Statistical Methods in Medical Research

http://smm.sagepub.com/

Spatiotemporal hurdle models for zero-inflated count data: Exploring trends in emergency department visits

Brian Neelon, Howard H Chang, Qiang Ling and Nicole S Hastings

Stat Methods Med Res published online 28 March 2014
DOI: 10.1177/0962280214527079

The online version of this article can be found at:
http://smm.sagepub.com/content/early/2014/03/27/0962280214527079

Published by:

SAGE

http://www.sagepublications.com

Additional services and information for *Statistical Methods in Medical Research* can be found at:

Email Alerts: http://smm.sagepub.com/cgi/alerts

Subscriptions: http://smm.sagepub.com/subscriptions

Reprints: http://www.sagepub.com/journalsReprints.nav

Permissions: http://www.sagepub.com/journalsPermissions.nav

>> OnlineFirst Version of Record - Mar 28, 2014

What is This?

Downloaded from smm.sagepub.com at EMORY UNIV on April 29, 2014

Article

SMMR
STATISTICAL METHODS IN MEDICAL RESEARCH

Statistical Methods in Medical Research
0(0) 1–19
© The Author(s) 2014
Reprints and permissions:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/0962280214527079
smm.sagepub.com
SAGE

Spatiotemporal hurdle models for zero-inflated count data: Exploring trends in emergency department visits

Brian Neelon,[1,2] Howard H Chang,[3] Qiang Ling[3] and Nicole S Hastings[1,4,5]

Abstract

Motivated by a study exploring spatiotemporal trends in emergency department use, we develop a class of two-part hurdle models for the analysis of zero-inflated areal count data. The models consist of two components—one for the probability of any emergency department use and one for the number of emergency department visits given use. Through a hierarchical structure, the models incorporate both patient- and region-level predictors, as well as spatially and temporally correlated random effects for each model component. The random effects are assigned multivariate conditionally autoregressive priors, which induce dependence between the components and provide spatial and temporal smoothing across adjacent spatial units and time periods, resulting in improved inferences. To accommodate potential overdispersion, we consider a range of parametric specifications for the positive counts, including truncated negative binomial and generalized Poisson distributions. We adopt a Bayesian inferential approach, and posterior computation is handled conveniently within standard Bayesian software. Our results indicate that the negative binomial and generalized Poisson hurdle models vastly outperform the Poisson hurdle model, demonstrating that overdispersed hurdle models provide a useful approach to analyzing zero-inflated spatiotemporal data.

Keywords

emergency department use, generalized Poisson distribution, hurdle model, overdispersion, spatiotemporal model, zero inflation

[1]Center for Health Services Research in Primary Care, Durham VAMC, USA
[2]Department of Biostatistics and Bioinformatics, Duke University, USA
[3]Department of Biostatistics, Emory University, USA
[4]Department of Medicine, Duke University School of Medicine, USA
[5]Geriatrics Research, Education and Clinical Center, Durham VAMC, USA

Corresponding author:
Brian Neelon, Department of Biostatistics and Bioinformatics, Box 2721, Duke University Medical Center, Durham, NC 27710-2721, USA.
Email: brian.neelon@duke.edu

Downloaded from smm.sagepub.com at EMORY UNIV on April 29, 2014

1 Introduction

Emergency department (ED) visits have been rising steadily in the United States for more than a decade. Between 1999 and 2009, ED visits increased 32%, to about 136 million per year.[1] Many of these visits were for nonurgent care and could have been managed in outpatient settings, saving an estimated $4.4 billion in health care costs annually.[2] Patient wait times in EDs also increased during this time, from an average of 47 min in 2003 to 58 min in 2009.[1] The use of EDs for routine care places undue strain on the health system, increasing health costs, limiting access to services, and reducing patient satisfaction.[3]

As with other health services, there is considerable geographic and temporal variation in ED use. This variability is due in part to differences in population demographics such as age and overall health status. Characteristics of neighborhoods themselves, such as lack of reliable public transportation to primary care facilities, also contribute to increased ED use.[4] In addition, availability of outpatient clinics varies at a local level, although the findings regarding the relationship between access to outpatient care and ED use are mixed.[5,6] While some policy makers have advocated measures such as neighborhood health centers to promote access to care, it is not possible to identify areas in need of such services without an effective approach to tracking ED trends spatially and over time. Methods for assessing spatiotemporal patterns in ED use are therefore needed as a first step toward optimal design and targeting of community-level interventions to reduce costly nonurgent ED care.

With these goals in mind, investigators at Duke University in Durham, North Carolina, recently reviewed annual ED records from the Duke University Decision Support Repository (DSR), a database containing demographic, diagnostic, and treatment information on over four million patients seen within the Duke University Health System. The review was part of an ongoing study examining patient- and neighborhood-level factors associated with ED use. By tracking spatial trends over time, the investigators sought to identify areas where ED use remained persistently high, fluctuated from year to year, or increased systematically over time.

From an analytic perspective, the DSR data posed a number of unique challenges. First, there was an abundance of zeros: in any given year, over 70% of the patients made no ED visits, while others made multiple visits ranging upward of 90 per year in extreme cases. Thus, one of our chief aims was to develop a set of spatiotemporal models that could account for zero inflation. Second, given the wide range of positive counts, the models needed to address potential overdispersion in the data. Third, in a previous cross-sectional study,[7] we found that the probability of ED use was correlated with the expected number of ED visits among users after adjusting for covariates; we therefore sought to develop models to accommodate this source of dependence. And finally, to improve small-area estimation, the models needed to provide adequate spatial and temporal smoothing.

This article describes a class of two-part hurdle models specifically designed to address these multiple aims. The models consist of two components: a binary component that models the probability of any ED use (i.e., at least one ED visit annually) and a truncated count component that models the number of visits among users. Together, these components accommodate both the high proportion of zeros and the right-skewness observed among the nonzero counts. To accurately model the dispersion in the positive counts, we consider three distributions for the nonzero observations: a truncated Poisson, a truncated negative binomial, and a truncated generalized Poisson distribution. Taking advantage of the unique hierarchical structure of the DSR data, our models incorporate both patient- and region-level predictors, as well as spatially and temporally correlated random effects for each model component. The random effects are assigned multivariate conditionally autoregressive priors that induce dependence between the components and provide smoothing across adjacent spatial units and time periods.

Downloaded from smm.sagepub.com at EMORY UNIV on April 29, 2014

Our approach builds on recent developments in spatial modeling of zero-inflated data. Agarwal et al.[8] developed a spatial zero-inflated Poisson (ZIP) model that introduced spatially correlated random effects into the Poisson component. Rathbun and Fei[9] proposed a similar model in which the "extra-Poisson" zeros were fitted using a spatial probit model. Gschlößl and Czado[10] developed a spatial zero-inflated generalized Poisson model to study the incidence of meningococcal disease. More recently, Recta et al.[11] proposed a correlated spatial Poisson hurdle model for point-referenced zero-inflated data. In the spatiotemporal setting, Fuentes et al.[12] developed a noninflated generalized Poisson model for fine particulate risk assessment, while Ver Hoef and Jansen[13] introduced spatiotemporal ZIP and Poisson hurdle models to investigate haulout patterns among harbor seals.

The approach taken here extends this prior work in several ways. First, we propose dynamic space-time models that provide temporal smoothing of each spatial unit through a set of autoregressive interactions, thereby improving small-area estimation. Second, we include multivariate space-time random effects linking the model components, which have been shown to improve inferences.[7] Third, we incorporate individual- and region-level information through a hierarchical structure to better explain spatiotemporal trends. And finally, we consider a range of distributional specifications designed to address overdispersion. The models have the additional advantage of being easy to implement in standard Bayesian software packages, such as WinBUGS.[14]

The remainder of the paper is organized into five sections: the following section describes the DSR data, Section 3 outlines the proposed models, Section 4 details the Bayesian model-fitting approach, Section 5 applies the models to the DSR data, and Section 6 provides a discussion and points to areas for future research.

2 The DSR data

The Duke University DSR has been in existence for over a decade. Originally built as an administrative and financial database, the DSR holds 17 years of demographic, diagnostic, and billing data on over four million patients seen in the Duke University Health System. The data have been deployed for secondary use in numerous research studies and quality improvement initiatives.[15,7]

As part of an ongoing study exploring contributors to ED use, university researchers recently reviewed patient records for Durham County residents who were seen at either an ED or non-ED clinic at least once between 2007 and 2011, the most recent year for which records were available. The records were georeferenced by residential address and subsequently linked at the Census block group level to data from the 2005–2009 American Community Survey.[16] Block groups are collections of residential blocks, and as such, form the second-smallest geographic level set forth by the US Census Bureau.

The final dataset contained over 122,000 records for approximately 40,000 patients and included information on the annual number of ED visits for each patient; patient demographics, such as age, race, gender, and insurance status; and median household income of each of the 129 Census block groups in Durham County. Because patients could appear in the database in multiple years, the number of records per subject ranged from one to five, with most patients having three or fewer records over the course of the five-year study period.

Table 1 provides summary statistics for the DSR data. The majority of the sample was female and of non-Hispanic White or non-Hispanic Black race. The median age was 37 years. About 60% had

Downloaded from smm.sagepub.com at EMORY UNIV on April 29, 2014

Table 1. Summary of DSR patient records ($N = 122{,}273$).

Variables	n	%
Male	49,719	41
Race		
Non-Hispanic White	56,734	46
Non-Hispanic Black	51,528	42
Hispanic	7523	6
Asian	3165	3
Other	3323	3
Insurance		
Duke private insurance	13,932	12
Other private insurance	58,918	48
Medicaid	17,761	14
Medicare	19,493	16
Self-pay	12,169	10
	Median	Range
Age (years)	37	(< 1, 103)
Block group median household income ($)	45,330	(5980, 134K)
Block group sample size	776	(39, 3212)

private medical insurance, 12% as part of a University-sponsored plan. The median block group income was just over $45K, approximately $5000 below the national average.[17] The median block group sample size, combined over five years, was 776.

