J Chron Dis Vol. 40, Suppl. 2, pp. 139S–161S, 1987
Printed in Great Britain. All rights reserved

0021-9681/87 $3.00 + 0.00
Copyright © 1987 Pergamon Journals Ltd

A GRAPHICAL APPROACH TO THE IDENTIFICATION AND ESTIMATION OF CAUSAL PARAMETERS IN MORTALITY STUDIES WITH SUSTAINED EXPOSURE PERIODS

JAMES ROBINS

Harvard School of Public Health, Occupational Health Program, 665 Huntington Avenue, Boston, MA 02115, U.S.A.

Abstract—In observational cohort mortality studies with prolonged periods of exposure to the agent under study, independent risk factors for death commonly determine subsequent exposure to the study agent. For example, in occupational mortality studies, date of termination of employment is both a determinant of subsequent exposure to the chemical agent under study (since terminated individuals receive no further exposure) and an independent risk factor for death (since disabled individuals tend to leave employment). When a risk factor determines subsequent exposure and is determined by previous exposure, standard analyses that estimate age-specific mortality rates as a function of cumulative exposure can underestimate the *true effect of exposure* on mortality, whether or not one adjusts for the risk factor in the analysis. This observation raises the question, "Which, if any, empirical population parameter can be causally interpreted as the *true effect of exposure* in observational mortality studies?" In answer, we offer a graphical approach to the identification and estimation of causal parameters in mortality studies with sustained exposure periods. We reanalyze the mortality experience of a cohort of arsenic-exposed copper smelter workers using our approach and compare our results with those obtained using standard methods. We find an adverse effect of arsenic exposure on all cause and lung cancer mortality, which standard methods failed to detect. The analytic approach introduced in this paper may be necessary to control bias in any epidemiologic study in which there exists a risk factor which both determines subsequent exposure and is determined by previous exposure to the agent under study.

1. INTRODUCTION

In cohort mortality studies with prolonged periods of exposure to the agent under study, independent risk factors for death commonly determine later exposure to the study agent. For example, in occupational cohorts, one observes that unexposed individuals who terminate employment at any age (say, 40) prior to age 65 have higher subsequent age-specific mortality rates than unexposed individuals who continue at work past that age—at least in part because workers at increased risk (e.g. workers with chronic disabling illnesses) tend to take early retirement. This tendency for workers at increased risk to preferentially leave employment we shall refer to as *the healthy worker survivor effect*. It follows that termination status is both a determinant of future exposure to the chemical agent under study (since terminated individuals receive no more exposure) and an independent risk factor for death. As briefly noted by Gilbert [1] and Robins [2, 3], when risk factors for death preclude or diminish subsequent exposure, standard intracohort analyses that estimate age-specific mortality rates as a function of cumulative exposure tend to underestimate the *true effect of exposure* on mortality. (For example, even under the null hypothesis of no exposure effect, subjects at risk at age 60 who left employment at age 40 will tend to have lower cumulative exposures and, because of the healthy worker survivor effect, higher mortality rates than subjects remaining at work at age 60.) If, in addition, previous exposure history is a

determinant of subsequent risk factor status, the association of mortality with cumulative exposure can underestimate the *true exposure effect*, even when one adjusts for the risk factor in the analysis. [For example, if leaving employment is a proxy for the (unrecorded) onset of disabling cardiac disease, then, controlling for the risk factor "age at termination of employment" in the analysis may be tantamount to controlling for an intermediate variable (i.e. the onset of disabling cardiac disease) on the causal pathway from exposure to death, which produces a bias towards the null.]

These observations raise the question, "What do we mean by the true effect of exposure?" Putting it more formally, "Which, if any, population parameter can be causally interpreted as the *true effect of exposure* when the association of observed exposure history (however summarized) with mortality is not causal?" In answer, in Section 2, we develop a graphical approach to the identification and estimation of causal parameters in mortality studies with sustained exposure periods.

In Section 3, we present several paradigmatic examples to demonstrate that when there exists a risk factor whose current level both determines subsequent exposure to the study agent and is determined by previous exposure, the observed association of mortality with cumulative exposure can underestimate the true exposure effect, whether or not one adjusts for the risk factor in the analysis.

In Section 4 we apply the approach developed in Section 2 to the estimation of the causal effect of arsenic exposure on total mortality in a cohort of copper smelter workers. We find a small adverse effect of arsenic exposure on survival. In contrast, a standard analysis, which does not control for the healthy worker survivor effect, finds no relationship between mortality and cumulative exposure. Our estimate of the causal effect of arsenic on mortality is obtained only by fitting statistical models and thus may be significantly biased if the models are incorrect (i.e. misspecified). Fortunately, for many occupational exposures, a central first question is whether the exposure has any effect on mortality whatsoever. Conveniently, a nonparametric (i.e. distribution-free) test of the null hypothesis of no exposure effect can be constructed. This nonparametric test also suggests a statistically significant adverse effect of arsenic on survival.

Gilbert recognized that for chronic disabling illnesses such as arteriosclerotic cardiovascular disease, bias due to the healthy worker survivor effect could not be controlled by standard analytic methods. But she conjectured that, for diseases for which the interval between clinical manifestation and death is less than 10 years, such as lung cancer, any valid test for an association of age-specific mortality with cumulative exposure lagged 10 years (that is, for an individual at risk at age t, any exposure received after age t-10 is ignored for the purposes of analysis) would be a valid test of the null hypothesis of no exposure effect. In Section 4, we show that Gilbert's conjecture appears to be incorrect when applied to a subcohort of copper smelter workers hired after exposures to arsenic were reduced in the 1920s by the introduction of industrial hygiene controls and process changes. In particular, an intracohort analysis restricted to cohort members hired after 1935 shows no association between lung cancer mortality and cumulative exposure to arsenic, even when cumulative exposure was lagged some 15 years. In contrast, we find a marked association between lung cancer mortality and arsenic exposure ($\chi_1^2 = 18$) when we apply our new distribution-free test.

We conjecture that these results are at least partly the consequence of the fact that smokers leave employment at greater rates than nonsmokers. If so, individuals who terminate employment early will be increased risk for lung cancer (since smokers will be overrepresented among early terminees) and will tend to have low cumulative exposures, even when lagged 15 years. Since information on cigarette smoking was not obtained, the test proposed by Gilbert will be biased because any difference in cigarette smoking rates between individuals with high and low lagged cumulative exposures cannot be adjusted for (nor can our conjecture be directly tested in these data, although we present circumstantial evidence for it in Section 4). In contrast, the analytic method introduced in this paper is unbiased, even when cigarette smoking is a determinant of employment status. Since the post-1935 exposure levels found in the copper smelter are typical of the lower levels of exposures to adverse agents that are found in American industry today, it may prove to be generally important to analyze occupational cohort mortality data, even for lung cancer, using the proposed methods.

The exposition given in this paper is informal, relying heavily on graphical examples. The

reader may refer to reference [4] for a formal mathematical development.

2. A GRAPHICAL APPROACH TO CAUSAL INFERENCE IN AN OBSERVATIONAL SURVIVAL STUDY

In this section we set ourselves three tasks.

(1) To represent the observed study data by graphs that we shall call "measured partially interpreted structured tree graphs" or "measured graphs" for short.

(2) To represent the causal parameters of the study by graphs that we shall call "measured causally interpreted structured tree graphs" or "causal graphs."

(3) To represent the subset of identifiable (i.e. empirically estimable) causal parameters by graphs that we shall call "fully randomized measured causally interpreted structured tree graphs" or "randomized graphs."

We shall see that, in observational studies, investigators may disagree as to the causal parameters of the study. Even when they agree, they may still disagree as to which causal parameters are estimable.

Task 1. A Graphical Representation of the Observed Study Data

A measured graph represents the observed study data (i.e. the data recorded for data analysis). Measured graph 1 (Fig. 1) (where we label graphs by their figure number) represents the first 18 months of follow-up data from an occupational cohort mortality study on a subcohort of workers who were all hired at (calendar) time t_1, matched on age at hire, and followed for 40 years. Current covariate status (i.e. exposure concentration and employment status) and vital status are recorded only every 6 months at times $t_1, t_2, \ldots, t_{81}$ (where $\Delta t = 6$ months is chosen short enough that recorded covariate history adequately approximates true covariate history).

Measured graph 1 is interpreted as follows (where, for the moment, the reader should ignore both the highlighted lines and the fractions in parentheses): At t_1, 300 individuals were at work of whom 100 received high exposure and 200 zero exposure. Seventy five of the 100 receiving high exposure at t_1 survived until t_2, at which time 15 of the 75 remained at work at high exposure jobs, 30 remained at work at unexposed jobs, and 30 left work and thus were unexposed. Of the 30 who left work at t_2, 15 survived to t_3, at which point six returned to a high exposure job at work, six returned to an unexposed job and three remained off work and thus unexposed.

We shall refer to a circle in a measured graph as a node, to line segments within a circle as intranodal lines, and to line segments between circles as internodal lines. Then, by definition, a measured graph has the following structure. One or more nodes represent each time of observation t_l, $t_l \in (1, \ldots, t_S)$. Successive observation times are referred to as generations. Exactly one node is present at t_1. Each node (except that at t_1) receives on its left circumference exactly one internodal line from a node in the previous generation. This line splits into one or more intranodal lines that traverse the interior of the circle and terminate at distinct points on the right circumference. From each of the above points on the right circumference, one or more internodal lines originate and extend to a node in the next generation. In Fig. 1, a particular intranodal line at t_s (and the corresponding point on the right circumference on which it terminates) represents a unique covariate (i.e. employment and exposure) history through t_{s-1} plus survival to t_s. An internodal line connecting a node at t_s to one at t_{s+1} represents a unique exposure and employment history through t_s.

A study can be represented by several different measured graphs. For example, measured graph 2 (Fig. 2) represents the same occupational mortality study as measured graph 1. Measured graph 2 differs from measured graph 1 only in that, among individuals with identical work and exposure histories through t_{s-1}, the subset of individuals off work at t_s and the subset of individuals at work at t_s are on distinct intranodal lines (and right circumference points). Thus in Fig. 2 an intranodal line at t_s represents a unique history of exposure through t_{s-1} and employment through t_s. Measured graphs 1 and 2 have (1) the same number of nodes in each generation, (2) the same number of internodal lines arising from corresponding nodes, and (3) corresponding internodal lines representing the same covariate histories.

