

## Control Sampling Strategies for Case-Crossover Studies: An Assessment of Relative Efficiency

Murray A. Mittleman,<sup>1,2</sup> Malcolm Maclure,<sup>1,2</sup> and James M. Robins<sup>1</sup>

The case-crossover study design is a method to assess the effect of transient exposures on the risk of onset of acute events. Control information for each case is based on his/her past exposure experience, and a self-matched analysis is conducted. Empiric evaluation of five approaches to the analysis of case-crossover data from a study of heavy physical exertion and acute myocardial infarction onset is shown. The data presented are from the Onset Study, a case-crossover study of the determinants of myocardial infarction onset conducted in 45 centers from August 1989 to October 1992. In model 1, exactly one control period (matched on clock-time) was sampled per case. In models 2-4, up to 25 control periods were sampled, and the effect of clock-time on the baseline hazard of infarction was modeled. In model 5, a census of the person-time experienced by each subject over the year preceding the infarction was sampled. The 95% confidence interval for model 1 was 2.7 times wider, and the relative efficiency, defined as  $v_1/v_M$ , where  $v_M$  represents the asymptotic variance estimate of the estimated log relative risk with  $M$  control periods sampled per case, was only about 14% of model 5. In models 2-4, the efficiency increased with the number of control periods, regardless of the modeling assumptions. Even with many control periods sampled, models 2-4 achieved only half the efficiency of model 5. The control sampling strategy in any given case-crossover study should be selected with the trade-offs between precision and potential biases of the estimates in mind. *Am J Epidemiol* 1995;142:91-8.

case-control studies; epidemiologic methods; exercise; myocardial infarction; statistics

The case-crossover study design was developed by Maclure (1) as a method to assess the effect of transient exposures on the risk of onset of acute events. In that paper, two methods of sampling control person-time were proposed. It is now apparent that there are large trade-offs between efficiency and potential bias inherent in the choice of control sampling procedure.

In this report, we present an empiric evaluation of five approaches to the analysis of case-crossover data arising in a study of heavy physical exertion as a determinant of acute myocardial infarction onset.

The case-crossover design was developed as a tool to assess the change in risk of a rare acute event during a brief "hazard period" following transient exposure to a determinant of event onset. With this method, control information for each case is based on his/her past

exposure experience. Thus, the case-crossover design is, in many ways, analogous to a crossover experiment, except that 1) subjects, not the investigator, choose when they cross back and forth between periods of transient exposure or nonexposure to the determinant, and 2) these periods are ascertained retrospectively. For example, individuals who exercise intermittently pass into and out of periods of potentially elevated risk of experiencing the onset of an acute myocardial infarction.

The Determinants of Myocardial Infarction Onset Study (Onset Study) is a multicenter, interview-based, case-crossover study designed to quantify the relative risk of acute myocardial infarction onset following heavy physical exertion and other potential determinants.

### MATERIALS AND METHODS

The details of the methods of data collection for the Onset Study have been previously reported (2). In brief, a total of 1,271 patients were interviewed a median of 4 days (range, 0-30 days) following their myocardial infarction. Of these, 43 were unable to complete the interview and were excluded from this analysis. The remaining 1,228 patients (836 men and

Received for publication May 23, 1994, and in final form March 6, 1995.

Abbreviation: METS, metabolic equivalents.

<sup>1</sup> Department of Epidemiology, Harvard School of Public Health, Boston, MA.

<sup>2</sup> Institute for Prevention of Cardiovascular Disease, Cardiovascular Division, Deaconess Hospital, Harvard Medical School, Boston, MA.

Reprint requests to Dr. Murray A. Mittleman, Institute for Prevention of Cardiovascular Disease, Cardiovascular Division, Deaconess Hospital, 1 Autumn St., Fifth Floor, Boston, MA 02215.

392 women, aged 22–92 years) were interviewed between August 1989 and October 1992.

The onset interview was designed to elicit detailed exposure information for each of the 26 hours preceding the onset of infarction symptoms. Physical exertion was estimated on a scale between one and eight metabolic equivalents (METS) using generally accepted MET values (3–6) (table 1). Patients were asked to estimate their usual frequency of exertion at each level during the previous year and to state the timing, type, and level of exertion during each of the 26 hours before the onset of their myocardial infarction. Patients were considered exposed if they reported a peak exertion estimated to be at least six METS during the interval of interest.