Figure 1 presents a partial histogram of the annual number of ED visits up to 10 visits. Over 70% of the patient-year records were zero, meaning that a patient made no ED visits in a given year. The number of nonzero visits ranged from one to 91, with 5% of the patients having greater than six visits in any one year. Using Vuong's procedure,[18] we tested for zero inflation relative to standard Poisson and negative binomial regression models that included gender, race, insurance, age, and block group median income as predictors. The tests indicated significant zero inflation relative to both count distributions ($p < 0.0001$ in each case).

Figure 2 displays the percent of ED users (upper panel) and the mean number of ED visits among users (lower panel) by block group and year. Patients were considered ED users if they had at least one ED visit during the year. In each year, the percent of ED users within block groups ranged from under 10% to over 60%, with the highest rates occurring in the central part of the county. The spatial pattern in ED use was fairly stable across years. There was considerably more variability in the mean number of visits among users, although generally speaking there was an increase over time, particularly along the county's eastern border. The temporal variation seen here may be due in part to small sample sizes in certain block groups. We therefore expect the proposed models to induce a measure of spatial and temporal smoothing among these smaller areal units.

3 Proposed model

3.1 The hurdle model

For the analysis of the DSR data, we consider a broad class of two-part hurdle models to address both zero inflation and potential overdispersion of the nonzero counts. Hurdle models are two-part

Downloaded from smm.sagepub.com at EMORY UNIV on April 29, 2014

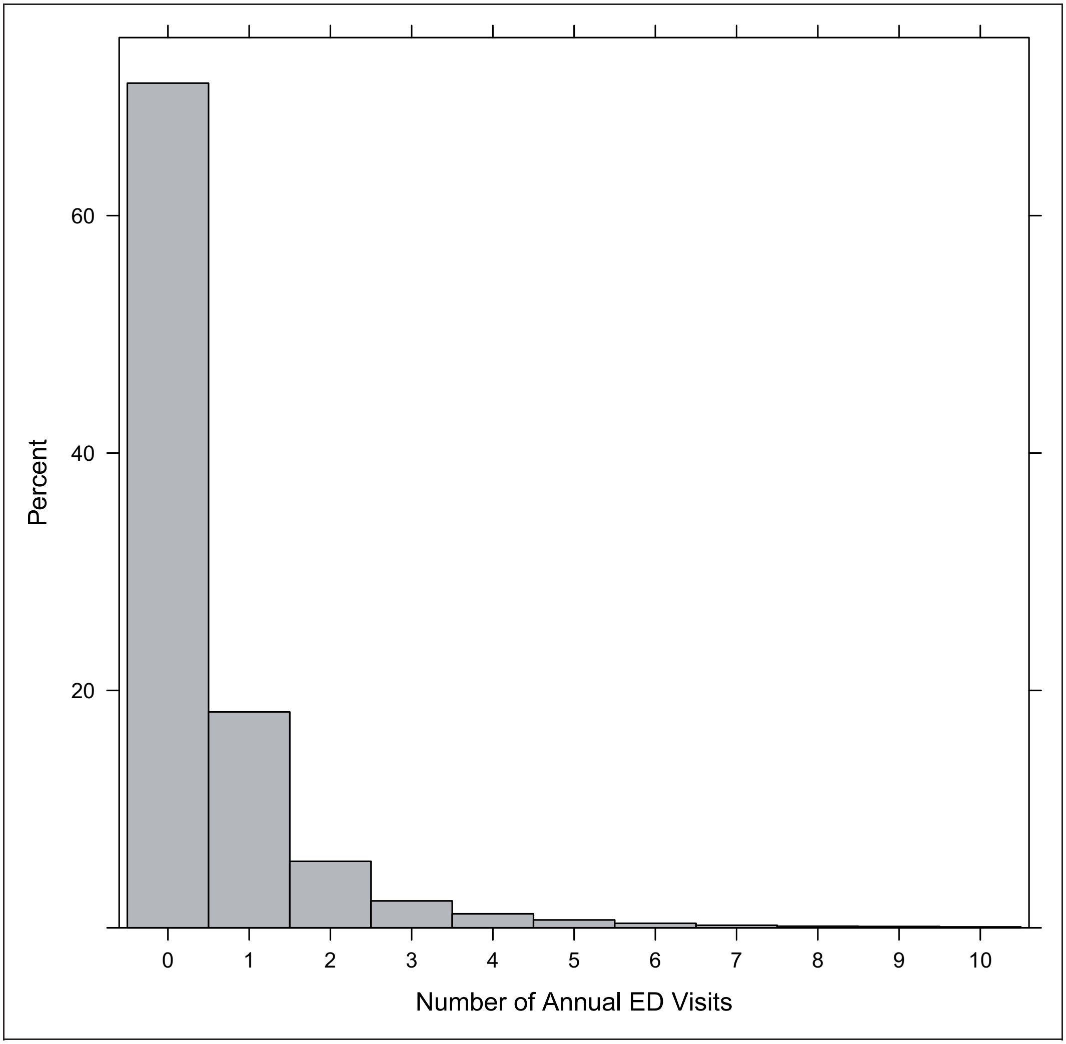

Figure 1. Partial histogram of annual ED visits, 2007–2011.
ED: emergency department.

mixtures consisting of a point mass at zero followed by a zero-truncated count distribution for the positive observations.[19,20] Letting Y denote a count-valued response, the generic structure of the hurdle model is given by

$$\Pr(Y=0)=1-\pi, 0 \leq \pi \leq 1$$
$$\Pr(Y=y)=\frac{\pi p(y;\mu,\alpha)}{1-p(0;\mu,\alpha)},\ \mu>0, \alpha>0, y=1,2,\ldots, \tag{1}$$

Downloaded from smm.sagepub.com at EMORY UNIV on April 29, 2014

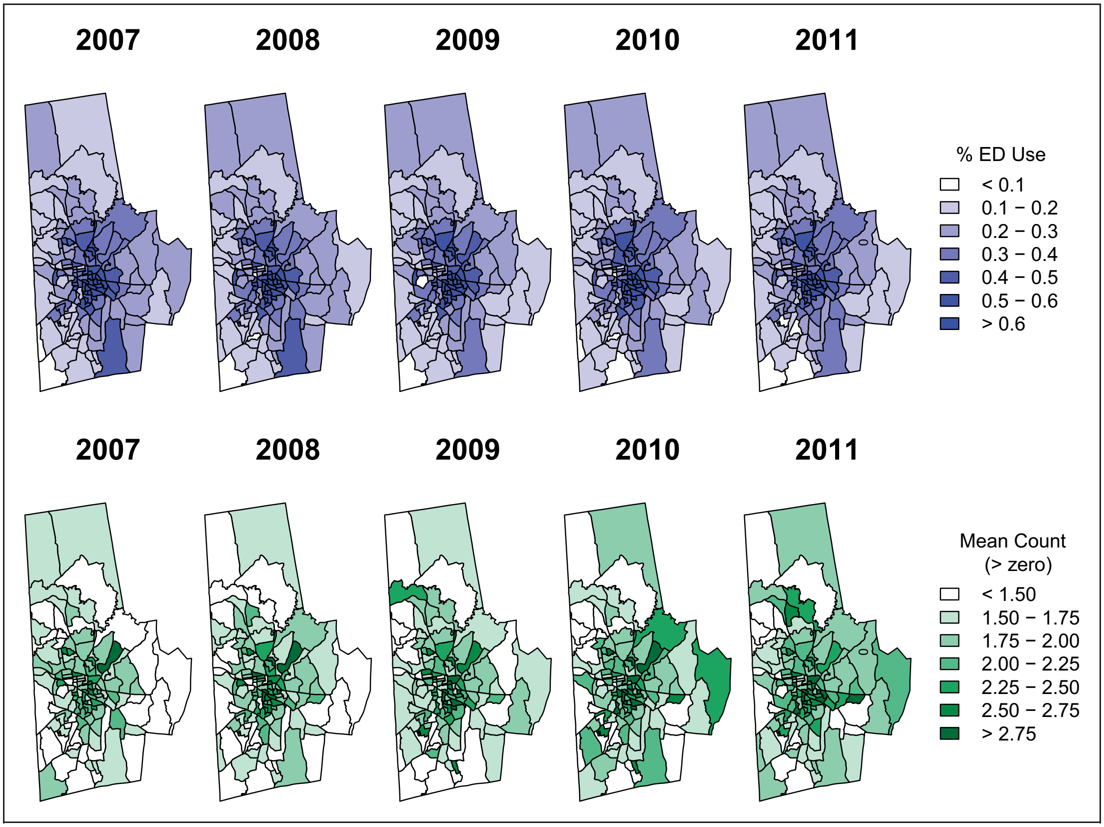

Figure 2. Percentage of ED users (upper panel) and mean number of ED visits among those with at least one ED visit (lower panel) for each of the 129 Durham County block groups from 2007 to 2011.
ED: emergency department.

where $\pi = \Pr(Y > 0)$ is the probability of a nonzero response; $p(y; \mu, \alpha)$ is an untruncated, or base, probability distribution with mean μ and dispersion parameter α; and $p(0; \mu, \alpha)$ is the base distribution evaluated at 0. This can be written more compactly as

$$Y \sim (1-\pi)1_{(y=0)} + \pi \frac{p(y; \mu, \alpha)}{1 - p(0; \mu, \alpha)} 1_{(y>0)}$$

where $1_{(\cdot)}$ denotes the indicator function. In health services research, π is known as the *utilization probability*—i.e., the probability of using services at least once. When $1 - \pi = p(0; \mu, \alpha)$, the hurdle model reduces to its base distribution; when $(1 - \pi) > p(0; \mu, \alpha)$, the zeros are inflated relative to the base distribution; and when $(1 - \pi) < p(0; \mu, \alpha)$, there is zero deflation. Typically, one assumes that π is strictly between 0 and 1, so that there is a nonzero utilization probability for all individuals under study.

