

Bayesian Inference for the Negative Binomial Distribution via Polynomial Expansions

Eric T. BRADLOW, Bruce G. S. HARDIE, and Peter S. FADER

To date, Bayesian inferences for the negative binomial distribution (NBD) have relied on computationally intensive numerical methods (e.g., Markov chain Monte Carlo) as it is thought that the posterior densities of interest are not amenable to closed-form integration. In this article, we present a “closed-form” solution to the Bayesian inference problem for the NBD that can be written as a sum of polynomial terms. The key insight is to approximate the ratio of two gamma functions using a polynomial expansion, which then allows for the use of a conjugate prior. Given this approximation, we arrive at closed-form expressions for the moments of both the marginal posterior densities and the predictive distribution by integrating the terms of the polynomial expansion in turn (now feasible due to conjugacy). We demonstrate via a large-scale simulation that this approach is very accurate and that the corresponding gains in computing time are quite substantial. Furthermore, even in cases where the computing gains are more modest our approach provides a method for obtaining starting values for other algorithms, and a method for data exploration.

Key Words: Beta-prime distribution; Empirical Bayes methods; Pearson Type VI distribution.

1. INTRODUCTION

For the past 80 years, numerous researchers have modeled count data assuming the individual-level count, X_i , is distributed $\text{Poisson}(\lambda_i)$ and the rate parameter λ_i is distributed $\text{gamma}(r, \alpha)$. In other words, the marginal distribution of X_i follows a negative binomial distribution (NBD):

$$\begin{aligned} p(x_i|r, \alpha) &= \int_0^\infty p(x_i|\lambda_i)p(\lambda_i|r, \alpha) d\lambda_i \\ &= \int_0^\infty \frac{\lambda_i^{x_i} e^{-\lambda_i}}{x_i!} \frac{\alpha^r \lambda_i^{r-1} e^{-\alpha\lambda_i}}{\Gamma(r)} d\lambda_i \\ &= \frac{\Gamma(r+x_i)}{\Gamma(r)x_i!} \left(\frac{\alpha}{\alpha+1}\right)^r \left(\frac{1}{\alpha+1}\right)^{x_i}. \end{aligned} \quad (1.1)$$

Eric T. Bradlow is Associate Professor of Marketing and Statistics (E-mail: ebradlow@wharton.upenn.edu), and Peter S. Fader is Associate Professor of Marketing, The Wharton School of the University of Pennsylvania, Philadelphia, PA 19104-6371. Bruce G. S. Hardie is Assistant Professor of Marketing, London Business School, London, NW1 4SA, UK.

©2002 American Statistical Association, Institute of Mathematical Statistics,
and Interface Foundation of North America
Journal of Computational and Graphical Statistics, Volume 11, Number 1, Pages 189–201

The first application of this characterization of the NBD was presented by Greenwood and Yule (1920) to model accident statistics. It has subsequently been used to model phenomena as diverse as the purchasing of consumer packaged goods (Ehrenberg 1959), salesperson productivity (Carroll, Lee, and Rao 1986), and library circulation (Burrell 1990).

The standard inference approach has been to estimate r and α using maximum likelihood from the marginal distribution of X_i given in (1.1). This so-called empirical Bayes method then bases inferences on the estimated marginal distribution $p(x_i|\hat{r}, \hat{\alpha})$. A common application is to make predictions of future counts, Y_i , conditional on past behavior; for example, $E(Y_i|x_i, \hat{r}, \hat{\alpha}) = (\hat{r} + x_i)/(\hat{\alpha} + 1)$. These traditional analyses treat r and α as fixed and known at \hat{r} and $\hat{\alpha}$. A key shortcoming of this approach is that it ignores the variability (estimation uncertainty) of \hat{r} and $\hat{\alpha}$ as estimates of r and α . Furthermore, it does not allow the incorporation of any prior information in the analysis, such as insights that may have been derived from previous studies or from subjective information.

A natural solution to these problems is to adopt a Bayesian approach in which prior distributions are specified for the NBD model parameters r and α . In recent years, several researchers have used Bayesian inference for the NBD (e.g., Deely and Smith 1998; Klugman 1992; Schlüter, Deely, and Nicholson 1997). However, none of these have obtained *closed-form inferences*, as we do here. To demonstrate these past shortcomings, and our approach, we set up the general problem formulation as follows.

Let the vector $\mathbf{x} = (x_1, \dots, x_n)$ represent the observed (count) data for n subjects measured across a time period of unit length. Assuming x_1, \dots, x_n are generated iid NBD(r, α), we have

$$\begin{aligned} p(\mathbf{x}|r, \alpha) &= \prod_{i=1}^n \left[\frac{\Gamma(r+x_i)}{\Gamma(r)x_i!} \left(\frac{\alpha}{\alpha+1} \right)^r \left(\frac{1}{\alpha+1} \right)^{x_i} \right] \\ &= \left[\prod_{i=1}^n \frac{\Gamma(r+x_i)}{\Gamma(r)x_i!} \right] \left(\frac{\alpha}{\alpha+1} \right)^{nr} \left(\frac{1}{\alpha+1} \right)^{\sum_{i=1}^n x_i}. \end{aligned} \quad (1.2)$$

Specifying a prior $p(r, \alpha)$, it follows that the k th marginal posterior moment of r given data \mathbf{x} is given by

$$\begin{aligned} E(r^k|\mathbf{x}) &= \int r^k p(r|\mathbf{x}) dr \\ &= \iint r^k p(r, \alpha|\mathbf{x}) d\alpha dr, \quad \text{and then from Bayes' rule} \\ &= \frac{\iint r^k p(\mathbf{x}|r, \alpha) p(r, \alpha) d\alpha dr}{\iint p(\mathbf{x}|r, \alpha) p(r, \alpha) d\alpha dr}. \end{aligned} \quad (1.3)$$

