ROYAL STATISTICAL SOCIETY
DATA | EVIDENCE | DECISIONS

Journal of the Royal Statistical Society
Statistics in Society
Series A

J. R. Statist. Soc. A (2018)
181, *Part* 3, *pp.* 803–823

Autologistic models for benchmark risk or vulnerability assessment of urban terrorism outcomes

Jingyu Liu and Walter W. Piegorsch,

University of Arizona, Tucson, USA

A. Grant Schissler

University of Nevada, Reno, USA

and Susan L. Cutter

University of South Carolina, Columbia, USA

[Received November 2016. Final revision August 2017]

Summary. We develop a quantitative methodology to characterize vulnerability among 132 US urban centres ('cities') to terrorist events, applying a place-based vulnerability index to a database of terrorist incidents and related human casualties. A centred autologistic regression model is employed to relate urban vulnerability to terrorist outcomes and also to adjust for auto-correlation in the geospatial data. Risk analytic 'benchmark' techniques are then incorporated in the modelling framework, wherein levels of high and low urban vulnerability to terrorism are identified. This new translational adaptation of the risk benchmark approach, including its ability to account for geospatial auto-correlation, is seen to operate quite flexibly in this sociogeographic setting.

Keywords: Benchmark dose; Centred autologistic model; Geospatial analysis; Maximum pseudolikelihood; Quantitative risk analysis; Spatial auto-correlation

1. Introduction: urban vulnerability assessment

Quantifying vulnerability to natural or human-induced hazardous events is a critical component in modern risk or hazard analysis. For example, multiple factors can affect a city's or other urban locale's response to a hazardous impact, and it is possible to quantify and/or index the city's vulnerabilities that are associated with that response. Work in this area often characterizes those populations, localities, ecosystems and other entities that are at risk to hazardous events, and such efforts have seen substantive development (Peng *et al.*, 2005; Mazzorana *et al.*, 2009; Santini *et al.*, 2010; MacKenzie, 2014). Many localities exhibit vulnerabilities due to their physical and geographic characteristics; others contain populations that are socially vulnerable, whereas still others have built-environment factors that make them vulnerable to certain hazards. In some locations, a combination of vulnerabilities affects a community's disaster resilience and response. Indeed, urban hazards are by their very nature multi-dimensional processes, with many factors underlying the occurrence and incidence of each adverse event. Although it is difficult to distil

Address for correspondence: Walter W. Piegorsch, Department of Mathematics, University of Arizona, Tucson, AZ 85721, USA.
E-mail: piegorsch@math.arizona.edu

 0964–1998/18/181803

such processes down to a few explanatory components, efforts to do so can provide important guidance to many interested parties—municipal planners, emergency managers, social services providers, disaster relief organizations, etc.—when assessing their locality's level of vulnerability to natural or human-induced hazards.

Towards this end, we previously exploited a series of vulnerability indices (Cutter *et al.*, 2003; Borden *et al.*, 2007) to benchmark urban susceptibility to adverse events, focusing on terrorism outcomes (Piegorsch *et al.*, 2007). Data from that study consisted of binary incidence indicators Y_i, indexed over 132 US metropolitan areas ('cities') that were deemed most susceptible to environmental and human-induced hazards by the US Federal Government ($i = 1, \ldots, 132$). Of interest with these data was the probability π_i that the ith urban centre would have experienced casualties due to terrorist incidents during the 35-year period (1970–2004) over which the terrorism outcomes were collected. Our goal was to study how place-based factors of urban vulnerability related to terrorist attacks, and to use statistical features of those relationships to indicate and characterize the associated levels of urban vulnerability. In particular, we translated a technique from toxicological risk assessment—called benchmark risk analysis (Crump, 2012)—to characterize an urban area's susceptibility to adverse events. Broadly speaking, the benchmark approach models π_i as a function of some target predictor variable x_i, quantifying the hazardous input, and builds a risk function $R(x)$ that models how the mean risk or vulnerability of the study subjects—here, the cities—changes over x. It then asks at which value of x does $R(x)$ achieve a prespecified *benchmark response* BMR? The solution is generically called the *benchmark dose*, or in our case the *benchmark index* BMI above which cities are at increased vulnerability to the hazard. The particular mathematical expression for BMI will depend on the statistical model that is chosen to represent π and R as functions of x.

Given the particular model, a point estimate $\widehat{\text{BMI}}_{\text{BMR}}$ for BMI, and corresponding $1-\alpha$ confidence bounds, can be calculated from the observed data pairs (x_i, Y_i) (see below). In most settings, the adverse outcome is detrimental, and thus only a lower confidence limit is required; this is denoted by BMIL_{BMR} (Crump, 1995). In both cases, where needed for clarity the subscript explicitly reminds us of the dependence on the BMR input value.

We should note that the terms 'risk' and 'vulnerability' carry very specific meanings in the literature, although at times different writers utilize them in different ways. Our use of vulnerability here is similar to that defined by Haimes (2006), the manifestation of a location's inherent condition(s) that can be exploited to adverse effect. We use the term risk in a more colloquial sense, however, suggesting some general source or instance of threat, danger or peril, rather than any technical definition or notion.

In our earlier work (Piegorsch *et al.*, 2007), we manipulated the benchmark paradigm to identify urban BMIs at varying levels of BMR, corresponding to increasing urban vulnerability to terrorism casualties based on our 132-cities data. The result provided a quantification of vulnerability in these various urban locations, based on their own particular socio-economic, physical and built-environment factors. We did so by constructing a single place-based vulnerability index PVI as a weighted average of these individual components; PVI served as our x-variable (see Piegorsch *et al.* (2007) for full details). The consequent estimates of BMI provided benchmark points on this PVI-scale at which urban vulnerability to terrorist casualties grew large, relative to lower values on the same scale.

A critical feature missing in Piegorsch *et al.* (2007), however, was recognition that one city's vulnerability to terrorism could be affected by its proximity to another city, i.e. potential spatial auto-correlation between the Y_is can affect calculation of $\widehat{\text{BMI}}_{\text{BMR}}$. Indeed, to our knowledge no approach currently exists that allows for calculation of benchmark doses or indices, etc. when potential auto-correlation exists in the data. Here, we propose an extension of the benchmark

dose approach for use with spatially auto-correlated observations and apply it to our 132-cities data. To do so, we expand our linear predictor to include a spatial autocovariate alongside the target benchmark predictor, which is a tactic that is often seen to assuage concerns with the effects of geospatial auto-correlation (Koutsias, 2003; de Frutos *et al.*, 2007; Vaughn *et al.*, 2015). Our goal is to develop a valid risk analytic technology that can relate urban vulnerability to terrorist activity while taking into account possible spatial auto-correlation after adjusting for place-based factors of the urban vulnerability. We then manipulate statistical features of the relationship to indicate and characterize the different levels of urban vulnerability to terrorist events in our data set.

In Section 2, we reintroduce the 132-cities data from Piegorsch *et al.* (2007), including an updated version of PVI that is more amenable to benchmark calculations with auto-correlated observations. Section 3 then describes a series of autologistic models that we employ to analyse the data, along with details on how the risk benchmark calculations can be incorporated in them. Section 4 follows with a short Monte Carlo simulation to study operating characteristics of the consequent BMI- and BMIL-estimates. Section 5 applies the new methods to the 132-cities data, whereas Section 6 ends with a brief discussion. All calculations that we present below are performed in the R programming environment (R Core Team, 2016).

Sample R code to fit the centred autologistic model can be obtained from

`http://wileyonlinelibrary.com/journal/rss-datasets`

2. Example: urban vulnerability data

To build a foundation for use of our auto-correlated risk-benchmark methodology, we employ data on urban vulnerability to terrorism events from our earlier work (Borden *et al.*, 2007; Piegorsch *et al.*, 2007). Therein, we examined 132 US metropolitan locations ('cities') that were deemed at greatest risk to terrorist hazards; see the display in Table 1. The cities are also listed in Piegorsch *et al.* (2007), where full details of the larger data set can be found. Summarizing from that source, the binary outcomes Y_i indicate whether or not the ith city experienced a terrorist-related event during the 35-year period 1970–2004, in which one or more human casualties (injuries or deaths) were recorded. These values also appear in Table 1. Note that, of the 132 cities in Table 1, $\Sigma Y_i = 36$, or 27.3%, reported a casualty event.

Associated with these adverse outcomes, each ith city is assigned a PVI that acts as our target predictor variable x_i. This PVI is a shifted version of the original place-based index built from separate socio-economic, physical and built-environment measures in Piegorsch *et al.* (2007): our construction requires the target predictor to be non-negative whereas the vulnerability index of Piegorsch *et al.* (2007) was manipulated to have roughly zero mean. We discuss our shifted PVI in more detail below.

Of interest is the probability π that a city experiences terrorist-related casualties during the time span (1970–2004) under study. Our goal is to relate π to the city's broader place-based vulnerability, as quantified by PVI.

