Environmental Research 110 (2010) 633–638

Contents lists available at ScienceDirect

Environmental Research

journal homepage: www.elsevier.com/locate/envres

Commentary

Time series analysis on the health effects of temperature: Advancements and limitations ☆

Antonio Gasparrini *, Ben Armstrong

Department of Social and Environmental Health Research, London School of Hygiene & Tropical Medicine, 15-17 Tavistock Place, London WC1H 9SH, UK

ARTICLE INFO

Article history:
Received 8 June 2010
Accepted 9 June 2010
Available online 23 June 2010

Keywords:
Temperature
Time series
Distributed lag
Harvesting
Multivariate meta-analysis

ABSTRACT

The association between extreme temperatures and health outcomes has been frequently investigated during the last few years. This assessment is usually based on a time series design, a framework which has gained a substantial development in the last two decades. In this contribution we offer an overview of the recent methodological advancements which provide new statistical tools to examine the health effects of temperature in a time series setting, highlighting at the same time the main limitations that still affect this research area.

1. Introduction

The increase in frequency and intensity of extreme weather events predicted in the near future (Luber and McGeehin, 2008) is arousing a growing interest, in the scientific and public health communities (Basu, 2009; Basu and Samet, 2002; Gosling et al., 2009). Several studies have investigated the association between mortality and both hot and cold temperatures, reporting increased risks in populations exposed to a wide range of climates (Analitis et al., 2008; Anderson and Bell, 2009; Baccini et al., 2008; McMichael et al., 2008). These studies are usually based on a time series design, where the series of daily counts of death or hospitalisations and ambient levels of temperature are compared, while controlling for potential confounding variables such as long-term and seasonal trends, air pollution and influenza epidemics. The purpose of these studies is to estimate the change in the counts of events associated with ambient temperature on the same day and on previous days (so-called *lagged effects*). Statistical approaches focus on regression methods within the generalized linear or additive modelling frameworks (GLM or GAM, respectively), assuming a Poisson distribution of the response (daily counts), and usually accounting for overdispersion (arising when the observed variance is greater than the expected number of daily events, differently from the standard Poisson assumption). Time series studies on temperature have benefitted from the remarkable statistical developments, achieved in the last two decades, to quantify the short-term effects of air pollution (Bell et al., 2004; Dominici, 2004; Schwartz et al., 1996; Touloumi et al., 2004).

In a paper published in this issue of the Journal, Barnett et al. (this issue) performed an analysis to examine the relationship between mortality and different temperature indexes, using a large dataset from 107 cities in the USA over a 14 years period. The analytical approach proposed by the authors highlights the flexibility and effectiveness of time series methods to attain sophisticated inferential deductions about complex associations. However, this complexity requires elaborate statistical tools that might appear obscure to many readers inexperienced with time series methods. In this contribution we attempt to review and elucidate recent advances in and limitations of this study design when applied to examine temperature–health associations, focussing mostly on the statistical issues.

2. Temporal decomposition

The time series design is characterized by a distinctive temporal structure of the data, with observations collected at ordered and equally spaced time points. In applications in environmental epidemiology, these time periods usually correspond to days, the smallest unit of time for which health outcome data are collected routinely. The main feature of the analytical methods is the *temporal decomposition* of the outcome and exposure series, where the variability is partitioned into contributions related to different timescales (Dominici et al.,

☆ *Funding*: Medical Research Council (UK); contract/grant number G0707030.
* Corresponding author. Fax: +44 20 75804524.
E-mail address: antonio.gasparrini@lshtm.ac.uk (A. Gasparrini).

0013-9351/$ - see front matter © 2010 Elsevier Inc. All rights reserved.
doi:10.1016/j.envres.2010.06.005

2003a; Zeger et al., 2006). From an epidemiological perspective, the temporal partition of contributions to the exposure–response function addresses different issues. First, an exposure may lead to multiple physiological mechanisms operating at different timescales, whose effects can be disentangled by decomposing the series. In addition, specific confounding factors can act at different temporal frames; hence, the decomposition may produce virtually unbiased estimates at specific timescales in the presence of unmeasured confounders, if such factors act on longer temporal frames.

