

Selection Bias and Confounding in Case-Crossover Analyses of Environmental Time-Series Data

Thomas F. Bateson¹ and Joel Schwartz²

The case-crossover study design is a popular analytic tool for estimating the effects of triggers of acute outcomes by environmental exposures. Although this approach controls for time-invariant confounders by design, it may allow for selection bias and confounding by time-varying factors. We conducted a simulation study of the sensitivity of the symmetric bidirectional case-crossover design to time-varying patterns in exposure and outcome. We identified the effects of selection bias and confounding on symmetric bidirectional case-crossover results and offer strategies to eliminate or reduce these biases. Selection bias results when exposure in the reference periods is not identically representative of exposure in the

hazard periods, even when the distribution of exposure is stationary. This bias can be estimated and removed. Selection bias also occurs when the distribution of exposure is nonstationary, but the adjusted symmetric bidirectional case-crossover methodology substantially controls for this. Confounding results from a common temporal pattern in the exposure and the outcome time series, but can also be the result of patterns in exposure and outcome that, although asymptotically uncorrelated, are correlated at finite series lengths. All three biases are reduced by choosing shorter referent-spacing lengths. This effect is illustrated using data on air pollution and daily deaths in Chicago. (EPIDEMIOLOGY 2001;12:654–661)

Keywords: epidemiologic research designs, epidemiologic methods, case-crossover studies, case-referent studies, environmental epidemiology, epidemiologic biases, selection bias, confounding variables

The case-crossover design of Maclure¹ represents an attractive approach to examining the impact of time-varying exposures that may be triggers of adverse health events. The design contrasts, for each individual who suffered an acute event, exposure during an at-risk “hazard period” just before the event to the individual’s exposure on 1 or more reference days when the event did not occur. This approach controls for all measured and unmeasured time-invariant confounders by design, because these are constant within each individual’s stratum. It has been applied^{2–6} to examine the effects of air pollution on mortality and morbidity.

Bateson and Schwartz⁷ examined how different referent selection strategies controlled for confounding by seasonal (365-day cycle) and long-term time trends that are usually present in environmental time series and may cause substantial bias. They found that symmetric bidi-

rectional (SBI) referent sampling with close referents reduced that bias. Recently, investigators^{8–15} have raised questions about the choice of reference periods that suggest that it is more difficult than initially expected to choose periods that avoid producing selection bias when there are temporal patterns in the exposure time series. We address questions of selection bias and confounding in SBI case-crossover study design and the implications for choice of reference periods in analyses of environmental time series in which bidirectional reference period sampling is justifiable. First, using simulations, we address the issue of selection bias caused by systematic and random patterns in the exposure time series alone. We then simulate common patterns in both the exposure and the outcome time series and examine the resulting confounding. Finally, we apply the lessons learned to air pollution and mortality data from Cook County, IL (Chicago area).

Selection Bias

“Avoiding selection bias through proper design depends on the extent to which the investigator is aware of potential sources of bias in the anticipated study.”¹⁶ Selection bias results when exposure in the reference periods is not identically representative of exposure in the hazard periods. In SBI case-crossover studies of environmental time-series data, selection bias arises when reference periods are selected on days that are not in the study base that produced the cases. This situation occurs when we try to match control days to the first and last

From the ¹Department of Epidemiology, Harvard School of Public Health, and ²Environmental Epidemiology Program, Harvard School of Public Health, and Channing Laboratory, Brigham and Women’s Hospital, Harvard Medical School, Boston, MA.

Address correspondence to Thomas F. Bateson, Environmental Epidemiology Program, Harvard School of Public Health, 677 Huntington Avenue, Boston, MA 02115.

T.F.B. is supported by National Research Service Award 72-060-6502-2. J.S. is supported in part by Environmental Protection Agency Grant R827353.

Submitted October 24, 2000; final version accepted April 9, 2001.

