

1. Introduction

A series of historical air pollution catastrophes in the mid twentieth century provided the first evidence that high levels of pollution can cause extraordinarily large increases in the death rate (1, 2). Today, air pollution levels are much lower, but numerous studies have shown pollution to be associated with small increases in the rates of morbidity and mortality in cities throughout the developed world (3-16) . Despite the fact that the associations observed are small, the burden of disease attributable to air pollution may be substantial when one considers the size of the population that is exposed to the pollution (17).

Study of the health effects of air pollution is fraught with difficulties, including error in the measurement of exposure, the discrepancy between a shared ambient exposure series and personal exposure, and confounding by time-independent and time dependent factors. Biases in the estimated association between exposure and outcome are especially troublesome due to the fact that the effects are usually very small, and the effects of confounders many times larger. The case-crossover design represents a novel approach to controlling confounders in a study of the effect of short-term air pollution exposure on the risk of various adverse health events. With a case-crossover design, only cases are required, and, for each case, exposure at the event time is compared with exposure at comparable control, or “referent” times (18). This is similar to a matched case-control study; here, the exposure at each event time is part of a “matched set” of exposures consisting of exposures for the same subject at his/her referent times. The case-crossover design is well suited to analysis of air pollution exposures because not only does it

control for fixed confounders by making within-person comparisons, but, with careful selection of referent times, it also controls for time dependent confounders by design, rather than by modeling.

In a case-crossover study, referents are matched on important confounders in order to control for their effects. Yet, there are other reasons to be careful about referent selection. The case-crossover design involves an implicit assumption that the distribution of exposure is constant across referent times, and that the referent exposures used are representative of the exposure that generated the event. Moreover, the statistical method typically used to analyze a case-crossover study, conditional logistic regression, is only unbiased with certain referent schemes. For these reasons, it is important to choose a valid referent selection strategy.

A wide variety of different referent schemes have been employed in air pollution exposure studies (19-25). For example, Lee and Schwartz (20) analyzed the association between SO₂ levels and mortality in Seoul, Korea. The pollution level on the day of death was compared to the levels on days seven and 14 days before and after death, an example of symmetric bidirectional referent selection. In another study, Levy et al. (21) considered the effect of particulate matter (PM) exposure on the risk of out-of-hospital primary cardiac arrest in subjects without a recognized history of heart disease. The referent selection strategy used was time-stratification; referents were all days on the same day of the week and in the same month and year as the event day. In a third example, Peters et al. (24) studied the association between elevated concentrations of fine

particles and the risk of acute myocardial infarction (MI). Their case-crossover study used a unidirectional referent sampling approach with referents three, four, and five days prior to the index day.

Quite a few different methods papers have appeared which consider referent selection in the case-crossover design (22, 26-35), but as of yet they have not presented a cohesive set of guidelines for choosing and evaluating a referent selection strategy. The purpose of this paper is to review the case-crossover literature in the air pollution context, as well as to clarify key referent selection issues. We begin in Section 2 with an overview of the case-crossover method. We present the underlying statistical model, the estimation method commonly used, and discuss aspects of the model that are specific to air pollution studies. In Section 3, we describe the various referent selection strategies that have been used, and discuss how each of the strategies relates to the issues of confounding, bias due to trend in the exposure series over time, and a specific source of bias in the estimating equations, usually called “overlap bias”. Section 4 extends our review to situations in which the exposure is not treated as fixed (in contrast to all previous sections). In Section 5, we discuss other important aspects of the case-crossover design. Finally, we conclude with recommendations for choosing a referent selection strategy for a case-crossover analysis of air pollution exposure data.

2. The Case-Crossover Method

The case-crossover design requires a set of subjects who have experienced an event. For each case, the exposure at the index time, when the event was triggered, is compared with

exposure at the referent times, when the event was not triggered. The design is appropriate for assessing the association between a short-term exposure and the risk of an acute event. It is also most suited to exposures that are transient and do not have carryover effects. In this review, we focus on the case of a rare event (see Section 6 for a discussion of methods developed for more common outcomes).

