Markov counting models for correlated binary responses

FORREST W. CRAWFORD*, DANIEL ZELTERMAN

Yale School of Public Health, Biostatistics, PO Box 208034, New Haven, CT 06510, USA

forrest.crawford@yale.edu

SUMMARY

We propose a class of continuous-time Markov counting processes for analyzing correlated binary data and establish a correspondence between these models and sums of exchangeable Bernoulli random variables. Our approach generalizes many previous models for correlated outcomes, admits easily interpretable parameterizations, allows different cluster sizes, and incorporates ascertainment bias in a natural way. We demonstrate several new models for dependent outcomes and provide algorithms for computing maximum likelihood estimates. We show how to incorporate cluster-specific covariates in a regression setting and demonstrate improved fits to well-known datasets from familial disease epidemiology and developmental toxicology.

Keywords: Markov process; Bernoulli trials; Developmental toxicity; Familial disease; Teratology.

1. INTRODUCTION

The simplest statistical model for a collection of n binary outcomes is the binomial distribution, which assumes that responses are independent and identically distributed. However, many investigations have found that the binomial distribution sometimes gives a poor fit to certain types of data (Greenwood and Yule, 1920; Haseman and Soares, 1976; Altham, 1978). This empirical observation, along with suspicions that the mechanism generating the outcomes might induce dependencies, has encouraged development of more flexible models that account for correlations in responses. Dependent or correlated binary data arise commonly in studies of developmental toxicology and litter size (Williams, 1975; Kupper and Haseman, 1978; Altham, 1978), familial disease aggregation (Liang and others, 1992; Yu and Zelterman, 2002), or when ascertainment considerations necessitate a biased approach to sampling (Matthews and others, 2008). Groups of dependent responses are often called “clusters”, and in many applications the response of interest is the number of affected units in a cluster with n members.

When individual unit-level data are available, mixed-effects logistic regression approaches (e.g., Stiratelli and others, 1984) can model correlation using cluster-specific effects; marginal models posit a population-averaged mean and a working covariance structure (Zeger and Liang, 1986). These approaches depend on the access to individual-level outcomes, which are not always available. Mixed-effects and marginal models allow specification of pairwise covariances, but may be unable to provide higher-order dependency between outcomes. This has led researchers to study models for the sum of dependent

*To whom correspondence should be addressed.

Bernoulli variables. One of the simplest is the beta-binomial model, used to account for extra-binomial variation in clustered counts (Moore and others, 2001; Yu and Zelterman, 2002). George and Bowman (1995) and Bowman and George (1995) present general expressions for the likelihood of a sum of exchangeable Bernoulli variables via a combinatorial argument. In this context, exchangeability means that the joint probability of all the outcomes in a cluster is invariant to permutation of the responses, a notion we define more formally in Section 2. Kuk (2004) uses the George and Bowman (1995) framework to define families of power functions that show superior fit in developmental toxicity studies and Pang and Kuk (2005) give a model that allows a random subset of responses to share their response. Yu and Zelterman (2002, 2008) derive the beta-binomial distribution and other models under the George and Bowman (1995) framework. Several authors describe methods to fit data consisting of observations on clusters of different sizes: Stefanescu and Turnbull (2003) interpret different cluster sizes in a missing data framework and derive EM algorithms for fitting. Xu and Prorok (2003) and Pang and Kuk (2007) deal with this issue by assuming that the marginal distributions of the first k responses in different cluster sizes are equal.

In this work, we take a very different approach: we show that sums of exchangeable Bernoulli random variables can be represented as continuous-time Markov counting processes via a technique called probabilistic embedding (Blom and Holst, 1991). By introducing an auxiliary variable, the binary responses are made to depend on the arrival times of points in a Markov counting process. This formulation provides a flexible way to parameterize and fit models of correlated binary outcomes, and accommodates different cluster sizes and ascertainment schemes. We review basic results for exchangeable Bernoulli variables and give examples of models derived under this framework. We then describe a class of Markov counting process and give five examples inspired by principles from infectious disease epidemiology. Next, we show that any Markov counting process can be expressed as a sum of exchangeable Bernoulli variables. We apply our approach to three datasets in which outcomes cluster in families and one developmental toxicology experiment. Appendices of Supplementary material available at *Biostatistics* online provide simulation results, algorithms for maximum likelihood estimation, regression with covariates, and numerical evaluation of likelihoods.

2. SUMS OF EXCHANGEABLE BERNOULLI VARIABLES

George and Bowman (1995) and Bowman and George (1995) describe a likelihood framework for sums of exchangeable Bernoulli random variables that depends on knowledge of joint probabilities of subsets of variables taking value 1. Consider a sequence of n exchangeable Bernoulli variables $Z_1, \dots, Z_n$. By exchangeability, we mean that the joint probability of a collection of variables taking certain values is invariant to reordering. More formally, $\Pr(Z_1 = z_1, \dots, Z_n = z_n) = \Pr(Z_{\pi(1)} = z_{\pi(1)}, \dots, Z_{\pi(n)} = z_{\pi(n)})$ for any permutation π of the indices $1, 2, \dots, n$ (De Finetti, 1931). Now consider the probability that r of the Z_i 's take value 1 and $n - r$ take value 0. By exchangeability, we can express this as the joint probability that the first r take value 1 and the remainder are 0. Let $\lambda_j = \Pr(Z_{i_1} = Z_{i_2} = \dots = Z_{i_j} = 1)$ be the joint probability that every Z_i for $i \in I_j$ is 1, where the cardinality of the set I_j is j . Now letting $Y_n = \sum_{i=1}^n Z_i$, application of the inclusion-exclusion formula gives

$$\Pr(Y_n = r) = \binom{n}{r} \sum_{j=0}^{n-r} (-1)^j \binom{n-r}{j} \lambda_{r+j}. \quad (2.1)$$

A derivation of (2.1) is given by George and Bowman (1995, p. 513). By specifying the joint probabilities λ_j for $j = 0, \dots, n$, the distribution of any sum of exchangeable Bernoulli variables can be represented. In particular, setting $\lambda_{r+j} = p^{r+j}$ recovers the binomial distribution. The λ_j 's are sometimes called "marginal" probabilities (Dang and others, 2009), since they express the joint probability of j successes,

summed over all possible outcomes of the remaining $n - j$ variables. This model is called “saturated” when all the λ_j ’s are allowed to be non-zero.