Similarly, the k th marginal posterior moment of α is given by

$$\begin{aligned} E(\alpha^k|\mathbf{x}) &= \int \alpha^k p(\alpha|\mathbf{x}) d\alpha \\ &= \iint \alpha^k p(r, \alpha|\mathbf{x}) dr d\alpha \\ &= \frac{\iint \alpha^k p(\mathbf{x}|r, \alpha) p(r, \alpha) dr d\alpha}{\iint p(\mathbf{x}|r, \alpha) p(r, \alpha) d\alpha dr}. \end{aligned} \quad (1.4)$$

(An equivalent expression can be written for $E(Y_i^k|\mathbf{x})$, the k th moment of the predictive distribution.) Performing the requisite integrations allows the analyst to make the inferences of interest; that is, to obtain parameter estimates, predictions, and so on, from the appropriate marginal posterior distributions.

To highlight the challenges associated with a Bayesian analysis of the NBD, we substitute (1.2) into (1.3), which yields

$$E(r^k|\mathbf{x}) = \frac{\iint r^k \left[\prod_{i=1}^n \frac{\Gamma(r+x_i)}{\Gamma(r)x_i!} \right] \left(\frac{\alpha}{\alpha+1} \right)^{nr} \left(\frac{1}{\alpha+1} \right)^{\sum_{i=1}^n x_i} p(r, \alpha) d\alpha dr}{\iint \left[\prod_{i=1}^n \frac{\Gamma(r+x_i)}{\Gamma(r)x_i!} \right] \left(\frac{\alpha}{\alpha+1} \right)^{nr} \left(\frac{1}{\alpha+1} \right)^{\sum_{i=1}^n x_i} p(r, \alpha) d\alpha dr}. \quad (1.5)$$

It is easy to specify a prior for α such that we can integrate α out of the above expression; such a distribution is well understood and appears in many places including some of the earlier mentioned Bayesian NBD references. To date, however, no researcher has been able to specify a prior for r such that there exists a closed-form solution to the above expression, and likewise for (1.4). [A closed-form solution trivially exists for the degenerate case of a point-mass prior for r , e.g., Klugman (1992).] Therefore, researchers have resorted to computationally intensive numerical methods (e.g., Markov chain Monte Carlo) in order to make the inferences of interest.

In this article, we present a closed-form solution to the Bayesian inference problem for the NBD. We do this by specifying a prior for the model parameters r and α which, after using a polynomial expansion to approximate the ratio of two gamma functions (details provided in Sections 2.1 and 2.2), leads to a closed-form solution for the relevant marginal posterior moments. The remainder of this article is organized as follows. Section 2 develops closed-form expressions for the moments of the marginal posterior densities for both r and α . In Section 3, we present a simulation study that demonstrates the accuracy of our approach and the substantial savings in computing time associated with its use. Specifically, we compare computing time for our approach with that from various MCMC runs for a Bayesian NBD model. Section 4 shows the corresponding moments (using the same method as in Section 2) for the evaluation of the moments of the predictive distribution. In particular, we derive closed-form expressions for the conditional mean and variance. We conclude (Section 5) with a few brief summary comments.

2. MOMENTS OF THE MARGINAL POSTERIOR DENSITIES

Our goal is to come up with closed-form expressions for $E(r^k|\mathbf{x})$ and $E(\alpha^k|\mathbf{x})$. We start by assuming independent priors for r and α with the following marginals:

1. *Marginal prior for α :* Consider a prior of the form

$$p(\alpha|\delta_1, \delta_2) \propto \left(\frac{\alpha}{\alpha+1} \right)^{\delta_1-1} \left(\frac{1}{\alpha+1} \right)^{\delta_2+1}.$$

When $\delta_1 = 1$ and $\delta_2 = -1$, we have an improper uniform prior on α . Specifying the proportionality constant as $1/B(\delta_1, \delta_2)$, we have a beta-prime prior on α (Johnson, Kotz,

and Balakrishnan (JKB) 1995, p. 248):

$$p(\alpha|\delta_1, \delta_2) = \frac{1}{B(\delta_1, \delta_2)} \left(\frac{\alpha}{\alpha+1} \right)^{\delta_1-1} \left(\frac{1}{\alpha+1} \right)^{\delta_2+1}, \quad \delta_1, \delta_2 > 0 \quad (2.1)$$

which is equivalent to assuming $\alpha/(\alpha+1) \sim \text{beta}(\delta_1, \delta_2)$. The mean and variance of the beta-prime distribution are

$$\begin{aligned} E(\alpha) &= \frac{\delta_1}{\delta_2 - 1} \\ \text{var}(\alpha) &= \frac{\delta_1(\delta_1 + \delta_2 - 1)}{(\delta_2 - 1)^2(\delta_2 - 2)}. \end{aligned}$$

If prior information is available for α (from past data or otherwise), then values of δ_1 and δ_2 can be determined by moment matching to arrive at an informative prior.

