

A Pólya-Gamma Sampler for a Generalized Logistic Regression

Luciana Dalla Valle*

University of Plymouth, Plymouth, UK

Fabrizio Leisen

University of Nottingham, Nottingham, UK

Luca Rossini

Queen Mary University of London, London, UK

and

Weixuan Zhu

Xiamen University, Xiamen, People's Republic of China

December 20, 2020

Abstract

In this paper we introduce a novel Bayesian data augmentation approach for estimating the parameters of the generalised logistic regression model. We propose a Pólya-Gamma sampler algorithm that allows us to sample from the exact posterior distribution, rather than relying on approximations. A simulation study illustrates the flexibility and accuracy of the proposed approach to capture heavy and light tails in binary response data of different dimensions. The methodology is applied to two different real datasets, where we demonstrate that the Pólya-Gamma sampler provides more precise estimates than the empirical likelihood method, outperforming approximate approaches.

Keywords: Bayesian inference; generalized logistic regression; Pólya-Gamma sampler; recidivism data.

*The authors gratefully acknowledge Alex Diana for advising on the PGdraw sampler and Michael Wögerer for useful feedback on a previous version of the paper. The fourth author is supported by the Chinese Fundamental Research Funds for the Central Universities No 20720181062.

1 Introduction

Data augmentation (DA) is undoubtedly one of the most popular Markov Chain Monte Carlo (MCMC) methods. The idea is to consider the target distribution as the marginal of a joint distribution in an augmented space. Assuming it is easy to sample from the conditional distributions, the algorithm is a straightforward two steps procedure which makes use of a simple Gibbs sampling. Since the seminal work of Tanner and Wong (1987), data augmentation has been extensively studied, developed and used in different contexts. For instance, Meng and van Dyk (1999), Liu and Wu (1999), van Dyk and Meng (2001) developed strategies to speed up the basic DA algorithm. Hobert and Marchev (2008) studied the efficiency of DA algorithms. Recently, Leisen et al. (2017) use a DA approach to perform Bayesian inference for the Yule-Simon distribution. The idea is used in the context of infinite hidden Markov models by Hensley and Djurić (2017). The literature on DA is vast and it is difficult to list all the contributions to the field. A review of DA algorithms can be found in Brooks et al. (2011), Chapter 10.

In the context of binary regression, the article of Albert and Chib (1993) sets a milestone in the use of DA algorithms. The authors develop exact Bayesian methods for modeling categorical response data by introducing suitable auxiliary variables. In particular, they use a DA approach to estimate the parameters of the probit regression model. However, although some valid approaches have been proposed (Holmes et al., 2006), Bayesian inference for the logistic regression was not satisfactorily addressed until the paper of Polson et al. (2013). The authors proposed an elegant DA approach which makes use of the following identity:

$$\frac{(e^\psi)^a}{(1 + e^\psi)^b} = 2^{-b} e^{\kappa\psi} \int e^{-\omega\psi^2/2} p(\omega) d\omega \quad (1)$$

where $\kappa = a - b/2$, $a > 0$, $b > 0$ and $p(\omega)$ is the density of a $\text{PG}(b,0)$ distribution. $\text{PG}(b,0)$ is a Pólya-Gamma distribution with parameters b and 0; we refer to the article of Polson et al. (2013) for details about Pólya-Gamma distributions. When $\psi = \mathbf{x}^T \boldsymbol{\beta}$ is a linear function of

predictors, the integrand is the kernel of a Gaussian likelihood in β . This naturally suggests a DA scheme which simply requires to be able to sample from the multivariate normal and the Pólya-Gamma distributions. This algorithm is called the *Pólya-Gamma sampler*. A theoretical study of the algorithm can be found in Choi and Hobert (2013).

In this paper, we propose a Pólya-Gamma sampler to estimate the parameters of the generalized logistic regression model introduced in Dalla Valle et al. (2019). The authors studied immigration in Europe employing a novel regression model which makes use of a generalized logistic distribution. Their work is motivated by the need of a logistic model with heavy tails. The distribution considered in the paper has the following density function:

$$f(x) = \frac{1}{B(p, p)} \frac{e^{px}}{(1 + e^x)^{2p}} \quad x \in \mathbb{R}, \quad (2)$$

where the parameter p controls the tails of the distribution. The standard logistic density can be recovered with $p = 1$. Compared with the standard logistic distribution, $p > 1$ means lighter tails and $0 < p < 1$ means heavier tails, as shown in Figure 1. Unfortunately, the distribution function is not explicit:

$$F(x) = \frac{1}{B(p, p)} B\left(\frac{e^x}{1 + e^x}; p, p\right) \quad x \in \mathbb{R} \quad (3)$$

where $B(t; p, q) = \int_0^t x^{p-1}(1-x)^{q-1}dx$, with $0 < t < 1$, is the incomplete Beta function. Therefore, a straightforward use of the methodology of Polson et al. (2013) is not possible. To overcome the problem, Dalla Valle et al. (2019) use an approximate Bayesian computational method that relies on the empirical likelihood (see Mengersen et al. (2013) and Karabatsos and Leisen (2018)). However, since the empirical likelihood approach is based on an approximation of the posterior distribution, this method may lead to wide credible intervals and sometimes ambiguous estimates, especially for the tail parameter.

In this paper, we propose a novel approach for estimating the generalised logistic regression, that allows us to draw samples from the exact posterior distribution and is able to overcome

the drawbacks of the empirical likelihood method. We show how to set a DA scheme for the generalized logistic regression which makes use of the Pólya-Gamma identity in equation (1). We test the performance of the proposed approach on simulated data and on two different real dataset: the first dataset includes people’s opinions towards immigration in Europe; the second contains information about the recidivism of criminals detained in Iowa. We show that the Pólya-Gamma sampler yields accurate results in high-dimension and outperforms the empirical likelihood algorithm providing a more precise estimation of the model parameters.

The rest of the paper is organised as follows. In Section 2, we introduce the generalized logistic regression framework in a Bayesian setting. In Section 3, we provide the details of the Pólya-Gamma sampler for this model. Section 4 and Section 5 illustrate the algorithm performance with simulated and real data. Section 6 concludes.




Figure 1: Probability density function of the generalized logistic with $p = 1$, $p = 0, 3$, $p = 0, 7$, $p = 2$ and $p = 5$.

2 The Generalized Bayesian Logistic Regression

Consider a binary regression set-up in which Y_1, \dots, Y_n are independent Bernoulli random variables such that $\Pr(Y_i = 1|\boldsymbol{\beta}) = H(x_i^T \boldsymbol{\beta})$ where x_i is a $k \times 1$ vector of known covariates associated with Y_i , $i = 1, \dots, n$, $\boldsymbol{\beta}$ is a $k \times 1$ vector of unknown regression coefficients, and $H : \mathbb{R} \rightarrow (0, 1)$ is a distribution function. In this case, the likelihood function is given by

$$L(\boldsymbol{\beta}) = \prod_{i=1}^n [H(x_i^T \boldsymbol{\beta})]^{y_i} [1 - H(x_i^T \boldsymbol{\beta})]^{1-y_i} \quad (4)$$

If $H(x) = \Phi(x)$ is the distribution function of a Gaussian, then we are in the probit regression framework. If $H(x) = e^x(1 + e^x)^{-1}$, then we are in the special case of the standard logistic regression. Albert and Chib (1993) proposed a DA approach for sampling from the posterior distribution of the probit regression model. Their method makes use of auxiliary variables which allow an easy implementation of the Gibbs sampler. Polson et al. (2013) proposed an elegant algorithm for tackling the logistic regression case.