Zero-inflated models[21] offer an alternative to hurdle models for the analysis of zero-adjusted count data. Unlike hurdle models, zero-inflated models partition zeros into two types: "structural" zeros (e.g., those that occur because a patient is ineligible for health services) and

Downloaded from smm.sagepub.com at EMORY UNIV on April 29, 2014

"chance" zeros (those that occur by chance among eligible patients). Zero-inflated models are thus an ideal choice when one hypothesizes the existence of a latent population of individuals with structural zeros. For the DSR analysis, we relied on hurdle rather than zero-inflated models because we regarded all patients as eligible for ED care, and hence there was no need to separately model the structural zeros.

3.2 Spatiotemporal hurdle model

To extend model (1) to the spatiotemporal setting, we propose the following space-time hurdle model

$$\begin{aligned}
\Pr(Y_{ijk} = y_{ijk}|\boldsymbol{\phi}_i, \boldsymbol{v}_j, \boldsymbol{\delta}_{ij}) &= (1-\pi_{ijk})1_{(y_{ijk}=0)} + \frac{\pi_{ijk}p(y_{ijk}; \mu_{ijk}, \alpha)}{1-p(0; \mu_{ijk}, \alpha)}1_{(y_{ijk}>0)} \\
\text{logit}(\pi_{ijk}) &= \boldsymbol{x}'_{ijk}\boldsymbol{\beta}_1 + f_1(z_{ijk}) + \phi_{1i} + \nu_{1j} + \delta_{1ij} \\
\ln(\mu_{ijk}) &= \boldsymbol{x}'_{ijk}\boldsymbol{\beta}_2 + f_2(z_{ijk}) + \phi_{2i} + \nu_{2j} + \delta_{2ij}, \\
i &= 1, \ldots, 129;\ j = 1, \ldots, 5;\ k = 1, \ldots, n_{ij}
\end{aligned} \tag{2}$$

where Y_{ijk} denotes the number of annual ED visits for the kth patient in block group i and year j; $\pi_{ijk} = \Pr(Y_{ijk} > 0)$ is the probability of a positive response; μ_{ijk} is the conditional mean of the base distribution given a set of spatiotemporal random effects; $\boldsymbol{x}_{ijk}$ is a $p \times 1$ vector of fixed-effect, individual- and region-level predictors; $\boldsymbol{\beta}_l$ is a $p \times 1$ vector of fixed-effect regression coefficients for component $l (l = 1, 2)$; $f_1(z_{ijk})$ and $f_2(z_{ijk})$ are optional smooth functions of a continuous predictor z_{ijk} (e.g., patient age) to be modeled via splines; $\boldsymbol{\phi}_i = (\phi_{1i}, \phi_{2i})'$ is a vector of spatially dependent "main effects" for the ith block group; $\boldsymbol{v}_j = (\nu_{1j}, \nu_{2j})'$ is a vector of temporal main effects for year j; and $\boldsymbol{\delta}_{ij} = (\delta_{1ij}, \delta_{2ij})'$ denotes a vector of space-time interactions.

Thus, we partition the spatiotemporal effects into three parts: a purely spatial component, represented by $\boldsymbol{\phi}_i$; a purely temporal component, represented by $\boldsymbol{v}_j$; and a residual interaction term, $\boldsymbol{\delta}_{ij}$. Together, these parameters capture unobserved block group effects over time. For example, $\boldsymbol{\phi}_i + \boldsymbol{v}_1 + \boldsymbol{\delta}_{i1}$ represents the unobserved effect for block group i in year 1; $\boldsymbol{\phi}_i + \boldsymbol{v}_2 + \boldsymbol{\delta}_{i2}$ denotes the effect for block group i in year 2; and so on. Autoregressive priors are subsequently used to provide spatiotemporal smoothing and pooling of information across neighboring block groups and adjacent years. Prior specification is discussed in greater detail in Section 4. For the analysis of the DSR data, we assume identical sets of predictors for the two components. In general, however, one might allow for unique predictors for the two components if the goal is to obtain a parsimonious model by removing extraneous variables from one component, or if there is a priori scientific reason to believe that the two components are associated with unique sets of predictors.

3.3 Choice of base distribution

To accurately capture the variability in the positive counts, we consider three choices for the base distribution: the Poisson, negative binomial, and generalized Poisson distribution.[22] The *spatiotemporal Poisson hurdle model* is expressed as

$$\begin{aligned}
\Pr(Y_{ijk} = y_{ijk}|\boldsymbol{\phi}_i, \boldsymbol{v}_j, \boldsymbol{\delta}_{ij}) &= (1-\pi_{ijk})1_{(y_{ijk}=0)} + \pi_{ijk}\text{TPois}(y_{ijk}; \mu_{ijk})1_{(y_{ijk}>0)} \\
&= (1-\pi_{ijk})1_{(y_{ijk}=0)} + \frac{\pi_{ijk}\mu_{ijk}^{y_{ijk}}\mathrm{e}^{-\mu_{ijk}}}{y_{ijk}!(1-\mathrm{e}^{-\mu_{ijk}})}1_{(y_{ijk}>0)}
\end{aligned} \tag{3}$$

Downloaded from smm.sagepub.com at EMORY UNIV on April 29, 2014

where $\mathrm{TPois}(\cdot;\mu_{ijk})$ denotes a truncated Poisson distribution, μ_{ijk} is the conditional mean of the Poisson base distribution given the random effects, and π_{ijk} and μ_{ijk} are modeled as in equation (2). The Poisson distribution implies equidispersion, or equivalence of the conditional mean and variance (hence no additional dispersion parameter, α). In many applications, this assumption is restrictive and can result in poor model fit.

An alternative is to select a negative binomial base distribution, giving rise to the *spatiotemporal negative binomial hurdle model*

$$\begin{aligned}\Pr(Y_{ijk}=y_{ijk}|\boldsymbol{\phi}_i,\boldsymbol{v}_j,\boldsymbol{\delta}_{ij}) &= (1-\pi_{ijk})1_{(y_{ijk}=0)}+\pi_{ijk}\mathrm{TNegBin}(y_{ijk};\mu_{ijk},\alpha)1_{(y_{ijk}>0)}\\ &= (1-\pi_{ijk})1_{(y_{ijk}=0)}+\frac{\pi_{ijk}}{1-\left(\frac{\alpha}{\mu_{ijk}+\alpha}\right)^{\alpha}}\frac{\Gamma(y_{ijk}+\alpha)}{\Gamma(\alpha)\,y_{ijk}!}\\ &\quad\times\left(\frac{\mu_{ijk}}{\mu_{ijk}+\alpha}\right)^{y_{ijk}}\left(\frac{\alpha}{\mu_{ijk}+\alpha}\right)^{\alpha}1_{(y_{ijk}>0)},\quad \alpha>0 \end{aligned}\tag{4}$$

where $\mathrm{TNegBin}(\cdot;\mu_{ijk},\alpha)$ denotes a truncated type-2 negative binomial distribution, μ_{ijk} is the conditional mean of the negative binomial base distribution, and π_{ijk} and μ_{ijk} are modeled as in equation (2). The negative binomial base distribution is appealing if there is evidence of overdispersion relative to the Poisson—i.e., a variance exceeding the mean. In particular, if $X\sim\mathrm{NegBin}\,(\mu,\alpha)$, then $\mathrm{E}(X)=\mu$ and $\mathrm{V}(X)=\mu(1+\mu/\alpha)$, and hence $(1+\mu/\alpha)$ is a measure of overdispersion. As $\alpha\to\infty$, the negative binomial converges to a Poisson distribution with mean and variance equal to μ. The added flexibility of the negative binomial in accommodating heterogeneity can yield improved model fit for highly dispersed count data.

Lastly, we consider the spatiotemporal generalized Poisson hurdle model

$$\begin{aligned}\Pr(Y_{ijk}=y_{ijk}|\boldsymbol{\phi}_i,\boldsymbol{v}_j,\boldsymbol{\delta}_{ij}) &= (1-\pi_{ijk})1_{(y_{ijk}=0)}+\pi_{ijk}\mathrm{TGPois}(y_{ijk};\mu_{ijk},\alpha)1_{(y_{ijk}>0)}\\ &= (1-\pi_{ijk})1_{(y_{ijk}=0)}+\frac{\pi_{ijk}}{1-\exp\left(-\frac{\mu_{ijk}}{1+\alpha\mu_{ijk}}\right)}\left(\frac{\mu_{ijk}}{1+\alpha\mu_{ijk}}\right)^{y_{ijk}}\\ &\quad\times\frac{(1+\alpha y_{ijk})^{y_{ijk}-1}}{y_{ijk}!}\exp\left[-\frac{\mu_{ijk}(1+\alpha y_{ijk})}{1+\alpha\mu_{ijk}}\right]1_{(y_{ijk}>0)}\end{aligned}\tag{5}$$

for $y_{ijk}=0,1,\ldots,C(\alpha)$. Here, $\mathrm{TGPois}(\cdot;\mu_{ijk},\alpha)$ denotes a truncated generalized Poisson distribution with (untruncated) conditional mean μ_{ijk} and dispersion parameter $\alpha\in(-1/y_{\max},\infty)$, where $y_{\max}$ is the maximum observed response. To ensure a valid probability distribution, $C(\alpha)=\lfloor -1/\alpha\rfloor$ for $\alpha<0$, and $C(\alpha)=\infty$ otherwise, where $\lfloor\cdot\rfloor$ denotes the floor function. Finally, π_{ijk} and μ_{ijk} are modeled as in equation (2).