The interview treated the entire 26-hour period prior to infarction as one long hazard period and did not draw special attention to the 1 hour immediately prior to myocardial infarction onset. This was intended to decrease differential reporting of exposure information for the hazard and control periods.

Because of the known circadian variation of myocardial infarction, with the incidence of onset approximately 30 percent higher in the period from 6 a.m. to noon compared with the remaining 18 hours of the day (7, 8) and the likely circadian variation of exposure to heavy exertion, clock-time was considered a potential confounder in this study. In the patients studied in this report, a total of 457 episodes of heavy exertion were reported. Of these, 44 percent occurred between 6 a.m. and noon, 43 percent between noon and 6 p.m., 9 percent between 6 p.m. and midnight, and the remaining 4 percent between midnight and 6 a.m.


In this report, we present a set of five alternate approaches to modeling the data arising in the Onset Study (figure 1). In all analyses, the hazard period was restricted to the 1-hour period immediately preceding myocardial infarction onset, since we have previously shown that the induction time for myocardial infarction onset following exposure to heavy exertion is less than 1 hour (2). Table 2 lists the parameters estimated in each of the five models. The models

**TABLE 1. Physical activity rating scale used to evaluate physical exertion in the Onset Study, a case-crossover study of the determinants of myocardial infarction onset conducted in 45 centers from August 1989 to October 1992**

METS*	Description	Physical activity list
1	Sleeping, reclining	Sunbathing, lying on couch watching television
2	Sitting	Eating, reading, desk work, sitting watching television, highway driving
3	Very light exertion	Office work, city driving, personal care, standing in line, strolling in park
4	Light exertion	Mopping, slow walking (e.g., shopping), bowling, sweeping, golfing with cart
5	Moderate exertion (deep breathing)	Normal walking, golfing on foot, slow biking, downhill skiing, calisthenics, raking leaves, cleaning windows, gardening with power tools, hanging wallpaper, interior painting, hunting, fishing, slow dancing, light restaurant work (e.g., waiting tables, serving drinks)
6	Vigorous exertion (panting/overheating)	Slow jogging, speed walking, tennis, swimming, cross-country skiing, shoveling snow, fast biking, mowing with a push mower, pruning trees or shrubs, heavy gardening, factory assembly work, heavy household repairs, climbing up and down a ladder, overhead work, ice hockey, drills, softball, picking up garbage, brick laying, hurried heavy restaurant work
7	Heavy exertion (gasping/much sweating)	Running, fast jogging, nonstop racquetball, pushing car stuck in snow, moving boulders, changing tires, heavy or deep snow shoveling, mixing cement, competitive basketball, touch football, hanging drywall, putting down wall-to-wall carpet, ladder or stair climbing with 50-lb† load, using jackhammer
8	Extreme/peak exertion	Sprinting, fast running, jogging uphill, aggressive sports with frequent sprinting and no rest, pushing or pulling with all your might, unusually extreme work

\* METS, metabolic equivalents.

† Metric equivalent is 22.7 kg.



**FIGURE 1.** Schematic representation of five modeling approaches to data arising in case-crossover studies: A, Pair-matched Interval Approach; B, Multiple Intervals Approach; C, Usual Frequency Approach. MI, myocardial infarction. See the text for a full description of each approach.

are listed in order of increasingly restrictive modeling assumptions.

#### Model 1: Pair-matched Interval Approach

In model 1 (figure 1A), exposure in the hazard period, defined as the 1-hour period immediately preceding myocardial infarction onset, was contrasted with exposure in the comparable 1-hour control period at the same time of day, on the day preceding the infarction. This approach can be thought of as analogous to a one-to-one matched (pair-matched) case-control study and can be analyzed using standard matched pair methods, such as the Mantel-Haenszel estimator or a conditional logistic regression model (1, 9). Instead of case and control subjects, the contrast is between a pair of hazard and control intervals contributed by the same subject. Model 1 allows for the possibility that time of day is a confounder. In modeling the effect of time of day, we assumed nothing about the effect of clock-time on the baseline hazard of infarction onset. This approach allows for the baseline hazard to vary with clock-time in a way that is unique to each individual.