Copyright © 2001 by Lippincott Williams & Wilkins, Inc.

cases in a time series with daily events using an SBI design. One solution is to restrict the cases to those that allow matching both the lag and the lead reference periods. For a study design choosing referents 7 days before and after each hazard day, the study base for a 1,096-day time series would be all person-time giving rise to the cases that occurred on or between days 8 and 1,089, yet reference days are chosen between days 1 and 1,096. Reference periods are used that are not eligible to be hazard periods, resulting in differential sampling of exposure. Those first and last days only contribute information as referents whereas the central days serve as both hazard days and reference days for other hazard days. If the distribution of exposure is nonstationary, there will clearly be selection bias. Surprisingly, this bias will exist even when the distribution of exposure is stationary because there will likely be unbalanced high or low values of the exposure on those days. Only when all days in a time series serve equally as both hazard and reference days (that is, the sampling ratio for hazard and reference days is equal for each day) does this selection bias disappear. This situation is the inverse of the issue of overlapping strata of control groups described by Austin *et al.*¹⁷ and again by Levy *et al.*¹⁵ in which choosing friend controls in a matched case-control study can lead to bias if each case's stratum of friends is not mutually exclusive. A popular friend becomes differentially represented causing selection bias. We investigated how this selection bias depends on the choice of SBI reference periods and the length of the study time series and we propose a method to estimate and remove it.

DEPENDENCE ON PERIOD OF PATTERN IN THE EXPOSURE SERIES

We simulated an exposure series as a cosinusoidal term with an arbitrary period of 28 days superimposed on a standard random normal variable. We simulated 3 years (1,096 days) of data and assigned each day an event count from a Poisson distribution ($\lambda = 100$). Because the number of events on any particular day was independent of exposure, there should be no effect of exposure on the frequency of event occurrence except any attributable to the study design. We chose referents spaced 1 day before and after each hazard day and increased the length of the spacing between the hazard day and the reference days by 1 day at a time up to 28 days. Each of the analyses was done on the same number of day-matched strata so that only the cases occurring on hazard days between days 29 and 1,068 were eligible for inclusion (1,040 days). For each spacing length, we simulated 1,000 time series and calculated the mean regression coefficient. Figure 1 shows a plot of the referent-spacing length *vs* the mean regression coefficient under this null situation. The pattern of selection bias is a function of the spacing length and the underlying pattern in the exposure time series. It is minimized when the spacing length is exactly equal to the pattern's period or when the spacing length is short.




FIGURE 1. Selection bias by spacing length between the hazard day and the two symmetric bidirectional reference days under a simulated null association between exposure and daily event count. The bias shown at each referent-spacing length is the mean of 1,000 regression coefficients calculated from 1,000 different time series each with a 28-day cosine pattern in exposure superimposed on a standard random normal.

DEPENDENCE ON LENGTH OF THE TIME SERIES

We examined whether this bias was a function of the length of the study period. We call the string of hazard days, which form the basis of the day-matched strata, the hazard-day series. Using a single time series, we chose reference days with spacing lengths of 14 days before and after each hazard day. We estimated the effect at series lengths of 7 days up to 1,040 days long (equivalent to one quarter the length of the period of the pattern up to more than 37 periods). The results were 1,034 regression coefficients of the effect of exposure under the null, each based on a different hazard-day series length. Figure 2 plots the hazard-day series length in cycles of the pattern




FIGURE 2. Regression coefficients for three simulated effect sizes of exposure plotted against the length of the hazard-day time series, in cycles of the 28-day pattern. The pattern in exposure is a 28-day cosine pattern superimposed on a standard random normal. The bottom series reflects a simulated $\ln(RR) = 0$, the middle series reflects a $\ln(RR) = 0.1$, and the top series a $\ln(RR) = 0.5$.

against the regression coefficients for exposure. The bias fluctuates between positive and negative with a 28-day period and is inversely proportional to the length of the series.