As in the negative binomial case, α functions as a heterogeneity parameter accommodating departures from equidispersion. In particular, if $X\sim\mathrm{GPois}\;(\mu,\alpha)$, then $\mathrm{E}(X)=\mu$ and $\mathrm{V}(Y)=\mu(1+\alpha\mu)^2$. When $\alpha=0$, the generalized Poisson reduces to the Poisson distribution; when $\alpha>0$, $\mathrm{V}(X)>\mathrm{E}(X)$ and there is overdispersion; and when $\alpha<0$, $\mathrm{V}(X)<\mathrm{E}(X)$ and there is underdispersion. Thus, unlike the negative binomial, the generalized Poisson allows for underdispersion. Moreover, while both distributions accommodate overdispersion, the generalized Poisson has a heavier tail compared to a negative binomial with the same first two moments and is therefore well suited for highly skewed data such as ours.[23]

Downloaded from smm.sagepub.com at EMORY UNIV on April 29, 2014

4 Bayesian inference

We adopt a Bayesian modeling approach and assign prior distributions to all model parameters. In previous work, we found that the probability of ED use was positively associated with the expected number of visits given use (i.e., the model components were correlated) after accounting for both individual- and cluster-level covariates. We showed that explicitly modeling this between-component correlation improved inferences. To accommodate this association in the current study, and to provide adequate spatial and temporal smoothing, we assume bivariate intrinsic conditional autoregressive model (BICAR) priors for the spatiotemporal random effects.[24] For example, for the spatial main effects, we assume the following BICAR prior

$$\boldsymbol{\phi}_i|\boldsymbol{\phi}_{(-i)}, \Sigma_\phi \sim \mathrm{N}_2\left(\frac{1}{m_i}\sum_{l\in\partial_i}\boldsymbol{\phi}_l, \frac{1}{m_i}\Sigma_\phi\right) \tag{6}$$

where ∂_i is the set of neighbors sharing a geographic border with block group i, m_i is the number of neighbors, and Σ_ϕ denotes the 2×2 covariance of $\boldsymbol{\phi}_i$ conditional on the remaining spatial random effects, $\boldsymbol{\phi}_{(-i)}$. To ensure proper posteriors, sum-to-zero constraints are applied to the random effects. The correlation between ϕ_{1i} and ϕ_{2i}, denoted ρ_ϕ, is easily derived from Σ_ϕ. This parameter is of interest because it captures, in part, the within-block-group association between the model components. When $\rho_\phi > 0$, for example, block groups with higher rates of ED use tend to have higher mean counts among users, after adjusting for observed factors such as median household income.

For the temporal main effects, we consider two specifications. First, we assign a BICAR prior to $\boldsymbol{\nu}_j\,(j = 1, \ldots, 5)$ analogous to the prior for the spatial effects. This choice is particularly beneficial when the temporal units are sparse, because it allows adjacent time periods to pool information to improve efficiency. However, given that we have sufficient sample size per year in the current study for estimating the temporal effects, any smoothing imposed by the BICAR prior is likely to be minimal. Therefore, in addition to BICAR temporal effects, we consider models with fixed annual effects; here, we assign independent normal priors to ν_{1j} and $\nu_{2j}\,(j = 2, \ldots, 5)$, with ν_{11} and ν_{21} set to 0 in correspondence with the reference year 2007.

For the space-time interactions, we assume a first-order dynamic BICAR prior,[25] whereby $\boldsymbol{\delta}_{ij} = \rho\boldsymbol{\delta}_{i(j-1)} + \boldsymbol{\psi}_{ij}\,(j = 2, \ldots, 5)$, $\boldsymbol{\psi}_{ij}$ is BICAR, and ρ is a temporal smoothing parameter common to both components. For identifiability, we set $\boldsymbol{\delta}_{i1} = 0\ \forall i$. Unlike with the annual main effects, temporal smoothing is needed here to improve small-area estimation, particularly when one considers that the minimum block group sample size in a given year is five, occurring in 2011. Potentially, one might introduce separate temporal smoothing parameters (ρ_1 and ρ_2, say) for the two model components. However, with only five years of data in our application, we assumed similar levels of smoothing across components.

To complete the prior specification, we assign improper priors to the fixed-effect intercepts, diffuse normal priors to the remaining fixed effects and spline coefficients, inverse-Wishart (IW) priors to covariance matrices, and a uniform U(0, 1) prior to the temporal autoregressive parameter, ρ. For the negative binomial hurdle model, we assign a Gamma prior to α, and for the generalized Poisson hurdle model, we assume $\alpha \sim \mathrm{U}(-1/y_{\max}, M)$ where $M > 0$ is chosen large enough to account for possible overdispersion. For example, with $M = 10$, the maximum allowable overdispersion under the generalized Poisson with mean μ is $(1 + 10\mu)^2$, which generally exceeds values one might expect in practice. Detailed prior specification for the DSR analysis is provided in Section 5.

Posterior computation proceeds via Markov chain Monte Carlo (MCMC), which can be implemented easily within WinBUGS. Although WinBUGS does not have a pre-designated function for truncated count distributions, one can apply the "zeros trick" to explicitly define the

Downloaded from smm.sagepub.com at EMORY UNIV on April 29, 2014

hurdle likelihood.[26] The BICAR prior can be specified with the mv.car function, and the remaining MCMC steps are readily coded using standard WinBUGS syntax.

We monitor MCMC convergence using trace plots and Geweke's z-test,[27] which assesses the compatibility of disjoint portions of the sampler. For model comparison, we adopt the deviance information criterion (DIC) proposed by Spiegelhalter et al.[28] DIC is defined as $\text{DIC} = \bar{D} + p_D$, where $\bar{D}$ is a measure of current model fit and p_D is a penalty for model complexity. For fixed effect models, the complexity—as measured by the number of model parameters—is easily determined. For random effect models, the dimension of the parameter space is less clear and depends on the degree of heterogeneity between subjects (more heterogeneity implies more "effective" parameters). The penalty p_D was designed as a way to estimate the number of effective parameters in Bayesian hierarchical models. As with other penalized selection criteria, smaller values of DIC are considered preferable.

To further evaluate model fit, we propose a series of posterior predictive assessments, whereby the observed data are compared to data replicated from the posterior predictive distribution.[29] If the model fits well, the replicated data should resemble the observed data. To quantify the degree of similarity, one typically chooses a "discrepancy statistic," such as a sample moment or quantile that captures some important aspect of the data. For the DSR analysis, we adopt three discrepancy measures: the sample proportion of zeros and the sample mean and variance among the positive observations. For each measure, we compute the posterior predictive mean and 95% credible interval (CrI). A 95% CrI that includes the observed sample value suggests adequate model fit. In addition to the above measures, for the final model we also produce a histogram comparing the observed and posterior-predictive counts of ED visits.

5 Analysis of the DSR data

We fit the Poisson, negative binomial, and generalized Poisson versions of the spatiotemporal hurdle model to the DSR data, with a logit link for the binary component. For each choice of base distribution, we considered two sub-models: a model with fixed temporal main effects and one with BICAR temporal effects. The models also included patient race, gender, age, and insurance, and block group median income as predictors. Since previous studies have suggested a nonlinear effect for patient age,[30] we modeled age using cubic B-splines with interior knots at the first, second, and third quartiles of the age distribution (20, 37, and 55 years, respectively).

For the model with fixed annual effects, we assumed independent $\text{N}(0, 100)$ priors for the temporal main effects; for the model with random temporal effects, we assumed a BICAR prior with an $\text{IW}(3, \boldsymbol{I}_2)$ prior for the conditional covariance Σ_ν, where $\boldsymbol{I}_2$ denotes the two-dimensional identity matrix. For both sub-models, we assigned improper priors to the fixed intercept parameters and $\text{N}(0, 100)$ priors to the remaining regression and spline coefficients. In addition, we assumed a BICAR prior for $\boldsymbol{\phi}_i$, an $\text{IW}(3, \boldsymbol{I}_2)$ distribution for Σ_ϕ, a dynamic BICAR prior for $\boldsymbol{\delta}_{ij} > (j = 2, \ldots, 5)$, a U(0,1) prior for ρ, and a BICAR for $\boldsymbol{\psi}_{ij}$ with an $\text{IW}(3, \boldsymbol{I}_2)$ prior for Σ_ψ. For the negative binomial sub-models, we assumed $\alpha \sim \text{Ga}(0.01, 0.01)$; for the generalized Poisson sub-models, we assumed $\alpha \sim \text{U}(-1/91, 10)$, where 91 represents the maximum observed count and 10 was chosen as a conservative upper bound.

The models were fit in WinBUGS 1.4.3 and called into R[31] using the function R2WinBUGS.[32] We ran the sampler for 15,000 iterations, discarding the first 5000 as burn-in. Trace plots and Geweke diagnostics indicated rapid convergence and efficient mixing of the chains.

Table 2 presents the model comparison results for the various models. The negative binomial and generalized Poisson hurdle models substantially outperformed the Poisson models with respect to DIC. Overall, the generalized Poisson model with fixed temporal effects had the lowest DIC value.

Downloaded from smm.sagepub.com at EMORY UNIV on April 29, 2014

Table 2. Model comparison results.