The main strengths of the case-crossover design are that it does not require a control sample (and hence avoids bias associated with improper control selection), makes effect modification assessment relatively simple, controls for fixed confounders, and, with proper selection of referents, can control for time dependent confounders by design. Time dependent confounding is a huge concern in air pollution studies, since there are strong effects such as season and day of the week that influence the exposure and are also associated with many outcomes of interest. Indeed, the effect of exposure is thought to be orders of magnitude smaller than the effects of confounding factors (Dominici F, McDermott A, Hastie T; Issues in semi-parametric regression: A case study of time series models for air pollution and mortality; Department of Biostatistics; Johns Hopkins University; unpublished manuscript), and thus there is always a concern that observed associations merely represent residual confounding. Restricting referents to the same day of the week and season as the index time facilitates control over time dependent confounders by design. Controlling for confounders in this way means that they no longer have to be included in the statistical model.

The underlying statistical model for the case-crossover design is typically the proportional hazards model for a rare disease with a constant baseline hazard for each individual. The hazard rate of person i at time t given time-varying covariates x_{it} is given by $\lambda_i(t; x_{it}) = \lambda_i \exp(x_{it}\beta)$. Over a short time period, the assumption of a constant baseline hazard is often reasonable, and is equivalent to assuming smooth seasonal effects in a time series analysis. We can interpret the parameter β in this model as the change in the risk of an event associated with a short-term unit increase in exposure.

Conditional logistic regression methods have typically been used for analysis of the case-crossover design. The use of this method has been motivated by the analogy to matched case-control designs, where conditional logistic regression maximizes the true conditional log-likelihood of the data. Computationally, this is achieved by finding the value of the parameter β that makes the derivative of the log-likelihood equal to zero; hence, the derivative of the log-likelihood is called an estimating equation. Yet, this equation can be used to obtain an estimate of β even when it is not the derivative of the log-likelihood.

When the estimating equation is the derivative of the true log-likelihood, we call it a score equation, and we know that it will have mean zero and lead to valid estimates.

Estimating equations that are not score equations will typically still be valid, if they have mean zero (used, for example, in analyses of case-cohort and two-phase case-control designs), but if they don't have mean zero, their estimates will be biased. The bias in the parameter estimates is proportional to the mean of the estimating functions. For this reason, we refer to the mean of the estimating functions for various approaches to designing and analyzing case-crossover studies. If, and only if, the estimating equations

have mean zero (i.e. they are unbiased) we expect the β estimates to be valid. In situations in which the estimating equations are not score equations, further analysis is needed to determine whether they are unbiased.

With a case-crossover design, the underlying likelihood of the data depends on the referent selection strategy. Only with certain referent selection strategies are the conditional logistic regression estimating equations score equations, and hence unbiased. For most of the commonly used referent selection strategies (e.g. symmetric bidirectional referents), the conditional logistic regression estimating equations are *not* score equations, and can be shown to have nonzero mean (see (31, 36)). We call this bias in the estimating equations “overlap bias”; in Section 3, we discuss which referent strategies are subject to overlap bias.

The case-crossover design assumes that the referent exposures are representative of the distribution of exposure that generated the event. Greenland (28) and Navidi (22) also showed that the design depends on the assumption that the distribution of exposure is stationary, or constant, across the referent times. We discuss the merits of the various referent strategies with respect to this assumption in Section 3.

Exposure data for air pollution studies frequently come from centrally located monitors of ambient pollution. In this case ambient air pollution is common to all individuals. Thus, research on the case-crossover design in the air pollution exposure context should focus on properties of the effect estimate conditional on a known and shared exposure

series. We want to know that the estimate and the corresponding estimating equations are unbiased for *a given* exposure series; knowing that they are unbiased when averaged over all possible exposure series is of little value when there is a single exposure in any given study. With the exception of a review of case-crossover methods that do not condition on exposure in Section 4, we restrict our attention to models conditional on exposure.

3. Referent Selection Strategies

Figure 1 displays two different, commonly used referent sampling strategies (see (30, 36) for diagrams of a wider variety of possible referent selection strategies). Figure 1A shows an example of a restricted unidirectional referent sampling design. With this design, referents are at particular lags pre-event. Restricting referents to be close to the index day reduces confounding due to season and ensures that the exposure distribution is approximately stationary across the referent times. Figure 1A shows referents placed seven, 14, and 21 days prior to the index day in order to additionally control for day-of-week confounding.

