

5 Application: Detecting Structural Breaks in the COVID-19 Pandemic

5.1 Rationale and Background

The COVID-19 pandemic, caused by the SARS-CoV-2 virus, represents a textbook example of a non-stationary dynamic process. Since its declaration as a global pandemic

by the World Health Organization (WHO) on March 11, 2020, its trajectory has been defined by a series of distinct phases driven by viral evolution, public health interventions, and the rollout of vaccination campaigns. The clinical presentation, ranging from asymptomatic infection to severe respiratory illness and death, has also shifted in response to these factors. This dynamic nature makes time-series data from the pandemic, particularly zero-inflated count data like daily deaths, an ideal testbed for change-point detection methodologies. Segmentation of this data is not merely a statistical exercise; it allows researchers to pinpoint critical moments where pandemic dynamics altered, thereby providing invaluable insights for evaluating policy effectiveness and informing future public health preparedness.

5.2 Case Study Design: Cyprus and Qatar

To demonstrate the real-world utility of our RJMCMC-ZIP framework, we analyze COVID-19 mortality data from two nations with differing timelines and epidemiological contexts: **Cyprus** and **Qatar**.

- **Cyprus:** The dataset covers a 117-day period from February 13, 2022, to June 10, 2022. This timeframe was largely dominated by the global surge of the highly transmissible **Omicron variant**. Official reports from the Cyprus Ministry of Health during this period highlighted concerns around case surges, balanced by a mature vaccination program.

- **Qatar:** The dataset spans a 195-day period from March 12, 2021, to September 23, 2021. This earlier phase of the pandemic was characterized by the prevalence of variants like **Delta** and, critically, the implementation of Qatar's aggressive **mass vaccination campaign**, as documented by the Qatar Ministry of Public Health.

For both countries, the daily number of deaths served as our zero-inflated response variable. We incorporated two crucial covariates: the logarithm of cumulative confirmed cases (a proxy for disease prevalence) and the logarithm of cumulative vaccination doses administered (a proxy for population immunity). A visual synopsis of these time series is presented in Figures 8 (Deaths), 9 (Cases), and 10 (Vaccinations), which illustrate the distinct patterns in each country.

5.3 Model Results and Epidemiological Interpretation

Our model was applied to each dataset to identify the most probable location of a single structural break and to estimate the regression coefficients for the segments before and after this point.

5.3.1 Analysis of Cyprus (The Omicron Wave)

For Cyprus, our model detected a significant change-point, clearly visualized by the vertical dashed line in the right panel of Figure 11. An examination of the estimated

parameters in **Table 2** reveals a fascinating story about the evolving dynamics of the Omicron wave:

- **Before the change-point (Segment 1):** An increase in confirmed cases was strongly associated with higher expected deaths ($\hat{\beta}_1 = 2.087$), while vaccination showed a significant protective association ($\hat{\beta}_2 = -1.101$).
- **After the change-point (Segment 2):** The magnitude of both effects intensified dramatically. The association between cases and deaths more than doubled ($\hat{\beta}_1 = 4.943$), and the protective association of vaccination nearly quadrupled ($\hat{\beta}_2 = -4.252$).

Interpretation: This "amplification of effects" after the change-point suggests that as the Omicron wave matured, the healthcare system may have been under increasing strain, making each new case statistically more impactful on mortality. Concurrently, the powerfully enhanced protective association of vaccination (β_2) strongly indicates that vaccines and boosters were exceptionally effective at preventing death, even amidst massive case numbers, effectively "decoupling" infections from severe outcomes for the vaccinated population.

5.3.2 Analysis of Qatar (The Mass Vaccination Campaign)

In Qatar, despite a consistently low number of daily deaths as shown in Figure 8 (left panel), our model identified a distinct structural break (Figure 11, left panel). This change-point appears to mark not a crisis, but a turning point of success. The parameter estimates in **Table 2** support this narrative:

- The positive link between confirmed cases and deaths remained strong and stable across both segments ($\hat{\beta}_1 = 3.942$ and 4.674), reflecting the persistent threat of the Delta variant.
- The most striking result is the evolution of the vaccination coefficient (β_2). Its protective association, already very strong in the first segment ($\hat{\beta}_2 = -7.348$), became profoundly more powerful in the second segment ($\hat{\beta}_2 = -12.347$).

Interpretation: This substantial increase in the magnitude of the vaccination coefficient provides compelling statistical evidence for the success of Qatar's national vaccination strategy. The change-point likely marks the moment the campaign reached a critical immunity threshold, fundamentally strengthening the population's resilience and severely blunting the mortality impact of ongoing viral transmission.

5.4 Critical Discussion: Association vs. Causation and Confounding Factors

It is imperative to interpret these compelling findings with appropriate scientific caution. Our model identifies and quantifies **strong statistical associations**, not definitive

causal relationships. While the negative coefficient for vaccinations ($\beta_2 < 0$) aligns perfectly with public health expectations and biological plausibility, attributing the observed mortality trends solely to the included covariates would be an oversimplification. The real-world dynamics of a pandemic are influenced by a complex web of **confounding variables**, which were not included in our model. Acknowledging these is crucial for a nuanced interpretation. Key potential confounders include:

- **Viral Variant Dynamics:** The intrinsic severity of the dominant variant (e.g., Delta vs. Omicron) changed over time.
- **Improvements in Clinical Care:** Advances in treatment protocols and hospital management likely reduced case fatality rates independently.
- **Public Health and Social Measures (PHSMs):** The implementation or relaxation of measures like lockdowns and mask mandates, as noted by the WHO, directly impact transmission.
- **Changes in Population Behavior and Immunity:** Natural immunity from prior infections and shifts in public adherence to guidelines also play a significant role.
- **Shifts in Testing Strategies:** Changes in testing availability can alter the number of confirmed cases, thus affecting estimated coefficients.

Therefore, our analysis should be viewed as providing strong, data-driven hypotheses that highlight critical shifts in pandemic trends, rather than final causal proof. The identified change-points offer valuable insights into the *effectiveness* of the overall combination of strategies implemented in both countries.

