Article

Scand J Work Environ Health 1999;25(4):36-42

Methodology for analyzing episodic events

by Eisen EA

The following article refers to this text: 2002;28(5):293-303

Key terms: case-crossover study; correlated outcomes; epidemiology; exposure-response models; longitudinal study; occupational health; repeated measure

This article in PubMed: www.ncbi.nlm.nih.gov/pubmed/10628439

Scand J Work Environ Health 1999;25 suppl 4:36—42

Methodology for analyzing episodic events

by Ellen A Eisen, ScD[1]

Eisen EA. Methodology for analyzing episodic events. *Scandinavian J Work Environ Health* 1999;25 suppl 4:36—42.

Many health outcomes of interest to occupational epidemiologists are common recurrent health events. Epidemiologic approaches to the study of such health events are reviewed. Episodic events are considered to be events that occur at a distinct point in time - either with sudden onset or the sudden crossing of a threshold of detection - and they must be reversible events that can recur in the same person in response to a proximate trigger. Studying such health events poses 4 challenges to existing methods: (i) key epidemiologic concepts, such as incidence, do not naturally accommodate recurrent events, (ii) study designs must capture time-varying exposures, (iii) statistical models must be able to handle correlated outcomes, and (iv) feedback bias must be addressed. In response, methods such as longitudinal studies, case-crossover designs and generalized estimating equations are identified as appropriate tools.

Key terms case-crossover studies, correlated outcomes, epidemiology, exposure-response models, longitudinal studies, occupational health, repeated measures.

Traditionally, epidemiologic research has focused on causes of low probability health outcomes — health events that can occur only once to any individual and only happen to a relatively small proportion of the population over a fixed interval of time. But consider outcomes such as respiratory symptoms, nonfatal injuries, or musculoskeletal disorders. These health problems are not only common in the population, but can recur over and over again to the same person. Are existing methods suitable for the study of such recurrent problems?

To answer this question, I will first define the term "episodic event" and identify the characteristics that distinguish this class of outcomes. I will then show why these characteristics pose challenges to existing epidemiologic methods — methods that were developed to study rare outcomes. For those interested in studying these common workplace health problems, however, the situation is far from bleak. Examples from the literature will illustrate several modeling approaches that can be effectively applied to episodic events.

Defining episodic events

What exactly is an episodic event? According to ordinary usage, an event is "something that happens". We would all agree that car accidents and heart attacks are events. These things happen quickly at a well-defined point in time. But what about periods of low-back pain? Must the onset of an event be sudden, or can it be gradual? Must an event be instantaneous, or can it extend over a prolonged period? Some events, such as nasal irritation, are an immediate response to the environment. Others, such as a heart attack, have both immediate and long-term causes. For an event to be episodic, must at least one cause be proximate?

Asthma attacks are a classic example of recurrent events. These events have sudden onset, brief duration, and proximate causes and are reversible. However, many health outcomes of current interest are less well defined, such as recurrent bouts of low-back pain or respiratory symptoms. The following discussion will be relevant for

1 Department of Work Environment, University of Massachusetts Lowell, Lowell, Massachusetts, United States.

Reprint requests to: Professor Ellen A Eisen, Department of Work Environment, University of Massachusetts Lowell, One University Avenue, Lowell, MA 01864, Unites States. [E-mail: ellen.eisen@uml.edu]

any recurrent health outcome, regardless of duration, as long as distinct episodes can be defined. For example, as illustrated in figure 1, we could define a discrete beginning and end point to a recurrent episode of back pain as pain above a detectable threshold. This use leads us to a precise definition: *Episodic events can occur suddenly or suddenly cross a threshold, but onset must occur at a distinct time point. They must also be reversible health events that can recur in the same person in response to a proximate exposure.*

The unit of analysis

Note that the terms "acute" and "chronic" do not appear in this definition of episodic events. The terminology was deliberately avoided because it is so often ambiguous. The hallmark of many chronic diseases is precisely the recurrence of short-term events, for example, asthmatics have asthma attacks. Why might we want to study asthma attacks rather than asthma? By studying events we can identify the immediate triggers of recurrent episodes of wheezing. But subjects are the familiar object of epidemiologic studies — not events. Studies of asthma in working populations can identify workplace exposures that cause adult-onset asthma. The appropriate unit of analysis then depends on the scientific question of interest. Whether we study subjects or events will, in turn, determine the choice of study design, the exposure metric, and the statistical model.