We simulated an effect of exposure equivalent to an $\ln[\text{relative risk (RR)}] = 0.1$ or $\ln(\text{RR}) = 0.5$ per unit increase in exposure and repeated the analyses. Figure 2 also shows the observed regression coefficients by the length of the hazard-day series for the two non-null RRs. The pattern of bias is similar and does not appear to be affected by the underlying RR of exposure, suggesting that the bias is additive on this scale. Hence, subtracting the estimated bias calculated from a null effect simulation from the effect estimate seen for a given series length and referent-spacing length might eliminate this bias.

Fixing the daily event count, rather than simulating a random count, will force the outcome to be independent of exposure and remove variability associated with the Poisson process. We used the same exposure data and recalculated the results. The pattern of bias was nearly identical to that in Figure 2. We subtracted this series of coefficients from those in Figure 2. Figure 3 shows that once the bias under the fixed-count null was removed, the short-term fluctuations in the observed regression coefficients were removed and the adjusted estimates of effect are converging on the true simulated effect of exposure.

Even without systematic patterns in the exposure series, patterns will occur randomly. Figure 4 shows three sets of observed regression coefficients plotted against the length of the hazard-day series for a 14-day referent-spacing length when we simulated a random exposure variable. The simulations are for $\ln(\text{RR})$ values of 0, 0.1,




FIGURE 3. Adjusted regression coefficients for three simulated effect sizes of exposure plotted against the length of the hazard-day time series, in cycles of the 28-day pattern. The pattern in exposure is a 28-day cosine pattern superimposed on a standard random normal. Coefficients are adjusted by subtracting the estimate of effect under the fixed-count null simulation from the estimates of effect under the naive analyses. The bottom series reflects a simulated $\ln(\text{RR}) = 0$, the middle series reflects a $\ln(\text{RR}) = 0.1$, and the top series a $\ln(\text{RR}) = 0.5$.

and 0.5. As the length of the hazard-day series is increased, the three sets of coefficients rattle around their true RRs. Again, we found that subtracting the estimates produced under the fixed-count null analysis of these exposure data from those when there was a true effect of exposure produced nearly unbiased and less noisy estimates of the simulated effect of exposure.

Confounding Revisited

Bateson and Schwartz⁷ demonstrated that confounding by seasonal and long-term time trends in environmental time-series data could be largely controlled by using the SBI case-crossover design. We have reexamined that work to consider confounding by patterns with less than a year-long period.

In the following simulations, we have induced patterns in the outcome owing to an omitted covariate that varies in conjunction with the pattern in exposure (that is, a confounder). We simulated an exposure series that was random normal with a standard deviation of 0.15 plus a cosinusoidal term with a 28-day period. We simulated 3 years of exposure data and generated daily event counts that were Poisson distributed with a mean of 100 multiplied by the exposure, the same cosine term, and the RR under study.

We analyzed these data using referent-spacing lengths from 1 to 28 days. Figure 5 shows a plot of the spacing length between the hazard and reference days against the mean of 1,000 regression coefficients under the null. The bias appears to be a sinusoidal function of the ratio between the spacing length to the periodicity of the pattern associated with the exposure and the daily event count. The confounding is minimized when the spacing is exactly equal to the period of the pattern or when the length is very short. We plotted the confounded estimates, adjusted for selection bias, against the series




FIGURE 4. Regression coefficients for three simulated effect sizes of exposure plotted against the length of the hazard-day time series in days. The exposure is random normal with a standard deviation of 0.15. The bottom series reflects a simulated $\ln(\text{RR}) = 0$, the middle series reflects a $\ln(\text{RR}) = 0.1$, and the top series a $\ln(\text{RR}) = 0.5$.




FIGURE 5. Confounding by spacing length between the hazard day and the two symmetric bidirectional reference days under a simulated null association between exposure and daily event count. The bias at each referent-spacing length is the mean of 1,000 regression coefficients calculated from 1,000 different time series, each with a 28-day cosine pattern in the exposure and event count time series.

length (not shown) and found that this confounding does not depend on the length of the time series.

