



Education corner

Exchangeability in the case-crossover design

Murray A Mittleman^{1,2*} and Elizabeth Mostofsky^{1,2}

¹Cardiovascular Epidemiology Research Unit, Department of Medicine, Beth Israel Deaconess Medical Center, Harvard Medical School, Boston, MA, USA and ²Department of Epidemiology, Harvard School of Public Health, Boston, MA, USA

*Corresponding author. Cardiovascular Epidemiology Research Unit, Beth Israel Deaconess Medical Center, 375 Longwood Avenue, Room 423, Boston, MA 02215, USA. E-mail: mmittlem@bidmc.harvard.edu

Accepted 13 March 2014

Abstract

In cohort and case-control studies, confounding that arises as a result of differences in the distribution of determinants of the outcome between exposure groups leading to non-exchangeability are addressed by restriction, matching or with statistical models. In case-only studies, this issue is addressed by comparing each individual with his/herself. Although case-only designs use self-matching and only include individuals who develop the outcome of interest, issues of non-exchangeability are identical to those that arise in traditional case-control and cohort studies. In this review, we describe one type of case-only design, the case-crossover design, and discuss how the concept of exchangeability can be used to understand issues of confounding, carryover effects, period effects and selection bias in case-crossover studies.

Key words: Epidemiology, case-crossover

Key Messages

- Several case-only study designs have been proposed, in which each individual is compared with his/herself, such as the case-crossover, case-time control, fixed-effects case-time control and the self-controlled case-series design.
- The hallmark of case-only designs is the use of self-matching in the design and analysis.
- Similar to other observational study designs, causal inference in case-only designs requires the assumption of exchangeability between exposure groups.
- The concept of non-exchangeability can be used to understand issues of confounding, selection bias, information bias, autocorrelation and carryover effects in case-only studies, and to identify ways to eliminate or minimize these potential sources of bias in the design and or/analysis.

Background

In observational studies examining whether an exposure causes an outcome, differences between the exposed and unexposed individuals with respect to their risk of developing the outcome may lead to confounding and a biased estimate of the association. In order to address this issue, several case-only study designs^{1–3} have been proposed, in which each individual is compared with his/herself, such as the case-crossover,⁴ case-time control,⁵ fixed-effects case-time control⁶ and the self-controlled case-series design.⁷ Although these approaches often use different statistical analyses to compute measures of association, they share similar assumptions about exchangeability.^{7,8}

Since the case-crossover design was first presented 20 years ago,⁴ it has been used in several fields to examine acute risks from transient exposures. The design is based on the theory behind prospective cohort studies (Figure 1, Panel A) and, since case-control studies use efficient sampling from the study base (Figure 1, Panel B),⁹ the case-crossover design also has a direct relationship to case-control studies. Therefore, the concerns about causality in case-crossover studies are identical to those of the more commonly used designs. However, since the case-crossover and other case-only designs compare the same person (or other unit of observation) at different times rather than different people at the same time, there has been some confusion about issues of causal inference in this setting. Several papers have reviewed the theory and practical issues that arise in case-crossover studies^{4,10,11} including issues of relative efficiency^{12,13} and sensitivity to exposure misclassification⁸ in comparison with standard observational study designs. In this paper, we explain the relation between cohort, case-control and case-crossover studies and discuss how the concept of exchangeability can be used to understand issues of confounding, carryover effects, period effects and selection and information bias in case-crossover studies. We present directed acyclic graphs¹⁴ to highlight the structure of non-exchangeability in these examples and we describe the analytical methods for estimating valid measures of association.

Counterfactual theory and exchangeability

In order to make causal inferences about the effect of an exposure in observational epidemiology, we wish to compare the risk of the outcome among the exposed with the risk of the outcome among those same people had they been unexposed. However, we usually cannot observe the same person under both exposed and unexposed scenarios; one outcome is factual (observed) and one is counterfactual (unobserved). In the real world, we compare exposed

individuals with those who are not exposed and, to make causal inference about the effect of the exposure on the outcome, we assume that the outcome among the unexposed represents the outcome that would have occurred among the exposed if exposure were removed. This assumption is called ‘exchangeability’; we must assume that people receiving a given level of exposure are exchangeable with those receiving other exposure levels. The findings are not considered causal if there is a lack of exchangeability, either because there are common causes of the treatment and the outcome (confounding) or if selection of the people included in the analysis is affected by both the exposure and the outcome or by causes or correlates of the exposure and the outcome (selection bias).¹⁵

In both cohort and case-control studies, we account for common causes of exposure and outcome (confounding) by restriction, matching and/or with statistical adjustment to achieve conditional exchangeability within levels of these measured confounders. By accounting for these measured confounders in the analysis, the findings may be causal with the assumption that there is no unmeasured-confounding. The greatest weakness of causal inference in observational studies is that it is impossible to be sure that all between-person differences have been balanced; since we cannot observe each participant’s counterfactual outcome, we cannot verify that the comparison groups are conditionally exchangeable.