A discussion by Cumming et al (1), about falls among the elderly, provides a fine illustration of the importance of the unit of analysis. We can think of falls among 2 groups of workers. Hypothetically, consider 10 subjects — 5 exposed construction workers and 5 unexposed clerical workers. Over a 1-year study period, 1 construction worker fell 3 times, 1 fell once, and 3 did not fall at all. During the same time period, 2 of the 5 unexposed clerical workers fell and each fell once. The object of the study could be either falls or fallers, where a faller is defined to be a subject who falls. If faller is the unit of analysis then 2/5 of the exposed and 2/5 of the unexposed were fallers and there was no apparent effect of exposure on the risk of falling. However, if the fall itself is the unit of analysis, then there were twice as many falls among the exposed (4 versus 2), suggesting that exposure *is* associated with falling. Thus different units of analysis can lead to opposite inferences about the primary etiologic question. Are construction workers at greater risk of falling, or not? How does one decide?

The answer depends on the scientific question of interest. If we want to know why some people fall more than others, then "faller" is the appropriate unit of study. But if we are interested in why a person falls *now* rather than *later*, then we need to study falling and its proximate causes. Restricting our attention to events still leaves open the question of whether to study the first fall, the most recent fall, or all falls that occur within a specified time period. Putting this question to the side for a moment, let us consider the implications of studying events rather than subjects for defining exposure metrics.

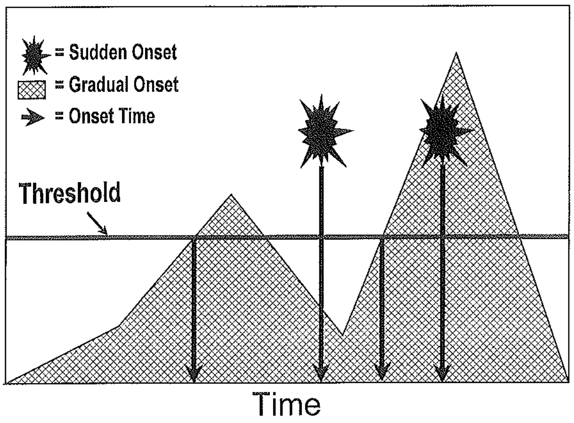

Figure 1. Patterns of recurrent events. Time of onset is clearly defined for sudden onset events. For more gradually occurring events, onset can be defined as the time when the pain crosses a threshold of detectability.

Time-varying exposure

The study of events allows us to examine exposures that vary over time. By locating an event at a distinct time point, we can define the biologically relevant exposure window at some fixed interval of time relative to the event. The length of the exposure window and its temporal position relative to the event will depend on the causal mechanism of interest. For example, in a recent study by my co-workers and I, on the role of metalworking fluid exposure in the initiation of adult-onset asthma, year of diagnosis was used to locate the *initial event* and an exposure window of 2 years just prior to the date of diagnosis was defined (2). The basis for selecting a 2-year window was clinical data suggesting that, on the average, symptoms occurred *2 years* prior to the diagnosis of occupational asthma (3). Risk sets, that is, subjects at risk of diagnosis at the date of diagnosis of each case, were defined and metalworking fluid exposure in the 2-year window was compared between the cases and the noncases.

As another example of a time varying exposure, recall the recently published study of the role of cellular phones in automobile collisions (4). The risk of a collision was found to be 4 times higher when a cellular phone was in use than when it was not. A series of 5-minute exposure windows was examined. The association was strongest when the exposure was defined as a cellular phone call initiated in the 5 minutes *just prior* to the collision. The relative risk decreased from almost 5 to 3 for calls placed 6 to 10 minutes prior to the accident and then fell to just above 1 for calls made 16 to 20 minutes before.

Methodologic challenges

When the outcome is a common recurrent condition, measuring multiple outcomes per subject will be more efficient than restricting the analysis to a single event. Therefore I will restrict our attention to methods suitable for studying multiple events. There are 4 methodological difficulties that arise when outcomes can recur for a person during the study period. *First*, some of the most basic epidemiologic concepts, measures of disease, frequency and association, do not naturally accommodate recurrent events. *Second,* studying repeated events, rather than subjects, requires different study designs. *Third*, and perhaps most important, outcomes measured repeatedly for the same person are correlated. This correlation violates the independence assumption required by the standard risk and rate models. *Finally*, repeated measures of both outcome and exposure allow both variables to vary over time, and therefore a potential feedback bias is introduced in which earlier outcomes may affect subsequent exposures.