In the first methods that were proposed, the partition of both the response and exposure series was obtained by Fourier series decomposition (Zeger et al., 1999) or seasonal-trend decomposition using LOESS functions (Schwartz, 2000b; Schwartz, 2001), and then the correlations between components at corresponding timescales were estimated. In current applications, the decomposition is directly achieved through regression models, applying functions to describe seasonality and long-term trends, thus *filtering out* the effects of unmeasured factors that change slowly in time (Peng and Dominici, 2008). This approach leaves only the residual shorter-term variation to be explained by other factors that have day-to-day variability, like temperature. Originally, harmonic functions based on pairs of sine-cosine terms of day of the year were used to model the cyclic seasonal component (Hunsberger et al., 2002; Stolwijk et al., 1999), with non-linear functions of time like polynomial terms to describe the long-term trend. Recently, the use of a single spline function of time has been favoured, producing an irregular seasonal trend which is believed to control for additional confounding effects operating at medium timescales. The main choices are based on regression splines within GLM or penalized splines within GAM (Ruppert et al., 2003; Schimek, 2009).

Independently of the type of spline and modelling framework, the main concern is the selection of the appropriate amount of smoothing in order to avoid residual confounding, but at the same time leaving a temporal window with enough variability to be explained by temperature. This choice corresponds to the selection of the optimal number of (effective) degrees of freedom (df) per year in the spline for time.

3. Exposure–response relationship

Temperature usually shows a typical association with health outcomes, characterized by *non-linear* and *delayed* effects. Empirically, risk may increase at both hot and cold temperatures, with the exposure–response relationship being described as *U*, *V* or *J*-shaped (Curriero et al., 2002; Hajat et al., 2007; Pattenden et al., 2003). In addition, the effect of an exposure to extreme temperatures is not limited to the same day, but persists for a period of time, typically from a few days for heat to some weeks for cold (Anderson and Bell, 2009; Braga et al., 2001). When assessing non-recurrent outcomes, an additional complexity is given by the *harvesting* effect: if temperature mainly affects a pool of susceptible individuals who would have otherwise experienced the outcome a few days later, the depletion of the pool after an extreme event will result in a decreased occurrence in those days (Hajat et al., 2005; Schwartz, 2000b). This anticipation (*displacement*) of the outcome will be measured as an increase in risk in the very first days, followed by a decrease some time later, with a smaller net effect.

These aspects require the application of subtle statistical approaches to accurately express in a regression model the exposure–response relationship for temperature effects, whose estimates usually require careful interpretation. The issue of non-linearity has been addressed in different ways, using a threshold parameterization to describe linear effects of cold and heat below and above specific cut-off temperatures, or alternatively relying on spline functions within GLM or GAM, as those previously described (Armstrong, 2006). The problem of delayed and harvesting effects has been tackled in air pollution studies, proposing the so-called *distributed lag models* (DLMs), where the linear but delayed effects were modelled including multiple lagged exposures (Schwartz, 2000a). In practice, the health effect in day t of the series is explained in terms of exposures at days $t-\ell$, with ℓ as the lag in the interval $0,\ldots,L$, and L as the maximum lag period. If the lag period considered is long, the distribution of effects can be modelled through a mathematical function; for example, strata (Welty and Zeger, 2005), polynomials (Goodman et al., 2004) or splines (Zanobetti et al., 2000) can be used to avoid collinearity in the model.

Despite the availability of well-developed methods to describe flexible but un-lagged exposure–response relationships, or alternatively flexible distributed lag models for simple linear dependencies, these two issues have been rarely addressed together. Extensions to DLMs have been proposed, applying distributed lag functions to each term of polynomial (Braga et al., 2001), linear piecewise (Roberts and Martin, 2007) or threshold functions (Muggeo, 2008). Nonetheless, these methods remain somewhat limited in their ability to describe complex dependencies. Recently, we have proposed a new modelling framework which can describe flexible relationships both in the space of the predictor and the lag dimension, leading to the family of *distributed lag non-linear models* (DLNMs) (Armstrong, 2006; Gasparrini et al., 2010). The core of this methodology is the specification of two independent functions to model the relationship along the two dimensions of predictor (temperature) and lags, respectively, given a menu of available choices. These two functions are then combined to form *cross-basis* variables to be included in the regression model, whose estimated parameters describe the bi-dimensional effect.