The responses in Table 1 are dichotomous and can be viewed as binomial: $Y_i \sim \text{Bin}(1, \pi_i)$. To model the relationship between π and PVI, Piegorsch *et al.* (2007) employed a binary generalized linear model (GLM):

$$\pi_i = P(Y_i = 1) = g^{-1}(\eta_i),$$

where the linear predictor η_i was taken as the simple linear form $\beta_0 + \beta_1 x_i$ and the single covariate x_i was PVI. The β-values were unknown regression parameters, and $g(\cdot)$ was the link function

Table 1. Terrorism casualty indicator *Y* and place-based vulnerability measure PVI for 132 US metropolitan areas

Metropolitan area	*Y*	*PVI*	*Metropolitan area*	*Y*	*PVI*
Albany, New York	1	4.199	Huntsville, Alabama	0	3.932
Albuquerque, New Mexico	0	2.942	Indianapolis, Indiana	0	3.340
Allentown–Bethlehem, Pennsylvania	0	3.933	Jackson, Missouri	0	4.279
Amarillo, Texas	0	3.419	Jacksonville, Florida	0	3.810
Anchorage, Alaska	0	2.757	Jefferson City, Missouri	0	3.029
Atlanta, Georgia	1	5.402	Juneau, Alaska	0	1.298
Augusta, Georgia	0	4.450	Kansas City, Missouri	0	3.979
Augusta, Maine	0	2.909	Knoxville, Tennessee	0	3.690
Austin, Texas	0	3.611	Lancaster, Pennsylvania	0	3.604
Bakersfield, California	0	2.880	Lansing, Mississippi	0	3.086
Baltimore–Annapolis, Maryland	0	4.942	Las Vegas, Nevada	0	3.329
Barre–Montpelier, Vermont	0	2.308	Lexington, Kentucky	0	3.340
Baton Rouge, Louisiana	0	6.735	Lincoln, Nebraska	0	3.174
Birmingham, Alabama	1	3.905	Little Rock, Arkansas	0	3.701
Bismarck, North Dakota	0	3.101	Los Angeles, California†	1	4.140
Boise, Idaho	0	5.415	Louisville, Kentucky	0	3.841
Boston, Massachusetts	1	4.323	Lubbock, Texas	0	3.110
Bridgeport–Stamford, Connecticut	1	2.700	Madison, Wisconsin	1	3.303
Buffalo, New York	1	3.998	McAllen, Texas	0	4.396
Cape Coral, Florida	0	4.262	Memphis, Tennessee	0	4.632
Carson City, Nevada	0	2.580	Miami–Fort Lauderdale, Florida	1	4.314
Charleston, South Carolina	0	6.262	Milwaukee, Wisconsin	0	4.080
Charleston, West Virginia	0	3.557	Minneapolis–St Paul, Minnesota	0	4.059
Charlotte, North Carolina	0	4.765	Mission Viejo, California	1	3.873
Chattanooga, Tennessee	0	4.183	Mobile, Alabama	0	4.144
Cheyenne, Wyoming	0	2.691	Modesto, California	0	2.911
Chicago, Illinois	1	5.123	Montgomery, Alabama	0	3.824
Cincinnati, Ohio	0	3.941	Nashville, Tennessee	0	3.368
Cleveland–Akron, Ohio	1	4.911	New Haven, Connecticut	1	3.821
Colorado Springs, Colorado	0	2.457	New Orleans, Louisiana	1	6.838
Columbia, South Carolina	0	4.856	New York–Newark, New York	1	5.873
Columbus, Georgia	1	3.614	Norfolk, Virginia‡	0	6.045
Columbus, Ohio	0	3.684	Oklahoma City, Oklahoma	1	3.768
Concord, New Hampshire	0	2.235	Olympia, Washington	0	2.638
Corpus Christi, Texas	0	3.949	Omaha, Nebraska	1	4.050
Dallas–Fort Worth, Texas	1	4.166	Orlando, Florida	0	4.466
Dayton, Ohio	0	3.498	Oxnard, California	0	2.512
Denver, Colorado	1	3.890	Palm Bay–Melbourne, Florida	0	3.764
Des Moines, Iowa	0	3.877	Pensacola, Florida	1	4.337
Detroit–Warren, Michigan	1	3.907	Philadelphia, Pennsylvania	1	5.456
Dover, Delaware	0	3.664	Phoenix–Glendale–Mesa, Arizona	1	2.815
El Paso, Texas	0	3.908	Pierre, South Dakota	0	2.211
Flint, Michigan	0	3.373	Pittsburgh, Pennsylvania	1	4.386
Fort Wayne, Indiana	0	2.780	Portland, Oregon	0	3.730
Frankfort, Kentucky	0	2.739	Poughkeepsie–Newburgh, New York	0	3.921
Fresno, California	0	3.322	Providence, Rhode Island	0	3.439
Grand Rapids, Michigan	0	3.083	Raleigh–Durham, North Carolina	0	4.736
Greensboro–Winston Salem, North Carolina	1	4.533	Reno, Nevada	0	2.752
			Richmond, Virginia	0	5.655
Harrisburg, Pennsylvania	0	3.937	Riverside–San Bernardino, California	0	3.544
Hartford, Connecticut	0	3.378	Rochester, New York	0	3.456
Helena, Montana	0	1.990	Sacramento, California	1	3.465
Honolulu, Hawaii	0	3.482	Salem, Oregon	0	2.705
Houston, Texas	0	5.563	Salt Lake City–Ogden, Utah	1	3.910

(*continued*)

Table 1 (*continued*)

Metropolitan area	*Y*	*PVI*	*Metropolitan area*	*Y*	*PVI*
San Antonio, Texas	0	4.634	Syracuse, New York	0	4.130
San Diego, California	1	3.200	Tallahassee, Florida	0	3.722
San Francisco, California§	1	3.769	Tampa–St Petersburg, Florida	1	4.869
Santa Fe, New Mexico	0	2.506	Toledo, Ohio	0	4.436
Sarasota–Bradenton, Florida	0	4.345	Topeka, Kansas	1	3.003
Scranton, Pennsylvania	0	4.157	Trenton, New Jersey	1	4.129
Seattle–Tacoma, Washington	0	3.404	Tucson, Arizona	0	2.951
Shreveport, Louisiana	0	4.249	Tulsa, Oklahoma	0	3.695
Spokane, Wyoming	0	2.939	Washington DC	1	5.697
Springfield, Illinois	0	3.012	Wichita, Kansas	1	3.089
Springfield, Massachusetts	0	3.279	Worcester, Massachusetts	0	2.867
St Louis, Missouri	1	4.663	Youngstown, Ohio	0	4.143
Stockton, California	0	3.307			

†Los Angeles region also includes Glendale, Long Beach and Huntington Beach, California.
‡Norfolk region also includes Chesapeake, Newport News and Virginia Beach, Virginia.
§San Francisco region also includes Oakland, San Jose and Fremont, California.

of the GLM. They then translated this model into a form that was amenable to the sort of benchmark risk analysis that is seen in toxicological risk assessment (Crump, 1984): after fitting the GLM to the data, point estimates of the unknown β-parameters were manipulated into estimates of BMI for some fixed BMR. Associated lower confidence limits BMIL were then constructed under the model, to serve as benchmark points marking various levels of urban vulnerability to terrorist casualties. As noted above, however, missing from this analysis was the recognition that responses between cities that were close to each other could exhibit spatial auto-correlation. How to adjust the benchmark analysis to account for such auto-correlation is the focus of the methods that are presented in the next section.

3. Autologistic modelling

Motivated by the vulnerability data in Table 1, we consider how to model a set of binary observations as a function of a target regressor variable while adjusting for extant spatial auto-correlation. We assume that the relationship between the variables may be affected by the observations' proximities to each other (Chun and Griffith (2013), chapter 2) leading to the development of local area spatial statistics. We appeal to the well-known logistic regression model, which is widely applied with binary data (Hosmer *et al.*, 2013). A simple and direct way to account for spatial auto-correlation within a logistic model allows the response probability π_i at the ith spatial location—here, the ith city—to depend on the other observed binary responses in some predefined neighbourhood N_i around that city. Besag (1972) first proposed such a model to incorporate spatial auto-correlation: the conditional autologistic model

$$\pi_i = P(Y_i = 1 | x_i, y_j, j \neq i) = \frac{1}{1 + \exp\left\{ -\beta_0 - \beta_1 x_i - \beta_2 \sum_{j \in N_i} a_{ij} y_j \right\}}. \quad (1)$$

In our notation, Y_i is the outcome indicator for casualties, x_i is PVI, the a_{ij}s stipulate the ith neighbourhood N_i (see below) and $\Sigma_{j \in N_i} a_{ij} y_j$ is a spatial autocovariate constructed around each Y_i. We refer to β_2 as the spatial auto-correlation parameter. When $\beta_2 = 0$, no spatial auto-

correlation exists in the data, and the autologistic model then reduces to a standard logistic GLM (which we call the 'independence model').

To define the neighbourhoods N_i for our 132 cities, we consider a simple adjacency characterization: set $a_{ij} = 1$ if the jth metropolitan area is physically adjacent to the ith metropolitan area, i.e. it shares any edge boundary—a 'rook' connection, as opposed to 'queen' edge-and-corner connections—as delineated by the border definitions that were used in Borden *et al.* (2007). Otherwise set $a_{ij} = 0$. The corresponding adjacency matrix $\mathbf{A} = \{a_{ij}\}$ defines a simple neighbourhood structure among the 132 cities in the data set. Fig. S1 and Table S1 in the on-line supplemental document specify this adjacency structure across our 132 cities; Fig. S1 also maps the casualty incidence.

3.1. Centred autologistic model

The autologistic model (1) is a popular choice for modelling spatially dependent binary data. Unfortunately, Caragea and Kaiser (2009) identified some irregularities with the traditional autologistic form. In model (1), the autocovariate modelling the dependence between Y_i and the other Y_js is non-negative for any i. Consequently, the odds that a response equals 1 under the traditional model relative to the odds that the response equals 1 under the independence model (the standard logistic GLM) can change only in one direction for any non-zero neighbours. A simple argument shows that this leads to an anomaly: assume for convenience that there is positive auto-correlation in the data under a simple adjacency neighbourhood structure. (Similar arguments will hold under negative auto-correlation.) Then the regression coefficient for the autocovariate should be positive. In fact, the autocovariate under model (1) is the number of non-zero neighbours and, if the majority of a response's neighbours are 0, the odds that the response equals 1 under that model relative to the independence model is expected to decrease. Yet, as the autocovariate under the traditional model is always non-negative, that odds ratio can only increase given a positive regression coefficient for the autocovariate. Any meaningful autocovariate should be negative in this situation (Hughes *et al.*, 2011).

