



Title	Bayesian Inference, Model Selection and Likelihood Estimation using Fast Rejection Sampling: The Conway-Maxwell-Poisson Distribution
Authors(s)	Benson, Alan, Friel, Nial
Publication date	2021-09
Publication information	Benson, Alan, and Nial Friel. "Bayesian Inference, Model Selection and Likelihood Estimation Using Fast Rejection Sampling: The Conway-Maxwell-Poisson Distribution." International Society for Bayesian Analysis, September 2021. https://doi.org/10.1214/20-ba1230 .
Publisher	International Society for Bayesian Analysis
Item record/more information	http://hdl.handle.net/10197/12594
Publisher's version (DOI)	10.1214/20-ba1230

Downloaded 2026-05-01 23:46:51

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)



© Some rights reserved. For more information

Bayesian inference, model selection and likelihood estimation using fast rejection sampling: the Conway-Maxwell-Poisson distribution

Alan Benson^{1,2,*} and Nial Friel^{1,2}

¹School of Mathematics and Statistics, University College Dublin

²Insight Centre for Data Analytics

*alan.benson@insight-centre.org

July 10, 2020

Abstract

Bayesian inference for models with intractable likelihood functions represents a challenging suite of problems in modern statistics. In this work we analyse the Conway-Maxwell-Poisson (COM-Poisson) distribution, a two parameter generalisation of the Poisson distribution. COM-Poisson regression modelling allows the flexibility to model dispersed count data as part of a generalised linear model (GLM) with a COM-Poisson response, where exogenous covariates control the mean and dispersion level of the response. The major difficulty with COM-Poisson regression is that the likelihood function contains multiple intractable normalising constants and is not amenable to standard inference and MCMC techniques. Recent work by [Chanialidis et al. \(2018\)](#) has seen the development of a sampler to draw random variates from the COM-Poisson likelihood using a rejection sampling algorithm. We provide a new rejection sampler for the COM-Poisson distribution which significantly reduces the CPU time required to perform inference for COM-Poisson regression models. An extension of this work shows that for any intractable likelihood function with an associated rejection sampler it is possible to construct unbiased estimators of the intractable likelihood which proves useful for model selection or for use within pseudo-marginal MCMC algorithms ([Andrieu and Roberts, 2009](#)). We demonstrate all of these methods on a real-world dataset of takeover bids.

1 Introduction

Modelling count data continues to be an important area in the practice of statistics. Often covariate information is available which may prove useful in explaining the counts, for example, time of day influencing the number calls to a particular call centre or amount borrowed influencing the number of loan defaults by a consumer. Poisson regression is a standard framework for modelling covariate-dependent count data allowing the mean response to depend on the covariates through a log-linear link function. In practice, Poisson regression fails to capture the phenomenon of dispersion, where the mean and variance of the response differ significantly. Dispersion measures the spread of a random variable and is quantified through the ratio of the variance of a random variable to its mean. Count data often exhibits underdispersion (variance less than the mean), overdispersion (variance greater than the mean) or equidispersion (equality of mean and variance) but Poisson regression is only capable of modelling equidispersed data. Many alternative approaches have been suggested to capture dispersion, such as using quasilielihood or the popular negative binomial regression ([Hilbe, 2011](#)) which allows modelling of overdispersed data only. The COM-Poisson distribution ([Conway and Maxwell, 1962](#)) was first introduced to model queuing systems and was then revived as a statistical model for dispersed

count data by [Shmueli et al. \(2005\)](#). This revival led to the introduction of COM-Poisson regression models by [Guikema and Goffelt \(2008\)](#) as a flexible regression model for dispersed count data with covariates. However, this flexibility comes with the caveat that the COM-Poisson likelihood function contains an intractable normalising term which requires the need for non-standard methods to conduct parameter inference. This intractability has restricted the popularity of COM-Poisson regression as a viable solution to modelling dispersion. Recent work by [Chaniyalidis et al. \(2018\)](#) has seen a Bayesian analysis of COM-Poisson regression under the framework of doubly-intractable Bayesian inference, where the authors overcome the intractability through the use of a rejection sampling algorithm for the COM-Poisson distribution. This rejection sampler is then used within the exchange algorithm of [Murray et al. \(2006\)](#) to perform Bayesian parameter inference.

Our novel contribution in this paper is to provide a more efficient approach to that of [Chaniyalidis et al. \(2018\)](#) by designing a simpler rejection sampler for the COM-Poisson distribution which is computationally less intensive to sample from and leads to faster inference within COM-Poisson regression models. In one example in this paper, we observe a greater than threefold reduction in CPU time by using our rejection sampler. The rejection sampler of [Chaniyalidis et al. \(2018\)](#) is derived from a discrete version of adaptive rejection sampling ([Gilks and Wild, 1992](#)) and requires the costly construction of a piecewise truncated geometric enveloping distribution to tightly encapsulate the COM-Poisson probability density function. Our rejection sampler requires only a single envelope distribution depending on the dispersion parameter of the COM-Poisson distribution. This single envelope distribution does come with a higher rejection rate but completely bypasses the costly setup and sampling cost involved in piecewise truncated geometric envelope distributions. A further and important novelty of our work is to show that when a rejection sampling algorithm is available to sample from an intractable likelihood function of interest, then it is possible to construct an unbiased estimate of this intractable likelihood function using the rejection sampler. Our idea works by exploiting a link between rejection sampling efficiency and its relationship with the reciprocal intractable normalising constant of the likelihood function.

The remainder of this paper is organised as follows: Section 2 reviews the COM-Poisson distribution and COM-Poisson regression models under the framework of doubly-intractable Bayesian inference. Section 3 reviews rejection sampling for intractable likelihoods and our new COM-Poisson rejection sampler is introduced in Section 3.1. Section 3.3 describes the novel method of constructing unbiased estimators of intractable likelihood functions through rejection sampling. Section 4 shows how pseudo-marginal MCMC can be performed using this unbiased intractable likelihood estimator. Section 5 shows results of these methods applied firstly to a dataset with no covariates in Section 5.1 and then to a real-world COM-Poisson regression dataset of company takeover bids in Section 5.2. Finally, Section 6 gives discussion and some thoughts for future work.

2 Bayesian doubly-intractable inference for GLMs

The COM-Poisson distribution ([Conway and Maxwell, 1962](#)) is a two-parameter discrete probability distribution that allows the flexibility to model count data with over-, under- and equi-dispersion. The probability mass function for a COM-Poisson random variable Y with parameters $\mu > 0$ and $\nu \geq 0$ is defined over the non-negative integers as

$$f(y|\mu, \nu) = \left(\frac{\mu^y}{y!}\right)^\nu \frac{1}{\mathcal{Z}_f(\mu, \nu)}.$$

The unnormalised component of the mass function is denoted $q_f(y|\mu, \nu) = \left(\frac{\mu^y}{y!}\right)^\nu$ and $\mathcal{Z}_f(\mu, \nu) = \sum_{y=0}^{\infty} q_f(y|\mu, \nu)$ is an intractable normalising constant, having no closed form representation for this model. The mode of the COM-Poisson distribution is $\lfloor \mu \rfloor$, with two consecutive modes at μ and $\mu - 1$ in

the case where μ is an integer. The moments of the COM-Poisson distribution are unavailable directly due to the intractable normalising constant. The moments can however be approximated through the use of an asymptotic representation of $\mathcal{Z}_f(\mu, \nu)$ derived by [Shmueli et al. \(2005\)](#). The approximations for the mean and variance using this asymptotic representation are, respectively,

$$E(Y) \approx \mu + \frac{1}{2\nu} - \frac{1}{2}, \quad \text{Var}(Y) \approx \frac{\mu}{\nu}. \quad (1)$$

These approximations are quite accurate for a wide range of μ and ν , except for small μ or small ν for the case of $E(Y)$ and for $\nu > 1$, regardless of the value of μ for the case of $\text{Var}(Y)$. The purpose of the parameter ν is to control dispersion through this parameter's inverse relationship with the variance as seen in (1). When $\nu = 1$ the distribution exhibits equidispersion and the probability mass function reduces to that of a Poisson random variable with $\mathcal{Z}_f(\mu, 1) = \exp(\mu)$. Alternatively setting $\nu < 1$ gives overdispersion and setting $\nu > 1$ gives underdispersion.

The main application of the COM-Poisson distribution is to GLMs, referred to as COM-Poisson regression models, where n responses $y_{1:n} = (y_1, \dots, y_n)$ are assumed to follow a COM-Poisson distribution with the parameters μ and ν of each response being conditional on b available covariates $x_i = (x_{i1}, \dots, x_{ib})$ for $i = 1, \dots, n$. Due to this conditioning we now index the parameters for observation y_i as μ_i and ν_i . COM-Poisson regression modelling was first introduced by [Guikema and Goffelt \(2008\)](#) and [Sellers and Shmueli \(2010\)](#) in an effort to extend Poisson regression by allowing covariates to influence not only the location parameter μ_i but also the dispersion parameter ν_i . For notational simplicity we now let $\theta_i = (\mu_i, \nu_i)^T$ be the parameter for observation y_i with θ_i derived from setting $\theta_i = \eta(\beta, x_i)$ where $\eta(\beta, x_i)$ is a link function of the b covariates and β are the parameters of this link function. The major difficulty is that the likelihood of the COM-Poisson regression model involves multiple intractable normalising constants,

$$f(y_{1:n}|\theta_{1:n}) = \prod_{i=1}^n f(y_i|\theta_i) = \prod_{i=1}^n \frac{q_f(y_i|\theta_i)}{\mathcal{Z}_f(\theta_i)}, \quad \theta_i = \eta(\beta, x_i). \quad (2)$$

The particular link function $\eta(\beta, x_i)$ for the COM-Poisson was specified by [Guikema and Goffelt \(2008\)](#) as a dual-link function to allow the available covariates to enter both the μ_i and ν_i parameter. In this dual-link function scenario the parameter β is split into two sub-vectors as $\beta = (\beta_\mu, \beta_\nu)$, where β_μ are the coefficients for the μ link component and β_ν are the coefficients for the ν link component. The dual-link function is then

$$\theta_i = \begin{bmatrix} \mu_i \\ \nu_i \end{bmatrix} = \eta(\beta, x_i) = \begin{bmatrix} \eta_1(\beta_\mu, x_i) \\ \eta_2(\beta_\nu, x_i) \end{bmatrix} = \begin{bmatrix} \exp\left(\beta_{\mu,0} + \sum_{j=1}^b \beta_{\mu,j} x_{ij}\right) \\ \exp\left(\beta_{\nu,0} + \sum_{j=1}^b \beta_{\nu,j} x_{ij}\right) \end{bmatrix},$$

where some β_μ are set fixed to 0 when the covariate is not included in $\eta_1(\beta_\mu, x_i)$, similarly for $\eta_2(\beta_\nu, x_i)$.