If two non-identical measured graphs have properties (1–3) and if one of the measured graphs (e.g. Fig. 2) always has at least as many intranodal lines per node as the other measured graph (e.g. Fig. 1) we say that the former measured graph (e.g. Fig. 2) is *coarser* than the

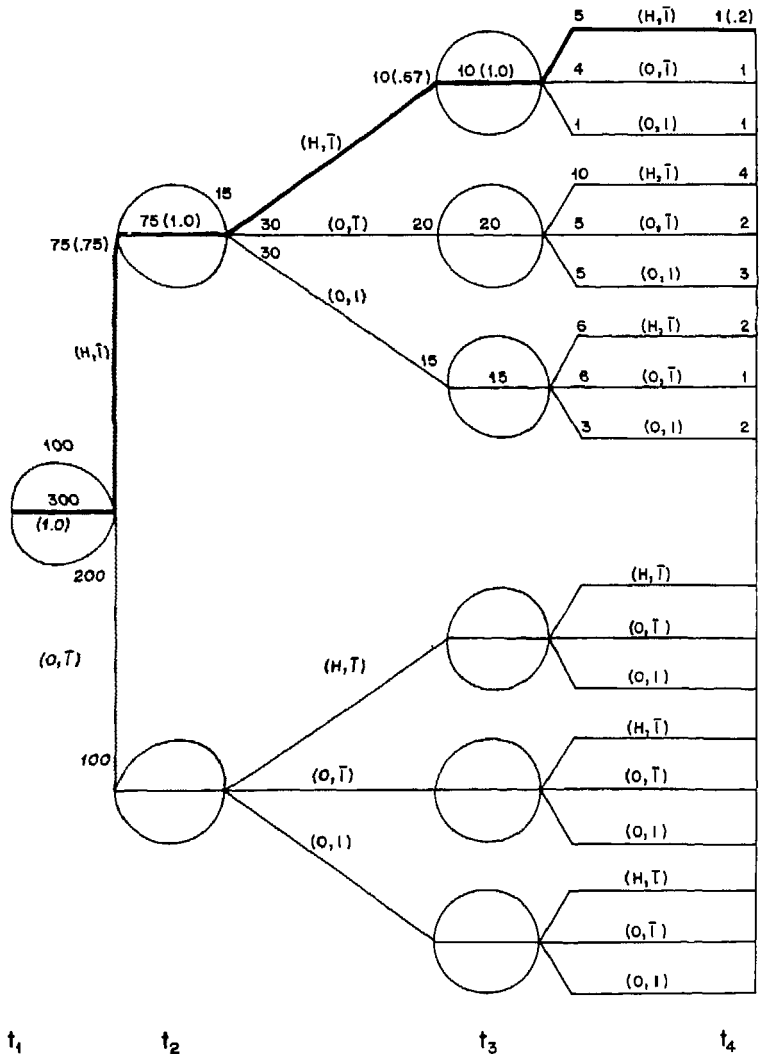


Fig. 1

Legend for Figs 1–9. A Measured Graph. H = high exposure concentration; M = moderate exposure concentration; O = unexposed; Ī = at work; I = off work; whole numbers = numbers of subjects surviving at the given time with a given covariate history. Fractions in parentheses are conditional probabilities defined in Section 2 of text. Covariates on a given internodal line linking nodes at times t_s and t_{s+1} refer to measurements made at t_s. Highlighted subgraphs represent generalized treatments as defined in the text.

latter measured graph (e.g. Fig. 1) and the *latter* is *finer* than the former.

Task 2. Identifying the Causal Parameters of the Study

An observational study differs from a controlled trial in that no investigator influenced the exposures of the study subjects. Rubin [5] has developed a theory of causal inference for point exposure observational studies. We shall extend his theory to include studies with sustained exposure periods. We shall need to define treatments, causal graphs, and the causal parameters of the study. In our causal theory we shall suppose that nature (our name for the "actor" in an observational study) deterministically decides each individual's covariate and survival status at every time.

Definition of a treatment

A particular (possibly joint) covariate level at time t_s [e.g. at work and receiving high exposure, i.e. (H, Ī) in Fig. 1] is a *treatment* for a particular individual alive at t_s if that individual could, at least conceptually, have either received or not

Figure 2

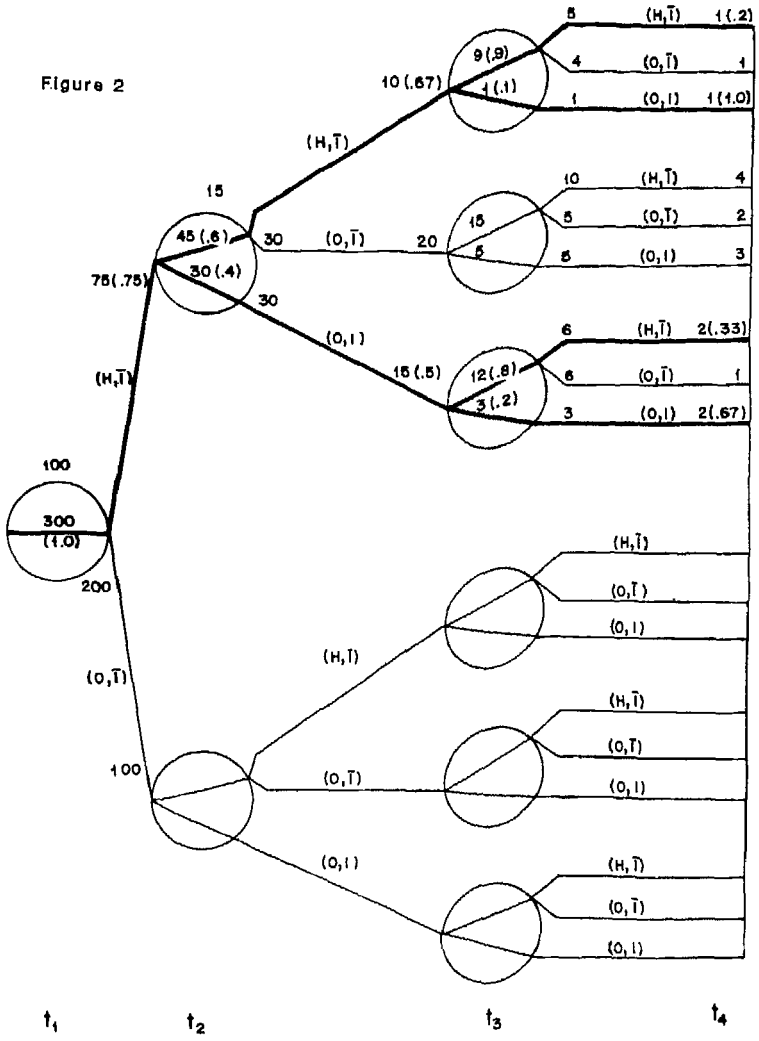


Fig. 2

received that covariate level at t_s. A necessary condition for the existence of a treatment at t_s is that, at least conceptually, a controlled trial could have been performed in which at t_s an investigator superseded nature's deterministic choice of covariate level and gave the individual the covariate level under consideration. [We shall assume that under any treatment the subsequent covariate and survival history for each individual is uninfluenced by the treatments received by any other individual at any time [5].]

Definition

A *measured graph* is a *causal graph* if, for any right circumference point, (a) the covariate levels determining membership on the internodal lines arising from that point are each treatments for any individual with the covariate and survival history represented by the right circumference point, and (b) the subsequent covariate history for any such individual given any one of these "treatments" can be represented by some path of intra- and internodal lines that lie on the measured graph. If, as we shall henceforth assume, exposure refers to exposure to an industrial chemical, it would be reasonable to assume that measured graph 2 is a causal graph. [Measured graph 2 would not be a causal graph if (1) our study was an investigation of the effect of routine physical exertion on mortality in which high exposure had been defined as "walks two or more miles daily at work" and (2) at some time t_s, there were individuals at work whose legs were paralyzed.] Furthermore, for convenience, we shall also interpret measured graph

1 to be a causal graph by making the assumption that for any individual off work at t_s, an investigator could have intervened and placed them at work at either exposure level. (See Section 3E of [4] for a further discussion of this assumption.)

The causal parameters of a causal graph

Figure 3 is a causal graph which represents the data from a "point exposure study" in which at t_1 (start of follow-up) each member of the population receives either high (H) or zero (0) exposure and, at t_2, vital status data are recorded. Following Rubin [5] we define the population causal effect of exposure on mortality to be the difference between the proportion surviving to t_2 in a controlled study in which the entire population received high exposure at t_1 and the proportion surviving in a study in which the entire population had received zero exposure at t_1. This definition applies to a population comprising a single individual, in which case these proportions must be either 1 or 0 because we have assumed deterministic outcomes. Since in the study that actually occurred each subject is either exposed or unexposed (but not both), at least one (and usually both) of these two controlled studies will be hypothetical (in the sense that the study did not actually occur).

If survival status had been ascertained at multiple times (e.g. at $t_2, t_3, \ldots, t_{S+1}$) we would then have compared the proportion surviving at each such time (i.e. we would have compared survival curves). We have defined causal parameters in terms of comparisons of survival curves rather than of incidence rates since, even in a large randomized trial, it is possible that each individual's life may be shortened by exposure and yet, at certain times, the incidence rate (i.e. the hazard rate) in the unexposed population may exceed that in the exposed. This reflects the fact that exposed and unexposed survivors at some times t may be noncomparable (on unmeasured risk factors) due to selective survival.

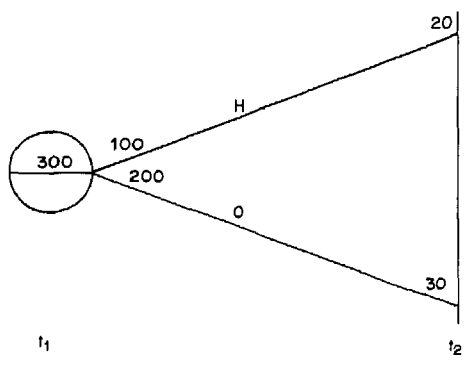


Fig. 3

A natural generalization of the above definition of a causal parameter that applies equally to studies with sustained exposure periods is the following. *A population (individual) causal parameter associated with a particular study population (subject) is the difference in the population (individual) survival curves in two different, usually hypothetical, controlled studies, each with well-defined outcomes.* To determine the causal parameters of a study, we must decide how, on the basis of the observed study data, to characterize the set of controlled studies whose outcomes are believed to be well defined. To do so, we first show that every causal graph has an associated *set of generalized treatments* each of which determines a unique controlled study. A particular generalized treatment, G, of a causal graph is represented by a highlighted subgraph of the causal graph constructed by use of the following.

