Environment International 37 (2011) 204–209

Contents lists available at ScienceDirect

Environment International

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

Constrained confidence intervals in time series studies of mortality and air pollution

Borek Puza, Steven Roberts *, Mo Yang

School of Finance, Actuarial Studies and Applied Statistics, College of Business and Economics, Australian National University, Canberra ACT 0200, Australia

ARTICLE INFO

Article history:
Received 13 June 2010
Accepted 24 September 2010
Available online 15 October 2010

Keywords:
Air pollution
Confidence interval
Constraint
Mortality
Particulate matter
Time series

ABSTRACT

This paper focuses on constrained confidence intervals in the context of environmental time series studies where one seeks to ascertain the effects of ambient air pollution on human mortality. If the regression parameter representing such effects is non-negative, corresponding to a belief that more pollution cannot be beneficial, a desirable goal is to produce a constrained confidence interval for the parameter which is entirely non-negative. We show how this goal can be achieved using the method of tail functions. The proposed methodology is illustrated by the application to an environmental study of 100 cities in the United States involving regressions of mortality counts on levels of particulate matter air pollution. The large number of constrained CIs that contain zero is an indication that for the majority of the 100 cities there is not enough evidence to conclude a positive association between air pollution and mortality.

1. Introduction

The impact on human mortality and/or morbidity of exposure to ambient air pollution is a topic that has received extensive attention in the scientific literature (Bell et al., 2008; Liang et al., 2009; Moolgavkar, 2000; Pope et al., 1999). Studies investigating the association between daily time series of mortality and/or morbidity and daily time series of ambient air pollution concentrations have been at the forefront of this research. These studies have contributed to a better understanding of the health effects of exposure to ambient air pollution and added to the weight of evidence that has lead to stricter regulations (Bell et al., 2004). Amongst these studies some of the most influential have been the recent multi-city studies conducted in North America and Europe (Bell et al., 2006; Peng et al., 2005; Samoli et al., 2008, 2005). By combining estimates such studies have produced pooled estimates of the effect of pollution, both at national and regional levels. The pollutants of interest in these investigations have primarily been particulate matter air pollution and ozone.

Our focus here is on the use of time series studies to obtain estimates of the adverse health effects of ambient air pollution at the *individual* city level. The individual city-level effect estimates are the quantities that are combined in multi-city studies. Estimation of the adverse health effects of air pollution at the city level is a complex issue that involves disentangling the relatively small effect of air pollution from those of other confounding variables including temperature, humidity, and seasonality. Complicating the estimation process further are the myriad choices that need to be made regarding what confounders to include and how to include them. Some multi-city studies avoid these problems by fitting the same model relating air pollution and confounding variables to mortality within each city (Dominici et al., 2007; Roberts and Martin, 2006).

One aspect of the complicated process of estimating air pollution effects is that a 95% confidence interval (CI) for the "true" city-level effect often contains negative values and in some cases is entirely negative (Roberts and Martin, 2006). A problem of interpretation can arise if a CI containing negative values leads people to believe that air pollution might be beneficial to health. One way to address this problem is to exclude negative values and only report the truncated CI. However, this solution fails when the original CI is entirely negative. This motivates the need for a procedure that is guaranteed to lead to a CI which is entirely non-negative and non-empty.

A number of studies support the belief that air pollution cannot be beneficial to health. For example, in Vedal et al. (2003) it has been shown that low levels of air pollution are associated with increases in mortality, and other studies have stated that negative values are spurious or difficult to interpret (HEI, 2001; Murray and Nelson, 2000; Smith et al., 2000). The finding that even low levels of air pollution pose health risks would, for example, mean that hormesis (low concentrations of air pollution being beneficial to health) is unlikely. A few previous studies of air pollution and mortality have used constrained maximum likelihood estimation to ensure that the estimated effects of air pollution are non-negative (Roberts, 2006, 2004). In contrast, our focus here is on the construction of confidence intervals rather than point estimates.

Abbreviations: CI, confidence interval; HEI, Health Effects Institute; MCMC, Markov chain Monte Carlo; NMMAPS, National Morbidity, Mortality, and Air Pollution Study; PEL, prior expected length; PM, particulate matter air pollution; SE, standard error; US, United States.

* Corresponding author. School of Finance, Actuarial Studies and Applied Statistics, Australian National University, Canberra ACT 0200, Australia. Tel.: +61 2 6125 3470; fax: +61 2 6125 0087.