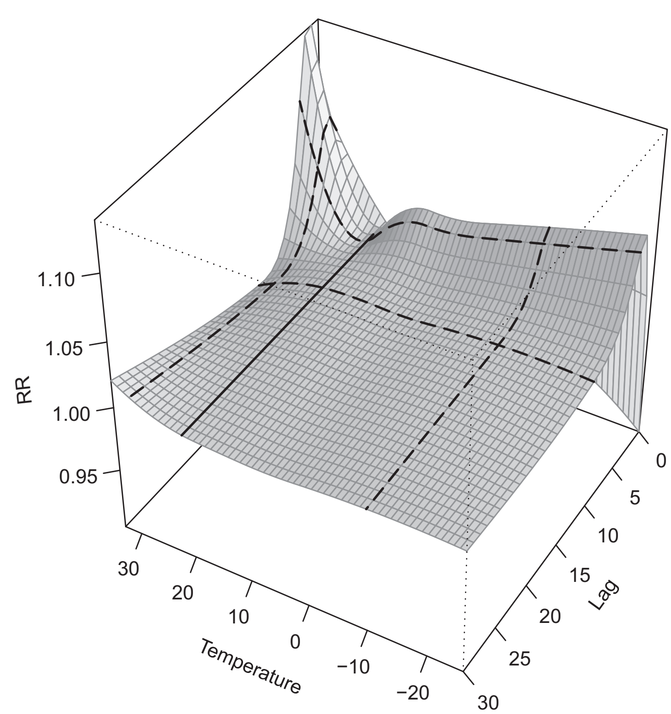

Fig. 1. 3-D plot of RR surface by temperature and lags. Highlighted are the reference at 21 °C (continuous line), the RR by lag at −10 and 30 °C, and the RR by temperature at lag 3 and 10 (dashed lines). Chicago 1987–2000.

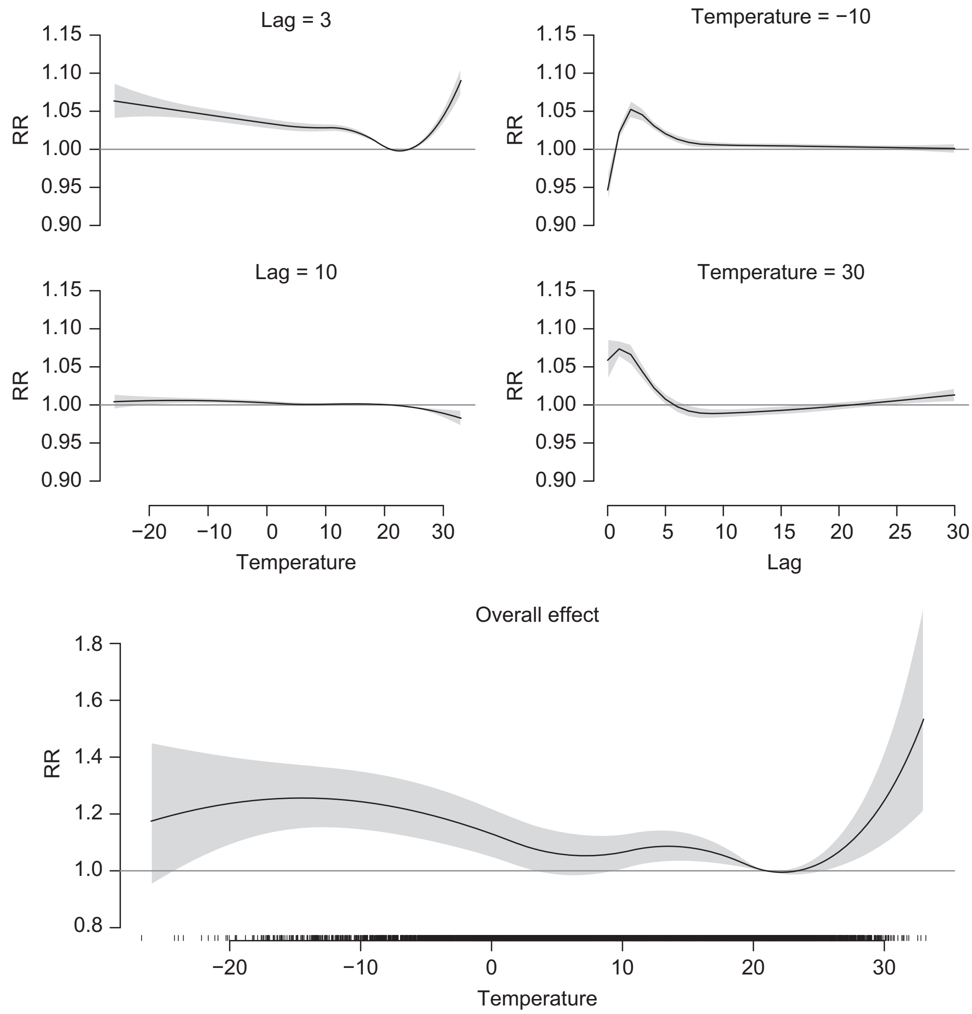

Fig. 2. Plot of RR by temperature at lag 3 and 10 (top left), RR by lag at −10 and 30 °C (top right) and overall RR (below). Reference at 21 °C. Chicago 1987–2000.