The generalized logistic distribution in equation (2) has mean zero and scale one. Consider now a generalized logistic with mean $x_i^T \boldsymbol{\beta}$ and scale one, i.e. with probability density function:

$$f(x) = \frac{1}{B(p, p)} \frac{e^{p(x-x_i^T \boldsymbol{\beta})}}{(1 + e^{(x-x_i^T \boldsymbol{\beta})})^{2p}} \quad x \in \mathbb{R} \quad (5)$$

We denote the above distribution with $\text{GLog}(x_i^T \boldsymbol{\beta}, 1)$. Mimicking Albert and Chib (1993), suppose to sample Z_i from a $\text{GLog}(x_i^T \boldsymbol{\beta}, 1)$, $i = 1, \dots, n$, and set $Y_i = 1$ if $Z_i > 0$. Otherwise, if $Z_i \leq 0$, then set $Y_i = 0$. It is easy to see that

$$P(Y_i = 1) = \frac{1}{B(p, p)} B\left(\frac{e^{x_i^T \boldsymbol{\beta}}}{1 + e^{x_i^T \boldsymbol{\beta}}}; p, p\right).$$

The joint posterior distribution is

$$\pi(\boldsymbol{\beta}, p, \mathbf{Z}|\mathbf{y}) \propto \pi(\boldsymbol{\beta})\pi(p) \prod_{i=1}^n \{\mathbb{I}(Z_i > 0) \mathbb{I}(y_i = 1) + \mathbb{I}(Z_i \leq 0) \mathbb{I}(y_i = 0)\}$$

$$\times \frac{1}{B(p, p)} \frac{e^{p(Z_i - x_i^T \boldsymbol{\beta})}}{(1 + e^{(Z_i - x_i^T \boldsymbol{\beta})})^{2p}} \quad (6)$$

where $\pi(\boldsymbol{\beta})$ and $\pi(p)$ are the prior distributions of $\boldsymbol{\beta}$ and p , respectively. With the probit model, the technique of Albert and Chib (1993) works because the full conditional distributions can be resorted to normal distributions (or truncated normal distributions) which are easy to sample from. In our case, the usual normal prior would lead to a full conditional on $\boldsymbol{\beta}$ which is not explicit, requiring the use of a Metropolis step (or alternative algorithms). The Pólya-Gamma identity in (1) allows us to overcome this problem. In particular, we get

$$\frac{e^{p(Z_i - x_i^T \boldsymbol{\beta})}}{(1 + e^{(Z_i - x_i^T \boldsymbol{\beta})})^{2p}} = 2^{-2p} \int_0^{+\infty} e^{-\omega_i(Z_i - x_i^T \boldsymbol{\beta})^2/2} p(\omega_i) d\omega_i,$$

where $p(\omega)$ is a Pólya-Gamma, $\text{PG}(2p, 0)$, distribution. Note that $e^{-\omega(x_i^T \boldsymbol{\beta})^2/2} p(\omega)$ is the unnormalized density of a $\text{PG}(2p, x_i^T \boldsymbol{\beta})$ random variable. Therefore, the augmented posterior distribution is

$$\begin{aligned} \pi(\boldsymbol{\beta}, p, \mathbf{Z}, \boldsymbol{\omega} | \mathbf{y}) &\propto \pi(\boldsymbol{\beta}) \pi(p) \prod_{i=1}^n \{ \mathbb{I}(Z_i > 0) \mathbb{I}(y_i = 1) + \mathbb{I}(Z_i \leq 0) \mathbb{I}(y_i = 0) \} \\ &\times 2^{-2p} e^{-\omega_i(Z_i - x_i^T \boldsymbol{\beta})^2/2} p(\omega_i). \end{aligned} \quad (7)$$

3 Inference Sampling Strategy

The algorithm proposed in Dalla Valle et al. (2019) makes use of an approximate likelihood, which relies on an approximation of the posterior distribution (see Mengersen et al. (2013) and Zhu and Leisen (2016)). The algorithm presented in this paper allows us to sample from the true posterior distribution, making it more appealing to perform the Bayesian inferential exercise.

Starting from the augmented posterior distribution in (7), this Section introduces the full conditional distributions required to run the algorithm. The variables involved are

$$\mathbf{Z}, \boldsymbol{\omega}, \boldsymbol{\beta}, p,$$

where $\mathbf{Z} = (Z_1, \dots, Z_n)$, $\boldsymbol{\omega} = (\omega_1, \dots, \omega_n)$ and $\boldsymbol{\beta} = (\beta_1, \dots, \beta_k)$.

We will focus on two cases: 1) the parameter p is known, 2) the parameter p is unknown. In the first case, the algorithm performance is very good and accurate. In the second case, the estimation of the parameter p introduces more noise in the algorithm leading to a less accurate, but still acceptable, estimation. Section 4 will provide a simulation study for both scenarios. The real data analysis assumes that p is unknown.

The parameter p is known

In this case, the posterior distribution in (7) reduces to:

$$\begin{aligned} \pi(\boldsymbol{\beta}, \mathbf{Z}, \boldsymbol{\omega} | \mathbf{y}) &\propto \pi(\boldsymbol{\beta}) \prod_{i=1}^n \{\mathbb{I}(Z_i > 0) \mathbb{I}(y_i = 1) + \mathbb{I}(Z_i \leq 0) \mathbb{I}(y_i = 0)\} \\ &\quad \times 2^{-2p} e^{-\omega_i(Z_i - x_i^T \boldsymbol{\beta})^2 / 2} p(\omega_i). \end{aligned} \quad (8)$$

The full conditional distributions are derived below and the variables of interest are:

$$\mathbf{Z}, \boldsymbol{\omega}, \boldsymbol{\beta}.$$

Full conditional for Z_i . The random variables Z_1, \dots, Z_n are independent with full conditional distributions:

$$\begin{aligned} \pi(Z_i | \mathbf{y}, \boldsymbol{\beta}, \boldsymbol{\omega}) &\propto \{\mathbb{I}(Z_i > 0) \mathbb{I}(y_i = 1) + \mathbb{I}(Z_i \leq 0) \mathbb{I}(y_i = 0)\} e^{-\omega_i(Z_i - x_i^T \boldsymbol{\beta})^2 / 2} \\ &\propto \{\mathbb{I}(Z_i > 0) \mathbb{I}(y_i = 1) + \mathbb{I}(Z_i \leq 0) \mathbb{I}(y_i = 0)\} e^{-\frac{\omega_i}{2}(Z_i^2 - 2Z_i x_i^T \boldsymbol{\beta})} \end{aligned}$$

which are truncated normal distributions. More precisely,

- if $y_i = 1$, then $\pi(Z_i|\mathbf{y}, \boldsymbol{\beta}, \boldsymbol{\omega})$ is distributed as a $\mathcal{N}(x_i^T \boldsymbol{\beta}, \omega_i^{-1})$ truncated to the left by 0;
- if $y_i = 0$, then $\pi(Z_i|\mathbf{y}, \boldsymbol{\beta}, \boldsymbol{\omega})$ is distributed as a $\mathcal{N}(x_i^T \boldsymbol{\beta}, \omega_i^{-1})$ truncated to the right by 0.

Full conditional for ω_i . Following Polson et al. (2013), it is easy to see that the full conditional distribution of ω_i is:

$$\pi(\omega_i|\boldsymbol{\beta}, \mathbf{Z}, \mathbf{y}) \propto e^{-\omega_i(Z_i - x_i^T \boldsymbol{\beta})^2/2} p(\omega_i) \sim \text{PG}(2p, Z_i - x_i^T \boldsymbol{\beta})$$

An efficient method for sampling Pólya-Gamma random variables is implemented in a modified version of the R package `BayesLogit` (Windle et al., 2016).

Full conditional for $\boldsymbol{\beta}$. Let $\Omega = \text{diag}(\omega_1, \dots, \omega_n)$. Assuming a multivariate $\mathcal{N}(\boldsymbol{\beta}^*, B^*)$ prior for $\boldsymbol{\beta}$, the full conditional distribution of $\boldsymbol{\beta}$ is

$$\begin{aligned} \pi(\boldsymbol{\beta}|\mathbf{Z}, \mathbf{y}, \boldsymbol{\omega}) &\propto \pi(\boldsymbol{\beta}) \prod_{i=1}^n \exp \left\{ -\frac{\omega_i}{2} (Z_i - x_i^T \boldsymbol{\beta})^2 \right\} \\ &= \exp \left\{ -\frac{1}{2} [\boldsymbol{\beta}^T (X^T \Omega X + B^{*-1}) \boldsymbol{\beta} - 2\boldsymbol{\beta}^T (X^T \Omega \mathbf{Z} + B^{*-1} \boldsymbol{\beta}^*)] \right\} \\ &\sim \mathcal{N}(\tilde{\boldsymbol{\beta}}^*, \tilde{V}_{\boldsymbol{\beta}}^*) \end{aligned}$$

where

$$\begin{aligned} \tilde{\boldsymbol{\beta}}^* &= (X^T \Omega X + B^{*-1})^{-1} (X^T \Omega \mathbf{Z} + B^{*-1} \boldsymbol{\beta}^*) \\ \tilde{V}_{\boldsymbol{\beta}}^* &= (X^T \Omega X + B^{*-1})^{-1}. \end{aligned}$$

The parameter p is unknown

In this case, the full conditional for p is:

$$\pi(p|\boldsymbol{\beta}, \mathbf{Z}, \boldsymbol{\omega}, \mathbf{y}) \propto \pi(p) \prod_{i=1}^n \frac{2^{-2p}}{B(p, p)} p(\omega_i)$$

where $p(\omega_i)$ is a $\text{PG}(2p, 0)$ and we assume a Gamma prior for p . Sampling from this full conditional distribution is not straightforward and alternative strategies must be employed. One can marginalize the above full conditional by integrating out the ω 's and applying the slice sampling of Neal (2003). Alternatively, one can further integrate out the \mathbf{Z} and obtain the following conditional

$$\pi(p|\boldsymbol{\beta}, \mathbf{y}) \propto \pi(p) \prod_{i=1}^n \left[\frac{1}{B(p, p)} B\left(\frac{e^{x_i^T \boldsymbol{\beta}}}{1 + e^{x_i^T \boldsymbol{\beta}}}; p, p\right) \right]^{y_i} \left[1 - \frac{1}{B(p, p)} B\left(\frac{e^{x_i^T \boldsymbol{\beta}}}{1 + e^{x_i^T \boldsymbol{\beta}}}; p, p\right) \right]^{1-y_i} \quad (9)$$

The slice sampling algorithm requires the computation of the inverse of the above full conditional distribution. We used the function `uniroot` available in R to compute the numerical inverse.

