American Journal of Epidemiology
Vol. 133, No. 2

Printed in U S A

The Case-Crossover Design: A Method for Studying Transient Effects on the Risk of Acute Events

Malcolm Maclure

A case-control design involving only cases may be used when brief exposure causes a transient change in risk of a rare acute-onset disease. The design resembles a retrospective nonrandomized crossover study but differs in having only a sample of the base population-time. The average incidence rate ratio for a hypothesized effect period following the exposure is estimable using the Mantel-Haenszel estimator. The duration of the effect period is assumed to be that which maximizes the rate ratio estimate. Self-matching of cases eliminates the threat of control-selection bias and increases efficiency. Pilot data from a study of myocardial infarction onset illustrate the control of within-individual confounding due to temporal association of exposures. *Am J Epidemiol* 1991:133:144–53.

case-control studies; crossover studies; epidemiologic methods; statistics

To test hypotheses concerning the immediate determinants of myocardial infarction, my colleagues and I launched a case-control study ("The Onset Study"). We were prompted by confirmation of the morning peak in the circadian pattern of myocardial infarction incidence (1), which cast doubt on the widespread belief that the vast majority of infarctions occur purely at random (2). We hypothesized that many myocardial infarctions are triggered by activities during the hours or days before (3).

The choice of control group was not straightforward. Healthy representatives of the general population are no longer easy to recruit in the Boston area. Participation rates have dropped from 90 percent in the late 1960s (4) to 60 percent in the 1980s (5). Moreover, interviews with healthy subjects would tend to be scheduled at less hectic, less active times. On the other hand, almost any hospital control group would also suffer selection bias: many emergency hospitalizations (fractures, lacerations, acute gallbladder disease) are due to risk factors that are of interest to the Onset Study, such as physical exertion, alcohol intake, anger, and heavy eating; hospitalizations for chronic illnesses would be preceded by periods of atypical activities and exposures.

Returning to first principles, as defined by the case-base paradigm (6), we asked the question, who would be the best representatives of the population base that produced the cases? A simple answer was the cases themselves. This led us to develop the case-crossover design.

METHOD

Relation to other crossover studies

According to the case-base paradigm, a case-control study is a retrospective follow-

From the Department of Epidemiology, Harvard School of Public Health, Boston, MA.

Reprint requests to Dr. Malcolm Maclure, Department of Epidemiology, Harvard School of Public Health, 677 Huntington Avenue, Boston, MA 02115.

Supported by grant no. HL 41016 from the National Heart Lung and Blood Institute.

The author is grateful to Drs. James Muller and Chung Hsieh for their helpful comments on the manuscript and to Susan Hankinson for help in preparing the grant proposal.

TABLE 1. Association of theophylline use with shortness of breath on exertion, after 10 weeks of treatment or placebo, in a randomized double-blind crossover experiment with one subject, McMaster University, 1985 (adapted from reference 8)

	Theophylline	Placebo
Severe shortness of breath on exertion	2	0
No. of 10-week periods of pill use	2	2
Risk	1.0	0.0
Risk ratio	∞	

up study employing a sample, rather than a full census, of the population base. It follows that most types of follow-up study should have case-control counterparts. The case-crossover design is the counterpart to a cohort study with crossover of subjects between periods of exposure and nonexposure.

The term "crossover" is mainly used to describe experiments in which all subjects pass through both the treatment and placebo phases. In these studies, each subject serves as his or her own control—"the ultimate form of statistical adjustment" (7) for confounding by constant subject characteristics. However, "crossover" need not imply randomization. It is an apt term for intermittent exposure to factors with transient effects, because the subjects cross back and forth between periods of different risk.