To account for this anomaly, Caragea and Kaiser (2009) offered a correction that extends the model into a centred autologistic form, by redefining the spatial autocovariate (essentially, by centring it):

$$\pi_i = P(Y_i = 1 | x_i, y_j, j \neq i) = \frac{1}{1 + \exp\left\{ -\beta_0 - \beta_1 x_i - \beta_2 \sum_{j \in N_i} a_{ij}(y_j - \mu_j) \right\}}, \tag{2}$$

where the new quantity $\mu_j = E[Y_j | \beta_2 = 0] = \{1 + \exp(-\beta_0 - \beta_1 x_j)\}^{-1}$ is the expected value of Y_j under an independence model with no spatial auto-correlation. Centring makes the estimates more stable and interpretable, and as such the centred autologistic construction in equation (2) provides a direct and simple way to build spatial dependence into a binary regression GLM (Hughes *et al.* (2011) and Kolaczyk and Csárdi (2014), section 8.3). We therefore turn to equation (2) for constructing quantitative risk or vulnerability assessments with such data.

3.2. Maximum pseudolikelihood estimation

Allowing for spatial auto-correlation induces dependences between the observations, complicating standard likelihood analysis with our centred autologistic model. Instead, we estimate the unknown β-parameters in equation (2) via maximization of the pseudolikelihood function that was proposed by Besag (1975). Besag found that, by multiplying together the conditional probability distributions of the Y_is given their neighbours, the resulting pseudolikelihood

$$L_{\mathrm{P}}(\boldsymbol{\beta}) = \prod_{i=1}^{n} P(Y_i = 1 | x_i, y_j, j \neq i) = \prod_{i=1}^{n} \pi_i^{y_i} (1 - \pi_i)^{1 - y_i}$$

has many of the same features as the usual likelihood function, despite any existing dependences between the observations. In particular, maximum pseudolikelihood estimates of the β-parameters are consistent and asymptotically normal under typical regularity conditions (Arnold and Strauss, 1991).

3.3. Autologistic benchmark risk analysis

To our knowledge, no previous work has attempted to incorporate spatial correlation—or any other form of auto-correlation—into a benchmark risk analysis. As we show below, doing so with the autologistic model requires non-trivial manipulation of the standard benchmark approach. For the binary data setting that is considered here, we continue to relate a prespecified level of outcome response—BMR—to the response probability $\pi(x)$ viewed as a function of $x =$ PVI. The goal is to find the smallest positive x at which some background-adjusted function of $\pi(x)$ equals BMR. The solution is the benchmark index, past which a city's vulnerability to terrorist casualties is elevated to unacceptable levels for that BMR. The 'background' adjustment is included to account for extemporaneous or spontaneous factors out of the risk assessor's control; a prototypical example is correction for spontaneous (zero-exposure) tumour incidence when assessing the risk of a carcinogenic exposure. We discuss this adjustment further below.

An added complexity with the centred autologistic model (2) is the presence of the centred autocovariate $\Sigma_{j \in N_i} a_{ij}(y_j - \mu_j)$. The covariate is designed to improve precision in the model by accounting for the recognized source of spatial auto-correlation; however, it also introduces additional terms that prevent us from simply inverting the estimated regression model to find $\widehat{\mathrm{BMI}}_{\mathrm{BMR}}$. A similar challenge occurs when calculating benchmark doses in epidemiological studies, where along with the exposure or dose variable of interest, x, investigators often record secondary covariates that are felt to describe known sources of variation for the disease outcome under study. To overcome this, Budtz-Jørgensen *et al.* (2001) proposed a clever definition for the benchmark response: view BMR as a specified proportional increase over zero-level background in the odds of an adverse event, evaluated at the same constellation of secondary covariates. At least for logistic-type models, the functional dependence on secondary covariates then cancels out in the various ratio operations. Budtz-Jørgensen *et al.* (2001) provided no guidance on selection of BMR here, but we might imagine in our terrorism vulnerability application taking BMR = 10 or BMR = 25, say, since an increase of this magnitude in the odds could be a useful marker for practical use. In any case, we shall follow standard practice and avoid selecting a value for BMR that deviates drastically from the range of corresponding responses that are seen in the data.

Employed with the centred autologistic regression model (2), we therefore define BMI as the dose resulting in a prespecified increase (of BMR-multiples) in the odds for an abnormal response, i.e. solve for x in

$$\mathrm{BMR} = \frac{\pi(x|\boldsymbol{\beta})/\{1 - \pi(x|\boldsymbol{\beta})\}}{\pi(0|\boldsymbol{\beta})/\{1 - \pi(0|\boldsymbol{\beta})\}}, \tag{3}$$

where $\pi(\cdot|\boldsymbol{\beta})$ is the centred autologistic response probability based on model (2) and $\boldsymbol{\beta} = (\beta_0\ \beta_1\ \beta_2)^{\mathrm{T}}$ is the vector of unknown autologistic coefficients. Note for adverse outcomes that we expect that $\beta_1 > 0$, representing an underlying increase in the probability of an adverse event as x increases. This typically defines a proportional rise in odds; thus we operate with BMR > 1.

We must be careful, however, when employing relationship (3) within our spatial setting.

The target predictor x is not an isolated value: for any x, definitive or abstract, there is some corresponding location, say s, with some neighbourhood N_s underlying it. Write this value of x as x_s. The odds ratio (3) is not calculated across locations; rather, it is calculated with regard to the location s that is associated with the chosen x_s. The numerator odds of observing an adverse event for the arbitrary location s are $\pi(x_s|\boldsymbol{\beta})/\{1-\pi(x_s|\boldsymbol{\beta})\}$. The denominator odds against which the numerator is compared are calculated at $x_s = 0$, but as if they occur at the same spatial location s with the same neighbourhood N_s. Thus the odds ratio (3) simplifies to

$$\frac{\pi(x_s|\boldsymbol{\beta})/\{1-\pi(x_s|\boldsymbol{\beta})\}}{\pi(x_s=0|\boldsymbol{\beta})/\{1-\pi(x_s=0|\boldsymbol{\beta})\}} = \frac{\exp\left\{\beta_0+\beta_1 x_s+\beta_2\sum_{t\in N_s} a_t(y_t-\mu_t)\right\}}{\exp\left\{\beta_0+\beta_2\sum_{t\in N_s} a_t(y_t-\mu_t)\right\}} = \exp(\beta_1 x_s).$$

Setting this equal to BMR then yields the unique BMI at the given BMR:

$$\mathrm{BMI}_{\mathrm{BMR}} = \log(\mathrm{BMR})/\beta_1. \tag{4}$$

This is the vulnerability index resulting in a (BMR − 1)-fold increase in odds of an adverse event, relative to a zero value of the index.

Before continuing, we should clarify what exactly is meant by a zero-valued PVI. To give it a standardized flavour, but for no other compelling reason, the placed-based index that was originally presented by Piegorsch *et al.* (2007), say, PVI07, had roughly zero mean and unit variance. For use in odds ratio (3), however, setting $x = 0$ has a specific interpretation: a PVI of 0 is the baseline index against which a city's PVI is compared, i.e. a city whose PVI is 0 exhibits the smallest, 'spontaneous' level of vulnerability imaginable for the population of cities under consideration. A natural tactic then would take the original placed-based values in Piegorsch *et al.* (2007) and add $|\min\{\mathrm{PVI07}\}|$ to each. In the 132-city data set, $\min\{\mathrm{PVI07}\} = -2.4208$ belongs to Juneau, Alaska. We felt that this would be unwise, however, since it gives Juneau a PVI of exactly 0, and this has the effect of comparing all other cities in the study with Juneau's particular vulnerability status. Instead we shifted the original index of Piegorsch *et al.* (2007) by adding a suitably large positive constant to it; we chose $\mathrm{PVI} = \mathrm{PVI07} + |\Phi^{-1}(0.0001)|$, where $\Phi^{-1}(0.0001) = -3.719$ from the standard normal cumulative distribution function. In effect, we might say that this marks an abstract baseline for odds ratio comparisons in equation (3) as a similarly located 'city' whose vulnerability characteristics place it in the lower 0.01% of all urban centres of interest. Although this simple shift has its own arbitrary aspects, we found that it produced a reasonable pattern of spread and separation for the purposes of our analysis. Note that we do not attempt to assign any formal stochastic model to our PVI; we view the index simply as a fixed constant derived from each city's underlying socio-economic, physical and built-environment characteristics, as outlined by Piegorsch *et al.* (2007).

We explored other transformations for the original PVI07-index but found that none gave a more conveniently spaced and sufficiently separated pattern of positive values as these simple, shifted PVIs. It is the shifted values that appear in Table 1.

3.4. BMI-estimation and inference

Given data pairs (x_i, Y_i) as the PVI and binary outcome for each ith city ($i = 1, \ldots, n$), we fitted the centred autologistic regression model from equation (2) via maximum pseudolikelihood, as in Section 3.2. This produces maximum pseudolikelihood estimators $\hat{\boldsymbol{\beta}} = (\hat{\beta}_0\ \hat{\beta}_1\ \hat{\beta}_2)^{\mathrm{T}}$. Unfortunately, the calculations are not trivial and require computer iteration. We employ the R package `ngspatial`; see `https://CRAN.R-project.org/package=ngspatial`. Given the

maximum pseudolikelihood estimator $\hat{\beta}_1$ (and a predetermined level for BMR), from equation (4) the corresponding point estimate for BMI is then simply

$$\widehat{\text{BMI}}_{\text{BMR}} = \log(\text{BMR})/\hat{\beta}_1. \tag{5}$$

As discussed above, we further desire a lower $1-\alpha$ confidence limit on BMI, denoted as BMIL_{BMR}. The straightforward form for BMI in equation (4) enables a particularly simple construction, at least conceptually: if $b_{1\text{U}}$ is an upper $1-\alpha$ confidence limit on β_1 such that $P(\beta_1 < b_{1\text{U}}) = 1-\alpha$ at least approximately in large samples, then clearly

$$P\left\{\frac{\log(\text{BMR})}{b_{1\text{U}}} < \frac{\log(\text{BMR})}{\beta_1}\right\} = P\left\{\frac{\log(\text{BMR})}{b_{1\text{U}}} < \text{BMI}_{\text{BMR}}\right\} = 1-\alpha \tag{6}$$

if $\text{BMR} > 1$. This defines a $1-\alpha$ lower confidence limit on BMI; thus we take $\text{BMIL}_{\text{BMR}} = \log(\text{BMR})/b_{1\text{U}}$. We note in passing that equation (6) can be written as $P\{b_{1\text{U}}^{-1}\log(\text{BMR}) < \text{BMI}_{\text{BMR}}, \forall\ \text{BMR} > 1\} = 1-\alpha$, which provides a simultaneous lower confidence statement on BMI, in the sense of Piegorsch and West (2005).