We also explored whether patterns in the exposure and outcome series that are similar in periodicity, but not identical, give rise to confounding. Patterns with different periods in the exposure and event-count series (owing to an omitted covariate) will be, in principle, uncorrelated. However, the correlation is only zero when the length of the time series is a multiple of the two periods or the series is infinitely long. At all other series lengths, there is some correlation and hence confounding. Figure 6 shows the confounding caused by the

correlation between a 28-day pattern in the exposure series and a 29-day pattern in the event count time series for SBI referent sampling with lengths of 14 days. The selection bias is superimposed on another pattern. We subtracted the estimated selection bias calculated with a fixed-count analysis from these results. The pattern of the confounding is a damped harmonic with a period of 812 days, the product of the periodic patterns in exposure and event-count series. A plot of the confounding from a subsequent analysis with SBI reference days spaced 7 days instead of 14 days from the hazard day (not shown) shows that the period of the confounding remains 812 days but that the amplitude of the bias is smaller.

Adjusting for Selection Bias in Actual Air Pollution Data

The simulations were based on trigonometric patterns in exposure, whereas real life is more complicated. To understand the potential for selection bias and confounding in the real world, we must use real-world data. We used air pollution data on particulate matter less than 10 μm in aerodynamic diameter (PM_{10}) (the average of today's and yesterday's PM_{10}), and a daily count of total mortality in Cook County, IL, from 1988 through 1993. Figure 7 shows an abbreviated periodogram for the time series of PM_{10} . The highest spikes indicate that there are relatively strong patterns in the PM_{10} data with those periods.

We ran SBI case-crossover analyses with a range of referent-spacing lengths. To avoid issues surrounding distributed lag effects¹⁸ and autocorrelation,¹⁴ we chose our reference days no closer than 6 days. We wanted to compare the estimated effect of PM_{10} across different spacing lengths from 6 to 14 days, so we excluded the




FIGURE 6. Confounding by length of the hazard-day time series under a simulated null association between exposure and daily event counts. The thin line is confounding because of the correlation between a 28-day pattern in the exposure time series and a 29-day pattern in the event-count time series, which included the selection bias. The bold line is the confounding adjusted for the selection bias.




FIGURE 7. Abbreviated periodogram of the PM_{10} exposure time series in Cook County, IL, from 1988–1993. Periods shown are restricted to those less than 60 days.

first and last 14 days from the hazard-day series so each analysis was done on the same number of cases. There were 2,162 day-matched strata available. We controlled for day-of-the-week effects and for temperature, relative humidity, and barometric pressure measured on the event day and temperature measured the day before. We modeled the weather covariates using penalized splines with approximately 4 degrees of freedom per term.

We estimated the PM_{10} effect using conditional logistic regression where each day-matched stratum was weighted by that day's event count and estimated the selection bias for each of the same referent spacings, with the daily event counts fixed at 1. The magnitude of the selection bias ranged from -0.0461 to -0.0181 per $100 \mu g/m^3$ increase in PM_{10} , substantial compared with the estimated effect size. The bias was greatest at a referent-spacing length of 14 days. Figure 8 shows the estimated selection bias for a spacing of 14 days at hazard-day series lengths of 2,143 to 2,162 days. The bias ranges from -0.0461 to -0.0013 over a span of 20 days.

We subtracted the spacing-specific estimates of the selection bias from the estimated PM_{10} effect calculated using the naive approach. For comparison, we analyzed the same time series using a robust generalized additive Poisson model. We modeled weather using natural splines with 4 degrees of freedom for each term. We controlled for day-of-the-week effects and controlled for seasonality and long-term time trends using a natural spline with 4 degrees of freedom per year of data. Table 1 shows the results of the adjusted SBI case-crossover analyses and the Poisson regression analysis.