6 Discussion

In this paper, we have introduced a robust Bayesian framework for detecting and estimating a single change-point in ZIP regression models. Our primary contribution lies in the successful adaptation and implementation of a Reversible Jump Markov Chain Monte Carlo (RJMCMC) algorithm to this class of models, which are ubiquitous in fields like epidemiology, ecology, and economics where count data exhibit an excess of zeros. The proposed methodology demonstrated strong performance, accurately identifying the absence of a change-point when none existed and precisely estimating its location and the corresponding parameter shifts in both simulated and real-world datasets.

The choice of a sophisticated Bayesian approach over more traditional frequentist methods, such as Cumulative Sum (CUSUM) control charts or likelihood ratio tests, was deliberate. Classical methods often rely on assumptions of normality or are not directly tailored for the unique challenges of zero-inflated and overdispersed count data. Ignoring the zero-inflation component can lead to biased parameter estimates and inaccurate change-point detection. Our framework, by contrast, natively models

the dual-state process (zero vs. count generation) inherent in ZIP data. A key advantage of our RJMCMC approach is that it provides a full posterior distribution for the change-point location (τ), offering a natural quantification of uncertainty that is not readily available from point-estimate methods. This probabilistic inference is crucial for interpreting the timing of a structural break with appropriate confidence.

The most significant limitation of our current model is its restriction to a single change-point. This was a strategic choice to first establish a solid theoretical and computational foundation for applying RJMCMC to ZIP models. However, we recognize that real-world time series, such as the COVID-19 pandemic data analyzed in Section 4, often exhibit multiple structural breaks driven by sequential events like new viral variants, policy interventions, and vaccination campaigns. A single change-point model may, therefore, capture only the most dominant shift in the data, potentially oversimplifying the underlying dynamics.

A natural and crucial extension is the development of a model that can estimate an unknown number of change-points. This can be achieved by expanding our RJMCMC algorithm to include "birth" and "death" moves for change-points, allowing the sampler to traverse models of varying dimensions, as originally proposed by Green (1995). Such an extension could be further enhanced by employing hierarchical priors on the segment-specific parameters, a technique that allows the model to "borrow strength" across segments and can improve estimation, particularly for shorter segments. The framework developed in this paper serves as an essential building block for these more complex, multi-change-point analyses.

Beyond the model's scope, practical considerations introduce other limitations. When segments between change-points are short (i.e., in data-sparse scenarios), the model faces an increased risk of overfitting, where it may capture random noise rather than a true shift in parameters. In such cases, posterior estimates can also become highly sensitive to the choice of prior distributions, as the likelihood provides less information to dominate the prior. While we used weakly informative priors to allow the data to drive the inference, we strongly recommend that future applications include a sensitivity analysis to assess the robustness of the findings to different prior specifications. Furthermore, while powerful, the RJMCMC algorithm is computationally more intensive than its frequentist counterparts. Users must ensure that convergence has been properly assessed using diagnostic tools, such as the trace plots provided in our simulation study, before drawing substantive conclusions.

In conclusion, this paper provides a novel and necessary tool for researchers analyzing complex count data characterized by structural breaks and zero-inflation. By integrating the RJMCMC algorithm with ZIP regression models, we offer a principled and flexible approach that honors the specific data-generating process, quantifies uncertainty robustly, and provides interpretable results. This work opens the door for more nuanced and accurate modeling of dynamic processes in fields ranging from public health surveillance to financial econometrics, laying the groundwork for future extensions to handle even more complex real-world phenomena.

Algorithm 1 A Modular RJMCMC for ZIP Change-Point Detection

Input: Initial state $S^{(0)} = (k^{(0)}, \boldsymbol{\Theta}_k^{(0)})$, total iterations N, jump probability p_{jump}.
Output: A sequence of posterior samples $\{S^{(i)}\}_{i=1}^{N}$.

procedure UpdateWithinModel($k, \boldsymbol{\Theta}_k, \boldsymbol{V}, \boldsymbol{B}$) ▷ Updates parameters for a fixed model k
  if $k = 2$ **then** ▷ Two-segment model $\mathcal{M}_2$
    Update Change-point m**:** Propose $m' \sim q(m'|m)$ via the hybrid scheme.
    Accept m' with Metropolis-Hastings probability

$$\alpha_m = \min\left(1, \frac{\pi(\boldsymbol{Y}|\boldsymbol{\Theta}_2{}', \boldsymbol{V}, \boldsymbol{B})\, p(\boldsymbol{\Theta}_2{}')\, q(m|m')}{\pi(\boldsymbol{Y}|\boldsymbol{\Theta}_2, \boldsymbol{V}, \boldsymbol{B})\, p(\boldsymbol{\Theta}_2)\, q(m'|m)}\right).$$

    Update Coefficients: For $s \in \{1, 2\}$, sample $\boldsymbol{\beta}_s' \sim N(\hat{\boldsymbol{\beta}}_s, \boldsymbol{\Sigma}_{\beta_s})$ and $\boldsymbol{\eta}_s' \sim N(\hat{\boldsymbol{\eta}}_s, \boldsymbol{\Sigma}_{\eta_s})$ per Eqs. (2.1) and (2.2).
    return updated $\boldsymbol{\Theta}_2'$
  else ▷ One-segment model $\mathcal{M}_1$
    Update Coefficients: Sample $\boldsymbol{\beta}_1' \sim N(\hat{\boldsymbol{\beta}}_1, \boldsymbol{\Sigma}_{\beta_1})$ and $\boldsymbol{\eta}_1' \sim N(\hat{\boldsymbol{\eta}}_1, \boldsymbol{\Sigma}_{\eta_1})$ using data from the full series.
    return updated $\boldsymbol{\Theta}_1'$
  end if
end procedure