#### Models 2-4: Multiple Intervals Approach

Models 2-4 (figure 1B) contrast exposure in the hazard period with a variable number of control periods sampled from the 25 (1-hour) periods preceding myocardial infarction onset, for which exposure data were obtained. These models are analogous to case-control studies in which a variable number of controls are matched to each case (*M*-to-one matched). These models differ from each other in the assumptions that are made about the effect of time of day on the baseline hazard of infarction onset. All of these models are fit using conditional logistic regression.

In model 2 (Nonparametric Multiple Intervals Approach), time of day is assumed to be a confounder and is statistically modeled in a conditional logistic regression with a series of 23 indicator variables coding for the 24 hours in the day. In this, and subsequent models, time of day is assumed to have a common effect on the baseline hazard in all individuals. This model assumes that the baseline hazard of infarction onset varies with each hour of the day and that the hour-specific odds ratios are common to all patients. However, no assumption is made about the functional

TABLE 2. Parameters estimated in each of the five models evaluated: Onset Study, August 1989 to October 1992

Model	Descriptor	Parameters to be estimated*
1	Pair-matched Interval Approach	$\beta E$
2	Nonparametric Multiple Intervals Approach	$\beta E + (\gamma_1 I_1 + \gamma_2 I_2 + \gamma_3 I_3 \dots + \gamma_{23} I_{23})$
3	Parametric Multiple Intervals Approach	$\beta E + (\gamma_1 \sin(2\pi t/24) + \gamma_2 \cos(2\pi t/24))$
4	Parsimonious Multiple Intervals Approach	$\beta E$
5	Usual Frequency Approach	$\beta E$

\*  $\beta$ , the coefficient for the main effect of exertion;  $\gamma_1$ - $\gamma_n$ , the coefficients for the effects of the confounders;  $E$ , an indicator variable for exposure to heavy exertion;  $I_1$ - $I_{23}$ , a set of 23 indicator variables coding for the 24 hours of the day;  $t$ , a continuous variable, representing time (in hours) on a 24-hour clock.

form of the variation of the baseline hazard over the course of the day.

Model 3 (Parametric Multiple Intervals Approach) is similar to model 2, except for the added constraint that the functional form of time of day is assumed to be estimable by two unknown parameters. As in model 2, time of day is assumed to be a confounder with a common effect on the baseline hazard in all patients and is modeled in a conditional logistic regression by entering terms for the first harmonic of time on a 24-hour clock. This parametric form was selected, since it is consistent with prior reports of the circadian pattern of infarction incidence (7, 8). With this approach, only two terms are needed to model the sine and cosine of the time effect. However, there is an added assumption that the effect of time of day can be adequately modeled using a harmonic regression model.

Model 4 (Parsimonious Multiple Intervals Approach) involves the still more restrictive assumption that time of day is not a confounder. Thus, this model differs from models 2 and 3 only by the added constraint, which is known to be incorrect, that all of the parameters for the time effect are equal to zero. This model implicitly assumes that the occurrence relation between heavy exertion and myocardial infarction onset can be estimated without knowledge of clock time. Thus, no terms are entered into the conditional logistic model to account for the effect of time of day.

#### Model 5: Usual Frequency Approach

Model 5 (figure 1C) is in many ways an extension of the Parsimonious Multiple Intervals Approach to a

much longer, possibly open-ended interval. This model contrasts exposure in the hazard period with the expected exposure, based on each individual's usual frequency of exposure over the year preceding his/her infarction. This approach is analogous to a highly stratified retrospective cohort study. In this paradigm, within each stratum there is exactly one case event, and all of the person-time is contributed by a single individual. The case event may either have been exposed or unexposed in the hour before. Within each stratum, the amount of person-time considered exposed can be estimated by multiplying the reported usual frequency of exposure by its reported usual duration. Unexposed person-time can then be calculated by subtracting the exposed person-time in hours from the number of hours in a year. The data can then be analyzed using standard Mantel-Haenszel methods for follow-up studies with sparse data in each stratum (1, 10) or with maximum likelihood methods (11). This model, like the Parsimonious Multiple Intervals Approach, assumes no confounding by time and that the effect of heavy exertion on myocardial infarction onset can be estimated without knowledge of clock-time.