Figs. 1 and 2 show the application of the DLNM methodology to a time series of non-accidental deaths in Chicago, Illinois, during the period 1987–2000, using the same database that Barnett and colleagues analysed (Samet et al., 2000a, 2000b). In this example, we defined a cross-basis for temperature choosing a quadratic spline with 5 df for the space of temperature, and a natural cubic spline with 5 df for the space of lags, with 25 total parameters included. Fig. 1 shows the bi-dimensional relative risk (RR) surface for non-accidental mortality using a reference value of 21 °C, the empirical point of minimum mortality: the dashed lines represent the effects by lags for specific temperatures (−10 and 30 °C), and conversely the effects by temperature at specific lags (3 and 10 days). These effects are also reported, together with confidence intervals, in the top of Fig. 2. Lag-specific effects have a two-fold interpretation: each of them represents the increase in risk in a day t given a unit increase in temperature at day $t-\ell$ (*backward* interpretation, from outcome to exposure), or alternatively the increase in risk related to a unit increase in temperature at the day t during the following $t+\ell$ day (*forward* interpretation, from exposure to outcome). The overall effects are computed by the sum of lag contributions, and are illustrated in the bottom of Fig. 2. The results from this model suggest a strong and immediate effect of heat in the first 5 days, followed by a decrease after around 10 days, potentially interpreted as harvesting; cold temperatures display a more delayed effect, lasting up to 15 days.

The DLNM modelling framework is implemented within the software R (R Development Core Team, 2010) in the package 'dlnm' (Gasparrini and Armstrong, 2010). Further information about the analysis for Chicago and the package can be found at http://cran.r-project.org/web/packages/dlnm/vignettes/dlnmOverview.pdf.

4. Pooling the results

The health effects of environmental factors are assessed often through multi-site studies, using a two-stage hierarchical analysis with a common site-specific model and then the application of meta-analytic techniques to pool the results (Dominici et al., 2000; Samoli et al., 2008), the same strategy used by Barnett and colleagues. This approach ensures that the heterogeneity between different locations is properly accounted for, allowing model parameters to vary across sites, but at the same time avoiding additional variability and potential biases due to differential modelling choices (Dominici et al., 2003b; Touloumi et al., 2004). Meta-regression methods are commonly applied to assess the effect modification of site-level characteristics.

Air pollution studies are consistent with a linear exposure–response relationship, summarizing the effect with a single coefficient estimating the log-RR for a unit increase in exposure. The non-linear effect of temperature poses additional challenges, and several solutions have been proposed. First, the exposure–response relationship may be simplified assuming linear dependencies beyond site-specific thresholds (Baccini et al., 2008;

McMichael et al., 2008), or alternatively restricting the analysis to specific seasons, where strong deviations from linearity are not expected (Analitis et al., 2008; Zanobetti and Schwartz, 2008). An alternative solution is to produce a summary measure of the estimated non-linear relationship, for example computing average slopes (Curriero et al., 2002), or estimating a single RR for specific absolute or relative (distribution percentiles) temperatures (Anderson and Bell, 2009; Stafoggia et al., 2006). The use of site-specific thresholds or relative temperatures is usually preferred, in order to take into account the adaptation of populations to their own climate. A standard meta-analysis is then carried out for these single parameters.

The methods illustrated above have limitations: constraining the exposure–response to a simple shape could generate biased results, especially when assessing lag-specific effects. Even if strong assumptions are not formulated on the single-site models, pooling only simple summary measures of these might lose important features of a complex association. More sophisticated approaches rely on multivariate meta-analytical techniques, applied to relationships described by multiple parameters (Jackson et al., 2010; van Houwelingen et al., 2002), which are simultaneously pooled while accounting for their correlations within each site-specific model. These methods have been used to investigate dose–response functions (Baccini et al., 2008; Dominici et al., 2002; Samoli et al., 2005) or distributed lag curves (Analitis et al., 2008; Samoli et al., 2009) in multi-site studies. However, this approach is suitable only for associations expressed by a limited number of parameters. It is currently computationally infeasible, for example, to pool surfaces as the one illustrated in Figs. 1 and 2. An interesting compromise is the *meta-smoothing* method (Schwartz and Zanobetti, 2000), where simple univariate meta-analyses can be applied to pool the effects for any combinations of temperature values and lags, without accounting for correlations. Further research is needed to assess the presence and extent of biases in this type of approach for point estimates and standard errors if compared to proper multivariate methods (Riley, 2009), and to develop approaches to investigating heterogeneity in temperature–health association over sites (multivariate meta-regression).