Actually finding such a $b_{1\text{U}}$ confidence limit is a more difficult task, however, since it is inappropriate under the centred autologistic model to imitate standard practice and to estimate the standard error of $\hat{\beta}_1$ by inverting the maximum pseudolikelihood information matrix (Varin *et al.*, 2011). Instead, we follow a suggestion by Hughes (2014), who called for a computer-intensive, parallel, parametric bootstrap approach. The method returns a bootstrap distribution of $B > 0$ resampled β_1-values, based on the original data. An approximate upper $1-\alpha$ confidence limit can then be taken as the $(1-\alpha)B$-quantile from this bootstrap distribution, denoted as $b_{1\text{B}}$. We found that at least $B = 5000$ bootstrap replicates were required for the bootstrap procedure to stabilize and produce viable confidence limits, at least for sample sizes at or above $n = 100$ as seen with the 132-cities data.

From this, a bootstrap-based lower confidence limit for BMI becomes

$$\text{BMIL}_{\text{BMR}} = \log(\text{BMR})/b_{1\text{B}}. \tag{7}$$

4. Monte Carlo evaluations

The BMIL that was constructed in Section 3.4 relies on large sample approximating arguments to be valid. To investigate the performance of the confidence limits in practice, we conducted a series of Monte Carlo evaluations. We examined the coverage properties of BMIL, along with selected characteristics of the underlying $\widehat{\text{BMI}}$-estimator.

4.1. Simulation settings

We began by selecting a series of three different scenarios under which a set of PVIs would be constructed:

(a) $n = 100$ values individually sampled from a uniform distribution whose lower limit is the minimum of the PVIs in Table 1 and whose upper limit is the maximum of those values, and distributed these 100 values randomly and independently—i.e. the PVIs had no embedded spatial correlation—onto a 10×10 lattice,
(b) $n = 100$ values individually sampled from the empirical cumulative distribution function of the actual PVI-values in Table 1 again distributed randomly and independently onto a 10×10 lattice or
(c) a 10×10 rectangular lattice of PVIs from Table 1 configured by manipulating $n = 100$ of

the 132 cities so that their positions on the lattice were roughly equivalent to their relative positions on the US map; see Table S4 in the on-line supplemental document.

For the first two ('uniform' and 'empirical') settings, we also expanded the sample size to $n = 400$ and $n = 900$ 'cities' by creating 20×20 and 30×30 lattices respectively.

Using each constructed lattice of 'cities,' we defined their neighbourhood structure and corresponding adjacency matrix $\mathbf{A}$ for the autologistic model by using the same simple adjacency as described earlier: $a_{ij} = 1$ if city j is physically adjacent (i.e. connects directly on the lattice) to city i, and $a_{ij} = 0$ otherwise.

Given values for the n PVIs and their adjacency or neighbourhood structure on the simulation lattice, we generated binary responses Y_i with response probabilities defined via equation (2), using the perfect sampler given by Hughes (2014). We selected a series of different configurations for the three unknown parameters in $\boldsymbol{\beta}$; we began with $\boldsymbol{\beta} = (-2.5, 0.5, -1)^{\mathrm{T}}$ which are roughly the values of the maximum pseudolikelihood estimates from fitting the centred autologistic model to the 132-cities data; see Section 5. We then manipulated the value of each parameter to vary somewhat from this base such that the response probability under the independence model at PVIs near 0 took small values between about 5% and 20%, and where the corresponding BMIs at BMR = 10 and BMR = 25 were between 2.5 and 6.5. These values bound roughly 95% of the calculated PVIs from Table 1. This led to $\beta_0 = -2.5, -1.5$, $\beta_1 = 0.5, 1$ and $\beta_2 = -1, 0, 1$. Coupled with the seven different sample size–lattice combinations above, this produced a total of 84 different design–parameter configurations for study.

For each simulation configuration, 2000 simulated data sets were generated. In each simulated data set, we calculated the consequent maximum pseudolikelihood estimates and, from these, the corresponding $\widehat{\mathrm{BMI}}$ from equation (5) and bootstrap-based BMIL from equation (7). We then assessed whether or not BMIL correctly covered (i.e. remained below) the true value of BMI for that parameter configuration. We set our nominal confidence level to $1 - \alpha = 0.95$ and recorded the empirical proportion of cases where BMIL did in fact express correct coverage. Note that, with 2000 simulations per configuration, the approximate standard error of our empirical coverage rates at the nominal 95% level is $\sqrt{\{(0.05)(0.95)/2000\}} = 0.005$ and it never exceeds $\sqrt{\{(0.5)(0.5)/2000\}} = 0.011$.

Under our construction equation (7), BMIL will correctly cover the true BMI from below if and only if the upper bootstrap limit $b_{1\mathrm{B}}$ correctly covers β_1 from above. For simplicity, we assessed whether only the latter condition held. In fact, the operation is independent of BMR, so our convergence assessments hold for any choice of BMR > 1.

We also collected summary statistics on the performance of the various point estimators under the centred autologistic model. For $\widehat{\mathrm{BMI}}$, we found the average relative bias $(1/2000)\Sigma_{t=1}^{2000}(\widehat{\mathrm{BMI}} - \mathrm{BMI})/\mathrm{BMI}$. Note from the form of BMI in equation (4) that any terms involving BMR cancel in this relative bias calculation, so the result will again be independent of the numerical value for the BMR.

Similarly, for the autologistic parameters β_0 and β_1 we found the average relative biases $(1/2000)\Sigma_{t=1}^{2000}(\hat{\beta}_j - \beta_j)/|\beta_j|$ $(j = 0, 1)$. For the autocovariate parameter we simply calculated the average bias $(1/2000)\Sigma_{t=1}^{2000}(\hat{\beta}_2 - \beta_2)$, since the true values of $|\beta_2|$ were set to either 0 or 1 under our configurations.

4.2. Simulation results

4.2.1. Coverage rates

Empirical coverage rates of our bootstrap-based lower 95% confidence limit on BMI under our centred autologistic model appear in Table 2. As noted above, these also represent empirical

Table 2. Empirical coverage rates for the centred autologistic benchmark index lower confidence limit BMIL based on 2000 simulated data sets, stratified by the PVI generation method (top; see the text), true autologistic regression parameter configuration (left column) and sample size *n*†

Coefficients $(\beta_0, \beta_1, \beta_2)$	*Uniform PVI*			*Empirical PVI*			*Map PVI,* $n = 100$
	$n = 100$	$n = 400$	$n = 900$	$n = 100$	$n = 400$	$n = 900$	
$(-1.5, 0.5, -1)$	0.972	0.957	0.962	0.958	0.958	0.960	0.964
$(-1.5, 0.5, 0)$	0.974	0.958	0.956	0.966	0.958	0.959	0.963
$(-1.5, 0.5, 1)$	0.972	0.959	0.967	0.970	0.965	0.961	0.968
$(-1.5, 1, -1)$	0.984	0.971	0.965	0.974	0.961	0.957	0.951
$(-1.5, 1, 0)$	0.986	0.969	0.968	0.974	0.960	0.966	0.969
$(-1.5, 1, 1)$	0.987	0.967	0.973	0.979	0.963	0.963	0.979
$(-2.5, 0.5, -1)$	0.978	0.960	0.954	0.957	0.959	0.959	0.961
$(-2.5, 0.5, 0)$	0.979	0.960	0.964	0.953	0.957	0.961	0.963
$(-2.5, 0.5, 1)$	0.974	0.967	0.963	0.964	0.961	0.960	0.967
$(-2.5, 1, -1)$	0.986	0.971	0.960	0.970	0.964	0.967	0.974
$(-2.5, 1, 0)$	0.991	0.976	0.968	0.969	0.967	0.960	0.980
$(-2.5, 1, 1)$	0.990	0.973	0.970	0.967	0.968	0.969	0.978
Average	0.981	0.966	0.964	0.967	0.962	0.962	0.968

†The nominal coverage level is set to 95%.