Unidirectional sampling has a major disadvantage in the air pollution context. If the exposure series has a long-term time trend, choosing referents only prior to the index time may lead to bias. As the referents become further removed in time from the index time, there is more opportunity for the assumption that the exposure distribution is stable over time to fail. Greenland (28) was the first to recognize this pitfall with unidirectional referent selection for air pollution exposures. Navidi (22) proposed that the time trend

bias could be eliminated by choosing referents both before and after the index time, a strategy called *bidirectional* referent selection (sometimes called *ambidirectional* selection). Technically, bidirectional sampling is only valid when cases are still at risk after an event, an assumption that is certainly not valid when the event is death. Navidi justified bidirectional sampling by noting that, with air pollution data, the exposure is exogenous (i.e. not dependent upon individual qualities or behaviors) and is available for all cases both before and after the event. A more rigorous justification was given by Lumley and Levy (31) who showed that, with a rare event, the bias due to sampling referents after the at-risk period is small. Even more important is the fact that the bias associated with sampling referents outside of the time at risk is smaller than the bias that would be incurred with unidirectional referent selection in the presence of a time trend. Despite the fact that this result is not widely appreciated, concerns over long term time trends have led to almost universal use of bidirectional referent strategies in air pollution studies.

Navidi introduced full stratum bidirectional referent selection, where the referents are all days in the exposure series other than the index day (22). It can be shown that a case-crossover analysis with full stratum bidirectional referents is equivalent to a Poisson regression analysis (31). Thus, with this strategy, season and day of the week must be modeled or the air pollution effect estimates will be confounded. A popular alternative to modeling is symmetric bidirectional sampling (26) where referents are selected at fixed intervals before and after the index time (see e.g. (19, 20, 23, 25) for applications). This referent selection strategy controls for bias due to time trend, season, and day of the week

(if referents are multiples of seven) (26). Simulation studies have shown that shorter time periods for the referents (e.g. seven or seven and 14 days before and after the index day) result in less confounding bias, and that confounding is not as well controlled if the seasonal pattern of exposure is not symmetric (26, 29). Figure 1B shows the symmetric bidirectional design with four referents at seven and 14 days before and after the index time.

Both of the referent strategies shown in Figure 1 are examples of *event-dependent referent window schemes* (36). For each of these strategies, the referent windows are completely determined by the event times. Two referent selection strategies have been proposed which differ in this respect; for the time stratified (31) and semi-symmetric bidirectional (34) designs, the referent windows are no longer completely determined by the event times. These two designs are examples of *event-independent referent window schemes*.

With time stratified referent selection, time is broken into disjoint strata *a priori*, and, for each event, all or a sample of the days in the stratum in which the event falls serve as referents (31). In this way, the referents are specified *a priori*. The time stratified design is not subject to bias due to time trend since there is no pattern in the placement of referents relative to the index time. This strategy maintains control over confounders by design by ensuring that referents are matched on important confounders; restricting referents to the same day of the week and in the same month and year as the index day controls for season and day of the week effects. A time stratified analysis with this

stratification and a shared exposure series is the same as a Poisson regression model with dummy variables to adjust for month within each year, and day of the week (30). The full stratum bidirectional design is a special case of a time stratified design in which there is one stratum, although in this case confounders cannot be assumed to be constant across the stratum (31).

The semi-symmetric bidirectional design was proposed by Navidi and Weinhandl (34). With this strategy, for each case, one referent is randomly chosen from days at a fixed lag pre- and post-event; if only one is available (e.g. at the beginning or end of the series), it serves as the referent. The random selection of referents ensures that referent windows are not completely determined by the event times, and hence this is an **event-independent** referent design. If the lag is small and a multiple of seven, confounding effects of season and day of the week can be controlled, and there is no bias due to time trend.

The distinction between **event-independent and event-dependent** referent strategies is an important one. It relates to the idea of overlap bias (see (36) for a more rigorous derivation of overlap bias). The likelihood of the index times conditional on the referent windows is typically used to motivate the use of the conditional logistic regression estimating equations. In the case of event-dependent referent windows, the referents and index times are simple functions of one another (e.g., with symmetric bidirectional referents at lag λ , the referent window is $W_i = \{t_i - \lambda, t_i, t_i + \lambda\}$). Hence the likelihood of the index times conditional on the referent windows is uninformative. The likelihood that is informative about β , in the case of event-dependent referent windows, is the

(unconditional) likelihood of the referent windows. The derivative of this log-likelihood is not the conditional logistic regression estimating equations, and hence the conditional logistic estimating equations are not score equations. It can be shown that conditional logistic regression estimating equations do not have mean zero. Given an exposure series, the magnitude of the overlap bias associated with a specific referent selection strategy can be calculated (see (36)). The bias is generally small, but it is highly unpredictable; simulation studies have shown that it depends on the particular exposure series (36). It can exist even for small β (the magnitude typically seen in air pollution studies), and, indeed, even when $\beta = 0$. For a given exposure series, there may be bias with some referent strategies and not with others. There is no way to predict in advance the magnitude of the overlap bias.