As it is, investigators must balance the advantages of keeping information from the site-specific model with the need to reduce the number of parameters (summary measures) to make second-stage meta-analytical methods feasible. This choice also depends on the aim of the investigation and the availability of data.

5. Model selection

In contrast to analyses performed in many other subject-areas, the regression models applied in time series data for environmental factors are based on a limited number of predictors, such as day of the week, indicators for holiday periods and influenza epidemics, weather and pollution variables. The predictors to be included in the model are typically defined a priori, in particular in multi-site studies. In air pollution research, the critical choice to select the final model thus commonly focuses on the specification of the functions to account for seasonal and long-term trend, as discussed above in Section 2. Several contributions have addressed this issue, comparing alternative selection criteria (Baccini et al., 2007; Burnett et al., 1997; Peng et al., 2006; Touloumi et al., 2006): the main options are based on Akaike or Bayesian information criteria (AIC and BIC, respectively), (generalized) cross-validation techniques, minimization of the partial auto-correlation function of the residuals (PACF) or the related white noise test. The first 3 statistics aim to maximize the ability of the model to predict new observations arising from the same phenomenon which produced the data, while the last two intend to minimize the correlation between residuals from proximate observations in the series, to match the standard assumption of uncorrelated residuals. While these models fit statistics and residual analyses provide helpful insight, as criteria none can guarantee control of confounding (Peng et al., 2006). More complex approaches remain in the domain of statistical theory (Crainiceanu et al., 2008; Dominici et al., 2004). This leaves this aspect of model choice controversial, and makes analyses of sensitivity of key findings to variations in model choices very important.

When temperature is the focus of the analysis, similar considerations apply but additional issues should be taken into account. First, given the stronger association of temperature than of pollution with season, the optimal amount of smoothing to control for time may not be the same as that applied when air pollution is the exposure of interest. Second, given the longer lag often suggested for temperature (Anderson and Bell, 2009; Braga et al., 2001) there is more tension between the need to control confounding by unmeasured factors causing medium-term fluctuations in mortality (favoring many degrees of freedom in the time smooth) and the need to leave variation from which effects of interest can be estimated (favoring fewer degrees of freedom). Finally, when the objective is to investigate temperature–mortality relationships in their own right, there is usually a trade-off between complete description of all patterns not explicable by noise (many criteria often select quite complex models) and simplicity of interpretation. Specific study purposes may suggest different trade-offs.

The analysis performed by Barnett and colleagues illustrates some of these issues. A cross-validation procedure is specifically justified for the purpose of comparing the predictive ability of different temperature indexes, which might turn out to be useful, for example, to assess the future burden of climate change. In their analysis, the performance of the selection criteria is thus consistent with the research question. In other circumstances, for example when the goal is to obtain unbiased estimates of the exposure–response relationship or to control for confounding effects, the choice of selection criteria may be different (Dominici et al., 2008; Peng et al., 2006).

6. Implication of the ecological design

Ecological studies are defined as those in which the unit of the analysis is represented by aggregated or grouped observations (Last, 2001): the evidence from these research designs is interpreted with caution, given the inherent risk of biases due to the lack of information about individual characteristics (Greenland and Robins, 1994). Two of the main limitations that have been emphasized are the presence of unmeasured confounders, and error due to measures being collected from monitors at central sites, which do not represent personal exposures, which vary.

However, for investigating acute effects of environmental stressors (ambient temperature or air pollution) that change over time, the limitations inherent in ecological designs are offset by advantages. Many individual factors such as genetic make-up do not vary over time so cannot confound. Others, such as diet or smoking, vary only slowly so their effects are filtered out by the smooth function of time as discussed in Section 2. Moreover, variation in exposure across individuals is not the problem that it first appears to be. While there may be large variation in temperature or air pollution over a city, changes in daily population averages are usually much better reflected by the central monitor. This implies that the error in assigning individuals to central site levels is primarily of a Berkson-type, which in linear models does not lead to bias in estimates of effect in linear models, though it reduces the precision (Armstrong, 1998; Zeger et al., 2000). Thus, results from time series studies in this context are considered more robust than those achieved from other ecological designs.