E-mail address: steven.roberts@anu.edu.au (S. Roberts).

0160-4120/$ – see front matter
doi:10.1016/j.envint.2010.09.004

In this paper we construct constrained CIs using the tail functions approach of Puza and O'Neill (2006a,b, 2008, 2009). This approach is attractive in that it can be "engineered" to provide an optimal CI in some regard, such as minimising prior expected length. Other methods could be used, but these have disadvantages which will be discussed below. The proposed method will be illustrated by application to environmental data from the United States (US). This illustration will highlight that constrained CIs provide interpretable results when commonly used (or "standard") methods do not. Finally, the method will be applied to 100 cities contained in the National Morbidity, Mortality, and Air Pollution Study (NMMAPS) database.

2. Materials and methods

2.1. Materials

The data used in this paper were obtained from the freely available NMMAPS database (http://www.ihapss.jhsph.edu/data/data.htm). This database contains daily mortality, air pollution and weather time series data for over 100 cities of the US for the period 1987–2000. For each city used in this investigation we extracted, where available, daily time series of mortality counts for individuals aged 65 years and over, the average concentration of particulate matter (PM) air pollution of less than 10 μm in diameter, measured in units of $\mu g/m^3$, and daily 24-hour average measures of temperature and dew point temperature.

2.2. Single-city estimation

To estimate the effect of PM on mortality within a single-city we utilize a Poisson log-linear model relating the mean number of daily deaths to daily PM and other confounding variables. The exact specification of the model used in this study is:

$$\begin{aligned}\log(\mu_t) = {} & \beta \mathrm{PM}_{t-2} + (\gamma_1 D_{1t} + \ldots + \gamma_6 D_{6t}) + s(\mathrm{temp}_t, df = 6) \\ & + s\left(\mathrm{temp}_{t,1-3}, df = 6\right) + s(\mathrm{dew}_t, df = 3) \\ & + s\left(\mathrm{dew}_{t,1-3}, df = 3\right) + s(t, df = 6 \times \text{ years}),\end{aligned} \tag{1.1}$$

where μ_t is the mean mortality count for day t, PM_{t-2} is the PM concentration on day $t-2$, D_{it} is the indicator variable for day t being day i of the week (e.g., $D_{5t}=1$ if day t is a Friday, and $D_{5t}=0$ otherwise), temp_t and dew_t are, respectively, the temperature and dew point temperature on day t,$\mathrm{temp}_{t,1-3}$ and $\mathrm{dew}_{t,1-3}$ are, respectively, the average temperature and average dew point temperature over the three days t, $t-1$ and $t-2$, and the functions $s()$ are natural cubic splines with the indicated degrees of freedom. The function $s(t, df = 6\times \text{ years})$ fits a smooth function of time that allows for slow-changing time trends in the mean mortality counts. Models of the same or similar form of Eq. (1.1) have been fitted in many studies of the association between PM and mortality (Dominici et al., 2007; Roberts and Martin, 2006). It is important to note that due to the frequency at which PM concentrations are recorded in most US cities the measure of PM used in Eq. (1.1) is restricted to a single-day's PM. In most US cities PM concentrations are recorded only "once every six days" meaning that if a PM concentration is recorded today that the next recorded PM concentration will not be for another six days. This "missingness" in the PM time series data means that, for example, fitting a single model that includes the lag-0, lag-1, and lag-2 PM concentrations simultaneously is not possible because there are typically "gaps" of six days between recorded PM concentrations. For reasons of illustration we focus on the lag-2 PM concentration here because it provides examples where the standard CIs are entirely negative.

The estimated effect of PM on mortality obtained from Eq. (1.1) for the city in question, $\hat{\beta}$, and its associated standard error, $\mathrm{SE}(\hat{\beta})$, are then used for inference on the effect of PM on mortality in that city, β. There are a number of papers that produce figures depicting the value of $\hat{\beta}$, along with an associated 95% CI, $\left[\hat{\beta} \pm 1.96 \times \mathrm{SE}\left(\hat{\beta}\right)\right]$ (which may be termed the "standard" CI), for each of a number of cities (Dominici et al., 2002; Roberts and Martin, 2006). As mentioned above, it is common for these 95% CIs to contain negative values. If we believe that β must be non-negative then it may be desirable that the intervals created based on $\hat{\beta}$ only contain plausible values, that is, only contain values greater than or equal to zero. In the next section we describe a method that will produce 95% CIs for β that by construction contain only plausible values. Note: The value of 1000β may be interpreted as approximately the percentage increase in the number of persons who will die on average per day if there is an increase of 10 $\mu g/m^3$ in PM two days earlier.