for $i = 1$ **to** N **do**
  `// - Step 1: Data Augmentation -`
  Based on current state $(k^{(i-1)}, \boldsymbol{\Theta}_k^{(i-1)})$, sample latent variables $(\boldsymbol{V}^{(i)}, \boldsymbol{B}^{(i)})$ for all $t = 1, \ldots, n$:
  If $Y_t > 0$, set $V_t \leftarrow Y_t, B_t \leftarrow 0$.
  If $Y_t = 0$, sample $B_t \in \{0, 1\}$ from its conditional posterior. If $B_t = 1$, sample $V_t \sim \text{TruncatedPoisson}(\theta_t; V_t > 0)$.
  `// - Step 2: Model Update/Jump -`
  Draw $u_{move} \sim U(0, 1)$.
  Let $(k^{(i)}, \boldsymbol{\Theta}_k^{(i)}) \leftarrow (k^{(i-1)}, \boldsymbol{\Theta}_k^{(i-1)})$ ▷ Initialize current state as previous state
  if $u_{move} > p_{jump}$ **then** ▷ Perform a within-model update
    $\boldsymbol{\Theta}_k^{(i)} \leftarrow \text{UpdateWithinModel}(k^{(i-1)}, \boldsymbol{\Theta}_k^{(i-1)}, \boldsymbol{V}^{(i)}, \boldsymbol{B}^{(i)})$
  else ▷ Attempt a between-model jump
    if $k^{(i-1)} = 1$ **then** ▷ Propose Birth: $\mathcal{M}_1 \to \mathcal{M}_2$
      Propose new state $(2, \boldsymbol{\Theta}_2')$ from $\boldsymbol{\Theta}_1^{(i-1)}$ via Eq. (3.1).
      Calculate acceptance ratio $A_{1\to2}$ via Eq. (3.3).
      Draw $u_{accept} \sim U(0, 1)$.
      if $u_{accept} < \min(1, A_{1\to2})$ **then**
        $(k^{(i)}, \boldsymbol{\Theta}_k^{(i)}) \leftarrow (2, \boldsymbol{\Theta}_2')$ ▷ Accept the move
      end if
    else ▷ Propose Death: $\mathcal{M}_2 \to \mathcal{M}_1$
      Propose new state $(1, \boldsymbol{\Theta}_1')$ from $\boldsymbol{\Theta}_2^{(i-1)}$ via Eq. (3.2).
      Calculate acceptance ratio $A_{2\to1} = A_{1\to2}^{-1}$.
      Draw $u_{accept} \sim U(0, 1)$.
      if $u_{accept} < \min(1, A_{2\to1})$ **then**
        $(k^{(i)}, \boldsymbol{\Theta}_k^{(i)}) \leftarrow (1, \boldsymbol{\Theta}_1')$ ▷ Accept the move
      end if
    end if
  end if
  Store final state: $S^{(i)} = (k^{(i)}, \boldsymbol{\Theta}_k^{(i)})$.
end for

Data Availability Statement

The data used in this study are publicly available and can be accessed from Our World in Data. The dataset contains COVID-19 and vaccination data compiled by Our World in Data and is updated regularly to reflect the latest information on COVID-19 vaccinations worldwide.

Code and Reproducibility Information

All numerical experiments, simulations, and data analyses presented in this manuscript were conducted using the R statistical computing environment (version 4.3.2). Key packages utilized include `MASS` for generalized linear model fitting and `mvtnorm` for multivariate normal distribution sampling. To facilitate full reproducibility of the reported results, the random number generator seed was fixed at `423` for all stochastic processes. Furthermore, the complete R source code covering data generation, model fitting, and the RJMCMC algorithm implementation is publicly available in our GitHub repository.

References

Agarwal DK, Gelfand AE, Citron-Pousty S. Zero-inflated models with application to spatial count data. Environmental and Ecological Statistics. 2002;9:341–355.

Angers JF, Biswas A. A Bayesian analysis of zero-inflated generalized Poisson model. Computational Statistics & Data Analysis. 2003;42(1-2):37–46.

Bhattacharyya A, Mitra R, Rai S, Pal S. Bayesian shrinkage priors in zero-inflated and negative binomial regression models with real world data applications of COVID-19 vaccine, and RNA-Seq. medRxiv. 2022;p. 2022–07.

Böhning D. Zero-inflated Poisson models and CA MAN: A tutorial collection of evidence. Biometrical Journal: Journal of Mathematical Methods in Biosciences. 1998;40(7):833–843.

Van den Broek J. A score test for zero inflation in a Poisson distribution. Biometrics. 1995;51(2):738–743.

Brooks ME, Kristensen K, van Benthem KJ, Magnusson A, Berg CW, Nielsen A, et al. Modeling zero-inflated count data with glmmTMB. BioRxiv. 2017;p. 132753.

Brooks SP, Giudici P, Roberts GO. Efficient construction of reversible jump Markov chain Monte Carlo proposal distributions. Journal of the Royal Statistical Society Series B: Statistical Methodology. 2003;65(1):3–39.

Cheung YB. Zero-inflated models for regression analysis of count data: a study of growth and development. Statistics in Medicine. 2002;21(10):1461–1469.

Fearnhead P. Exact and efficient Bayesian inference for multiple changepoint problems. Statistics and Computing. 2006;16(2):203–213.

Gelman A, Carlin JB, Stern HS, Rubin DB. Bayesian data analysis. Chapman and Hall/CRC; 1995.

Ghosh SK, Mukhopadhyay P, Lu JCJ. Bayesian analysis of zero-inflated regression models. Journal of Statistical Planning and Inference. 2006;136(4):1360–1375.

Gilks WR, Wild P. Adaptive rejection sampling for Gibbs sampling. Journal of the Royal Statistical Society: Series C (Applied Statistics). 1992;41(2):337–348.

Green PJ. Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. Biometrika. 1995;82(4):711–732.

Grimm KJ, Stegmann G. Modeling change trajectories with count and zero-inflated outcomes: Challenges and recommendations. Addictive Behaviors. 2019;94:4–15.

Hall DB. Zero-inflated Poisson and binomial regression with random effects: a case study. Biometrics. 2000;56(4):1030–1039.