We note three major issues with the model of [George and Bowman \(1995\)](#) given by (2.1). First, it is unclear how to interpret the joint probabilities λ_j or correlations when analyzing data from clusters of different sizes since the number of unknown parameters for each observation is equal to the cluster size. [Xu and Prorok \(2003\)](#) and [Pang and Kuk \(2007\)](#) deal with this problem by assuming that the marginal probability of r responses having value 1 in a family of size $n \geq r$ is equal to the probability of r responses having value 1 in a family of size $n' > n$, but this assumes that response probabilities do not depend on cluster size. Second, it can be difficult to specify joint probabilities λ_j for $j = 0, \dots, n$ that result in a well-defined probability mass function ([George and Bowman, 1995](#); [Stefanescu and Turnbull, 2003](#)). Often one must solve a non-trivial combinatorial problem in order to specify the λ_j ’s (see, e.g., [Kuk, 2004](#); [Pang and Kuk, 2005](#)). Third, sampling or ascertainment of clusters can sometimes depend on the responses; for example, often families in epidemiological studies are selected via a single-affected member. The likelihood of observing r affected individuals in a family of size n must then be computed conditional on having at least one response having value 1, which may be a function of family size n . The interaction of ascertainment conditions and varying cluster sizes can substantially complicate inference for dependent counts.

2.1 Examples of models for λ_j

2.1.1 Binomial. When the Bernoulli variables are independent with probability p of success, $\lambda_j = p^j$ and (2.1) reduces to the binomial probability $\Pr(Y_n = r) = \binom{n}{r} p^r (1 - p)^{n-r}$.

2.1.2 Beta-Binomial. [Yu and Zelterman \(2008\)](#) show that setting $\lambda_1 = p$, $\lambda_2 = p(p + \alpha)/(1 + \alpha)$, and $\lambda_k = p(p + \alpha) \cdots (p + (k - 1)\alpha)/(1 + \alpha) \cdots (1 + (k - 1)\alpha)$ for $3 \leq k \leq n$ gives the beta-binomial distribution

$$\Pr(Y_n = r) = \binom{n}{r} \prod_{k=0}^{r-1} (k\alpha + p) \prod_{s=0}^{n-r-1} (s\alpha + 1 - p) \bigg/ \prod_{j=0}^{n-1} (j\alpha + 1),$$

when $\alpha > -\min\{p, 1 - p\}/(n - 1)$. Here, p is the marginal success probability, and α is a measure of correlation. Setting $\alpha = 0$ recovers the binomial distribution.

2.1.3 q -Power. Consider the family of distributions in which $\lambda_j = p^{j^\gamma}$, where p is called the marginal response probability. When $0 \leq \gamma \leq 1$, the probability distribution (2.1) is well defined. [Kuk \(2004\)](#) proposes to set $q = 1 - p$ and model the number of zero outcomes, $n - Y_n$. Then (2.1) becomes $\Pr(n - Y_n = r) = \binom{n}{r} \sum_{j=1}^r (-1)^j \binom{r}{k} q^{(n-r+j)^\gamma}$. Here, γ is a measure of positive intra-cluster correlation: setting $\gamma = 1$ results in no correlation between responses.

3. MARKOV COUNTING PROCESSES

There is an important correspondence between the [George and Bowman \(1995\)](#) representation (2.1) and continuous-time Markov counting models. To make this clear, we formally define this class of processes and show how to calculate their transition probabilities. In the next section, we construct an equivalence between Markov counting processes and sums of exchangeable Bernoulli random variables. Consider a continuous-time Markov process $X(t)$ that counts the number of arrivals (or points) before time t . When

k points have arrived, the rate of arrival of the next point is μ_k . Let $P_{mr}(t) = \Pr(X(t) = r \mid X(0) = m)$ be the probability that at time t there have been r arrivals, given that there were m already at time 0. This probability obeys the forward equation

$$\frac{dP_{mr}(t)}{dt} = \mu_{r-1}P_{m,r-1}(t) - \mu_r P_{mr}(t), \quad (3.1)$$

where $\mu_r > 0$ is the instantaneous rate of the $r + 1$ st arrival, given that r have already arrived (Karlin and Taylor, 1975, p. 119). This counting model is also known as the “generalized Yule” or “pure birth” process. The homogeneous Poisson process with $\mu_r = \mu$ is the best-known counting process, with transition probability $P_{mr}(t) = (\mu t)^{r-m} e^{-\mu t} / (r - m)!$. For a general Markov counting process with rates μ_k , $k = 0, 1, \dots$, the transition probability is

$$P_{mr}(t) = \left(\prod_{k=m}^{r-1} \mu_k \right) \sum_{k=m}^r \left(\prod_{\ell \neq k} (\mu_\ell - \mu_k) \right)^{-1} \exp[-\mu_k t], \quad (3.2)$$

for $0 \leq m \leq r$ and $t > 0$ when $\mu_k \neq \mu_r$ for all m and r (Renshaw, 2011, p. 65). For a given set of rates $\{\mu_k\}$, simpler representations of the likelihood (3.2) are often available, as we show in Section 3.1. When $\mu_k = \mu_\ell$ for some k and ℓ , it can be more difficult to derive likelihood expressions. Fortunately, computational evaluation of the likelihood is straightforward and robust via numerical methods. We give a general method for numerically evaluating $P_{mr}(t)$ in Appendix of Supplementary material available at *Biostatistics* online.