Redefining incidence

Beginning with the first challenge — How do we measure the frequency of a recurrent event? If we are studying recurrent episodes of wheeze at work, then we want to measure the rate of wheeze. But are we measuring *an incidence rate*? An incidence rate measures the occurrence of *new* events per unit of person-time. Generally, recurrent events will not be purely incident except for the first one, and so the mathematical links between incidence rates and risk do not hold.

Since the incidence of a recurrent event, like falling, is a little vague, maybe prevalence is more appropriate. Prevalence is defined as a measure of status (ie, the proportion of subjects who have the disease of interest at a specific time). Whereas the proportion of subjects with low-back pain can be measured, the proportion of subjects falling cannot be because falls are instantaneous. Without measures of incidence or prevalence, we cannot estimate associations with incidence rate ratios. The point is that the basic definitions and measures of association do not make sense for many types of recurrent events.

For the study of episodic events we need to modify our definition of incidence slightly. In several studies of recurrent symptoms, an incident symptom was defined as a *recurrent event that is preceded by a period without an event*. For example, in our study of nasal irritation among workers exposed to sodium borate, we studied workers over a 4-day period with an hourly symptom questionnaire (5). In that analysis, an incident symptom was defined as a report of irritation that was not reported in the previous survey. Incidence was similarly redefined in 2 recent studies of the acute effects of air pollution. In both, daily diaries were used to collect information on respiratory symptoms — over a 6-week period for a cohort of Swiss children (6) and a 3-year period for a population of nursing students in Los Angeles (7). In both analyses, the investigators defined an incident symptom as the presence of a symptom when the previous day was symptom-free.

Study designs for recurrent outcomes

Let us turn now to the second challenge — finding a study design suitable for studying multiple events. I will run through the list of familiar study designs and point out the key shortcomings and advantages of each. Most can be used to study a single episodic event per subject. Only a few, however, can be applied in a straightforward manner to the study of multiple events. To be appropriate, a design must permit us to capture information about time co-varying outcomes and exposures.

Cross-sectional studies. Cross-sectional studies measure prevalence rather than incidence. Recall that cross-sectional studies will overrepresent prevalent conditions of longer duration and underrepresent conditions of shorter duration. In the case of events, such as falls, however, the problem is exaggerated because the outcome is essentially instantaneous. For such outcomes, prevalence studies are inappropriate. For longer lasting outcomes, such as low-back pain, a cross-sectional study *can* be designed to measure prevalence. However, one still needs to be attentive to the possibility of length-biased sampling and the healthy worker effect.

Case-referent studies. Case-referent studies can be designed with cases identified as persons with an event (eg, subjects with a fall resulting in lost workdays). However, the case-referent design cannot take full advantage of information about multiple events in a repeated-measures study. Thus this design has limited applicability to the study of episodic events.

Case-crossover studies. A variant of the case-referent design, called case-crossover may be more appropriate. The classic crossover study is an experimental design in which two or more interventions are compared by treating each subject with each intervention, in a random sequence. Maclure (8) adapted this experimental design to the observational study of acute health events with sudden onset. For each subject with an event, the event is treated as the "case" and one or more earlier time periods are selected as "matched" reference periods. Exposure at the time of the event is then compared with exposure in a previous time interval for the same person. A major advantage of this design is that, because subjects serve as their own matched referents, the design perfectly controls for confounders that vary between subjects but remain constant for a specific person over time. The case-crossover design is appropriate for outcomes that occur with sudden onset shortly after exposure. Moreover, the design is particularly well suited to evaluating risk factors that change over the study period, that is, for time-varying exposures. The design has been effectively used in recent studies of myocardial infarction (9) and in the study of cellular car phones and automobile collisions already mentioned. A primary disadvantage is that, because the design is limited to the evaluation of time-varying exposures, it is not appropriate for studies of workplace or subject characteristics that do not change for a person over the study period. For example, in a case-crossover study of injuries, fixed exposures such as work organization or job demand and control can only be studied as effect modifiers, and not as primary exposures.

To date, case-crossover designs have been applied only to the study of a single event. Moreover, the outcomes have been rare where the probability of a recurrence within the study period is vanishingly small. Can the method be applied to more common outcomes, where multiple events are observed per subject? By taking advantage of statistical methods for repeated measures developed over the last few years, we are able to apply this design to the study of episodic events. An example follows.