Nevertheless, there are some reasons for caution in interpreting ecological time series studies of the effects of temperature. Berkson error may be a more relevant problem for studies of temperature than of air pollution effects, because of the non-linearity of most temperature–health relationships. For example, if a threshold-linear model pertained at the individual level, variation of temperatures across the city would blur the threshold, even if the between-day time fluctuations in individuals were perfectly correlated with those of the central site measurements. There are also implications of the ecological design for assessment of harvesting and, more broadly, of the lag pattern of the effect. Elegant conceptual frameworks have been proposed to characterize this phenomena (Rabl, 2005; Schwartz, 2000b), but it may not be appreciated that the apparent post-exposure protective effect in the presence of harvesting is not a real effect acting at the individual level, but an artefact of the ecological nature of the design. The measured decrease in risk is explained by the change in the structure of the underlying population after the depletion of the pool of frail individuals. More generally, lag curves such as those depicted in Fig. 2 are likely to be the results of the sum of delayed positive effects and harvesting, and should not be automatically interpreted as the temporal representation of some physiological mechanism linking temperature and mortality.

7. Conclusions

Time series analysis represents a valuable tool to assess the acute health effects of environmental factors that fluctuate over time. The recent developments described above address some of the main problems regarding its application in temperature–health studies, providing flexible methods to investigate this complex association. These investigations are relatively simple to conduct because of the routinely collected data, available in most locations. However, it is important to consider that this ecological design still has some limitations, which need to be kept in mind when planning a study or interpreting analytical results.

In addition, these new approaches face new challenges related to the complexity of the analytical methods, mainly due to the need to select a model from a large number of alternatives. As highlighted by Barnett and colleagues, the estimate of the association is particularly sensitive to the choice of functions, lag period and other model parameters, and available selection criteria are still limited to reliably identify a ‘‘best’’ model. An extensive sensitivity analysis on the various modelling choices is therefore always recommended.

References

Analitis, A., Katsouyanni, K., Biggeri, A., Baccini, M., Forsberg, B., Bisanti, L., et al., 2008. Effects of cold weather on mortality: results from 15 European cities within the PHEWE Project. American Journal of Epidemiology. 168, 1397.

Anderson, B.G., Bell, M.L., 2009. Weather-related mortality: how heat, cold, and heat waves affect mortality in the United States. Epidemiology 20, 205–213.

Armstrong, B., 1998. Effect of measurement error on epidemiological studies of environmental and occupational exposures. Occupational and Environmental Medicine 55, 651.

Armstrong, B., 2006. Models for the relationship between ambient temperature and daily mortality. Epidemiology 17, 624–631.

Baccini, M., Biggeri, A., Accetta, G., Kosatsky, T., Katsouyanni, K., Analitis, A., et al., 2008. Heat effects on mortality in 15 European cities. Epidemiology 19, 711–719.

Baccini, M., Biggeri, A., Lagazio, C., Lertxundi, A., Saez, M., 2007. Parametric and semi-parametric approaches in the analysis of short-term effects of air pollution on health. Computational Statistics and Data Analysis 51, 4324–4336.

Barnett, A.G., Tong, S., Clements, A.C.A. What measure of temperature is the best predictor of mortality? Environmental Research, this issue, doi:10.1016/j.envres.2010.05.006.

Basu, R., 2009. High ambient temperature and mortality: a review of epidemiological studies from 2001 to 2008. Environmental Health 8, 40.

Basu, R., Samet, J.M., 2002. Relation between elevated ambient temperature and mortality: a review of the epidemiologic evidence. Epidemiologic Reviews 24, 190–202.

Bell, M.L., Samet, J.M., Dominici, F., 2004. Time-series studies of particulate matter. Annual Review of Public Health 25, 247–280.

Braga, A.L., Zanobetti, A., Schwartz, J., 2001. The time course of weather-related deaths. Epidemiology 12, 662–667.

Burnett, R.T., Brook, J.R., Yung, W.T., Dales, R.E., Krewski, D., 1997. Association between ozone and hospitalization for respiratory diseases in 16 Canadian cities. Environmental Research 72, 24–31.

Crainiceanu, C.M., Dominici, F., Parmigiani, G., 2008. Adjustment uncertainty in effect estimation. Biometrika 95, 635.

Curriero, F.C., Heiner, K.S., Samet, J.M., Zeger, S.L., Strug, L., Patz, J.A., 2002. Temperature and mortality in 11 cities of the eastern United States. American Journal of Epidemiology 155, 80–87.

Dominici, F., 2004. Time-series analysis of air pollution and mortality: a statistical review. Research report—Health Effects Institute, 3–27; discussion, pp. 29–33.

Dominici, F., Daniels, M.J., Zeger, S.L., Samet, J.M., 2002. Air pollution and mortality: estimating regional and national dose-response relationships. Journal of the American Statistical Association 97, 100–111.

Dominici, F., McDermott, A., Hastie, T.J., 2004. Improved semiparametric time series models of air pollution and mortality. Journal of the American Statistical Association 99, 938–949.