In contrast, with a time stratified design (of which the full stratum bidirectional design is a special case), referents are specified *a priori*, and are not functions of the index times. The likelihood of the index times conditional on the referent windows is informative for β , and the derivative of this log-likelihood is exactly the conditional logistic regression estimating equations. Hence, these are score equations and have mean zero. There is no overlap bias associated with this design.

In the case of semi-symmetric bidirectional referent selection, the likelihood of the index times conditional on the referent windows is also informative, but the derivative of this log-likelihood is not the conditional logistic regression estimating equations. Since the mean of the conditional logistic regression estimating equations is not zero, there is

overlap bias associated with a standard conditional logistic regression analysis. The likelihood specified by Navidi and Weinhandl must be used to obtain an unbiased estimate of β . It turns out, however, that using conditional logistic regression with an offset term which takes a value of $\log 2$ for cases with only one possible referent and zero otherwise yields the same estimates (36).

Another common characteristic of the **event-dependent** referent window designs is that they are subject to edge effects (27). Edge effects occur when cases at the beginning or end of the exposure series have a different referent selection strategy than others; this occurs when referents at the edges are not available. For example, with the symmetric bidirectional design, cases at the beginning of the series will not have pre-event referents, and cases at the end will not have post-event referents. Such cases could be dropped from the analysis, or all available referents could be used. Edge effects also contribute to bias in the conditional logistic regression estimating equations. Referent strategies based on *a priori* referents, such as the full stratum bidirectional and time stratified designs, are not subject to edge effects since all cases have the same referent scheme. **In the semi-symmetric bidirectional design, the offset term in the conditional logistic regression analysis is needed to compensate for the edge effects** (36). For event-dependent referent schemes, the magnitude of the bias due to edge effects has not been studied independently, but it is likely that the impact of a few cases at the ends of the exposure series is small, especially for long exposure series. **When $\beta = 0$, all of the overlap bias associated with these designs is due to edge effects; this is not true for other values of β .**

Bateson and Schwartz (28) suggested subtracting off the edge effects at $\beta = 0$ to correct for overlap bias. This correction just accounts for overlap bias at $\beta = 0$.

Though bias is the dominant concern in air pollution studies, efficiency is also relevant. Mittleman et al. (33) assessed the statistical efficiency of a variety of referent strategies. They showed that, as expected, increasing the number of referents leads to more precise estimation, and that, in contrast to matched case-control studies, this pattern continues with greater than five referents. Autocorrelation in the exposure series also affects efficiency. It is clear that referents close to the index time will be correlated with the index exposure, and hence there will be less power to detect an exposure effect. Thus, it is advisable to choose referents that are at least, say, six days removed from the index time in order to maximize the information contributed by each referent. It should be noted that the issue with autocorrelation is not “overmatching” in the traditional epidemiology sense, which induces bias rather than a loss of efficiency.

4. Case-Crossover Designs that Do Not Condition on Exposure

This paper has focused on the case of a fixed, shared exposure series. Several of the case-crossover methods papers have not conditioned on exposure (32, 35). An unconditional formulation is appropriate for analyses of unshared exposures; if exposures can be assumed to be independent across individuals, observed exposures can be thought of as realizations of a random variable with a given distribution. Therefore, it would be appropriate to examine the properties of the effect estimate and estimating equations averaged over all possible exposure series.

In an unconditional setup, exposures are random, so we can consider the likelihood of the exposures conditional on the index times. Marshall and Jackson (32) presented such a setup, and Navidi (22) briefly discussed this situation. Vines and Farrington (35) examined the properties of the conditional logistic regression estimate averaged over exposure series. They show that the likelihood of the exposures conditional on the index times matches the conditional logistic regression likelihood if the exposures are exchangeable within subject (i.e. if the distribution of the exposure is constant over time and there is a constant correlation between pairs of exposures over time). Yet, if $\beta = 0$, the assumption of a constant distribution of exposure over time (i.e., stationarity, rather than exchangeability) is enough to guarantee that conditional logistic regression yields unbiased estimates averaged over exposure series.