Case-crossover design

The case-crossover design⁴ is best suited for studies of intermittent exposures with short induction times and transient acute effects. Similar to a crossover experiment, the same person provides information on outcome risk under both exposed and unexposed states, except that nature or the individual determines the times under exposure to a potential trigger rather than the investigator assigning exposure. In the case-crossover design, individuals within a cohort who experience the outcome event (cases) are identified, and information about each subject’s exposure during a hazard period prior to the event is compared with that individual’s exposure distribution at other times (control periods). This self-matching eliminates confounding (non-exchangeability) by determinants that are constant within individuals over the sampling period but often differ between study subjects. For instance in a case-crossover study, there is no confounding by sex because each person is compared with himself/herself. Therefore each stratum comprises only a male or female study participant. This characteristic of self-matched designs prevents the study of exposures that do not change within an individual over the study period. For instance, one cannot use self-matched

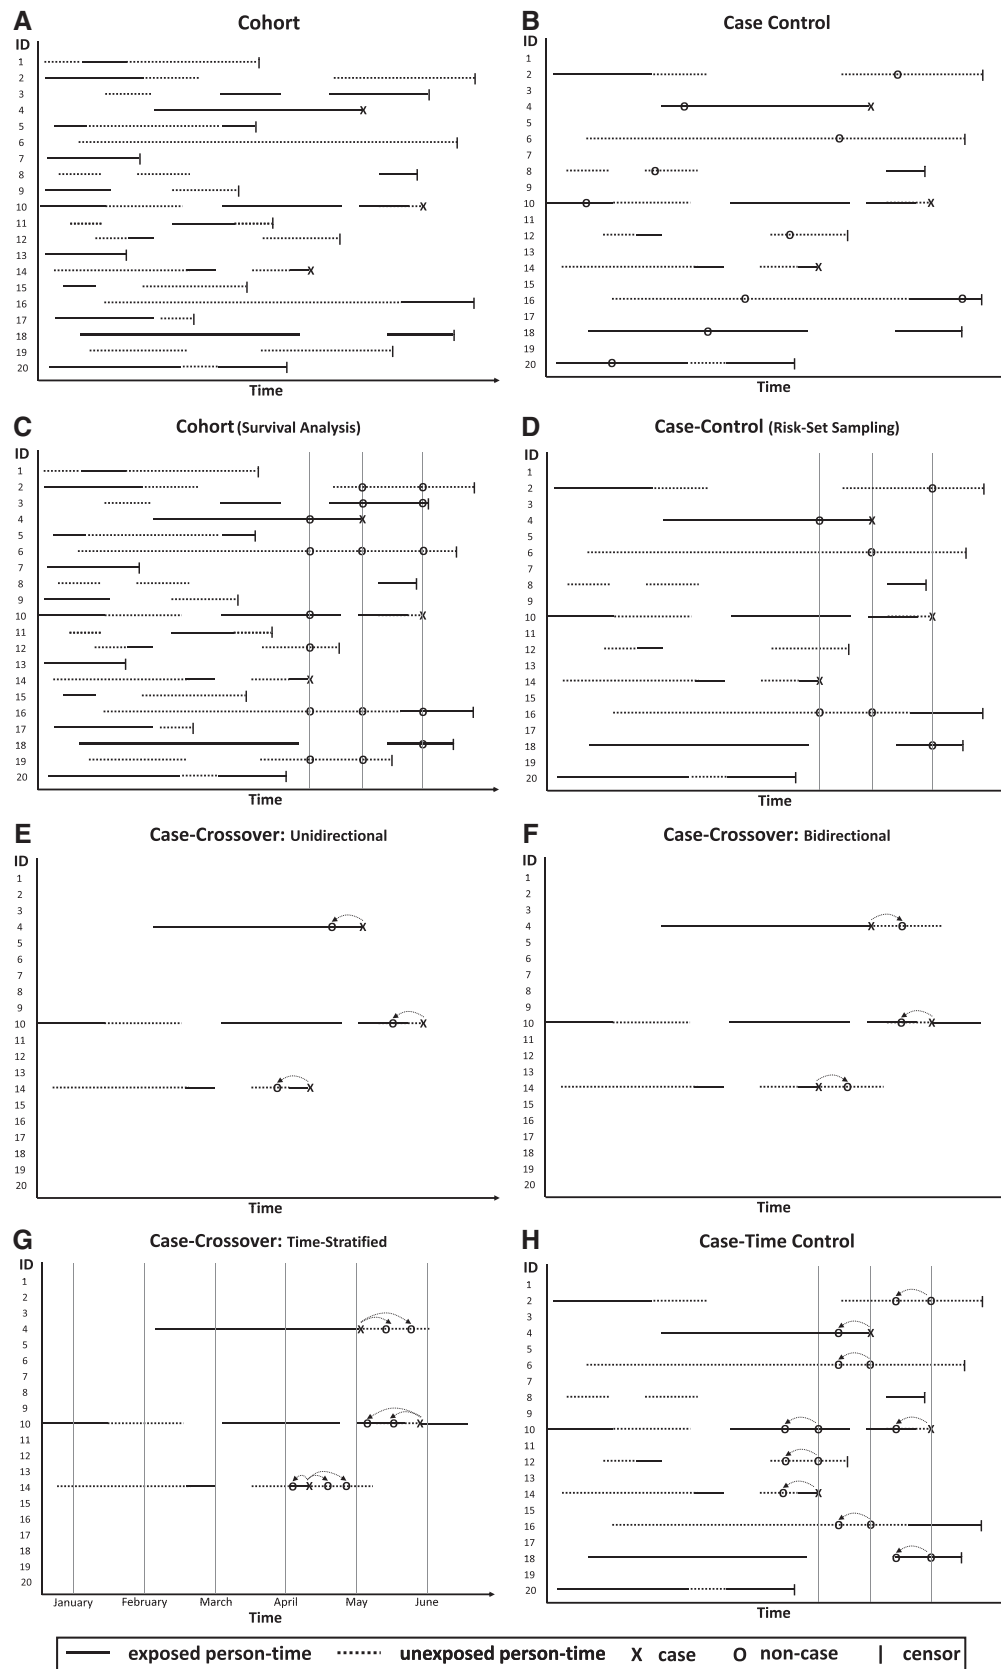


Figure 1. Study designs to examine the causal effect of an exposure on an outcome