Bayesian parameter inference for β in COM-Poisson regression models is a doubly-intractable problem as we now demonstrate. Placing a prior on β , the posterior distribution for β is formed with this prior and the complete likelihood (2) as

$$\begin{aligned} \pi(\beta|y_{1:n}) &\propto f(y_{1:n}|\theta_{1:n})\pi(\beta) \\ &= \prod_{i=1}^n \frac{q_f(y_i|\theta_i)}{\mathcal{Z}_f(\theta_i)}\pi(\beta), \quad \theta_i = \eta(\beta, x_i). \end{aligned} \quad (3)$$

This posterior is termed doubly-intractable since the likelihood is intractable, containing n intractable normalising constants, and the posterior itself cannot be normalised. Without the ability to evaluate $f(y_{1:n}|\theta_{1:n})$ pointwise, the direct application of standard MCMC methods is infeasible. For example,

a Metropolis-Hastings algorithm proposing a move in the link function from β to β' using a proposal distribution $h(\beta, \beta')$ requires evaluation of the intractable ratios $\left\{ \frac{\mathcal{Z}_f(\theta_i)}{\mathcal{Z}_f(\theta'_i)} \right\}_{i=1}^n$ to compute the acceptance ratio

$$\alpha(\beta, \beta') = \min \left\{ 1, \frac{\prod_{i=1}^n \frac{q_f(y_i|\theta'_i)}{\mathcal{Z}_f(\theta'_i)} h(\beta', \beta) \pi(\beta')}{\prod_{i=1}^n \frac{q_f(y_i|\theta_i)}{\mathcal{Z}_f(\theta_i)} h(\beta, \beta') \pi(\beta)} \right\}. \quad (4)$$

Possible methods to compute this acceptance ratio include replacing the intractable ratios with approximations using the asymptotic approximation from [Shmueli et al. \(2007\)](#) or by a truncated sum $\mathcal{Z}_f(\theta_i) = \sum_{i=0}^k q_f(y|\theta_i)$ but both of these methods introduce some bias into the acceptance ratio.

An ingenious solution to the issue of computing the acceptance ratio (4) was proposed by [Murray et al. \(2006\)](#) drawing on earlier work by [Møller et al. \(2006\)](#). The solution, known as the exchange algorithm, augments the doubly-intractable posterior (3) with auxiliary data, provided that it is possible to sample data exactly from the intractable likelihood. The original exchange algorithm overcomes a single intractable ratio i.e. $n = 1$ however the algorithm can easily be extended to the case of $n > 1$ intractable ratios as we now show. Starting from the posterior (3) which contains the product of n independent non-identical intractable likelihood functions, the exchange algorithm proposes, as in standard Metropolis-Hastings, to update the parameter from the current state β to proposed state β' using $h(\beta, \beta')$ but in addition the posterior is augmented with n auxiliary draws $y'_{1:n} = (y'_1, \dots, y'_n)$ generated from the likelihood at the proposed parameters $\theta'_i = \eta(\beta', x_i)$. The augmented posterior is written

$$\pi(\beta, \beta', y'_{1:n} | y_{1:n}) \propto \underbrace{f(y_{1:n} | \theta_{1:n})}_{\text{likelihood}} \pi(\beta) h(\beta, \beta') \underbrace{f(y'_{1:n} | \theta'_{1:n})}_{\text{auxiliary draws}} \quad (5)$$

and the marginal of this augmented posterior for β is the target posterior of interest (3). The algorithm can be seen as a Markov chain over an augmented parameter $(\beta, y_{1:n})$. The acceptance ratio for the augmented posterior (5) is now tractable due to the cancellation of all intractable normalising constants

$$\begin{aligned} \alpha_{\text{exchange}, n}(\beta, \beta') &= \min \left\{ 1, \frac{\prod_{i=1}^n \frac{q_f(y_i|\theta'_i)}{\mathcal{Z}_f(\theta'_i)} h(\beta', \beta) \pi(\beta') \prod_{i=1}^n \frac{q_f(y'_i|\theta_i)}{\mathcal{Z}_f(\theta_i)}}{\prod_{i=1}^n \frac{q_f(y_i|\theta_i)}{\mathcal{Z}_f(\theta_i)} h(\beta, \beta') \pi(\beta) \prod_{i=1}^n \frac{q_f(y'_i|\theta'_i)}{\mathcal{Z}_f(\theta'_i)}} \right\}, \quad y'_i \sim f(\cdot | \theta'_i), \forall i \\ &= \min \left\{ 1, \frac{\prod_{i=1}^n q_f(y_i|\theta'_i) h(\beta', \beta) \pi(\beta') \prod_{i=1}^n q_f(y'_i|\theta_i) \prod_{i=1}^n \frac{1}{\mathcal{Z}_f(\theta'_i)} \prod_{i=1}^n \frac{1}{\mathcal{Z}_f(\theta_i)}}{\prod_{i=1}^n q_f(y_i|\theta_i) h(\beta, \beta') \pi(\beta) \prod_{i=1}^n q_f(y'_i|\theta'_i) \prod_{i=1}^n \frac{1}{\mathcal{Z}_f(\theta'_i)} \prod_{i=1}^n \frac{1}{\mathcal{Z}_f(\theta_i)}} \right\} \quad (6) \end{aligned}$$

The clever cancellation of the normalising constants in (6) is due to the *exchange* of parameters (θ_i, θ'_i) associated with the observed data $(y_{1:n})$ and the auxiliary data $(y'_{1:n})$, the auxiliary being data discarded after each move. The major requirement in the exchange algorithm is the ability to generate the auxiliary data, requiring exact draws from the likelihood at each proposed parameter. Perfect sampling ([Propp and Wilson, 1996](#)) from the likelihood is one possible method to do this but can be prohibitively expensive for many models. In certain models sampling from an intractable likelihood may even be as difficult as computing the intractable normalising constant itself and instead one could resort to sampling using an approximate likelihood sampler such as in the case of the exponential random graph model ([Caimo and Friel, 2011](#)). The COM-Poisson distribution was shown by [Chaniavidis et al. \(2018\)](#) to be amenable to rejection sampling, allowing the auxiliary data to be generated. In Section 3.1 we present our faster rejection sampler having a cheaper sampling cost than the sampler of [Chaniavidis et al. \(2018\)](#). Rejection sampling is not applicable to all intractable likelihoods but when rejection sampling is available we show in Section 3.3 that the intractable likelihood can be estimated without

bias through an unbiased estimate of the reciprocal normalising constant. This unbiased intractable likelihood estimator is also guaranteed to be positive which differs from other estimates using Russian roulette stochastic truncation (Lyne et al., 2015) or truncated Markov chains (Wei and Murray, 2017).

3 Rejection sampling to an intractable likelihood estimator

Rejection sampling (von Neumann, 1951) is a sampling scheme to generate statistically independent samples from a target probability distribution of interest. We denote the target density (likelihood) as

$$f(y|\theta) = \frac{q_f(y|\theta)}{\mathcal{Z}_f(\theta)}$$

where in this section we drop the observation index i for ease of notation. This target density may be tractable or intractable but is assumed in any case to be difficult to sample from. The rejection sampling method works by constructing an envelope distribution over the target density, proposing samples from this envelope distribution and choosing to accept a portion of these samples. The envelope distribution should be chosen so that it is both computationally efficient to sample from and matches closely in shape the target density. The portion of samples that will be accepted depends on the similarity of the envelope and the target distribution. The closer the envelope distribution matches the target distribution, the higher the number of samples that will be accepted. To generate a single draw y from $f(y|\theta)$ using rejection sampling, we consider an envelope distribution $g(y|\gamma)$ where $\gamma \in \Gamma$ is its parameter. We assume that the envelope density can be written in a similar form to $f(y|\theta)$ as

$$g(y|\gamma) = \frac{q_g(y|\gamma)}{\mathcal{Z}_g(\gamma)}, \quad \text{with } \mathcal{Z}_g(\gamma) = \int_y q_g(y|\gamma) dy.$$

The envelope distribution's normalising constant $\mathcal{Z}_g(\gamma)$ may be either tractable or intractable. A necessary condition for $g(\cdot|\gamma)$ is that it dominates the support of $f(\cdot|\theta)$ i.e. $g(y|\gamma) = 0 \implies f(y|\theta) = 0$. It is also necessary to bound $f(y|\theta)$ using $g(y|\gamma)$ and a positive finite bounding constant M that satisfies the envelope inequality $Mg(y|\gamma) > f(y|\theta), \forall y$. Finding such a constant can be a difficult task even for tractable densities. The optimal value of M is denoted as $M_{f/g} = \sup_y \left\{ \frac{f(y|\theta)}{g(y|\gamma)} \right\}$. In the case of intractable likelihoods, this optimal constant is impossible to evaluate since at least one of either $\mathcal{Z}_f(\theta)$ or $\mathcal{Z}_g(\gamma)$ is unknown leaving $M_{f/g}$ intractable. In the case where both $f(y|\theta)$ and $g(y|\gamma)$ are tractable distributions and $M_{f/g}$ can be found, the vanilla rejection sampling algorithm to obtain a single draw from $f(y|\theta)$ proceeds as in Algorithm 1.

Algorithm 1: Rejection sampling

Input: Target distribution $f(\cdot|\theta)$, envelope distribution $g(\cdot|\gamma)$ and $M_{f/g} = \sup_y \left\{ \frac{f(y|\theta)}{g(y|\gamma)} \right\}$

- 1 Sample a proposal y^* from $g(\cdot|\gamma)$
- 2 Accept y^* as a draw from $f(\cdot|\theta)$ by a Bernoulli trial with acceptance probability

$$\alpha(y^*) = \frac{f(y^*|\theta)}{M_{f/g}g(y^*|\gamma)}, \tag{7}$$

otherwise repeat.

Algorithm 1 introduces the rejection sampling Bernoulli trial with the outcome conditional on the proposed y^* through the acceptance probability $\alpha(y^*)$. This Bernoulli trial is at the core of the rejection sampling method as it decides for a given y^* what proportion of samples from $g(\cdot|\gamma)$ to accept. Once intractability is introduced into this Bernoulli trial through $M_{f/g}$, it would appear that $\alpha(y^*)$

becomes intractable but this is fortunately not the case. The intractable bounding constant $M_{f/g}$ can be decomposed into tractable and intractable components by extracting the reciprocal normalising constants from $M_{f/g}$ as follows

$$M_{f/g} = \sup_y \left\{ \frac{f(y|\theta)}{g(y|\gamma)} \right\} = \frac{1/\mathcal{Z}_f(\theta)}{1/\mathcal{Z}_g(\gamma)} \sup_y \left\{ \frac{q_f(y|\theta)}{q_g(y|\gamma)} \right\} = \frac{\mathcal{Z}_g(\gamma)}{\mathcal{Z}_f(\theta)} B_{f/g} \quad (8)$$

and this introduces the tractable bounding constant $B_{f/g} = \sup_y \left\{ \frac{q_f(y|\theta)}{q_g(y|\gamma)} \right\}$. $B_{f/g}$ can be viewed as the upper bound on the ratio of the unnormalised densities $q_f(\cdot|\theta)$ and $q_g(\cdot|\gamma)$, whereas $M_{f/g}$ is the upper bound on the associated normalised densities $f(\cdot|\theta)$ and $g(\cdot|\gamma)$. The advantage of rejection sampling algorithms is that they can be run without full knowledge of $M_{f/g}$, it is sufficient to know $B_{f/g}$ which is usually more readily available. Substituting the expression $M_{f/g} = \frac{\mathcal{Z}_g(\gamma)}{\mathcal{Z}_f(\theta)} B_{f/g}$ into the acceptance probability (7) gives the tractable acceptance probability for the conditional Bernoulli trial as

$$\alpha(y^*) = \frac{f(y^*|\theta)}{\left(\frac{\mathcal{Z}_g(\gamma)}{\mathcal{Z}_f(\theta)} B_{f/g}\right) g(y^*|\gamma)} = \frac{q_f(y^*|\theta) \frac{1}{\mathcal{Z}_f(\theta)}}{\left(\frac{\mathcal{Z}_g(\gamma)}{\mathcal{Z}_f(\theta)} B_{f/g}\right) q_g(y^*|\gamma) \frac{1}{\mathcal{Z}_g(\gamma)}} = \frac{q_f(y^*|\theta)}{B_{f/g} q_g(y^*|\gamma)}.$$

It is important to note that $\mathcal{Z}_g(\gamma)$ is also not required to be tractable, although it typically is for simple envelope distributions. A case where $\mathcal{Z}_g(\gamma)$ is intractable would be when the envelope distribution itself requires sampling by another rejection sampler and this opens the possibility of chaining many rejection samplers together to sample from evermore complex target distributions. The main purpose of rejection sampling in this paper is to provide a means to generate the auxiliary draws required within the exchange algorithm acceptance ratio (6) allowing the doubly-intractable inference to be performed but Section 3.3 will show another novel usage for rejection sampling. Before discussing this, we demonstrate our COM-Poisson rejection sampler in the next section.