Model 1 (the Pair-matched Interval Approach) involves sampling exactly one control period per case. In models 2-4 (the Multiple Intervals Approach), we sampled between 2 and 25 control periods for each case. The control periods were selected to be equally spaced over the 25 hours before infarction onset. Model 5 (the Usual Frequency Approach) uses all of the person-time experienced by each patient over the year preceding the infarction to estimate the expected exposure odds. We compared the results of the regression output for models 1-5.

## RESULTS

The relative risk (odds ratio) with 95 percent confidence intervals and the standard errors of the log relative risks are shown in table 3. For models 2-4, results for several matching ratios are shown.

The 95 percent confidence interval for model 1 is 2.7 times as wide as that for model 5, while those for models 2-4 are intermediate and similar to each other. Increasing the number of control periods sampled had little effect on the point estimates. However, the standard error decreased, and the 95 percent confidence intervals narrowed as more control periods were sampled. The effect of increasing the number of control periods sampled on the width of the 95 percent confidence interval was about the same in all of the multiple interval models.

Figure 2 shows the empiric relative efficiency of the relative risk estimators for model 1 and the effect of

TABLE 3. Effect of varying the number of control periods sampled per case on the relative risk, 95% confidence intervals, and the standard error of the log relative risk for heavy exertion as a determinant of myocardial infarction in the Onset Study, a case-crossover study of the determinants of myocardial infarction onset conducted in 45 centers from August 1989 to October 1992

Control periods sampled	Model														
	1		2			3			4			5			
	RR*	95% CI*	SE*	RR	95% CI	SE	RR	95% CI	SE	RR	95% CI	SE	RR	95% CI	SE
1	5.44	2.67-11.09	0.363												
3				6.33	3.87-10.36	0.251	5.90	3.63-9.61	0.248	6.20	3.82-10.08	0.248			
4				5.93	3.79-9.27	0.228	5.52	3.55-8.58	0.225	5.68	3.66-8.82	0.224			
6				5.76	3.81-8.72	0.211	5.51	3.66-8.32	0.209	5.70	3.79-8.59	0.209			
8				5.80	3.68-8.73	0.208	5.73	3.81-8.60	0.208	5.95	3.97-8.94	0.207			
12				5.94	4.00-8.83	0.202	5.84	3.95-8.65	0.200	5.97	4.04-8.84	0.200			
25				6.94	4.72-10.23	0.197	6.83	4.65-10.04	0.197	6.97	4.79-10.23	0.196			
$\infty$										5.85	4.48-7.63	0.136			

\* RR, relative risk; CI, confidence interval; SE, standard error of the log relative risk.

varying the number of control periods sampled on models 2-4 relative to model 5. Empiric relative efficiency was defined as  $\nu_{\infty}/\nu_M$ , where  $\nu_M$  represents the asymptotic variance estimate of the estimated log relative risk with  $M$  control periods sampled per case (12). In model 1, with one control period sampled for each case, the empiric relative efficiency was only 14 percent. In models 2-4, as the number of control periods sampled was increased, up to 25 control periods per case, the empiric relative efficiency increased, but only to a maximum of 48 percent. The modeling assumptions made in models 2-4 had little effect on the relative efficiency, whereas the number of control periods sampled was an important determinant of the relative efficiency.

## DISCUSSION

The theory behind the case-crossover design is grounded in the case-base paradigm (1, 13). As in a conventional case-control study conducted on a dynamic study base, only a fraction of the person-time in the study base is sampled. In the absence of selection and information bias, the estimated exposure odds ratio is an unbiased estimate of the underlying incidence rate ratio that would have been observed if a census of the person-time had been ascertained as in a cohort study. The key difference in case-crossover studies is that the pool of person-time from which the control data are sampled is highly stratified. Each case forms its own stratum; thus, there is only one patient's person-time in each stratum. Strata that do not contain a case event are excluded, since they contribute no information. Within each stratum, all person-time under follow-up experienced by the case before the onset of the outcome is eligible for sampling, but which intervals of person-time are actually sampled depends mainly on the availability and representativeness of exposure data for that period.

In this paper, we have shown empirically that the efficiency of relative risk estimators in case-crossover studies varies greatly depending upon the strategy used in control sampling. In the case of heavy exertion as a determinant of myocardial infarction onset, we further found that the variance was affected more by the number of control periods sampled than by the modeling assumptions made about the effect of confounding by clock-time.