(Continued)

designs to study smoking or drugs that are used continuously because there is no variation in exposure within strata (the individual), and the association of interest is not estimable. This is referred to in the literature as a lack of positivity necessary for causal inference.¹⁶

In a case-crossover study, only strata with variation in exposure contribute information; there must be some difference in exposure between the hazard and control period(s). This is analogous to a matched case-control study in which only matched pairs that are discordant for exposure contribute information. It is also analogous to a crossover experiment. If a subject drops out of the experiment before crossing to another treatment regime, data on a single exposure would not contribute information to the analysis. Since a person only contributes information to the estimated effect of exposure if there is variation in exposure over time, a case-crossover study isolates short-term effects of intermittent exposures from long-term effects of a constant exposure.

Different sampling mechanisms reflect how self-matched designs are analogous to either a highly stratified cohort or a case-control study.¹² Conceptualized as a cohort study, a case-crossover study can utilize all of the person-time experienced by each subject over a specified time period preceding the event (Figure 1, Panel E). For instance, to examine whether a cancer diagnosis triggers cardiovascular deaths,¹⁷ administrative data were obtained to estimate exposed and unexposed person-time over follow-up for each individual who died of a cardiovascular event during the follow-up period. Alternatively, exposed and unexposed person-time can be estimated using

questionnaires on the usual frequency of exposure. For example, to examine whether alcohol consumption triggers ischaemic stroke,¹⁸ we interviewed people upon hospital admission for stroke and asked about the usual frequency of alcohol consumption in the previous year to estimate the expected exposure frequency. The amount of exposed person-time was estimated by multiplying the usual frequency of exposure by the estimated usual duration, and unexposed person-time was estimated by subtracting the exposed person-time from the total duration of the specified time period. If an at-risk cohort is studied, this information could be collected at the start of follow-up, before any outcome events occur. To adjust for the self-matching in the case-crossover design, data from each individual are treated as if they were from a matched set. An incidence rate ratio can be calculated using Mantel-Haenszel estimation appropriate for sparse follow-up data,¹⁹ assuming pairwise exchangeability between the case period and the control period, i.e. under the null, the marginal exposure probability is the same in the case period and the control period.⁷ These data can also be analysed using Poisson^{20,21} or fixed-effects models.⁶

Alternatively, instead of using all of the available person-time for an individual, one could match on time. This is conceptually similar to the formation of risk sets at each time a case occurs in survival analysis of the full cohort (Figure 1, Panel C). Analogous to a highly stratified case-control study (Figure 1, Panel D), an efficient way to estimate the exposure distribution in the person-time giving rise to the case is to select the hazard period and matched control interval(s) for each individual (Figure 1, Panel E).

Figure 1. Continued

- In a full cohort study (Panel A), the entire sample is followed for event occurrence; assumption of exchangeability between exposed and unexposed person-time in the entire study base.
- In a case-control study with traditional incidence density sampling (Panel B), all or a random sample of the cases are identified and controls are sampled at random from the person-time at risk; assumption of exchangeability between exposed and unexposed person-time in the study base represented by the cases and sample of controls. Note that an individual can be randomly selected as a control and later become a case (ID#4) and an individual can be randomly selected as a control more than once (ID#16).
- Non-exchangeability arising from changes over time can be addressed by matching on time using survival analysis methods in a cohort study (Panel C) or risk-set sampling in a case-control study (Panel D); assumption of exchangeability between exposed and unexposed person-time at the time the case occurred.
- In a case-crossover study, non-exchangeability arising from slowly varying characteristics is eliminated by matching each case to himself/herself. If event occurrence affects subsequent exposure, a unidirectional case-crossover design (Panel E) is appropriate. Otherwise, a bidirectional case-crossover design (Panel F) can reduce exposure time trends by stratifying on the calendar month and then sampling person-time before and after the event (ID#4), before the event (ID#10) or after the event (ID#14). Assumption of exchangeability between exposed and unexposed person-time within each individual. Note that in the unidirectional design, individuals with no variation in exposure between the hazard and control periods do not contribute information to the estimator (ID#4).
- In a case-crossover study, non-exchangeability arising from changes over time can be addressed by matching on time or by conducting a case-time control study (Panel G) that incorporates a control group. At the time the case occurs, obtain exposure information on cases and on controls for hazard periods and for control periods; exchangeability between exposed and unexposed person-time within each individual after accounting for a time-trend that is assumed to be the same for cases and controls.