The next sections illustrate the performance of the algorithm with simulated and real data.

4 Simulation Studies

In this section, we assess the performance of the algorithm by implementing different simulation experiments. In particular, we use four different values of the tail parameter p : $p = 0.3$ and $p = 0.7$ (heavy tails) and $p = 1.5$ and $p = 3$ (light tails). We focus on two scenarios:

1. $\boldsymbol{\beta} = (1, -1, -3, 1, 3)$, with $k = 5$;
2. $\boldsymbol{\beta} = (2.3, 1, -2, 1.5, -2.7, 0.2, -1.4, 3, -0.6, -1.2)$, with $k = 10$.

We assume a standard normal distribution for the explanatory variables. For each scenario and each value of p , we generate 100 different datasets of sample sizes $n = 100$ and $n = 250$, respectively. We analyse two different cases: the case with p known and the case with p unknown.

4.1 Case with p known

We considered a vague prior for β , such as the multivariate normal distribution, $\mathcal{N}(\boldsymbol{\nu}, B)$, with prior mean vector $\boldsymbol{\nu} = 0$ and prior covariance matrix $B = 100 \cdot I_k$. For each of the 100 different simulated datasets, we run 20,000 iterations of the MCMC algorithm and discarded the first 5,000 iterations as burn-in period. Table 1 lists the posterior means of the five-dimensional β of scenario 1 for the different values of p and n , averaged over the 100 simulated datasets. The values in brackets show the standard deviations over the 100 simulations. We notice that the posterior means of β are very close to the true values, with a better estimation performance as n increases from 100 to 250. Table 3 shows the posterior means and, in brackets, the standard deviations over the 100 simulations for the ten-dimensional scenario. These results are in line with the five-dimensional case, thus leading us to conclude that increasing the number of covariates is not an issue in terms of estimation accuracy. Moreover, the standard deviations in Tables 1 and 3 substantially decrease for both scenarios when we move from $n = 100$ to $n = 250$.

We compare the performance of our Pólya-Gamma sampler to the one of the approximate Bayesian computation with empirical likelihood (BC_{el}) adopted in Dalla Valle et al. (2019). Table 2 and Table 4 display the posterior means and the standard deviations over 100 simulation datasets estimated with BC_{el} . The sampling efficiency of BC_{el} depends heavily on the prior distribution. So we consider a normal prior distribution $\mathcal{N}(\boldsymbol{\nu} = 0, B = 9 \cdot I_k)$ for β . The results are based on 20,000 posterior samples obtained with BC_{el} . From the table, one can see that both methods achieve similar accuracy level in terms of the estimation of the posterior means. However, the standard deviations given by BC_{el} is larger than the ones obtained by Pólya-Gamma sampler, especially when p parameter is large. Finally, consider that the prior distribution used by Pólya-Gamma sampler is much vaguer than the one used by BC_{el} , we conclude that Pólya-Gamma sampler outperforms BC_{el} .

In order to illustrate the complete inferential procedure, we show how the model parameters are estimated. Figure 2 displays the posterior chains for a specific dataset with sample size $n = 250$, where $p = 0.3$ and $\boldsymbol{\beta} = (\beta_0, \beta_1, \beta_2, \beta_3, \beta_4) = (1, -1, -3, 1, 3)$. As shown in Figure 2, the chains, after a burn-in of 5,000 iterations, converge quickly to the true values of all parameters and give an excellent representation of the real parameter values.

We also performed a convergence analysis for scenario 1, with five-dimensional $\boldsymbol{\beta}^1$. The analysis was carried out using the R `coda` package (Plummer et al., 2006). In particular, we used single datasets of sample sizes $n = 100$ and $n = 250$, respectively, where we fixed $p = 0.3$. We computed Geweke’s convergence tests and the autocorrelation and partial autocorrelation functions for the values of $\boldsymbol{\beta}$. Table 5 reports the results of Geweke’s convergence test (Geweke, 1992), indicating no issues since the absolute values of the test statistics are all lower than 1.96.

4.2 Case with p unknown

We considered vague priors for $\boldsymbol{\beta}$ and p . In particular, a multivariate normal distribution, $\mathcal{N}(\boldsymbol{\nu}, B)$, is chosen for $\boldsymbol{\beta}$ with prior mean vector $\boldsymbol{\nu} = 0$ and prior covariance matrix $B = 100 \cdot I_k$. For p we chose a gamma prior, $\mathcal{Ga}(a, b)$ with hyperparameters $a = b = 1$. For each of the 100 different simulated datasets, we run 6,000 iterations of the MCMC algorithm and discard the first 1,000 iterations as burn-in period. Table 6 lists the posterior means of p and the five-dimensional $\boldsymbol{\beta}$ of scenario 1 for the different values of n and the true p , averaged over the 100 simulated datasets. The figures in brackets show the standard deviations over the 100

¹We performed other convergence tests for the ten-dimensional case. These results are omitted for lack of space but are available on request.

	p	β_0	β_1	β_2	β_3	β_4
<i>true values</i>	0.3	1	-1	-3	1	3
$n = 100$	–	1.113	-1.128	-3.191	1.056	3.396
	–	(0.795)	(0.666)	(0.937)	(0.696)	(0.904)
$n = 250$	–	1.041	-1.061	-3.173	1.027	3.273
	–	(0.411)	(0.412)	(0.548)	(0.451)	(0.555)
<i>true values</i>	0.7	1	-1	-3	1	3
$n = 100$	–	1.179	-1.094	-3.282	1.148	3.309
	–	(0.448)	(0.545)	(0.640)	(0.558)	(0.659)
$n = 250$	–	1.029	-1.110	-3.187	1.085	3.206
	–	(0.304)	(0.295)	(0.503)	(0.312)	(0.479)
<i>true values</i>	1.5	1	-1	-3	1	3
$n = 100$	–	1.169	-1.124	-3.605	1.158	3.582
	–	(0.426)	(0.430)	(0.867)	(0.470)	(0.782)
$n = 250$	–	1.096	-1.068	-3.277	1.111	3.326
	–	(0.258)	(0.238)	(0.541)	(0.298)	(0.518)
<i>true values</i>	3	1	-1	-3	1	3
$n = 100$	–	1.183	-1.214	-3.681	1.239	3.729
	–	(0.417)	(0.391)	(0.853)	(0.478)	(0.984)
$n = 250$	–	1.130	-1.154	-3.431	1.119	3.408
	–	(0.253)	(0.284)	(0.653)	(0.239)	(0.662)

Table 1: PG sampler. Case with p known: posterior means over the 100 different simulated datasets for n equal to 100 and 250 with true values of $p = 0.3$ (first panel), $p = 0.7$ (second panel), $p = 1.5$ (third panel) and $p = 3$ (fourth panel) and $(\beta_0, \beta_1, \beta_2, \beta_3, \beta_4) = (1, -1, -3, 1, 3)$. The values in brackets are the standard deviations over the 100 different simulations.

	p	β_0	β_1	β_2	β_3	β_4
<i>true values</i>	0.3	1	-1	-3	1	3
$n = 100$	–	1.185	-1.136	-3.356	1.159	3.403
	–	(0.773)	(0.708)	(0.900)	(0.785)	(0.844)
$n = 250$	–	1.052	-1.067	-3.185	1.045	3.292
	–	(0.506)	(0.435)	(0.600)	(0.488)	(0.630)
<i>true values</i>	0.7	1	-1	-3	1	3
$n = 100$	–	1.146	-1.177	-3.751	1.331	3.65
	–	(0.859)	(0.711)	(0.750)	(0.778)	(0.819)
$n = 250$	–	1.315	-1.214	-3.289	1.251	3.304
	–	(0.511)	(0.412)	(0.614)	(0.407)	(0.687)
<i>true values</i>	1.5	1	-1	-3	1	3
$n = 100$	–	1.326	-1.086	-3.747	1.187	3.714
	–	(1.214)	(0.618)	(1.382)	(0.600)	(0.997)
$n = 250$	–	1.064	-1.019	-3.068	1.010	3.643
	–	(0.513)	(0.376)	(0.598)	(0.399)	(0.736)
<i>true values</i>	3	1	-1	-3	1	3
$n = 100$	–	1.473	-1.373	-3.761	1.373	3.982
	–	(1.044)	(1.049)	(1.556)	(0.913)	(1.739)
$n = 250$	–	1.180	-1.256	-3.601	1.186	3.534
	–	(0.513)	(0.591)	(1.210)	(0.539)	(1.111)