The relation of the case-crossover design to a prospective crossover study can be explained most simply by considering a real crossover experiment involving only one subject—an asthmatic given theophylline or placebo according to a randomized, double-blind regimen (8). Table 1 (adapted from table 1 in reference 8) shows that severe shortness of breath on exertion was actually more common in this patient after 10 weeks of using theophylline than after 10 weeks of placebo. Suppose the outcome were myocardial infarction instead of shortness of breath. The experiment would be infeasible because it would be unethical and because the probability of an infarction would be extremely low.

Similar data to those in table 1 could be assembled retrospectively, however, as shown in tables 2 and 3. If the same patient were to suffer an infarction and be interviewed as a subject in a case-crossover study, he might report taking theophylline twice per day during the past month, the last time being less than 2 hours before the infarction. To estimate the relative risk of infarction over a 2-hour period following use of theophylline, his data would be tabulated as in table 2. Table 3 shows how the data would appear if his last dose of theophylline had been more than 2 hours before infarction onset.

These must be regarded as case-control data because information on the patient's usual frequency of theophylline intake are

TABLE 2. Hypothetical retrospective data from a patient who reported using theophylline twice per day and suffered a myocardial infarction less than 2 hours after last theophylline use

	Within 2 hours of last theophylline use*	
	Yes	No
Myocardial infarction	1	0
Person-hours	4	20
Rate ratio	∞	

* It is assumed arbitrarily that the theophylline effect on risk of infarction lasts only 2 hours.

TABLE 3. Hypothetical retrospective data from a patient who reported using theophylline twice per day and suffered a myocardial infarction more than 2 hours after last theophylline use

	Within 2 hours of last theophylline use*	
	Yes	No
Myocardial infarction	0	1
Person-hours	4	20
Rate ratio	0	

* It is assumed arbitrarily that the theophylline effect on risk of infarction lasts only 2 hours.

based on only a truncated follow-up interval—just a sample of the much longer time he would have been followed in a prospective cohort study. On the other hand, they resemble cohort data because the control data have units of person-time, not counts, as in the usual case-control study. The parameter estimated in the analysis of case-crossover data is therefore the average incidence rate ratio. To adjust for the self-matching of cases, data from each subject are treated as if they were from a matched pair. Methods of rate ratio estimation appropriate for sparse follow-up data are used (9). Thus, the analysis of a case-crossover study with 100 cases may be viewed as a pooled analysis of 100 retrospective cohort studies each with a sample size of one.

Mantel-Haenszel estimator

Table 2 can be modeled as an independent binomial observation with a fixed "success" probability equal to the proportion of time the person usually spends in the "effect-period" following recurrent exposure episodes. (In tables 2 and 3, the effect-period is arbitrarily assumed to be 2 hours.) Greenland and Robins (9) have shown that the Mantel-Haenszel estimator of the rate ratio (RR_{MH}), and a corresponding estimator of the variance of its logarithm, are unbiased for sparse person-time data.

Calculation of the rate ratio can be illustrated using the example of the asthmatic patient. If the coincidence of theophylline intake and myocardial infarction (table 2) were purely due to chance, then the asthmatic patient would have been about five times more likely to say his last use of theophylline was more than 2 hours before his infarction (table 3). These odds are reflected in the control data in both tables: the total time per day within 2 hours after theophylline intake is 4 hours, and the remainder equals 20 hours. His after-the-fact odds of exposure during the 2 hours before the infarction would be either infinity (table 2) or zero (table 3). These two tables are equivalent to the stratified tables for discordant pairs in a matched-pair cohort study, in each of which the rate ratios are either infinity or zero.

Under the null hypothesis, table 3 is expected to occur five times more often than table 2. Table 2 would add 20 to the numerator of the rate ratio ($\sum_i a_i N_{0i}/T_i$, in Rothman's notation (10)) and zero to its denominator ($\sum_i b_i N_{1i}/T_i$). Five tables like table 3 would add five zeros to the numerator and five fours to the denominator. Thus, under the null hypothesis, the rate ratio would equal 1.0. This is directly analogous to the Mantel-Haenszel odds ratio for a matched-pair case-control study, which is algebraically equivalent to the McNemar estimate—the number of discordant pairs with exposed cases divided by the number of discordant pairs with nonexposed cases (10). Table 2 would add five, ($N_{1i}N_{0i}/T_i$), to the numerator of the variance and 20, ($a_i N_{0i}/T_i$), to its denominator, while each table 3 would add five to the numerator and four, ($b_i N_{1i}/T_i$), to the denominator (9, 10).