3.1 Examples of models for μ_j

It can be challenging to translate informal ideas about dependency into parametric models for dependent count data in the framework of George and Bowman (1995). However, counting process rates are often easy to specify; usually a consideration of the conditional risk of a new event, given the number that have already occurred, is enough to express the μ_k ’s in a useful form. Modelers do not need to accommodate awkward constraints on the rates, such as monotonicity, that might make them difficult to specify jointly or interpret (see Stefanescu and Turnbull, 2003, for example). Here we present five simple counting processes derived from basic principles of infectious disease epidemiology. We imagine a household of size n with k members already affected by the disease. Transmissibility of disease status induces dependency in the outcomes of individual family members; households are “clusters” and individuals are “units”. We distinguish between two sources of risk to members of a cluster of size n : exogenous or extra-cluster risk to which all unaffected units are subject, and infectivity, or risk experienced by each susceptible member in proportion to the number already affected. Table 1 shows a summary of the counting processes we consider in what follows.

3.1.1 Susceptible. Consider a cluster of size n in which each unaffected (susceptible) unit experiences the same exogenous risk $\alpha > 0$. When there are k affected units, the number of unaffected units is $n - k$ and the risk to the cluster is $\mu_k = \alpha(n - k)$. This formulation produces a counting process with a familiar epidemiological interpretation corresponding to constant per-unaffected-unit risk and no infectivity between units. In fact, this model is formally equivalent to the binomial model with success probability $1 - e^{-\alpha}$. The likelihood for this “susceptible-1” model is $P_{mr} = \binom{n-m}{r-m} e^{-\alpha(n-r)} (1 - e^{-\alpha})^{r-m}$. We report this fact here to show that the susceptible counting process model, which has a traditional epidemiological interpretation, corresponds exactly to the simplest model for n binary outcomes. One straightforward extension of the susceptible-1 model is to allow the cluster risk to be a non-negative power function of the number

Table 1. *Illustration of the proposed counting process models. Model name and arrival rate μ_k are given in the first two columns. A stochastic realization of the counting process is shown, where a vertical line represents the time of each arrival and the gray step function represents the rate μ_k . A schematic diagram of a household is given for each type of model. Filled gray circles represent affected family members and white circles represent unaffected members; in each diagram, there are $n = 6$ family members with three affected and three unaffected. Exogenous or extra-household risk per unaffected member is α , and the risk per potential contact between affected and unaffected members is β*

Model	Rate μ_k	Counting process example	Risk schematic
Susceptible-1	$\alpha(n - k)$
Susceptible-2	$\alpha(n - k)^\gamma$
Infectivity-1	$\beta k(n - k)$
Infectivity-2	$\beta k^\eta(n - k)^\gamma$
Combined	$\alpha(n - k) + \beta k(n - k)$

of susceptibles, $\mu_k = \alpha(n - k)^\gamma$, where $\alpha > 0$ and $\gamma > 0$. If $0 < \gamma < 1$, the cluster experiences risk smaller than that obtained by the susceptible-1 model, and if $\gamma > 1$, the cluster experiences greater risk.

3.1.2 Infectivity. In contrast to the susceptible models, the infectivity-1 model considers only risk due to affected cluster members. Each potential contact between susceptible and affected units presents an opportunity for a new case. When there are k affected units, the number of ways one affected and one susceptible unit can come into contact is $k(n - k)$, so $\mu_k = \beta k(n - k)$ where $\beta > 0$ is the per-contact infectivity. This model formalizes the epidemiological notion of infectivity or contagion in a closed community (Britton, 1997). Since $\mu_0 = 0$, this model is most useful when ascertainment is of clusters with at least one affected member. The infectivity-2 model extends the infectivity-1 model to allow the cluster risk to vary as a power of the number of affected and susceptible members, $\mu_k = \beta k^\eta(n - k)^\gamma$, where β , η , and γ are non-negative.

3.1.3 Combined. Now we combine the susceptible-1 model with the infectivity-1 model. The per-susceptible risk from extra-cluster sources is α , and the risk contributed by one affected member to each susceptible is β . These assumptions entail the cluster risk $\mu_k = \alpha(n - k) + \beta k(n - k)$. The susceptible-1 model results from $\beta = 0$, and infectivity-1 model is obtained by setting $\alpha = 0$. Testing whether the

outcome (positive disease status) clusters in families is equivalent to asking whether β is non-zero. Finding $\beta > 0$ might indicate a genetic or household component to disease risk. The parameterization separates the effect of per-susceptible risk (α) from within-cluster infectivity (β). In regression analyses, it is possible to assess how much of the infectivity is due to cluster-level covariates, as we show below in Section 4.4.

3.1.4 Regression and relative risk for the combined model. Suppose we observe N clusters, where n_i is the number of units in cluster i and r_i is the number of affected units in cluster i . In the i th cluster, we model the counting process rate as $\mu_k = \alpha_i(n_i - k) + \beta_i k(n_i - k)$ for $k = 0, \dots, n_i$. Let d_i be a covariate for the i th cluster and let $\phi = (\phi_0, \phi_1)$ and $\psi = (\psi_0, \psi_1)$ be covariates. In toxicology experiments, d_i might correspond to the dose of toxin received by units in cluster i . We use a log-linear parameterization for the counting process rates, $\log \alpha_i = \phi_0 + \phi_1 d_i$ and $\log \beta_i = \psi_0 + \psi_1 d_i$. We employ a gradient ascent EM algorithm derived in Supplementary material available at *Biostatistics* online to estimate the parameters and standard errors in regression models.