In the Pair-matched Interval Approach, with one control period sampled for each case period, the 95 percent confidence interval was 2.7 times wider, and the relative efficiency was only about 14 percent of that observed with the Usual Frequency Approach. The efficiency greatly increased as the number of control periods sampled in each matched set increased.

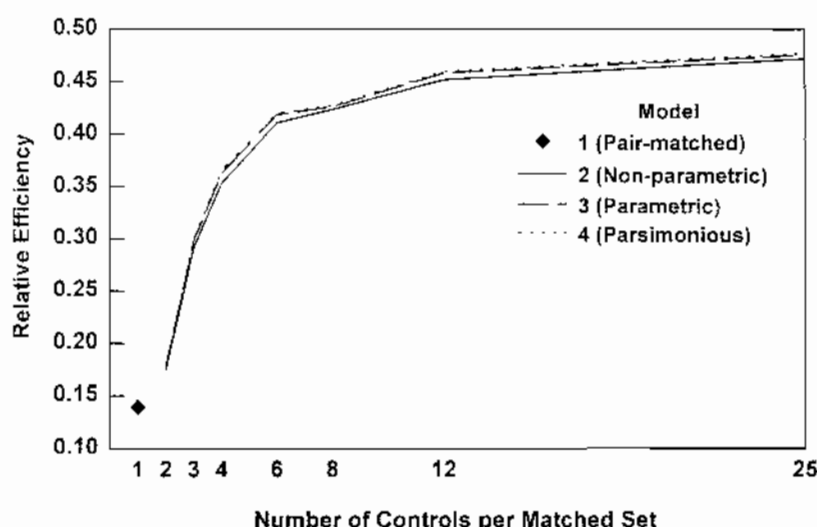


FIGURE 2. Empiric relative efficiency of relative risk estimators from a case-crossover study of heavy exertion as a determinant of myocardial infarction onset.

regardless of the modeling assumptions made about the effect of time of day. However, even with multiple control periods sampled for each case, the Multiple Intervals Approach achieved only half the efficiency of the Usual Frequency Approach.

#### Expected relative efficiency

The low relative efficiency that we observed even with multiple control periods sampled for each case with the Multiple Intervals Approach seems to conflict with the rule of thumb for case-control studies, which states that, upon sampling four or five controls for each case, there is little further gain in efficiency (9, 12). When the relative risk is null (equal to one), the relative efficiency is  $M/(M + 1)$  percent of that which would have been observed if a cohort study had been conducted, with a census of the person-time in the study base accounted for. Under the null, the relative efficiency is not influenced by the prevalence of exposure in the study base. Thus, a one-to-one matched sampling plan (e.g., the Pair-matched Interval Approach) would have a relative efficiency of 50 percent, and a matching ratio of four to one would have 80 percent efficiency.

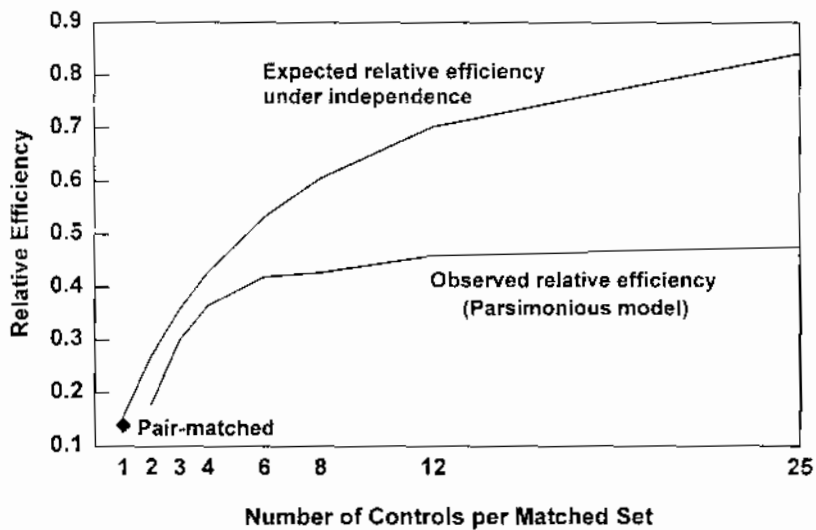
Generalizing to situations in which the relative risk is not null, we find that the expected relative efficiency depends on the magnitude of the relative risk and the prevalence of exposure in both the cases and the study base (12). The further the relative risk is from the null, the lower is the expected efficiency. Furthermore, when the relative risk is far from the null, estimators of the relative risk are least efficient when the prevalence of exposure in the study base is low (12). The details

required to calculate the expected relative efficiency are given in the Appendix.