3.1 COM-Poisson rejection sampler

For the COM-Poisson distribution, our sampler uses a choice of two envelope distributions dependent on the value of ν . This choice is motivated by Figure 1, where a Poisson distribution is used in the case of $\nu \geq 1$ and a geometric distribution in the case of $\nu < 1$. Theorem 3.1 demonstrates and proves why this works. The COM-Poisson rejection sampling algorithm is given in pseudocode after this theorem.

Theorem 3.1 (COM-Poisson intractable rejection sampler). *Suppose that $Y \sim g_1(y|\gamma)$, a Poisson random variable with parameter $\gamma = \mu$ and $Y \sim g_2(y|\gamma)$, a geometric random variable with parameter $\gamma = p$, for some $0 < p < 1$, are used as two enveloping distributions as follows,*

$$g(y|\gamma) = \begin{cases} g_1(y|\gamma = p) = p(1-p)^y, & \text{if } \nu < 1 \text{ (geometric envelope),} \\ g_2(y|\gamma = \mu) = \frac{\mu^y}{e^{\mu} y!}, & \text{if } \nu \geq 1 \text{ (Poisson envelope),} \end{cases} \quad (9)$$

together with the following tractable enveloping bounds

$$B_{f/g} = \begin{cases} \frac{1}{p} \frac{\mu \binom{\nu \lfloor \frac{\mu}{(1-p)^{1/\nu}} \rfloor}{\lfloor \frac{\mu}{(1-p)^{1/\nu}} \rfloor}}{(1-p)^{\binom{\nu \lfloor \frac{\mu}{(1-p)^{1/\nu}} \rfloor}} \left(\lfloor \frac{\mu}{(1-p)^{1/\nu}} \rfloor!\right)^{\nu}}, & \text{if } \nu < 1, \\ \left(\frac{\mu \lfloor \mu \rfloor}{\lfloor \mu \rfloor!}\right)^{\nu-1}, & \text{if } \nu \geq 1. \end{cases} \quad (10)$$

then a sample from the COM-Poisson distribution can be drawn at any parameter value (μ, ν) .

Proof. See Appendix A. □

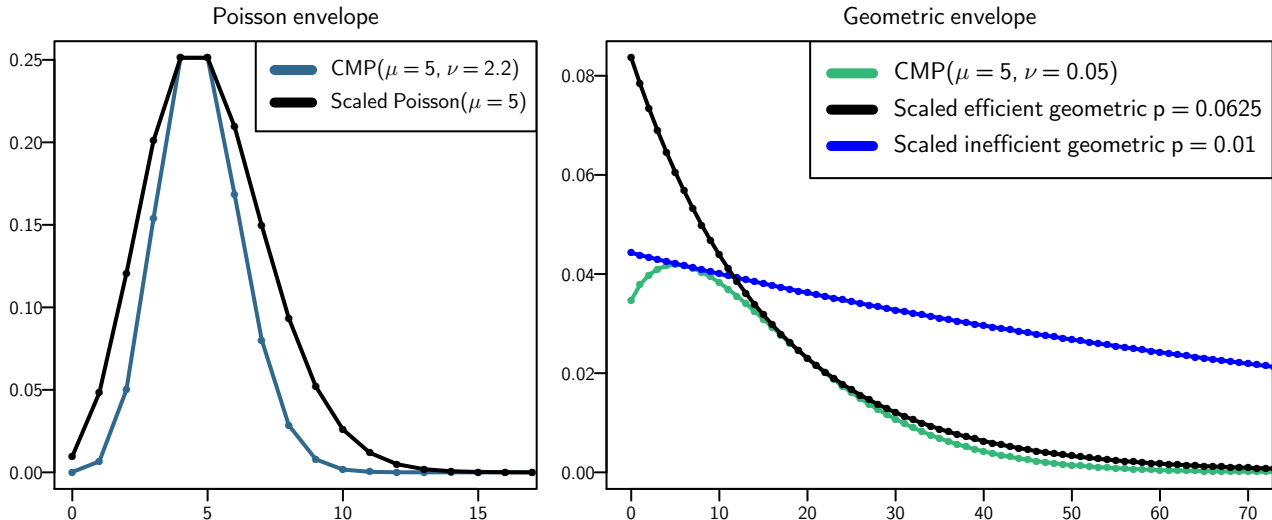


Figure 1: Enveloping the COM-Poisson distribution for rejection sampling: In the underdispersed ($\nu \geq 1$) case (left) a Poisson distribution, having a higher relative dispersion is scaled to envelope the COM-Poisson distribution. In the overdispersed ($\nu < 1$) case (right) a geometric distribution, having a higher relative dispersion is used. The parameter p of the geometric distribution in the overdispersed case controls the efficiency of the enveloping. To choose p we match the first moment of the geometric distribution with the COM-Poisson distribution (11), shown as the black curve (right). The blue curve corresponds to an alternative choice of p , showing how an inefficient choice of p results in a geometric envelope with a high rejection region (area between blue and green curves).

Before we present the practical implementation of the COM-Poisson sampling algorithm developed from Theorem 3.1, it is necessary to provide some brief discussion on the choice of p in the geometric envelope. The optimal p parameter, p^* say, would be one for which the acceptance rate of the rejection sampling is maximised. However it is not possible to compute p^* since the acceptance rate, as we will show in Section 3.3, involves intractable normalising constants. Nevertheless any choice of p will result in a valid rejection sampler and clearly the closer p is to p^* , the more efficient the sampler is. We have chosen to select p by matching the first moment of the geometric to the approximate expected value of the COM-Poisson distribution using (1). Setting the first moment of the geometric distribution equal to the expected value approximation (1) implies,

$$\frac{1-p}{p} = \mu + \frac{1}{2\nu} - \frac{1}{2} \quad \iff \quad p = \frac{2\nu}{2\mu\nu + 1 + \nu}. \quad (11)$$

Algorithm 2 presents pseudocode of the COM-Poisson sampler for coding on a computer.

3.2 Efficiency and comparison of the rejection sampling algorithm

To analyse the efficiency of the COM-Poisson sampler we conducted an experiment in which the sampler was run for different values of μ each over a fine grid of ν values $\in [0, 6]$, as outlined in Figure 2. Here we monitored the acceptance rate of the COM-Poisson rejection sampler, based on the average of the reciprocal of the number of draws from the envelope distribution needed to yield a single draw from the COM-Poisson distribution over 100,000 independent runs at each grid point. This illustrates

Algorithm 2: Sampler for the COM-Poisson(μ, ν) distribution

Input: Parameter $\theta = (\mu, \nu)$

```
1 START
2 if  $\nu \geq 1$  then
3   Sample  $y' \sim \text{Poisson}(\mu)$  using the algorithm from Ahrens and Dieter (1982).
4   Calculate  $B_{f/g}^{[\nu \geq 1]}$  using (10) and set  $\alpha = \frac{(\mu^{y'}/y'!)^\nu}{B_{f/g}^{[\nu \geq 1]}(\mu^{y'}/y'!)}$ .
5   Generate  $u \sim \text{Uniform}(0, 1)$ .
6   if  $u \leq \alpha$  then
7     return  $y'$ 
8   else
9     GOTO START
10  end
11 end
12 if  $\nu < 1$  then
13   Compute  $p = \frac{2\nu}{2\mu\nu + 1 + \nu}$ .
14   Sample  $y' \sim \text{geometric}(p)$  by first sampling  $u_0 \sim \text{Uniform}(0, 1)$  and returning  $\lfloor \frac{\log(u_0)}{\log(1-p)} \rfloor$ .
15   Calculate  $B_{f/g}^{[\nu < 1]}$  using (10) and set  $\alpha = \frac{(\mu^{y'}/y'!)^\nu}{B_{f/g}^{[\nu < 1]}(1-p)^{y'} p}$ .
16   Generate  $u \sim \text{Uniform}(0, 1)$ .
17   if  $u \leq \alpha$  then
18     return  $y'$ 
19   else
20     GOTO START
21  end
22 end
23 Note: For multiple draws from a COM-Poisson( $\mu, \nu$ ) distribution, calculate the bound  $B_{f/g}^{[\nu \geq 1]}$ 
or  $B_{f/g}^{[\nu < 1]}$  once and then use it as an input to the algorithm.
```

that the sampler can achieve acceptance rates between 50% and 100% for $\nu > 1$ (corresponding to a Poisson envelope distribution) and acceptance rates between 30% and 80% for $\nu < 1$ (corresponding to a Geometric envelope distribution). For $\nu < 1$, the acceptance of the sampler increases as $\nu \rightarrow 0$. This occurs because the COM-Poisson distribution, with parameter $\nu = 0$, is exactly a geometric distribution with parameter $p = \mu^\nu$ provided $\mu^\nu < 1$. Thus the geometric proposal matches closely the COM-Poisson distribution for small ν .

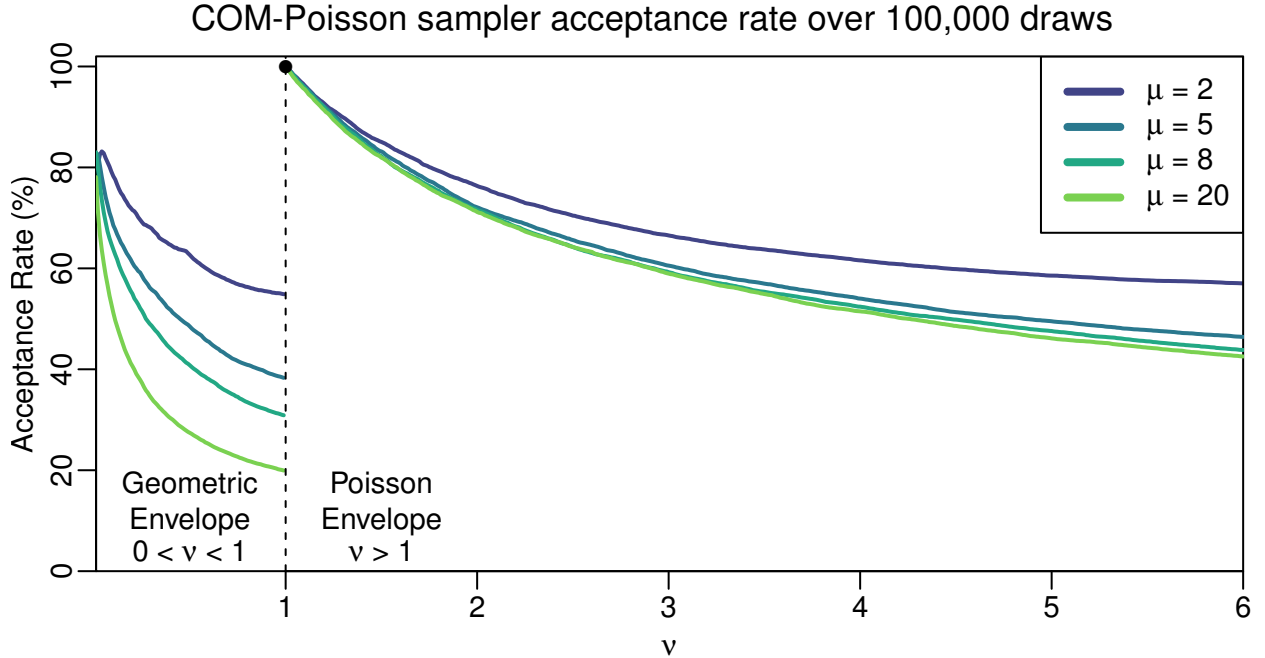


Figure 2: COM-Poisson sampler efficiency: The acceptance rate is the inverse of the number of draws required from either a geometric or a Poisson to obtain one draw from a COM-Poisson distribution for varying values of ν and μ .