Table 4 displays selected data from the pilot phase of the Onset Study. For the one-hour period after sexual activity, the rate ratio is 22 (95 percent confidence interval (CI) 3.2–160). For the one-hour period after drinking coffee, the rate ratio is 1.8 (95 percent CI 0.35–9.8). The table is limited to 10 subjects for simplicity; after 300 interviews, the rate ratio for sexual activity within 2 hours of infarction was 2.6 (95 percent CI 0.92–7.1).

The effect-period

If an effect is transient, the risk interval after an exposure may be divided into periods of normal background risk and the period of excess or reduced risk due to the exposure. If there is a delay before impact of a point exposure, and/or there is "carry-over"—a delay before recovery from the impact—then the period of altered risk is not exactly concurrent with the period of exposure. We therefore define the effect-period, the period of altered risk in a population, to be the difference between the min-

TABLE 4. Selected data from the pilot phase of the Onset Study, a case-crossover study of the determinants of myocardial infarction (MI) onset, Harvard University, 1989

Subject	Sexual activity during past year				Coffee drinking during past year			
	Last time before MI	Usual frequency	Concurrence odds*		Last time before MI	Usual frequency	Concurrence odds	
			Observed	Expected			Observed	Expected
1	5 minutes	1/year	1:0	1:8,765	9 hours	2/day	0:1	730:8,036
2	90 minutes	2/week	0:1	104:8,662	20 minutes	1/day	1:0	365:8,401
3	2 days	2/month	0:1	24:8,742	3 hours	3/month	0:1	36:8,736
4	3 days	1/week	0:1	52:8,714	22 hours	5/day	0:1	1,820:6,946
5	7 days	2/week	0:1	104:8,662	6 hours	8/day	0:1	2,920:5,846
6	11 days	3/month	0:1	36:8,730	7 hours	2/month	0:1	24:8,742
7	14 days	2/month	0:1	24:8,742	12 hours	2/day	0:1	730:8,036
8	21 days	2/month	0:1	24:8,742	5 hours	2/day	0:1	730:8,036
9	35 days	2/month	0:1	24:8,742	<1 hour	10/day	1:0	3,650:5,116
10	20 years	0/year	0:1	0:8,766	24 hours	1/day	0:1	365:8,401
Mantel-Haenszel estimate of relative risk			Numerator:	8,765			Numerator:	13,517
			Denominator:	392			Denominator:	7,355
			Ratio:	= 22			Ratio:	= 1.8
(95 percent CI†)				(3.2–160)				(0.35–9.8)

* The observed concurrence odds (1:0 or 0:1) are the odds that exposure was less than one hour before onset of myocardial infarction. The expected concurrence odds are the odds that a random event during the past year would have fallen within one hour after an episode of exposure (assuming any two exposure episodes are more than one hour apart). The effect-period after the hypothesized trigger is here assumed to be one hour long, with a minimum induction time of zero

† CI, confidence interval.

imum delay before impact and the maximum carry-over time.

One hypothesized mechanism of myocardial infarction onset, for example, begins with a short-lived triggering event (perhaps a burst of anger) which causes a surge in blood pressure and rupture of an atherosclerotic plaque in the heart (2, 11). A delay follows during which a clot grows. At the point when the clot begins to cause ischemia in the surrounding tissue, the subject enters a period of excess risk of infarction. This continues until the clot has dissolved, or the ischemia is relieved by increased blood supply through neighboring arteries, or an infarction occurs. An individual at risk of myocardial infarction may pass through many periods of excess risk of varying severity and duration without infarcting. The infarction occurs at the confluence of additional stressors, such as increased platelet aggregability, and decreased fibrinolysis causing occlusive thrombosis, or vasoconstriction causing complete occlusion of an artery partially obstructed by plaque and thrombosis. These and other physiologic factors are hypothesized to cause individual induction times (the time from exposure to infarction) to range from a few minutes to several days.