In the Onset Study, the relative risk for myocardial infarction onset following heavy exertion was close to six and thus was far from the null. In addition, the exposure probability in any given hour was low (about 0.8 percent). In fact, even in the hazard period 1 hour before onset, the prevalence of exposure was only 4.4 percent (2). Under the conditions observed in this study, we would expect the one-to-one matched analysis to have a relative efficiency of only 15 percent, very close to that observed empirically.

Even taking into account the extreme relative risk and the low prevalence of exposure, we would expect a larger increase in the relative efficiency as the number of control periods sampled per case was increased than was observed empirically in this study (figure 3). An explanation for this discrepancy is that the multiple control periods sampled were not independent (9, 12). In the Onset Study, the control periods for which information was available were the 25 consecutive 1-hour periods preceding infarction onset. Thus, the probabilities of exposure in the control periods sampled were not independent of each other. The hour following heavy exertion is more likely to include additional heavy exertion, but soon thereafter the individual is likely to be resting. It is therefore not surprising that the gain in efficiency was not as dramatic as predicted, since there was less additional information with each additional contiguous control period sampled.

The empiric relative efficiencies reported here are specific to this particular data set. However, the for-



**FIGURE 3.** Observed relative efficiency of relative risk estimators for a case-crossover study with up to 25 control periods sampled for each case compared with expected under the assumption of independence of exposure among control intervals. This example assumes the following:  $\psi$  (relative risk) = 5.8;  $p_1$  (average probability of case exposure) = 4.4%;  $p_2$  (average probability of control exposure) = 0.8%;  $N$  (total number of matched sets) = 1,228.

mulae in the Appendix may be used to obtain the expected relative efficiency for any case-crossover study given assumptions about the magnitude of the relative risk, the prevalence of exposure in the study base, and the matching ratio.

#### Within-person confounding

An important consequence of matching on the individual is that all characteristics of an individual that remain constant over time do not vary within strata. Thus, there can be no confounding by these characteristics. This consequence can be stated as freedom from between-person confounding. However, there can be within-person confounding. This problem arises when multiple transient exposures are correlated in time within an individual. Within-person confounding can be modeled in case-crossover studies as long as data regarding the temporal correlation between multiple exposures are collected. For example, in the analysis of heavy physical exertion, it is possible that the association is confounded by the concurrence of episodes of anger with exertion (14). If this were the case, models 1–4 could adjust for the confounding using conditional logistic regression with a term entered for outbursts of anger. This can be easily accomplished as long as exposure information is collected for each of the potential within-person confounders for each hour sampled as a control period.

In order to control for within-person confounding with the Usual Frequency Approach, it is necessary to collect information on the usual frequency of exposure

to each factor conditional on exposure to all other potential confounders. For example, the following question may be asked: "How often during the past year have you been very angry during a period of heavy exertion?" The complexity of such a question will usually make it impractical to ask. Therefore, the inability to control within-person confounding is a limitation of the Usual Frequency Approach. There will, however, be exceptions. For example, it is feasible to ask for the usual timing of heavy exertion during the day.

#### Bias versus precision

None of the models presented is clearly superior in all cases. There will often be a trade-off between precision of the estimates and bias that may be introduced by making modeling assumptions that may not represent the true state of nature. In this example, the Usual Frequency Approach was preferable as a final model, because it had the most efficient estimators and because within-person confounding by clock-time was negligible. In other studies, controlling within-person confounding may be paramount and the Usual Frequency Approach inadequate.

The preferred approach to control period sampling in any given case-crossover study depends upon the length of the hazard period, the induction time from exposure to outcome onset, the degree of within-person confounding, and the quality of data available. Therefore, whenever possible, it is desirable to collect exposure data that allow all modeling approaches to be

used. Moreover, data from the different sampling periods can be used for cross-validation. For example, the possibility of recall bias in reporting usual frequency can be assessed by comparing the observed occurrence in the discrete control period with the expected occurrence based on usual frequency. In this example, comparing exposure in the 1-hour period 24 hours preceding infarction onset with the expected exposure based on the Usual Frequency Approach (model 5) gave a relative risk of 1.20 (95 percent confidence interval 0.70–2.06).