This allows for matching on temporal factors such as day of week or time of day. These data can be analysed using a Mantel-Haenszel estimator, assuming pairwise exchangeability between the case and each control period. With one control period per case, they can be analysed using conditional logistic regression. With two or more control periods per case, one must also assume global exchangeability between all control periods within matched sets;⁷ the data can be analysed using conditional logistic regression with certain case-crossover sampling techniques described below.^{22,23}

In a case-crossover study, the observed exposure frequency immediately prior to the event (hazard period) is compared with the expected exposure frequency based on the exposure distribution at times when the acute event did not occur (control period). This exposure odds ratio is an unbiased estimate of the incidence rate ratio that would have been observed if a full cohort study had been conducted.^{9,24} In order for this measure to validly estimate a causal effect of exposure on the outcome, the exchangeability assumption must be met. One must assume that the probability of the outcome is the same over time and that the probability of exposure in the case and control periods are conditionally independent apart from any causal effect of exposure. If someone were more or less likely to be exposed at times when he/she is more or less likely to experience the outcome, the analysis is not valid. Just as restriction, stratification and matching are used to attain conditional exchangeability between individuals in cohort and case-control studies, these methods can be used in case-crossover studies so that within-person comparisons are assumed to be conditionally exchangeable and causal inference about the short-term effect of triggers can be examined.

Confounding in case-crossover studies

Since each case serves as his or her own control, the self-matched design eliminates confounding by stable and slow-varying characteristics, whether measured or not. Therefore, over short time windows, the person-time during the hazard period is assumed to be exchangeable with that individual's person-time during the control period since the baseline risk is assumed to be constant (i.e. conditional exchangeability is achieved as long as the analysis accounts for within-person matching). However, confounding in a case-crossover study is still possible.

Confounding by transient co-exposures

During a given window of time, an individual may engage in several behaviours that potentially trigger the

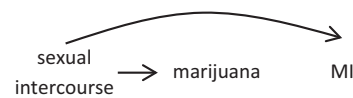


Figure 2. Compared with the frequency of exposure during the control period(s), there is a higher frequency of exposure to both sexual intercourse and marijuana during the hazard period. Assuming that sexual intercourse precedes marijuana use, the association between marijuana use and myocardial infarction may be biased.

outcome. For instance, in the Myocardial Infarction Onset study,²⁵ investigators examined whether there is an increased risk of myocardial infarction (MI) in the hour following marijuana use, and noted that, ‘three patients who smoked marijuana in the hour before their infarction also engaged in other triggering behaviours in that hour; one patient reported using cocaine in addition to smoking marijuana, another reported sexual intercourse and a third patient reported both sexual intercourse and cocaine use’. Therefore, the comparison will yield biased estimates if there are factors that increase MI risk and occur at different levels/frequencies between the hazard and control periods; instead of comparing the risk of MI following marijuana with the risk under an exchangeable time window had the person not smoked marijuana, other co-exposures may lead to a biased estimate (Figure 2).

To attain conditional exchangeability in the presence of co-exposures, one must stratify jointly on both the exposure and co-exposures. If data are available on the timing of each relevant co-exposure for each time period, exposure in the case period may be matched to control interval(s) with similar co-exposures, or statistical adjustment may be used. However, it is more difficult when information on exposure is obtained with questionnaires that ask about the usual frequency of exposure since it is often cumbersome to collect information on the usual frequency of engaging in several behaviours simultaneously. On the other hand, for studies examining exposures that are passively recorded for administrative or other purposes, such as vaccination records or environmental exposures, it may be feasible to have access to such data for longer follow-up periods.

Confounding by acute indication

In pharmacoepidemiological research, the greatest concern in cohort and case-control studies is that the indication for treatment may be related to the risk of future health outcomes; the resulting imbalance in the underlying risk between the treatment groups under comparison can result in biased estimates. Therefore, investigators attempt to control for confounding by indication as best as possible

by restriction, matching or using statistical models to adjust for measured factors that are related to both medication use and the outcome of interest; since measurement error or unmeasured or unknown risk factors often confound these studies, residual and unmeasured confounding often remain a concern.

In case-crossover studies examining the acute impact of transient changes in drug use, the self-matching eliminates confounding by indication by long-term health factors, including factors that cannot be measured. However, a lack of exchangeability in a case-crossover study can occur when there are transient indications for a drug. For instance, beta adrenergic agonist drugs are potent bronchodilators that are prescribed for acute treatment of asthma symptoms and exacerbations. In a study on the risk of death immediately following beta agonist use, it may be the indication for the drug rather than (or in addition to) the drug that poses the health risk. Therefore, for each individual, the person-time represented by the hazard and control periods are not exchangeable; an individual is more likely to be exposed (or exposed to higher levels) during periods of higher risk of the outcome than during other times, resulting in a biased estimate of the association (Figure 3). Conversely, individuals may be less likely to engage in strenuous physical activity when they experience asthma symptoms and therefore, in a study of physical activity as a trigger of death, the probability of exposure would be lower during the hazard period than during comparison times, resulting in a biased estimate (Figure 4).

To address the concern of transient changes in triggers in response to causes (or correlates) of the outcome, one could conduct analyses stratified by the cause (or correlate) of the outcome (e.g. asthma symptoms), but intractable confounding may arise.




Figure 3. Rather than the beta agonist acting alone to increase death risk, asthma exacerbation may lead to higher beta agonist use and it also increases the risk of death.




Figure 4. Rather than a causal effect of physical activity triggering death, asthma exacerbation may confound the association by resulting in both reduced physical activity and higher death risk soon after.

Selection bias in case-crossover studies

Selection biases are distortions that occur when the distribution of exposure and disease is different for the people and times included in the study than for all people and times that should have been theoretically eligible to participate. In case-control studies, selection bias occurs when the distribution of exposure and outcome is different for the people available and willing to participate compared with that for the underlying cohort of interest; in cohort studies, selection bias occurs when the people who are lost to follow-up or who experience a competing risk are different from those who remain in the cohort with respect to exposure and risk of the outcome. The self-matching in the case-crossover design eliminates concerns of selecting individuals who are not representative of the population that produced the cases, because the control periods reflect the cases' exposures themselves.⁴ However, selection bias in a case-crossover study is still possible.