Table 3. Empirical rates of convergence failure for the maximum pseudolikelihood algorithm, each across 10^7 fitting attempts (5000 bootstrap resamples from 2000 simulated data sets), stratified by the PVI generation method (top; see the text), true autologistic regression parameter configuration (left column) and sample size *n*

Coefficients $(\beta_0, \beta_1, \beta_2)$	*Uniform PVI*			*Empirical PVI*			*Map PVI,* $n = 100$
	$n = 100$	$n = 400$	$n = 900$	$n = 100$	$n = 400$	$n = 900$	
$(-1.5, 0.5, -1)$	—†	—†	—†	—†	—†	—†	—†
$(-1.5, 0.5, 0)$	—†	—†	—†	—†	—†	—†	—†
$(-1.5, 0.5, 1)$	0.0001	—†	—†	0.0002	—†	—†	0.0001
$(-1.5, 1, -1)$	0.0027	—†	—†	0.0311	—†	—†	0.0185
$(-1.5, 1, 0)$	0.0019	—†	—†	0.0280	—†	—†	0.0145
$(-1.5, 1, 1)$	0.0019	—†	—†	0.0133	—†	—†	0.0067
$(-2.5, 0.5, -1)$	—†	—†	—†	—†	—†	—†	0.0001
$(-2.5, 0.5, 0)$	—†	—†	—†	—†	—†	—†	—†
$(-2.5, 0.5, 1)$	—†	—†	—†	0.0003	—†	—†	0.0001
$(-2.5, 1, -1)$	0.0001	—†	—†	0.0003	—†	—†	0.0002
$(-2.5, 1, 0)$	—†	—†	—†	0.0002	—†	—†	0.0001
$(-2.5, 1, 1)$	0.0001	—†	—†	0.0001	—†	—†	0.0003

†No convergence failures.

coverage rates of the bootstrap-based 95% upper confidence limit on β_1. As can be seen, all values rest above the nominal 95% confidence coefficient; in some cases the excess is significantly above the target level. As the sample size increases, the rates drop towards the nominal level, on average. These results suggest that the bootstrap-based limits operate in a reasonable, if slightly conservative, manner.

We should mention that we experienced some minor instabilities with the maximum pseudo-

Table 4. Empirical biases for the centred autologistic maximum pseudolikelihood estimates based on 2000 simulated data sets, stratified by PVI generation method (top; see the text), true autologistic regression parameter configuration (left column) and sample size n

Coefficients $(\beta_0,\beta_1,\beta_2)$	*Uniform PVI*			*Empirical PVI*			*Map PVI,* $n = 100$
	$n = 100$	$n = 400$	$n = 900$	$n = 100$	$n = 400$	$n = 900$	
Relative bias of β_0							
$(-1.5, 0.5, -1)$	−0.053	−0.011	−0.009	−0.095	−0.020	−0.011	−0.070
$(-1.5, 0.5, 0)$	−0.070	−0.014	−0.007	−0.092	−0.019	−0.006	−0.083
$(-1.5, 0.5, 1)$	−0.152	−0.053	−0.024	−0.181	−0.087	−0.075	−0.170
$(-1.5, 1, -1)$	−0.144	−0.024	−0.012	−0.030	−0.002	−0.015	0.007
$(-1.5, 1, 0)$	−0.191	−0.030	−0.012	−0.122	0.012	−0.015	−0.093
$(-1.5, 1, 1)$	−0.207	−0.024	−0.017	−0.225	0.000	−0.014	−0.211
$(-2.5, 0.5, -1)$	−0.052	−0.013	−0.004	−0.053	−0.010	−0.009	−0.047
$(-2.5, 0.5, 0)$	−0.058	−0.013	−0.008	−0.071	−0.012	−0.010	−0.063
$(-2.5, 0.5, 1)$	−0.024	−0.006	0.004	−0.009	0.033	0.031	−0.006
$(-2.5, 1, -1)$	−0.096	−0.021	−0.006	−0.108	−0.012	−0.007	−0.087
$(-2.5, 1, 0)$	−0.103	−0.023	−0.009	−0.111	−0.014	−0.007	−0.112
$(-2.5, 1, 1)$	−0.111	−0.016	−0.008	−0.154	−0.027	−0.014	−0.163
Average	−0.105	−0.021	−0.009	−0.104	−0.013	−0.013	−0.092
Relative bias of β_1							
$(-1.5, 0.5, -1)$	0.050	0.010	0.008	0.083	0.017	0.010	0.061
$(-1.5, 0.5, 0)$	0.063	0.013	0.006	0.085	0.018	0.007	0.075
$(-1.5, 0.5, 1)$	0.074	0.016	0.008	0.100	0.027	0.017	0.062
$(-1.5, 1, -1)$	0.130	0.024	0.011	0.096	0.013	0.012	0.070
$(-1.5, 1, 0)$	0.153	0.029	0.012	0.123	0.006	0.012	0.103
$(-1.5, 1, 1)$	0.150	0.024	0.014	0.144	0.012	0.011	0.125
$(-2.5, 0.5, -1)$	0.051	0.013	0.005	0.059	0.010	0.011	0.047
$(-2.5, 0.5, 0)$	0.056	0.012	0.008	0.074	0.012	0.012	0.067
$(-2.5, 0.5, 1)$	0.063	0.020	0.009	0.088	0.028	0.013	0.067
$(-2.5, 1, -1)$	0.089	0.020	0.006	0.093	0.012	0.007	0.076
$(-2.5, 1, 0)$	0.093	0.022	0.009	0.094	0.014	0.007	0.096
$(-2.5, 1, 1)$	0.090	0.019	0.009	0.084	0.018	0.012	0.094
Average	0.089	0.019	0.009	0.094	0.016	0.011	0.079
Bias of β_2							
$(-1.5, 0.5, -1)$	−0.129	−0.034	−0.019	−0.115	−0.032	−0.017	−0.137
$(-1.5, 0.5, 0)$	−0.068	−0.011	−0.011	−0.046	−0.005	−0.012	−0.074
$(-1.5, 0.5, 1)$	0.058	0.024	0.004	0.054	0.020	0.006	0.045
$(-1.5, 1, -1)$	−0.525	−0.111	−0.043	−3.142	−0.140	−0.053	−2.942
$(-1.5, 1, 0)$	−0.228	−0.076	−0.033	−1.553	−0.061	−0.033	−1.673
$(-1.5, 1, 1)$	−0.011	−0.024	−0.008	−0.276	−0.006	−0.007	−0.149
$(-2.5, 0.5, -1)$	−0.125	−0.028	−0.014	−0.128	−0.038	−0.015	−0.139
$(-2.5, 0.5, 0)$	−0.064	−0.014	−0.006	−0.057	−0.015	−0.009	−0.085
$(-2.5, 0.5, 1)$	0.047	0.025	0.006	0.048	0.024	0.007	0.054
$(-2.5, 1, -1)$	−0.207	−0.064	−0.024	−0.382	−0.041	−0.022	−0.216
$(-2.5, 1, 0)$	−0.123	−0.031	−0.019	−0.101	−0.020	−0.017	−0.111
$(-2.5, 1, 1)$	0.064	0.010	−0.001	0.033	0.013	0.001	0.040
Average	−0.109	−0.028	−0.014	−0.472	−0.025	−0.014	−0.449

Table 5. Empirical relative bias for the centred autologistic BMI estimator (5) based on 2000 simulated data sets, stratified by PVI generation method (top; see the text), true autologistic regression parameter configuration (left column) and sample size n

Coefficients $(\beta_0, \beta_1, \beta_2)$	*Uniform PVI*			*Empirical PVI*			*Map PVI,* $n = 100$
	$n = 100$	$n = 400$	$n = 900$	$n = 100$	$n = 400$	$n = 900$	
$(-1.5, 0.5, -1)$	0.072	0.019	0.005	−0.132	0.085	0.022	0.387
$(-1.5, 0.5, 0)$	0.055	0.011	0.004	−2.334	0.062	0.021	0.512
$(-1.5, 0.5, 1)$	0.073	0.015	0.004	−0.781	0.063	0.018	−0.125
$(-1.5, 1, -1)$	−0.011	−0.001	−0.001	0.885	0.053	0.012	1.158
$(-1.5, 1, 0)$	−0.020	−0.004	−0.002	0.141	0.064	0.011	0.392
$(-1.5, 1, 1)$	−0.019	0.001	−0.004	2.820	0.047	0.011	0.063
$(-2.5, 0.5, -1)$	0.024	0.016	0.008	0.373	0.086	0.019	0.158
$(-2.5, 0.5, 0)$	0.089	0.012	0.001	0.900	0.056	0.013	0.161
$(-2.5, 0.5, 1)$	0.099	0.009	0.003	−0.023	0.056	0.021	0.499
$(-2.5, 1, -1)$	−0.022	−0.006	0.000	0.173	0.029	0.008	0.047
$(-2.5, 1, 0)$	−0.027	−0.008	−0.003	0.149	0.022	0.007	0.059
$(-2.5, 1, 1)$	−0.020	−0.004	−0.003	0.101	0.016	0.002	0.162
Average	0.024	0.005	0.001	0.189	0.053	0.014	0.289

likelihood fitting algorithm. Cases occurred where the algorithm failed to converge with some bootstrapped resamples, so that `ngspatial` reported only '`NA`' for $\hat{\beta}_1$. The resulting bootstrap distribution therefore contained fewer than the desired $B = 5000$ resampled $\hat{\beta}_1$-values. Table 3 gives the rates of how often this occurred; each value is the number of reported `NA` values out of 10^7 fitting attempts (5000 bootstrap samples times 2000 simulated data sets) at each simulation configuration. As can be seen, the failure rates are quite low and occur only at the smallest sample size of $n = 100$. When this occurs in practice, users may simply increase B to perhaps 7500 or 10 000, to produce sufficient bootstrap resamples for BMIL. We view this as a slight inconvenience and a tolerable consequence of employing such a complex model or fitting procedure.

4.2.2. Estimator bias

Empirical biases of the maximum pseudolikelihood estimators under our centred logistic model appear in Table 4 (relative biases for β_0 and β_1). Some general patterns are clear: $\hat{\beta}_0$ and $\hat{\beta}_2$ typically exhibit small negative bias, whereas $\hat{\beta}_1$ exhibits small positive bias. The absolute relative biases are usually of the order of 10^{-1} or smaller, and all biases trend towards 0 as the sample size increases.

Similar values for $\widehat{\mathrm{BMI}}$ appear in Table 5. We see a more variable pattern among these relative bias terms, however, with a mix of both positive and negative values. Most remain small, although a few relative biases past ± 2 do appear with $n = 100$ by using the empirical generator for $x = \mathrm{PVI}$. As the sample size increases, however, the relative biases do mitigate and move towards 0 on average.