Base distribution	Temporal effects	DIC	p_D	Posterior predictive checks $\Pr(Y=0)$	$E(Y\|Y>0)$	$V(Y\|Y>0)$
Poisson	Fixed	232,158	566	0.709 (0.705, 0.714)[a]	1.93 (1.90, 1.95)	1.48 (1.43, 1.55)
Poisson	Bivariate CAR	232,171	574	0.709 (0.705, 0.714)	1.93 (1.90, 1.96)	1.49 (1.42, 1.56)
Negative binomial	Fixed	211,198	367	0.709 (0.705, 0.713)	1.93 (1.89, 1.96)	3.76 (3.42, 4.20)
Negative binomial	Bivariate CAR	211,209	377	0.709 (0.705, 0.714)	1.93 (1.89, 1.97)	3.78 (3.42, 4.20)
Generalized Poisson	Fixed	211,035	367	0.709 (0.706, 0.713)	1.93 (1.89, 1.97)	4.66 (4.10, 5.46)
Generalized Poisson	Bivariate CAR	211,046	374	0.710 (0.705, 0.714)	1.93 (1.89, 1.97)	4.67 (4.05, 5.45)
				Observed: 0.709	Observed: 1.94	Observed: 5.89

CAR: conditional autoregressive model.
[a]Posterior median and 95% credible interval.

In terms of posterior predictions, all models accurately reproduced the observed proportion of zeros and the conditional mean among the positive counts. While none of the models did especially well in capturing the observed conditional variance, the ordinary Poisson models showed the poorest fit, further supporting the need to model overdispersion in the counts.

Table 3 presents the posterior means and 95% CrIs for the three hurdle models with fixed annual effects. The effect estimates and intervals for the binary component were similar across models, which is expected since this component has the same structure in all three models. Adjusting for other predictors, male gender, non-Hispanic Black and Hispanic race/ethnicity, and nonprivate insurance (i.e., federally subsidized or self-paid) were associated with increased ED use, while Asian race and Duke insurance were associated with decreased rates of use. Median household income had minimal impact on ED use.

In contrast to the binary component, the parameter estimates in the count component varied substantially across the models, indicating that the choice of base distribution had a significant impact on covariate effects. For example, non-Hispanic Black race showed a much stronger effect for the negative binomial and generalized Poisson models than for the ordinary Poisson models. A similar—albeit less transparent—phenomenon occurred for the federal and self-pay insurance categories: while all models showed a positive effect, the effect was most pronounced in the two overdispersed models.

Interestingly, for all models, the adjusted estimates for male gender and Hispanic race reversed direction between the binary and count components. Hispanics, for example, were more likely than non-Hispanic Whites to visit the ED at least once; however, among ED users, they tended to make fewer visits than Whites. This points to a potential difference between the way Hispanics and non-Hispanic Whites use ED services. In particular, although modest ED use seems to be more ubiquitous among Hispanics, they are less inclined than Whites to use EDs repeatedly.

And finally, there was moderate positive correlation between spatial main effects for the two components, with ρ_ϕ ranging from 0.46 (95% CrI $=[0.29, 0.64]$) for the Poisson hurdle model to 0.56 (95% CrI $=[0.37, 0.72]$) for the negative binomial and generalized Poisson models. This suggests a modest benefit to modeling the spatial association between the components. In particular, we find that block groups with a high proportion of ED users also tended to have higher mean counts among users after adjusting for covariates.

Using the above parameter estimates, it is straightforward to make clinically meaningful inferences of interest. For example, suppose one wishes to compute the incidence density ratio

Downloaded from smm.sagepub.com at EMORY UNIV on April 29, 2014

Table 3. Posterior summaries for the Poisson, negative binomial, and generalized Poisson hurdle models.

	Poisson			Negative binomial			Generalized Poisson		
	Mean	2.5%	97.5%	Mean	2.5%	97.5%	Mean	2.5%	97.5%
Logit(π)									
Intercept	−1.21	−1.31	−1.12	−1.21	−1.33	−1.09	−1.20	−1.36	−1.09
Year[a]									
2008	−0.00	−0.05	0.04	−0.01	−0.05	0.04	−0.00	−0.05	0.04
2009	0.00	−0.04	0.05	0.00	−0.04	0.05	0.01	−0.04	0.05
2010	−0.02	−0.07	0.02	−0.03	−0.07	0.02	−0.02	−0.07	0.03
2011	−0.02	−0.06	0.03	−0.02	−0.06	0.03	−0.02	−0.06	0.03
Male	0.19	0.16	0.21	0.19	0.16	0.22	0.19	0.16	0.22
Race[b]									
Black	0.69	0.65	0.72	0.69	0.65	0.72	0.69	0.65	0.73
Hispanic	0.49	0.43	0.55	0.49	0.42	0.55	0.48	0.42	0.54
Asian	−0.42	−0.54	−0.31	−0.43	−0.54	−0.31	−0.42	−0.55	−0.31
Other	0.11	0.02	0.19	0.11	0.02	0.19	0.11	0.01	0.20
Insurance[c]									
Duke	−0.29	−0.35	−0.24	−0.29	−0.35	−0.24	−0.29	−0.35	−0.24
Medicaid	1.19	1.15	1.24	1.19	1.14	1.24	1.19	1.15	1.24
Medicare	0.81	0.75	0.87	0.81	0.74	0.87	0.81	0.75	0.87
Self-pay	1.54	1.49	1.58	1.54	1.49	1.58	1.54	1.49	1.59
Median income	−0.01	−0.01	0.00	−0.01	−0.01	0.00	−0.01	−0.01	0.00
Log(μ)									
Intercept	0.31	0.23	0.39	−4.02	−4.70	−3.39	−1.35	−1.68	−1.13
Year[a]									
2008	0.05	0.01	0.08	0.08	0.02	0.15	0.09	0.01	0.16
2009	0.08	0.04	0.12	0.09	0.03	0.16	0.10	0.03	0.17
2010	0.10	0.06	0.14	0.12	0.05	0.19	0.14	0.06	0.22
2011	0.16	0.12	0.20	0.20	0.13	0.27	0.22	0.15	0.30
Male	−0.07	−0.09	−0.05	−0.10	−0.14	−0.05	−0.08	−0.13	−0.04
Race[b]									
Black	0.04	0.01	0.07	0.19	0.13	0.25	0.18	0.12	0.24
Hispanic	−0.61	−0.66	−0.55	−0.61	−0.70	−0.51	−0.52	−0.61	−0.44
Asian	−0.77	−0.96	−0.57	−0.77	−1.06	−0.51	−0.70	−0.97	−0.44
Other	−0.44	−0.54	−0.36	−0.44	−0.61	−0.28	−0.37	−0.53	−0.22
Insurance[c]									
Duke	−0.15	−0.22	−0.09	−0.16	−0.27	−0.04	−0.14	−0.24	−0.03
Medicaid	0.67	0.64	0.70	0.91	0.85	0.98	0.80	0.74	0.86
Medicare	0.74	0.70	0.78	0.89	0.80	0.97	0.77	0.69	0.84
Self-pay	0.41	0.38	0.44	0.58	0.51	0.64	0.51	0.46	0.57
Median income	0.00	−0.01	0.00	0.00	−0.01	0.00	0.00	−0.01	0.00
Spatial CAR covariance									
Sigma.phi[1,1]	0.22	0.16	0.29	0.22	0.16	0.29	0.22	0.16	0.29
Sigma.phi[1,2]	0.12	0.07	0.19	0.16	0.10	0.24	0.14	0.09	0.21
Sigma.phi[2,2]	0.31	0.22	0.42	0.39	0.26	0.56	0.29	0.20	0.41
rho.phi	0.46	0.29	0.64	0.56	0.37	0.72	0.56	0.37	0.72

(continued)

Downloaded from smm.sagepub.com at EMORY UNIV on April 29, 2014

Table 3. Continued.

	Poisson			Negative binomial			Generalized Poisson		
	Mean	2.5%	97.5%	Mean	2.5%	97.5%	Mean	2.5%	97.5%
Dynamic CAR covariance									
Sigma.psi[1,1]	0.03	0.02	0.04	0.03	0.02	0.04	0.03	0.02	0.04
Sigma.psi[1,2]	−0.01	−0.03	0.02	0.00	−0.01	0.01	0.00	−0.01	0.01
Sigma.psi[2,2]	0.16	0.12	0.20	0.07	0.04	0.12	0.06	0.04	0.09
rho.psi	−0.07	−0.35	0.24	−0.02	−0.30	0.25	−0.03	−0.28	0.25
Dynamic CAR									
Autoregressive parameter									
rho	0.64	0.45	0.78	0.62	0.28	0.88	0.59	0.25	0.84
Overdispersion parameter									
alpha	–	–	–	0.01	0.01	0.02	2.75	2.30	3.41

CAR: conditional autoregressive model.
[a]Reference: Calendar Year 2007.
[b]Reference: Non-Hispanic White.
[c]Reference: Other Private Insurance.

(IDR) comparing the mean number of visits over five years for self-pay patients to those with non-Duke private insurance. The IDR in this case is defined as E_1/E_2, where $E_1 = \pi\mu/[1 - p(0; \mu, \alpha)]$ is the expected count for the self-pay group and E_2 is defined similarly for the private group. The expected counts, E_1 and E_2, are in turn derived using predictive margins,[33] in which all patients are assigned first to self-insurance to compute E_1 and subsequently to private insurance to compute E_2. Applying this approach to the generalized Poisson model, the IDR is estimated to be 1.69 (1.62, 1.78), indicating that self-pay patients averaged nearly 1.7 times more ED visits during the study period than those on non-Duke private insurance. Similar calculations can be made for other comparisons of interest.