Table 2: BC_{el} sampler. Case with p known: posterior means over the 100 different simulated datasets for n equal to 100 and 250 with true values of $p = 0.3$ (first panel), $p = 0.7$ (second panel), $p = 1.5$ (third panel) and $p = 3$ (fourth panel) and $(\beta_0, \beta_1, \beta_2, \beta_3, \beta_4) = (1, -1, -3, 1, 3)$. The values in brackets are the standard deviations over the 100 different simulations.

p	β_0	β_1	β_2	β_3	β_4	β_5	β_6	β_7	β_8	β_9
<i>true values</i>	0.3	1	-2	1.5	-2.7	0.2	-1.4	3	-0.6	-1.2
$n = 100$	- 2.597 (0.760)	1.049 (0.894)	-2.401 (0.993)	1.737 (0.766)	-2.945 (0.836)	0.063 (1.001)	-1.617 (0.806)	3.366 (0.960)	-0.785 (0.762)	-1.369 (0.800)
$n = 250$	- 2.479 (0.569)	1.098 (0.445)	-2.243 (0.481)	1.733 (0.539)	-3.035 (0.650)	0.312 (0.507)	-1.563 (0.527)	3.374 (0.627)	-0.709 (0.542)	-1.259 (0.486)
<i>true values</i>	0.7	1	-2	1.5	-2.7	0.2	-1.4	3	-0.6	-1.2
$n = 100$	- 2.925 (0.719)	1.361 (0.658)	-2.386 (0.754)	1.914 (0.619)	-3.466 (0.861)	0.287 (0.570)	-1.736 (0.697)	3.788 (0.832)	-0.752 (0.606)	-1.539 (0.680)
$n = 250$	- 2.573 (0.475)	1.098 (0.370)	-2.230 (0.398)	1.686 (0.459)	-3.072 (0.528)	0.179 (0.310)	-1.574 (0.385)	3.369 (0.505)	-0.667 (0.313)	-1.348 (0.384)
<i>true values</i>	1.5	1	-2	1.5	-2.7	0.2	-1.4	3	-0.6	-1.2
$n = 100$	- 3.126 (0.678)	1.401 (0.583)	-2.903 (0.691)	2.049 (0.657)	-3.784 (0.767)	0.311 (0.528)	-1.977 (0.671)	4.220 (0.918)	-0.874 (0.554)	-1.679 (0.552)
$n = 250$	- 2.816 (0.549)	1.180 (0.321)	-2.440 (0.517)	1.810 (0.390)	-3.265 (0.635)	0.248 (0.264)	-1.735 (0.376)	3.623 (0.620)	-0.680 (0.276)	-1.428 (0.378)
<i>true values</i>	3	1	-2	1.5	-2.7	0.2	-1.4	3	-0.6	-1.2
$n = 100$	- 3.442 (0.672)	1.600 (0.607)	-3.048 (0.739)	2.171 (0.577)	-4.101 (0.764)	0.227 (0.554)	-2.225 (0.673)	4.558 (0.870)	-1.021 (0.517)	-1.779 (0.517)
$n = 250$	- 3.005 (0.609)	1.236 (0.297)	-2.594 (0.546)	1.954 (0.455)	-3.468 (0.750)	0.291 (0.263)	-1.833 (0.483)	3.864 (0.734)	-0.795 (0.309)	-1.536 (0.359)

Table 3: PG sampler. Case with p known: posterior means over the 100 different simulated datasets for n equal to 100 and 250 with true values of $p = 0.3$ (first panel), $p = 0.7$ (second panel), $p = 1.5$ (third panel) and $p = 3$ (fourth panel) and $(\beta_0, \beta_1, \beta_2, \beta_3, \beta_4, \beta_5, \beta_6, \beta_7, \beta_8, \beta_9) = (2.3, 1, -2, 1.5, -2.7, 0.2, -1.4, 3, -0.6, -1.2)$. The values in brackets are the standard deviations over the 100 different simulations.

p	β_0	β_1	β_2	β_3	β_4	β_5	β_6	β_7	β_8	β_9
<i>true values</i>	0.3	1	-2	1.5	-2.7	0.2	-1.4	3	-0.6	-1.2
$n = 100$	- 2.868 (1.046)	1.058 (1.282)	-2.617 (1.527)	1.379 (1.381)	-3.035 (1.500)	-0.178 (1.488)	-1.687 (1.493)	3.656 (1.619)	-0.920 (1.261)	-0.980 (1.463)
$n = 250$	- 2.875 (1.239)	1.221 (1.206)	-2.495 (1.137)	1.843 (1.333)	-3.235 (1.579)	0.068 (1.199)	-1.880 (1.343)	3.693 (1.231)	-0.772 (1.145)	-1.116 (1.313)
<i>true values</i>	0.7	1	-2	1.5	-2.7	0.2	-1.4	3	-0.6	-1.2
$n = 100$	- 3.145 (1.819)	1.089 (1.674)	-2.975 (1.816)	1.874 (1.898)	-3.429 (1.901)	0.168 (1.753)	-1.879 (1.852)	3.598 (2.011)	-0.818 (1.844)	-1.497 (1.910)
$n = 250$	- 2.855 (1.508)	1.150 (1.501)	-2.331 (1.389)	1.716 (1.585)	-3.227 (1.783)	0.149 (1.501)	-1.674 (1.441)	3.445 (1.612)	-0.750 (1.431)	-1.421 (1.451)
<i>true values</i>	1.5	1	-2	1.5	-2.7	0.2	-1.4	3	-0.6	-1.2
$n = 100$	- 4.270 (4.033)	0.988 (2.160)	-1.849 (2.320)	1.139 (2.242)	-1.806 (2.781)	-0.181 (2.282)	-0.929 (2.552)	2.260 (2.663)	-0.662 (2.032)	-0.826 (1.611)
$n = 250$	- 3.111 (1.083)	1.244 (1.159)	-2.424 (1.272)	1.595 (1.441)	-3.861 (1.571)	0.409 (1.407)	-1.883 (1.422)	4.051 (1.422)	-0.184 (1.264)	-1.593 (1.186)
<i>true values</i>	3	1	-2	1.5	-2.7	0.2	-1.4	3	-0.6	-1.2
$n = 100$	- 2.823 (3.488)	0.301 (2.364)	-0.607 (2.492)	0.790 (2.319)	-0.966 (2.964)	0.211 (2.425)	-0.270 (2.510)	0.907 (2.449)	0.015 (2.365)	-0.125 (2.791)
$n = 250$	- 2.933 (1.576)	1.079 (1.431)	-1.769 (1.929)	1.670 (1.660)	-2.906 (1.684)	0.271 (1.687)	-1.350 (1.740)	3.035 (1.968)	-0.830 (1.352)	-1.172 (1.888)

Table 4: BC_{el} sampler. Case with p known: posterior means over the 100 different simulated datasets for n equal to 100 and 250 with true values of $p = 0.3$ (first panel), $p = 0.7$ (second panel), $p = 1.5$ (third panel) and $p = 3$ (fourth panel) and $(\beta_0, \beta_1, \beta_2, \beta_3, \beta_4, \beta_5, \beta_6, \beta_7, \beta_8, \beta_9) = (2.3, 1, -2, 1.5, -2.7, 0.2, -1.4, 3, -0.6, -1.2)$. The values in brackets are the standard deviations over the 100 different simulations.