The combined regression model offers an appealing benefit related to the interpretation of risk. Suppose we estimate α and β as in Section 4.4 under different levels of a dose/exposure d for clustered units. Then a natural comparison of dose-dependent risk that controls for infectivity of the outcome is the ratio of the per-susceptible risks α , $RR = e^{\phi_0 + \phi_1 d} / e^{\phi_0} = e^{\phi_1 d}$. This is an analog of the *relative risk* often reported in epidemiological studies under the binomial or Poisson models (McNutt and others, 2003; Zou, 2004). The difference is that RR controls for risk attributable to the interaction of already affected units with susceptible units—infectivity. We apply this regression approach in Section 4.4.

3.2 The connection

Now we show how to construct a sequence of exchangeable dependent Bernoulli variables from a Markov counting process. The Bernoulli trials are “embedded” in the counting process in the following way using probabilistic arguments introduced by Blom and Holst (1991) and Blom and others (1994, p. 186). To each Bernoulli variable Z_i we associate a latent value T_i . If $\sum_{j=0}^i T_j < t$, where $t > 0$ has been chosen in advance, then $Z_i = 1$ and otherwise 0. The T_i ’s are shown to be equivalent to exponential waiting times in a Markov counting process. The relationship between the counting process rates μ_k and the joint probabilities λ_j in the model of George and Bowman (1995) is derived.

Consider a set of n units and fix $t > 0$ and $\mu_k > 0$ for $k = 0, \dots, n-1$ with $\mu_n = 0$. Label the binary response of the i th unit Z_i . We construct the responses in n steps. Let $S_0 = \{1, \dots, n\}$ represent the indices of the n units initially at risk.

Step 1: For each $i \in S_0$, let $W_i \sim \text{Exponential}(\mu_0/n)$ independently and $T_0 = \min\{W_i; i \in S_0\}$. Let i_1^* be the index that achieves this minimum. Let $Z_{i_1^*} = \mathbb{1}\{T_0 < t\}$ and $S_1 = S_0 \setminus \{i_1^*\}$.

Step k : For each $i \in S_{k-1}$, let $W_i \sim \text{Exponential}(\mu_{k-1}/(n-k+1))$ independently and $T_{k-1} = \min\{W_i; i \in S_{k-1}\}$. Let i_k^* be the index that achieves this minimum. Let $Z_{i_k^*} = \mathbb{1}\left\{\sum_{j=0}^{k-1} T_j < t\right\}$, and $S_k = S_{k-1} \setminus \{i_k^*\}$.

Step n : Now S_{n-1} has only one element. Let $T_{n-1} \sim \text{Exponential}(\mu_{n-1})$ and let i_n^* be the remaining unit. Let $Z_{i_n^*} = \mathbb{1}\left\{\sum_{j=0}^{n-1} T_j < t\right\}$.

This procedure produces a set of n exchangeable Bernoulli variables $Z_1, \dots, Z_n$ whose joint probability is given by the transition probability of a counting process. To see why this is so, recall that since the W_i ’s at each step are independent, their minimum has exponential distribution with rate equal to the sum of the rates of the W_i ’s. At step k we have $T_{k-1} = \min\{W_i; i \in S_{k-1}\}$. Since the W_i ’s are independent, it follows that $T_{k-1} \sim \text{Exponential}(\sum_{j=0}^{n-k} (\mu_{k-1}/(n-k+1))) = \text{Exponential}(\mu_{k-1})$.

Now consider a Markov counting process $X(t)$ starting at $X(0) = 0$. We can interpret T_{k-1} as the dwell time of the counting process in state $k-1$ before jumping to k , so $\sum_{j=0}^{k-1} T_j$ is the time at which the process

jumps to state k . Then the probability of r successes is

$$\begin{aligned}\Pr(Y_n = r) &= \Pr(Z_{i_1^*} = \cdots = Z_{i_r^*} = 1, Z_{i_{r+1}^*} = \cdots = Z_{i_n^*} = 0) \\ &= \Pr\left(\sum_{j=0}^{r-1} T_j < t, \sum_{j=0}^r T_j > t\right) \\ &= P_{0r}(t),\end{aligned}$$

by construction. In the second line of (3.2), we have replaced the Bernoulli variables Z_j by their corresponding latent variables T_j . In the third line, we have replaced the statements about the sum of waiting times with equivalent statements about the value of the corresponding Markov process $X(t)$ at time t .

To show that the Z_i 's thus defined are exchangeable, it suffices to demonstrate that the index i_k^* at each step is chosen uniformly at random from the elements of S_k . We appeal to the notion of competing risks: the waiting time $\min\{W_i; i \in S_k\}$ is independent of the particular index i_k^* that achieves this minimum (Lange, 2010, p. 188). Therefore, the probability of choosing any particular i_k^* is given by $\Pr(i_k^*) = \mu_k / (n - k) / \sum_{j=1}^{n-k} (\mu_k / (n - k)) = 1 / (n - k)$. Then the probability of any particular sequence is $\Pr(i_1^*, \dots, i_n^*) = 1 / n!$ and so the i_k^* 's constitute a random permutation of the integers $1, \dots, n$. It follows that the count $X(t)$ corresponds to a sum of exchangeable Bernoulli variables. We emphasize that the times T_k in the counting process representation are auxiliary variables whose purpose is to aid in construction of the equivalence. It is not necessary to consider T_k to be the waiting time until infection of the $(k + 1)$ th individual in a familial disease model. By exchangeability, the order in which the subjects attained their response is irrelevant. Likewise, the time t is meaningless since scaling t by a constant c and dividing each μ_k by c does not alter the transition probability. We henceforth set $t = 1$ and write the counting process probability as $P_{0r} = P_{0r}(1)$.