Heilbron D, Gibson D. Shared needle use and health beliefs concerning AIDS: Regression modeling of zero-heavy count data. Poster session. In: Proceedings of the sixth international conference on AIDS, San Francisco, CA; 1990. .

Jansakul N, Hinde J. Score tests for zero-inflated Poisson models. Computational Statistics & Data Analysis. 2002;40(1):75–96.

Junnumtuam S, Niwitpong SA, Niwitpong S. Bayesian Computation for the Parameters of a Zero-Inflated Cosine Geometric Distribution with Application to COVID-19 Pandemic Data. CMES-Computer Modeling in Engineering & Sciences. 2023;135(2):1229–1254.

Khedhiri S. Statistical modeling of COVID-19 deaths with excess zero counts. Epidemiologic Methods. 2021;10(s1):20210007.

Klein N, Kneib T, Lang S. Bayesian generalized additive models for location, scale, and shape for zero-inflated and overdispersed count data. Journal of the American Statistical Association. 2015;110(509):405–419.

Lambert D. Zero-inflated Poisson regression, with an application to defects in manufacturing. Technometrics. 1992;34(1):1–14.

Lee S, Lee Y, Chen CW. Parameter change test for zero-inflated generalized Poisson autoregressive models. Statistics. 2016;50(3):540–557.

Loeys T, Moerkerke B, De Smet O, Buysse A. The analysis of zero-inflated count data: Beyond zero-inflated Poisson regression. British Journal of Mathematical and Statistical Psychology. 2012;65(1):163–180.

Majidizadeh M. Bayesian analysis of the COVID-19 pandemic using a Poisson process with change-points. Monte Carlo Methods and Applications. 2024;30(4):449–465.

Majidizadeh M, Taheriyoun AR. Bayesian multiple change-points detection in autocorrelated binary process with application to COVID-19 infection pattern. Journal of Statistical Computation and Simulation. 2024;94(17):3723–3749.

Neelon B. Bayesian zero-inflated negative binomial regression based on pólya-gamma mixtures. Bayesian Analysis. 2019;14(3):829.

Özmen İ, Demirhan H. A Bayesian approach for zero-inflated count regression models by using the reversible jump Markov chain Monte Carlo method and an application. Communications in Statistics—Theory and Methods. 2010;39(12):2109–2127.

Perumean-Chaney SE, Morgan C, McDowall D, Aban I. Zero-inflated and overdispersed: what's one to do? Journal of Statistical Computation and Simulation. 2013;83(9):1671–1683.

Ridout M, Hinde J, Demétrio CG. A score test for testing a zero-inflated Poisson regression model against zero-inflated negative binomial alternatives. Biometrics. 2001;57(1):219–223.

Rodrigues J. Bayesian analysis of zero-inflated distributions. Communications in Statistics-Theory and Methods. 2003;32(2):281–289.

Scott AJ, Knott M. A cluster analysis method for grouping means in the analysis of variance. Biometrics. 1974;p. 507–512.

Tanner MA, Wong WH. The calculation of posterior distributions by data augmentation. Journal of the American Statistical Association. 1987;82(398):528–540.

Tawiah K, Iddrisu WA, Asampana Asosega K. Zero-Inflated Time Series Modelling of COVID-19 Deaths in Ghana. Journal of Environmental and Public Health. 2021;2021(1):5543977.

Vostrikova LY; Russian Academy of Sciences. Detecting "disorder" in multidimensional random processes. 1981;259(2):270–274.

Wen CC, Baker N, Paul R, Hill E, Hunt K, Li H, et al. A Bayesian zero-inflated beta-binomial model for longitudinal data with group-specific changepoints. Statistics in Medicine. 2024;43(1):125–140.

Zhang NR, Siegmund DO. A modified Bayes information criterion with applications to the analysis of comparative genomic hybridization data. Biometrics. 2007;63(1):22–32, 309.

Zhou W, Huang D, Liang Q, Huang T, Wang X, Pei H, et al. Early warning and predicting of COVID-19 using zero-inflated negative binomial regression model and negative binomial regression model. BMC Infectious Diseases. 2024;24(1):1006.

A Exploratory Analysis for Multiple Change-Points

The primary contribution of this paper is a principled RJMCMC framework for identifying a single change-point in zero-inflated time series. However, in many real-world applications, a series may exhibit multiple structural breaks. As a proof-of-concept for future extensions, we demonstrate here how our developed algorithm can be applied iteratively to explore series with multiple change-points. This heuristic, which we term Iterative Bayesian Segmentation, preserves the core benefits of the Bayesian paradigm, particularly in quantifying uncertainty for each detected change-point.

This approach stands in contrast to common frequentist methods like Binary Segmentation (Scott and Knott, 1974; Vostrikova, 1981), which, while algorithmically similar, typically provide only point estimates for change-point locations and lack a cohesive framework for uncertainty propagation. Our iterative method, while not a fully integrated joint model, maintains methodological consistency with the main body of this work. A formal Bayesian treatment for an unknown number of change-points would necessitate a more complex RJMCMC sampler with trans-dimensional moves that "split" or "merge" adjacent segments, a topic of advanced research (Green, 1995; Fearnhead, 2006). The following simulation illustrates the viability of our iterative heuristic.

A.1 Simulation Design

We simulated a single time series of length $n = 300$ from a ZIP regression model with two known change-points at $m_1 = 100$ and $m_2 = 200$. These points divide the data into three distinct regimes. A simple linear covariate, $x_t = t/n$, was included. The parameters for the Poisson mean ($\boldsymbol{\beta}$) and zero-inflation probability ($\boldsymbol{\eta}$) for each segment were defined as follows:

- **Segment 1 (**$1 \leq t \leq 100$**):** Baseline regime with moderate counts and moderate zero-inflation.

$$\boldsymbol{\beta}_1 = [\beta_{10}, \beta_{11}]^T = [0.8, 0.1]^T, \quad \boldsymbol{\eta}_1 = [\eta_{10}, \eta_{11}]^T = [0.5, -0.1]^T$$

- **Segment 2 (**$101 \leq t \leq 200$**):** Outbreak regime with a sharp increase in counts and a significant decrease in zero-inflation.

$$\boldsymbol{\beta}_2 = [\beta_{20}, \beta_{21}]^T = [2.0, 0.1]^T, \quad \boldsymbol{\eta}_2 = [\eta_{20}, \eta_{21}]^T = [-1.0, -0.1]^T$$

- **Segment 3 (**$201 \leq t \leq 300$**):** Post-intervention regime where counts decrease but remain above baseline, and zero-inflation partially returns.

$$\boldsymbol{\beta}_3 = [\beta_{30}, \beta_{31}]^T = [1.2, 0.1]^T, \quad \boldsymbol{\eta}_3 = [\eta_{30}, \eta_{31}]^T = [0.0, -0.1]^T$$

The parameter shifts were designed to be detectable but challenging enough to produce realistic estimation errors.