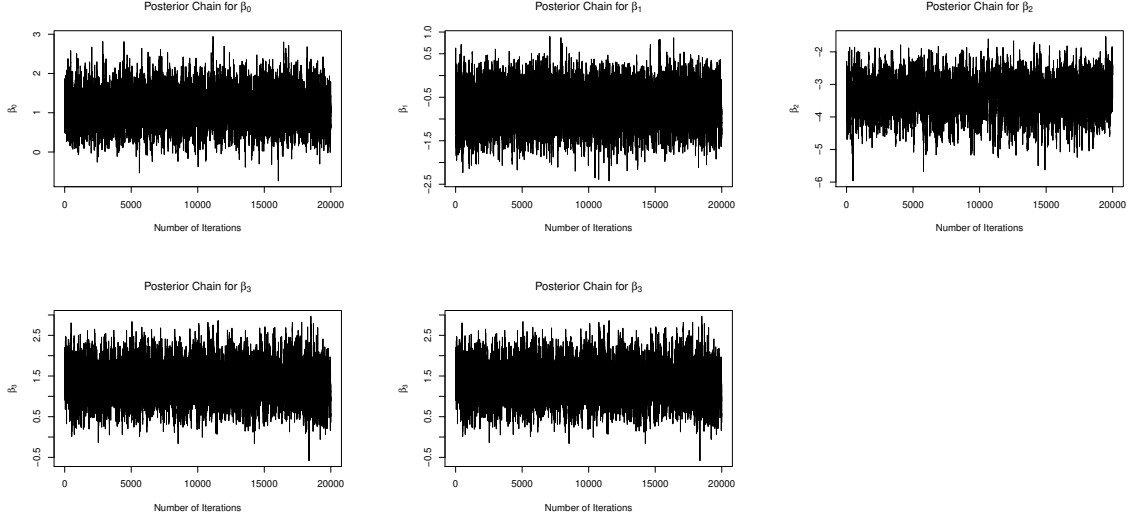


Figure 2: Case with p known: sample chains of the parameters posterior distributions for the simulated data with sample size $n = 250$ and with true parameter values $p = 0.3$ and $\boldsymbol{\beta} = (\beta_0, \beta_1, \beta_2, \beta_3, \beta_4) = (1, -1, -3, 1, 3)$.

simulations. We notice that the posterior means of p are close to their true values when p is less than 1, corresponding to the “heavy tail” situation, but the estimation accuracy drops when p is greater than 1. This similar performance is also confirmed for the posterior means of the vector of unknown coefficients $\boldsymbol{\beta} = (\beta_0, \beta_1, \beta_2, \beta_3, \beta_4)$. Generally, increasing the sample size from 100 to 250 leads to better estimations of the model parameters. Table 8 shows the posterior means and, in brackets, the standard deviations over the 100 different simulated datasets, for the ten-dimensional scenario. These results are in line with the five-dimensional case, thus leading us to conclude that increasing the number of covariates is not an issue in terms of estimation accuracy. Moreover, most standard deviations in Tables 6 and 8 decrease for the different scenarios and p values when we move from $n = 100$ to $n = 250$.

We further compare the performance of our Pólya-Gamma sampler to the one of BC_{el} . Table 7 and Table 9 shows the posterior means of p and $\boldsymbol{\beta}$, and the standard deviations over 100 different

	Variables	Test Statistic		Variables	Test Statistic
$n = 100$	$\beta_0 = 1$	-0.8422	$n = 250$	$\beta_0 = 1$	-0.0932
	$\beta_1 = -1$	-0.6155		$\beta_1 = -1$	0.9662
	$\beta_2 = -3$	0.2379		$\beta_2 = -3$	0.4544
	$\beta_3 = 1$	-0.2488		$\beta_3 = 1$	0.4355
	$\beta_4 = 3$	0.4159		$\beta_4 = 3$	-0.3423

Table 5: Case with p known: Geweke’s test statistics for the posterior chain of the unknown coefficients, $\boldsymbol{\beta} = (\beta_0, \beta_1, \beta_2, \beta_3, \beta_4)$ for a simulated dataset of sample size $n = 100$ (left panel) and $n = 250$ (right panel) with $p = 0.3$.

simulated datasets. As in the scenario with p known, we choose a much less vaguer normal prior distribution $\mathcal{N}(\boldsymbol{\nu} = 0, B = 9 \cdot I_k)$ for $\boldsymbol{\beta}$. Results are based on 20,000 posterior samples obtained with BC_{el} . It can be easily seen that when $p < 1$, the Pólya-Gamma sampler outperforms the BC_{el} substantially, in terms of both the posterior means and standard deviations across different datasets. The advantage of Pólya-Gamma sampler against BC_{el} reduces when $p > 1$, however, the standard deviations still suggest that Pólya-Gamma sampler is more consistent. As done previously with p known, here, with p unknown, we display in Figure 3 the posterior chains for a specific dataset with sample size $n = 250$, where $p = 0.3$ and $\boldsymbol{\beta} = (\beta_0, \beta_1, \beta_2, \beta_3, \beta_4) = (1, -1, -3, 1, 3)$. Despite exhibiting slightly slower mixing behaviours compared to Figure 2, the chains, after a burn-in of 1,000 iterations, converge relatively quickly to the true values of all parameters. We also computed Geweke’s convergence tests for the values of p and five-dimensional $\boldsymbol{\beta}$, with $n = 100$ and $n = 250$ and $p = 0.3^2$. Table 10 reports the results of Geweke’s convergence test (Geweke, 1992), that indicate no issues since the absolute values of the test statistics are all lower than 1.96.

	p	β_0	β_1	β_2	β_3	β_4
<i>true values</i>	0.3	1	-1	-3	1	3
$n = 100$	0.271 (0.096)	1.400 (0.815)	-1.311 (0.667)	-3.882 (0.629)	1.330 (0.828)	3.946 (0.689)
$n = 250$	0.230 (0.056)	1.439 (0.553)	-1.480 (0.583)	-4.341 (0.565)	1.442 (0.599)	4.489 (0.635)
<i>true values</i>	0.7	1	-1	-3	1	3
$n = 100$	0.798 (0.174)	0.995 (0.320)	-0.976 (0.394)	-2.991 (0.383)	1.028 (0.393)	2.884 (0.384)
$n = 250$	0.821 (0.214)	0.927 (0.271)	-0.959 (0.283)	-3.026 (0.345)	0.966 (0.287)	3.117 (0.345)
<i>true values</i>	1.5	1	-1	-3	1	3
$n = 100$	1.129 (0.351)	1.388 (0.382)	-1.390 (0.405)	-4.335 (0.526)	1.427 (0.528)	4.351 (0.490)
$n = 250$	1.164 (0.265)	1.465 (0.318)	-1.436 (0.307)	-4.381 (0.579)	1.490 (0.382)	4.444 (0.557)
<i>true values</i>	3	1	-1	-3	1	3
$n = 100$	1.593 (0.387)	1.506 (0.415)	-1.576 (0.424)	-4.697 (0.563)	1.607 (0.528)	4.736 (0.608)
$n = 250$	1.691 (0.412)	1.570 (0.267)	-1.607 (0.324)	-4.778 (0.587)	1.560 (0.289)	4.740 (0.586)

Table 6: PG sampler. Case with p unknown: posterior means over the 100 different simulated datasets for n equal to 100 and 250 compared with the true values of $p = 0.3$ (first panel), $p = 0.7$ (second panel), $p = 1.5$ (third panel) and $p = 3$ (fourth panel) and $(\beta_0, \beta_1, \beta_2, \beta_3, \beta_4) = (1, -1, -3, 1, 3)$. The values in brackets are the standard deviations over the 100 different simulations.

	p	β_0	β_1	β_2	β_3	β_4
<i>true values</i>	0.3	1	-1	-3	1	3
$n = 100$	0.559 (0.520)	0.258 (0.440)	-0.264 (0.394)	-0.868 (0.577)	0.345 (0.480)	0.930 (0.631)
$n = 250$	0.149 (0.269)	0.166 (0.507)	-0.165 (0.445)	-0.427 (0.603)	0.139 (0.484)	0.398 (0.548)
<i>true values</i>	0.7	1	-1	-3	1	3
$n = 100$	0.807 (0.698)	0.435 (0.507)	-0.576 (0.438)	-1.991 (0.570)	0.498 (0.491)	1.579 (0.581)
$n = 250$	0.754 (0.214)	0.483 (0.471)	-0.591 (0.414)	-2.076 (0.610)	0.514 (0.437)	1.618 (0.519)
<i>true values</i>	1.5	1	-1	-3	1	3
$n = 100$	2.464 (1.378)	0.563 (0.558)	-0.617 (0.494)	-1.626 (0.489)	0.595 (0.441)	1.716 (0.556)
$n = 250$	1.877 (1.687)	0.271 (1.013)	-0.483 (0.629)	-1.619 (0.774)	0.377 (0.583)	1.518 (0.787)
<i>true values</i>	3	1	-1	-3	1	3
$n = 100$	3.056 (1.204)	0.632 (0.630)	-0.391 (0.601)	-1.385 (0.857)	0.508 (0.551)	1.433 (0.945)
$n = 250$	2.641 (1.759)	0.478 (0.750)	-0.551 (0.705)	-1.570 (0.728)	0.598 (0.594)	1.634 (0.723)

Table 7: BC_{el} sampler. Case with p unknown: posterior means over the 100 different simulated datasets for n equal to 100 and 250 compared with the true values of $p = 0.3$ (first panel), $p = 0.7$ (second panel), $p = 1.5$ (third panel) and $p = 3$ (fourth panel) and $(\beta_0, \beta_1, \beta_2, \beta_3, \beta_4) = (1, -1, -3, 1, 3)$. The values in brackets are the standard deviations over the 100 different simulations.