In a population, the times of myocardial infarction onset after a triggering event would be distributed as an epidemic curve (figure 1). The shape of the curve would be determined by the distribution of individual induction times. Although the epidemic curve would rise and fall smoothly in an infinitely large population, its shape would be less distinct in a study population of finite size. It might therefore be described roughly as a step-function: After the exposure (x), the background incidence rate prevails during what is commonly called the induction time (I_x). (This should be considered the minimum induction time because it corresponds to the minimum individual induction time in the population, not the average individual induction time.) In figure 1, the induction period is followed by an extremely hazardous effect-period (duration = E_{x1}) when the incidence rate is highly elevated. Next there is a moderately hazardous effect-period (duration = E_{x2}) when the incidence rate is slightly elevated. Finally, there follows a return to the background incidence rate. This step-function is illustrated by the broken line in figure 1.

Simplifying the epidemic curve to a step-function is appropriate when the follow-up interval is short. For example, if a man is told he has twice the risk of an infarction on the day after sex, he is unlikely to care whether the increase in risk is concentrated in the morning or the afternoon. The shape of the epidemic curve is of little interest. For the same reason, the "attack rate"—the cumulative incidence—is of greater interest when studying the epidemiology of acute disease, whereas in chronic disease epidemiology the incidence rate is preferred because the epidemic curve extends over many years (12). Therefore, although the measure of association in the case-crossover design is computed as an average incidence rate ratio, we will hereafter interpret it as estimating the relative risk for a particular effect-period.

The height of each step in figure 1 clearly depends on the width of each step, which is chosen arbitrarily. Thus, the estimate of the relative risk depends directly on the assumed duration of the effect-period. In a re-analysis of the pilot data in table 4, after assuming an effect-period of 2 hours instead of one hour, the relative risk of infarction after drinking coffee went from 1.8 to 0.65, because the ratio of affected to unaffected population-time was halved. The relative risk of infarction after sexual activity went from 22 to 30 despite the decrease in the ratio of affected to unaffected population-time. This was because subject 2, who had his infarction 90 minutes after sexual activity, was now assigned an after-the-fact "concurrence odds" of 1:0, instead of 0:1 (i.e., the odds of occurrence of sexual activity within 2 hours before the infarction.)

The actual duration of the effect-period can be inferred empirically in the same way the minimum induction period is normally inferred in cancer epidemiology—by examining the change in magnitude of the relative risk under different assumptions about du-