In Figure 3 we compare the efficiency of our COM-Poisson rejection sampler to the rejection sampler of Chaniavidis et al. (2018). The sampler developed in Chaniavidis et al. (2018) involves constructing a tight envelope distribution around the COM-Poisson distribution using a piecewise envelope distribution constructed from four truncated geometric distributions. This piecewise geometric distribution is then sampled using the inversion method for truncated geometric distributions. These truncated geometric distributions are designed to match closely to the COM-Poisson density function, however the computational cost of constructing this envelope is significant and can impact on the CPU time. Our COM-Poisson sampler described in Algorithm 2 is much simpler in construction and requires only a single envelope component depending on the value of ν . Although the single envelope can result in a moderately higher rejection rate compared to Chaniavidis et al. (2018) and as illustrated in Figure 2, we describe an experiment to illustrate that a potentially higher rejection rate is offset, in CPU time, by the simplicity of constructing and sampling from the single envelope distribution. For example, spending time searching for p^* may improve the acceptance rate; but this search is costly and will reduce the number of effective draws per unit time.

In this experiment we compare the computational run time of the sampler in Algorithm 2 to that presented in Chaniavidis et al. (2018) over a wide range of μ and ν values. To do this, a 128×128 grid of equally-spaced $\mu \in [1, 25]$ and $\nu \in [0.01, 10.0]$ values was constructed. At each site 2500 values were drawn from each sampler and the average time taken to draw from each sampler is computed. The plot in Figure 3 illustrates that Algorithm 2 is between 1.01 and 8.01 times faster. From this one can observe that our sampler represents a significant speedup when used in the context of an MCMC algorithm, as outlined in Section 5.2.1. This analysis outlines an important aspect of rejection

sampling in that the choice of using an efficient envelope must be considered simultaneously with the computational cost of sampling from and constructing this efficient envelope.

Sampler Comparison – Scale Factor

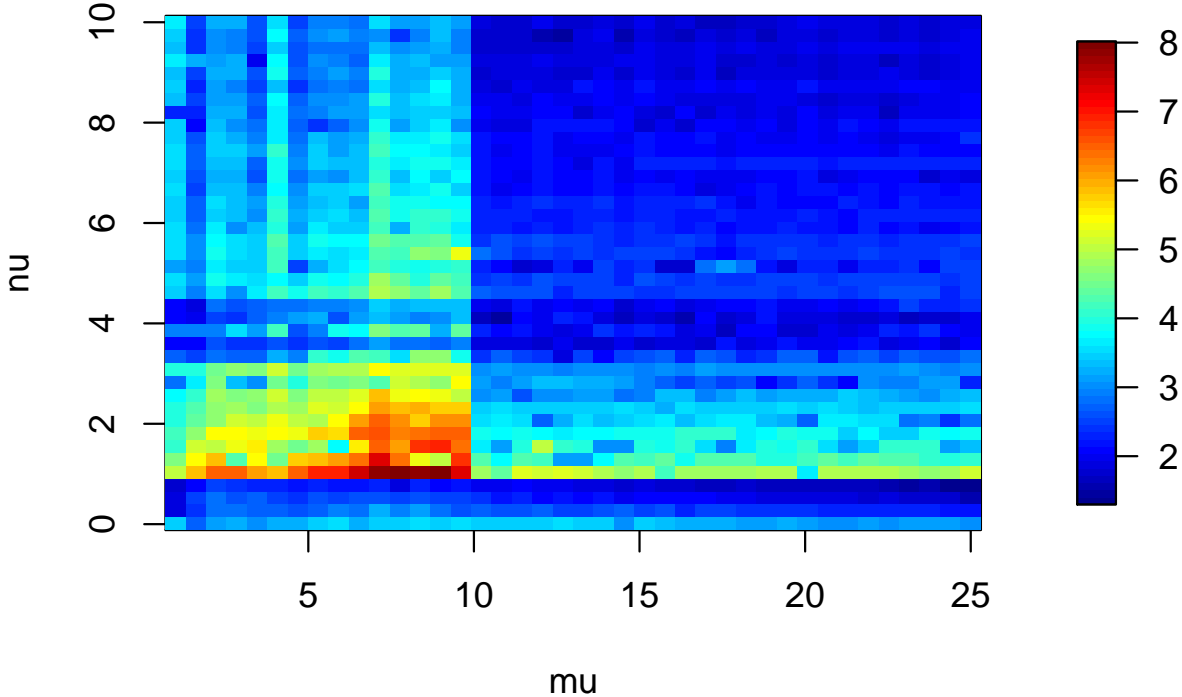


Figure 3: This plot shows an experiment to compare the sampler in Algorithm 2 to that in Chaniyalidis et al. (2018) over a wide range of μ and ν values. A 128×128 grid of $\mu \in [1, 25]$ and $\nu \in [0.01, 10.0]$ values is constructed. At each grid point 2500 values are drawn from each sampler and the average time per sample is computed. The above plot displays the ratio of average sample time from both sampler, where algorithm 2 is in the dominator of the ratio. At its fastest Algorithm 2 is 8.01 times faster and at its slowest 1.01 times faster.

The speedup of our algorithm in Figure 3 is notably reduced for $\mu \geq 10$. The reason for this reduction is a subtle artefact of computer generation of Poisson samples for large μ . As discussed in Ahrens and Dieter (1982), computer implementations of Poisson samplers generally switch to a different sampler for values of $\mu \geq 10$. Further work could examine methods to overcome this restriction in order to improve our sampler in this region. Moreover, the results of the experiment in Figure 3 can be viewed as a lower bound to the potential speedup of our algorithm.

Another benefit of rejection sampling algorithms is that they can be implemented in parallel thus speeding up the computation, however care is needed when implementing parallel random number generators with potential issues discussed in Mascagni and Srinivasan (2000). Algorithm 2 gives the pseudocode for the COM-Poisson rejection sampler.

3.3 Unbiased likelihood estimation by monitoring rejection sampling efficiency

Note: In this section we reintroduce the index parameter i to discuss rejection sampling at different parameters θ_i , each θ_i necessarily introduces indexed tractable and intractable sampling bounds and envelope parameters which we denote as $B_{f/g,i}$, $M_{f/g,i}$ and γ_i respectively.

Literature on rejection sampling has explored some methods for recycling the rejected proposals from the envelope distribution. Casella and Robert (1996) consider reducing the variance of rejection

sampling estimates by incorporating the rejected samples in a post-processing step using the Rao-Blackwell theorem. Rao et al. (2016) uses the rejected proposals from a rejection sampler in a different way by considering the joint distribution of rejected and accepted draws leading to a method to perform doubly-intractable inference in the case where the likelihood has an associated rejection sampler. Our approach considers yet another use for the rejected proposals by examining the relationship of rejected proposals with the efficiency of the rejection sampler. The efficiency is used to provide an unbiased estimator of the reciprocal normalising constant $\frac{1}{\mathcal{Z}_f(\theta_i)}$ for each parameter θ_i . An unbiased estimator of the reciprocal normalising constant at each θ_i will lead directly to an unbiased estimator of the complete likelihood (2).

The efficiency of the rejection sampler is critical to the overall performance of the exchange algorithm. The efficiency is defined as the total number of draws N_1 required from the envelope distribution until the first acceptance of a draw from $f(y|\theta_i)$ or in general the total number of draws N_r required until r acceptances from $f(y|\theta_i)$. The number of draws determines how closely the envelope distribution matches the target distribution. The ideal envelope distribution would exactly match the target distribution and reject no proposals, however in this scenario we could simply circumvent rejection sampling and sample from the envelope distribution directly. In Algorithm 1, run at parameter θ_i , each of the samples proposed from the envelope is followed by a Bernoulli trial with conditional acceptance probability $\alpha_i(y^*) = \frac{f(y^*|\theta_i)}{M_{f/g,i}g(y^*|\gamma_i)}$. To remove the conditional dependence on y^* , consider the unconditional acceptance probability $\bar{\alpha}_i$ over all envelope proposals by integrating $\alpha_i(y^*)$ with respect to the envelope,

$$\bar{\alpha}_i = \int_{y^*} \alpha_i(y^*)g(y^*|\gamma_i)dy^* = \int_{y^*} \frac{f(y^*|\theta_i)}{M_{f/g,i}} dy^* = \frac{1}{M_{f/g,i}}.$$

It is now obvious that the unconditional acceptance probability and from this the rejection sampler efficiency are controlled directly by the magnitude of the intractable bound $M_{f/g,i}$. On average, one expects to accept a portion $\frac{1}{M_{f/g,i}}$ of proposals from $g(\cdot|\gamma_i)$ as being from $f(\cdot|\theta_i)$ and to discard the remainder. Equivalently, the total number of draws until the first accepted draw follows a geometric distribution with success probability $\frac{1}{M_{f/g,i}}$ and mean $M_{f/g,i}$ i.e. $N_1 \sim \text{Geometric}\left(\frac{1}{M_{f/g,i}}\right)$. In general the total number of draws (N_r) until r acceptances follows a negative binomial distribution with mean $rM_{f/g,i}$. A natural method to estimate $M_{f/g,i}$ is to run rejection sampling at the parameter θ_i and record the observed total number of draws $n_{r,i}$ required for r acceptances. Then an unbiased estimate of $M_{f/g,i}$ based on r acceptances is given as

$$\widehat{M}_{f/g,i}^{(r)} = \frac{n_{r,i}}{r}.$$

By the central limit theorem increasing r will lead to lower variance estimates of $M_{f/g,i}$ since $\widehat{M}_{f/g,i}^{(r)}$ can be viewed as the average of r independent geometric trials each having variance $M_{f/g,i}(M_{f/g,i} - 1)$. To use $\widehat{M}_{f/g,i}^{(r)}$ to give an unbiased estimate of the likelihood, we have the final assumption that $Z_g(\gamma_i)$ is known i.e. the envelope distribution is tractable. Then by rewriting (8) a link emerges between the inverse normalising constants of $g(\cdot|\gamma_i)$ and $f(\cdot|\theta_i)$ as

$$\frac{1}{\mathcal{Z}_f(\theta_i)} = \frac{1}{\mathcal{Z}_g(\gamma_i)} \frac{M_{f/g,i}}{B_{f/g,i}}. \quad (12)$$

By replacing $M_{f/g,i}$ with its unbiased estimate $\widehat{M}_{f/g,i}^{(r)}$ and multiplying both sides by $q_f(y|\theta_i)$ an unbi-

ased estimator of the likelihood for observation y_i is given as

$$\widehat{f}^{(r)}(y_i|\theta_i) = \frac{q_f(y_i|\theta_i)}{\mathcal{Z}_g(\gamma_i)} \frac{\widehat{M}_{f/g,i}^{(r)}}{B_{f/g,i}} \quad (13)$$

with all terms on the RHS of (13) known. The estimate of the complete likelihood follows as

$$\widehat{f}^{(r)}(y_{1:n}|\theta_{1:n}) = \prod_{i=1}^n \frac{q_f(y_i|\theta_i)}{\mathcal{Z}_g(\gamma_i)} \frac{\widehat{M}_{f/g,i}^{(r)}}{B_{f/g,i}}. \quad (14)$$

by estimating $\widehat{M}_{f/g,i}^{(r)}$ at each θ_i . We point out here that the idea to use the acceptance probability of a rejection sampling algorithm to develop an unbiased estimate of the intractable normalising constant has previously explored, in the context of inference for stochastic differential equations. We refer the reader to Section 5 of Beskos et al. (2006).