Generalized treatment algorithm

Beginning at the left circumference of the t_1 node, highlight *all* intranodal lines in that node. At each point on the right circumference of the t_1 node at which a highlighted intranodal line terminates, highlight one internodal line. At t_2 highlight all intranodal lines in the nodes on which the highlighted internodal lines (originating from the t_1 node) terminated. From each point on the right circumference of a t_2 node at which a highlighted intranodal line terminates, highlight one internodal line leading to a node in generation t_3. Continue in this manner through generation t_{S+1}.

The hypothetical controlled study associated with the highlighted subgraph of causal graph 2 is as follows. At t_1, an investigator intervenes and gives all individuals high exposure. Nature determines survival and employment status through t_2. For those individuals still at work at t_2 in this hypothetical study, an investigator again intervenes and gives each high exposure; nature then determines their survival and employment status through t_3. For individuals off work at t_2, nature gives each zero exposure at t_2, and then determines their survival and employment status through t_3. For those at work at t_3, an investigator gives each high exposure, etc.

That is, at each highlighted right circumference point from which two or more internodal lines originate, the investigator intervenes and gives to all subjects (with the history represented by that right circumference point) the highlighted treatment originating from it.

A generalized treatment, therefore, is the treatment "regime" or "protocol" of a controlled study. To verbally characterize this treatment regime, one must characterize the internodal line that is highlighted at each right circumference point at which two or more internodal lines were eligible for highlighting. For example, the generalized treatment highlighted in Fig. 2 can be characterized by, "if at work at any t_s, receive high exposure" and that of causal graph 1 can be characterized by "if alive at t_s, remain at work and receive high exposure." As causal graphs, Figs 1 and 2 differ in that the treatment regimes represented by the generalized treatments of Fig. 1 require that the investigator supersede nature and determine each subject's employment status at each time.

Definition

The *G-causal parameter* (written for a particular individual i as $S(t, G_1, G_2, i)$ while for the population parameter the i is dropped) comparing two generalized treatments G_1 and G_2 of a causal graph is the difference between the population (individual) survival curve of the controlled study defined by G_1 (written as $S(t \mid G_1)$ and $S(t \mid G_1, i)$ respectively) and the survival curve of the study defined by G_2.

Example. The G-causal parameter comparing the generalized treatment of causal graph 1 "if alive, stay at work and receive high exposure" with the generalized treatment "if alive, stay at work and receive zero exposure" represents the direct effect of exposure on mortality controlling for (the possible intermediate variable) employment history. In contrast, the G-causal parameter comparing the generalized treatment of causal graph 2 "if at work, receive high exposure" with the generalized treatment "if at work, receive zero exposure" is a measure of the overall effect of exposure on mortality when not controlling for employment history (since in these two controlled studies the investigator does not control (i.e. determine) employment history.) [Note that the controlled trial represented by the generalized treatment of causal graph 2 "if at work, receive zero exposure" could also be expressed as "if alive, receive zero exposure" since all subjects are unexposed when off work.] The best measure of the overall effect of exposure might be to compare the survival curve of the controlled trial "if alive, receive high exposure" (whether on or off work) with the survival curve of the trial "if alive, receive zero exposure". Unfortunately, as discussed below, we will be unable to estimate the survival curve of the trial "if alive, receive high exposure" without making strong, untestable modelling assumptions. The need for modelling assumptions arises because, in the observed study, no subject received high exposure while off work and, thus, the trial "if alive, receive high exposure" is not represented by a generalized treatment of causal graph 2.

We will assume that in our study, there exists a finest causal graph (e.g., causal graph 1 in our occupational mortality study) [4].

Definition

The *essential causal parameters* of a study are the set of G-causal parameters of the finest causal graph and of all coarser causal graphs. (Note it follows from the definition of a causal graph that any measured graph coarser than a causal graph is itself a causal graph.) Henceforth, we shall restrict attention to the essential causal parameters.

Task 3. Determining the Estimable Causal Parameters

Definition 1

A causal graph is a *randomized* graph if the *subsets of the population represented, in the observed study, by the internodal lines arising from any given right circumference point* have exactly the same distribution of subsequent covariate and survival histories as one another in any study defined by a generalized treatment of the causal graph (or of a coarser causal graph) whose highlighted subgraph passes through that right circumference point. Informally, the causal graph is randomized if the above subsets are "comparable" in the sense that just before receiving their observed treatments at t_s, the subsets had identical distributions on all "risk factors" predicting subsequent covariate and survival history. It follows from definition 1 that any measured graph B coarser than a randomized graph A is itself a randomized graph.

Example. Suppose among any group of subjects with identical employment and work his-

tories through t_{s-1}, (1) the subgroup who received high exposure at work at t_s is comparable to the subgroup who received zero exposure at work, but (2) because of the healthy worker survivor effect, the subgroup that was "off work at t_s" differs on unmeasured risk factors for death from those subjects "at work at t_s." Then causal graph 2 is randomized but causal graph 1 is not randomized.

If a causal graph is randomized, we can compute the probability of survival to any time t_s in the hypothetical controlled study defined by any particular generalized treatment G_1 by the following algorithm.

G-Computation algorithm

1. On each intranodal line on the highlighted subgraph G_1 write (as we have done in parentheses on the highlighted graphs in Figs 1 and 2), the conditional probability in the observed study of being in the subset defined by that intranodal line, given that one is a member of the subset defined by the node in which the line lies. On the right end of each internodal line on the highlighted subgraph G_1, write, in parentheses, the conditional probability of surviving to the next node (time) given that (in the observed study) one was in the subset represented by that internodal line.

2. For each highlighted path of intra- and internodal lines that connects a node at t_s to the left circumference of the t_1 node form the product of all conditional probabilities entered in Step 1. (This product is the probability of surviving to t_s in the hypothetical study defined by G_1 with the covariate history represented by that path of intra- and internodal lines, since, for a randomized graph, the conditional probabilities of the observed study defined in Step 1 are also the conditional probabilities of the hypothetical study.) The sum of these products is the desired survival probability.

Example. For the generalized treatment G_1 highlighted on randomized graph 2

$$\begin{aligned} S(t_4|G_1) = &(0.75)(0.6)(0.67)(0.9)(0.2) \\ &+ (0.75)(0.6)(0.67)(0.1)(1.0) \\ &+ (0.75)(0.4)(0.5)(0.8)(0.33) \\ &+ (0.75)(0.4)(0.5)(0.2)(0.67). \end{aligned}$$

When causal graph 2 is extended to include the entire 40 year follow-up period, there would be 2^{80} terms in the final sum in step 2 of the G-computation algorithm for t_{81}. Thus, the sum could not be evaluated even with the aid of a high speed computer. Nonetheless, we can accurately evaluate the sum by using the Monte Carlo G-computation algorithm given in the Appendix.

In order to allow for sampling variability we shall henceforth suppose that the observed study population has been randomly sampled from a near-infinite hypothetical superpopulation. The causal parameters of interest will be those of the superpopulation. We redefine a causal graph to be a randomized graph if and only if Definition 1 holds for the superpopulation. Thus, even for a randomized graph, chance association of treatments with unmeasured risk factors may exist in the observed study population due to sampling variability.

We shall hereafter assume that causal graph 2 is a randomized graph. (We would not make such an assumption in a study of the mining industry since miners in ill health are often selectively transferred to less exposed surface jobs.) We shall use the two following equivalent terminologies to characterize the fact that causal graph 2 is a randomized graph: conditional on past employment and exposure history, (1) the subgroups receiving high exposure at work and zero exposure at work at t_s are comparable at t_s, (2) exposure at work was received *at random* at t_s.

The "nonparametric maximum likelihood estimator" (NPMLE) of the superpopulation survival curve associated with a given generalized treatment of randomized graph 2 is obtained by first estimating the unknown superpopulation proportions (i.e. conditional probabilities) necessary to apply the G-computation algorithm by the corresponding sample proportions and then applying the algorithm with the sample proportions in place of the superpopulation proportions.

The following example shows that it is not, in general, possible to empirically estimate the survival curve of a hypothetical study defined by a generalized treatment of a causal graph that is *not* randomized.

Example 2.1. Suppose that causal graph 1 was not a randomized graph, because, among the set of individuals with continuous high exposure at work through t_{s-1}, the subset "off work" at t_s was less healthy on average than the subset still at work at t_s (i.e. the healthy worker survivor effect is operating). Then, for the subset who were "off work" at t_s (in the observed study) we

could not estimate what their mortality would have been had they instead remained at work at t_s in high exposure jobs (since there is no *comparable* group of individuals in the observed study who did remain at work at high exposure). Thus, we could not estimate the population survival curve for the study defined by the generalized treatment "if alive, remain at work and receive high exposure." [It follows, by a similar argument, that for a subset who were observed to be "off work" (and thus unexposed) at t_s, we could not estimate what their mortality would have been if, at t_s, they had instead received high exposure while "off work". It follows that, without further untestable assumptions, we would be unable to estimate the population survival curve of the controlled study "if alive, receive high exposure".]

We stress that in an observational study, the assumption that a measured graph is a randomized graph is always subjective since no amount of empirical evidence can determine whether a measured graph is a randomized graph (or for that matter whether it is a causal graph).

In occupational mortality studies the typical size of the cohort is approximately 10,000 members. Subcohorts defined by a particular calendar date and age at hire will be smaller still. Thus, we cannot hope to consistently estimate the $2^{81} + 2^{80}$-2 conditional probabilities associated with the intra- and internodal lines of the highlighted subgraph of randomized graph 2 (when extended to t_{81}) without specifying parsimonious parametric or semiparametric (e.g. Cox) models. An example is given in Section 4. Unfortunately, in our experience, survival curve estimates tend to be highly sensitive to model choice. Thus, although, in principle, the G-causal parameters of any randomized graph can be precisely estimated in sufficiently large samples, in practice we must face the possibility of unavoidable bias in our estimation of the G-causal parameters of randomized graph 2.

Fortunately, with many occupational exposures a central first question is whether the exposure has any effect on mortality at all, i.e. whether, for randomized graph 2, $S(t, G_1, G_2) = 0$ for all t, and for all possible generalized treatments G_1 and G_2 (which we call the G-null hypothesis for causal graph 2). [Note that if each individual's survival is uninfluenced by his exposure history, then the G-null hypothesis holds for randomized graph 2.] Conveniently, for any randomized graph a completely nonparametric test of the *G-null hypothesis that all G-causal parameters of the randomized graph are identically zero* can be constructed based on the following.