A.2 Methodology and Results

The Iterative Bayesian Segmentation procedure was applied in stages:

1. **Initial Run on Full Series:** The single change-point RJMCMC algorithm was first run on the entire series $\boldsymbol{Y} = \{y_t\}_{t=1}^{300}$. The algorithm strongly favored the two-segment model, with an estimated posterior probability $\hat{P}(\mathcal{M}_{00}|\boldsymbol{Y}) > 0.79$. The posterior distribution for the change-point location, m, was unimodal, with a posterior mode at $\hat{m}^{(1)} = 112$. This represents an estimation error of 8 time units relative to the first true change-point ($m_1 = 100$), which exhibited the larger parameter shift.

2. **Iterative Decomposition and Analysis:** The series was then split into two sub-series based on the detected change-point: $\boldsymbol{Y}_A = \{y_t\}_{t=1}^{112}$ and $\boldsymbol{Y}_B = \{y_t\}_{t=113}^{300}$. Our RJMCMC algorithm was subsequently run on each sub-series independently.

   - For the first sub-series, $\boldsymbol{Y}_A$, which contained no other true change-points, the algorithm correctly favored the one-segment (no change-point) model, with $\hat{P}(\mathcal{M}_0|\boldsymbol{Y}_A) \approx 0.74$. This correctly indicated the absence of further structural breaks in this segment.
   - For the second sub-series, $\boldsymbol{Y}_B$, which contained the second true change-point at $t = 200$, the algorithm once again detected a change-point with high posterior probability, $\hat{P}(\mathcal{M}_{00}|\boldsymbol{Y}_B) > 0.72$. The posterior mode for the location of this change-point within the sub-series was at $t' = 77$. Mapping this back to the original time scale gives a location of $\hat{m}^{(2)} = 112 + 77 = 189$. This estimate has an error of 11 time units relative to the second true change-point ($m_2 = 200$).

To assess the general reliability of this iterative procedure, we performed a small-scale simulation study with 50 independent datasets generated from the same design. The iterative method successfully identified the presence of two distinct change-points in 41 of the 50 runs, yielding a success rate of 82%. In the remaining cases, the procedure typically identified only the first, more prominent change-point at $m_1 = 100$ and failed to detect the second, more subtle change at $m_2 = 200$ in the subsequent iterative step.

A.3 Discussion and Limitations

This exploratory analysis demonstrates that our single change-point model can be adapted into a practical iterative heuristic for detecting multiple change-points. The method successfully located both structural breaks in our representative example with realistic estimation errors, and it maintained a fully Bayesian inferential framework, providing posterior distributions for the location of each change-point.

However, we must acknowledge the limitations of this approach. First, as a greedy, sequential method, it is not guaranteed to find the globally optimal segmentation. Second, the uncertainty from the estimation of the first change-point's location is not

propagated into the analysis of subsequent segments, which a fully joint model would naturally handle. Despite these limitations, this iterative procedure serves as a powerful and computationally efficient proof-of-concept, establishing a strong foundation for future research into a fully integrated RJMCMC model for multiple change-points in zero-inflated regression contexts.

B Figures and Tables

Table 1: Estimated parameters and standard deviations for two model specifications. The upper table reports parameter estimates for the model with a single change-point (two-segment dataset; see Dataset 4.2), while the lower table presents the corresponding estimates for the model without a change-point (one-segment dataset; see Dataset 4.1).

Parameter	**Model 2 (Segment 1)**	**Model 2 (Segment 2)**
β_0	0.5315 (0.1120)	-0.4501 (0.1991)
β_1	0.8528 (0.0951)	-0.5908 (0.1824)
β_2	-1.2767 (0.1899)	2.1455 (0.2973)
η_0	-9.7449 (2.8781)	-5.1565 (4.0903)
η_1	-0.3616 (3.4460)	5.3955 (3.7102)
η_2	-0.1615 (3.704)	9.2580 (4.4904)
	Model 1 (Segment 1)	
β_0	0.4544 (0.0859)	
β_1	0.7616 (0.0538)	
β_2	-0.9634 (0.1325)	
η_0	-7.8570 (3.0709)	
η_1	1.5260 (1.2407)	
η_2	2.7830 (4.5212)	

Table 2: Estimated Parameters and Their Standard Deviations for Cyprus and Qatar. The upper table shows the results for Cyprus, while the lower table presents those for Qatar.