If each of the nine case-crossover studies are measuring the same underlying effect of PM_{10} on mortality, then all of the reference pairs could be combined in a single analysis. The result was an estimated coefficient of PM_{10} effect of 0.0576 per $100 \mu g/m^3$ increase [95% confidence interval (CI) = 0.0233 – 0.0920]. The Poisson analysis estimated

the coefficient of PM_{10} effect at 0.0559 per $100 \mu g/m^3$ increase (95% CI = 0.0238 – 0.0880).

We have shown that a combination of patterns in the exposure and event count time series can produce some confounding even when the patterns are asymptotically uncorrelated. We examined the adjusted SBI case-crossover results over a whole year's change in the hazard-day series length to look for any pattern that might indicate the presence of this confounding. Figure 9 shows the adjusted regression coefficients for a lag length of 14 days for series lengths of 1,800 days up to 2,160 days.

Discussion

The SBI case-crossover approach is an attractive alternative to Poisson time-series analyses using generalized additive models, but we are still learning about its strengths and sensitivities. The greatest strength of the case-crossover design is that it controls for all time-invariant confounding by design.¹ The addition of a second reference period symmetrically spaced in time about the hazard day substantially controls for time-varying confounders such as seasonal variation and long-term time trends.⁷ When individuals' personal covariate information is available, the design allows for the modeling of effect modification.

Alongside these strengths, we have identified a sensitivity of the SBI case-crossover approach to selection bias from both systematic and random patterns in the exposure time series and proposed a method for removing it. Even a small bias could be substantial in comparison with some effect size estimates. We found selection bias, on the log scale, as large as -0.0461 , where Poisson regression yields a PM_{10} effect estimate of 0.0559 . Moreover, the size of this bias varied substantially with small changes in series length.

The SBI case-crossover analyses of patterned exposure data in Figure 1 show that selection bias is




FIGURE 8. Selection bias associated with a 14-day symmetric bidirectional spacing of reference days by length of the hazard-day series for PM₁₀ effect on total mortality in Cook County, IL, 1988–1993. PM2davg is the average of today's PM₁₀ and yesterday's PM₁₀. All models control for temperature, barometric pressure, and relative humidity on the hazard day and for temperature on the day before the hazard day as well as day of the week.

minimized by choosing referents at spacing lengths equal to the period of the pattern in exposure, but this period may not be known *a priori* or the pattern may be multiperiodic. Hence, a strategy of choosing reference days closer to the event day should tend to minimize bias from patterns with any periodicity and, therefore, be preferable. Even when the exposure is random, there will be bias under the null that follows a pattern specific to that particular exposure time series. Figures 2 and 4 showed that the bias is a function of the length of the hazard-day time series

with much bias at short series lengths relative to the period in the exposure. As the length of the study increases, the bias decreases. This phenomenon may explain the bias Lee *et al*¹³ observed in the presence of skewed and incomplete waves. Figure 2 also showed that the selection bias appears to be independent of any true RR of exposure and to be dependent on patterns and random fluctuations in the exposure time series.

Lumley and Levy¹⁴ and Levy *et al* argue¹⁵ that designs selecting referents from categories that are not

TABLE 1. Results Are the Coefficients of Effect of Particulate Matter <10 μm in Aerodynamic Diameter per 100 $\mu\text{g}/\text{m}^3$ on Total Mortality in Cook County (1988–1993) from a Generalized Additive Poisson Regression Using 24 Degrees of Freedom to Control for Season and from Symmetric Bidirectional (SBI) Case-Crossover Analyses with Different Lag Lengths from 6 to 14 Days

Analytic Method	Spacing between the Hazard and Reference Days	Beta	Standard Error	95% CL
Poisson Regression		0.0559	0.0164	0.0238, 0.0880
SBI case-crossover	6 days	0.0477	0.0204	0.0076, 0.0877
SBI case-crossover	7 days	0.0462	0.0203	0.0064, 0.0860
SBI case-crossover	8 days	0.0383	0.0203	−0.0014, 0.0780
SBI case-crossover	9 days	0.0301	0.0201	−0.0093, 0.0695
SBI case-crossover	10 days	0.0501	0.0202	0.0105, 0.0896
SBI case-crossover	11 days	0.0660	0.0199	0.0270, 0.1050
SBI case-crossover	12 days	0.0661	0.0198	0.0272, 0.1049
SBI case-crossover	13 days	0.0738	0.0197	0.0352, 0.1123
SBI case-crossover	14 days	0.0732	0.0198	0.0344, 0.1121
SBI case-crossover	6–14 days Inclusive*	0.0576	0.0175	0.0233, 0.0920