It is also clear that, for any exposure distribution, there is no large sample bias in the conditional logistic regression effect estimate when averaged over all possible exposure series if full stratum bidirectional or time stratified referents are used. This follows from the fact that, for a given exposure series, we know that the expected value of the effect estimate is β , using these referent strategies. Thus, the average over all exposure series of this expected value is also β (i.e. for exposures

$X, E(\hat{\beta} - \beta) = E_X(E(\hat{\beta} - \beta | X)) = E_X(0) = 0$). Consequently, use of the time stratified design also ensures that the conditional logistic regression estimating equations are unbiased when averaged over all possible exposures. With other referent strategies, strong assumptions about the distribution of exposure are needed to ensure unbiased

estimating equations. In particular, the assumption of a constant correlation between pairs of referent exposures over time is almost surely untrue.

5. Other Considerations Involved in the Case-Crossover Design

There are many important decisions to be made with respect to the case-crossover design which have yet to be discussed. The most fundamental design decisions concern the choice of study population and definition of the event. Maclure and Mittleman (37) discuss these issues as well as the determination of the induction period, hazard period, and latency period. They define the hazard period as the at-risk time period prior to the event during which the occurrence of an exposure is hypothesized to elicit a biological response. The induction, hazard, and latency periods affect the choice of referent selection strategy and index time. For example, the index time should fall within defined the hazard period. In addition, exclusion periods (times when exposures cannot be used as referents) should be defined if there is any risk of carry-over effects of the exposure into times after the hazard period.

Other case-crossover review papers have summarized many of the important strengths and weaknesses of the design. Maclure and Mittleman (37) provide a good discussion of the situations in which use of the case-crossover design is appropriate. We also note that the design is most useful for exposures that have a short latency. Redelmeier and Tibshirani (38) discuss bias and interpretation issues related to the available data, unavailable data, analytic technique, quantitative statistics, and etiologic model. Of particular importance with air pollution data, they point out that the estimates obtained

from a case-crossover study always represent short-term, rather than cumulative, effects of exposure. In addition, the case-crossover design (and many other common designs) cannot distinguish between exposures that trigger additional events or exposures that hasten inevitable events.

Endogenous exposures are those generated or influenced by the individual under study. They can only be observed while the individual is observed, and can often be assumed to change as a result of a major event. Examples of endogenous exposures include coffee consumption and episodes of anger. While ambient air pollution is exogenous, since a subject's past event cannot influence his future exposure, personal air pollution, and possibly other important covariates, are endogenous. Study of endogenous exposures can be problematic because the inclusion of an endogenous exposure requires a unidirectional referent selection strategy. In addition, the effect of overlap bias is not yet understood for endogenous exposures.

We have focused our review on using the case-crossover design to assess the effect of exposure on the risk of a *rare* event. For more common events, there is a different underlying likelihood. A hazard model is no longer appropriate. Assuming a logistic regression model, Navidi (22) derived the likelihood of the index times conditional on a fixed exposure series. However, his derivation assumes that the index times are independent within subject over time, after conditioning on exposure. Navidi suggested that, if index times are in fact dependent, covariates such as the number of previous events or the time since the last event should be included in the model. Navidi's

likelihood is incorrect in this case, however. Lumley and Levy (31) derived the true likelihood when the probability of an event at a given time depends on the subject's history. This likelihood could be used to generate score equations to estimate β ; however, finding the solution is computationally difficult. A simplification occurs if, as in Navidi's formulation, index times are independent, and, in addition, the length of the exposure series is large relative to the number of subjects. In this case, one may perform an unconditional logistic regression analysis. See Navidi (22) for other computational alternatives. More research is needed to determine the appropriate analysis for a repeated events application of the case-crossover design. See Dewanji and Moolgavkar (39) for a Poisson process approach to the analysis of recurrent event data.

6. Discussion and Conclusions

The case-crossover design is well suited to the study of the association between short-term air pollution exposure and the risk of an acute adverse health event. By making within-person comparisons, confounding by time independent confounders is eliminated. If referents are matched on important time dependent confounders, these effects are also controlled. Effect modification is easily assessed, and standard conditional logistic regression methods can be used for analysis.