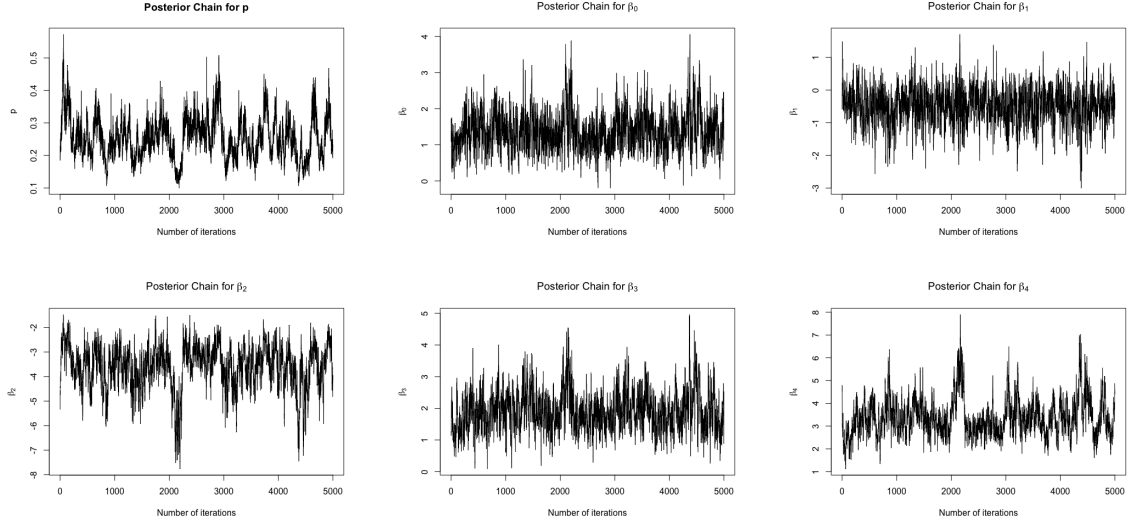


Figure 3: Case with p unknown: sample chains of the posterior distribution of the parameters for the simulated data with sample size $n = 250$ and with true parameter values $p = 0.3$ and $\boldsymbol{\beta} = (\beta_0, \beta_1, \beta_2, \beta_3, \beta_4) = (1, -1, -3, 1, 3)$.

5 Real Data Applications

In this Section, we present two different real data applications to illustrate the accurate and efficient performance of the Pólya-Gamma sampler for estimating the generalized logistic regression. The applications analyse novel datasets of different dimensions, which are related to current and topical problems in the social sciences. In the first application, we consider data selected from the European Social Survey (ESS), that were previously analysed in Dalla Valle et al. (2019) to identify the determinants of public opinions towards immigration. The second application analyses the motives behind the recidivism of offenders released from prison

²Results for the ten-dimensional case are omitted for lack of space but are available on request.

p	β_0	β_1	β_2	β_3	β_4	β_5	β_6	β_7	β_8	β_9
<i>true values</i>	0.3	1	-2	1.5	-2.7	0.2	-1.4	3	-0.6	-1.2
$n = 100$	0.197 (0.063)	3.440 (0.735)	1.395 (1.221)	-3.179 (1.128)	2.271 (0.903)	-3.667 (0.868)	0.097 (1.126)	-2.093 (1.025)	4.356 (1.031)	-1.044 (0.948)
$n = 250$	0.186 (0.040)	3.685 (0.696)	1.656 (0.648)	-3.341 (0.657)	2.563 (0.706)	-4.497 (0.678)	0.447 (0.757)	-2.309 (0.667)	4.982 (0.722)	-1.070 (0.839)
<i>true values</i>	0.7	2.3	1	-2	1.5	-2.7	0.2	-1.4	3	-0.6
$n = 100$	0.658 (0.197)	2.517 (0.644)	1.132 (0.673)	-2.213 (0.550)	1.739 (0.543)	-3.016 (0.465)	0.202 (0.575)	-1.608 (0.617)	3.264 (0.654)	-0.713 (0.642)
$n = 250$	0.657 (0.138)	2.653 (0.306)	1.186 (0.443)	-2.286 (0.448)	1.883 (0.435)	-3.172 (0.396)	0.248 (0.375)	-1.593 (0.289)	3.410 (0.494)	-0.693 (0.285)
<i>true values</i>	1.5	2.3	1	-2	1.5	-2.7	0.2	-1.4	3	-0.6
$n = 100$	0.956 (0.329)	3.678 (0.561)	1.657 (0.611)	-3.351 (0.548)	2.397 (0.665)	-4.363 (0.672)	0.316 (0.650)	-2.286 (0.697)	4.876 (0.680)	-1.014 (0.645)
$n = 250$	0.986 (0.263)	3.818 (0.430)	1.611 (0.372)	-3.306 (0.407)	2.460 (0.408)	-4.423 (0.486)	0.339 (0.358)	-2.362 (0.385)	4.913 (0.428)	-0.921 (0.361)
<i>true values</i>	3	2.3	1	-2	1.5	-2.7	0.2	-1.4	3	-0.6
$n = 100$	1.471 (0.378)	3.804 (0.533)	1.731 (0.590)	-3.314 (0.595)	2.389 (0.646)	-4.510 (0.603)	0.308 (0.568)	-2.497 (0.595)	5.029 (0.674)	-1.156 (0.517)
$n = 250$	1.488 (0.430)	4.008 (0.540)	1.656 (0.358)	-3.449 (0.480)	2.595 (0.441)	-4.600 (0.603)	0.387 (0.338)	-2.426 (0.434)	5.146 (0.655)	-1.042 (0.326)

Table 8: PG sampler. Case with p unknown: posterior means over the 100 different simulated datasets for n equal to 100 and 250 compared with the true values of $p = 0.3$ (first panel), $p = 0.7$ (second panel), $p = 1.5$ (third panel) and $p = 3$ (fourth panel) and $(\beta_0, \beta_1, \beta_2, \beta_3, \beta_4, \beta_5, \beta_6, \beta_7, \beta_8, \beta_9) = (2.3, 1, -2, 1.5, -2.7, 0.2, -1.4, 3, -0.6, -1.2)$. The values in brackets are the standard deviations over the 100 different simulations.

	p	β_0	β_1	β_2	β_3	β_4	β_5	β_6	β_7	β_8	β_9
<i>true values</i>	0.3	2.3	1	-2	1.5	-2.7	0.2	-1.4	3	-0.6	-1.2
$n = 100$	0.149 (0.164)	0.554 (0.590)	0.300 (0.608)	-0.697 (0.831)	0.321 (0.568)	-0.690 (0.613)	0.207 (0.582)	-0.412 (0.657)	0.844 (0.629)	-0.232 (0.671)	-0.348 (0.609)
$n = 250$	0.056 (0.085)	0.409 (0.535)	0.128 (0.402)	-0.334 (0.470)	0.242 (0.497)	-0.478 (0.560)	0.040 (0.502)	-0.184 (0.583)	0.493 (0.566)	-0.208 (0.663)	-0.178 (0.485)
<i>true values</i>	0.7	2.3	1	-2	1.5	-2.7	0.2	-1.4	3	-0.6	-1.2
$n = 100$	0.885 (0.286)	0.847 (0.604)	0.385 (0.649)	-0.573 (0.541)	0.418 (0.589)	-0.650 (0.494)	0.120 (0.580)	-0.401 (0.682)	0.710 (0.697)	-0.258 (0.657)	-0.307 (0.662)
$n = 250$	0.641 (0.288)	0.896 (0.708)	0.401 (0.749)	-0.580 (0.701)	0.515 (0.794)	-0.671 (0.755)	0.158 (0.801)	-0.534 (0.583)	0.790 (0.651)	-0.277 (0.497)	-0.401 (0.724)
<i>true values</i>	1.5	2.3	1	-2	1.5	-2.7	0.2	-1.4	3	-0.6	-1.2
$n = 100$	2.228 (2.090)	0.955 (0.811)	0.397 (0.773)	-0.454 (1.079)	0.434 (0.983)	-0.631 (1.121)	-0.025 (0.980)	-0.382 (0.926)	0.617 (1.363)	-0.257 (0.743)	-0.147 (0.986)
$n = 250$	1.118 (1.543)	1.191 (0.758)	0.423 (0.885)	-0.738 (0.948)	0.595 (0.939)	-0.773 (1.030)	0.085 (0.970)	-0.525 (0.976)	1.284 (1.111)	-0.303 (0.958)	-0.368 (0.925)
<i>true values</i>	3	2.3	1	-2	1.5	-2.7	0.2	-1.4	3	-0.6	-1.2
$n = 100$	3.337 (1.909)	0.567 (1.020)	0.055 (1.182)	-0.004 (1.252)	0.149 (1.053)	0.002 (1.269)	-0.028 (0.983)	-0.198 (1.246)	0.005 (1.276)	-0.018 (0.976)	-0.009 (1.210)
$n = 250$	2.041 (1.912)	0.995 (0.827)	0.148 (1.118)	-0.376 (1.227)	0.258 (1.087)	-0.379 (1.294)	0.186 (0.891)	-0.249 (1.070)	0.617 (1.231)	-0.110 (0.932)	-0.216 (1.144)