This unbiased estimator of the intractable likelihood is very useful especially for model selection as it can be used to construct an estimate of the Bayesian information criterion (BIC) (Schwarz, 1978). Bayesian model choice is a challenging problem for models with intractable likelihood functions. The intractability prevents calculation of the likelihood which is required for most model selection criteria. The BIC estimate from the unbiased likelihood estimator (14) is

$$\widehat{\text{BIC}}^{(r)} = k \log(n) - 2 \log \left(\widehat{f}^{(r)}(y_{1:n}|\hat{\theta}_{1:n}) \right) \quad (15)$$

where $\hat{\theta}_{1:n}$ maximises the unbiased estimate. The variance of this BIC estimate depends on r and we find in experiments that r needs to be high (> 1000) to achieve an accurate estimate of the BIC. We note that an interesting future research question is to understand the statistical properties, such as consistency of the maximum likelihood estimator $\hat{\theta}_{1:n}$. A useful starting point might be to consider the literature where this issue has been explored for different intractable likelihood problems including Geyer and Thompson (1992) and Beskos et al. (2006).

It may also be possible to use the unbiased likelihood estimator (13) as part of other Bayesian model selection frameworks, for example, in the calculation of the marginal likelihood (or evidence) and Bayes factors. This could be a focus of future work. In addition to BIC estimation we consider plugging the unbiased estimator directly into the intractable acceptance ratio (4) which is the basis of pseudo-marginal MCMC algorithms discussed in the next section.

Finally, we note that Chanialidis et al. (2018) presented results comparing models based on the deviance information criterion (DIC) (Spiegelhalter et al., 2002), without explaining how they overcame the need to evaluate the intractable likelihood for each draw from the MCMC sample. Following personal communication with the authors it turns out that they approximated the normalising constant, $\mathcal{Z}_f(\mu, \nu)$, using a truncated summation involving $k_2 - k_1 + 1$ terms, for positive integers, k_1, k_2 , where $k_1 < k_2$, as follows,

$$\mathcal{Z}_f(\mu, \nu) \approx \hat{\mathcal{Z}}_f(\mu, \nu) = \sum_{y=k_1}^{k_2} \left(\frac{\mu^y}{y!} \right)^\nu. \quad (16)$$

Here k_1 and k_2 are selected to satisfy the inequality $k_1 < \lfloor \mu \rfloor < k_2$, where, as before, $\lfloor \mu \rfloor$ is the mode of the COM-Poisson distribution. In turn, they approximated the likelihood for each draw from the sample generated from their MCMC algorithm by

$$\hat{f}(y_{1:n}|\mu_{1:n}, \nu_{1:n}) = \prod_{i=1}^n \left(\frac{\mu_i^{y_i}}{y_i!} \right)^{\nu_i} \frac{1}{\hat{\mathcal{Z}}_f(\mu_i, \nu_i)}, \quad (17)$$

which is then used, for example, to estimate the posterior expected (deviance) log-likelihood which is required for calculation of the DIC. Of course, this introduces an approximation of the likelihood into the estimation of the DIC. This raises questions as to how many terms $k_2 - k_1 + 1$ are needed in order to get an accurate approximation of the likelihood as well as how to choose both k_1 and k_2 and finally raises the issue of the computational cost involved in evaluating the finite truncation for the n terms in the product in (17). We return to this issue in Section 5.1.1, where we find that the number of terms required to get an accurate approximation of the likelihood depends considerably on the values of μ and ν .

4 Pseudo-marginal MCMC with the intractable likelihood estimator

MCMC methods for Bayesian inference require the ability to evaluate the unnormalised posterior distribution. For doubly-intractable problems where the posterior distribution cannot be evaluated as such, there exist pseudo-marginal MCMC algorithms (Andrieu and Roberts, 2009) which require only an unbiased positive estimate of the posterior distribution. Pseudo-marginal algorithms have been applied to a range of problems such as genetic modelling (Beaumont, 2003) and Markov jump processes (Georgoulas et al., 2017). This pseudo-marginal approach can be seen as an alternative to the exchange algorithm which uses exact sampling in place of unbiased posterior estimators.

We now describe the pseudo-marginal framework and show how our unbiased estimator of the likelihood (14) can be used. Consider a target distribution $\pi(\theta)$ which cannot be pointwise evaluated, but assume that it is possible to construct an unbiased estimate $\hat{\pi}(\theta, z)$ of the target through the use of auxiliary variables $z \sim w(\cdot|\theta)$, $z \in Z$. Using $\hat{\pi}(\theta)$ in place of the true target in a Metropolis-Hastings algorithm leads to the approximate acceptance ratio

$$\min \left\{ 1, \frac{\hat{\pi}(\theta')h(\theta', \theta)}{\hat{\pi}(\theta)h(\theta, \theta')} \right\}$$

Pseudo-marginal MCMC algorithms can be viewed as Metropolis-Hastings algorithms over the joint parameter and auxiliary variable state space $\theta \cup Z$. Andrieu and Roberts (2009) describe two alternative forms: the ‘‘Monte Carlo within Metropolis’’ or MCWM algorithm and the ‘‘grouped independence Metropolis-Hastings’’ or GIMH algorithm, with only the latter algorithm having the correct target distribution $\pi(\theta)$. In MCWM, both the current $\hat{\pi}(\theta)$ and proposed $\hat{\pi}(\theta')$ estimate of the target are refreshed every iteration using new auxiliary draws z and z' . In GIMH, only the proposed estimate $\hat{\pi}(\theta')$ is refreshed and $\hat{\pi}(\theta)$ is fixed to the estimate used when θ was last accepted.

The difficulty with pseudo-marginal algorithms is in constructing an unbiased estimate of the target. Using the rejection sampling unbiased estimator (14) the pseudo-marginal acceptance ratio for the β to β' move in the COM-Poisson model is

$$\begin{aligned} \alpha_{PM,n}(\beta, \beta') &= \min \left\{ 1, \frac{\widehat{f}^{(r)}(y_{1:n}|\theta'_{1:n}) h(\beta', \beta) \pi(\beta')}{\widehat{f}^{(r)}(y_{1:n}|\theta_{1:n}) h(\beta, \beta') \pi(\beta)} \right\} \\ &= \min \left\{ 1, \frac{\prod_{i=1}^n \frac{q_f(y_i|\theta'_i) \widehat{M}'_{f/g,i}{}^{(r)}}{\mathcal{Z}_g(\gamma'_i) B'_{f/g,i}} h(\beta', \beta) \pi(\beta')}{\prod_{i=1}^n \frac{q_f(y_i|\theta_i) \widehat{M}_{f/g,i}{}^{(r)}}{\mathcal{Z}_g(\gamma_i) B_{f/g,i}} h(\beta, \beta') \pi(\beta)} \right\} \\ &= \min \left\{ 1, \frac{\prod_{i=1}^n q_f(y_i|\theta'_i) \mathcal{Z}_g(\gamma_i) \widehat{M}'_{f/g,i}{}^{(r)} B_{f/g,i} h(\beta', \beta) \pi(\beta')}{\prod_{i=1}^n q_f(y_i|\theta_i) \mathcal{Z}_g(\gamma'_i) \widehat{M}_{f/g,i}{}^{(r)} B'_{f/g,i} h(\beta, \beta') \pi(\beta)} \right\} \end{aligned} \quad (18)$$

This acceptance ratio is entirely tractable once the estimates $\widehat{M}_{f/g,i}^{(r)}$ have been computed by running rejection sampling at each θ'_i for the GIMH algorithm or at both θ'_i and θ''_i for the MCWM algorithm.

5 Results

We present two examples of COM-Poisson modelling. The first example is a retail inventory dataset from [Shmueli et al. \(2005\)](#) which contains no covariate information. This first example is a toy example that demonstrates how the BIC can be approximated using the unbiased likelihood estimate. The second example is a takeover bids dataset which contains covariates which we model using a COM-Poisson regression model. The second example will firstly show the reduction in computation time available using our sampler and also the ability to choose different regression models using the unbiased likelihood estimate.

5.1 Inventory Data

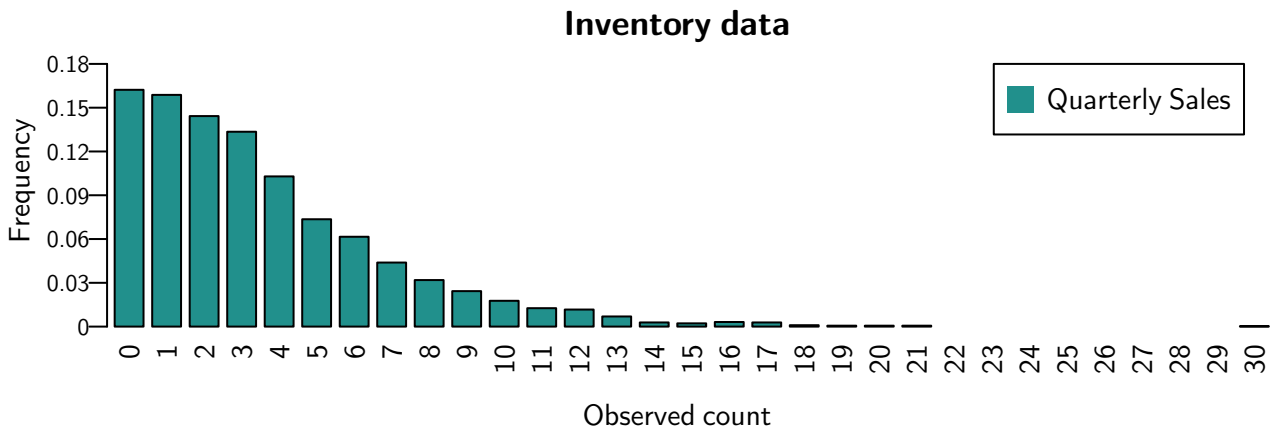


Figure 4: Inventory data: Sales of a particular item of clothing from a well-known clothing brand. The sample mean sales is 3.56 with sample variance 11.31.