2. *Marginal prior for r* : Consider a prior of the form

$$p(r|a, b, z_1, z_2) \propto \frac{(r - z_1)^a}{(r - z_2)^b}.$$

When $a = b = 0$, we have an improper uniform prior on r . Specifying the proportionality constant as $[\Gamma(b)(z_1 - z_2)^{b-a-1}]/[\Gamma(b-a-1)\Gamma(a+1)]$, we have a Pearson Type VI prior on r (JKB, p. 344):

$$p(r|a, b, z_1, z_2) = L_1 \frac{(r - z_1)^a}{(r - z_2)^b} \quad \text{where} \quad L_1 = \frac{\Gamma(b)(z_1 - z_2)^{b-a-1}}{\Gamma(b-a-1)\Gamma(a+1)}, \quad (2.2)$$

where $b > a > -1$ and $r \geq z_1 > z_2$. Note that the beta-prime density (2.1) is a Pearson Type VI density with $a = \delta_1 - 1$, $b = \delta_1 + \delta_2$, $z_1 = 0$ and $z_2 = -1$. The mean and variance of this distribution are

$$\begin{aligned} E(r) &= z_1 + \frac{(z_1 - z_2)(a+1)}{b-a-2} \\ \text{var}(r) &= \frac{(z_1 - z_2)^2(a+1)(b-1)}{(b-a-2)^2(b-a-3)}. \end{aligned}$$

(Note: there are errors in JKB for these equations which are corrected here). The four parameters z_1 , z_2 , a , and b can be expressed in terms of the first four moments of r to arrive at an informative prior. The rationale for this choice of marginal prior will be discussed in Section 2.1.

2.1 MOMENTS OF THE MARGINAL POSTERIOR OF r

To derive an expression for the k th moment of the posterior density of r , we first substitute (2.1) into (1.5) and perform the integration over α . As can be seen by the forms in (1.5) and (2.1), the term $(\frac{\alpha}{\alpha+1})^{nr}$ from the NBD combines readily with $(\frac{\alpha}{\alpha+1})^{\delta_1-1}$ from the beta-prime prior for α (likewise for the $\frac{1}{\alpha+1}$ terms). This then yields

$$E(r^k|\mathbf{x}) = \frac{\int r^k \left[\prod_{i=1}^n \frac{\Gamma(r+x_i)}{\Gamma(r)x_i!} \right] \frac{\Gamma(nr+\delta_1)}{\Gamma(nr+\delta_1+C_1-1)} p(r) dr}{\int \left[\prod_{i=1}^n \frac{\Gamma(r+x_i)}{\Gamma(r)x_i!} \right] \frac{\Gamma(nr+\delta_1)}{\Gamma(nr+\delta_1+C_1-1)} p(r) dr}, \quad (2.3)$$

where $C_1 = \sum_{i=1}^n x_i + \delta_2 + 1$.

The task now facing us is to substitute (2.2) into (2.3) and perform the integrations over r . The problem here is that r appears in two ratios of gamma functions in each integral—a form that is not analytically tractable. To get past this hurdle, we will first express each ratio of gamma functions as a polynomial in r , thus making it easy to integrate over the prior on r by integrating the product of the polynomials term by term. The first ratio of gamma functions involving $\Gamma(r + x_i)/\Gamma(r)$ can be expressed exactly as a polynomial in r . The second ratio $\Gamma(nr + \delta_1)/\Gamma(nr + \delta_1 + C_1 - 1)$ is approximated by a polynomial in r . As we describe, by including a sufficiently large number of terms, this approximation can be made as close as necessary.

Consider the first ratio of gamma functions. Letting $x^* = \max(x_1, \dots, x_n)$ and n_j the number of observations (e.g., people) with count $x_i = j$, we have

$$\begin{aligned} \prod_{i=1}^n \frac{\Gamma(r + x_i)}{\Gamma(r)} &= \left(\frac{\Gamma(r+1)}{\Gamma(r)} \right)^{n_1} \left(\frac{\Gamma(r+2)}{\Gamma(r)} \right)^{n_2} \cdots \left(\frac{\Gamma(r+x^*)}{\Gamma(r)} \right)^{n_{x^*}} \\ &= r^{n_1} [r(r+1)]^{n_2} \cdots [r(r+1) \cdots (r+x^*-1)]^{n_{x^*}} \\ &= r^{s_1} (r+1)^{s_2} \cdots (r+x^*-1)^{s_{x^*}}, \end{aligned} \quad (2.4)$$

where $s_i = n_i + n_{i+1} + \cdots + n_{x^*}$.

We expand the second ratio of gamma functions as follows:

$$\begin{aligned} \frac{\Gamma(nr + \delta_1)}{\Gamma(nr + \delta_1 + C_1 - 1)} &= \prod_{j=2}^{C_1} \frac{1}{nr + \delta_1 + C_1 - j} \\ &= \prod_{j=2}^{C_1} \frac{1}{(nr + \delta_1 + C_1)(1 - j/(nr + \delta_1 + C_1))} \\ &= y^{1-C_1} \prod_{j=2}^{C_1} \frac{1}{1 - j/y} \\ &= y^{1-C_1} \prod_{j=2}^{C_1} \left(1 + \frac{j}{y} + \frac{j^2}{y^2} + \cdots \right), \end{aligned}$$

where $y = nr + \delta_1 + C_1$. Taking the first m terms of the infinite series $(1 + \frac{j}{y} + \frac{j^2}{y^2} + \cdots)$, and then the first m terms of the corresponding product of these polynomials, we obtain the key approximation:

$$\frac{\Gamma(nr + \delta_1)}{\Gamma(nr + \delta_1 + C_1 - 1)} \approx y^{1-C_1} \left(1 + \frac{U_1}{y} + \frac{U_2}{y^2} + \cdots + \frac{U_m}{y^m} \right),$$

where the polynomial coefficients U_1, \dots, U_m can be calculated by recursive polynomial multiplication. (Note: computing the U 's is simply done by multiplying each of the polynomial terms involving y in turn and collecting terms. The four lines of code in Fortran to do this are available from the authors upon request. In addition, standard software packages such as Mathematica will perform this calculation directly.) Letting $K_1 = (\delta_1 + C_1)/n$, we

have

$$\frac{\Gamma(nr + \delta_1)}{\Gamma(nr + \delta_1 + C_1 - 1)} \approx [n(r + K_1)]^{1-C_1} \sum_{l=0}^m \frac{U_l}{[n(r + K_1)]^l}, \quad (2.5)$$

where $U_0 = 1$. [This generalizes to m terms the one-term approximation for the ratio of two gamma functions given in Johnson, Kotz, and Kemp (1992, equation 1.33) as well as the two-term expansion given in Abramowitz and Stegun (1972, equation 6.1.47).]