Table 9: BC_{el} sampler. Case with p unknown: posterior means over the 100 different simulated datasets for n equal to 100 and 250 compared with the true values of $p = 0.3$ (first panel), $p = 0.7$ (second panel), $p = 1.5$ (third panel) and $p = 3$ (fourth panel) and $(\beta_0, \beta_1, \beta_2, \beta_3, \beta_4, \beta_5, \beta_6, \beta_7, \beta_8, \beta_9) = (2.3, 1, -2, 1.5, -2.7, 0.2, -1.4, 3, -0.6, -1.2)$. The values in brackets are the standard deviations over the 100 different simulations.

	Variables	Test Statistic		Variables	Test Statistic
$n = 100$	$p = 0.3$	-0.507	$n = 250$	$p = 0.3$	0.914
	$\beta_0 = 1$	1.721		$\beta_0 = 1$	-1.394
	$\beta_1 = -1$	-1.308		$\beta_1 = -1$	1.320
	$\beta_2 = -3$	-1.293		$\beta_2 = -3$	1.321
	$\beta_3 = 1$	1.413		$\beta_3 = 1$	-1.545
	$\beta_4 = 3$	0.539		$\beta_4 = 3$	-1.259

Table 10: Case with p unknown: Geweke’s test statistics for the posterior chain of the tail parameter p and of the unknown coefficients, $\boldsymbol{\beta} = (\beta_0, \beta_1, \beta_2, \beta_3, \beta_4)$ for a simulated dataset of sample size $n = 100$ (left panel) and $n = 250$ (right panel).

between 2016 and 2018 in the State of Iowa, USA.

5.1 Immigration in European Countries

The first dataset focuses on people’s attitudes towards immigration. The data was selected from the ESS and was collected following hour-long face-to-face interviews. We focused on questions regarding immigration from the ESS8 edition 1.0 published in October 2017 and collected between August 2016 and March 2017 (ESS8, 2016, 2018) for Great Britain (GB), Germany (DE) and France (FR), respectively. For these three countries, the total number of observations is 5,354, of which 1,419 are for GB; 2,284 for DE and 1,651 for FR.

The dependent variable is `immig`, that indicates whether the respondent would allow immigrants from poorer countries outside Europe. More precisely, if `immig` = 1 the respondent is against immigration, if `immig` = 0 the respondent is in favor of immigration. Following Dalla Valle et al. (2019), we considered as covariates the variables: `pplfair`, `trstep`, `trstun`, `happy`, `agea`, `edulvlb` and `hinctnta`. The first three variables include answers to question

ranging from 0 (most negative) to 10 (most positive): do you think that most people try to take advantage of you (**pplfair**)? Do you trust the European parliament (**trstep**)? Do you trust the United Nations (**trstun**)? The **happy** variable takes values from 0 (unhappy) to 10 (extremely happy). The last three variables are subject-specific and include the age, ranging from 15 to 100 (**agea**); the highest level of education, from primary education to doctoral degree (**edulvlb**) and the household’s total net income (**hinctnta**).

For this dataset, we run 5,000 iterations of the Pólya-Gamma sampler algorithm and we discarded the first 1,000 iterations as burn-in. We used the MLE estimates of β and p as starting points of the algorithm and we employed a normal prior with mean 0, standard deviation 5 times greater than the one estimated with a standard logistic regression.

In order to assess the performance of the Pólya-Gamma sampler, we report the results obtained using the empirical likelihood approach on the same dataset, adopting a similar setting. Table 11 lists the results of the Pólya-Gamma sampler (left-hand side) and the empirical likelihood approach (right-hand side). For each European country (GB, DE and FR) and for each parameter, Table 11 shows the posterior means, the 2.5% and 97.5% quantiles of the posterior distributions of the Pólya-Gamma and the empirical likelihood methods. The effects of the parameters are similar across countries and suggest that the higher the level of education, the higher the probability for people to be favourable towards immigration. Also, trustful people and individuals with a strong confidence in the European parliament tend to be in favour of immigration. The posterior mean estimates obtained with the Pólya-Gamma method are in line with those obtained with the empirical likelihood. For some of the covariates, such as **pplfair** and **trstep**, the empirical likelihood approach assigns posterior support to the value zero, while the corresponding Pólya-Gamma credible intervals (CIs) do not include zero. In addition, both methods agree on the estimation of the p parameter for all countries, suggesting heavy tails for the generalised logistic regression. Therefore, the Pólya-Gamma sampler outperforms the

empirical likelihood method for the estimation of the generalised logistic regression, since it provides more precise parameter estimates as for the β .

5.2 Recidivism of Offenders Released from Prison in Iowa

The second dataset contains information about the recidivism of offenders released from prison between 2016 and 2018 in Iowa³. Recidivism happens when a prisoner is released from jail and he or she relapses back into criminal behavior and returns to jail. We applied the generalised logistic regression to determine whether a prisoner is likely to recidivate based on his or her characteristics.

After removing missing values, the data consists of 13,644 records corresponding to offenders detained in prison and released from 2016 to 2018. The dependent variable is **recidivism**, which is equal to 1 if the prisoner is arrested again within a 3-year tracking period after being released, and it is equal to 0 if he or she has not returned to prison within the tracking period.

We model the probability of being recidivist using a subsample of all the explanatory variables present in the original database: the sex of the offender (**sex**, which is equal to 1 for men and 0 for women); the age when released from prison (**age**); whether an offender committed a felony or misdemeanor (**felony**, which is equal to 1 for felony and 0 for misdemeanor); whether an offender was charged with a drug crime (**drug**); whether an offender was charged with a public order crime (**puborder**); whether an offender was charged with a violent crime (**violent**); and whether an offender was released on discharge or parole (**discharge**, which is equal to 1 for discharge and 0 for parole).

In this second example, we run the MCMC algorithm using 5,000 iterations and discarding

³See the following website for details: <https://data.iowa.gov/Correctional-System/3-Year-Recidivism-for-Offenders-Released-from-Pris/mw8r-vqy4>.

Country	Parameters	Pólya-Gamma Sampler			Empirical Likelihood		
		post. mean	2.50%	97.50%	post. mean	2.50%	97.50%
GB	const	-0.8	-2.441	0.541	-0.4202	-2.4377	1.4487
	pplfair	-0.223	-0.552	-0.068	-0.1822	-0.3176	0.0113
	trstep	-0.362	-0.888	-0.123	-0.1731	-0.3796	0.0447
	trstun	-0.075	-0.302	0.046	-0.0372	-0.2354	0.1409
	happy	0.047	-0.099	0.241	-0.0133	-0.2892	0.1209
	agea	0.015	-0.001	0.04	0.0094	-0.0139	0.0358
	edulvlb	-0.004	-0.009	-0.001	-0.0021	-0.0039	-0.0006
	hinctnta	0.017	-0.105	0.148	-0.044	-0.1856	0.1228
	p	0.815	0.2	1.944	0.8218	0.4284	1.5696
DE	const	-1.152	-2.648	0.29	-1.1452	-2.4617	0.6941
	pplfair	-0.301	-0.616	-0.114	-0.1612	-0.3291	-0.0093
	trstep	-0.429	-1	-0.148	-0.1524	-0.2839	-0.0107
	trstun	-0.091	-0.33	0.085	-0.0112	-0.2115	0.1847
	happy	-0.169	-0.46	0.003	-0.0873	-0.2438	0.0659
	agea	0.036	0.013	0.077	0.0124	-0.0061	0.0375
	edulvlb	0	-0.004	0.002	-0.0007	-0.0026	0.0016
	hinctnta	-0.2	-0.533	-0.042	-0.1171	-0.2352	0.0414
	p	0.608	0.155	1.814	0.7338	0.2955	1.7036
FR	const	1.054	-0.35	3.361	0.8442	-0.3443	1.4314
	pplfair	-0.111	-0.347	0.01	-0.0921	-0.1898	0.0193
	trstep	-0.364	-0.872	-0.107	-0.1912	-0.3592	-0.0114
	trstun	-0.122	-0.353	0.004	-0.0514	-0.2047	0.1417
	happy	0.029	-0.129	0.227	0.0191	-0.0609	0.1545
	agea	0.001	-0.013	0.021	-0.0043	-0.0155	0.0091
	edulvlb	-0.005	-0.012	-0.002	-0.0038	-0.0056	-0.0027
	hinctnta	-0.092	-0.275	0.016	-0.1061	-0.239	0.0286
	p	0.827	0.159	2.055	0.7845	0.3991	1.2809

Table 11: EU immigration data results for Great Britain (GB), Germany (DE) and France (FR), obtained with the Pólya-Gamma (left-hand side) and the empirical likelihood (right-hand side) methods. For each parameter, columns 3 and 6 show the posterior means, columns 4 and 7 show the 2.5% quantiles of the posterior distribution, columns 5 and 8 show the 97.5% quantiles of the posterior distribution.

the first 1,000 iterations. As in the immigration example, we adopted vague priors and we chose the initial values of β and p based on the corresponding MLE estimators.