Parameter	**Cyprus (Segment 1)**	**Cyprus (Segment 2)**
β_0	2.845 (4.927)	15.085 (6.323)
β_1	2.087 (1.415)	4.943 (1.107)
β_2	-1.101 (2.319)	-4.252 (2.190)
η_0	4.085 (6.774)	-10.86 (7.67)
η_1	1.476 (1.847)	4.08 (4.07)
η_2	-2.825 (3.858)	2.46 (6.64)
Parameter	**Qatar (Segment 1)**	**Qatar (Segment 2)**
β_0	12.518 (6.794)	17.616 (9.971)
β_1	3.942 (0.624)	4.674 (1.360)
β_2	-7.348 (0.1899)	-12.347 (3.368)
η_0	-9.927 (2.8781)	-11.92 (13.98)
η_1	5.081 (5.4460)	6.90 (6.97)
η_2	4.622 (9.2704)	6.19 (9.50)

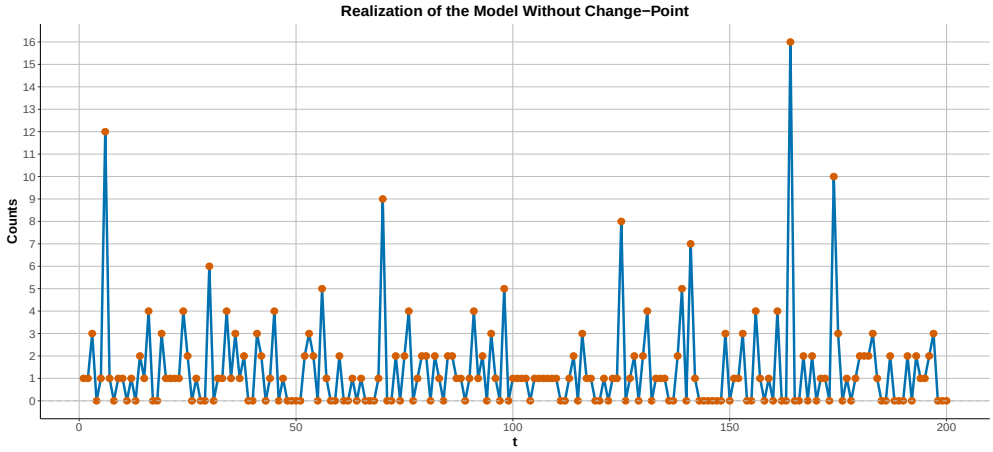

Figure 1: Model Realization Without Change-Point (4.1): This plot shows a single realization of the model with no change-point, generated from 200 observations.

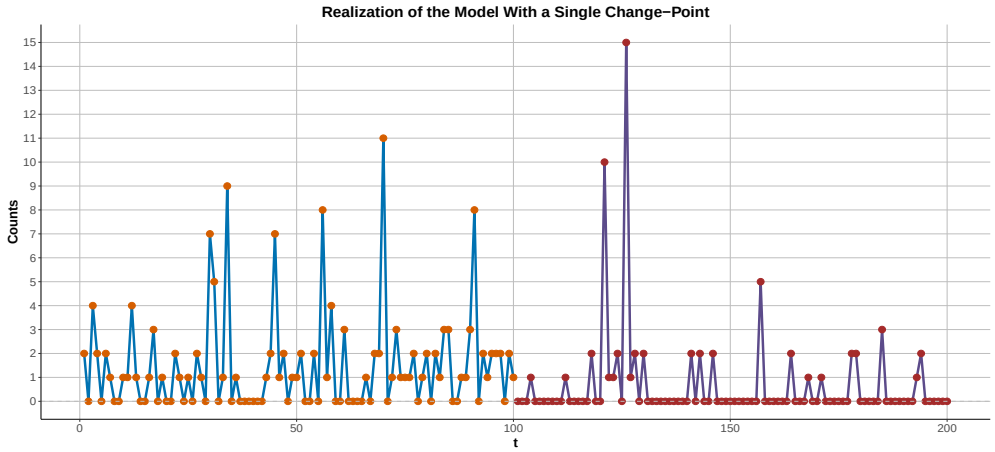

Figure 2: Model Realization with a Single Change-Point (4.2): This plot displays a single realization of the model with one change-point. The first 100 observations belong to the initial segment, and the next 100 to the second segment, distinguished by different colors.

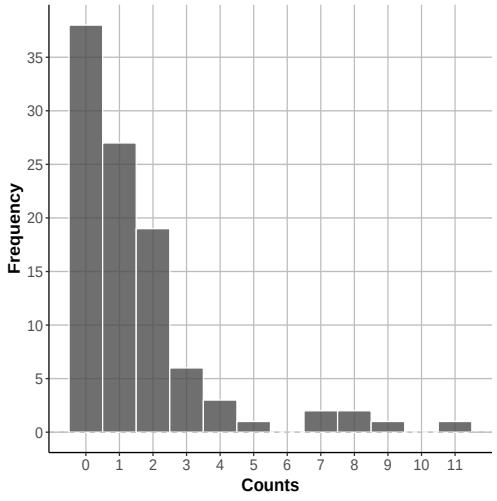

(a) Histogram of Count Data for the First Segment

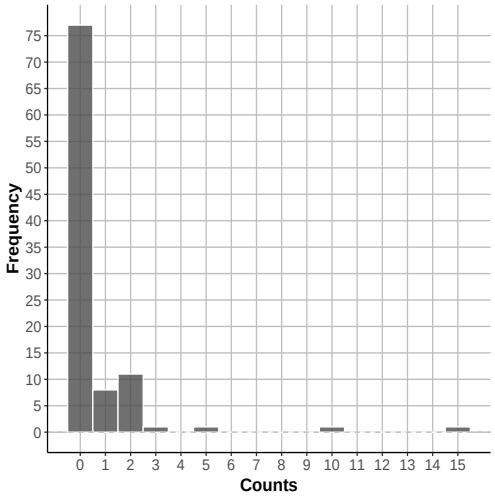

(b) Histogram of Count Data for the Second Segment

Figure 3: Histograms of Count Data for the Single Change-Point Model (4.2). **Panel (a)** shows the histogram for the first segment, and **Panel (b)** displays the histogram for the second segment. These plots clearly illustrate the differences in the count data between the two segments, particularly regarding the frequency of zero and non-zero values.