The G-null test algorithm

Step (1) At each time t_s, select the n_s cases who died in the interval (t_{s-1}, t_s). Sample m_s controls without replacement from individuals surviving past t_s. Note that each case will have a covariate history represented by a unique path of intra- and internodal lines extending from t_1 to t_s. (2) For each right circumference point through which at least one of the n_s-case paths pass, determine (a) the subset of the n_s cases and m_s controls who had the covariate history represented by that right circumference point, and (b) enter each subset member into a $2 \times K$ table of case–control status by the K-treatment levels (internodal lines) associated with the right circumference point. (3) If the treatment levels have been quantitatively scored and $K \geqslant 2$, compute the numerator of the Mantel (i.e. log-rank) test for trend for the $2 \times K$ table and the appropriate null variance [6]. (If $K = 1$, the table does not contribute to the test statistics. If $K > 2$ and the treatments have no natural ordering, one could compute the $K - 1$ dimensional numerator of the $K - 1$ sample log rank test). (4) Sum the numerators from each table, possibly after multiplication by a table-specific weight chosen so as to increase power against alternatives felt to be *a priori* likely. (5) Multiply the variance associated with each table by the square of the table weight and sum over tables. (6) Divide the square of the numerator sum by the variance sum and compare to a χ^2_1 distribution. A program "G-null test" that implements the above algorithm, written in Pascal by Donald Blevins and the author, is available upon request. The proof, given in Appendix E of [4], that the G-null test will reject at the nominal level under the G-null hypothesis of a randomized graph is based on the following theorem and the fact that the table-specific numerators are uncorrelated.

G-Null test theorem

The G-null hypothesis holds for a randomized graph if and only if, in the observed study, for any right circumference point the (subsequent) survival curve is the same for each of the treatment groups defined by the internodal lines arising from that right circumference point.

Example. Each 2 × 2 table contributing to the G-null test for randomized graph 2 compares the (case–control) exposure concentrations received at work at time t_k among a subset (of a matched set) of n_s cases and m_s controls all of whose members have the same exposure history through t_{k-1} and employment history through t_k. For example, let $t_s = t_4$ and $t_k = t_2$ and consider the table corresponding to cases (who died between t_3 and t_4) who received high exposure at t_1 and remained at work at t_2. If the control sampling fraction is 1, we see from Fig. 2 that for these cases there are $1 + 1 + 1 = 3$ controls who received high exposure at t_2 and $4 + 2 + 3 = 9$ controls who received zero exposure at t_2. By subtraction, $(5 + 4 + 1) - 3 = 7$ cases dying between t_3 and t_4 received high exposure at t_2 and $(10 + 5 + 5) - 9 = 11$ cases received zero exposure at t_2. This table gives a contribution of $O - E = 7 - [(11 + 7)/(11 + 7 + 3 + 9)] \cdot (3 + 7)$ to the Mantel (i.e. log-rank) numerator. An investigator might give zero weight to tables comparing exposures received less than 10 years prior to the death time of the case if he thought that there was a 10-year biologic latent period.

Although in Section 4 we demonstrate the utility of the G-null test in detecting an adverse exposure effect in a cohort of copper smelter workers, we show, in Example 2.2 below and Example 3.3 in the next section, that an investigator may, at times, be unable to use the G-null test either because (1) the test has poor power or (2) the null hypothesis of interest is not the G-null hypothesis of any randomized graph.

Example 2.2. As an extreme example, suppose for randomized graph 2, (1) at each time t_s, there are 100,000 rather than 2 possible exposure levels at work, and that no two study subjects received identical exposure at t_1. (This might be the case if separate exposure measurements were made on each individual with a precise measuring instrument); (2) among individuals remaining at work, initial exposure level at t_1 is negatively correlated with exposure levels exprienced 10 or more years later; and (3) the true state of nature is such that only exposure experienced within 10 years prior to any time t_k adversely influences mortality at t_k. Then the power of a one-sided G-null test will be less than the alpha level of the test regardless of sample size. This reflects the fact that, in this setting, the G-null test is the Mantel (i.e. log rank) trend test for the association of initial exposure with subsequent mortality (since only at t_1 will any subset of cases and controls have identical previous exposure histories). In realistic examples as well, only exposures experienced soon after time of hire will contribute to the G-null test unless one uses many controls per case (e.g. 50–100), increases Δt from 6 months to 1, 2, or 3 years, and groups exposure into only 3 or 4 levels (say, high, medium, low). Unfortunately, when, for example, high exposure tends to make less healthy individuals preferentially leave work, even under the G-null hypothesis a causal graph formed from randomized graph 2 by either grouping exposure levels or increasing Δt will not itself be a randomized graph. Therefore, the G-null test based on such a causal graph may fail to reject at the nominal rate. Thus, we are faced with the usual trade-off between power and bias.

Fortunately, if, among subjects at work at any t_k, the probability of receiving any particular exposure level at t_k depends only on employment and exposure history in the past m (e.g. 2) years, then a valid, and often much more powerful, test of the G-null hypothesis of randomized graph 2 can be obtained based on the following.

m-Modified G-null test algorithm

Remembering that for randomized graph 2 each table constructed in step 2b of the G-null test algorithm is determined by a death time t_s, an exposure (treatment) time t_k, and a unique exposure history through t_{k-1} and employment history through t_k, we add a step 2c to the algorithm in which we combine into a single table of exposure level at t_k vs case–control status at t_s, those tables that had the same exposure history in the interval $[t_{k-m}, t_{k-1}]$ and employment history in the interval $[t_{k-m}, t_k]$. A program that both implements this modification to the G-null test algorithm and implements a test of the assumption that the probability of receiving exposure at t_k is conditionally independent of exposure and employment history experienced more than m years previously is available from the author.

3. PARADIGMATIC EXAMPLES

In this section we present the examples mentioned in the Introduction. We shall need the following definitions.

Definition

A covariate (e.g. employment history) is a *population (or empirical) risk factor for death*

controlling for exposure history if and only if there exists a time t such that the mortality rate at t differs among subgroups with different covariate but identical exposure histories up to t.

Definition

A covariate (e.g. employment history) is a population (or empirical) *predictor* of *future exposure* if, among subjects alive at some t_{s+1}, the probability of receiving a particular exposure level at t_{s+1} differs among subgroups of the population with different covariate histories but identical exposure histories through t_s.

Definition

Current covariate status (or level) is a predictor of exposure if, among subjects alive at some t_{s+1}, the probability of receiving a particular exposure level of t_{s+1} differs among subgroups with different covariate status (level) at t_{s+1} but identical covariate and exposure histories through t_s.

Definition

A covariate is a population (or empirical) predictor of exposure if it is either (1) a predictor of future exposure or (2) current covariate status is a predictor of exposure.

Definition

A covariate is said to be a *population (or empirical) risk factor for exposure* if it is a population predictor of *future* exposure.

Definition

A covariate is a *causal risk factor for a particular outcome* (e.g. death or exposure history) if and only if there exists some individual and some time t such that the outcome for that individual at t (e.g. vital status or exposure level at t) depends on the individual's covariate history up to t. (Since, for a covariate to be a causal risk factor, a subject's outcome must be defined under at least two different covariate histories, a necessary condition for a covariate to be a causal risk factor is that an investigator could, in principle, supersede nature and manipulate (i.e. control) a subject's covariate status.)

Definition

A covariate is a *causal risk factor for an outcome (e.g. death) controlling for exposure history* if and only if there exists some individual and some time t such that, given some fixed exposure history for the individual, the outcome for that individual at t depends on the individual's covariate history up to t.

Definition

If a covariate is a causal risk factor for an outcome controlling for exposure, we say that the covariate has a *direct* effect on outcome when controlling for exposure.

The above definitions have symmetric counterparts with the role of the covariate and exposure interchanged. [When exposure is a causal risk factor for a particular covariate (as outcome variable) then, in general, a necessary condition for exposure to be a causal risk factor for death controlling for covariate history is that an investigator could, in principle, supersede nature and control both the subject's exposure and covariate status. Without this condition one could not, conceptually, fix a subject's covariate history while manipulating their exposure history.]

Definition

A population risk factor is a *pure causal risk factor for death* (controlling for exposure) in a given study if it is a causal risk factor for death (controlling for exposure) and is always received *at random* at each time t conditional on past exposure and covariate history.

Definition

A covariate is an *intermediate variable* on the causal pathway from exposure to death if (1) it is a causal risk factor for death controlling for exposure and (2) exposure is a causal risk factor for the covariate (as outcome variable).

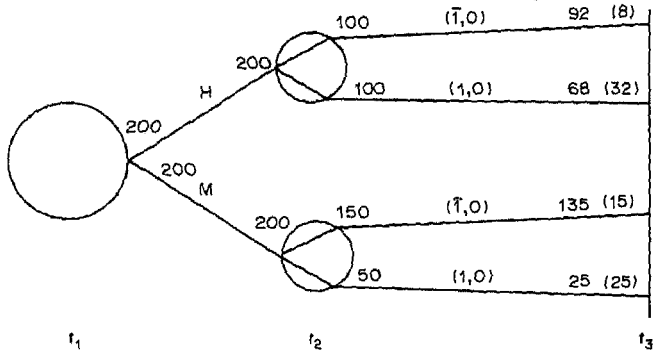
In the examples in this section, we suppose the observed data are at their (large sample) expectations.