3.2.1 The relationship between μ_k and λ_k in the George and Bowman (1995) model. The joint success probabilities λ_k in the model of George and Bowman (1995) can be derived recursively from the counting process transition probabilities, which are functions of the arrival rates μ_k . First, note that the probability of n successes in n exchangeable Bernoulli trials is given by $\Pr(Y_n = n) = \lambda_n = P_{0n}$ in the counting process model. Likewise, the probability of $n - 1$ successes is given by $\Pr(Y_n = n - 1) = n[\lambda_{n-1} - \lambda_n] = P_{0,n-1}$. Rearranging, we find that $\lambda_{n-1} = (1/n)P_{0,n-1} + \lambda_n$, and so on until we reach $\lambda_0 = P_{00} - \sum_{j=1}^n (-1)^j \binom{n}{j} \lambda_j$, recovering each joint probability λ_k from the collection of arrival rates in the counting process representation. Unlike the formulation of George and Bowman (1995), in which the relationships between the λ_k 's is complicated, there are no conditions on the rates μ_k in the Markov process, other than positivity: when all $\mu_k > 0$ for $k = 0, \dots, n - 1$ and $\mu_n = 0$, P_{mr} is always a valid probability distribution on $r \in \{m, \dots, n\}$.

3.3 Ascertainment and different cluster sizes

The counting process framework can accommodate data in which clusters are only observed if they meet some condition on the outcome of interest. For example, in some observational epidemiological studies, only families with one or more affected children are available for study. When observation is conditional on the outcome of interest, *ascertainment bias* may result. If only families with m affected members can be studied, the probability of r affected members in a family of size n must be evaluated conditional on having at least m affected members, P_{mr} . In the same way, we can account for clusters of different sizes. Let n_i be the size of the i th cluster and let r_i be the number of units affected. By specifying the dependence of μ_k on n_i for $k = 0, \dots, n_i - 1$ and letting $\mu_{n_i} = 0$, the relevant likelihood is P_{0r_i} , evaluated using rates

Table 2. *Results for the IPF dataset*

Model		Estimate	SE	L	AIC	BIC	χ^2
Binomial	p	0.296	0.032	−93.0	188.1	191.4	312.3
Beta-Binomial	p	0.238	0.031	−101.6	207.3	213.9	220.7
	a	0.086	0.057				
Altham	p	0.334	0.037	−91.3	186.5	193.1	49.4
	θ	0.793	0.093				
q -Power	q	0.720	0.036	−87.9	179.9	186.5	12.0
	γ	0.835	0.087				
Shared	p	0.282	0.036	−89.0	182.0	188.6	21.8
	π	0.439	0.098				
FH	p	0.177	0.032	−24.0	52.0	58.6	52.1
	p'	0.549	0.111				
IR	a	−1.533	0.215	−22.1	48.1	54.8	32.2
	b	1.222	0.414				
Susceptible-1	α	0.350	0.045	−93.0	188.1	191.4	312.3
Susceptible-2	α	0.308	0.071	−92.8	189.6	196.2	258.1
	γ	1.163	0.233				
Combined	α	0.275	0.044	−87.4	178.8	185.4	9.6
	β	0.300	0.124				

μ_k that depend on n_i . This is an improvement over previous models, which have generally required that either all clusters be of the same size or that one assume marginal compatibility (Pang and Kuk, 2007).

4. APPLICATIONS

Supplementary material available at *Biostatistics* online shows validation results obtained by fitting the proposed models to simulated data. In this section, we analyze four datasets that appear to exhibit clustering of responses and compare our results to those obtained using other models, with emphasis on interpretation of estimated parameters. In each case, we compare our results to previous studies using several goodness-of-fit summaries: maximum log-likelihood value (L), Akaike information criterion (AIC), Bayesian information criterion (BIC), and χ^2 statistic. In addition to the standard binomial model, we analyze each dataset using several other models that have shown good performance in previous research on dependent count outcomes: the beta-binomial model (Moore and others, 2001), which models overdispersion with respect to binomial outcomes; the Altham (1978) model for positive and negative association between outcomes; the q -power model, introduced by Kuk (2004); the shared response model of Pang and Kuk (2005) in which a random subset of responses in each cluster are shared; the family history (FH) model of Yu and Zelterman (2002) in which the first positive outcome happens with a different probability than subsequent outcomes; and the incremental risk (IR) model of Yu and Zelterman (2002). However, we caution against direct comparison of summaries based on the maximum likelihood value—the fitted models are quite different and the AIC and BIC may not be suitable for comparison between non-nested models (Dang and others, 2009).

4.1 IPF in families with COPD

Liang and others (1992) present observed frequencies of 60 cases of interstitial pulmonary fibrosis (IPF) in the siblings of families with at least one case of chronic obstructive pulmonary disease (COPD). Table 2

presents results. The FH and IR models of [Yu and Zelterman \(2002\)](#) show good performance in the likelihood-based measures (L , AIC, and BIC). The q -power and combined models are superior in their χ^2 statistics, with the combined model achieving the lowest value. The binomial, beta-binomial, Altham, q -power, and shared response models all indicate that the marginal probability of IPF in a single sibling is ~ 0.3 (the first estimated parameter in the q -power model is the marginal probability of “failure”—no IPF). Each of these indicates positive correlation of IPF cases within families. Under the FH model, the first affected sibling occurs with low probability, and subsequent siblings are affected with much greater probability. In the IR model, the risk to unaffected siblings increases monotonically with the number of affected siblings; while baseline risk of IPF is low, each affected sibling substantially increases risk to unaffected siblings. The Susceptible-1 and -2 models show moderate-positive association of IPF cases. The combined model separates the marginal per-unaffected risk α from the per-contact infectivity β , indicating substantial contributions of risk from each.