Dominici, F., McDermott, A., Zeger, S.L., Samet, J.M., 2003a. Airborne particulate matter and mortality: timescale effects in four US cities. American Journal of Epidemiology 157, 1055–1065.

Dominici, F., Samet, J.M., Zeger, S.L., 2000. Combining evidence on air pollution and daily mortality from the 20 largest US cities: a hierarchical modelling strategy. Journal of the Royal Statistical Society: Series A 163, 263–302.

Dominici, F., Sheppard, L., Clyde, M., 2003b. Health effects of air pollution: a statistical review. International Statistical Review 71, 243–276.

Dominici, F., Wang, C., Crainiceanu, C., Parmigiani, G., 2008. Model selection and health effect estimation in environmental epidemiology. Epidemiology 19, 558–560.

Gasparrini, A., Armstrong, B., 2010. dlnm: Distributed Lag Non-linear Models in R. R package version 1.2.3.

Gasparrini, A., Armstrong, B., Kenward, M.G., 2010. Distributed lag non-linear models. Statistics in Medicine. Epub ahead of print May 7. 10.1002/sim.3940.

Goodman, P.G., Dockery, D.W., Clancy, L., 2004. Cause-specific mortality and the extended effects of particulate pollution and temperature exposure. Environmental Health Perspectives 112, 179–185.

Gosling, S.N., Lowe, J.A., McGregor, G.R., Pelling, M., Malamud, B.D., 2009. Association between elevated athmospheric temperatures and human mortality: a critical review of the literature. Climate Change 92, 299–341.

Greenland, S., Robins, J., 1994. Invited commentary: ecologic studies—biases, misconceptions, and counterexamples. American Journal of Epidemiology 139, 747.

Hajat, S., Armstrong, B.G., Gouveia, N., Wilkinson, P., 2005. Mortality displacement of heat-related deaths: a comparison of Delhi, Sao Paulo, and London. Epidemiology 16, 613–620.

Hajat, S., Kovats, R.S., Lachowycz, K., 2007. Heat-related and cold-related deaths in England and Wales: Who is at risk? Occupational and Environmental Medicine 64, 93–100.

Hunsberger, S., Albert, P.S., Follmann, D.A., Suh, E., 2002. Parametric and semiparametric approaches to testing for seasonal trend in serial count data. Biostatistics 3, 289–298.

Jackson, D., White, I.R., Thompson, S.G., 2010. Extending DerSimonian and Laird's methodology to perform multivariate random effects meta-analyses. Statistics in Medicine 29, 1282–1297.

Last, J.M., 2001. A dictionary of epidemiology. Oxford University Press, New York.

Luber, G., McGeehin, M., 2008. Climate change and extreme heat events. American Journal of Preventive Medicine 35, 429–435.

McMichael, A.J., Wilkinson, P., Kovats, R.S., Pattenden, S., Hajat, S., Armstrong, B., et al., 2008. International study of temperature, heat and urban mortality: the 'ISOTHURM' project. International Journal of Epidemiology 37, 1121.

Muggeo, V.M., 2008. Modeling temperature effects on mortality: multiple segmented relationships with common break points. Biostatistics 9, 613–620.

Pattenden, S., Nikiforov, B., Armstrong, B.G., 2003. Mortality and temperature in Sofia and London. Journal of Epidemiology and Community Health 57, 628–633.

Peng, R.D., Dominici, F., 2008. Statistical Methods for Environmental Epidemiology with R—A Case Study in Air Pollutioon and Health. Springer-Verlag, New York.

Peng, R.D., Dominici, F., Louis, T.A., 2006. Model choice in time series studies of air pollution and mortality. Journal of the Royal Statistical Society: Series A 169, 179–203.

R Development Core Team, 2010. R: A Language and Environment for Statistical Computing.