Figure 3 displays the age trends on the linear-predictor scale for the two components of the generalized Poisson hurdle model. The figure shows a multimodal effect for age, with ED use decreasing during the first decade of life, increasing steadily until around age 30, and then declining until age 75 before a final upswing to its peak level late in life. A similar pattern is observed for the number of visits among users, but here the mode occurs during infancy rather than among the elderly.

Figure 4 presents yearly maps of the predicted spatiotemporal effects, $\boldsymbol{\eta}_1$ and $\boldsymbol{\eta}_2$, from the generalized Poisson hurdle model, where $\eta_{1ij} = \nu_{1j} + \phi_{1i} + \delta_{1ij}$ and $\eta_{2ij} = \nu_{2j} + \phi_{2i} + \delta_{2ij}$ $(i = 1, \ldots, 129;\ j = 1, \ldots, 5)$. The highest ED activity occurred among the central block groups and the lowest among the block groups in the southwest corner of the county. Across years, the most significant change took place between 2007 and 2008, with several central block groups transitioning into the highest ED category (represented by the darkest shade), and the southwestern block groups transitioning into the lowest category (represented by lightest shade). The spatial pattern stabilized following 2008 with only minor fluctuations in select block groups.

The shift in spatial pattern from 2007 to 2008 is even more evident in Figure 5, which presents caterpillar plots of the spatiotemporal effects for years 2007–2011 ranked in decreasing order on the log-odds scale. Positive (negative) effects indicate above (below) average ED activity, adjusting for other factors. Error bars denote 95% CrIs; intervals that exclude zero indicate significant effects. Overall, there was a noticeable change in the magnitude of the effects following 2007, although the

Downloaded from smm.sagepub.com at EMORY UNIV on April 29, 2014

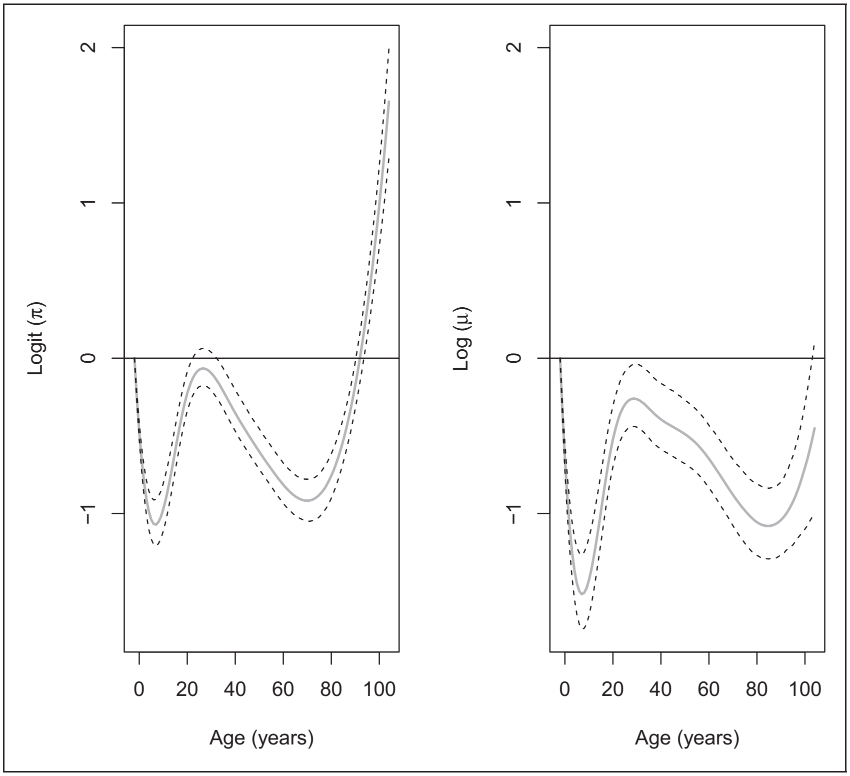

Figure 3. Age effect on linear-predictor scale for the binary component (left panel) and count component (right panel) for the generalized Poisson hurdle model. Horizontal lines denote effect at age = 0. Dashed lines denote 95% posterior intervals.

direction of the effects generally remained the same. The change in magnitude was not uniform across block groups, however, with several block groups showing large increases and others showing little or no change. This highlights the importance of modeling the space-time interactions in order to capture block-group–specific temporal trends.

Consistent with Figure 4, the "high-ED" block groups (e.g., those indexed 1–20) were generally situated in the center of the county, whereas the "low-ED" block groups (110 and above) tended to be located on the outskirts of the county, particularly along the southwest border. The patterns for 2009–2011 are generally similar to the 2008 plot, suggesting that the main change in ED visits occurred between the end of 2007 and the end of 2008, remaining relatively stable thereafter. This is again consistent with the results from Figure 4.

As a final check of model fit, we compared histograms of the observed and posterior-predictive counts generated from the generalized Poisson hurdle model (Figure 6). The generalized Poisson model showed excellent fit, reproducing almost exactly the observed distribution of counts.

6 Discussion

We have introduced a series of two-part hurdle models for the spatiotemporal analysis of zero-inflated count data. The proposed models have several attractive features: they provide spatial and

Downloaded from smm.sagepub.com at EMORY UNIV on April 29, 2014

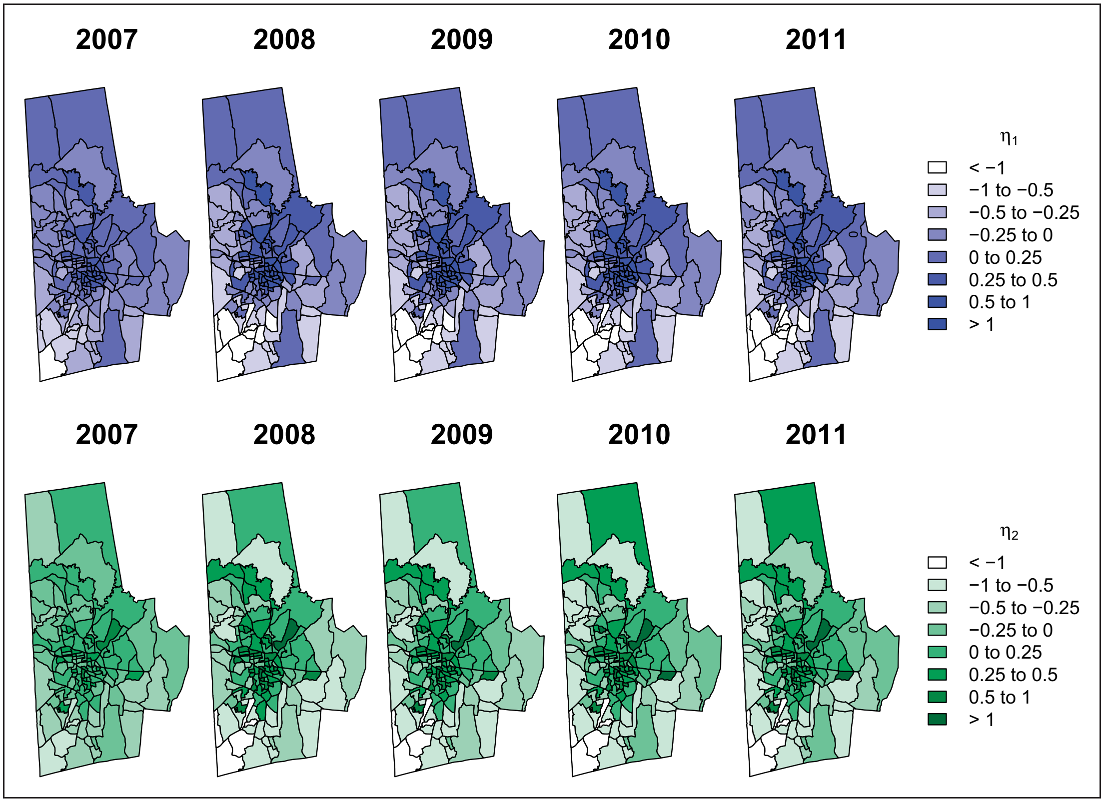

Figure 4. Spatiotemporal random effects for the binary (η_1) and count (η_2) components of the generalized Poisson hurdle model.

temporal smoothing to improve small-area estimation; they use multivariate space-time random effects to link model components; they incorporate individual- and region-level information to help explain spatiotemporal trends; and, depending on the choice of base distribution, they address potential over- or underdispersion in the counts. In addition, the models can be conveniently implemented in freely available packages such as WinBUGS.

In our application, models accommodating overdispersion, and in particular the generalized Poisson hurdle model, substantially outperformed the ordinary Poisson hurdle model. Because the negative binomial and generalized Poisson base distributions include the Poisson as either a limiting distribution (in the case of the negative binomial) or as a specific sub-model (in the case of the generalized Poisson), their performance should be comparable to the Poisson for equidispersed data while providing a distinct advantage for overdispersed data. Indeed, since both distributions arise as mixtures of ordinary Poissons,[23] they reduce to the Poisson in the case of a degenerate mixture. In the overdispersed (or nondegenerate) setting, the choice between the negative binomial and generalized Poisson base distributions will depend on the structure of the data, with the generalized Poisson typically providing better fit for highly skewed data.[23]

Our analysis of the DSR data yielded several relevant public health findings. Nonprivate insurance, male gender, and non-Hispanic Black race were associated with increased ED use.