2.3. The tail functions approach

The tail functions approach is a recently developed methodology that can be used for constructing CIs (Puza and O'Neill, 2009, 2008, 2006a,b). Briefly, this approach involves a "twisting" of the confidence bounds obtained using the standard approach, whilst ensuring that the desired coverage properties of the CI are preserved. To apply this approach in the present context, we begin with the fact that $\hat{\beta}$ is approximately $N(\beta, \delta^2)$, where $\delta = \mathrm{SE}(\hat{\beta})$. This fact leads to the standard 95% CI, $\left[\hat{\beta} - 1.96\delta, \hat{\beta} + 1.96\delta\right]$, via the identity

$$0.95 = P\left(0.025 \le \Phi\left(\frac{\hat{\beta}-\beta}{\delta}\right) \le 0.975\right) = P\left(\hat{\beta} - 1.96\delta \le \beta \le \hat{\beta} + 1.96\delta\right),$$

where $Z = \left(\hat{\beta}-\beta\right)/\delta \sim N(0,1)$ and $\Phi(z) = P(Z \le z)$, so that $\Phi(Z) \sim U(0,1)$. The lower and upper bounds of the standard CI are the solutions in β of the equations $\Phi\left(\left(\hat{\beta}-\beta\right)/\delta\right) = 0.975$ and $0.025 = \Phi\left(\left(\hat{\beta}-\beta\right)/\delta\right)$, respectively.

Next, consider any non-decreasing function, $\tau(\beta)$, which has a range in the interval [0,1]. Then observe that for any value of β it is also true that

$$0.95 = P\left(0.05\tau(\beta) \le \Phi\left(\frac{\hat{\beta}-\beta}{\delta}\right) \le 0.95 + 0.05\tau(\beta)\right) = P\left(L\left(\hat{\beta}\right) \le \beta \le U\left(\hat{\beta}\right)\right),$$

where the bounds $L\left(\hat{\beta}\right)$ and $U\left(\hat{\beta}\right)$ are the solutions in β of the two equations $\Phi\left(\left(\hat{\beta}-\beta\right)/\delta\right) = 0.95 + 0.05\tau(\beta)$ and $0.05\tau(\beta) = \Phi\left(\left(\hat{\beta}-\beta\right)/\delta\right)$, respectively. We call $\tau(\beta)$ the "tail function." Note that the associated CI, $\left[L\left(\hat{\beta}\right), U\left(\hat{\beta}\right)\right]$, is a generalization of the standard CI as defined by the "standard" tail function, $\tau(\beta) = 1/2, -\infty < \beta < \infty$. Typically, the calculations required to obtain the bounds defined by a particular tail function $\tau(\beta)$ can be performed easily using the Newton–Raphson algorithm.

In the context where β is known to be non-negative, one suitable tail function $\tau(\beta)$ has the form:

$$\tau(\beta) = \begin{cases} 0, & \beta < 0 \\ \beta/(2c), & 0 \le \beta \le c \\ 1/2, & \beta > c. \end{cases} \tag{1.2}$$

Fig. 1a shows the standard tail function, $\tau(\beta) = 1/2, -\infty < \beta < \infty$, and the alternative tail function (Eq. (1.2)) for $c=1$ and $c=2$, respectively. Fig. 1b shows the associated 95% confidence bounds, calculated taking $\delta = 1$ (without loss of generality). It will be observed that, after elimination of all negative values of β (or "truncation"), the alternative CI exists (i.e. is non-empty) for all possible values of $\hat{\beta}$, whereas the standard CI is empty for values less than -1.96.

For example, if $\hat{\beta}=-3$, the standard CI is $[-3-1.96,-3+1.96]=[-4.96,-1.04]$, which finally becomes empty after truncation. By contrast, the alternative CI works out as $[-4.645,0.046]$ for $c=1$ and $[-4.645,0.082]$ for $c=2$. These two CIs finally become $[0,0.046]$ and $[0,0.082]$, respectively, after truncation. We see that the truncated standard CI is empty for all $\hat{\beta}<-1.96\delta$. Also, for any chosen value of c, the truncated alternative CI has the attractive properties of: (a) never being empty, (b) becoming narrower and more focused near zero as $\hat{\beta}$ decreases, and (c) approaching the standard CI as $\hat{\beta}$ increases. In more detail, the lower bound of the truncated alternative CI is zero for $\hat{\beta}<1.645\delta$ and identical to the lower bound of the standard CI for $\hat{\beta}>c+1.96\delta$. Also, the upper bound of the truncated alternative CI is identical to the upper bound of the standard CI for $\hat{\beta}>c-1.96\delta$.