4.2 Childhood cancer syndrome

[Li and others \(1988\)](#) report the incidence of cancer in siblings of childhood cancer victims with Li-Fraumeni syndrome from a review of the Cancer Family Registry. [Yu and Zelterman \(2002\)](#) present a summary of the data consisting of counts of siblings of children with cancer. In our analysis, we account for ascertainment of families via a one affected child by the conditioning argument outlined in Section 3.3. Therefore, the dataset we analyze here is the same as that presented in [Yu and Zelterman \(2002\)](#), but adjusted to include the affected children. Table 3 shows the results. The IR, susceptible, and combined models achieve the best likelihood-based scores, with the Infective-2 and Susceptible-2 models having the lowest χ^2 value. The first models in Table 3 indicate that the marginal probability of childhood cancer in already-affected families is large, between 0.4 and 0.5. There may be correlation in the outcomes of individuals in these families, but the considered models disagree about its sign. The beta-binomial, q -power, and IR models indicate negative correlation, but the Altham model (and the shared response model, by design) indicates positive association. The infective, susceptible, and combined models offer an alternative explanation: each affected sibling increases the risk to others, but this increase diminishes as more siblings are affected. Notably, there is little evidence from these models of increased per-contact risk due to infectivity. We do not fit the FH model of [Yu and Zelterman \(2002\)](#) to the cancer dataset since only families with one affected child were ascertained.

4.3 Childhood mortality in Brazil

[Yu and Zelterman \(2002\)](#) summarize data first reported by [Sastry \(1997\)](#) on deaths of children in families of various sizes in a study of childhood mortality in impoverished areas of Brazil. [Yu and Zelterman \(2002\)](#) note that family size appears to correlate with mortality and show that the FH and IR models fit the data well. Table 4 shows the results, with the FH and IR models showing the best likelihood-based measures, and the combined model clearly outperforming the others in its χ^2 statistic. The marginal probability of death of a single child is estimated to be slightly > 0.1 in this population, and the correlation of responses is estimated by most models to be positive, with the exception of the Altham model, where $\theta > 1$; the large standard error and χ^2 value here suggest that the Altham model fits these data poorly. The susceptible models offer little insight, but the combined model tells a fuller story: baseline risk to a given child is low, but the risk to the family depends both on the number of children who have died, and the number remaining. This suggests that the childhood mortality may have a “contagious” component within families in this community.

Table 3. *Results for the childhood cancer data*

Model		Estimate	SE	L	AIC	BIC	χ^2
Binomial	p	0.487	0.047	-34.5	71.1	73.8	39.5
Beta-Binomial	p	0.436	0.043	-40.5	85.0	90.5	99.6
	a	-0.043	0.046				
Altham	p	0.488	0.045	-34.5	73.0	78.5	38.5
	θ	0.970	0.105				
q -Power	q	0.493	0.058	-33.9	71.9	77.3	35.6
	γ	0.911	0.088				
Shared	p	0.494	0.059	-34.5	73.1	78.5	39.0
	π	0.135	0.325				
IR	a	1.403	0.624	-27.5	59.0	64.5	37.6
	b	-1.132	0.390				
Infective-1	β	0.275	0.051	-35.1	72.2	74.9	45.3
Infective-2	β	0.739	0.222	-27.8	61.5	69.7	22.9
	η	< 0.001	0.536				
	γ	0.434	0.246				
Susceptible-1	α	0.428	0.078	-29.5	60.9	63.7	29.6
Susceptible-2	α	0.904	0.321	-27.2	58.4	63.8	21.5
	γ	0.433	0.257				
Combined	α	0.384	0.219	-29.7	63.3	68.8	31.0
	β	< 0.001	0.139				

Table 4. *Results for the Brazilian childhood mortality data*

Model		Estimate	SE	L	AIC	BIC	χ^2
Binomial	p	0.146	0.007	-791.9	1585.7	1591.7	2300.4
Beta-Binomial	p	0.134	0.007	-773.1	1550.1	1562.1	135.5
	a	0.115	0.023				
Altham	p	0.123	0.010	-788.0	1579.9	1591.9	8788.0
	θ	1.105	0.040				
q -Power	q	0.859	0.007	-774.3	1552.7	1564.7	135.5
	γ	0.915	0.023				
Shared	p	0.137	0.007	-766.6	1537.2	1549.2	124.6
	π	0.323	0.031				
FH	p	0.111	0.007	-459.2	922.5	934.5	338.9
	p'	0.300	0.024				
IR	a	-2.043	0.064	-458.8	921.6	933.6	271.2
	b	0.813	0.101				
Susceptible-1	α	0.158	0.008	-791.9	1585.7	1591.7	2300.5
Susceptible-2	α	0.066	0.010	-764.9	1533.8	1545.7	3847.2
	γ	1.716	0.104				
Combined	α	0.123	0.007	-750.3	1504.6	1516.6	67.4
	β	0.159	0.023				

4.4 Developmental toxicity of an herbicide

Researchers exposed pregnant mice to different doses of the herbicide 2,4,5-trichlorophenoxyacetic acid (2,4,5-T) during gestation and recorded the number of implanted fetuses and the number of fetuses

Table 5. Combined model regression estimates and standard errors for the developmental toxicity data in Table 1 of [George and Bowman \(1995\)](#). The overall results for the parameters ϕ_0 , ϕ_1 , ψ_0 , and ψ_1 are given in the first two lines. Below, exogenous risk (α) and infectivity (β) parameters are given for each dose level, where $\alpha = \exp[z'_i \phi]$, $\beta = \exp[z'_i \psi]$, and $z_i = (1, \text{dose})$. Standard errors of α and β for the different dose levels were obtained by the delta method