We speculate that the mixed bias patterns for $\widehat{\mathrm{BMI}}$ may be a function of the reciprocal operation that was used in equation (5) when defining the BMI-estimator. The relative bias for $\hat{\beta}_1$ appears stable; however, using $1/\hat{\beta}_1$ in $\widehat{\mathrm{BMI}}$ induces greater instability. This is perhaps not surprising, as it is quite common to find that reciprocal random variables are often less stable than their original counterparts.

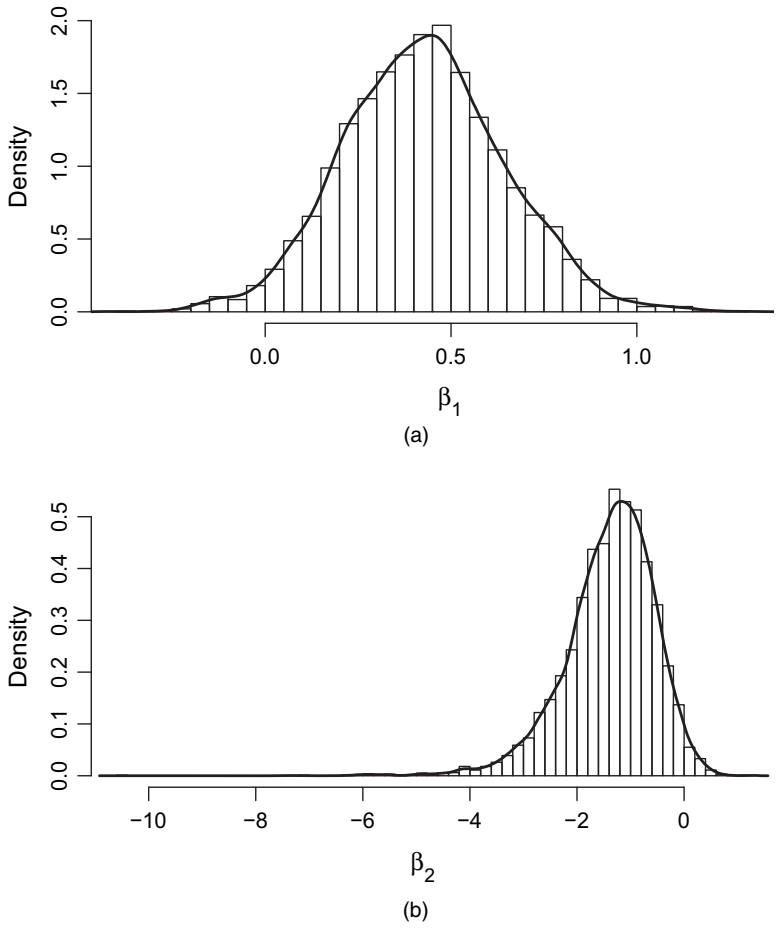

Fig. 1. Histograms and overlaid kernel density estimates (———) of the bootstrap distributions for (a) $\hat{\beta}_1$ and (b) $\hat{\beta}_2$, based on 5000 bootstrap replicates from the centred autologistic fit of the 132-cities data: the histograms employ Scott's rule for bin selection

5. Urban vulnerability example, revisited

Returning to our example data on placed-based urban vulnerability to terrorism, we apply the centred autologistic model to the data in Table 1. Our goal is to update the earlier study in Piegorsch *et al.* (2007) by incorporating potential spatial auto-correlation when studying the relationship between our PVI-measure and the terrorist casualty data, and also to explore how this auto-correlated model operates in a real, risk analytic, benchmarking application.

Regressing Y on $x = \text{PVI}$ via the centred autologistic model (2) produces the maximum pseudolikelihood estimates $\hat{\beta}_0 = -2.4460$, $\hat{\beta}_1 = 0.4111$ and $\hat{\beta}_2 = -1.0990$. Fig. 1 presents histograms (using Scott's rule for bin selection) and overlaid kernel density estimates of the corresponding bootstrap distributions for $\hat{\beta}_1$ and $\hat{\beta}_2$, based on the 5000 bootstrap values from `ngspatial`. (No `NA` instabilities were observed with this bootstrap resample.) The former appears roughly bell shaped with perhaps a slight right skew, whereas the latter displays a left skew towards a few extreme negative values. For this sample size, at least, the large sample features of the maximum pseudolikelihood estimates may not be fully realized with this data set.

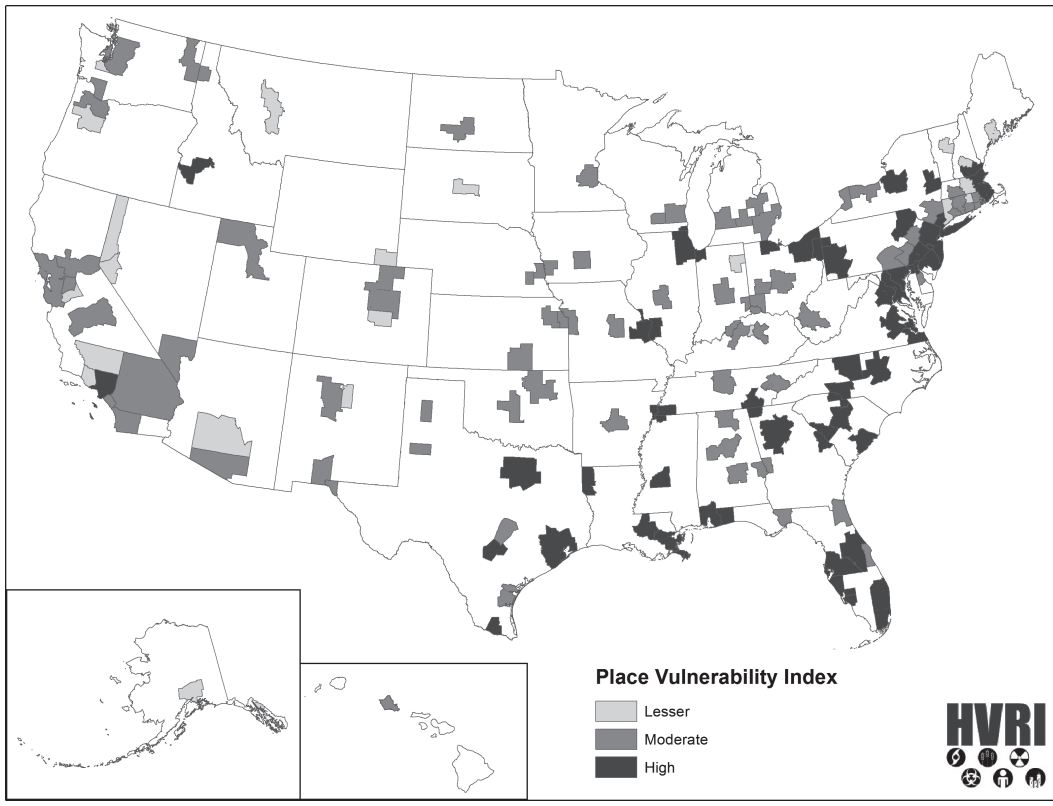

Fig. 2. Map of the 132 cities from Table 1, coded to indicate vulnerability status from our autologistic analysis: ▫, city PVI < $BMIL_{10}$; ■, $BMIL_{10} \leqslant$ PVI < $BMIL_{25}$; ■, PVI $\geqslant BMIL_{25}$

The `ngspatial` program provides pointwise 95% bootstrap confidence intervals for the non-intercept parameters as $0.0122 < \beta_1 < 0.8574$ and $-3.2750 < \beta_2 < -0.0436$. Since the first interval fails to contain $\beta_1 = 0$, we can conclude that our non-negative PVI-measure significantly affects the response probability π; from the strictly positive values in the interval it appears to do so via an increasing relationship. Similarly, since the second interval fails to contain $\beta_2 = 0$, we can conclude that the autocovariate is important for describing π. And, since the β_2-interval presents only negative values, it indicates that the spatial correlation appears negative. This result is intriguing: it suggests that, when a city exhibits a positive terrorist casualty result, an adjacent city would expect *not* to experience such a result, and vice versa.

5.1. Benchmark analysis

Moving to the benchmark calculations, we refer to equation (5) for a point estimate of BMI_{BMR}. As noted above, the choice for BMR here is an open issue, since the benchmark paradigm has not previously been applied with models of this sort. Thus a certain portion of this analysis must be viewed as exploratory in nature. After examining a variety of possible BMRs, we settled on BMR = 10 as a first choice for the relative increase in odds over background. Using equation (5), this led to $\widehat{BMI}_{10} = \log(10)/\hat{\beta}_1 = 5.6011$. The upper, one-sided, 95% bootstrap confidence bound on β_1 from the maximum pseudolikelihood fit (above) is $b_{1B} = 0.7882$, producing a 95% lower bootstrap limit of $BMIL_{10} = \log(10)/b_{1B} = 2.9213$.

For comparison, we also considered a higher benchmark response at BMR = 25, expecting that an increase of this magnitude in the odds could be an informative marker for practical use. This gave $\widehat{\text{BMI}}_{25} = \log(25)/\hat{\beta}_1 = 7.8300$, with 95% lower bootstrap limit of $\text{BMIL}_{25} = \log(25)/b_{1\text{B}} = 4.0838$. Urban localities with PVIs larger than this value might be viewed as exhibiting excess vulnerability to terrorism casualties, at least on the basis of our analysis.