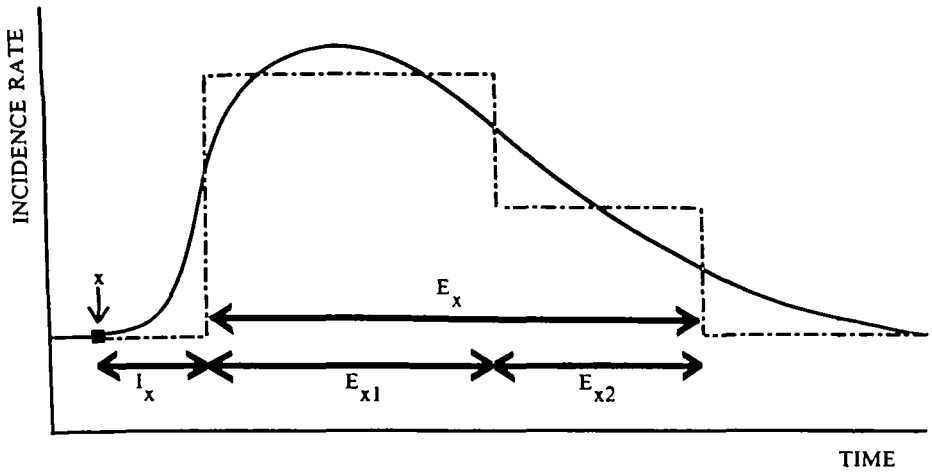

FIGURE 1. An epidemic curve (solid line) for acute-onset disease following a point exposure (x) with a transient effect. The step-function (dotted line) is a simplification of the curve showing an estimate of the population induction time (I_x) and the effect-period (E_x), comprising a high-risk effect-period (E_{x1}) and a moderate-risk effect-period (E_{x2}).

ration. Overestimation or underestimation of the duration results in nondifferential exposure misclassification, only diluting the association. The best estimate of duration is the one with minimal nondifferential misclassification, i.e., the one that maximizes the relative risk estimate (13). In the Appendix, the role of the assumed duration of the effect-period and the induction period in the estimation of the relative risk is presented algebraically.

DISCUSSION

Threats to validity

Five factors have been identified as threats to the validity of crossover studies: 1) carryover and period effects, 2) treatment sequencing and patient assignment, 3) crossover rules and timing, 4) dropouts and faulty or ou lying data, and 5) inappropriate statistical analyses for repeated outcomes (7). In the case-crossover design, the first threat is manifest as uncertainty about the duration of the effect-period as discussed above. The second and third threats are manifest as within-individual confounding, because the timing and frequency of crossover are not under the investigator's control. The fourth threat is manifest as selection or information bias. The last threat does not apply to outcomes that are rare, such as myocardial infarction.

Within-individual confounding

Use of subjects as their own controls eliminates confounding by subject characteristics that remain constant, but not by those that change over time. For example, a person who normally drinks coffee after sexual activity exhibits within-individual confounding. This can be controlled by further stratification, as long as the temporal correlation among exposure episodes in the study population is not too high.

The data in table 4 illustrate the problem and solution. If the effect-period is assumed to be 2 hours, then the second subject's infarction could be attributed either to sexual activity or to coffee drinking. Thus, the relative risk of 0.65 for coffee drinking within 2 hours is potentially confounded by sexual activity. Stratification of the data by sexual activity within 2 hours would produce one stratum containing subjects 1 and 2, and another containing subjects 3–10. In each stratum, the observed odds of coffee drinking concurrent with infarction onset would be unchanged for each subject (1:0 or 0:1). However, the expected odds of concurrence would change if episodes of sexual

activity were temporally associated, either directly or inversely, with coffee drinking.

For example, about half of subject 2's episodes of sexual activity were in the morning and half were in the evening; he drank coffee only in the mornings. The expected odds of concurrence of coffee drinking and a random event, within any 2-hour period after sexual activity, were therefore about one-to-one, not 1:23 (i.e., 365:8,401). The relative risk for coffee drinking in the 2 hours after sexual activity, i.e., in the stratum containing only subjects 1 and 2, is estimated to be 5.5 (95 percent CI 0.5–56). When many more than two patients are in this stratum, and a stable relative risk is calculable, it will be possible to quantify modification of the coffee effect measure by sexual activity. This will be done by contrasting the relative risk with that found in the stratum containing subjects who were sexually inactive immediately before their infarctions.

This example shows how it is possible to control within-individual confounding, and to examine modification, provided there exist good data on the temporal association of exposures both immediately before the infarction and in the subject's usual pattern of daily life. In the Onset Study, temporal association can be ascertained precisely during the 26-hour period before the infarction, but only imprecisely during preceding weeks. Therefore, in the analysis, full control of confounding and assessment of effect modification will be restricted by the need to assume relatively short effect-periods (less than 2 hours). To control several confounders simultaneously, conditional logistic regression (14) will be used. For example, the 26-hour period before the infarction will be divided into one 2-hour case-interval and 12 2-hour control-intervals. The analysis will then proceed similarly to the conditional logistic analysis of a case-control study, with one-to-12 matching. The point estimates are expected to be unbiased but the usual confidence intervals would be too narrow because of the temporal correlation of exposures among the intervals (Greenland, personal communication, 1990).