Dose (mg/kg)	Extra-cluster risk			Infectivity		
	Parameter	Estimate	SE	Parameter	Estimate	SE
All	ϕ_0	-2.760	0.122	ψ_0	-3.453	0.177
	ϕ_1	0.016	0.003	ψ_1	0.042	0.003
0	$\alpha = e^{\phi_0}$	0.063	0.122	$\beta = e^{\psi_0}$	0.032	0.177
30	$\alpha = e^{\phi_0+30\phi_1}$	0.103	0.144	$\beta = e^{\psi_0+30\psi_1}$	0.113	0.203
45	$\alpha = e^{\phi_0+45\phi_1}$	0.132	0.168	$\beta = e^{\psi_0+45\psi_1}$	0.214	0.231
60	$\alpha = e^{\phi_0+60\phi_1}$	0.168	0.196	$\beta = e^{\psi_0+60\psi_1}$	0.404	0.265
75	$\alpha = e^{\phi_0+75\phi_1}$	0.214	0.227	$\beta = e^{\psi_0+75\psi_1}$	0.764	0.304
90	$\alpha = e^{\phi_0+90\phi_1}$	0.273	0.260	$\beta = e^{\psi_0+90\psi_1}$	1.444	0.345

Table 6. Relative risk (RR) estimates and standard errors for the combined model in the developmental toxicology data. Standard errors were obtained by the delta method

Dose	Expression	RR	SE
0	e^0	1	
30	$e^{30\phi_1}$	1.628	0.0178
45	$e^{45\phi_1}$	2.078	0.0246
60	$e^{60\phi_1}$	2.652	0.0321
75	$e^{75\phi_1}$	3.384	0.0405
90	$e^{90\phi_1}$	4.318	0.0502

that died, were resorbed, or had a cleft palate ([Holson and others, 1992](#); [Chen and Gaylor, 1992](#)). They observed the number of implanted fetuses, number of “affected” fetuses, and the dose of 2,4,5-T for each mouse in the experiment and these are given in Table 1 of [George and Bowman \(1995\)](#). The mice were grouped into six levels, receiving doses of 0, 30, 45, 60, 75, or 90 mg/kg of 2,4,5-T. The responses of littermates are correlated because the fetuses gestate in the same mother. Let n_i be the number of implanted fetuses (cluster size) in dam i , let d_i be the dose, and let r_i be the number of fetuses affected. We fit the combined model with covariate vector $z_i = (1, d_i)$.

The results of the regression are given in Table 5. The first two lines give estimates and standard errors for the elements of ϕ and ψ . The next lines give α and β , stratified by different dose level, where the standard errors were obtained by the delta method. Both α and β increase with dose level, and β increases much more quickly than α . Therefore, both exogenous risk and within-cluster effects appear to be significantly related to the number of affected fetuses—and litter size—in this experiment. The baseline risk and infectivity are very small in the absence of 2,4,5-T, and the “infectivity” of each affected fetus increases with dose. We obtain $L = -753.0$ and $\chi^2 = 1044.4$ for the fitted model. In toxicity trials, the relationship between dose and risk for individual units is often of greatest interest. Letting d be the dose of toxin delivered, $e^{\phi_1 d}$ is an estimate of the dose-dependent relative risk to unaffected units, not due to contagion. Table 6 gives estimates and standard errors for the RR in this experiment. For example, at dose 90 mg/kg, 2,4,5-T delivers a more than four-fold increase in the risk to an individual fetus, over that to which a fetus gestating in a control ($d = 0$) mouse is subject.

5. DISCUSSION

The paradigm of George and Bowman (1995) is useful because the likelihood for any dependency model of exchangeable Bernoulli variables can be expressed simply. However, it can be difficult to translate knowledge of the dependency pattern into the joint outcome probabilities necessary to write the likelihood. In this work, we have developed a flexible class of Markov counting models for analyzing clustered binary data. We have established a correspondence between these models and sums of dependent Bernoulli variables under the framework of George and Bowman (1995). We believe the combined model outlined in Section 3.1.3 is most useful. Inference under this model addresses a fundamental question in infectious disease epidemiology: estimating $\beta > 0$ means that some disease risk is due to infectivity or interaction between affected or unaffected units in a cluster.

SUPPLEMENTARY MATERIAL

Supplementary Material is available at <http://biostatistics.oxfordjournals.org>.

ACKNOWLEDGMENTS

We thank Theodore R. Holford and Hongyu Zhao for helpful comments on the manuscript. Forest W. Crawford was funded by NIH/NCATS grant KL2TR000140. *Conflict of Interest*: None declared.

FUNDING

D.Z. was supported, in part, by grants R01 CA168733, P30 CA16359, R01 CA177719, awarded by NIH/NCI, R01 ES005775 awarded by NIH/NIEHS, P30 MH 06229407 awarded by NIH/NIMH, and P01-NS047399 awarded by NIH/NINDS.

REFERENCES

- ALTHAM, P. M. E. (1978). Two generalizations of the binomial distribution. *Applied Statistics* **27**(2), 162–167.
- BLOM, G. AND HOLST, L. (1991). Embedding procedures for discrete problems in probability. *Mathematical Scientist* **16**, 29–40.
- BLOM, G., HOLST, L. AND SANDELL, D. (1994). *Problems and Snapshots from the World of Probability*. New York: Springer.
- BOWMAN, D. AND GEORGE, E. O. (1995). A saturated model for analyzing exchangeable binary data: Applications to clinical and developmental toxicity studies. *Journal of the American Statistical Association* **90**(431), 871–879.
- BRITTON, T. (1997). Tests to detect clustering of infected individuals within families. *Biometrics* **53**, 98–109.
- CHEN, J. J. AND GAYLOR, D. W. (1992). Correlations of developmental end points observed after 2,4,5-trichlorophenoxyacetic acid exposure in mice. *Teratology* **45**(3), 241–246.
- DANG, X., KEETON, S. L. AND PENG, H. (2009). A unified approach for analyzing exchangeable binary data with applications to developmental toxicity studies. *Statistics in Medicine* **28**(20), 2580–2604.
- DEFINETTI, B. (1931). *Funzione caratteristica di un fenomeno aleatorio*. Academia Nazionale del Linceo.