Example 3.1. Consider a study represented by the causal graphs shown in Figs 4a and b. Suppose an Investigator A believes that causal graph 4a (but not 4b) is a randomized graph and an Investigator B believes both are randomized graphs. Since in both Figs 4a and b, $p(D \mid \bar{1}, E_{HO}) = 0.08 \neq p(D \mid 1, E_{HO}) = 0.32$ and

$$\tfrac{100}{200} = p(1 \mid E_H) \neq p(1 \mid E_M) = \tfrac{50}{200},$$

employment history is an independent population risk factor for death and exposure history is a population risk factor for employment status at t_2 [where, only in this section, $p(D \mid \bar{1}, E_{HO})$ is the mortality rate at t_3 of the subgroup $(\bar{1}, E_{HO})$; $1(\bar{1})$ refers to employment

(a) A causal graph: Number of deaths at t_3 in parentheses



(b) A causal graph: Number of deaths at t_3 in parentheses

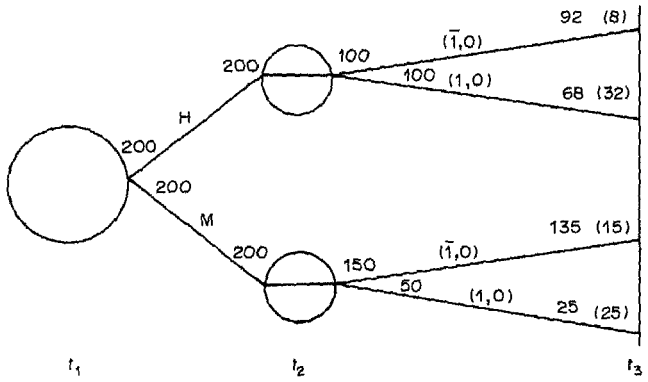


Fig. 4

status at t_2; E_H is high exposure at t_1; and E_{HO} is the exposure history-high exposure at t_1 and zero exposure at t_2]. Since both investigators believe that exposure at t_1 was given at random, both would conclude that exposure was a causal risk factor for employment status at t_2. Investigator B, if correct in his belief, would correctly conclude that employment history is a pure causal risk factor for death (and thus is an intermediate variable on the causal pathway from exposure to death) and, therefore, that the observed mortality difference at t_3 between $(\bar{\mathbf{1}}, E_{HO})$ and $(\mathbf{1}, E_{HO})$ subjects must be due to an effect of the terminated (i.e. unemployed) state *per se*. (Any adverse effect of unemployment *per se* would presumably have been mediated through (unmeasured) factors such as cancellation of health insurance, intensified financial and psychological stresses, and/or increases in cigarette consumption that might result from loss of employment.) Similarly, Investigator B would conclude that the mortality difference between $(\bar{\mathbf{1}}, E_{HO})$ and $(\bar{\mathbf{1}}, E_{MO})$ subjects $[p(D \mid E_{MO}, \bar{\mathbf{1}}) - p(D \mid E_{HO}, \bar{\mathbf{1}}) = 0.10 - 0.08 = 0.02]$ was due to a direct beneficial effect of high exposure when controlling for the intermediate variable employment history. Note it follows from the definition given above that if the direct effect of exposure controlling for employment history is *beneficial* to some study subject (and detrimental to none) then (a) the G-causal parameter of causal graph 4b comparing the controlled study "receive high exposure at t_1 and leave work at t_2" with the study" receive moderate exposure at t_1 and leave work at t_2" *is greater than zero* at some time t or (b) the G-causal parameter obtained by replacing "leave work" with "remain at work" in the above description is greater than zero. If graph 4b is randomized, it follows from the G-computation algorithm that at t_3 this latter G-causal parameter equals

$$p(D \mid E_{MO}, \bar{\mathbf{1}}) - p(D \mid E_{HO}, \bar{\mathbf{1}}) = 0.02.$$

Suppose Investigator B is incorrect and the true state of nature is that (1) $\mathbf{1}$ is not a causal risk factor (in which case, we might call $\mathbf{1}$ a *pure selection factor*), (2) causal graph 4a (but not

Table 1. Analysis of study in Fig. 4a

	1		$\bar{1}$		Crude	
	E_{HO}	E_{MO}	E_{HO}	E_{MO}	E_{HO}	E_{MO}
D	32	25	8	15	40	40
Total	100	50	100	150	200	200
RR	0.64		0.80		1	

RR is stratum specific risk ratio.
$E_{HO}(E_{MO})$ is high (moderate) exposure at t_1, zero exposure at t_2.
1(1) is on (off) work at t_2.

4b) is a randomized graph and (3) exposure has no effect (direct or indirect) on any individual's mortality. *Table 1 shows that the association of exposure history with mortality, controlling for employment history, falsely suggests a beneficial effect of high exposure.* This bias occurs because initial exposure is a population (and causal) risk factor for employment status at t_2, and, because of the healthy worker survivor effect, employment status is an independent population risk factor for death. See Lemma 8.3 of Ref. [4] as corrected in Ref. [7]. (A similar bias can arise in a point exposure trial.) As is well known, when, as in this instance, a dichotomous covariate such as **1** is both associated with exposure and is an independent population risk factor then, when the crude exposure–disease association is null (as it is in Table 1), the stratified associations cannot both be null. By supposition, in our example, it is the crude and not the stratified association that is causal. Stratifying on a covariate **1**, measured after start of follow-up (i.e. at t_2), has induced rather than controlled causal confounding. [Note **1** is not an intermediate variable on the causal pathway from exposure to death since **1** is not a causal risk factor for death. Furthermore, **1** is not a proxy for an (unrecorded) intermediate variable, since, because we have supposed the hypothesis of no exposure effect holds, there is no causal pathway.] Given that the true state of nature is as in (1)–(3) above, the data in Fig. 4a would be as shown if (1) high exposure were an irritant whose only effect was to cause 50 individuals, who would have stayed at work if given moderate exposure at t_1, to leave work at t_2; and (2) these 50 individuals (of whom seven were to die) had on average (a) a poorer prognosis than the 100 individuals (of whom eight were to die) who remain at work when given high exposure at t_1; but (b) a better prognosis than the 50 individuals (of whom 25 were to die) who would leave work regardless of their exposure. It follows from the G-computation algorithm that the crude risk difference, $p(D \mid E_{HO}) - p(D \mid E_{MO}) = 40/200 - 40/200 = 0$, equals the G-causal parameter of randomized graph 4a, $S(t_3, G^{4a} = [M], G^{4a} = [H])$, which measures the overall effect of exposure (and thus, given that **1** is not a causal risk factor, the direct effect of exposure as well).

Example 3.2. Table 2a represents an analysis of a study represented by randomized graph 5. Suppose that the true state of nature is as described in the preceding paragraph. Then the study in Fig. 5 differs from that in Fig. 4a only in that nonzero exposures were given at t_2. A non-causal association between exposure history and death is evident in Table 2a. Suppose that the difference between high and moderate exposure concentrations was less than the difference between moderate and zero exposure concentrations so that the cumulative exposure at t_3 will be greater in individuals with exposure history E_{MM} than E_{HO}. Then, in Table 2a, increasing levels of cumulative exposure would falsely appear beneficial. This bias results from the fact that (1) because of the healthy worker survivor effect, **1** is both an independent population risk factor

(since $p(D \mid \mathbf{1}, E_{MO}) = \frac{25}{50} \neq p(D \mid \bar{\mathbf{1}}, E_{MO}) = \frac{1}{10}$))

and (2) **1** is a population predictor of exposure

(i.e. $1 = p(E_{HH} \mid E_H, \bar{1}) \neq p(E_{HH} \mid E_H, \mathbf{1}) = 0$.)

See Lemmas 8.1 and 8.2 of Ref. [4]. (In Example 3.1, the crude analysis was unbiased because **1** was not a predictor of exposure.) Furthermore, because, as in Example 3.1, exposure at t_1 is a

Table 2(A). Analysis of study in Fig. 5

	E_{HH}	E_{MM}	E_{HO}	E_{MO}
D	8	14	32	26
Total	100	140	100	60
RR	0.18	0.23	0.74	

RR are risk ratios with column E_{MO} as the reference category.
E_{HO} is high exposure at t_1, zero exposure at t_2.

Table 2(B). Stratified analysis of study in Fig. 5

	1		$\bar{1}$		
	E_{HO}	E_{MO}	E_{HH}	E_{MM}	E_{MO}
D	32	25	8	14	1
Total	100	50	100	140	10
RR	0.64		0.8	1.0	

RR are stratum specific risk ratios using, in both strata, column E_{MO} as the reference category.

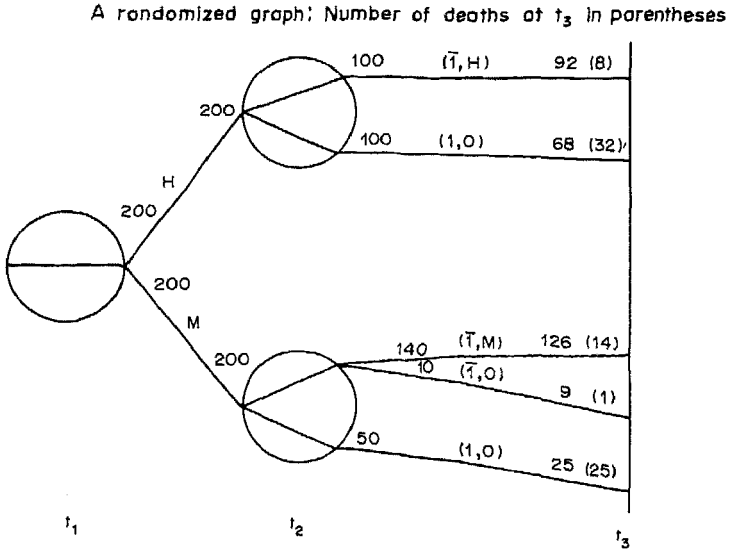


Fig. 5

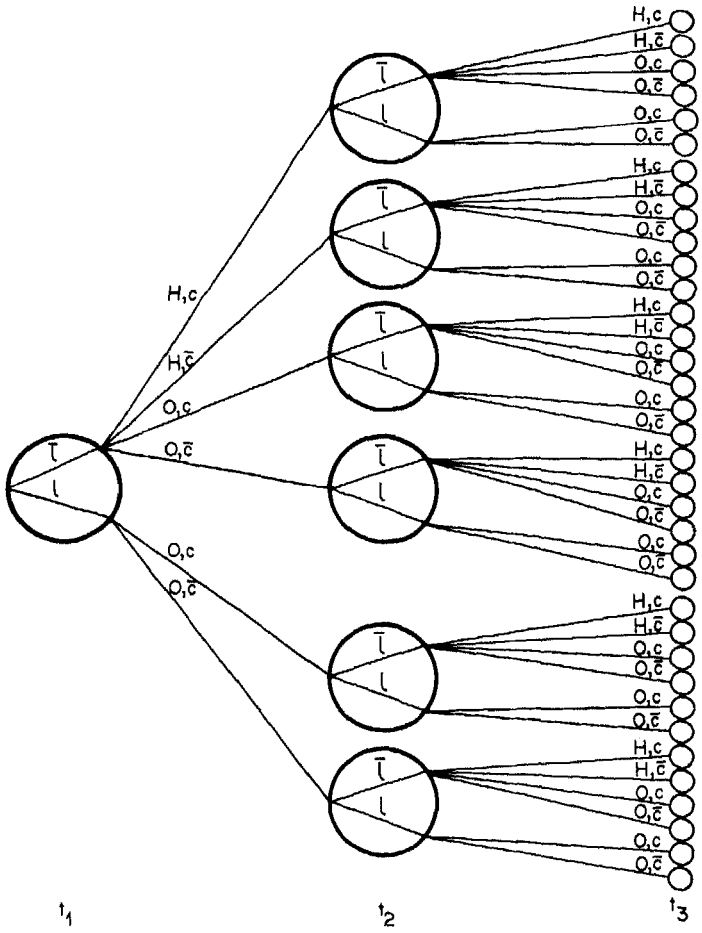


Fig. 6. A randomized graph. c = current smoker; c̄ = current non-smoker.