All models controlled for temperature, barometric pressure, and relative humidity on the hazard day and temperature on the day before the hazard day as well as day of the week. 95% CL = 95% confidence limits.

* This model contains nine pairs of symmetric bidirectional controls, one for each referent-spacing length from 6 to 14 days.




FIGURE 9. Adjusted regression coefficients by length of the hazard-day series for PM_{10} effect per 100 58 g/m^3 on total mortality in Cook County, IL, 1988–1993. PM_{2davg} is the average of today's PM_{10} and yesterday's PM_{10} . All models control for temperature, barometric pressure, and relative humidity on the hazard day and for temperature on the day before the hazard day as well as day of the week. Coefficients are adjusted by subtracting the estimate of effect from the fixed-count null analysis from the estimates of effect under the naive analyses.

mutually exclusive can be biased owing to violations of the study base principle. This issue may be especially significant in case-crossover analyses of time-series data with sparse events. In the Cook County data, deaths occurred daily, so each day serves equally as both hazard and reference day except for the first and last days. We show how the selection bias due to this differential sampling of exposure can be estimated and removed while preserving the SBI design that has been shown to control for seasonal and long-term confounding.⁷ Our method of subtracting off the estimated selection bias is similar to Suissa's^{8,10} case-time-control design. His idea was to use a control-crossover study to estimate the bias in a case-crossover study attributable to trends in exposure over the intervening time between a prior control period and the case period. We used a fixed-count null analysis of the exposure to estimate the selection bias attributable to the SBI control sampling methodology and then subtracted it from the naive effect estimates. Our method could be used to remove the selection bias whether it arose from an unobserved pattern in the data or simply from random variation in the exposure measure.

Bateson and Schwartz⁷ simulated various patterns that caused confounding and showed that choosing SBI referents minimized confounding compared with other referent sampling methods. Here we show that confounding is clearly a function of the specific common pattern in the exposure and event-count time series. When there is confounding by an omitted covariate in the SBI case-crossover approach, the bias is generally minimized

by choosing referent-spacing lengths as short as possible without sampling reference periods so close that they would be highly autocorrelated with exposure in the hazard period. Unlike the selection bias we observed from patterns and random fluctuations in exposure, which diminishes as the length of the series increases, confounding by a common pattern does not appear to be a function of series length.

We have also shown that even when there are asymptotically uncorrelated patterns in the exposure and event count time series, confounding can arise, at least in a simulated setting. The bias appears to be a dampened harmonic function with a period equal to the product of the patterns in the exposure and event-count time series. This finding implies that SBI case-crossover studies of relatively short time series may be especially sensitive to this effect. In long time series, this confounding may be dampened down. One way to detect this confounding would be to see whether the effect of exposure, adjusted for selection bias, varies systematically over time.

In our analysis of the PM_{10} effect on total mortality in Cook County, we examined how the adjusted effect of PM_{10} changed as the length of the hazard-day time series varied between 1,800 days and 2,160 days. Figure 9 did not show the kind of confounding pattern that we saw in Figure 6. This result implies that either there was no confounding of that type, or that if present its magnitude had been sufficiently dampened that its effect was unappreciable.

The results in Table 1 show the adjusted effect of PM_{10} on mortality in Cook County using the SBI case-crossover approach with referents chosen at spacings of 6–14 days. Although there is variability between the estimates, all nine of the adjusted estimates fall within the 95% CI of the generalized additive Poisson model analysis of the same data.