In the application below, a suitable value of c is chosen, utilizing information from all 100 cities, so as to approximately minimize the prior expected length (PEL) of the proposed alternative CI. See the Supplemental material for details. Note that the exact value of c used is not critical, as any value will lead to a "proper" CI, meaning one with the desired coverage probabilities. Also see the Supplemental material for additional notes on the tail functions method and how it preserves coverage.

2.4. Illustration

Here we highlight the similarities and differences between the standard 95% CI and the alternative 95% CI to three cities of the US. The cities Los Angeles, Tacoma, and Syracuse were selected because they correspond to cases where the standard 95% CI contains: only non-negative values (Los Angeles); both negative and non-negatives values (Tacoma); and only negative values (Syracuse). The standard CI for these three cities spans the range of possibilities in terms of the coverage of positive and negative values.

Using Eq. (1.1) and the tail functions methodology with $c=0.0003$ (see the Supplemental material for details of how this value was obtained), 95% CIs for the lag-2 effect of PM on mortality were produced using both the standard and alternative (tail functions) approaches discussed above. Fig. 2 shows the CIs produced by each method for each city. It will be noted that for Los Angeles the standard and alternative intervals are identical. This is a desirable property of the tail functions approach. It is reasonable that when the standard CI contains only plausible values that the alternative CI will be similar or the same. This property is evident in Fig. 1b which shows that the standard and alternative CIs coincide exactly when the estimate is sufficiently large (at least $c+1.96\delta$). In the case of Los Angeles, $\hat{\beta}=0.0009173$, which is larger than $c+1.96\delta=0.0003+1.96\times0.0002054=0.0007026$.

In the case of Tacoma, both the standard and alternative CIs straddle zero, so that both intervals are finally from zero up to an upper bound, following elimination of all implausible values. The utility of the tail functions approach is made most clear by comparing the standard and alternative 95% CIs for Syracuse. Here the standard CI is entirely negative and the alternative CI straddles zero. Thus, the final alternative CI is from zero up to an upper bound, whereas the final standard CI is empty. The results for Syracuse clearly illustrate why the process of truncating the standard CI is untenable.

The interpretation of the truncated alternative intervals is intuitive. A truncated CI that contains zero as its lower bound suggests that there is not enough evidence to conclude a positive association between air pollution and mortality, whilst an entirely positive interval constitutes evidence for a positive association. The advantage of truncating a CI is that it avoids the problems of interpretation that may arise if negative values are taken literally. The main advantage of the truncated alternative CI over the truncated standard CI is that it guarantees that the final interval will not be empty, whilst preserving the desired 95% coverage probability.