We have seen that several study designs work for examining a single event per subject — cross-sectional designs for prolonged events and case-crossover designs for instantaneous events. Let us now turn to cohort studies. Cohort studies are typically time-to-event studies designed to examine mortality or the incidence of disease. Time until the first onset of an episodic event can certainly be studied in this framework. Recurrent episodes, however, cannot be so easily handled.

Longitudinal studies. A cohort study can also be designed as a longitudinal study in which repeated health measurements are collected for a cohort followed over time. Typically longitudinal designs are used to study the development of chronic conditions by observing changes over time in a continuous variable. For example, pulmonary function collected annually over many years can provide evidence about the role of environmental factors in the development of chronic lung disease. A longitudinal study, however, with the simultaneous collection of exposure and health data, is also the most appropriate design for studying recurrent binary events.

Intervention studies. When the induction time between exposure and disease is short, as is the case for many episodic events, there is the opportunity to study the effectiveness of an intervention by measuring short-term change in health status. For example, Moulton recently reanalyzed data from a randomized trial of the effectiveness of vitamin A supplementation on recurrent childhood diseases (10). This was a double blinded intervention study of 1405 preschool age children randomized to receive either vitamin A or a placebo every 4 months for 2 years.

While epidemiologists prefer to talk about incidence rates and cohort studies, there is a enormous amount of statistical literature on longitudinal data analysis that is directly relevant to the study of episodic events (11, 12). I will now turn to the 3rd challenge — finding statistical methods available for modeling longitudinal data.

Modeling approaches for correlated outcomes

If only a single event is observed per subject, no special statistical models are needed to analyze event data. Depending on the study design, one of the standard rate or risk models can be applied. If, however, we allow events to recur, then we introduce a substantial statistical problem, namely, that repeated observations on the same person are not independent of one another. This is the 3rd major challenge posed by recurrent events.

Incidence rate models, such as proportional hazards or Poisson regression, can be adapted to deal with correlated dichotomous outcomes. Because proportional hazards are better suited to handle time varying covariates, it is more attractive for analyzing episodic events. Several modifications have been developed within the general framework of proportional hazards to address correlated outcomes. These approaches can be broadly divided into 2 types, the subject-specific and the population-averaged approaches (13). The subject-specific approaches yield specific parameter estimates for each person in the study and provide an explicit description of the heterogeneity in responsiveness across the population. By contrast, population-averaged models provide an average exposure-response parameter for the entire study population.

Subject-specific approaches. There are 2 models that can be used to estimate subject-specific parameters — a 2-stage model and a random effects model. The same picture can be used to illustrate both. A separate dose-response curve is fit to the data for each person. In a random effects model only the intercept can vary by subject. In the 2-stage model both intercept and slope can vary across subjects. This approach makes the notion of heterogeneous risk explicit and is appropriate if one is interested in studying differences in individual susceptibility.

In their classic 1979 paper, Korn & Whittemore (14) proposed a 2-stage model for handling repeated categorical outcome data in a study of asthma by the

Environmental Protection Agency in the United States. Groups of asthmatics, selected in several different communities, kept daily symptom diaries. Asthma attacks were examined in relation to daily air pollution measurements collected in each community. In stage I, a separate logistic model was fit to the data for each subject. This approach permits an individual exposure-response slope to be estimated for each subject. In stage II, the host factors that influence responsiveness were examined.

Population-averaged approaches. There are several modeling approaches that have been developed to provide population-averaged parameter estimates for repeated measures studies. I will focus on the most widely applied — generalized estimating equations (GEE). Liang & Zeger (15) proposed the GEE method as an extension of the generalized linear model that accounts for intrasubject correlations within the familiar framework. Thus the GEE method can be applied to a wide range of familiar model forms, including logistic, linear, and Poisson regression.

With the GEE, a particular form for a "working" correlation matrix is chosen to account for the *average* dependency of the repeated measures for each subject. Typical forms include an autoregressive, exchangeable, or independent correlation structure. Estimates of the variances of the regression parameters are computed using the matrix, and the estimates are claimed to be robust, that is, not too sensitive to changes in the matrix. The primary disadvantage of the GEE approach is that it is essentially a cross-sectional, rather than a longitudinal, approach.