With air pollution exposures, confounding is of particular concern, since confounding tends to dominate the exposure effects. Due to confounding, time trends, non-stationarity, and autocorrelation in the exposure series, proper referent selection is particularly important. Referents should be matched on the most dominant time-varying

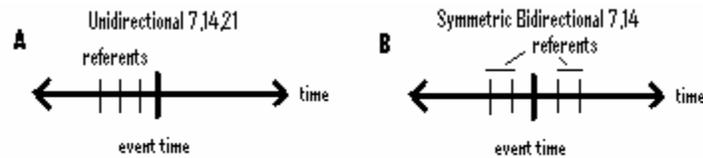
confounders, and should be sampled both pre- and post- event. Referent windows should be small enough to legitimize the stationarity assumption. Sampling referents too close to the index day will result in a loss of power due to autocorrelation in the exposure series. Finally, the analysis should condition on the fixed and known exposure series.

If an **event-dependent** referent scheme is used, the conditional logistic regression estimating equations are biased. This bias is usually small, though it can exist for effect estimates in the range of those typically seen in air pollution studies, and even when there is no exposure effect. The magnitude of the overlap bias depends on the particular exposure series, and hence model shopping on referent strategies will tend to exacerbate the bias. (Model shopping occurs when a referent selection strategy is chosen because it estimates larger effects than other candidate referent selection strategies). Therefore, it seems prudent that overlap bias be avoided entirely.

We conclude with a set of specific recommendations for choosing a referent selection strategy in an air pollution exposure study. We assume that the exposure is exogenous and that the outcome of interest is rare. We recommend that the time stratified design be used for referent selection. This strategy avoids overlap bias as well as bias due to time trend and edge effects. The stratification can be tailored to control for important time dependent confounders by design. A stratification based on year, month, and day of the week is adequate for most studies. (Alternatively, day-of-week could be dropped in favor of “similar temperature”). Though the semi-symmetric bidirectional design also achieves many of these goals, it requires some modification of the traditional conditional logistic

regression analysis (see Section 3). In addition, a time stratified design is likely to be more efficient due to its ability to use of a larger number of referents.

Figure 1 Two referent selection strategies commonly used in air pollution exposure studies



References

1. Ciocco A, Thompson D. A follow-up of Donora ten years after: methodology and findings. *Am J Public Health* 1961;51:155-64.
2. Logan W, Glasg M. Mortality in London fog incidents. *Lancet* 1953;1:336-8.
3. American Thoracic Society BR. Health effects of outdoor air pollution, Part 1. *Am J Respir Crit Care Med* 1996;153:3-50.
4. American Thoracic Society BR. Health effects of outdoor air pollution, Part 2. *Am J Respir Crit Care Med* 1996;153:477-98.
5. Burnett RT, Goldberg MS. Size-fractionated particulate mass and daily mortality in eight Canadian cities. In: Revised analyses of time-series studies of air pollution and health, Special Report. Boston, MA: Health Effects Institute, 2003:85-90.

6. Klemm RJ, Mason R. Replication of re-analysis of Harvard Six-City mortality study. In: Revised analyses of time-series studies of air pollution and health, Special Report. Boston, MA: Health Effects Institute, 2003:165-72.
7. Lipfert F, Wyzga R. Air pollution and mortality: issues and uncertainty. J Air Waste Manag Asssoc 1993;45:949-66.
8. Moolgavkar SH. Air pollution and daily deaths and hospital admissions in Los Angeles and Cook counties. In: Revised analyses of time-series studies of air pollution and health, Special Report. Boston, MA: Health Effects Institute, 2003:183-98.
9. Pope CA, Dockery D, Schwartz J. Review of epidemiological evidence of health effects of particulate air pollution. Inhal Toxicol 1995;47:1-18.
10. Samet JM, Zeger SL, Dominici F, et al. The national morbidity, mortality, and air pollution study. Part II: Morbidity, mortality, and air pollution in the United States. Res Rep Health Eff Inst 2000;94:5-70.
11. Samet JM, Zeger SL, Dominici F, et al. National morbidity, mortality, and air pollution study. Part I: Methods and methodologic issues. Res Rep Health Eff Inst 2000;94:5-14.
12. Samet JM, Dominici F, Curriero FC, et al. Fine particulate air pollution and mortality in 20 U.S. cities, 1987-1994. N Engl J Med 2000;343:1742-9.