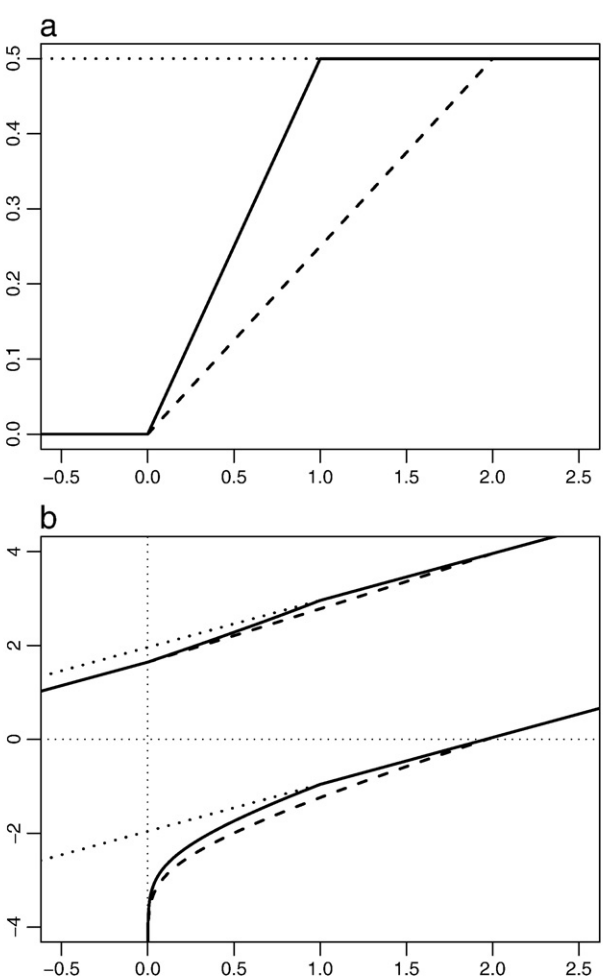

Fig. 1. Examples of (a) tail functions and (b) the confidence bounds implied by the tail functions in a. The solid lines correspond to tail function (Eq. (1.2)) with $c=1$, the dashed lines to Eq. (1.2) with $c=2$, and the dotted lines to the standard tail function, namely the constant (1/2).

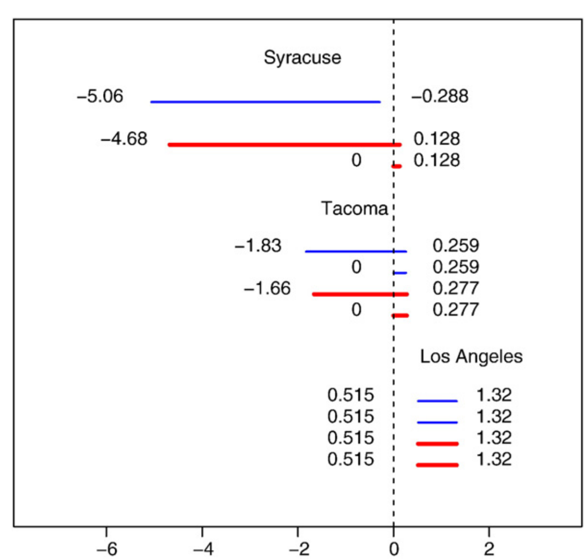

Fig. 2. Confidence intervals for the effect of lag-2 PM on mortality (1000β) for three selected cities. For each city the four CIs, top to bottom are: untruncated standard, truncated standard, untruncated alternative and truncated alternative. The blue and red intervals correspond to the standard and alternative intervals, respectively. For Syracuse the truncated standard CI is empty.

3. Application

Here we apply the methods illustrated in the previous section to 100 of the US cities contained in the NMMAPS database so as to obtain both an untruncated standard CI and a truncated alternative CI for each city. In Figs. 3 and 4 it can be seen that many of the standard CIs contain negative values. This could lead to problems of interpretation, such as the belief that increases in air pollution might be beneficial to health. This problem is avoided by the truncated alternative CIs. A natural interpretation of the large number of truncated alternative CIs containing zero is that for a majority of cities there is not enough evidence to conclude a positive association between air pollution and mortality.

Fig. 4 also displays a pooled estimate and 95% CIs (standard and truncated alternative) for the "overall" effect of air pollution on mortality, where this effect is defined as the weighted average of the 100 regression parameters obtained from Eq. (1.1) applied to each city, $\hat{\beta}_1, \ldots, \hat{\beta}_{100}$, with the weights being proportional to the cities' populations. These CIs for the overall effect were obtained using exactly the same theory as for each individual city effect, β_i, making use of the fact that the weighted average of the 100 regression parameter estimates is normal with mean equal to the overall effect and a known variance. As might be expected, the pooled standard CI (0.0000476, 0.0003352) and alternative CI (0.0000666, 0.0003352) are similar, with an identical upper bound. (Here there is no truncation, since both lower bounds 0.0000476 and 0.0000666 are already positive.)