To illustrate how the GEE approach differs from a simple cross-sectional analysis, consider Schwartz & Zeger's study (7) of passive smoking and air pollution, in which 100 nursing students recorded acute respiratory symptoms in daily diaries for 3 years, producing 1000 binary observations per subject. Symptom incidence was defined as the presence of a symptom following a symptom-free day. Schwartz & Zeger modeled incidence by applying GEE to logistic regression. But suppose, instead, the correlations for individual subjects were ignored, and all of the daily diary data were combined in a simple logistic regression model. That is, suppose the 100 000 observations were treated as independent outcomes in a single cross-sectional survey. The regression slopes in the 2 different models would be virtually identical. The confidence intervals would be different, however, and correct only in the GEE model. The disadvantage of the GEE approach is that it treats longitudinal data as though they were cross-sectional. The advantage is that the confidence intervals correctly reflect the repeated-measures nature of the data.

Repeated measures are usually collected for 1 of 2 reasons. In situations in which the study population is limited, adding follow-up surveys is an effective way to improve efficiency. Efficiency motivated the design in both Korn & Whittemore's asthma study and Schwartz & Zeger's nursing study. In both cases, the questions were essentially cross-sectional, and applying a GEE approach would be appropriate. Repeated data can also be collected to examine incidence, persistence, and remission, that is, change in health status over time. In such cases, the GEE method is inappropriate and a time-to-event approach based on the proportional hazards model is more suitable. For example, one can fit a series of separate proportional hazards models to the 1st, 2nd, and 3rd events of each subject. The 1st models time to the 1st event; the 2nd includes only the subjects with at least 1 event, and models time from the 1st to the 2nd; and so on. Several more complicated modifications permit the modeling of recurrent events in a single proportional hazards model (16, 17, 10).

An illustration. To illustrate further the alternative statistical approaches, I will use the Borax Acute Health Study — a study of respiratory irritation and exposure to sodium borate dust (5). Irritation is a common, transient event with a brief latency period (measured in minutes) and a short duration of response. We designed a longitudinal study of a cohort of 106 workers in a borax plant in the United States. Field technicians administered a brief symptom survey *every hour* during the first *6 hours* of the workshift for *4 consecutive workdays*. The analysis was based on the 24 repeated-symptom measurements for each subject and the parallel exposures. Each subject wore a real-time personal exposure monitor for the full shift. Using the nature of the irritant response as a guide, we decided to define 15 minutes as the biologically relevant exposure window.

We used a 2-stage model to estimate subject-specific measures of association. In stage I we fit a separate logistic regression model to the data for each subject who reported at least 3 events and estimated a subject-specific odds ratio. We then examined the distribution of the exposure-response parameters in stage II and tried to identify host characteristics that might explain differences in the odds ratios. The results of stage I suggested a range in the dose-response parameters, most being close to 1.0 and a small subset of subjects having an odds ratio greater than 4-fold. The mean odds ratio of 1.3 was significantly larger than 1.0 and therefore provided evidence for an overall association between nasal irritation and dust exposure.

This analysis, however, was very inefficient. Only the 25% of the cohort who had multiple responses contributed to the exposure-response modeling. The data for the majority of the cohort was ignored. If we were to reanalyze the data today, how would we do it? We could recast our design as case-crossover. Cases of nasal irritation could be identified. The hazard period could be defined as the 15 minutes prior to the irritation. Matched reference periods could be selected as the 15 minutes prior to another survey for the same person. With a single case and reference period per subject, we could fit a conditional logistic model for all subjects combined. But recall that workers could potentially report irritation in multiple surveys. Multiple case and control periods per subject could be included in the analysis if GEE were applied to account for the correlation between repeated measures for a specific person. The result would be a population-averaged odds ratio rather than a set of subject-specific odds ratios and would take greater advantage of the available data.

Feedback bias

We have seen that, to study episodic events effectively, we need estimates of associations between exposures and health events that vary over time. The time varying nature of both variables, however, poses the 4th challenge by introducing a new potential bias. We hypothesize that past or current exposure will influence current health status. We also acknowledge the likely influence of past health status on current status (ie, a positive correlation between repeated outcomes). However, there is also the opportunity for causality to go in the reverse direction (ie, for health status to have an impact on exposure, that is, for *past* health status to have an impact on *subsequent* exposure). This reversal in the direction of a causal relationship has been referred to as feedback (18).

For example, suppose, in our study of acute irritation, a worker with a high 15-minute exposure has severe nasal irritation. If, as a result, the worker moves further away from the source of exposure, he or she will reduce exposure in the next time window. If only the more resistant persons continue to receive high exposure, we may have a bias caused by healthy-worker survivor selection. Including previous health status as a confounder is one way to address this potential bias. A 2nd problem arises when a covariate that varies over time acts as both a confounder *and* an intermediate variable. Such a situation can cause bias whether or not we control for it.