13. Schwartz J. Daily deaths associated with air pollution in six U.S. cities and short-term mortality displacement in Boston. In: Revised analyses of time-series studies of air pollution and health, Special Report. Boston, MA: Health Effects Institute, 2003:219-26.
14. Schwartz J. Airborne particles and daily deaths in 10 U.S. cities. In: Revised analyses of time-series studies of air pollution and health, Special Report. Boston, MA: Health Effects Institute, 2003:211-8.
15. Schwartz J. Air pollution and daily mortality in Birmingham, Alabama. *Am J Epidemiol* 1995;137:1136-47.
16. Sheppard L, Levy D, Norris G, et al. Effects of ambient air pollution on non-elderly asthma hospital admissions in Seattle, Washington, 1987-1994. *Epidemiology* 1999;10:23-30.
17. Murray CJ, Lopez AD. *The Global Burden of Disease*. MA, USA: Harvard School of Public Health and World Health Organization, 1996.
18. Maclure M. The case-crossover design: A method for studying transient effects on the risk of acute events. *Am J Epidemiol* 1991;133:144-53.
19. Kwon HJ, Cho SH, Nyberg F, et al. Effects of ambient air pollution on daily mortality in a cohort of patients with congestive heart failure. *Epidemiology* 2001;12:413-9.
20. Lee JT, Schwartz J. Reanalysis of the effects of air pollution on daily mortality in Seoul, Korea: A case-crossover design. *Environ Health Perspect* 1999;107:633-6.

21. Levy D, Sheppard L, Checkoway H, et al. A case-crossover analysis of particulate matter air pollution and out-of-hospital primary cardiac arrest. *Epidemiology* 2001;12:193-9.
22. Navidi W. Bidirectional case-crossover designs for exposures with time trends. *Biometrics* 1998;54:596-605.
23. Neas LM, Schwartz J, Dockery D. A case-crossover analysis of air pollution and mortality in Philadelphia. *Environ Health Perspect* 1999;107:629-31.
24. Peters A, Dockery DW, Muller JE, et al. Increased particulate air pollution and the triggering of myocardial infarction. *Circulation* 2001;103:2810-5.
25. Sunyer J, Schwartz J, Tobias A, et al. Patients with chronic obstructive pulmonary disease are at increased risk of death associated with urban particle air pollution: A case-crossover analysis. *Am J Epidemiol* 2000;151:50-6.
26. Bateson TF, Schwartz J. Control for seasonal variation and time trend in case crossover studies of acute effects of environmental exposures. *Epidemiology* 1999;10:539-44.
27. Bateson TF, Schwartz J. Selection bias and confounding in case-crossover analyses of environmental time-series data. *Epidemiology* 2001;12:654-61.
28. Greenland S. Confounding and exposure trends in case-crossover and case time-control designs. *Epidemiology* 1996;7:231-9.

29. Lee JT, Kim H, Schwartz J. Bidirectional case-crossover studies of air pollution: bias from skewed and incomplete waves. *Environ Health Perspect* 2000;108:1107-11.
30. Levy D, Lumley T, Sheppard L, et al. Referent selection in case-crossover analyses of acute health effects of air pollution. *Epidemiology* 2001;12:186-92.
31. Lumley T, Levy D. Bias in the case-crossover design: implications for studies of air pollution. *Environmetrics* 2000;11:689-704.
32. Marshall RJ, Jackson RT. Analysis of case-crossover designs. *Stat Med* 1993;12:2333-41.
33. Mittleman MA, Maclure M, Robins JM. Control sampling strategies for case-crossover studies: An assessment of relative efficiency. *Am J Epidemiol* 1995;142:91-8.
34. Navidi W, Weinhandl E. Risk set sampling for case-crossover designs. *Epidemiology* 2002;13:100-5.
35. Vines SK, Farrington CP. Within-subject exposure dependency in case-crossover studies. *Stat Med* 2001;20:3039-49.
36. Janes H, Sheppard L, Lumley T. Overlap bias in the case-crossover design, With application to air pollution exposures. UW Biostatistics Working Paper Series 2003. Working Paper 213. <http://www.bepress.com/uwbiostat/paper213>

37. Maclure M, Mittleman MA. Should we use a case-crossover design? *Annu Rev Public Health* 2000;21:193-221.
38. Redelmeier DA, Tibshirani RJ. Interpretation and bias in case-crossover studies. *J Clin Epidemiol* 1997;50:1281-7.
39. Dewanji A, Moolgavkar SH. A Poisson process approach for recurrent event data with environmental covariates. *Environmetrics* 2000;11:665-73.