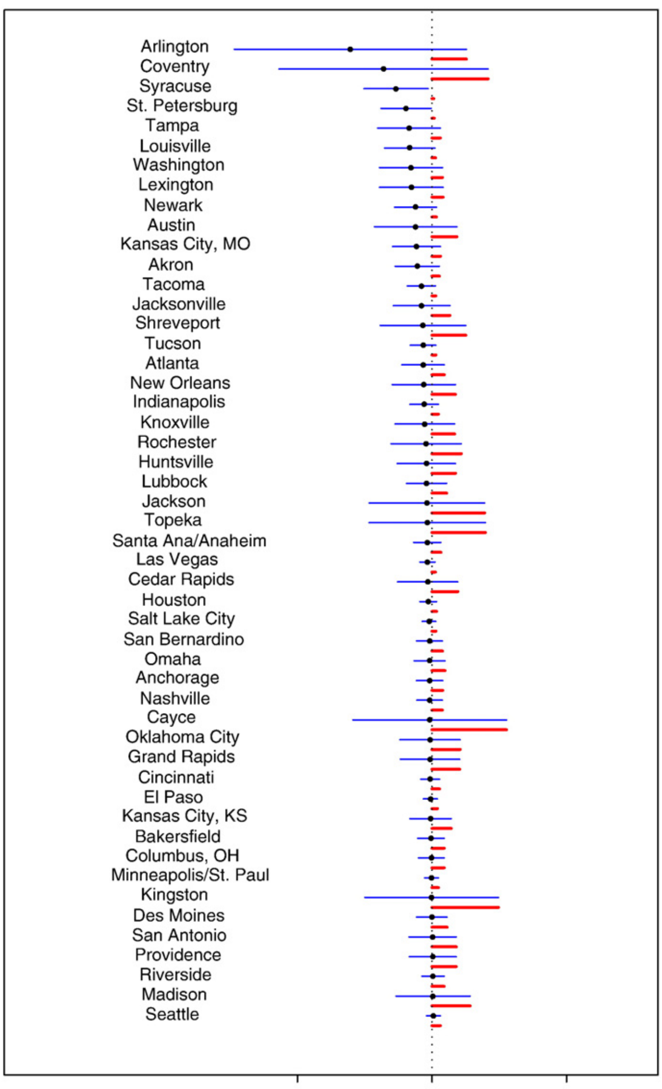

Fig. 3. Confidence intervals for the effect of lag-2 PM on mortality (1000β) for 50 of the 100 US cities considered. These 50 cities have the smallest values of $\hat{\beta}$. For each city, the two CIs, top to bottom, are: untruncated standard (blue interval) and truncated alternative (red interval).

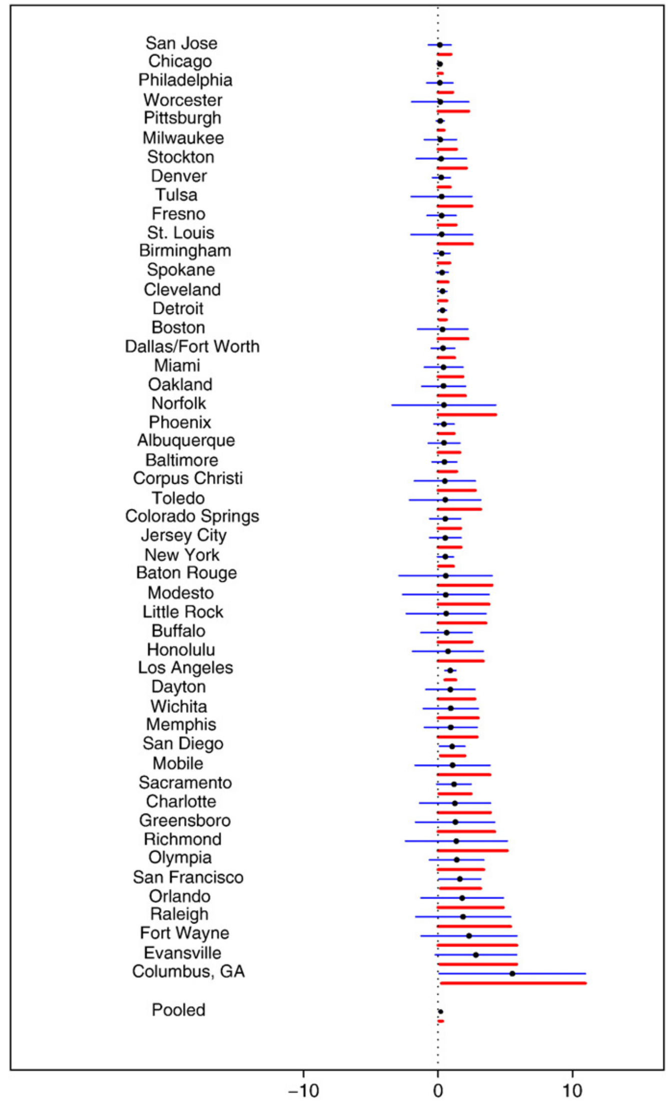

Fig. 4. Confidence intervals for the effect of lag-2 PM on mortality (1000β) for 50 of the 100 US cities considered. These 50 cities have the largest values of $\hat{\beta}$ For each city, the two CIs, top to bottom, are: untruncated standard (blue interval) and truncated alternative (red interval).