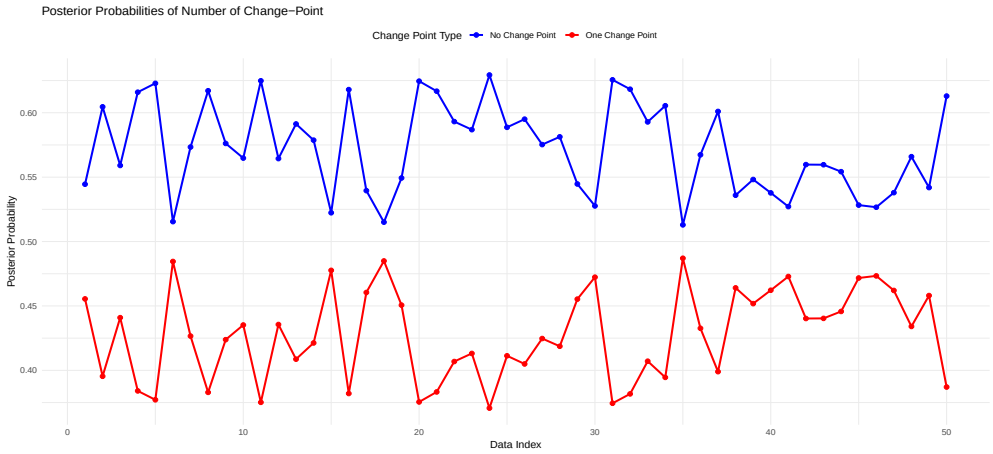

Figure 4: Posterior Probabilities for the Model Without a Change-Point (4.1). This plot presents the posterior probabilities for the no change-point model across 50 independent replications. The results indicate that the model without a change-point is correctly identified in all replications.

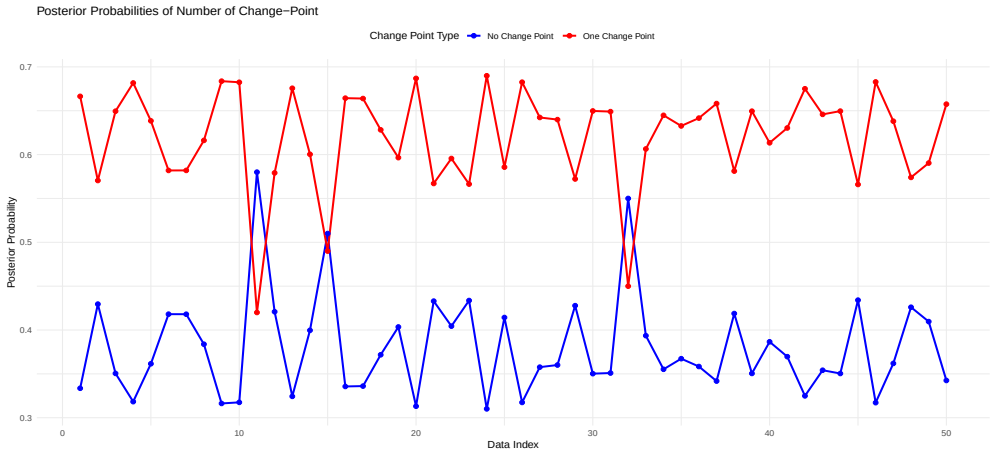

Figure 5: Posterior Probabilities for the Single Change-Point Model (4.2). This plot illustrates the posterior probabilities for the number of change-points (zero or one) across 50 independent replications. The results demonstrate the method's effectiveness, correctly identifying the single change-point in 47 of the 50 cases.

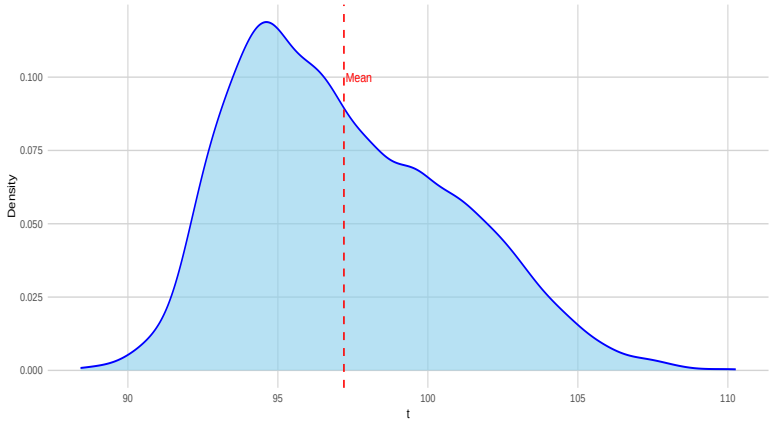

Figure 6: Posterior Density for the Change-Point Location (Single Change-Point Model) (4.2). This plot displays the aggregate posterior density for the location of the single change-point across 50 simulated realizations. The **dashed line** marks the estimated change-point location based on the posterior mean ($\hat{\tau}_{\text{mean}} = 97$), while the true change-point is located at $\tau = 100$. The posterior mode, which provides a maximum a posteriori (MAP) estimate, is at $\hat{\tau}_{\text{mode}} = 94$, indicating a slight bias towards an earlier location.

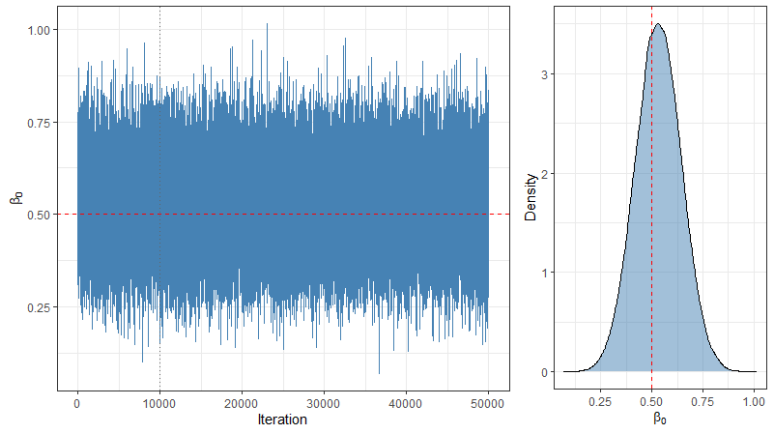

Figure 7: Diagnostic plots for the MCMC chain of the parameter β_0. The left panel shows the trace plot over 50,000 iterations, with the red dashed line indicating the end of the 10,000-iteration burn-in period. The chain appears stationary, resembling white noise, which suggests convergence. The right panel displays the posterior density of β_0 using samples after burn-in. The posterior mean (green dashed line) is very close to the true value (solid black line), indicating accurate estimation.