Now, if we let $z_1, z_2 \in \{-x^*, \dots, -1, 0\}$ and assume a, b integer, then for

$$v_t = \begin{cases} 1 & \text{if } |z_1| = t - 1 \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad w_t = \begin{cases} -1 & \text{if } |z_2| = t - 1 \\ 0 & \text{otherwise} \end{cases}$$

(2.2) and (2.4) combine to give us

$$\begin{aligned} \prod_{i=1}^n \frac{\Gamma(r + x_{1i})}{\Gamma(r)} p(r) &= L_1 [r^{s_1 + v_1 a + w_1 b} \\ &\quad \dots (r + x^* - 1)^{s_{x^*} + v_{x^*} a + w_{x^*} b} (r + x^*)^{s_{x^*+1} + v_{x^*+1} a + w_{x^*+1} b}] \\ &= L_1 [r^{h_1} (r+1)^{h_2} \dots (r+x^*)^{h_{x^*+1}}], \end{aligned}$$

where $h_t = s_t + v_t a + w_t b$ and $s_{x^*+1} = 0$. (While the restriction of a, b to integers is somewhat restrictive, it still leaves a very flexible class of priors for r .) Letting $h = h_1 + h_2 + \dots + h_{x^*+1}$, we can write this product as a polynomial in r :

$$\prod_{i=1}^n \frac{\Gamma(r + x_{1i})}{\Gamma(r)} p(r) = L_1 [a_0 + a_1 r + \dots + a_h r^h] = L_1 \sum_{j=0}^h a_j r^j, \quad (2.6)$$

where the polynomial coefficients can be computed for given h_1, \dots, h_{x^*+1} using the same algorithm used to compute the U 's in (2.5). (Note that $a_0 = a_1 = \dots = a_{h_1-1} = 0$, as there are no terms less than h_1 .)

Substituting (2.5) and (2.6) into (2.3), we get

$$\begin{aligned} E(r^k | \mathbf{x}) &\approx \frac{\int r^k \left(\sum_{j=0}^h a_j r^j \right) (r + K_1)^{1-C_1} \sum_{l=0}^m \frac{U_l}{[n(r + K_1)]^l} dr}{\int \left(\sum_{j=0}^h a_j r^j \right) (r + K_1)^{1-C_1} \sum_{l=0}^m \frac{U_l}{[n(r + K_1)]^l} dr} \\ &= \frac{\sum_{j=0}^h a_j Q_{j+k}}{\sum_{j=0}^h a_j Q_j}, \end{aligned} \quad (2.7)$$

where

$$\begin{aligned} Q_j &= \int r^j (r + K_1)^{1-C_1} \sum_{l=0}^m \frac{U_l}{[n(r + K_1)]^l} dr \\ &= \sum_{l=0}^m \frac{\Gamma(C_1 - j + l - 2)}{\Gamma(C_1 + l - 1)} \frac{\Gamma(j + 1)}{K_1^{C_1 - j + l - 2}} \frac{U_l}{n^l} \end{aligned} \quad (2.8)$$

which follows naturally (from the normalizing constant) from the definition of the Pearson Type VI density in (2.2). That is, the Pearson Type VI density is a conjugate family for the

NBD after employing the approximation to $\Gamma(nr + \delta_1)/\Gamma(nr + \delta_1 + C_1 - 1)$ given in (2.5). Thus, closed-form posterior inferences are now feasible for r .

In order for the integral in (2.8) to be the normalizing constant of the Pearson Type VI, and hence defined, we must have $C_1 - j + l - 2 > 0 \forall j, l$. The maximum value that j can take on is $h + k$ ((2.6) is of order h , and we are computing $E(r^k | \mathbf{x})$; thus $h + k$ terms at most) while the minimum value that l can take on is 0. Therefore, Q_j is defined $\forall j$ when $C_1 - h - k > 2$. Now $h = \sum_{i=1}^{x^*+1} h_i = \sum_{i=1}^{x^*+1} s_i + a - b = \sum_{i=1}^{x^*+1} \sum_{j=1}^{x^*} n_j + a - b = \sum_{i=1}^n x_i + a - b$. Recall $C_1 = \sum_{i=1}^n x_i + \delta_2 + 1$. Therefore $C_1 - h - k > 2 \Leftrightarrow \delta_2 > a - b + k + 1$. If we are simply interested in the mean of the marginal posterior of r ($k = 1$), the approximation presented in (2.7) holds, provided the parameters of the marginal prior distributions for r and α satisfy the condition $\delta_2 > a - b + 2$ (which is not very restrictive).