Case selection bias

A patient's recent behaviours may impact on his/her willingness to participate. For instance, in a study of MI triggers, if someone were concerned about reporting recent illicit drug use, it may appear as if recent use does not increase MI risk when in fact usual users but not recent users were willing to participate in the study. This would result in a biased estimate of the effect of recent drug use as a trigger of MI onset because the prevalence of exposure among the included cases is lower than the prevalence among all cases (Figure 5).

There is also a possibility of bias caused by differential survival of cases who had an MI triggered by different mechanisms. For example, if patients whose MIs were triggered by physical activity were less likely to die than those whose infarctions were unrelated to physical activity, the apparent relative risk may be biased because the exposure distribution among prevalent cases at the time of the study recruitment is greater than the distribution among all cases. Since recently exposed individuals are more likely to survive long enough to participate, the exposure will appear harmful when in fact it lowered the risk of a fatal event. On the other hand, if cases who were physically active immediately prior to the MI were more likely to die




Figure 5. Among all people theoretically eligible to be included in the study, cases who agree to participate are less likely to have recently used an illicit drug; and, by definition, recent MI affects participation in the study.

of arrhythmia before reaching the hospital, the apparent relative risk will be biased since physical activity will appear less harmful, or even protective (Figure 6).

This is analogous to a prospective cohort study examining the relationship between physical activity and cardiovascular outcomes. If the people who are alive and available to participate are more likely to be regularly active and therefore survive long enough to experience cardiovascular events than all people who are theoretically eligible to be included in the study, exposure will appear harmful even though it increased survival.

Control time selection bias by sampling person-time dependent on exposure

In case-control studies, it is important to select controls independently of exposure so that the exposure distribution among controls is representative of the exposure distribution in the person-time giving rise to the cases (within strata of any matching factors); a person is eligible to serve as a control as long as he/she is at risk for the event.²⁶ Similarly, in case-crossover studies, it is important that the control time windows are selected independently of exposure so that the exposure distribution in the control periods is representative of the exposure distribution in the person-time giving rise to the cases (within strata of any matching factors, i.e the individual); as long as it is possible for an event to occur, the person-time can be included in the control period. For example, in a study of cellphone use and the risk of car collisions, one should not select control times when individuals are more or less likely to use a phone. Rather, one should sample person-times (within individuals) that represent the probability of exposure at times when a collision may occur (Figure 7).^{11,27}




Figure 6. Among all people theoretically eligible to be included in the study, cases who were physically active immediately prior to the MI are more likely to survive and participate in the case-crossover study than cases who were not recently physically active; and, by definition, recent MI affects participation in the study.



Figure 7. Among all of the person-times theoretically eligible to be included in the study, exposed person-time is more likely to be selected for the control period (oversample times when people are more likely to speak on the phone) or unexposed person-time is more likely to be selected for the control period (e.g. oversample times when people are less likely to speak on the phone). Therefore, control time is related to the probability of exposure and, by definition, it is directly related to the outcome.

Control time selection bias by restricting to the first/last exposure

Some exposures of interest may only occur once during the follow-up period (e.g. death of a spouse), so a person is only exposed in the control period, only exposed in the hazard period or exposed in neither period. On the other hand, some exposures may occur repeatedly, such as episodes of anger, marijuana use or a diagnosis of cancer. In these situations, restricting to the first or last exposure and ignoring other exposure episodes by counting them as unexposed times would induce a bias; instead of representing a sample of the person-time giving rise to the event, selecting a control period when the person is more likely to be exposed may lead to a downward bias in the estimate of the effect of the trigger on the outcome, and selecting a control period when the person is less likely to be exposed may lead to an upward bias.

Information bias in case-crossover studies

In a retrospective cohort or case-control study, exposure information is obtained after the outcome has occurred. This raises concerns that the outcome may impact on reporting, resulting in differential recall between cases and non-cases and different interpretation of the questions by the cases and non-cases.

In a case-crossover study using self-reported exposure, the same person provides exposure information for both case and control periods, so the interpretation is the same, but the questions about exposure during case and control intervals may have different wording and require different methods of memory recall.⁴ Recall bias may occur if the case over- or under-reports recent exposure compared with recall of exposure during the control period(s), especially if the control period is long (Figure 8).²⁸ The direction and magnitude of the problem depends on the exposure under study. Some exposures are major life events and are therefore likely to be recalled for many months, such as the death of a significant person in one's life.²⁹ For certain behaviours, people have a typical routine, so it may be easier to accurately recall daily habits of coffee consumption than behaviours that fluctuate in a less routine manner, such as outbursts of anger. The accuracy of self-reported number of days with a flu or cold in the past year may differ between the winter and summer seasons. A subject's




Figure 8. The recall of exposure is affected by recent symptom onset; patients may over-estimate or underestimate exposure in the hazard period compared with their recall of exposure during control period(s).

recall is also likely to be coloured by his/her perceived life-style; people may over-report their habitual frequency of vigorous physical activity and underestimate unhealthy behaviours such as binge drinking or anger. Although case-crossover studies often rely on self-reported exposure information, some have used objectively recorded clinical or administrative data to examine acute triggers using the case-crossover design, eliminating concerns of recall bias.