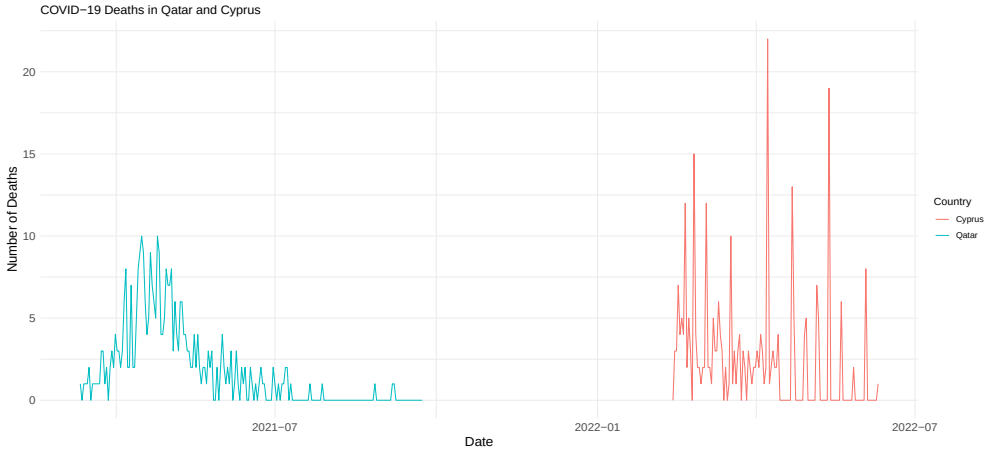

Figure 8: Comparative Visualization of COVID-19 Deaths in Qatar and Cyprus. The **left-hand plot** depicts the time series of COVID-19 fatalities in Qatar during the 195-day interval from March 12, 2021, through September 23, 2021. The **right-hand plot** illustrates the corresponding data for Cyprus, spanning 117 days, from February 13, 2022, to June 10, 2022.

Figure 9: Confirmed COVID-19 Cases in Qatar and Cyprus. The **left panel** displays the time series of confirmed COVID-19 cases in Qatar spanning the period from March 12, 2021, to September 23, 2021. The **right panel** presents the confirmed case data for Cyprus from February 13, 2022, to June 10, 2022.

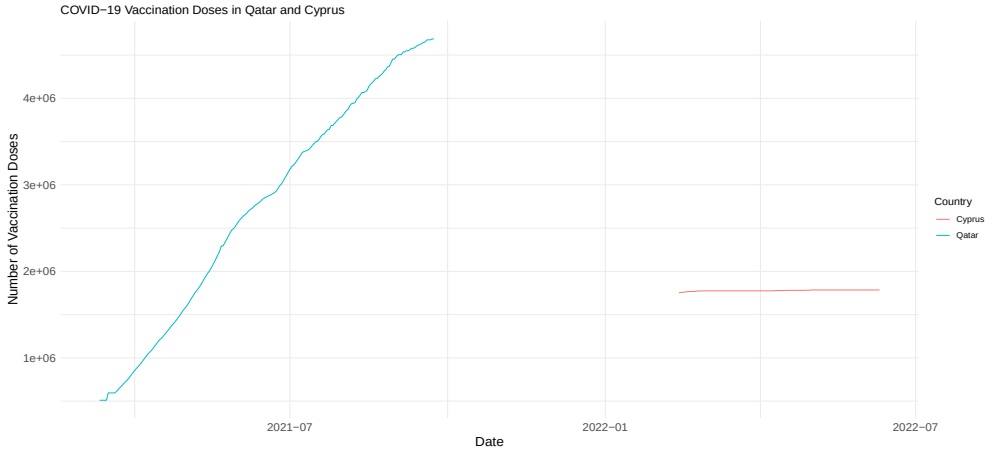

Figure 10: COVID-19 Vaccination Doses Administered in Qatar and Cyprus. The **left panel** displays the time series of total vaccination doses administered in Qatar from March 12, 2021, to September 23, 2021. The **right panel** presents the equivalent vaccination dose data for Cyprus, covering the period from February 13, 2022, to June 10, 2022.

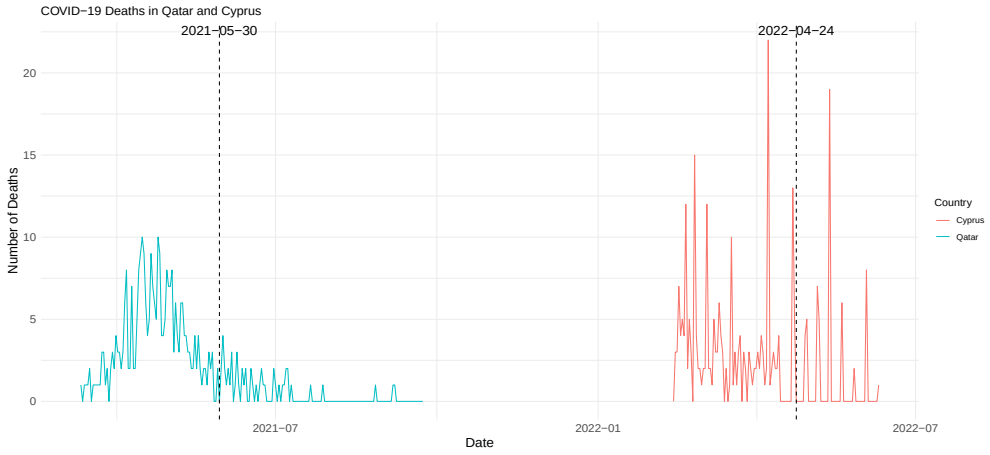

Figure 11: Detected Change-Points in COVID-19 Deaths for Cyprus and Qatar. The figure uses vertical dashed lines to denote the location of the statistically detected change-points, with the corresponding date of each change-point indicated directly above the line.