	Pólya-Gamma Sampler			Empirical Likelihood		
Parameters	post. mean	2.50%	97.50%	post. mean	2.50%	97.50%
const	-0.1443	-0.3811	0.0232	-0.4278	-0.9341	0.1719
sex	0.311	0.1468	0.5873	0.1371	-0.2583	0.4574
age	-0.0004	-0.0007	-0.0002	-0.0003	-0.0005	-0.0002
felony	-0.0941	-0.2521	0.0184	-0.0974	-0.4457	0.0999
drug	-0.0308	-0.1496	0.0774	-0.0996	-0.4009	0.1912
puborder	-0.2769	-0.5663	-0.116	-0.0694	-0.3384	0.2261
violent	-0.5759	-1.0713	-0.3271	-0.4743	-0.9779	-0.2638
discharge	-0.5736	-1.068	-0.3225	-0.4654	-0.6726	-0.1936
p	0.9343	0.3088	1.8439	0.5679	0.3409	1.0905

Table 12: Criminal recidivism data results obtained with the Pólya-Gamma (left-hand side) and the empirical likelihood (right-hand side) methods. For each parameter, columns 2 and 5 show the posterior means, columns 3 and 6 show the 2.5% quantiles of the posterior distribution, columns 4 and 7 show the 97.5% quantiles of the posterior distribution.

As we did with the immigration data, we compared the performances of the Pólya-Gamma sampler and the empirical likelihood approaches, using the same settings for both algorithms. The results for the parameters β and p are listed in Table 12, that displays the Pólya-Gamma results on the left-hand side and the empirical likelihood results on the right-hand side. Table 12 illustrates the posterior means, the 2.5% and 97.5% quantiles of the posterior distributions of the Pólya-Gamma and the empirical likelihood methods.

The results suggest that men are more likely to recidivate than women, younger prisoners are more likely to become repeat offenders than older prisoners, there is a lower probability that the criminal is a recidivist if he or she was charged with a public order or violent crime and offenders released on parole are more likely to return to jail than those released on discharge.

The sign of the posterior means for both approaches is the same, suggesting a similar direction for the effects obtained using the Pólya-Gamma and the empirical likelihood method. However, the CIs of the first approach are narrower than those of second approach. Moreover, the Pólya-Gamma method, unlike the empirical likelihood, does not give posterior support to zero to the variables **sex** and **puborder**, suggesting that the inclusion of these variables in the model is appropriate.

In agreement with our findings with the immigration data, the results obtained with the recidivism dataset demonstrate that the Pólya-Gamma sampler method yields a higher estimates precision and accuracy compared to the empirical likelihood method for the covariate parameters β .

6 Conclusions

This paper introduces a novel DA scheme for the generalized logistic regression model. This model is particularly flexible since it is able to accommodate both light and heavy tails in dichotomous response data. The proposed DA scheme makes use of the Pólya-Gamma identity and it is strongly related to the slice sampler algorithm. The Pólya-Gamma sampler allows us to implement a Bayesian method able to draw samples from the exact posterior distribution. On the contrary, other methods, such as the empirical likelihood approach, are based on approximations of the posterior and may lead to low precision in the estimation of the parameters. Our simulation study demonstrates the estimation accuracy of the newly proposed Pólya-Gamma sampler with datasets of different dimensions. We compared the performances of the Pólya-Gamma and the empirical likelihood methods for modelling new interesting datasets regarding the opinion on immigration in European countries and the probability of being recidivist for prisoners in Iowa, USA. Our results demonstrate the superiority of the Pólya-Gamma sampler

over the empirical likelihood in terms of parameter precision. The Pólya-Gamma method allows a more accurate estimation of the tail parameter for the generalised logistic regression compared to approximate methods.

References

- Albert, J. H. and S. Chib (1993, 06). Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association* 88(422), 669–679.
- Brooks, S., A. Gelman, G. Jones, and X.-L. Meng (2011, 01). Handbook of markov chain monte carlo.
- Choi, H. M. and J. P. Hobert (2013). The polya-gamma gibbs sampler for bayesian logistic regression is uniformly ergodic. *Electron. J. Statist.* 7, 2054–2064.
- Dalla Valle, L., F. Leisen, L. Rossini, and W. Zhu (2019). Bayesian analysis of immigration in europe with generalized logistic regression. *Journal of Applied Statistics* 0(0), 1–15.
- ESS8 (2016). ESS Round 8: European Social Survey Round 8 Data. *Data file edition 2.0. NSD - Norwegian Centre for Research Data, Norway ??? Data Archive and distributor of ESS data for ESS ERIC*.
- ESS8 (2018). ESS Round 8: European Social Survey. *ESS-8 2016 Documentation Report. Edition 2.0. Bergen, European Social Survey Data Archive, NSD - Norwegian Centre for Research Data for ESS ERIC*.
- Geweke, J. (1992). Evaluating the accuracy of sampling-based approaches to calculating posterior moments. In J. M. Bernardo, J. Berger, A. P. Dawid, and J. F. M. Smith (Eds.), *Bayesian Statistics 4*, pp. 169–193. Oxford: Oxford University Press.

- Hensley, A. A. and P. M. Djurić (2017, March). Nonparametric learning for hidden markov models with preferential attachment dynamics. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 3854–3858.
- Hobert, J. P. and D. Marchev (2008). A theoretical comparison of the data augmentation, marginal augmentation and px-da algorithms. *Ann. Statist.* *36*(2), 532–554.
- Holmes, C. C., L. Held, et al. (2006). Bayesian auxiliary variable models for binary and multinomial regression. *Bayesian analysis* *1*(1), 145–168.
- Karabatsos, G. and F. Leisen (2018). An approximate likelihood perspective on abc methods. *Statist. Surv.* *12*, 66–104.
- Leisen, F., L. Rossini, and C. Villa (2017, 04). A note on the posterior inference for the yule–simon distribution. *Journal of Statistical Computation and Simulation* *87*(6), 1179–1188.
- Liu, J. and Y. Wu (1999, 06). Parameter expansion for data augmentation. *Journal of the American Statistical Association* *94*.
- Meng, X.-L. and D. van Dyk (1999, 7/11/2019). Seeking efficient data augmentation schemes via conditional and marginal augmentation. *Biometrika* *86*(2), 301–320.
- Mengersen, K. L., P. Pudlo, and C. P. Robert (2013). Bayesian computation via empirical likelihood. *Proceedings of the National Academy of Sciences* *110*(4), 1321–1326.
- Neal, R. M. (2003, 06). Slice sampling. *Ann. Statist.* *31*(3), 705–767.
- Plummer, M., N. Best, K. Cowles, and K. Vines (2006). CODA: Convergence Diagnosis and Output Analysis for MCMC. *R News* *6*(1), 7–11.

- Polson, N. G., J. G. Scott, and J. Windle (2013, 12). Bayesian inference for logistic models using pólya–gamma latent variables. *Journal of the American Statistical Association* 108(504), 1339–1349.
- Tanner, T. A. and W. H. Wong (1987). The calculation of posterior distributions by data augmentation. *Journal of the American Statistical Association* 82(398), 528–540.
- van Dyk, D. A. and X.-L. Meng (2001, 03). The art of data augmentation. *Journal of Computational and Graphical Statistics* 10(1), 1–50.
- Windle, J., N. G. Polson, and J. G. Scott (2016). BayesLogit: Bayesian Logistic Regression. Technical report, R Package Version 0.6.
- Zhu, W. and F. Leisen (2016). A bootstrap likelihood approach to bayesian computation. *Australian and New Zealand Journal of Statistics* 58(1), 227–244.