Autocorrelation

Autocorrelation occurs when the exposure or outcome risk in one time period is correlated with exposure levels in other time periods. The non-independence violates a key assumption of most statistical models used to analyse case-crossover data, resulting in a loss of statistical efficiency and, at times, it may induce bias.^{12,30}

Autocorrelation between control periods

Since the correlated exposure levels are not related to the case period, non-independence between exposure levels of several control periods does not result in a lack of exchangeability but it results in lower efficiency of the estimator. In a matched cohort or case-control study, the efficiency of the estimator is defined by the number of matched sets with discrepant exposure levels; therefore, exposure autocorrelation will result in lower precision since there will be fewer individuals with discrepant exposure levels between periods, and therefore fewer individuals contributing to the measure of association.

Autocorrelation between case and control periods

In a case-crossover study, exposure during the hazard period must be independent of exposure during the control period(s). Otherwise, exposures in the distant past could be the cause of recent disease onset rather than the hypothesized exposure during the hazard period (carryover effect). For instance, in a study examining triggers of traumatic injuries in the work environment,³¹ investigators planned to compare exposure to several potential triggers in the 10 min before the injury with exposure during a control period 60–70 min before the injury. However, since the average duration of many of the potential triggers was greater than 90 min, exposure during the control period inherently impacts on the exposure during the hazard period. A similar issue occurs in studies of drugs that are used regularly, such as those that are used to treat frequently occurring symptoms, such as migraine or asthma. In the setting of case-crossover sampling with discrete control intervals, if there is a positive or inverse association between exposure in the control period and exposure in the hazard period, sampling control periods with

exposures that are correlated with exposure during the hazard period will result in selection bias and a lack of exchangeability, leading to an estimate that is biased upward or downward. As noted previously,³⁰ exposure to such a drug in one time window will likely be correlated with drug exposure in the next time window. Analogous to the importance of a sufficiently long wash-out period between treatments in a crossover experiment, autocorrelation in exposure between periods in case-crossover studies is avoided by allowing for sufficient spacing between the selected time windows. Expert subject knowledge about the dynamics of the exposure is needed to identify sufficiently long wash-out periods.

Autocorrelation between outcomes

In a case-crossover study, one assumes that within an individual, repeated events are independent of each other. If multiple outcome events can occur within an individual, bias may arise if a single exposure can trigger multiple outcomes during the hazard period and these outcomes are considered independent events. For example, a recent case-crossover study examined the association between a new prescription for a diuretic and the occurrence of falls among nursing-home residents.³² If an individual fell several times on the same day, it would be incorrect to consider each of these falls as an independent event since all may be due to one prescription change. If there were truly no association between beginning a diuretic and the occurrence of falls, this would lead to increased variance; if there were an association, assuming that all of these falls were independent events would lead to an estimate that is biased away from the null. To avoid bias in studies where multiple outcome events can occur within an individual, it is crucial to assure that outcome events are truly independent of one another by selecting sufficiently wide time windows between events.

Time trends in exposure and/or outcome

Since the measure of association in a case-crossover study involves the comparison of exposure during the hazard period with the exposure distribution during the control period(s), time trends may introduce non-exchangeability between the time periods under comparison. Whether the time trend induces confounding, selection bias or both depends on whether there are time trends in the exposure, the outcome, both exposure and outcome or factors related to exposure and/or the outcome. It also depends on the time-scale; systematic exposure patterns may vary over the hours in a day (e.g. coffee intake, circadian variation in cardiovascular risk), the days in a week (e.g. weekday traffic pollution), the season (e.g. influenza) or the year



Figure 9. Over time, there are changes in the probability of exposure, and in a unidirectional case-crossover study, the selected hazard period is always later in time. Therefore, the person-time selected for the control period has a different probability of exposure from that of the selected hazard period.

(e.g. adoption of a new drug or health practice, increasing age of a closed cohort). In most situations, we can attain conditional exchangeability by modelling, stratifying or matching on time to the extent that the temporal variation is captured.

Time trends in exposure

The case-crossover design was initially used to examine the short-term risk of MI following behavioural factors such as physical activity³³ and cocaine use.²⁵ Since experiencing an MI is likely to impact on subsequent behaviours, control periods cannot be selected after the event occurs since this may induce a reverse causation bias. Therefore, in this setting, the selected control periods must always precede the hazard period (unidirectional sampling; Figure 1, Panel E). If there is no change in outcome risk over time, there may be no confounding by time,³⁴ but changes in the probability of exposure between the hazard and control period(s) would induce non-exchangeability due to selection bias since the investigator selects control times that do not represent the exposure distribution in the underlying study base (Figure 9).

The case-time-control design⁵ is an attempt to account for non-exchangeability due to time trends in exposure using information on exposure trends obtained from a sample of conventionally sampled control individuals (Figure 1, Panel H). This design assumes that there are no unmeasured time-varying factors that are associated with both the exposure of interest and the outcome (confounding) and/or that modify the effect of the exposure on the outcome.^{34,35} A more commonly used approach is to sample matched pair intervals to address short-term time trends and select control periods close in time to case periods to address long-term time trends.

Time trends in disease risk

If there are changes over time in the risk of developing the outcome of interest but the probability of exposure remains stable, the results are valid. For instance, intermittent use of common non-prescription analgesics may be fairly stable over time within an individual. If a study were conducted with a long control period, the risk of the outcome may have increased, but the results would still be unbiased (Figure 10). This is analogous to using proxy con-



Figure 10. If there are no changes over time in the probability of exposure but there are changes over time in the probability of the outcome, the person-time selected for the control period has a different probability of outcome risk compared with the selected hazard period, but the exposure distribution is correctly represented.