Downloaded from smm.sagepub.com at EMORY UNIV on April 29, 2014

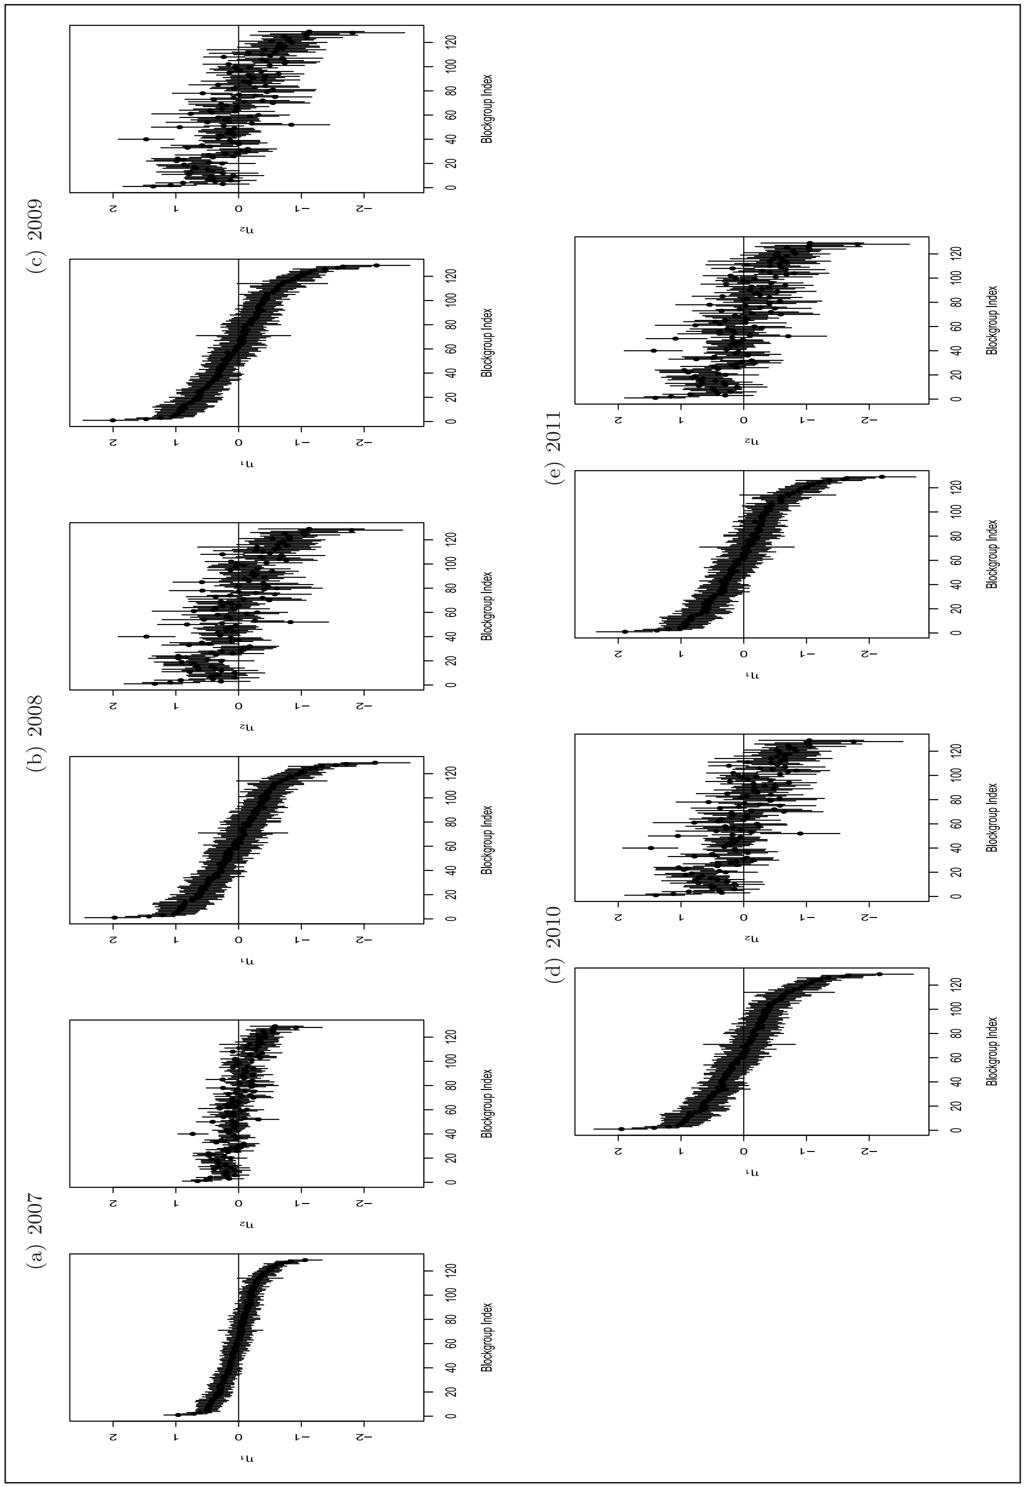

Figure 5. Spatiotemporal random effects for the binary (η_1) and count (η_2) components of the generalized Poisson hurdle model, ranked in order of decreasing log-odds. Error bars denote 95% credible intervals.

Downloaded from smm.sagepub.com at EMORY UNIV on April 29, 2014

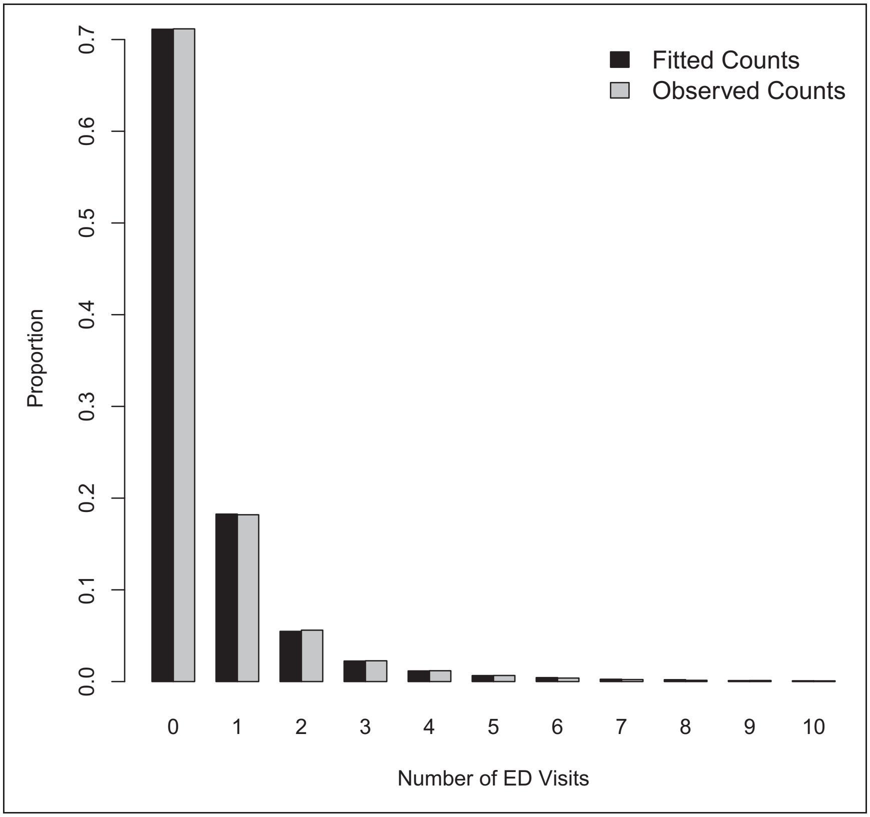

Figure 6. Partial histograms of observed and posterior-predicted ("fitted") counts from the generalized Poisson hurdle model.

Compared to non-Hispanic Whites, Hispanics were more likely to use the ED at least once but less inclined to make repeat visits. We also found a multimodal effect for age, with peak ED use occurring just after birth, around age 30, and after age 75. This trend may be due in part to increased rates of fever, cough, and upper respiratory infections among infants; injuries and spinal disorders among young adults; and falls, strokes, and cardiac events among the elderly.[30,34] In all years, block groups in the center of county had the highest rates of ED use while those in the southwest had the lowest. These geographic differences may reflect lack of community resources, such as outpatient clinics, in some areas. And finally, the spatial pattern across block groups changed most noticeably between 2007 and 2008 before stabilizing in the later years. The increased ED use observed in central Durham following 2007 may be a result of the recent economic downturn beginning in late 2007.[35,36]

These findings have important policy implications for the management of ED-related care. By monitoring spatial patterns in ED use over time, policy makers can target communities with continued need for services such as mobile clinics and community health centers to reduce nonurgent ED visits. To maximize benefit, these facilities should accommodate a variety of schedules through flexible evening and weekend hours in order to serve patients who require routine care during nonbusiness hours.[37] Local officials can also establish urgent care facilities

Downloaded from smm.sagepub.com at EMORY UNIV on April 29, 2014

that offer alternative outlets for primary medical, dental, and mental health care.[38] Community-based efforts such as these are essential to alleviating ED burden and reducing patient wait times.

As part of ongoing research, we are currently conducting an analysis of trends in ED-related costs by developing a related two-part spatiotemporal model for semicontinuous data. For more on semicontinuous models, see Olsen and Schafer[39] and Neelon et al.[40] Future work might also incorporate individual-level random effects into one or both components of the hurdle model to explicitly model between-subject heterogeneity. This would lead to a three-level model providing an additional, though computationally more intensive, approach to modeling overdispersion in the ED counts.

In general, the models developed here, particularly those accommodating overdispersion, should prove useful for the spatiotemporal analysis of zero-inflated count data. The proposed Bayesian approach provides a practical framework for fitting such models.

Ethical approval

This work was conducted in accordance with a human subjects research protocol approved by Duke University's Institutional Review Board.