In conclusion, the case-crossover design can be considered a family of related designs with different sampling and modeling strategies. This paper provides the beginnings of an understanding of their relative strengths and limitations.

#### ACKNOWLEDGMENTS

Supported by a grant (HL41016) from the National Heart, Lung, and Blood Institute. Dr. Mittleman was supported by a Bourse de formation en Recherche from the Fonds de la Recherche en Santé du Québec.

#### REFERENCES

1. Maclure M. The case-crossover design: a method for studying transient effects on the risk of acute events. *Am J Epidemiol* 1991;133:144–53.
2. Mittleman MA, Maclure M, Toffler GH, et al. Triggering of

- acute myocardial infarction by heavy exertion: protection against triggering by regular exertion. *N Engl J Med* 1993; 329:1677–83.
3. Sallis JF, Haskell WL, Wood PD, et al. Physical activity assessment methodology in the Five-City Project. *Am J Epidemiol* 1985;121:94–106.
4. Astrand PO, Rodahl K. *Textbook of work physiology*. 3rd ed. New York: McGraw-Hill Book Company, 1986.
5. Ainsworth BE, Haskell WL, Leon AS, et al. Compendium of physical activities: classification of energy costs of human physical activities. *Med Sci Sports Exerc* 1993;25:71–80.
6. Taylor HL, Jacobs DR, Schucker B, et al. A questionnaire for the assessment of leisure time physical activities. *J Chronic Dis* 1978;31:741–55.
7. 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.
8. Muller JE, Toffler GH, Stone PH. Circadian variation and triggers of onset of acute cardiovascular disease. *Circulation* 1989;79:733–43.
9. Breslow NE, Day NE, eds. *Statistical methods in cancer research. Vol 1. The analysis of case-control studies*. Lyon: International Agency for Research on Cancer, 1980. (IARC scientific publication no. 32).
10. Greenland S, Robins JM. Estimation of a common effect parameter from sparse follow-up data. *Biometrics* 1985;41: 55–68.
11. Marshall RJ, Jackson RT. Analysis of case-crossover designs. *Stat Med* 1993;12:2333–41.
12. Breslow NE, Day NE, eds. *Statistical methods in cancer research. Vol 2. The design and analysis of cohort studies*. Lyon: International Agency for Research on Cancer, 1987. (IARC scientific publication no. 82).
13. Miettinen OS. *Theoretical epidemiology: principles of occurrence research in medicine*. New York: John Wiley & Sons, Inc, 1985.
14. Mittleman MA, Maclure M, Sherwood JB, et al. Triggering of acute myocardial infarction by episodes of anger. *Circulation* (in press).

#### APPENDIX

This Appendix outlines the calculation of the expected relative efficiency in case-crossover studies. The derivation of these formulae has been presented elsewhere (12). With  $M$  control periods sampled in each matched set, the variance of the log relative risk,  $v$ , is given by

$$\text{Var}(\log(\psi)) = v_M = \left[ \sum_{m=1}^M \frac{T_m m \psi (M - m + 1)}{(m\psi + M - m + 1)^2} \right]^{-1} \quad (1)$$

where  $\psi$  is the odds ratio,  $M$  is the number of control periods sampled per case, and  $T_m$  is the number of matched sets with exactly  $m$  subjects exposed.  $T_m$  is approximated by  $NP_m$ , where  $N$  is the number of matched sets in the study.  $P_m$  is the probability that a matched set contributes to  $T_m$  and is given by

$$P_m = \binom{M}{m} (1 - p_1) p_2^m (1 - p_2)^{M-m} + \binom{M}{m-1} p_1 p_2^{m-1} (1 - p_2)^{M-m+1} \quad (2)$$

and is evaluated for the mean values of  $p_1$  and  $p_2$ , which are the probabilities that a case and control period are exposed, respectively. The expected relative efficiency is taken as  $v_\infty / v_M$ , where  $v_M$  is the expected variance of the log odds ratio with  $M$  controls per case. This equation simplifies to the  $M/(M+1)$  rule under the condition that the relative risk is null ( $\psi = 1$ ).