Figure 11. The person-time selected for the control period has a different probability of exposure and a different probability of outcome risk compared with the selected hazard period.

trols in a case-control study; the exposure distribution in the controls accurately represents the exposure distribution in the person-time giving rise to the cases.^{24,36}

Time trends in factors related to exposure and outcome

In a case-crossover study, non-exchangeability may arise if the distribution of exposure and the risk of the outcome changes between the hazard and control periods. For instance, in a study of coffee consumption and short-term risk of ischaemic stroke,¹⁸ individuals are more likely to be exposed at times when they are more likely to experience an event because there is a circadian peak of stroke onset in the morning hours and people tend to drink more coffee in the morning than during the rest of the day. Therefore, time is associated with both exposure and outcome, resulting in confounding, and selecting control periods in the presence of time trends related to both exposure and outcome induces a selection bias (Figure 11).

This non-exchangeability can be reduced or eliminated by conditioning on time. For example, the investigator can collect data on pair-matched intervals¹² to compare exposure at the same time of day as the hazard period.

If experiencing the outcome does not affect subsequent exposure—for instance, having an MI today does not affect tomorrow's air pollution levels—either pre-event times or post-event times can be selected as control periods. In a case-crossover study, control periods are selected to represent the exposure distribution in the person-time giving rise to the cases, but for events that can occur only once, such as death, control times selected after the event are by definition times when the individual is not at risk of the event. However, since individual events do not affect the distribution of future exposure in the overall study population, selecting post-event control times is acceptable. This is similar to selecting control individuals not at

risk of the event but with similar exposure distribution as the cases in a case-control study ('proxy controls')^{24,36} and, with minimal additional assumptions, allows for the estimation of the expected exposure distribution at the time the case occurred, accounting for short-term time trends in exposure.

In a case-crossover study with semi-symmetrical bidirectional sampling,³⁷ either the period before or the period after the event occurs is selected at random (Figure 1, Panel F), whereas in a typical bidirectional time-stratified case-crossover design,³⁰ the investigator may select control intervals on the same day of the week from the same calendar month as each case interval (Figure 1, Panel G). A special case of the time-stratified design is the full stratum bidirectional design,³⁸ where the control period includes all days in the exposure series other than the index day. These three designs avoid the overlap bias that results from selecting control periods as a function of the event times.^{21–23,39,40} Furthermore, this reduces non-exchangeability by seasonal and long-term time trends; since all times under comparison for each individual are within the same month, it minimizes trends in exposure, outcome risk or confounders that do not change greatly within a month. Also, if an investigator selects control periods that are 7 days apart, the time-stratified design eliminates exposure autocorrelation shorter than a week and it conditions for exposure and outcome patterns that may have a circaseptan periodicity. Poisson^{20,21,41} and fixed-effects⁶ approaches are free from overlap bias, but require modelling assumptions to account for confounding by time trends.

Conclusion

In all epidemiological studies, causal inference requires that within strata, the people or person-time under comparison are exchangeable. The case-crossover design provides a useful tool for examining triggers of acute outcomes since the self-matching assures that people are comparable with themselves with respect to static or slow-varying factors. To assure exchangeability between the time periods under comparison, it is important to consider issues of confounding, selection bias and autocorrelation. Conditional exchangeability in a case-crossover study can usually be obtained by properly matching or stratifying by time and conducting an analysis that takes the self-matching into account. Control periods must be selected close enough to the time of the case event so that the assumption of exchangeability is met but separate from the hazard period such that exposure is independent of the case event, to prevent short-term autocorrelation within an individual and carryover effects.

Funding

This work was supported by grants from the National Heart, Lung, and Blood Institute (HL-T32-098048 and HL-F32-120505), the US

Environmental Protection Agency (USEPA) (RD-83479801) and the National Institute of Environmental Health Sciences (P01-ES009825).

No funding organization had any role in the: design and conduct of the study; collection; management, analysis and interpretation of the data; or preparation of the manuscript. Its contents are solely the responsibility of the grantee and do not necessarily represent the official views of the funders. Further, USEPA does not endorse the purchase of any commercial products or services mentioned in the publication.

Conflict of interest: None declared.