Rabl, A., 2005. Air pollution mortality: harvesting and loss of life expectancy. Journal of Toxicology and Environmental Health: Part A 68, 1175–1180.
Riley, R.D., 2009. Multivariate meta-analysis: the effect of ignoring within-study correlation. Journal of the Royal Statistical Society: Series A 172, 789–811.
Roberts, S., Martin, M.A., 2007. A distributed lag approach to fitting non-linear dose–response models in particulate matter air pollution time series investigations. Environmental Research 104, 193–200.
Ruppert, D., Wand, M., Carroll, R., 2003. Semiparametric Regression. Cambridge University Press.
Samet, J.M., Zeger, S.L., Dominici, F., Curriero, F., Coursac, I., Dockery, D.W., 2000a. The National Morbidity, Mortality, and Air Pollution Study (NMMAPS). Part 2, Morbidity and Mortality from Air Pollution in the United States. Health Effects Institute.
Samet, J.M., Zeger, S.L., Dominici, F., Dockery, D., Schwartz, J., 2000b. The National Morbidity, Mortality, and Air Pollution Study (NMMAPS). Part 1, Methods and Methodological Issues. Health Effects Institute.
Samoli, E., Analitis, A., Touloumi, G., Schwartz, J., Anderson, H.R., Sunyer, J., et al., 2005. Estimating the exposure–response relationships between particulate matter and mortality within the APHEA multicity project. Environmental Health Perspectives 113, 88–97.
Samoli, E., Peng, R., Ramsay, T., Pipikou, M., Touloumi, G., Dominici, F., et al., 2008. Acute effects of ambient particulate matter on mortality in Europe and North America: results from the APHENA study. Environmental Health Perspectives 116, 1480–1486.
Samoli, E., Zanobetti, A., Schwartz, J., Atkinson, R., Le Tertre, A., Schindler, C., et al., 2009. The temporal pattern of mortality responses to ambient ozone in the APHEA project. Journal of Epidemiology and Community Health 63, 960–966.
Schimek, M., 2009. Semiparametric penalized generalized additive models for environmental research and epidemiology. EnvironMetrics 20, 699–717.
Schwartz, J., 2000a. The distributed lag between air pollution and daily deaths. Epidemiology 11, 320–326.
Schwartz, J., 2000b. Harvesting and long term exposure effects in the relation between air pollution and mortality. American Journal of Epidemiology 151, 440–448.
Schwartz, J., 2001. Is there harvesting in the association of airborne particles with daily deaths and hospital admissions? Epidemiology 12 55–61.
Schwartz, J., Spix, C., Touloumi, G., Bacharova, L., Barumamdzadeh, T., le Tertre, A., et al., 1996. Methodological issues in studies of air pollution and daily counts of deaths or hospital admissions. Journal of Epidemiology and Community Health 50 (Suppl 1), S3–S11.
Schwartz, J., Zanobetti, A., 2000. Using meta-smoothing to estimate dose–response trends across multiple studies, with application to air pollution and daily death. Epidemiology 11, 666–672.
Stafoggia, M., Forastiere, F., Agostini, D., Biggeri, A., Bisanti, L., Cadum, E., et al., 2006. Vulnerability to heat-related mortality: a multicity, population-based, case-crossover analysis. Epidemiology 17, 315–323.
Stolwijk, A., Straatman, H., Zielhuis, G., 1999. Studying seasonality by using sine and cosine functions in regression analysis. Journal of Epidemiology and Community Health 53, 235–238.
Touloumi, G., Atkinson, R., Le Tertre, A., Samoli, E., Schwartz, J., Schindler, C., et al., 2004. Analysis of health outcome time series data in epidemiological studies. EnvironMetrics 15, 101–117.
Touloumi, G., Samoli, E., Pipikou, M., Le Tertre, A., Atkinson, R., Katsouyanni, K., 2006. Seasonal confounding in air pollution and health time-series studies: effect on air pollution effect estimates. Statistics in Medicine 25, 4164–4178.
van Houwelingen, H.C., Arends, L.R., Stijnen, T., 2002. Advanced methods in meta-analysis: multivariate approach and meta-regression. Statistics in Medicine 21, 589–624.
Welty, L.J., Zeger, S.L., 2005. Are the acute effects of particulate matter on mortality in the National Morbidity, Mortality, and Air Pollution Study the result of inadequate control for weather and season? A sensitivity analysis using flexible distributed lag models. American Journal of Epidemiology 162, 80–88.
Zanobetti, A., Schwartz, J., 2008. Temperature and mortality in nine US cities. Epidemiology 19, 563–570.
Zanobetti, A., Wand, M.P., Schwartz, J., Ryan, L.M., 2000. Generalized additive distributed lag models: quantifying mortality displacement. Biostatistics 1, 279–292.
Zeger, S.L., Dominici, F., Samet, J., 1999. Harvesting-resistant estimates of air pollution effects on mortality. Epidemiology 10, 171–175.
Zeger, S.L., Irizarry, R., Peng, R.D., 2006. On time series analysis of public health and biomedical data. Annual Review of Public Health 27, 57–79.
Zeger, S.L., Thomas, D., Dominici, F., Samet, J.M., Schwartz, J., Dockery, D., et al., 2000. Exposure measurement error in time-series studies of air pollution: concepts and consequences. Environmental Health Perspectives 108, 419–426.