2.2 MOMENTS OF THE MARGINAL POSTERIOR OF α

Turning our attention to α , the moments of its marginal posterior can be computed in a similar manner. Substituting (1.2) and (2.1) into (1.4) and simplifying, we get

$$E(\alpha^k | \mathbf{x}) = \frac{1}{\prod_{l=1}^k (C_1 - l - 1)} \frac{\int \left[\prod_{i=1}^n \frac{\Gamma(r + x_i)}{\Gamma(r)} \right] \frac{\Gamma(nr + \delta_1 + k)}{\Gamma(nr + \delta_1 + C_1 - 1)} p(r) dr}{\int \left[\prod_{i=1}^n \frac{\Gamma(r + x_i)}{\Gamma(r)} \right] \frac{\Gamma(nr + \delta_1)}{\Gamma(nr + \delta_1 + C_1 - 1)} p(r) dr}, \quad (2.9)$$

Applying the same logic used to derive (2.5), we have

$$\begin{aligned} \frac{\Gamma(nr + \delta_1 + k)}{\Gamma(nr + \delta_1 + C_1 - 1)} &= \prod_{j=2}^{C_1-k} \frac{1}{nr + \delta_1 + C_1 - j} \\ &\approx [n(r + K_1)]^{1+k-C_1} \sum_{l=0}^m \frac{V_l^k}{[n(r + K_1)]^l}, \end{aligned} \quad (2.10)$$

where $V_0^k = 1$, and we have a separate set of polynomial coefficients, V_1^k, \dots, V_m^k for each of the k moments of the posterior of α . These can again be computed using the previously mentioned recursive algorithm.

Substituting (2.6) and (2.10) into (2.9), we get

$$\begin{aligned} E(\alpha^k | \mathbf{x}) &\approx \frac{n^k}{\prod_{l=1}^k (C_1 - l - 1)} \frac{\int \left(\sum_{j=0}^h a_j r^j \right) (r + K_1)^{1+k-C_1} \sum_{l=0}^m \frac{V_l^k}{[n(r + K_1)]^l} dr}{\int \left(\sum_{j=0}^h a_j r^j \right) (r + K_1)^{1-C_1} \sum_{l=0}^m \frac{U_l}{[n(r + K_1)]^l} dr} \\ &= \frac{n^k}{\prod_{l=1}^k (C_1 - l - 1)} \frac{\sum_{j=0}^h a_j R_j^k}{\sum_{j=0}^h a_j Q_j}, \end{aligned} \quad (2.11)$$

where

$$\begin{aligned} R_j^k &= \int r^j (r + K_1)^{1+k-C_1} \sum_{l=0}^m \frac{V_l^k}{[n(r + K_1)]^l} dr \\ &= \sum_{l=0}^m \frac{\Gamma(C_1 + l - j - k - 2)}{\Gamma(C_1 + l - k - 1)} \frac{\Gamma(j + 1)}{K_1^{C_1+l-j-k-2}} \frac{V_l^k}{n^l}, \end{aligned} \quad (2.12)$$

which again follows from the normalizing constant of the Pearson Type VI density in (2.2). Thus, closed-form posterior inferences are obtained for α . In order for (2.12) to be defined, we must have $C_1 + l - j - k - 2 > 0 \forall j, l$, which leads to the same condition as before ($\delta_2 > a - b + k + 1$).

3. A SIMULATION STUDY

To assess the accuracy and computational feasibility of our approach, we performed a large-scale simulation study using three main factors: r (3 levels), α (3 levels), and m , the number of polynomial expansion terms (5 levels) for the ratio of two gamma functions given in (2.5) and (2.10). We selected these as simulation factors as we wanted to know: (a) whether the quality of the estimates varied with the true (latent) values of r and α , and (b) whether the number of terms used in the approximating expansion played a significant role in determining accuracy. A fully crossed design was implemented yielding 45 simulation conditions.

For each condition, a given simulation consisted of $N = 500$ observations from a negative binomial distribution with the given r and α . Two hundred replicates were generated for each of the 45 conditions; that is, 200 datasets of size 500. The results reported for each condition are the average over these replicates.

To select simulation values for r and α , we used a reparameterization of the NBD values of r and α in terms of $\rho = \text{corr}(X_i, Y_i) = 1/(\alpha + 1)$, the correlation between two negative binomial draws in nonoverlapping periods of equal length for a given individual, and $P_0 = P(X_i = 0) = [\alpha/(\alpha + 1)]^r = (1 - \rho)^r$, the fraction of 0 values. (ρ also equals the square of the correlation between the observed count (X_i) and the individual's latent rate (λ_i .) This reparameterization is convenient as both ρ and P_0 are bounded between 0 and 1, allowing us to select a grid along the unit square. In addition, ρ and P_0 are common and intuitive quantities which can be selected by a researcher based on historical figures or expert judgment.

We selected the 3×3 grid corresponding to 0.25, 0.50, 0.75 for each of ρ and P_0 leading to the values of r and α given in columns 3 and 4 of Table 1. For m , the number of polynomial

Table 1. Mean Absolute Estimation Error for $(E(r|\mathbf{x}), E(\alpha|\mathbf{x}))$ with $X_i \sim \text{NBD}(r, \alpha)$. The simulation conditions are derived from $\rho = \text{corr}(X_i, Y_i) = 1/(\alpha + 1)$, $P_0 = P(X_i = 0) = (1 - \rho)^r$, and m = number of terms used in the approximation of the ratio of two gamma functions.