References

1. Becker NG, Salim A, Kelman CW. Analysis of a potential trigger of an acute illness. *Biostatistics* 2006;7:16–28.
2. Feldmann U. Design and analysis of drug safety studies, with special reference to sporadic drug use and acute adverse reactions. *J Clin Epidemiol* 1993;46:237–44.
3. Miettinen OS, Caro JJ. Principles of nonexperimental assessment of excess risk, with special reference to adverse drug reactions. *J Clin Epidemiol* 1989;42:325–31.
4. 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.
5. Suissa S. The case-time-control design. *Epidemiology* 1995;6:248–53.
6. Allison PD, Christakis NA. Fixed-effects methods for the analysis of nonrepeated events. *Sociological Methodology* 2006;36:155–72.
7. Vines SK, Farrington CP. Within-subject exposure dependency in case-crossover studies. *Stat Med* 2001;20:3039–49.
8. Greenland S. A unified approach to the analysis of case-distribution (case-only) studies. *Stat Med* 1999;18:1–15.
9. Miettinen O. Estimability and estimation in case-referent studies. *Am J Epidemiol* 1976;103:226–35.
10. Maclure M. 'Why me?' versus 'why now?' - differences between operational hypotheses in case-control versus case-crossover studies. *Pharmacoepidemiol Drug Saf* 2007;16:850–53.
11. Maclure M, Mittleman MA. Should we use a case-crossover design? *Annu Rev Public Health* 2000;21:193–221.
12. Mittleman MA, Maclure M, Robins JM. Control sampling strategies for case-crossover studies: an assessment of relative efficiency. *Am J Epidemiol* 1995;142:91–98.
13. Whitaker HJ, Hocine MN, Farrington CP. The methodology of self-controlled case series studies. *Stat Methods Med Res* 2009;18:7–26.
14. Hernán MA, Hernández-Díaz S, Werler MM, Mitchell AA. Causal knowledge as a prerequisite for confounding evaluation: an application to birth defects epidemiology. *Am J Epidemiol* 2002;155:176–84.
15. Greenland S, Robins JM. Identifiability, exchangeability, and epidemiological confounding. *Int J Epidemiol* 1986;15:413–19.
16. Westreich D, Cole SR. Invited commentary: positivity in practice. *Am J Epidemiol* 2010;171:674–77; discussion 78–81.
17. Fang F, Fall K, Mittleman MA *et al.* Suicide and cardiovascular death after a cancer diagnosis. *N Engl J Med* 2012;366:1310–18.

18. Mostofsky E, Burger MR, Schlaug G, Mukamal KJ, Rosamond WD, Mittleman MA. Alcohol and acute ischemic stroke onset: the stroke onset study. *Stroke* 2010;**41**:1845–49.
19. Greenland S, Robins JM. Estimation of a common effect parameter from sparse follow-up data. *Biometrics* 1985;**41**:55–68.
20. Farrington CP. Relative incidence estimation from case series for vaccine safety evaluation. *Biometrics* 1995;**51**:228–35.
21. 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**:186–92.
22. Janes H, Sheppard L, Lumley T. Case-crossover analyses of air pollution exposure data: referent selection strategies and their implications for bias. *Epidemiology* 2005;**16**:717–26.
23. Janes H, Sheppard L, Lumley T. Overlap bias in the case-crossover design, with application to air pollution exposures. *Stat Med* 2005;**24**:285–300.
24. Wacholder S, McLaughlin JK, Silverman DT, Mandel JS. Selection of controls in case-control studies. I. Principles. *Am J Epidemiol* 1992;**135**:1019–28.
25. Mittleman MA, Mintzer D, Maclure M, Tofler GH, Sherwood JB, Muller JE. Triggering of myocardial infarction by cocaine. *Circulation* 1999;**99**:2737–41.
26. Rothman KJ, Greenland S, Lash TL. *Modern Epidemiology*. 3rd edn. Philadelphia, PA: Lippincott Williams & Wilkins, 2008.
27. Maclure M, Mittleman MA. Cautions about car telephones and collisions. *N Engl J Med* 1997;**336**:501–02.
28. Moller J, Hessen-Soderman AC, Hallqvist J. Differential misclassification of exposure in case-crossover studies. *Epidemiology* 2004;**15**:589–96.
29. Mostofsky E, Maclure M, Sherwood JB, Tofler GH, Muller JE, Mittleman MA. Risk of acute myocardial infarction after the death of a significant person in one's life: the Determinants of Myocardial Infarction Onset Study. *Circulation* 2012;**125**:491–96.
30. Lumley T, Levy D. Bias in the case-crossover design: implications for studies of air pollution. *Environmetrics* 2000;**11**:689–704.
31. Sorock GS, Lombardi DA, Gabel CL, Smith GS, Mittleman MA. Case-crossover studies of occupational trauma: methodological caveats. *Inj Prev* 2001;**7**(Suppl 1):i38–42.
32. Berry SD, Mittleman MA, Zhang Y *et al*. New loop diuretic prescriptions may be an acute risk factor for falls in the nursing home. *Pharmacoepidemiol Drug Saf* 2012;**21**:560–63.
33. Mittleman MA, Maclure M, Tofler GH, Sherwood JB, Goldberg RJ, Muller JE. Triggering of acute myocardial infarction by heavy physical exertion. Protection against triggering by regular exertion. Determinants of Myocardial Infarction Onset Study Investigators. *N Engl J Med* 1993;**329**:1677–83.
34. Greenland S. Confounding and exposure trends in case-crossover and case-time-control designs. *Epidemiology* 1996;**7**:231–39.
35. Suissa S. The case-time-control design: further assumptions and conditions. *Epidemiology* 1998;**9**:441–45.
36. Miettinen OS, Cook EF. Confounding: essence and detection. *Am J Epidemiol* 1981;**114**:593–603.
37. Navidi W, Weinhandl E. Risk set sampling for case-crossover designs. *Epidemiology* 2002;**13**:100–05.
38. Navidi W. Bidirectional case-crossover designs for exposures with time trends. *Biometrics* 1998;**54**:596–605.
39. Austin H, Flanders WD, Rothman KJ. Bias arising in case-control studies from selection of controls from overlapping groups. *Int J Epidemiol* 1989;**18**:713–16.
40. Mittleman MA. Optimal referent selection strategies in case-crossover studies: a settled issue. *Epidemiology* 2005;**16**:715–16.
41. Navidi W. Poisson regression and the case-crossover design: similarities and differences. *Commun Stat Theory Methods* 2008;**37**:213–20.