Selection bias

The case-crossover design is immune to one of the main causes of bias in case-control studies—selection of controls that are unrepresentative of the population that produced the cases. Use of cases as their own self-matched controls guarantees representativeness, so long as the matching is preserved in the analysis. (If the matching is ignored, the relative risk will be biased toward the null.) For example, the best representative of the subset of the general population who have exactly the same risk profile as subject 2, is subject 2 himself—as he was a day or two before his infarction (assuming an effect period of less than 12 hours).

Matching in a case-crossover study is a special type of stratified sampling. In stratified sampling, the population is first divided into strata and then the sample is drawn so it is representative within each stratum (but not necessarily representative of the relative frequency of different strata.) In the case-crossover design, the population base is considered to be stratified in the extreme, so there is only one individual per stratum. The sample drawn from each stratum includes "all subjects" (i.e., the one and only subject), except for strata that lack any occurrence of disease. Strata with no cases need not be sampled because they contribute no information to a matched analysis. They each are equivalent to the uninformative, caseless concordant pairs in a matched cohort study.

Although control-selection bias is eliminated, biased case-selection is still possible. Some patients' recent exposures may influence their willingness to participate (e.g., patients with cocaine-induced myocardial infarction may decline to be interviewed). However, transient factors can influence selection in fewer ways than constant factors. For example, regular coffee drinkers may participate less than occasional coffee drinkers, but, among occasional coffee drinkers, it is unlikely that participation will be influenced by whether or not they happened to drink coffee on the day of the infarction. Therefore, self-selection bias by cases should

be less of a problem in a case-crossover study than in a traditional case-control study.

The possibility that interviewers will make more of an effort to interview patients consistent with hypotheses can be reduced by standardized procedures and training. For example, in the Onset Study, information on timing of pain onset (which influences the decision about patient eligibility) is collected before exposure information.

Information bias

In most epidemiologic studies, exposure information is collected as uniformly as possible from cases and noncases to avoid differential information bias. In a case-crossover study, however, questions about exposure during case and control intervals may differ in wording and require different methods of memory recall. For example, a spurious association could occur due to systematic error in reporting usual exposure frequency (the "control" information) compared with exposure during the short interval before the outcome (the "case" information). The direction and magnitude of such a bias would vary among exposures. For example, evidence suggests that usual frequency of sexual activity is slightly overreported (15) and food frequency questionnaires tend to give higher frequencies of specific foods than extrapolation of 24-hour recall (16). By contrast, the usual frequencies of unintentional and nonhabitual actions, such as bursts of anger or occasional heavy lifting, are likely to be underreported due to fading memory. There may also be exaggeration or denial of exposures on the day of the infarction.

The possibility, even the existence, of information bias does not eliminate the utility of the case-crossover design. Some exposures may be quite resistant to biased recall: for example, usual frequency of smoking or coffee drinking during the past month. For other exposures, the direction and magnitude of the bias may be assessed in a validation study, and sometimes in the case-crossover study itself. In the Onset Study, for example, we found the mean of reported "usual" frequency of sexual activity was slightly higher than the mean actual frequency during the 2 weeks before the infarction (excluding the final 24 hours). This is consistent with the cited validation study (15) and suggests that the relative risk of 2.6 is not due to information bias.

Generalizability

The case-crossover design is a scientific way to ask and answer the question clinicians so often ask patients: "Were you doing anything unusual just before the episode?" It is therefore generalizable, in principle, to all acute-onset outcomes hypothesized to be caused by brief exposures with transient effects.