population risk factor for **1**, Table 2b falsely suggests a beneficial effect of exposure on controlling for employment history. In contrast, if we use the G-computation algorithm applied to randomized graph 5, we correctly conclude that there is no overall exposure effect on mortality (and thus, given **1** is not a causal risk factor, no direct exposure effect as well). Since

$$160/200 = S(t_3|G^5 = ([M], [M]))$$
$$= S(t_3|G^5 = ([M], [0])) = S(t_3|G^5 = ([H], -))$$

[where here we describe a particular generalized treatment of the superscripted causal graph by first placing the treatment chosen at each t in brackets (provided a choice exists) and then listing these treatments in chronological order]. Also note that the G-null test for randomized graph 5 will be valid, since, as required by the G-null test therorem,

$$p(D|E_{MM}, \bar{1}) = p(D|E_{MO}, \bar{1}) = 0.1$$
$$\text{and} \quad p(D|E_H) = p(D|E_M) = 40/200.$$

Example 3.3. In occupational mortality studies in which data on cigarette smoking history has been obtained, interest often centers on assessing the direct effect of exposure controlling for the causal risk factor cigarette smoking history. Suppose that, because of the healthy worker survivor effect, we are only willing to assume that exposure and cigarette smoking at t_s are received at random conditional on cigarette smoking and exposure history through t_{s-1} and employment history through t_s (i.e. Fig. 6 is our finest randomized graph). (In Fig. 6, for simplicity, we have supposed that smoking status at t_s is recorded only as currently smoking and not currently smoking. In actual practice we would normally record at least three to four smoking levels.) If so, as we see below, standard methods [8, 9] of analysis can lead to bias. For example, even under the null hypothesis of no direct exposure effect controlling for (the possible intermediate variable) cigarette smoking, the observed mortality rate at t among individuals with identical cigarette smoking histories can depend on exposure history (whether or not one also adjusts for employment history).

On the other hand, if randomized graph 6 is our finest randomized graph, we can test, using the G-computation algorithm, the null hypothesis of no direct exposure effect on mortality controlling for smoking history from full cohort data. In contrast, we show that it may not be possible to validly test this null hypothesis using case–control data (when the control sampling fraction is unknown). This result implies that one cannot always validly test this null hypothesis using a G-null test (since the G-null test algorithm only requires case–control data). This reflects the fact that the G-null hypothesis for randomized graph 6 implies not only that exposure has no direct effect controlling for cigarette smoking, but also that cigarette smoking has no direct effect on mortality controlling for exposure.

Figure 7b represents data from a case–control study in which all deaths and a random sample of noncases were sampled. If the control sampling fraction were 0.1, Fig. 7a would represent the data on the full cohort. If the sampling fraction were 1, Figure 7b would represent the full cohort. We shall suppose that (1) all subjects were nonsmokers at t_1 and received no exposure at t_2 and (2) **1** is not a causal risk factor. (Both Figs 7a and b can be viewed as special cases of Fig. 6). Obviously, without knowledge of the sampling fraction, we cannot empirically determine from our case–control data which graph represents the full cohort data. But, if Fig. 7a were the randomized graph of the cohort, then we can show using the full cohort data that exposure has no effect on controlling for the intermediate variable smoking history [since, using the G-computation algorithm,

$$p(D > t_3|G_{H,\bar{c}}) = \tfrac{3000}{4000}\tfrac{8000}{16,000} + \tfrac{4000}{8000}\tfrac{8000}{16,000} = \tfrac{10,000}{16,000}$$

and

$$p(D > t_3|G_{0,\bar{c}}) = \tfrac{10,000}{16,000}.$$

$G_{H,\bar{c}}$ is the hypothetical study in which all individuals received high exposure at t_1 and do not smoke at t_2. $G_{0,\bar{c}}$ is similar, except all individuals received zero exposure at t_1.] If Fig. 7b is the randomized graph, then we can show using the full cohort data that exposure has a beneficial direct effect on mortality, controlling for smoking history, [since, then,

$$p(D > t_3|G_{H,\bar{c}}) = \tfrac{300}{1300}\tfrac{4400}{8800} + \tfrac{400}{4400}\tfrac{4400}{8800} = 0.16$$

and

$$p(D > t_3|G_{0,\bar{c}}) = \tfrac{1000}{7000} = 0.14.]$$

Suppose randomized graph 7a represents the full cohort data so that the null hypothesis of no direct exposure effect holds. Even so, the mortality rate at t_3 among subjects with identical smoking histories depends on exposure history (whether or not we adjust for employment history) since

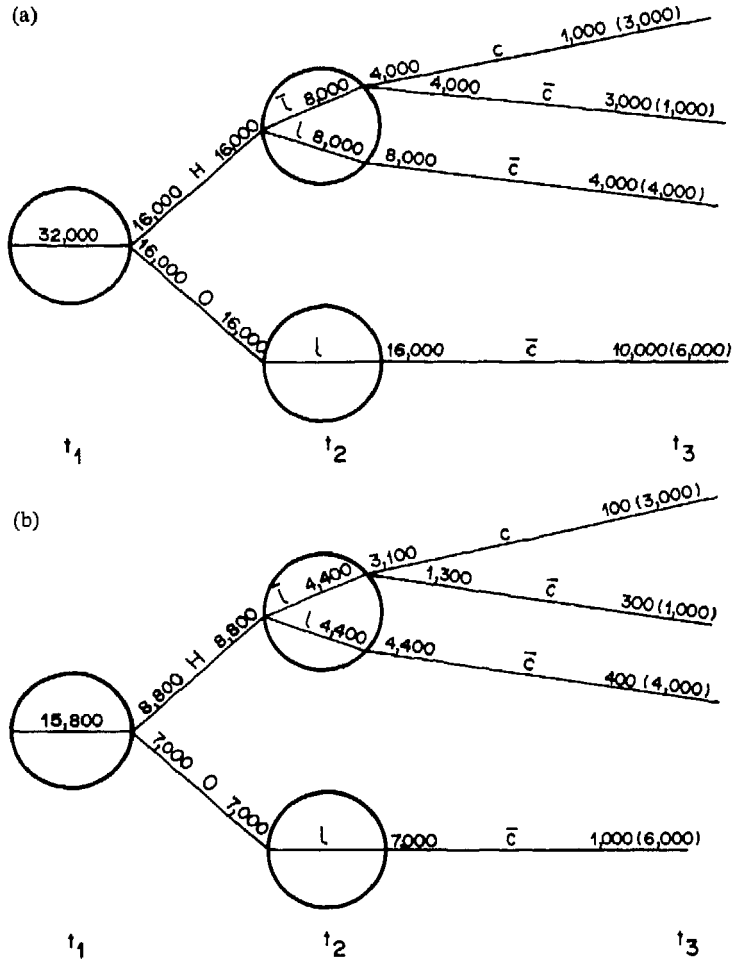


Fig. 7. (a) A randomized graph—number of deaths at t_3 in parentheses. (b) A case–control study—number of deaths at t_3 are shown in parentheses. Number of controls at t_3 are shown without parentheses.

$$p(D \mid E_H, \bar{c}) = 5000/12{,}000 \neq p(D \mid E_0, \bar{c}) = 6000/16{,}000,$$

and

$$p(D \mid E_H, \bar{c}, \mathbf{1}) = 4000/8000 \neq p(D \mid E_0, \bar{c}\,\mathbf{1}) = 6000/16{,}000.$$

Since Fig. 7a is assumed randomized, (1)

$$p(D \mid E_H, c, \bar{\mathbf{l}}) = 3000/4000 \neq p(D \mid E_H, \bar{c}, \bar{\mathbf{l}}) = \tfrac{1000}{4000}$$

implies cigarette smoking is a causal risk factor for death controlling for exposure, and (2)

$$1/2 = p(c \mid E_H) \neq p(c \mid E_0) = 0$$

implies exposure is a causal risk factor for smoking. Thus, cigarette smoking is an intermediate variable on the causal pathway from exposure to death. The parameter

$$p(D > t_3 \mid E_H) - p(D > t_3 \mid E_0) = \tfrac{8000}{16{,}000} - \tfrac{10{,}000}{16{,}000} = \tfrac{-2000}{16{,}000}$$

demonstrates an adverse overall effect of exposure (when not controlling for the intermediate variable cigarette smoking history) since, using the G-computation algorithm, we can compute that this parameter *equals* the unique G-causal parameter of the *coarser* randomized graph derived from randomized graph 7a by placing $(E_H, \bar{\mathbf{l}}, c)$ and $(E_H, \bar{\mathbf{l}}, \bar{c})$ subjects on separate intranodal lines at t_2. (Each generalized treatment of this *coarser* graph represents a study in which the investigator gives all subjects the same exposure at t_1 (either H or 0) and nature determines both cigarette smoking and employment status at t_2.) This equality holds because, in our example, neither cigarette smoking nor employment history are population predictors of exposure. Since exposure has no

direct effect, it follows that the overall exposure effect is completely mediated through its effect on the intermediate variable cigarette smoking. Since neither employment nor cigarette smoking history are predictors of exposure, even if data on cigarette smoking (and employment history) had not been obtained we could still test for (and find) an adverse overall exposure effect by computing $p(D > t_3|E_H) - p(D > t_3|E_0)$, although we would be unable to determine that there was no direct exposure effect.

Before proceeding to our reanalysis of the copper smelter cohort data in Section 4, we give two important caveats.