Inventory data or stock count data is of high importance to retailers as it facilitates how they can make decisions regarding future stock levels. The counts we observe in this data are the quarterly sales counts of a particular item of clothing in each store across $n = 3,168$ retail stores. A frequency barplot of the data is shown in Figure 4. The maximum quarterly sales of the clothing item sold was 30 in one store. No stores sold between 22 and 29 items and many stores had 0 sales of the clothing item. The sample mean sales is 3.56, sample variance 11.31 and the sample dispersion index is 3.18, thus overdispersed. This data was analysed by [Shmueli et al. \(2005\)](#) using a maximum likelihood approach.

Since this data contains no extra covariate information, we model μ and ν directly rather than considering them as functions of covariates. This could be viewed as using the link function with only intercept where $\mu = \exp(\beta_{\mu,0})$ and $\nu = \exp(\beta_{\nu,0})$ but we consider it easier to model μ and ν . We place $\text{Gamma}(1, 1)$ and $\text{Gamma}(0.0625, 0.25)$ priors on μ and ν , respectively. We run the exchange algorithm for 200,000 iterations and discard the first 5,000 as a burn-in. Each parameter μ and ν were updated using a single-site update. The results are shown in Table 1 with the target acceptance rate of 44% for single-site updates achieved.

We now compare the COM-Poisson distribution to the Poisson distribution using $\widehat{\text{BIC}}^{(r)}$ with $r = 5000$. The BIC for the Poisson distribution can be calculated exactly since the likelihood is tractable. The comparison of both models is shown in Table 2 where the COM-Poisson is shown to outperform the Poisson under BIC.

Table 1: Inventory data posterior estimates: Posterior estimates from the exchange algorithm. The posterior estimate for ν shows overdispersion in data.

	Posterior Mean	Posterior SD	Accept Rate
μ	0.8243	0.1444	44.39%
ν	0.1286	0.0119	41.68%

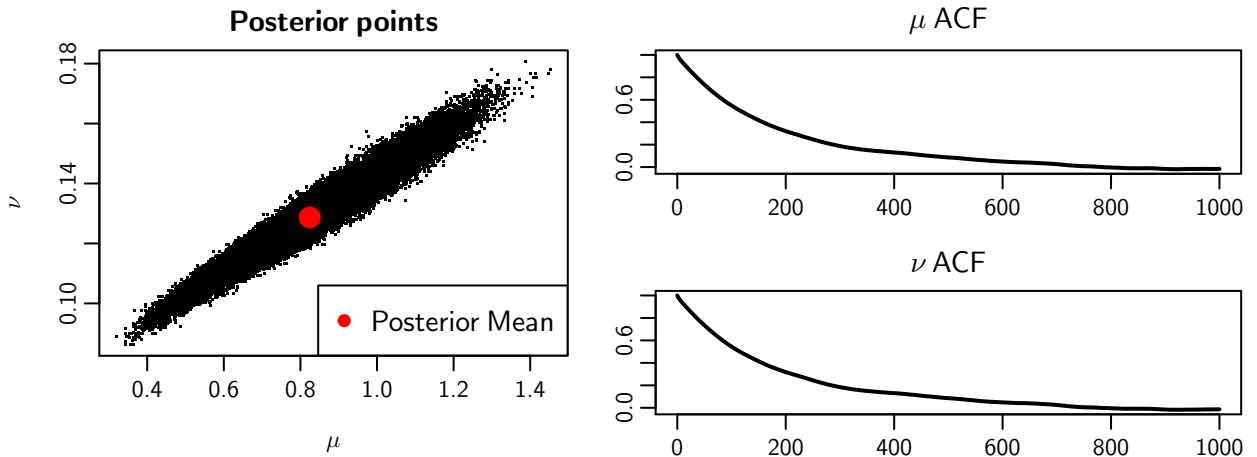


Figure 5: Inventory data posterior visualisation: The distribution of MCMC points for (μ, ν) is shown along with posterior mean (left). The estimated autocorrelation function for μ (top right) and ν (bottom right). There is high correlation between μ and ν in this example.

Table 2: Inventory data: The COM-Poisson model significantly outperforms the Poisson model using the BIC. The estimate BIC for the COM-Poisson was calculated using $\widehat{\text{BIC}}^{(r)}$ with $r = 5000$.

	Poisson	COM-Poisson
BIC	17927.68	15067.39

5.1.1 Comparing our method to an approximate (truncated) likelihood approach

We now compare our exact approach to instead carrying out MCMC with the true likelihood in the posterior distribution replaced by an approximate likelihood using a truncated normalising constant to m terms

$$\mathcal{Z}_f(\mu, \nu) \approx \hat{\mathcal{Z}}_f^m(\mu, \nu) = \sum_{y=0}^{m-1} \left(\frac{\mu^y}{y!} \right)^\nu. \quad (19)$$

Such a plug-in approximation of the likelihood leads to one carrying out MCMC on an approximate posterior distribution. In turn, this leads to the important issue of understanding the ergodicity of the resulting Markov chain, an issue which is attracting much interest in the literature, for example, the noisy MCMC approach of [Alquier et al. \(2016\)](#) and the Russian roulette approach of [Lyne et al. \(2015\)](#). Therefore, an immediate question arising from using an approximation of the normalising constant, as detailed above (19), would be to understand the convergence properties of the resulting noisy MCMC algorithm although it is beyond the scope of this paper. Of course, the exact sampling algorithm developed in this paper avoids such approximation.

Instead, we carry out a small study to explore empirically estimated posterior means and standard deviations arising from the noisy MCMC algorithm derived by replacing the true likelihood with the approximate likelihood with truncated normalising constant. To do this, we again analyse the Inventory dataset experimenting with two choices for the number of terms m in the truncated normalising constant (19). In particular, we set $m = 3,300$ as this lead to a similar computation time for one transition of the exact MCMC algorithm. Table 3 shows that the results of this experiment are in good agreement with the results of the exact MCMC algorithm detailed in Table 1. We also note that the truncation of the normalising constant to $m = 100$ terms yielded identical results to that of $m = 3,300$, but resulted in an 86% decrease in the time required by our exact sampler. However these results disguise the fact that $m = 100$ or even $m = 3,300$ terms are not sufficient to give an accurate likelihood approximation for all regions of the parameter space. To investigate, we began the algorithm at the starting point ($\mu = 500, \nu = 0.0001$) and found that for both cases $m = 100$ and $m = 3,300$ the resulting noisy MCMC algorithms failed to converge whereas our algorithm converged to the posterior estimates as in Table 1. The failure to converge is due to an insufficient number of terms in the truncation, meaning the truncation is only practical in some but not all areas of the parameter space. Of course, this has immediate implications when one implements this truncated normalised constant within an MCMC algorithm, as above, since the algorithm may visit areas of the parameter space for which the truncation is inaccurate. We also emphasise that one would not know in advance the appropriate value of m to use. This is illustrated in Figure 6 where we explore the accuracy of (19) for different values of μ and ν .

Table 3: Truncated normalising constant estimates: Posterior estimates for using a truncated normalising constant in the likelihood with $m = 3,300$ terms.

	Posterior Mean	Posterior SD	Accept Rate	m	Time
μ	0.8265	0.1448	44.05%	3,300	1.00
ν	0.1282	0.0119	45.71%	3,300	1.00

5.2 Takeover bids COM-Poisson regression model

We consider the dataset from [Jaggia and Thosar \(1993\)](#) on the number of bids received by U.S. firms that were targets of bid offers during the period 1978-85 and taken over within 1 year of the

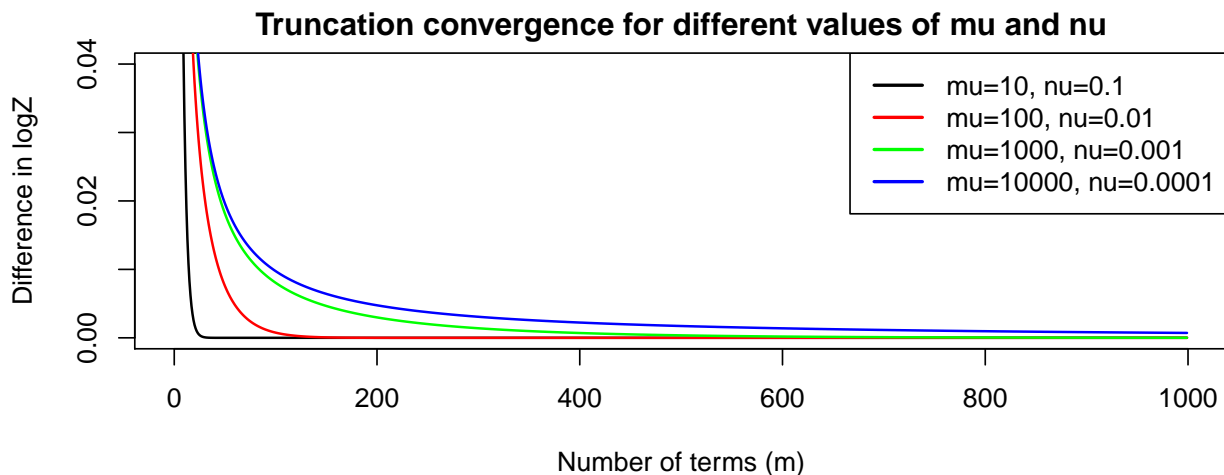


Figure 6: Truncated normalising constant convergence: This figure shows that the number of terms required for a truncated normalising constant to converge changes with the value of μ and ν . The y-axis represents the difference in $\log\left(\hat{Z}_f^m(\mu, \nu)\right)$ for two consecutive values of m . This should converge to 0 but for large μ and small ν it can take over 1,000 terms to converge.

initial offer. The response variable is the count of takeover bids excluding the initial bid (NUMBIDS) filed for the particular company. The covariates available include variables which capture the firm-specific characteristics and defensive action taken by the firm in response to their initial takeover bid. A detailed list of available predictors with is given in Appendix B. This data has previously been analysed using Poisson regression by Jaggia and Thosar (1993) and series expansions by Cameron and Johansson (1997). We reconsider this analysis by comparing two Poisson regression models to three COM-Poisson regression models. The three models considered in our analysis are listed in Table 4. We performed a prior variable selection by fitting a Poisson GLM by maximum likelihood to the data to establish which predictors to consider in the Bayesian GLM analysis. The first model and simplest considers NUMBIDS in a Poisson GLM where the linear predictor of $\log(\mu_i)$ contains an intercept and two predictors (BIDPREM and WHTKNIGHT). The second model adds an extra linear predictor, SIZE, to Model 1. For COM-Poisson regression we consider 3 models, Model 3 moves the covariate SIZE into the ν_i link function. Model 4 considers dropping BIDPREM from the μ_i link function and finally Model 5 considers adding another covariate FINREST into the ν_i link function. For simplicity we use β to denote coefficients within the μ_i link function and ρ to denote coefficients within the ν_i link function.

Priors for β and ρ were set as $\mathcal{N}(0, 5^2)$. The exchange algorithm was run for 100,000 observations with a burnin of 10,000 draws discarded. Each coefficient was updated using a single-site update with a random walk proposal. The proposal variance of the random walk was tuned to achieve an acceptance rate of 44% for each site. The posterior results for all models are presented in Table 5. The BIC was calculated using (15) with r set to 5000 to give an accurate estimate of the BIC. Model 5 was selected as the best model with the lowest BIC of 386.40. All 3 COM-Poisson regression models are selected as better than the Poisson models, Model 1 and Model 2 suggesting that including covariates in the dispersion link function is worthwhile.