To illustrate the benchmark delineations geospatially, Fig. 2 maps the 132 urban centres from Table 1, distinguishing those whose PVIs exceed our two benchmarks above: cities whose PVIs exceed BMIL_{25} are marked in black, whereas those whose PVIs exceed only BMIL_{10} are marked in dark grey. Locations with PVIs below both benchmarks are marked in light grey. Fig. 2 shows that urban locations exhibiting high relative benchmark vulnerability are primarily in the eastern half of the USA. There, the high vulnerability locations are generally concentrated along shorelines (eastern seaboard, Gulf Coast, Great Lakes and various rivers), and in the interior south. A few western locations also exceed BMIL_{25}, but not to the extent that is seen east of the mountain states. Lesser vulnerability urban centres appear almost exclusively in the west and through the northern tier of the map; no such lesser vulnerability cities appear in the deep south. This might be partly explained by patterns of urban development, where eastern urban centres are more dense, less separated (and often geographically continuous—the 'megalopolis' effect), more constrained by geophysical borders, such as rivers, and older than those in the west.

6. Discussion

We have described a statistical methodology that can characterize the vulnerability of US urban centres ('cities') to terrorist attack, using a place-based vulnerability index PVI and a database of terrorist incidents and related human casualties. We show how to incorporate potential auto-correlation in the geospatial data via a centred autologistic model and enhance a risk analytic 'benchmark' approach for identifying urban locations at increased risk or vulnerability to terrorist events. In fact, the methods are sufficiently extensible to apply in a variety of data scenarios where spatial auto-correlation may confound more simplistic models; the possibilities far exceed the specific urban vulnerability risk analysis that is studied here.

One can imagine an alternative construction to our autologistic strategy: fit a generalized linear mixed model with a logit link and spatially varying random intercepts. Similar approaches are common in disease mapping with count data, and they can account for spatial correlation across the mapped units in an effective manner (Waller and Carlin, 2010). Indeed, the two strategies could produce roughly similar results when the spatial correlation is positive. We are unsure how to implement a generalized linear mixed model with random intercepts to produce negative spatial correlation, however, as would be required with the data in our urban vulnerability analysis. Nonetheless, further manipulation of generalized linear mixed models with spatially varying coefficients for the risk analytic problems that were introduced here could prove quite fruitful for future research.

A critical component that is required in our analysis is the definition of the neighbourhood structure surrounding each metropolitan area. We approached this by using a simple, strict definition of adjacency, where one location's boundary—under the definitions that were employed in Borden *et al.* (2007)—must literally contact another's for them to be defined as neighbours; see Fig. 2 and also Fig. S1 in the on-line supplemental document. One can alternatively use these models to study spatial auto-correlation as a function of distance between cities, however. For example, we took paired longitudes and latitudes within each of the 132 Euclidean urban regions to define 'centres' for each city (see Table S3 in the supplemental document). From these, we

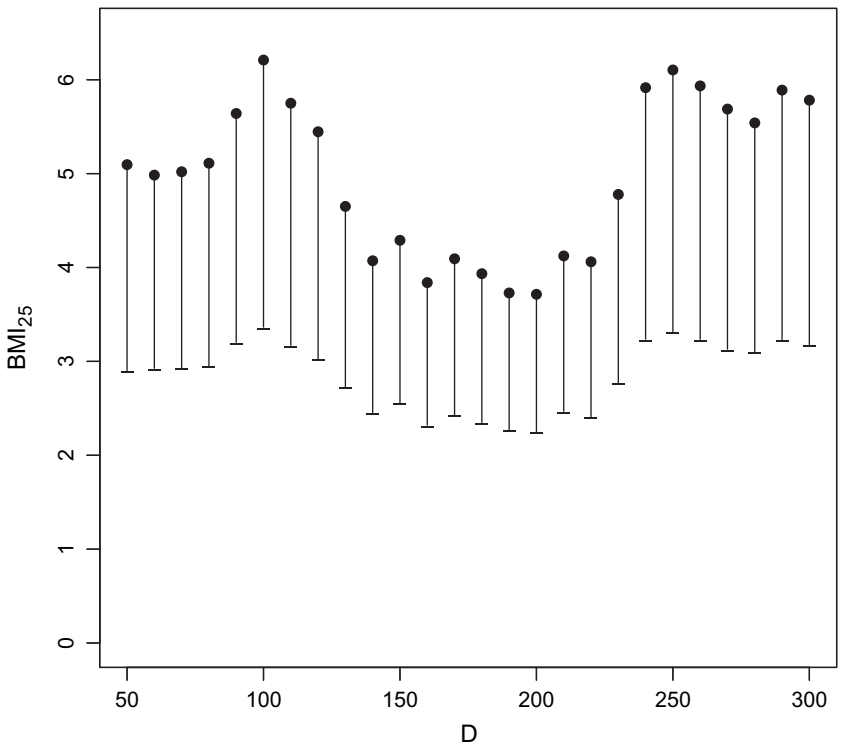

Fig. 3. Relationship between $\widehat{\text{BMI}}_{25}$-estimates (●) and separation distance D defining a neighbourhood within the centred autologistic fit for the 132-cities data: —, pointwise 95% BMIL_{25}-values

calculated the distance, say d_{ij}, between the centre points for the ith and jth city pairing. Then to define the neighbourhood N_i for city i in model (2) we populated that N_i with those cities j whose distance from city i rests below some fixed positive bound on the distance, say, $d_{ij} < D$. Fig. S2 in the supplemental document gives a visualization of the consequent neighbourhood structure at $D = 100$.

To illustrate, we returned to the data in Table 1 and calculated the point estimate and lower confidence limit for BMI_{25}. We applied a range of values for D, from 50 to 300 miles of city separation. The results appear in Fig. 3. We see that the point estimate oscillates somewhat, but in particular BMIL_{25} stays within a narrow band between about 2.25 and 3.65. Compare these with $\widehat{\text{BMI}}_{25} = 7.83$ and $\text{BMIL}_{25} = 4.06$ for the strict adjacency neighbourhoods in Section 5. Moving to different separation distances appears to lower the benchmark points to within a fairly stable range.

Some explanation for this may be available by conducting a similar distance analysis on the auto-correlation parameter estimate $\hat{\beta}_2$, as plotted in Fig. 4. There, we see that $\hat{\beta}_2$ also oscillates somewhat, but it generally stays near 0. In particular, all of its corresponding, pointwise, 95% bootstrap intervals contain $\beta_2 = 0$. The particularly wide limits at $D = 50$ are likely to be due to the small number of pairings that occur for this very tight neighbourhood definition: only 26 city pairings report as neighbours when D is set to 50. By contrast, 82 pairings report as neighbours under the adjacency structure. Thus, by moving to a distance-based neighbourhood definition, we find no significant indication of spatial auto-correlation under our model. The stricter adjacency definition does provide such an indication, however, and by accounting for

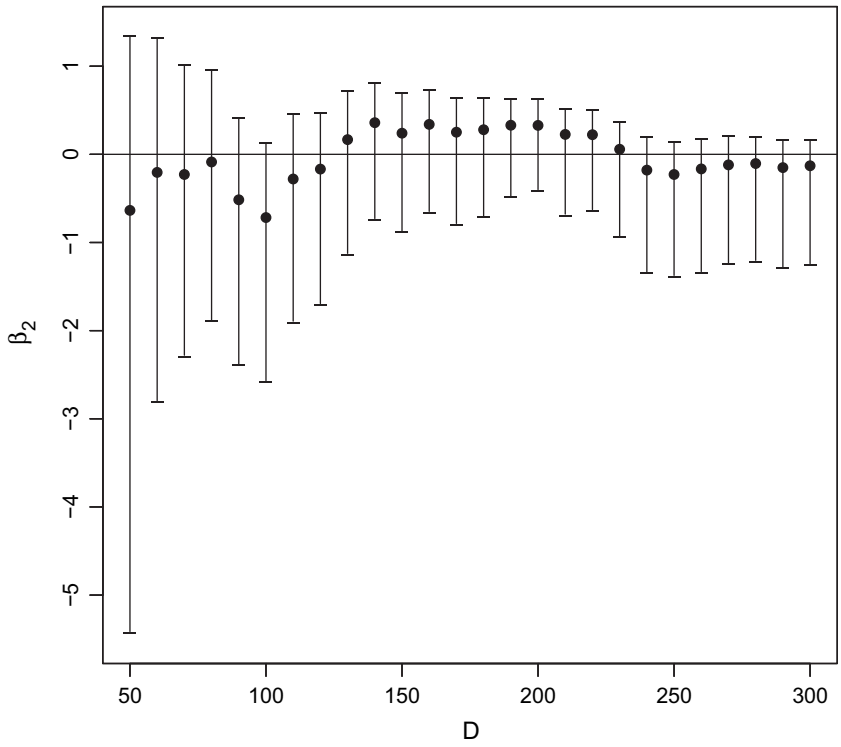

Fig. 4. Relationship between estimated autocovariance parameter $\hat{\beta}_2$ (●) and separation distance D defining a neighbourhood within the centred autologistic fit for the 132-cities data: ⌶, pointwise 95% confidence limits for β_2

this auto-correlation in the model the corresponding estimate and lower confidence limit on BMI_{25} become more sensitive and increase above those seen in Fig. 4. A similar phenomenon is exhibited at $\mathrm{BMR} = 10$ (the results are not shown).

If desired, one can alternatively incorporate other metrics to define connections between the cities, beyond those of simple adjacency or Euclidean distance that were employed here. Many forms are possible, including Manhattan ('city block') or Tchebychev ('chessboard') distance, Mahalanobis distance to incorporate non-spherical alignments in the separation dimensions or other constructs that define how two locations 'connect' to form the pertinent neighbourhoods (Liu (2017), section 3.2). The choice between these will depend on the particular needs of the risk analyst and on the geospatial nature of the incumbent data.

Our application to the US data on urban vulnerability enables interesting knowledge discovery: the indication of negative, adjacency-based, spatial correlation along the 132 cities in the database may seem counterintuitive and leads to intriguing speculation on its possible cause. Perhaps the occurrence of terrorist events in one city tends to increase emergency preparedness and urban response planning in adjacent cities, leading to fewer terrorism (or at least lowered casualty) events. Or, perhaps putative terrorists may not target nearby cities to maximize their desired impact across a wider geographic space.