Caveat 3.1. We show that when the healthy worker survivor effect is operating, one cannot empirically distinguish a direct biological effect of exposure on mortality from an indirect effect of exposure mediated via the intermediate variable employment status. Again, suppose, in the study represented in Fig. 4a, causal graph 4a (but not causal graph 4b) were a randomized graph, since the healthy worker survivor effect is operative. An investigator interested in detecting a biological effect of exposure on mortality would wish to test whether the G-causal parameters of causal graph 4b which represent the direct effect of exposure controlling for (the possible intermediate variable) employment history were zero, e.g. whether

$$S(t_3, G^{4b} = ([H], [\mathbf{1}]), G^{4b} = ([M], [\mathbf{1}])) = S(t_3, G^{4b} = ([H], [\bar{\mathbf{1}}]), G^{4b} = ([M], [\bar{\mathbf{1}}])) = 0. \quad (1)$$

[Of course, this is a test for a biological exposure effect only if there is not a further intermediate variable such as cigarette smoking on the causal pathway from exposure to death (see Example 3.3).] But, as we have seen in Example 2.1, if causal graph 4b is not a randomized graph, one cannot, in general, empirically test whether equation (1) holds. Fortunately, as we discussed in the last sentence of example 3.1, when **1** is not a causal risk factor for death, we can test whether a direct exposure effect exists, since equation (1) implies that the identifiable parameter of randomized graph 4a

$$S(t_3, G^{4a} = ([H]), G^{4a} = ([M])) = 0. \quad (2)$$

That is, if there is no direct effect of exposure (i.e. equation (1) holds) and **1** is not a causal risk factor, then there can be no overall (direct plus indirect) effect of exposure (i.e. equation (2) holds).

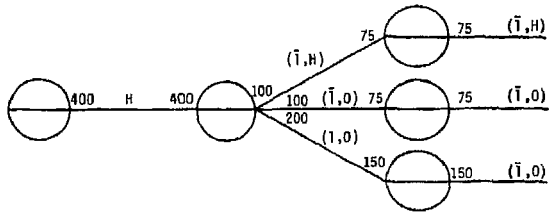


Fig. 8. A randomized graph in which all surviving laid off workers return at t_3.

On the other hand, if **1** is a causal risk factor and exposure is a causal (and population) risk factor for **1**, then equation (2) could be true and (1) false (or (1) false and (2) true) since, in this case, equation (2) implies only that there is no overall effect of exposure (when not controlling for the intermediate variable employment history). In this setting, if we are interested in the biological effect of exposure, we are in a bind: the causal parameters of randomized graph 4a (or, more generally, randomized graph 2) can be estimated (using the G-computation algorithm), but they have little substantive interest. On the other hand, we cannot empirically estimate those causal parameters of causal graph 4b (or, more generally, causal graph 1) that do have substantive interest. Furthermore, when the healthy worker survivor effect is operating and thus **1** is a population risk factor for death, we cannot empirically determine whether or not **1** is also a causal risk factor. *Henceforth, we shall assume that* **1** *is not a causal risk factor.*

Caveat 3.2. An artifactual healthy worker survivor effect due to "lay-off". Some investigators who have attempted to demonstrate the empirical existence of the "healthy worker survivor effect" have defined "date of termination" as "date of last employment" and ignored all other information on employment history. These same investigators then typically have defined a subject's "duration of employment" up to time t_s to be the absolute value of the difference between date of hire and the minimum of t_s and "date of termination." In studies of an industry in which workers frequently leave and then return to work (for example, due to economic layoff), such a practice will artifactually create the appearance that "the healthy worker survivor effect is operating" even when it is not. As an example, suppose Fig. 8 is a randomized graph representing the results of an occupational mortality study. Then (1) there is no healthy worker survivor effect (since, by definition of a random-

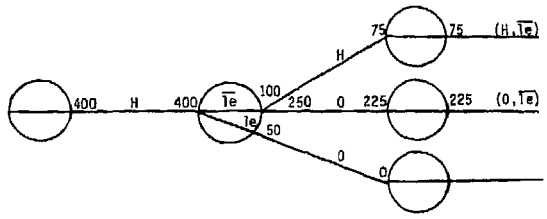


Fig. 9. A measured graph derived from the randomized graph in Fig. 8 where an individual is 1e after leaving work for the last time.

ized graph, the subsets of individuals leaving and continuing at work at t_2 are comparable); (2) neither employment nor exposure history is an independent population (or causal) risk factor for mortality, since

$$p(D \mid \bar{\mathbf{1}}, E_{HH}) = p(D \mid \bar{\mathbf{1}}, E_{HO}) = p(D \mid \mathbf{1}, E_{HO}) = 25/100; \quad (3)$$

and (3) since 1 is not an independent population risk factor, standard methods of analysis are unbiased whether one performs a crude analysis [since $p(D \mid E_{HH}) = p(D \mid E_{HO}) = 25/100$] or a stratified analysis (see equation (3)). Note that of the 200 individuals who left work at t_2, all 150 survivors return to work at t_3.

An investigator who defined "date of termination" as date of last employment would create the measured graph shown in Fig. 9 where an individual is **1e** at t_2 if he has "terminated employment" at a time less than or equal to t_2 and is $\bar{\mathbf{1}}\bar{\mathbf{e}}$ otherwise. 1e status (equivalently, "duration of employment" as defined above) is an enormously strong population risk factor for death on controlling for exposure history since

$$25/250 = p(D \mid E_{HO}, \bar{\mathbf{1}}\bar{\mathbf{e}}) \neq p(D \mid E_{HO}, \mathbf{1e}) = 50/50 = 1.$$

Many investigators have incorrectly interpreted this observation as evidence *that "unhealthy individuals tend to leave work"*. Of course, in our example **1e** status (and, equivalently, "duration of employment") is a population risk factor only because, when we defined "date of termination" as "date of last employment", the future event, death, determines the previous event, "date of termination". An investigator, on discovering that "date of termination" ("duration of employment") is a population risk factor, often then chooses to adjust for this risk factor in order to "control confounding." Such an investigator would falsely conclude that exposure is harmful, since

$$25/100 = p(D \mid E_{HH}, \bar{\mathbf{1}}\bar{\mathbf{e}}) > p(D \mid E_{HO}, \bar{\mathbf{1}}\bar{\mathbf{e}}) = 25/250$$

The bias associated with inappropriately defining "date of termination" as date of last employment cannot be overcome by using the analytic methods proposed in this paper. For example, if an investigator assumes that measured graph 9 were a randomized graph, he would also discover an artifactual adverse exposure effect since, using the G-computation algorithm,

$$S(t_3 \mid G^9 = ([H], [H])) = \tfrac{350}{400}\tfrac{75}{100} + (\tfrac{50}{400})(\tfrac{0}{50}) < S(t_3 \mid G^9 = ([H], [0])) = \tfrac{350}{400}\tfrac{225}{250} + (\tfrac{50}{400})(\tfrac{0}{50}).$$

The problem, of course, is that measured graph 9 is not a randomized graph and thus we cannot validly apply the G-computation algorithm, since the 100 individuals with covariates ($\bar{\mathbf{1}}$e, H) measured at t_2 are not comparable to the 250 ($\bar{\mathbf{1}}\bar{\mathbf{e}}$, 0) individuals. This follows because the 150 of the 250 ($\bar{\mathbf{1}}\bar{\mathbf{e}}$, 0) individuals who had been off work at t_2 are *guaranteed* to survive to t_3 *by the very fact that they are* $\bar{\mathbf{1}}\bar{\mathbf{e}}$ *rather than* **1e**. In summary, one should not use date of last employment to summarize a subject's employment history.

4. A WORKED EXAMPLE

Lee and Fraumeni[10] assembled data on 8047 arsenic-exposed white males who worked at a Montana copper smelter. A worker was entered into follow-up at the latest of 1 January 1938 and his date of first hire. Follow-up ended on 31 December 1977. Our analysis is based on the 5947 workers who were hired subsequent to 1 January 1935. We chose this date because, after 1935, exposures are known to have been less than earlier and we were interested in trying to detect small effects associated with lower doses. We abstracted information on vital status, employment status, and exposure status (coded for the purposes of this analysis as high = 3, medium = 2, low = 1, unexposed = 0) at 6 monthly intervals. We fit the following logistic models for the conditional probabilities of measured graph 2 (modified so that there are four internodal lines arising from the superior right circumference point in each node, representing high, medium, low and zero exposure at work).

$$\text{logit}[\gamma_{D,i}[t + \Delta t \mid E(t), L(t), Z(t_1)]] = \beta_{0,D,i,t} + \boldsymbol{\beta}_D \cdot \mathbf{X}_D \quad (4)$$

$$\text{logit}[\gamma_{L,i}[t+\Delta t \mid E(t), \mathbf{1}(t) = 0, D > t+\Delta t, L(t-\Delta t), Z(t_1)]] = \beta_{0,L,i,t} + \boldsymbol{\beta}_L \cdot \mathbf{X}_L \quad (5)$$

$$\text{logit}[\gamma_{R,i}[t+\Delta t \mid E(t), D > t + \Delta t, L(t-\Delta t), \mathbf{1}(t) = 1, Z(t_1)]] = \beta_{0,R,i,t} + \boldsymbol{\beta}_R \cdot \mathbf{X}_R \quad (6)$$

where (1) i is a stratum indicator and we have stratified by 5-year joint intervals of age at hire and calendar year of hire. (2) t is age where we allow t to take on only the values Δt, $2\Delta t$, $3\Delta t, \ldots,$ with $\Delta t = 0.5$ years. (3) $\gamma_{D,i}(t \mid \cdot)$ is the conditional probability of being dead by age t, given that one was alive at age $t - \Delta t$ in stratum i with covariate history represented by "$\cdot$". (4) $Z(t_1)$ is all information on the subject up to age at start of follow-up t_1. $E(t)$ and $L(t)$ are respectively exposure and employment history through t. (5) $\gamma_{L,i}(t \mid \cdot)$ is the conditional probability of leaving work at age t, $t > t_1$, given one is alive at t and was at work at $t - \Delta t$ (i.e. $\mathbf{1}(t - \Delta t) = 0$) in stratum i given "$\cdot$". (6) $\gamma_{R,i}(t \mid \cdot)$ is the conditional probability of returning to work at age t ($t > t_1$) given one was out of work at $t - \Delta t$ and alive at t in stratum i given "$\cdot$". For illustrative purposes we choose

$$\boldsymbol{\beta}_D \cdot \mathbf{X}_D = \beta_{1,D}[ce(t) \cdot [1 - \mathbf{1}(t)]] + \beta_{2,D}\mathbf{1}(t) + \beta_{3,D}[Cl(t) \cdot \mathbf{1}(t)] + \beta_{4,D}[ce(t) \cdot \mathbf{1}(t)] \quad (7)$$

$$\boldsymbol{\beta}_L \cdot \mathbf{X}_L = \beta_{1,L} ce(t) + \beta_{2,L}[I_R \cdot CR(t)] + \beta_{3,L}[ce(t) \cdot (t - t_h)] \quad (8)$$

$$\boldsymbol{\beta}_R \cdot \mathbf{X}_R = \beta_{1,R} ce(t) + \beta_{2,R}[ce(t) \cdot (t - t_h)] + \beta_{3,R}[Cl(t)] \quad (9)$$

where (1) $e(u)$ is a quantitative estimate of exposure concentration received at time u and $ce(t)$ is the sum of the biannual measurements, $e(u)$, taken on an individual up to age t. $ce(t)$ is a discretized version of lifetime cumulative exposure. (2) t_h is age at hire. (3) $\mathbf{1}(t) = 1$ if an individual is out of work at t and $\mathbf{1}(t)$ is zero otherwise. (4) $Cl(t)$ is the number of years since last at work. (5) I_R is an indicator variable taking the value 1 if an individual has ever been off work since time of hire and 0 otherwise. (6) $CR(t)$ is the number of years elapsed since an individual was last off work.