5.2.1 Comparing sampler efficiency

Following Section 3.2, here we compare the efficiency of Algorithm 2 to the rejection sampler of Chaniyalidis et al. (2018) in the context of the takeover bids dataset. To do so, we examine the number of MCMC draws per second resulting from Algorithm 2 and from the rejection sampler of Chaniyalidis et al. (2018) for each COM-Poisson models (Models 3, 4 and 5) outlined in Table 4. Our sampler provides a threefold increase in the number of MCMC draws per second in each of the COM-Poisson

Table 4: Takeover bids candidate models: The k column gives the number of parameters estimated which will be used directly in the calculation of the BIC. Descriptions of all covariates are given in Appendix B.

Model	Response (y_i)	Linear Predictor(s)	k
1 (Poisson)	NUMBIDS	$\log(\mu_i) = \beta_0 + \beta_1\text{BIDPREM} + \beta_2\text{WHTKNGHT}$	3
2 (Poisson)	NUMBIDS	$\log(\mu_i) = \beta_0 + \beta_1\text{BIDPREM} + \beta_2\text{WHTKNGHT} + \beta_3\text{SIZE}$	4
3 (COM-Poisson)	NUMBIDS	$\log(\mu_i) = \beta_0 + \beta_1\text{BIDPREM} + \beta_2\text{WHTKNGHT}$ $\log(\nu_i) = \rho_0 + \rho_1\text{SIZE}$	5
4 (COM-Poisson)	NUMBIDS	$\log(\mu_i) = \beta_0 + \beta_2\text{WHTKNGHT}$ $\log(\nu_i) = \rho_0 + \rho_1\text{SIZE}$	4
5 (COM-Poisson)	NUMBIDS	$\log(\mu_i) = \beta_0 + \beta_1\text{BIDPREM} + \beta_2\text{WHTKNGHT}$ $\log(\nu_i) = \rho_0 + \rho_1\text{SIZE} + \rho_2\text{FINREST}$	6

Table 5: Takeover bids: Posterior parameter estimates are shown for all 5 models. The lowest BIC is shown in **bold**. The COM-Poisson regression models provide the best fit to the data according to BIC with all COM-Poisson models having a BIC lower than the Poisson models.

Parameter	Model 1	Model 2	Model 3	Model 4	Model 5
	Mean (SD)	Mean (SD)	Mean (SD)	Mean (SD)	Mean (SD)
$\hat{\beta}_0$	1.130 (0.505)	1.063 (0.532)	1.077 (0.384)	0.329 (0.100)	0.354 (0.091)
$\hat{\beta}_1$	-0.728 (0.368)	-0.713 (0.382)	-0.553 (0.281)	-	-
$\hat{\beta}_2$	0.583 (0.152)	0.576 (0.152)	0.458 (0.110)	0.463 (0.111)	0.431 (0.103)
$\hat{\beta}_3$	-	0.035 (0.017)	-	-	-
$\hat{\rho}_0$	-	-	0.674 (0.175)	0.646 (0.175)	0.789 (0.179)
$\hat{\rho}_1$	-	-	-0.171 (0.051)	-0.174 (0.052)	-0.176 (0.049)
$\hat{\rho}_2$	-	-	-	-	-0.952 (0.448)
$\widehat{\text{BIC}}(r = 5000)$	397.49	398.32	386.89	386.98	386.40
Rank	4	5 (Worst)	2	3	1 (Best)

models agreeing with the experiment presented in Figure 3.

Table 6: Algorithm efficiency comparison: This table shows the number of MCMC draws per second for our rejection sampler (Algorithm 2) compared with the rejection sampler of Chaniavidis et al. (2018). Our algorithm is at least 3 times faster for all 3 models.

	Model 3	Model 4	Model 5
MCMC draws per sec. (Algorithm 2)	9,973	12,534	9,297
MCMC draws per sec. (Chaniavidis et al. (2018))	2,866	3,585	2,828
Relative efficiency	3.480	3.496	3.287

5.3 Pseudo-marginal MCMC results

We choose Model 5 in Table 4 to implement the GIMH pseudo-marginal MCMC algorithm using the unbiased likelihood estimator (14) in the acceptance ratio. The unbiased likelihood estimator was calculated for 5 scenarios with $r = 1, 5, 10, 50, 100$. The posterior density estimates for each of the pseudo-marginal MCMC runs are shown in Table 7 along with the CPU time and multivariate effective sample size (mESS) (Vats et al., 2019). The mESS is a generalisation of the effective sample size (ESS), which captures the cross-autocorrelation between parameters in the MCMC output.

$$\text{mESS} = n_{\text{mcmc}} \left(\frac{|\Lambda|}{|\Sigma|} \right)^{1/p}$$

where n_{mcmc} is the number of MCMC draws, Λ is the sample autocovariance matrix of the MCMC output and Σ is the true covariance of the posterior which can be estimated by a batch-means method described in Vats et al. (2019). Figure 7 compares the pseudo-marginal ($r = 100$) posterior MCMC chain, density and autocorrelation with the exchange algorithm. There is close agreement between the two methods, however the CPU time for the pseudo-marginal algorithm is dramatically increased.

Table 7: Pseudo-marginal MCMC results: The results for the pseudo-marginal MCMC algorithm run with the unbiased likelihood estimator.

	Exchange	Pseudo-marginal MCMC - GIMH				
		$r = 1$	$r = 5$	$r = 10$	$r = 50$	$r = 100$
	Mean (σ)	Mean (σ)	Mean (σ)	Mean (σ)	Mean (σ)	Mean (σ)
$\hat{\beta}_0$	0.354 (0.09)	0.337 (0.08)	0.361 (0.09)	0.354 (0.09)	0.354 (0.09)	0.356 (0.09)
$\hat{\beta}_2$	0.431 (0.10)	0.446 (0.10)	0.424 (0.10)	0.432 (0.10)	0.433 (0.10)	0.431 (0.11)
$\hat{\rho}_0$	0.789 (0.18)	0.785 (0.14)	0.796 (0.17)	0.790 (0.17)	0.794 (0.18)	0.793 (0.18)
$\hat{\rho}_1$	-0.176 (0.05)	-0.176 (0.05)	-0.172 (0.05)	-0.178 (0.05)	-0.175 (0.05)	-0.175 (0.05)
$\hat{\rho}_2$	-0.952 (0.45)	-0.817 (0.50)	-0.981 (0.44)	-0.944 (0.45)	-0.962 (0.45)	-0.955 (0.45)
CPU(s)	10.617	46.257	188.050	355.001	1668.14	3512.57
mESS	4,962	345	904	2701	8,820	10,922

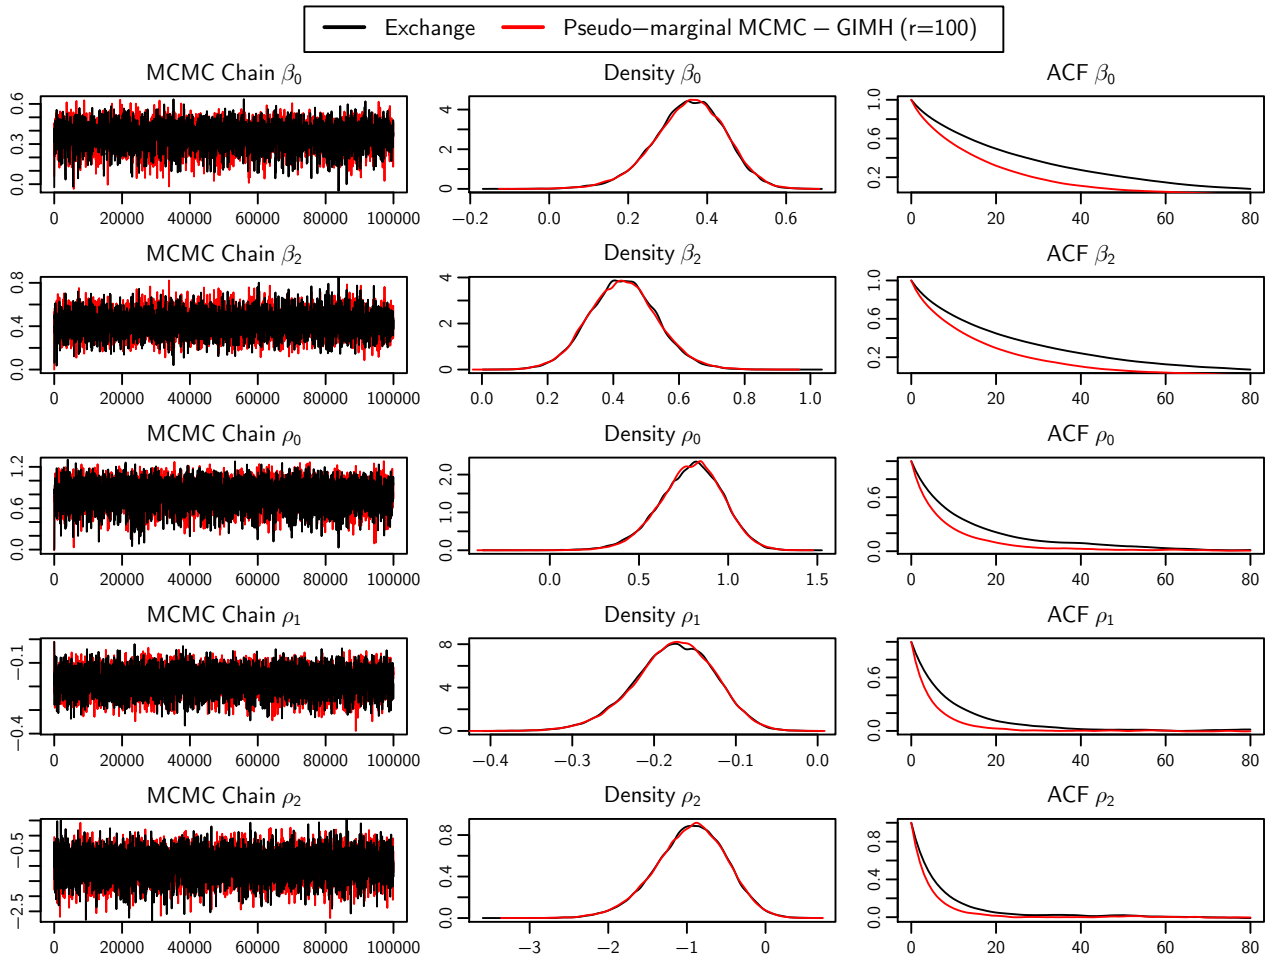


Figure 7: Comparison of the pseudo-marginal MCMC algorithm and the exchange algorithm: Both algorithms show agreement in posterior density estimates. Pseudo-marginal MCMC shows a reduction in autocorrelation compared to the exchange algorithm for all parameters, however, this is at the expense of increased computational time, as shown in Table 7.