In practice, its utility relative to the traditional case-control design will depend on the relative susceptibilities of each design to selection bias and information bias, which will depend on the particular exposures and diseases of interest. The case-crossover design may prove most useful in specialized clinical settings with unusual referral patterns, limited patient mixes, and no resources for collecting data from nonpatient controls.

REFERENCES

1. Muller JE, Stone PH, Turi ZG, et al. Circadian variation in the frequency of onset of acute myocardial infarction. N Engl J Med 1985;313:1315–22.
2. Muller JE, Toffler GH, Stone PH. Circadian variation and triggers of onset of acute cardiovascular disease. Circulation 1989;79:733–43.
3. Toffler GH, Stone PH, Maclure M, et al. Frequent reporting of possible triggers of acute myocardial infarction: experience in the MILIS Study. Am J Cardiol, in press.
4. Cole P. Coffee-drinking and cancer of the lower urinary tract. Lancet 1971;1:1335–7.
5. Maclure M. Asbestos and renal adenocarcinoma: a case-control study. Environ Res 1987;42:353–61.
6. Miettinen OS. Theoretical epidemiology: principles of occurrence research in medicine. New York: John Wiley and Sons, 1985.
7. Louis TA, Lavori PW, Bailar JC III, et al. Crossover and self-controlled designs in clinical research. N Engl J Med 1984;310:24–31.
8. Guyatt G, Sackett D, Taylor W, et al. Determining

optimal therapy—randomized trials in individual patients. N Engl J Med 1986;314:889–92.
9. Greenland S, Robins JM. Estimation of a common effect parameter from sparse follow-up data. Biometrics 1985;41:55–68.
10. Rothman KJ. Modern epidemiology. Boston: Little, Brown & Co, 1986:196, 213.
11. Toffler GH, Brezinski D, Schafer AI, et al. Concurrent morning increase in platelet aggregability and the risk of myocardial infarction and sudden cardiac death. N Engl J Med 1987;316:1514–18.
12. Kleinbaum DG, Kupper LL, Morgenstern H. Epidemiologic research: principles and quantitative methods. Belmont, CA: Lifetime Learning Publications, 1982.
13. Rothman KJ. Induction and latent periods. Am J Epidemiol 1981;114:253–9.
14. Breslow NE, Day NE. Statistical methods in cancer research. Vol 1. The analysis of case-control studies. IARC publication no. 32. Lyon: International Agency for Research on Cancer, 1980.
15. Hornsby PP, Wilcox AJ. Validity of questionnaire information on frequency of coitus. Am J Epidemiol 1989;130:94–9.
16. Willett WC. Nutritional epidemiology. New York: Oxford University Press, 1989:107.

APPENDIX

The Role of Assumed Durations of the Effect-Period and the Induction Period in Estimation of the Relative Risk

For rare exposures of short duration, with effect-periods of short duration, the proportion of the time spent by each subject (j) in a state of excess or reduced risk due to exposure (x) may be estimated simply by multiplying the number of episodes of exposure (n) by the assumed duration of the effect-period (E_x) and dividing by the duration between the time of the outcome event (T_0) and the start of the time window (T_{xw}) in which exposure was assessed ($nE_x/(T_0 - T_{xw})$). However, if episodes of an exposure are frequent and its effect-period is long, then overlap among effect-periods must be discounted. The following algebra may be used to discount overlap. (If the effect of exposure is hypothesized to compound when effect-periods overlap, then discounting should still be done, and the additional effect due to overlap should be evaluated as if it were due to a separate exposure.)