- GEORGE, E. O. AND BOWMAN, D. (1995). A full likelihood procedure for analysing exchangeable binary data. *Biometrics* **51**(2), 512–523.
- GREENWOOD, M. AND YULE, G. U. (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 repeated accidents. *Journal of the Royal Statistical Society* **83**(2), 255–279.
- HASEMAN, J. K. AND SOARES, E. R. (1976). The distribution of fetal death control mice and its implications on statistical tests for dominant lethal effects. *Mutation Research-Fundamental and Molecular Mechanisms of Mutagenesis* **41**(2), 277–287.
- HOLSON, J. F., GAINES, T. B., NELSON, C. J., LABORDE, J. B., GAYLOR, D. W., SHEEHAN, D. M. AND YOUNG, J. F. (1992). Developmental toxicity of 2,4,5-trichlorophenoxyacetic acid (2,4,5-t): I. multireplicated dose-response studies in four inbred strains and one outbred stock of mice. *Fundamental and Applied Toxicology* **19**(2), 286–297.
- KARLIN, S. AND TAYLOR, H. M. (1975). *A First Course in Stochastic Processes*. New York: Academic Press.
- KUK, A. Y. C. (2004). A litter-based approach to risk assessment in developmental toxicity studies via a power family of completely monotone functions. *Journal of the Royal Statistical Society C* **53**(2), 369–386.
- KUPPER, L. L. AND HASEMAN, J. K. (1978). The use of a correlated binomial model for the analysis of certain toxicological experiments. *Biometrics* **34**, 69–76.
- LANGE, K. (2010). *Applied Probability*, 2nd edn. Springer Texts in Statistics. New York: Springer.
- LI, F. P., FRAUMENI, J. F., MULVIHILL, J. J., BLATTNER, W. A., DREYFUS, M. G., TUCKER, M. A. AND MILLER, R. W. (1988). A cancer family syndrome in twenty-four kindreds. *Cancer Research* **48**(18), 5358–5362.
- LIANG, K. Y., ZEGER, S. L. AND QAQISH, B. (1992). Multivariate regression analyses for categorical data. *Journal of the Royal Statistical Society B* **54**(1), 3–40.
- MATTHEWS, A. G., FINKELSTEIN, D. M. AND BETENSKY, R. A. (2008). Analysis of familial aggregation studies with complex ascertainment schemes. *Statistics in Medicine* **27**(24), 5076–5092.
- MCNUTT, L. A., WU, C., XUE, X. AND HAFNER, J. P. (2003). Estimating the relative risk in cohort studies and clinical trials of common outcomes. *American Journal of Epidemiology* **157**(10), 940–943.
- MOORE, D. F., PARK, C. K. AND SMITH, W. (2001). Exploring extra-binomial variation in teratology data using continuous mixtures. *Biometrics* **57**(2), 490–494.
- PANG, Z. AND KUK, A. Y. C. (2005). A shared response model for clustered binary data in developmental toxicity studies. *Biometrics* **61**(4), 1076–1084.
- PANG, Z. AND KUK, A. Y. C. (2007). Test of marginal compatibility and smoothing methods for exchangeable binary data with unequal cluster sizes. *Biometrics* **63**(1), 218–227.
- RENSHAW, E. (2011). *Stochastic Population Processes: Analysis, Approximations, Simulations*. New York: Oxford University Press.
- SASTRY, N. (1997). A nested frailty model for survival data, with an application to the study of child survival in northeast Brazil. *Journal of the American Statistical Association* **92**(438), 426–435.
- STEFANESCU, C. AND TURNBULL, B. W. (2003). Likelihood inference for exchangeable binary data with varying cluster sizes. *Biometrics* **59**(1), 18–24.
- STIRATELLI, R., LAIRD, N. AND WARE, J. H. (1984). Random-effects models for serial observations with binary response. *Biometrics* **40**, 961–971.
- WILLIAMS, D. A. (1975). The analysis of binary responses from toxicological experiments involving reproduction and teratogenicity. *Biometrics* **31**(4), 949–952.

- XU, J. L. AND PROROK, P. C. (2003). Modelling and analysing exchangeable binary data with random cluster sizes. *Statistics in Medicine* **22**(15), 2401–2416.
- YU, C. AND ZELTERMAN, D. (2002a). Statistical inference for familial disease clusters. *Biometrics* **58**(3), 481–491.
- YU, C. AND ZELTERMAN, D. (2002b). Sums of dependent Bernoulli random variables and disease clustering. *Statistics and Probability Letters* **57**(4), 363–373.
- YU, C. AND ZELTERMAN, D. (2008). Sums of exchangeable Bernoulli random variables for family and litter frequency data. *Computational Statistics and Data Analysis* **52**(3), 1636–1649.
- ZEGER, S. L. AND LIANG, K. Y. (1986). Longitudinal data analysis for discrete and continuous outcomes. *Biometrics* **42**(1), 121–130.
- ZOU, G. (2004). A modified Poisson regression approach to prospective studies with binary data. *American Journal of Epidemiology* **159**(7), 702–706.

[Received March 31, 2014; revised February 11, 2015; accepted for publication February 12, 2015]