Also, well below a quarter (22/132) of the cities in Fig. 2 are labelled in lesser vulnerability light grey, suggesting a rather disquieting vulnerability landscape: one might be concerned that almost any large- and medium-sized urban centre in the 50 United States is likely to exceed this low level benchmark at $\mathrm{BMR} = 10$ and to exhibit moderate-to-high vulnerability to terrorist

casualties. This agrees with a similar concern in our previous study (Piegorsch *et al.*, 2007), although we caution against making any direct comparison between our earlier results and those seen herein. Although both presentations are rooted in geospatial comparisons between the same 132 US urban centres, the underlying models and forms of benchmark analyses that were constructed here differ in a substantive fashion from those earlier, more-simplistic results. Nonetheless, it is possible to attempt some basic comparisons between the two risk analytic outcomes. We give a visualization-based juxtaposition in the on-line supplemental document—see supplemental section S3. Therein, our autologistic approach offers a more-sensitive–less-conservative indication of the vulnerability palette across these 132 cities, at least under the settings that we employed.

It is worth re-emphasizing the risk analytic nature of our modelling approach. On its own, PVI stands as a perfectly useful metric for quantifying and distinguishing multisource urban vulnerability to hazardous events. By employing PVI as a benchmarking predictor and then relating it to data on terrorist causalities via the autologistic model (2), however, we build a specific focus into the risk assessment. Our results help to identify those values of PVI which indicate targeted, place-based vulnerability to terrorist impacts. For example, Boise, Idaho, exhibits a high predicted vulnerability to general hazards, on the basis of its large PVI of 5.415. This is largely due to its past experiences with natural and human-induced catastrophic events (Borden *et al.*, 2007). In effect, our model employs Boise's particular mix of socio-economic, physical and built-environment features, as encapsulated in its PVI, to provide risk analytic guidance on that city's specific vulnerability to causalities that might occur from a terrorist attack. Even though Boise did not experience such an attack over the 35-year period in our study's database, our analysis suggests that its terrorist casualty risk exists at a heightened level. (See the on-line supplemental section S3 for more on this particular example.)

From a policy-making perspective, our approach could be employed for potential allocation of national and regional funding to support urban preparedness and response to terrorism events. Suppose that officials in, for example Norfolk, Virginia, or Charleston, South Carolina, were considering new or updated forms of coastal antiterrorist protection. According to our analysis, their relatively high PVIs indicate heightened vulnerability to terrorist events at the community level, even though, as with Boise, above, neither city experienced terrorism casualties during the 35-year period that is represented by the data in Table 1. This could motivate increased funding allocation(s) or other enhanced efforts to connect urban risk management programmes with terrorism vulnerability assessments in these communities (Kunreuther, 2002).

As in Piegorsch *et al.* (2007), we find that not all urban areas are equally at risk; nor do they have similar underlying vulnerabilities. A location's capacity to prepare adequately for and to respond to hazard(s) varies widely across the USA, especially in urban areas. We see that 'place matters' and that policy makers cannot ignore geographic differences in socio-economic, built- and natural environments when allocating antiterrorism resources and training. Doing so will correspondingly limit urban areas' abilities to prepare for and to respond to terrorist events. Hence the query remains: do our findings affect future urban planning in vulnerability screening, disaster management and allocation of homeland security resources, and are these questions of valid public concern? We think so, and we believe that much more work is required to quantify and understand urban vulnerability to terrorist events and to broader scale hazards. We hope that our results will stimulate further research on these various issues.

7. Supplemental materials

The supplemental document contains additional materials that are associated with the results

that are presented here. These include further details and visualizations with the neighbourhood structures for the 132 cities in Table 1, the specific PVI-values that were used for the 10×10 'map' lattice in our simulation study from Section 4, a comparison of our results in Fig. 2 with those seen in Piegorsch *et al.* (2007), an additional study of the centred autologistic model and sample R code to fit the centred autologistic model.

Acknowledgements

Thanks are due to Rachel E. Reeves for cartographic assistance, to Grace Chiu, Stephan R. Sain and Hideyasu Shimadzu for helpful discussions during the development of this work, and to the Joint Editor, Associate Editor and a referee for insightful comments on an earlier draft of the manuscript. This material represents a portion of the first author's doctoral dissertation from the University of Arizona Graduate Interdisciplinary Program in Statistics. The research was supported in part by grant ES027394 from the US National Institutes of Health.

References

Arnold, B. C. and Strauss, D. J. (1991) Pseudolikelihood estimation: some examples. *Sankhya* B, **53**, 233–243.
Besag, J. (1975) Statistical analysis of non-lattice data. *Statistician*, **24**, 179–195.
Besag, J. E. (1972) Nearest-neighbour systems and the auto-logistic model for binary data. *J. R. Statist. Soc.* B, **34**, 75–83.
Borden, K. A., Schmidtlein, M. C., Emrich, C. T., Piegorsch, W. W. and Cutter, S. L. (2007) Vulnerability of US cities to environmental hazards. *J. Homlnd Secur. Emerg. Mangmnt*, **4**, no. 2, article 5.
Budtz-Jørgensen, E., Keiding, N. and Grandjean, P. (2001) Benchmark dose calculation from epidemiological data. *Biometrics*, **57**, 698–706.
Caragea, P. C. and Kaiser, M. S. (2009) Autologistic models with interpretable parameters. *J. Agric. Biol. Environ. Statist.*, **14**, 281–300.
Chun, Y. and Griffith, D. A. (2013) *Spatial Statistics and Geostatistics: Theory and Applications for Geographic Information Science and Technology*. Thousand Oaks: Sage.
Crump, K. S. (1984) A new method for determining allowable daily intakes. *Toxicol. Sci.*, **4**, 854–871.
Crump, K. S. (1995) Calculation of benchmark doses from continuous data. *Risk Anal.*, **15**, 79–89.
Crump, K. S. (2012) Benchmark analysis. In *Encyclopedia of Environmetrics*, vol. 1 (eds A. H. El-Shaarawi and W. W. Piegorsch), 2nd edn, pp. 191–198. Chichester: Wiley.
Cutter, S. L., Boruff, B. J. and Shirley, W. L. (2003) Social vulnerability to environmental hazards. *Socl Sci. Q.*, **84**, 242–261.
de Frutos, A., Olea, P. P. and Vera, R. (2007) Analyzing and modelling spatial distribution of summering lesser kestrel: the role of spatial autocorrelation. *Ecol. Modllng*, **200**, 33–44.
Haimes, Y. Y. (2006) On the definition of vulnerabilities in measuring risks to infrastructures. *Risk Anal.*, **26**, 293–296.
Hosmer, D. W., Lemeshow, S. and Sturdivant, R. X. (2013) *Applied Logistic Regression*, 3rd edn. New York: Wiley.
Hughes, J. (2014) *ngspatial*: a package for fitting the centered autologistic and sparse spatial generalized linear mixed models for areal data. *R J.*, **6**, no. 2, 81–95.
Hughes, J., Haran, M. and Caragea, P. C. (2011) Autologistic models for binary data on a lattice. *Environmetrics*, **22**, 857–871.
Kolaczyk, E. D. and Csárdi, G. (2014) *Statistical Analysis of Network Data with R*. New York: Springer.
Koutsias, N. (2003) An autologistic regression model for increasing the accuracy of burned surface mapping using landsat thematic mapper data. *Int. J. Remote Sens.*, **24**, 2199–2204.
Kunreuther, H. (2002) Risk analysis and risk management in an uncertain world. *Risk Anal.*, **22**, 655–664.
Liu, J. (2017) Autologistic modeling in benchmark risk analysis. *PhD Thesis.* Interdisciplinary Program in Statistics, University of Arizona, Tucson.
MacKenzie, C. A. (2014) Summarizing risk using risk measures and risk indices. *Risk Anal.*, **34**, 2143–2162.
Mazzorana, B., Zischg, A., Largiader, A. and Hübl, J. (2009) Hazard index maps for woody material recruitment and transport in alpine catchments. *Natrl Hazrds Earth Syst. Sci.*, **9**, 197–209.
Peng, R. D., Schoenberg, F. P. and Woods, J. A. (2005) A space–time conditional intensity model for evaluating a wildfire hazard index. *J. Am. Statist. Ass.*, **100**, 26–35.
Piegorsch, W. W., Cutter, S. L. and Hardisty, F. (2007) Benchmark analysis for quantifying urban vulnerability to terrorist incidents. *Risk Anal.*, **27**, 1411–1425.

Piegorsch, W. W. and West, R. W. (2005) Benchmark analysis: shopping with proper confidence. *Risk Anal.*, **25**, 913–920.
R Core Team (2016) *R: a Language and Environment for Statistical Computing*. Vienna: R Foundation for Statistical Computing.
Santini, M., Caccamo, G., Laurenti, A., Noce, S. and Valentini, R. (2010) A multi-component GIS framework for desertification risk assessment by an integrated index. *Appl. Geog.*, **30**, 394–415.
Varin, C., Reid, N. and Firth, D. (2011) An overview of composite likelihood methods. *Statist. Sin.*, **21**, 5–42.
Vaughn, N. R., Asner, G. P., Smit, I. P. J. and Riddel, E. S. (2015) Multiple scales of control on the structure and spatial distribution of woody vegetation in African savanna watersheds. *PLOS One*, **10**, no. 12, article e0145192.
Waller, L. A. and Carlin, B. P. (2010) Disease mapping. In *Handbook of Spatial Statistics* (eds A. E. Gelfand, P. J. Diggle, M. Fuentes and P. Guttorp), pp. 217–243. Boca Raton: CRC Press.

Supporting information

Additional 'supporting information' may be found in the on-line version of this article:

'Supplementary material to "Autologistic models for benchmark risk/vulnerability assessment of urban terrorism outcomes"'.