None of the methods I have described will adequately control for these more complicated interrelationships among the variables. One way to view these problems is to borrow some concepts from econometricians who have developed sophisticated methods for analyzing time series data (19). The method, called structural equations modeling, involves fitting a set of multiple model equations simultaneously. The results are unbiased even when exposure depends on previous health status. G-estimation, an approach developed by Robins to address the healthy worker effect, is an example of a structural model applied to epidemiologic data (20).

Concluding remarks

I have attempted to make only a few general points. First, many of the health outcomes of current interest to researchers in occupational health are episodic in nature. Second, it is generally more efficient to collect repeated measures if one is interested in examining the proximate causes of recurrent events. Third, correctly analyzing repeated-measures data requires statistical methods designed to handle correlated outcomes. Fortunately there are several different methods that have already been developed and are available for wider epidemiologic application. Not all of the methodologic problems have been resolved, however. It is important to develop further the statistical methods needed to handle recurrent events effectively and to make the methods more easily accessible to epidemiologists.

References

1. Cumming RG, Kelsey JL, Nevitt MC. Methodologic issues in the study of frequent and recurrent health problems: falls in the elderly. Ann Epidemiol 1990;1:49—56.
2. Eisen EA, Holcroft CR, Greaves IA, Wegman DH, Woskie SR, Monson RR. A Strategy to reduce healthy worker effect in a cross-sectional study of asthma and metalworking fluids. Am J Ind Med 1997;31:671—7.
3. Chen-Yeung M, Lam S, Koener S. Clinical features and natural history of occupational asthma due to western red cedar. Am J Ind Med 1982;72:411—6.
4. Redelmeier DA, Tibshirani RJ. Association between cellular-telephone calls and motor vehicle collisions. N Engl J Med 1997;336:453—8.
5. Eisen EA, Wegman DH, Kriebel D, Woskie SR, Hu X. An epidemiologic approach to the study of acute reversible health effects in the workplace. Epidemiology 1991;2:263—70.
6. Braun-Fahrlander C, Ackermann-Liebrich, Schwartz J, Gnehm HP, Rutishauser M, Wanner Hu. Air pollution and respiratory symptoms in preschool children. Am Rev Respir Dis 1992;145:42—7.
7. Schwartz J, Zeger SL. Passive smoking, air pollution, and acute respiratory symptoms in a diary study of student nurses. Am Rev Respir Dis 1990;141:62—7.
8. 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.
9. Mittleman MA, Maclure M, Tofler GH, Sherwood JB, Goldberg RJ, Muller JE. For the determinants of myocardial infarction onset study investigators: triggering of acute myocardial infarction by heavy exertion: protection against triggering by regular exertion. N Engl J Med 1993;329:1677—83.
10. Moulton HM, Dibley MJ. Multivariate time-to-event models for studies of recurrent childhood diseases. Int J Epidemiol 1997;26:1334—9.
11. Dwyer JH, Feinleib M, Lippert P, Hoffmeister H, editors. Statistical models for longitudinal studies of health. New York (NY): Oxford University Press, 1992.
12. Diggle PJ, Liang K-Y, Zeger SL. Analysis of longitudinal data. Oxford (UK): Clarendon Press, 1994

13. Hu FB, Goldberg J, Hedeker D, Flay BR, Pentz MA. Comparison of population averaged and subject specific approaches for analyzing repeated binary data. Am J Epidemiol 1998;147:694—703.
14. Korn EL, Whittemore AS. Methods of analyzing panel studies of acute health effects of air pollution. Biometrics 1979;35:795—802.
15. Liang K-Y and Zeger SL. Longitudinal data analysis using generalized linear models. Biometrika 1986;73:13—22.
16. Andersen PK, Gill RD. Cox's regression model for counting processes: a large sample study. Ann Stat 1982;10(4):1100—20.
17. Prentice RL, Williams BJ, Peterson AV. On the regression analysis of multivariate failure time data. Biometrika 1981;68:373—9.
18. Zeger SL, Liang KY. Feedback models for discrete and continuous time series. Stat Sin 1997;1:51—64.
19. Angrist JD, Imbens GW, Rubin DB. Identification of causal effects using instrumental variables. J Am Stat Assoc 1996;91(434):444—72.
20. Robins JM, Blevins D, Ritter G, Wulsohn M. G-estimation of the effect of prophylaxis therapy for pneumocystis carinii pneumonia on the survival of AIDS patients. Epidemiology 1992;3:319—36.