Consider n episodes of exposure x, each having duration D_{xnj} starting at time T_{xnj} in the time window $(T_{0j} - T_{xwj})$. For point exposures (with $D_{xnj} \ll E_x$), the amount of affected time between two consecutive exposures is either E_x or the lag time between the exposures $(T_{x(n-1)j} - T_{xnj})$, whichever is shorter. For continuous exposures, it is either $(E_x + D_{xnj})$ or $(T_{x(n-1)j} - T_{xnj})$, whichever is shorter. Therefore, the proportion of time (P_{xj}) that the jth subject was in a state of altered risk due to exposure x is

$$P_{xj} = \sum_n \min[E_x + D_{xnj}, (T_{x(n-1)j} - T_{xnj})]/(T_{0j} - T_{xwj}).$$

If the minimum individual induction period in the population is hypothesized to have duration I_x, then I_x would be added to T_{xnj} and $T_{x(n-1)j}$ but would cancel in the subtraction.

Appendix figure 1 illustrates the notation. It is a timeline for subject j (the j subscripts are omitted) who had two episodes of exposure ($n = 2$). The first episode (called x_2 because it is second when counted backwards from the outcome time T_0) is a point exposure with

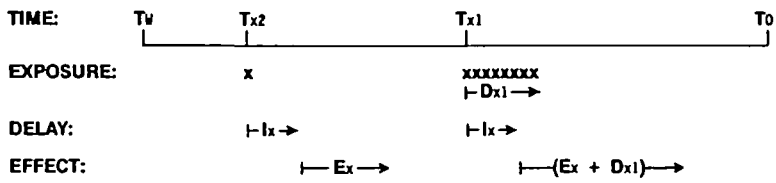

APPENDIX FIGURE 1. A hypothetical timeline of a subject from whom exposure information was available in the window from time T_w to the outcome event at time T_0. The subject was exposed twice: a point-exposure at the time T_{x2} and an interval-exposure beginning at time T_{x1} and ending after the duration D_{x1}. The effect of each exposure episode (increase in risk of outcome) begins after the induction time I_x and lasts for the duration of the effect-period, E_x for the point exposure and $E_x + D_{x1}$ for the interval exposure.

duration $D_{x2} = 0$. The other episode, x_1, has duration D_{x1}. The induction time I_x and the effect-period E_x follow consecutively after x_2. The next effect-period however is longer; it equals $E_x + D_{x1}$. It begins at time $T_{x1} + I_x$. The proportion of time the subject is affected by exposure, P_x, is $(2E_x + D_{x1})/(T_0 - T_w)$.

Cases are classified not as "exposed" or "unexposed," but as "concurrent" or "nonconcurrent" with the assumed effect-period. The numbers of concurrent cases (c_{xj}) and nonconcurrent cases ($1 - c_{xj}$) in the jth stratum (which equal 0 and 1, or 1 and 0) are a function of the assumed durations of the effect-period (E_x) after exposure x, and the lag time ($T_{0j} - T_{x1j} - D_{x1j}$) between the outcome event (T_{0j}) and the end of the last episode of exposure ($T_{x1j} + D_{x1j}$), where T_{x1j} is the start of the last exposure episode and D_{x1j} is its duration. Thus,

$$c_{xj} = \text{integer } \{\{E_x - \min[E_x, (T_{0j} - T_{x1j} - D_{x1j} - 0.5)]\}/(E_x - (T_{0j} - T_{x1j} - D_{x1j}))\}.$$

A more complicated algorithm is required if a minimum induction period I_x is hypothesized, because a case that occurs during the induction period of the last exposure, yet is concurrent with the effect-period of the second-to-last exposure, should be considered a concurrent case.

Each subject contributes $c_x(1 - P_{xj})$ to the numerator and $(1 - c_{xj})P_{xj}$ to the denominator of the Mantel-Haenszel estimator (RR_{MH}). Each subject contributes both these terms to the denominator of the estimator of the variance of the log of RR_{MH}, and $P_{xj}(1 - P_{xj})$ to the numerator. The equation for the variance (9, 10) in this notation is:

$$\text{Var}[\ln(RR_{MH})] = \frac{\sum_j P_{xj}(1 - P_{xj})}{\left[\sum_j c_{xj}(1 - P_{xj})\right]\left[\sum_j (1 - c_{xj})P_{xj}\right]}.$$