P_0	ρ	r	α	$E(X_i)$	m					Mean
					1	2	10	50	300	
0.25	0.25	4.82	3.00	1.61	(27.62, 17.54)	(22.24, 14.05)	(7.05, 4.51)	(0.98, 0.62)	(0.05, 0.07)	(11.59, 7.36)
0.25	0.50	2.00	1.00	2.00	(33.73, 17.31)	(27.57, 14.18)	(11.05, 5.78)	(1.68, 0.93)	(0.10, 0.06)	(14.81, 7.66)
0.25	0.75	1.00	0.33	3.00	(33.04, 11.91)	(27.39, 9.94)	(13.20, 5.03)	(3.77, 1.57)	(0.12, 0.08)	(15.49, 5.71)
0.50	0.25	2.41	3.00	0.83	(6.04, 7.78)	(4.80, 6.24)	(1.08, 1.33)	(0.71, 0.87)	(0.00, 0.04)	(2.53, 3.26)
0.50	0.50	1.00	1.00	1.00	(6.41, 7.11)	(5.33, 5.90)	(2.15, 2.49)	(0.31, 0.31)	(0.05, 0.06)	(2.85, 3.18)
0.50	0.75	0.50	0.33	1.50	(4.51, 3.97)	(4.16, 3.56)	(2.44, 2.25)	(0.70, 0.70)	(0.01, 0.04)	(2.36, 2.10)
0.75	0.25	1.00	3.00	0.33	(0.88, 2.98)	(0.72, 2.26)	(0.27, 0.57)	(0.32, 0.99)	(0.03, 0.00)	(0.44, 1.37)
0.75	0.50	0.42	1.00	0.42	(1.05, 3.22)	(0.87, 2.73)	(0.37, 1.05)	(0.15, 0.35)	(0.01, 0.02)	(0.49, 1.48)
0.75	0.75	0.21	0.33	0.60	(0.75, 2.20)	(0.71, 2.05)	(0.42, 1.23)	(0.11, 0.27)	(0.00, 0.01)	(0.40, 1.15)
Mean					(12.67, 8.22)	(10.42, 6.77)	(4.22, 2.69)	(0.97, 0.73)	(0.05, 0.05)	

expansion terms used, we selected five values: 1, 2, 10, 50, and 300. Experimentation with our approach indicated that these values would demonstrate the increased accuracy of the estimates as m increases. Computing time using an HP-UX 9000 server, and code written in Fortran, for generating 200 datasets of size 500 and estimating the corresponding moments using our approach was 2, 3, 7, 45, and 800 seconds, respectively, for $m = 1, 2, 10, 50, 300$. This clearly suggests very practical computing time for a given single dataset, especially when compared to the numerical methods (e.g., MCMC) traditionally used in the Bayesian setting. A comparison of computing times is given below.

To complete the simulation specification, we needed to select values of the hyperparameters (a, b, z_1, z_2) , the Pearson Type VI parameters for the prior on r , and (δ_1, δ_2) , the beta-prime parameters for the prior on α . These values were then inserted into Equations (2.7) and (2.11), along with the 500 negative binomial draws, to obtain estimates $E(r|\mathbf{x})$ and $E(\alpha|\mathbf{x})$ for each simulation.

We ran each of the 45 simulation conditions under three choices of these hyperparameters, all having $z_1 = 0$ and $z_2 = -1$ indicating a beta-prime prior on r :

	prior on r	prior on α
i)	improper uniform ($a = 0, b = 0$)	weakly informative with mean 1 ($\delta_1 = 2, \delta_2 = 3$)
ii)	weakly informative with mean 1 ($a = 1, b = 5$)	weakly informative with mean 1 ($\delta_1 = 2, \delta_2 = 3$)
iii)	weakly informative with mean 1 ($a = 1, b = 5$)	improper uniform ($\delta_1 = 1, \delta_2 = -1$)

We note that the condition with two uniform priors does not satisfy the constraint $\delta_2 > a - b + 2$, described earlier, and thus could not be explored. As our results indicated no significant differences across the three choices of the hyperpriors (determined by a repeated measures ANOVA), the results reported below are the average across the three sets. As with any Bayesian analysis, though, this is due, in part, to the large sample size selected ($N = 500$) in comparison to the small amount of information in the priors.

Reported in Table 1 are the mean absolute estimation errors (MAEE) for the first moments, $\text{avg}(|E(r|\mathbf{x}) - r|)$ and $\text{avg}(|E(\alpha|\mathbf{x}) - \alpha|)$, for each of the 45 simulation conditions and aggregated across various cells. Columns 1–5 convey the 3×3 grid for ρ and P_0 , and the corresponding values of r , α and $E(X_i) = r/\alpha$, respectively. The results indicate a number of interesting features. First, as expected, the estimation error decreases with an increasing number of expansion terms m . In fact, we observe a very poor fit for a small number of terms and a good fit for the $m = 300$ condition. Second, estimation error increases with $E(X)$. This suggests that when the sample average is large, many more terms in the expansion m are necessary to provide adequate estimates. But even in the worst case (row 3), the 300-term approximation is excellent and, if necessary, additional terms could be added with very little difficulty or computational burden.

As a baseline for comparison, we assessed computing time for standard Bayesian approaches as follows. Using the free and readily available software program BUGS (Bayesian inference using Gibbs sampling) we obtained simulates from three Markov chains (for each dataset) each of length 5,000 from overdispersed starting values. Each chain was

run for 5,000 iterations based on convergence diagnostics applied to the simulates. Each dataset was of size 500 (as in our simulation design). The results reported here are the average simulation time over 100 datasets of size 500 using various values of r and α from the simulation design described above (times did not vary significantly by choice of these parameters). Two different Bayesian models were fit to the data. The first had $X_i \sim \text{Poisson}(\lambda_i)$, $\lambda_i \sim \text{gamma}(r, \alpha)$, $r \sim \text{gamma}(0.1, 1.0)$, $\alpha \sim \text{exponential}(1)$. This model was chosen identically to that given in the hierarchical Poisson example in the BUGS documentation. The second model fit had the same model for X_i and λ_i but instead assumed a multivariate log-normal prior for r and α with mean 0 and diagonal variance-covariance matrix with 100 along the diagonal (fairly disperse). All of these computations were run on a DELL 933 MHZ PC with 256MB of RAM. As the clock-speed for this machine is comparable to that of the HP-UX 9000 used for our approach, we believe a comparison of these computing times to be reasonable.