Funding

The work described in this manuscript is supported in part by Grant Number 1C1CMS331018-01-00 from the Department of Health and Human Services, Centers for Medicare & Medicaid Services, and in part by the Bristol Myers Squibb Foundation Together on Diabetes program, respectively. The contents of this manuscript are solely the responsibility of the authors and have not been approved by the Department of Health and Human Services, Centers for Medicare & Medicaid Services. This work was also supported in part by a methodology grant from the Duke University Department of Biostatistics and Bioinformatics, and in part by USEPA grant R834799. Its contents are solely the responsibility of the grantee and do not necessarily represent the official views of the USEPA. Further, USEPA does not endorse the purchase of any commercial products or services mentioned in the publication. The views expressed in this article are those of the authors and do not necessarily reflect the position or policy of the Department of Veterans Affairs or the United States government.

References

1. Hing E and Bhuiya F. *Wait time for treatment in hospital emergency departments: 2009*. Hyattsville, MD: National Center for Health Statistics, 2012, p.102. http://www.cdc.gov/nchs/data/databriefs/db102.pdf (accessed 1 March 2014).
2. Weinick RM, Burns RM and Mehrotra A. Many emergency department visits could be managed at urgent care centers and retail clinics. *Health Aff* 2010; **29**: 1630–1636.
3. Cunningham PJ. What accounts for differences in the use of hospital emergency departments across U.S. communities? *Health Aff* 2006; **25**: 324–336.
4. McWilliams A, Tapp H, Barker J, et al. Cost analysis of the use of emergency departments for primary care services in Charlotte, North Carolina. *N C Med J* 2011; **72**: 265–271.
5. Horney C, Schmader K, Sanders LL, et al. Health care utilization before and after an outpatient ED visit in older people. *Am J Emerg Med* 2012; **30**: 135–142. http://www.sciencedirect.com/science/article/pii/S0735675710005516 (accessed 1 March 2014).
6. Doran KM, Colucci AC, Hessler RA, et al. An intervention connecting low-acuity emergency department patients with primary care: effect on future primary care linkage. *Ann Emerg Med* 2013; **61**: 312–321. e7, http://www.sciencedirect.com/science/article/pii/S0196064412016447 (accessed 1 March 2014).
7. Neelon B, Ghosh P and Loebs PF. A spatial Poisson hurdle model for exploring geographic variation in emergency department visits. *J R Stat Soc: Ser A* 2013; **176**: 389–413.
8. Agarwal DK, Gelfand AE and Citron-Pousty S. Zero-inflated models with application to spatial count data. *Environ Ecol Stat* 2002; **9**: 341–355.
9. Rathbun S and Fei S. A spatial zero-inflated Poisson regression model for oak regeneration. *Environ Ecol Stat* 2006; **13**: 409–426.
10. Gschlößl S and Czado C. Modelling count data with overdispersion and spatial effects. *Stat Pap* 2008; **49**: 531–552. http://dx.doi.org/10.1007/s00362-006-0031-6 (accessed 1 March 2014).

Downloaded from smm.sagepub.com at EMORY UNIV on April 29, 2014

11. Recta V, Haran M and Rosenberger JL. A two-stage model for incidence and prevalence in point-level spatial count data. *Environmetrics* 2012; **23**: 162–174. http://dx.doi.org/10.1002/env.1129 (accessed 1 March 2014).
12. Fuentes M, Song HR, Ghosh SK, et al. Spatial association between speciated fine particles and mortality. *Biometrics* 2006; **62**: 855–863. http://dx.doi.org/10.1111/j.1541-0420.2006.00526.x (accessed 1 March 2014).
13. Ver Hoef J and Jansen J. Space-time zero-inflated count models of harbor seals. *Environmetrics* 2007; **18**: 697–712.
14. Lunn DJ, Thomas A, Best N, et al. WinBUGS – a Bayesian modelling framework: concepts, structure, and extensibility. *Stat Comput* 2000; **10**: 325–337.
15. Horvath MM, Winfield S, Evans S, et al. The DEDUCE Guided Query tool: providing simplified access to clinical data for research and quality improvement. *J Biomed Inform* 2011; **44**: 266–276. http://www.sciencedirect.com/science/article/pii/S1532046410001747 (accessed 1 March 2014).
16. US Census Bureau. *American Community Survey 2005–2009*. Washington, DC: US Census Bureau, 2010. http://www.census.gov/acs/www/ (accessed 1 March 2014).
17. DeNavas-Walt C, Proctor BD and Smith JC. *Income, poverty, and health insurance coverage in the United States: 2011*. Washington, DC: U.S. Census Bureau, 2012, pp.60–243. http://www.census.gov/prod/2012pubs/p60-243.pdf (accessed 1 March 2014).
18. Vuong QH. Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica* 1989; **57**: 307–333. http://www.jstor.org/stable/1912557 (7 March 2014).
19. Mullahy J. Specification and testing of some modified count data models. *J Econom* 1986; **33**: 341–365.
20. Heilbron DC. Zero-altered and other regression models for count data with added zeros. *Biometrical J* 1994; **36**: 531–547. http://dx.doi.org/10.1002/bimj.4710360505 (accessed 1 March 2014).
21. Lambert D. Zero-inflated Poisson regression, with an application to defects in manufacturing. *Technometrics* 1992; **34**: 1–14.
22. Consul PC and Jain GC. A generalization of the Poisson distribution. *Technometrics* 1973; **15**: 791–799.
23. Joe H and Zhu R. Generalized Poisson distribution: the property of mixture of Poisson and comparison with negative binomial distribution. *Biometrical J* 2005; **47**: 219–229. http://dx.doi.org/10.1002/bimj.200410102 (accessed 1 March 2014).
24. Mardia KV. Multi-dimensional multivariate Gaussian Markov random fields with application to image processing. *J Multivariate Anal* 1988; **24**: 265–284.
25. Gelfand AE, Banerjee S and Gamerman D. Spatial process modelling for univariate and multivariate dynamic spatial data. *Environmetrics* 2005; **16**: 465–479. http://dx.doi.org/10.1002/env.715 (accessed 1 March 2014).
26. Spiegelhalter D, Thomas A, Best N, et al. *WinBUGS user manual, Version 1.4.3, 2007*. Cambridge, UK: MRC Biostatistics Unit, Institute of Public Health, http://www.mrc-bsu.cam.ac.uk/bugs/winbugs/contents.shtml (accessed 1 March 2014).
27. Geweke J. Evaluating the accuracy of sampling-based approaches to calculating posterior moments. In: Bernardo JM, Berger JO, Dawid AP and Smith AFM (eds) *Bayesian statistics 4*. Oxford: Clarendon Press, 1992, pp.169–193.
28. Spiegelhalter DJ, Best NG, Carlin BP, et al. Bayesian measures of model complexity and fit. *J R Stat Soc: Ser B (Stat Methodol)* 2002; **64**: 583–639. http://dx.doi.org/10.1111/1467-9868.00353 (accessed 1 March 2014).
29. Gelman A, Meng XL and Stern H. Posterior predictive assessment of model fitness via realized discrepancies. *Stat Sinica* 1996; **6**: 733–807.
30. Niska R, Bhuiya F and Xu J. *National hospital ambulatory medical care survey: 2007 emergency department summary*. Washington, DC: National Center for Health Statistics, 2010. http://www.cdc.gov/nchs/data/nhsr/nhsr026.pdf (accessed 1 March 2014).
31. R Core Team. *R: a language and environment for statistical computing*. Vienna, Austria: R Core Team, 2012. http://www.R-project.org/ (accessed 1 March 2014).
32. Sturtz S, Ligges U and Gelman A. R2WinBUGS: a package for running WinBUGS from R. *J Stat Softw* 2005; **12**: 1–16. http://www.jstatsoft.org (accessed 1 March 2014).
33. Graubard BI and Korn EL. Predictive margins with survey data. *Biometrics* 1999; **55**: 652–659. http://dx.doi.org/10.1111/j.0006-341X.1999.00652.x (accessed 1 March 2014).
34. Centers for Disease Control and Prevention. *National hospital ambulatory medical care survey: 2010 emergency department summary tables*. Atlanta, GA: Centers for Disease Control and Prevention, 2010http://www.cdc.gov/nchs/data/ahcd/nhamcs_emergency/2010_ed_web_tables.pdf (accessed 1 March 2014).
35. Paradise J and Dark C. *Emergency departments under growing pressures*. Menlo Park, CA: The Henry J. Kaiser Family Foundation, 2009, p.7960. http://kaiserfamilyfoundation.files.wordpress.com/2013/01/7960.pdf (accessed 1 March 2014).
36. Mortensen K and Chen J. The great recession and racial and ethnic disparities in health services use. *JAMA Intern Med* 2013; **173**: 315–317.
37. US Government Accountability Office. *Hospital emergency departments: health center strategies that may help reduce their use*. No. 11-414R. Washington, DC: U.S. Government Accountability Office, 2011. http://www.gao.gov/assets/100/97416.pdf (accessed 1 March 2014).
38. Roby DH, Pourat N, Pirritano MJ, et al. Impact of patient-centered medical home assignment on emergency room visits among uninsured patients in a county health system. *Med Care Res Rev* 2010; **67**(4): 412–430.
39. Olsen MK and Schafer JL. A two-part random-effects model for semicontinuous longitudinal data. *J Am Stat Assoc* 2001; **96**: 730–745.
40. Neelon B, O'Malley AJ and Normand SLT. A Bayesian two-part latent class model for longitudinal medical expenditure data: assessing the impact of mental health and substance abuse parity. *Biometrics* 2011; **67**: 280–289.

Downloaded from smm.sagepub.com at EMORY UNIV on April 29, 2014