4. Discussion

The tail functions approach to CI estimation is a generalisation of the standard approach, as implicitly defined by the constant tail function equal to 1/2. Two other constant tail functions, equal to 0 and 1, respectively, lead to one-sided CIs, and the tail functions approach may be thought of as a "graduated blend" of the intervals defined by these two extremes. The choice of tail function is to some extent arbitrary, since any tail function will lead to a "proper" CI, meaning one that has the desired frequentist coverage probability for all possible values of the target parameter. However, some tail functions may be preferable to others in the presence of prior information, since they lead to CIs with more attractive properties.

In the present context of inference on a regression parameter, β, the prior information comes in two forms. First, we believe that β is non-negative, and for this a suitable tail function, $\tau(\beta)$, is one which equals 0 for values less than 0 and which continuously approaches 1/2 as β tends to infinity. A convenient example of $\tau(\beta)$ with these properties is Eq. (1.2), namely the function that is 0 for $\beta<0$, 1/2 for $\beta>c$, and linear between the points (0,0) and (c,1/2). The question of how the tuning constant c should be chosen can be addressed using the remaining prior information, for example to select the value which minimises prior expected length (see the Supplemental material). Although other classes of tail functions could be considered, such as one with two tuning parameters defining a *curved* line between (0,0) and (c,1/2), it is not clear that any improvements would be worth the added computational burden.

The availability of prior information suggests that a Bayesian approach could also be adopted for the construction of constrained CIs. In the present context this approach leads to posterior intervals that are guaranteed to exist, contain no negative values, and moreover have a smaller prior expected length than the corresponding tail functions CIs. However, under the informative prior used in this paper, such posterior intervals have a frequentist coverage probability that falls severely below the desired 95% for extreme values of β. If the uninformative prior $f(\beta)\propto 1$, $\beta\geq 0$, is used instead, the coverage probability of the corresponding posterior intervals is greatly improved but still falls distinctly below 95% over a range of β values, and in that case the prior expected length of those intervals becomes very comparable to that of the tail functions CI. Moreover, under both priors the associated posterior intervals have an actual length which converges to zero at a much slower rate than the tail functions CI as $\hat{\beta}$ tends to minus infinity. For a more detailed account of these comparisons, see the Supplemental material. Also see Puza and O'Neill (2006b) for a related comparison regarding inference on an unconstrained normal mean. In any case, the tail functions approach provides an alternative to the Bayesian approach when prior information is available and it is deemed important that a CI be constructed with a frequentist coverage probability that is guaranteed to be at least 95%, as is the case in the present context.

Another way to perform CI estimation is via the unified approach of Feldman and Cousins (1998). This involves inverting a likelihood ratio test and leads to confidence bounds similar to those depicted in Fig. 1b. The unified approach may be thought of as another special case of the tail functions approach, with an implicit tail function that could be calculated and added to Fig. 1a (see Puza and O'Neill, 2008 for an example). However, the unified approach typically does not provide the optimal constrained CI, at least in terms of prior expected length. Yet other approaches have been considered, as in Mandelkern (2002), Silvapulle and Sen (2004) and Fraser et al. (2004), but each of these has its own problems, such as being only approximate or not leading to a “proper” confidence interval.

The approach proposed in this paper is equally applicable to other ambient air pollutants and/or to other end-points such as counts of morbidity or cause-specific mortality, the only difference being that the value of the tuning constant c would need to be re-calculated. In any particular situation, the benefits of the proposed approach will depend on the likelihood that the standard CIs for the location under investigation will contain negative values, remembering that the standard and alternative CIs are essentially the same when both intervals are strictly positive. The likelihood that a standard CI will contain negative values will be a function, among other things, of the number of days of data available and the size of the observed end-point counts.

The limitations of the proposed approach need to be acknowledged. Firstly, this approach is a more complicated process, both conceptually and computationally, than the straightforward standard approach. Secondly, the proposed approach may not “add value” if the only goal of an investigation is to produce a single pooled estimate of the effect of air pollution based on a number of locations. The reason for this is that the process of pooling information across cities results in a more powerful estimation process and as a result the CIs for the pooled estimate are typically much narrower than those for any individual city and strictly positive. Thirdly, the tail functions approach does not provide an avenue for improved point estimation under constraints, as discussed, for example, in Tripathi and Kumar (2007). Finally, the truncated tail functions approach should only be applied if there is strong justification for the assumption of the parameter of interest being non-negative.

Acknowledgement

Roberts was partially supported by the Australian Research Council [DP0878988].