The results of these simulations indicated that the average computing time for the BUGS NBD model was 105 seconds whereas that for the multivariate log-normal NBD model was 231 seconds. Now while the comparison of our computing time, with say $m = 300$ terms, at 4 seconds to these two times is an order of magnitude faster, it may not be one of great (or any) practical importance. However, with larger datasets the computing time of our approach will provide even greater benefits. Additionally, we believe that our approach can be used as a means to check MCMC output, as well as to provide starting values.

4. MOMENTS OF THE PREDICTIVE DISTRIBUTION

The approach developed above for evaluating the moments of the marginal posterior densities can be applied to the evaluation of the moments of the predictive distribution ($E(Y_i^k|\mathbf{x})$). As prediction is among the most common applications of the NBD model in practice, closed-form moments in this instance have great practical value. Let Y_i be the count variable for individual i in a nonoverlapping time period, where the two periods corresponding to X_i, Y_i are of equal length. In this section, we focus on evaluating the mean and variance of the predictive distribution: $E(Y_i|\mathbf{x})$ and $\text{var}(Y_i|\mathbf{x})$. (The method easily generalizes to higher moments.)

Given r and α , the mean and variance of Y_i , conditional on x_i , are:

$$E(Y_i|x_i, r, \alpha) = \frac{r + x_i}{\alpha + 1} \quad \text{and} \quad \text{var}(Y_i|x_i, r, \alpha) = \frac{r + x_i}{\alpha + 1} + \frac{r + x_i}{(\alpha + 1)^2}.$$

It follows that

$$E(Y_i|\mathbf{x}) = \iint \frac{r + x_i}{\alpha + 1} p(r, \alpha|\mathbf{x}) d\alpha dr, \quad (4.1)$$

and

$$\begin{aligned} \text{var}(Y_i|\mathbf{x}) &= \iint \left(\frac{r + x_i}{\alpha + 1} + \frac{r + x_i}{(\alpha + 1)^2} \right) p(r, \alpha|\mathbf{x}) d\alpha dr \\ &= E(Y_i|\mathbf{x}) + \iint \frac{r + x_i}{(\alpha + 1)^2} p(r, \alpha|\mathbf{x}) d\alpha dr. \end{aligned} \quad (4.2)$$

Central to evaluating (4.1) and (4.2) is the evaluation of the following integral:

$$\iint \frac{r + x_i}{(\alpha + 1)^k} p(r, \alpha | \mathbf{x}) d\alpha dr = \frac{\iint \frac{r + x_i}{(\alpha + 1)^k} p(\mathbf{x} | r, \alpha) p(r) p(\alpha) d\alpha dr}{\iint p(\mathbf{x} | r, \alpha) p(r) p(\alpha) d\alpha dr}. \quad (4.3)$$

Substituting (1.2) and (2.1) into (4.3) and simplifying, we get

$$\begin{aligned} & \iint \frac{r + x_i}{(\alpha + 1)^k} p(r, \alpha | \mathbf{x}) d\alpha dr \\ &= \prod_{l=1}^k (C_1 + l - 2) \frac{\int (r + x_i) \left[\prod_{i=1}^n \frac{\Gamma(r + x_i)}{\Gamma(r)} \right] \frac{\Gamma(nr + \delta_1)}{\Gamma(nr + \delta_1 + C_1 + k - 1)} p(r) dr}{\int \left[\prod_{i=1}^n \frac{\Gamma(r + x_i)}{\Gamma(r)} \right] \frac{\Gamma(nr + \delta_1)}{\Gamma(nr + \delta_1 + C_1 - 1)} p(r) dr}. \end{aligned} \quad (4.4)$$

Using the same logic as in (2.5), we have

$$\begin{aligned} \frac{\Gamma(nr + \delta_1)}{\Gamma(nr + \delta_1 + C_1 + k - 1)} &= \prod_{j=2-k}^{C_1} \frac{1}{nr + \delta_1 + C_1 - j} \\ &\approx [n(r + K_1)]^{1-C_1-k} \sum_{l=0}^m \frac{W_l^k}{[n(r + K_1)]^l}, \end{aligned} \quad (4.5)$$

where $W_0^k = 1$, and we have a separate set of polynomial coefficients, W_1^k, \dots, W_m^k for $k = 1, 2$, computed using the previously mentioned recursive algorithm.

Substituting (2.6) and (4.5) into (4.4), we get

$$\begin{aligned} & \iint \frac{r + x_i}{(\alpha + 1)^k} p(r, \alpha | \mathbf{x}) d\alpha dr \\ &\approx \frac{\prod_{l=1}^k (C_1 + l - 2)}{n^k} \frac{\int (r + x_i) \left(\sum_{j=0}^h a_j r^j \right) (r + K_1)^{1-C_1-k} \sum_{l=0}^m \frac{W_l^k}{[n(r + K_1)]^l} dr}{\int \left(\sum_{j=0}^h a_j r^j \right) (r + K_1)^{1-C_1} \sum_{l=0}^m \frac{U_l}{[n(r + K_1)]^l} dr} \\ &= \frac{\prod_{l=1}^k (C_1 + l - 2)}{n^k} \frac{\sum_{j=0}^h a_j [S_{j+1}^k + x_i S_j^k]}{\sum_{j=0}^h a_j Q_j}, \end{aligned} \quad (4.6)$$

where

$$\begin{aligned} S_j^k &= \int r^j (r + K_1)^{1-C_1-k} \sum_{l=0}^m \frac{W_l^k}{[n(r + K_1)]^l} dr \\ &= \sum_{l=0}^m \frac{\Gamma(C_1 + k + l - j - 2)}{\Gamma(C_1 + k + l - 1)} \frac{\Gamma(j + 1)}{K_1^{C_1 + k + l - j - 2}} \frac{W_l^k}{n^l}, \end{aligned} \quad (4.7)$$

which again follows from the normalizing constant of the Pearson Type VI density in (2.2).