The periodogram of PM_{10} shown in Figure 7 reveals that there are several strong periodicities present in these data. The differences in effect estimates across the nine referent-spacing lengths may be the result of separate biases that rise and fall according to how each referent-spacing length interacts with each of these patterns. If we assume the differences are due to random variability and are measuring the same underlying effect, then we could use all of the reference pairs in a single analysis. This analysis produced nearly identical results as the comparable Poisson result. The relative efficiency of this case-crossover result, with 18 controls, compared with the Poisson result is 87.5%.

CONCLUSION

We were surprised to find that, in a 6-year time series, the selection bias caused by differential sampling of a few reference days at either end of the series could induce such meaningful bias in the context of air pollution research. Although we have demonstrated how this bias may be estimated and adjusted for, its potential must be acknowledged and investigators should be careful when making assumptions about the reference exposure distribution, even when that distribution is stationary. Although our analysis was limited to environmental time series, the issue of comparability of the case and reference groups is a general one for case-crossover analyses.

With proper attention to bias and confounding, this approach to studying the triggers of acute events in environmental time series can be a viable alternative to

Poisson regression and may offer advantages by allowing for the facile examination of effect modification.

References

1. Maclure M. The case-crossover design: a method for studying transient effects on the risk of acute events. *Am J Epidemiol* 1991;133:144–153.
2. Navidi W, Thomas D, Langholz B, Stram D. Statistical methods for epidemiologic studies of the health effects of air pollution. *Res Rep Health Eff Inst* 1999;86:1–50.
3. Neas LM, Schwartz J, Dockery D. A case-crossover analysis of air pollution and mortality in Philadelphia. *Environ Health Perspect* 1999;107:629–631.
4. 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;107:633–636.
5. Sunyer J, Schwartz J, Tobias A, Macfarlane D, Garcia J, Anto JM. 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;151:50–56.
6. Levy D, Sheppard L, Checkoway H, Kaufman J, Lumley T, Koenig J, Siscovick D. A case-crossover analysis of particulate matter air pollution and out-of-hospital primary cardiac arrest. *Epidemiology* 2001;12:193–199.
7. Bateson TF, Schwartz J. Control for seasonal variation and time trends in case-crossover studies of acute effects of environmental exposures. *Epidemiology* 1999;10:539–544.
8. Suissa S. The case-time-control design. *Epidemiology* 1995;6:248–253.
9. Greenland S. Confounding and exposure trends in case-crossover and case-time-control designs. *Epidemiology* 1996;7:231–239.
10. Suissa S. The case-time-control design: further assumptions and conditions. *Epidemiology* 1998;9:441–445.
11. Navidi W. Bidirectional case-crossover designs for exposures with time trends. *Biometrics* 1998;54:596–605.
12. Hallqvist J, Möller J, Ahlbom A, Diderichsen F, Reuterwall C, de Faire U. Does heavy physical exertion trigger myocardial infarction? A case-crossover analysis nested in a population-based case-referent study. *Am J Epidemiol* 2000;151:459–467.
13. Lee J-T, Kim H, Schwartz J. Bidirectional case-crossover studies of air pollution: bias from skewed and incomplete waves. *Environ Health Perspect* 2000;108:1107–1112.
14. Lumley T, Levy D. Bias in the case-crossover design: implications for studies of air pollution. *Environmetrics* 2000;11:689–704.
15. Levy D, Lumley T, Sheppard L, Kaufman J, Checkoway H. Referent selection in case-crossover analyses of acute health effects of air pollution. *Epidemiology* 2001;12:186–192.
16. Kleinbaum DG, Kupper LL, Morgenstern H. *Epidemiologic Research: Principles and Quantitative Methods*. New York: Van Nostrand Reinhold, 1982;123.
17. Austin H, Flanders WD, Rothman KJ. Bias in case-control studies from selection of controls from overlapping groups. *Int J Epidemiol* 1989;18:712–716.
18. Schwartz J. The distributed lag between air pollution and daily deaths. *Epidemiology* 2000;11:320–326.