Appendix A. Supplementary data

Supplementary data to this article can be found online at doi:10.1016/j.envint.2010.09.004.

References

Bell ML, Samet JM, Dominici F. Time-series of particulate matter. Annu Rev Public Health 2004;25:247–80.
Bell ML, Peng RD, Dominici F. The exposure–response curve for ozone and risk of mortality and the adequacy of current ozone regulations. Environ Health Perspect 2006;114:532–6.
Bell ML, Ebisu K, Peng RD, Walker J, Samet JM, Zeger SL, et al. Seasonal and regional short-term effects of fine particles on hospital admissions in 202 US counties, 1999–2005. Am J Epidemiol 2008;168:1301–10.
Dominici F, Daniels M, Zeger SL, Samet JM. Air pollution and mortality: estimating regional and national dose–response relationships. J Am Stat Assoc 2002;97: 100–11.
Dominici F, Peng RD, Zeger SL, White RH, Samet JM. Particulate air pollution and mortality in the United States: did the risks change from 1987–2000. Am J Epidemiol 2007;166: 880–8.
Feldman GJ, Cousins RD. Unified approach to the classical statistical analysis of small signals. Phys Rev D 1998;57:3873–89.
Fraser DAS, Reid N, Wong ACM. Inference for bounded parameters. Phys Rev D 2004;69, doi:10.1103/PhysRevD.69.033002.
Health Effects Institute. Airborne particles and health: HEI epidemiologic evidence. Cambridge, MA: Health Effects Institute; 2001.
Liang WM, Wei HY, Kuo HW. Association between daily mortality from respiratory and cardiovascular diseases and air pollution in Taiwan. Environ Res 2009;109:51–8.
Mandelkern M. Setting confidence intervals for bounded parameters. Stat Sci 2002;17: 149–72 (with discussion).
Moolgavkar SH. Air pollution and daily mortality in three U.S. counties. Environ Health Perspect 2000;108:777–84.
Murray CJ, Nelson CR. State-space modelling of the relationship between air quality and mortality. J Air Waste Manage Assoc 2000;50:1075–80.
Peng RD, Dominic F, Pastor-Barriuso R, Zeger SL, Samet JM. Seasonal analyses of air pollution and mortality in 100 US cities. Am J Epidemiol 2005;161:585–94.
Pope III CA, Hill RW, Villegas GM. Particulate air pollution and daily mortality on Utah's Wasatch Front. Environ Health Perspect 1999;107:567–73.
Puza B, O'Neill T. Generalised Clopper–Pearson confidence intervals for the binomial proportion. J Stat Comput Simul 2006a;76:489–508.
Puza B, O'Neill T. Interval estimation via tail functions. Can J Stat 2006b;34:299–310.
Puza B, O'Neill T. Optimal constrained confidence estimation via tail functions. Math Sci 2008;33:134–40.
Puza B, O'Neill T. Constrained confidence estimation of the binomial p via tail functions. Math Sci 2009;34:43–8.
Roberts S. Biologically plausible particulate air pollution mortality concentration–response functions. Environ Health Perspect 2004;112:309–13.
Roberts S. A new model for investigating the mortality effects of multiple air pollutants in air pollution mortality time-series studies. J Toxicol Environ Health 2006;69: 417–35.
Roberts S, Martin MA. Applying a moving total mortality count to the cities in the NMMAPS database to estimate the mortality effects of particulate matter air pollution. Occup Environ Med 2006;63:193–7.

Samoli E, Analitis A, Touloumi G, Schwartz J, Anderson HR, Sunyer J, et al. Estimating the exposure–response relationships between particulate matter and mortality within the APHEA multicity project. Environ Health Perspect 2005;113:88–95.
Samoli E, Peng R, Ramsay T, Pipikou M, Touloumi G, Dominici F, et al. Acute effects of ambient particulate matter on mortality in Europe and North America: results from the APHENA study. Environ Health Perspect 2008;116:1480–6.
Silvapulle MJ, Sen PK. Constrained statistical inference. Hoboken, New Jersey: Wiley; 2004.
Smith RL, Davis JM, Sacks J, Speckman P, Styer P. Regression models for air pollution and daily mortality: analysis of data from Birmingham, Alabama. Environmetrics 2000;11:719–43.
Tripathi YM, Kumar S. Estimating a positive normal mean. Stat Pap 2007;48:609–29.
Vedal S, Brauer M, White R, Petkau J. Air pollution and daily mortality in a city with low levels of pollution. Environ Health Perspect 2003;111:45–51.