In order for (4.7) to be defined, we must have $C_1 + k + l - j - 2 > 0 \forall j, l$. Thus, S_j^k is defined $\forall j$ when $C_1 - h - k > 2 \Leftrightarrow \delta_2 > a - b + 1$. This is automatically satisfied when we constrain the parameters of the prior distributions such that the approximation of the moments of the marginal posterior distributions of r and α , (2.7) and (2.11), hold (i.e., $\delta_2 > a - b + 2$). Thus no new constraints need to be applied.

Substituting (4.6) into (4.1) and (4.2), we get the following expressions for the conditional mean and variance:

$$E(Y_i|\mathbf{x}) \approx \frac{C_1 - 1}{n} \frac{\sum_{j=0}^h a_j [S_{j+1}^1 + x_i S_j^1]}{\sum_{j=0}^h a_j Q_j}, \quad (4.8)$$

and

$$\text{var}(Y_i|\mathbf{x}) \approx E(Y_i|\mathbf{x}) + \frac{C_1(C_1 - 1)}{n^2} \frac{\sum_{j=0}^h a_j [S_{j+1}^2 + x_i S_j^2]}{\sum_{j=0}^h a_j Q_j}. \quad (4.9)$$

5. CONCLUSIONS

The negative binomial distribution is widely used to model count data. However, the standard inference approach estimates the model parameters using maximum likelihood, treating them fixed and known at $\hat{\tau}$ and $\hat{\alpha}$. In contrast, the Bayesian approach allows for uncertainty in parameter estimates and the incorporation of prior information.

It had previously been thought that Bayesian inference for the NBD must make use of numerical integration since expressions for the posterior densities of interest contain a model parameter within a ratio of gamma functions. However, in this article we have presented closed-form solutions to the Bayesian inference problem for the NBD. The key insight is to approximate the ratio of two gamma functions using a polynomial expansion, which then allows for the use of a conjugate prior. Given this approximation, we arrive at closed-form expressions for the moments of both the marginal posterior densities and the predictive distribution. In fact, the posterior densities themselves are a special case of the results given. We demonstrate that this approach is very accurate and that the corresponding gains in computing time are quite substantial.

The accuracy of our approach depends on the number of terms used in the approximation to the ratio of the two gamma functions. Even for a large number (i.e., 300) of terms, however, the computational time required to make the relevant inferences is minimal. When compared to standard numerical methods (e.g., MCMC), the gains in computing time are quite substantial. This is of great practical importance as it provides analysts with the flexibility of the Bayesian approach to inference for the NBD without the excessive computational burden traditionally associated with such an approach. Of course, as shown in Section 3, for small to moderate sized datasets the computational gains could be quite modest. However, the ability to generate starting values for other algorithms, and/or provide fast diagnostic information is also of value and is readily available using our approach.

The approach developed in this article need not be limited to the NBD. We hope that this work spurs on other researchers to search for closed-form solutions in other situations where the application of Bayesian techniques has forced the analyst to resort to computationally intensive methods. A natural starting point would be to examine other common distributions such as the beta-binomial and the Pareto. Some researchers have attempted to implement Bayesian approaches for these models—for example, Lee and Sabavala (1987) and Arnold and Press (1989)—but in both cases they had to resort to numerical methods of some sort.

[Received June 2000. Revised November 2000.]

REFERENCES

- Abramowitz, M., and Stegun, I. A. (eds.) (1972), *Handbook of Mathematical Functions*, New York: Dover Publications.
- Arnold, B. C., and Press, S. J. (1989), "Bayesian Estimation and Prediction for Pareto Data," *Journal of the American Statistical Association*, 84, 1079–1084.
- Burrell, Q. L. (1990), "Using the Gamma-Poisson Model to Predict Library Circulations," *Journal of the American Society for Information Science*, 41, 164–70.
- Carroll, V. P., Lee, H. L., and Rao, A. G. (1986), "Implications of Salesforce Productivity Heterogeneity and Demotivation: A Navy Recruiter Case Study," *Management Science*, 32, 1371–1388.
- Deely, J. J., and Smith, A. F. M. (1998), "Quantitative Refinements for Comparisons of Institutional Performance," *Journal of the Royal Statistical Society, Series A*, 161, 5–12.
- Ehrenberg, A. S. C. (1959), "The Pattern of Consumer Purchases," *Applied Statistics*, 8, 26–41.
- Greenwood, M., and Udney Yule, G. (1920), "An Inquiry into the Nature of Frequency Distributions Representative of Multiple Happenings with Particular Reference to the Occurrence of Multiple Attacks of Disease or of Repeat Accidents," *Journal of the Royal Statistical Society*, 83, 255–279.
- Johnson, N. L., Kotz, S., and Balakrishnan, N. (1995), *Continuous Univariate Distributions*, (Vol. 2, 2nd ed.), New York: Wiley.
- Johnson, N. L., Kotz, S., and Kemp, A. (1992), *Univariate Discrete Distributions* (2nd ed.), New York: Wiley.
- Klugman, S. A. (1992), *Bayesian Statistics in Actuarial Science*, Boston: Kluwer Academic Publishers.
- Lee, J. C., and Sabavala, D. J. (1987), "Bayesian Estimation and Prediction for the Beta-Binomial Model," *Journal of Business and Economic Statistics*, 5, 357–67.
- Schlüter, P. J., Deeley, J. J., and Nicholson, A. J. (1997), "Ranking and Selecting Motor Vehicle Accident Sites by Using a Hierarchical Bayesian Model," *The Statistician*, 46, 293–316.