6 Discussion

This paper provides a new rejection sampler to sample from the COM-Poisson distribution. This rejection sampler allows for faster parameter inference to be performed in COM-Poisson regression models compared with the sampler of [Chanialidis et al. \(2018\)](#). Our rejection sampling algorithm shows that one must consider both the rejection rate of the enveloping distribution and the sampling time from the enveloping distribution in order to construct an efficient rejection sampler. We have also shown how the number of rejected proposals within rejection sampling can be used to construct an unbiased intractable likelihood estimator. This estimator can be used to perform model selection using BIC but future work could include this estimator within other Bayesian model choice strategies. Model selection for doubly-intractable problems such as COM-Poisson regression is a difficult problem and our approach offers one solution to perform model selection. The unbiased likelihood estimator was shown to work well when used within a pseudo-marginal MCMC algorithm ([Andrieu and Roberts, 2009](#)) as it provided an unbiased estimate of the acceptance ratio. This unbiased likelihood estimator could be constructed for other intractable models where a rejection sampling algorithm exists to sample from the likelihood, presenting opportunities for future research.

Acknowledgements

The Insight Centre for Data Analytics is supported by Science Foundation Ireland under Grant Number 12/RC/2289_P2.

References

- Ahrens, J. H. and U. Dieter (1982). Computer generation of Poisson deviates from modified Normal distributions. *ACM Transactions on Mathematical Software* 8(2), 163–179.
- Alquier, P., N. Friel, R. Everitt, and A. Boland (2016). Noisy Monte Carlo: convergence of Markov chains with approximate transition kernels. *Statistics and Computing* 26(1), 29–47.
- Andrieu, C. and G. O. Roberts (2009). The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics* 37(2), 697–725.
- Beaumont, M. A. (2003). Estimation of population growth or decline in genetically monitored populations. *Genetics* 164(3), 1139–1160.
- Beskos, A., O. Papaspiliopoulos, G. O. Roberts, and P. Fearnhead (2006). Exact and computationally efficient likelihood-based estimation for discretely observed diffusion processes (with discussion). *Journal of the Royal Statistical Society, Series B* 68(3), 333–382.
- Caimo, A. and N. Friel (2011). Bayesian inference for exponential random graph models. *Social Networks* 33(1), 41–55.
- Cameron, A. C. and P. Johansson (1997). Count data regression using series expansions: With applications. *Journal of Applied Econometrics* 12(3), 203–223.
- Casella, G. and C. P. Robert (1996). Rao-blackwellisation of sampling schemes. *Biometrika* 83(1), 81–94.
- Chaniialidis, C., L. Evers, T. Neocleous, and A. Nobile (2018). Efficient Bayesian inference for COM-Poisson regression models. *Statistics and Computing* 28, 595–608.
- Conway, R. W. and W. L. Maxwell (1962). A queuing model with state dependent service rates. *Journal of Industrial Engineering* 12(2), 132–136.
- Georgoulas, A., J. Hillston, and G. Sanguinetti (2017). Unbiased Bayesian inference for population Markov jump processes via random truncations. *Statistics and Computing* 27(4), 991–1002.
- Geyer, C. and E. A. Thompson (1992). Constrained Monte Carlo maximum likelihood for dependent data. *Journal of the Royal Statistical Society, Series B* 54(3), 657 – 699.
- Gilks, W. R. and P. Wild (1992). Adaptive rejection sampling for Gibbs sampling. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 41(2), 337–348.
- Guikema, S. D. and J. P. Goffelt (2008). A flexible count data regression model for risk analysis. *Risk Analysis* 28(1), 213–223.
- Hilbe, J. M. (2011). *Negative binomial regression*. Cambridge University Press.
- Jaggia, S. and S. Thosar (1993). Multiple bids as a consequence of target management resistance: A count data approach. *Review of Quantitative Finance and Accounting* 3(4), 447–457.
- Lyne, A.-M., M. Girolami, Y. Atchadé, H. Strathmann, and D. Simpson (2015, 11). On Russian roulette estimates for Bayesian inference with doubly-intractable likelihoods. *Statistical Science* 30(4), 443–467.
- Mascagni, M. and A. Srinivasan (2000). Algorithm 806: SPRNG: A scalable library for pseudorandom number generation. *ACM Transactions on Mathematical Software* 26(3), 436–461.

- Møller, J., A. N. Pettitt, R. Reeves, and K. K. Berthelsen (2006). An efficient Markov chain Monte Carlo method for distributions with intractable normalising constants. *Biometrika* 93(2), 451–458.
- Murray, I., Z. Ghahramani, and D. J. C. MacKay (2006). MCMC for doubly-intractable distributions. In *Proceedings of the 22nd Annual Conference on Uncertainty in Artificial Intelligence (UAI-06)*, pp. 359–366. AUAI Press.
- Propp, J. G. and D. B. Wilson (1996). Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Structures and Algorithms* 9(1&2), 223–252.
- Rao, V., L. Lin, and D. B. Dunson (2016). Data augmentation for models based on rejection sampling. *Biometrika* 103(2), 319–335.
- Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics* 6(2), 461–464.
- Sellers, K. F. and G. Shmueli (2010). A flexible regression model for count data. *The Annals of Applied Statistics* 4(2), 943–961.
- Shmueli, G., T. P. Minka, J. B. Kadane, S. Borle, and P. Boatwright (2005). A useful distribution for fitting discrete data: Revival of the Conway-Maxwell-Poisson distribution. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 54(1), 127–142.
- Shmueli, G., R. P. Russo, and W. Jank (2007). The BARISTA: A model for bid arrivals in online auctions. *The Annals of Applied Statistics* 1(2), 412–441.
- Spiegelhalter, D. J., N. G. Best, B. P. Carlin, and A. Van Der Linde (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 64(4), 583–639.
- Vats, D., J. M. Flegal, and G. L. Jones (2019). Multivariate Output Analysis for Markov chain Monte Carlo. *Biometrika* 106(2), 321–337.
- von Neumann, J. (1951). Various Techniques Used in Connection with Random Digits. *National Bureau of Standards Applied Mathematics Series 12*, 36–38.
- Wei, C. and I. Murray (2017, 20–22 Apr). Markov Chain Truncation for Doubly-Intractable Inference. In A. Singh and J. Zhu (Eds.), *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, Volume 54 of *Proceedings of Machine Learning Research*, Fort Lauderdale, FL, USA, pp. 776–784. PMLR.

Appendices

A COM-Poisson rejection sampler Proof

Proof. There are two cases to consider depending on the value of ν .

(In this proof we use the Iverson bracket $[\cdot]$, where $[x] = 1$ if the enclosed statement x is true and 0 if the statement x is false.)

Case 1: When $[\nu \geq 1] = 1$

The enveloping bound as illustrated in Figure 1 (left) is the ratio of the COM-Poisson mass function to a Poisson mass function

$$B_{f/g}^{[\nu \geq 1]} = \sup_y \left\{ \frac{\left(\frac{\mu^y}{y!}\right)^\nu}{\frac{\mu^y}{y!}} \right\} = \sup_y \left\{ \left(\frac{\mu^y}{y!}\right)^{\nu-1} \right\}.$$

To find the supremum, assume that the supremum occurs at the point y_m and by the unimodality of the COM-Poisson distribution, it will satisfy

$$\begin{aligned} \left(\frac{\mu^{y_m}}{y_m!}\right)^{\nu-1} &\geq \left(\frac{\mu^{y_m+h}}{(y_m+h)!}\right)^{\nu-1} \\ \iff \frac{\mu^{y_m}}{y_m!} &\geq \frac{\mu^{y_m+h}}{(y_m+h)!}, \end{aligned} \quad (20)$$

for $h \in \{-y_m, -y_m + 1, \dots\}$. To find the value of y_m that will satisfy this collection of inequalities over h it is enough to consider the dominant inequalities at $h = \pm 1$,

$$\mu - 1 \leq y_m \leq \mu \quad \implies y_m = \lfloor \mu \rfloor. \quad (21)$$

Therefore if y_m satisfies (21), it will also satisfy (20) for all $h \in \{-y_m, -y_m + 1, \dots\}$. Thus the bound becomes

$$B_{f/g}^{[\nu \geq 1]} = \left(\frac{\mu^{\lfloor \mu \rfloor}}{\lfloor \mu \rfloor!}\right)^{\nu-1},$$

which is a tractable bound and for the special case of Poisson ($\nu = 1$), then $B_{f/g}^{[\nu=1]} = 1$. In the case of integer μ this bound will also ensure coverage of the dual mode at μ and $\mu - 1$ (see Figure 1).

Case 2: When $[\nu < 1] = 1$

The enveloping bound as illustrated in Figure 1 (right) is the ratio of the COM-Poisson mass function to a geometric mass function

$$B_{f/g}^{[\nu < 1]} = \sup_y \left\{ \frac{\left(\frac{\mu^y}{y!}\right)^\nu}{(1-p)^y p} \right\} = \frac{1}{p} \sup_y \left\{ \frac{(\mu^\nu)^y}{(1-p)^y y!^\nu} \right\}.$$

For now we assume p is unknown and we prove the general bound for any p . A discussion of which value of p to choose in practice will be given after this proof. As in **Case 1**, assume that the supremum occurs at the point y_m , which by the unimodality of the COM-Poisson distribution, will satisfy

$$\frac{1}{p} \frac{(\mu^\nu)^{y_m}}{(1-p)^{y_m} y_m!^\nu} \geq \frac{1}{p} \frac{(\mu^\nu)^{y_m+h}}{(1-p)^{y_m+h} (y_m+h)!^\nu}$$

$$\iff \frac{1}{y_m!} \geq \frac{\mu^h}{(1-p)^{h/v}(y_m+h)!}, \quad (22)$$

for $h \in \{-y_m, -y_m + 1, \dots\}$. To find the value of y_m that will satisfy this collection of inequalities over h it is enough to consider again only the inequalities at $h = \pm 1$. This leads to

$$\frac{\mu}{(1-p)^{1/\nu}} - 1 \leq y_m \leq \frac{\mu}{(1-p)^{1/\nu}} \implies y_m = \left\lfloor \frac{\mu}{(1-p)^{1/\nu}} \right\rfloor. \quad (23)$$

As for **Case 1**, if y_m satisfies (23) it will also satisfy (22) for all $h \in \{-y_m, -y_m + 1, \dots\}$. Thus the bound becomes

$$B_{f/g}^{[\nu < 1]} = \frac{1}{p} \frac{\mu \binom{\nu \left\lfloor \frac{\mu}{(1-p)^{1/\nu}} \right\rfloor}{\left\lfloor \frac{\mu}{(1-p)^{1/\nu}} \right\rfloor}}{(1-p)^{\left\lfloor \frac{\mu}{(1-p)^{1/\nu}} \right\rfloor} \left(\left\lfloor \frac{\mu}{(1-p)^{1/\nu}} \right\rfloor ! \right)^\nu}.$$

□

B Takeover bids data

The variables available for the takeover bids dataset are show in Table 8 along with a brief description and possible values. This data was first analysed in [Jaggia and Thosar \(1993\)](#).

Table 8: Takeover bids: Description of variables and possible values.

Name	Description	Values	Type
NUMBIDS	Number of takeover bids received after initial bid	0, 1, 2, ...	Response
LEGLREST	Did management try legal defense by lawsuit?	1=YES, 0=NO	Predictor
REALREST	Did management propose changes in asset structure?	1=YES, 0=NO	Predictor
FINREST	Did management propose changes in ownership structure?	1=YES, 0=NO	Predictor
WHTKNIGHT	Did management invite friendly third-party bid?	1=YES, 0=NO	Predictor
REGULATN	Did federal regulators intervene?	1=YES, 0=NO	Predictor
BIDPREM	Bid price divided by price 14 working days before bid.	> 0	Predictor
INSTHOLD	Percentage of stock held by institutions.	[0, 1]	Predictor
SIZE	Total book value of assets	USD billions	Predictor