With J labeling failure type (i.e. $J \in \{L, D, R\}$) we estimated the β_J by "conditional logistic regression". We estimated the age–disease–calendar period specific parameters $\beta_{0,J,i,t}$ by logit $d_{J,i,t}/\Sigma_k e^{\beta_J \cdot \mathbf{X}_{J,k}}$ where $d_{J,i,t}$ is the number of failures of type J in stratum i at age t and k indexes the individuals in the risk set of the $d_{J,i,t}$ cases [11]. The results are given in Table 3. The modelling assumptions made in equations (4–9) allow us to estimate all the conditional probabilities of measured graph 2 needed to apply the G-computation algorithm. [In fact, these modelling assumptions would allow us to use the G-computation algorithm to estimate the survival curve of the controlled trial "if alive, receive high exposure". Nonetheless, we choose not to estimate this survival curve, since, even when causal graph 2 is randomized, we have no way of empirically testing whether these models accurately reflect the risk of death at time t for subjects receiving high exposure while off work at $t - \Delta t$.]

Table 3. Maximum conditional likelihood estimates of coefficients of equations (4)–(6) with covariate specification given in equations (7)–(9)

	Failure type		
	Death (D)	Leaving (L)	Return (R)
β_1	−0.003 (0.0015)*	−0.049 (0.011)	0.023 (0.025)
β_2	0.16 (0.06)	−0.008 (0.013)	−0.0001 (0.0070)
β_3	−0.008 (0.003)	0.0016 (0.0004)	−0.367 (0.031)
β_4	0.002 (0.0013)	—	—

*Standard errors in parenthesis.

For individuals who, in the observed trial, had $Z(t_1)$ defined by being born in 1902, beginning work in 1935, and staying at work continually at high exposure jobs until start of follow-up in 1938, Table 4 gives estimates, based on the Monte Carlo G-computation algorithm of the Appendix, of the survival curves up to age 35, 55, and 75 that would have been observed in hypothetical studies beginning in 1938 defined by the generalized treatment of randomized graph 2 "if at work, receive high exposure" and "if at work receive zero exposure". We cannot determine whether the apparent adverse effect of exposure in Table 4 was statistically significant without performing bootstrap replications of our model fitting procedures. Even if we had, the p value obtained may be very dependent on our model choice. Instead, we computed the G-null test for randomized graph 2 based on a 1:50 case–control ratio with controls matched to cases on time since hire and to within 6 months on age at hire and within three years on calendar period of hire. The

Table 4. Comparison of the estimated survival probabilities in hypothetical studies of a cohort of arsenic-exposed copper smelter workers

Study	Age 35	55	75
G_H[1]	0.992[3] (0.0005)[4]	0.873 (0.013)	0.325 (0.005)
G_0[2]	0.992 (0.0005)	0.874 (0.004)	0.370 (0.002)

[1]Generalized treatment of "if at work past start of follow-up, receive high exposure" for the subset of the observed study population hired at age 33 in 1935 remaining on work at high exposure until start of follow-up in 1938.
[2]Generalized treatment of "if at work past start of follow-up, receive zero exposure" for the same subset of the population.
[3]Estimated survival probability.
[4]Empirical Monte Carlo standard error.

G-null test χ^2_1 values (when using equal weights for all tables) was 9.0 ($p < 0.003$), suggesting a significant effect of exposure on mortality. In addition, a summary Mantel–Haenszel odds ratio of 1.16 was estimated (using all tables that contributed to the G-null test) with high or medium exposure constituting the "exposed" and low or zero exposure "the unexposed". We report this Mantel–Haenszel odds ratio to facilitate the interpretation of our test of the null hypothesis. Even if the null hypothesis of no exposure effect is true, in a sufficiently large sample, any valid test will reject the null hypothesis with near certainty due to residual confounding (i.e. due to the fact that causal graph 2 is not exactly a randomized graph). Our p value of 0.003 is more likely to be due to residual confounding when associated with a Mantel–Haenszel odds ratio of 1.16 than had it been associated with a Mantel–Haenszel odds ratio of, say, 2. This reflects the fact that in the latter case a strong confounder must have been overlooked, which is *a priori* less likely.

In order to compare our results to standard methods of analysis, we refit equation (4) with

$$\beta_D \cdot X_D = \beta_1 ce(t) \qquad (10)$$

We used several scoring systems for exposure concentration, $e(u)$. The χ^2_1 values for testing $\beta_1 = 0$ in equation (10) were 0.00001; 1.50; 2.25 based on scoring $e(u)$ as 1, 2, 3; 0, 1, 2; 0, 1, 4, respectively, for low, medium, and high exposure. In addition, we replaced cumulative exposure in equation (10) with cumulative exposure lagged 9 and 15 years (i.e. $ce(t-9)$, $ce(t-15)$) as well. With cumulative exposure lagged 9 years, the χ^2 values were 0.4 and 2.2 for the first two of the three scoring systems given above. With cumulative exposure lagged 15 years, these χ^2 values remained insignificant at 0.8 and 2.6.

To see whether our inability to detect an exposure effect using standard methods was due to the healthy worker survivor effect, we used the Monte Carlo G-computation algorithm to estimate for the subset of workers who remained at work at high exposure through age 56 (with the $Z(t_1)$ mentioned above), the probability of survival past age 75 for those who left work at age 57 compared to those who remained at work at 57 when both groups were treated, beginning at age 57, with the generalized treatment "if at work receive zero exposure." The survival probability for the "off work at 57 group" was 0.26, and for the "at work at 57 group" 0.31, suggesting, that the healthy worker survivor effect is operative (under the assumption that **1** is not a causal risk factor—see Caveat 3.1).

Finally, from Table 3, we can see that exposure is a population risk factor for leaving work since $\beta_{1L} = -0.049$ and leaving work is an independent population risk factor for death ($\beta_{2D} = 0.016$). Thus, as shown in Examples 3.1 and 3.2, a statistical association between cumulative exposure and mortality when controlling for employment history may not be causal when causal graph 2 (but not causal graph 1) is randomized. In particular, the apparently beneficial effect of exposure among individuals at work ($\beta_{1D} = -0.003$) presumably reflects selection effects rather than true causality. (It could, of course, also represent sampling variability or model misspecification.)

Even more surprisingly, the G-null test for lung cancer (i.e. the G-null test modified so that only lung cancer cases are admitted as cases) gives a χ^2 value of 18 and summary Mantel–Haenszel odds ratio of 2.0 while the χ^2 values for β_1 in equation (10) (when modified so that only lung cancer deaths are considered) is not greater than 3.0 for any of the exposure scoring systems and lag periods discussed above. Empirical evidence that this result might be due in part to a healthy worker survivor effect lasting for more than nine years was obtained by estimating equation (4) with $\beta_D \cdot \mathbf{X}_D = \beta_1 \text{dur}(t-9) + \beta_2 \text{rate}(t-9)$ where $\text{dur}(t-9)$ is number of years at work until 9 years before age at risk, t, and $\text{rate}(t-9) = ce(t-9)/\text{dur}(t-9)$. The coefficient for β_1 was negative ($\chi^2_1 = 3.5$) and for β_2 positive ($\chi^2_1 = 12$). (Note, since we have matched on age at hire, duration is inversely correlated with

total time off work.) Since it is implausible that, at the individual level, increased duration of exposure would diminish one's risk of lung cancer (when controlling for exposure rate), the negative β_1 presumably reflects the fact that, as discussed in the introduction, individuals at increased risk (i.e. cigarette smokers) tend to preferentially leave work.

The above G-null test for lung cancer is valid only if causal graph 2 is randomized. A nearly necessary condition (condition 1) for causal graph 2 to be randomized is that the exposure level received at work at t_s among subjects with identical exposure and employment histories through t_{s-1} does not depend on cigarette smoking history (see Lemma 8.6 of Ref. [4]). Since data on cigarette smoking was not obtained, we cannot empirically test whether condition 1 holds. If causal graph 2 is randomized, the above χ^2 value of 18 is valid evidence against the null hypothesis of no overall effect of exposure on lung cancer when not controlling for (the possible intermediate variable) cigarette smoking history. But, as discussed in Example 3.3, the causal question of interest in most occupational studies would be whether exposure has a direct (biological) effect on lung cancer when controlling for cigarette smoking. When, as in our smelter worker cohort, exposure is a risk factor for leaving work and employment history is a population risk factor for lung cancer mortality, the χ^2 value of 18 from the G-null test for lung cancer (based on randomized graph 2) is valid evidence against the null hypothesis of no direct exposure effect controlling for cigarette smoking history only if, at every time, neither past exposure nor past employment history is a causal risk factor for current smoking status on controlling for past smoking history (condition 2) (see Section 8D.3 of Ref. [4]). This reflects the fact that cigarette smoking will be an intermediate variable on the causal pathway from exposure to death when condition 2 does not hold. If cigarette smoking were an intermediate variable, the evidence cited above against the null hypothesis of no overall exposure effect would not count as valid evidence against the null hypothesis of no direct exposure effect. [Note in Example 3.3 we were unable to use a test of the hypothesis of no overall effect of exposure (when not controlling for cigarette smoking history) to test for a direct exposure effect because condition 2 was false.] We argue in sections 8D.3 and 12 of Ref. [4] that in the substantive context of our mortality study of copper smelter workers, it might be reasonable to assume that conditions (1)–(2) are adequate approximations to the truth. (Technically, for reasons discussed in Section 12 of [4], the results given in this paragraph are valid only if the effect of exposure on causes of death other than lung cancer is negligible.)

5. SUMMARY AND DISCUSSION

In Examples 3.1 and 3.2, we have shown that when there exists a time at which current risk factor status is an empirical predictor of future or current exposure to the agent under study and is determined by previous exposure, the association of mortality with observed exposure history may have no causal interpretation whether or not one adjusts for the risk factor in the analysis (provided the risk factor is not a pure causal risk factor). It follows that if, in analyzing an occupational mortality study, one finds that (1) employment history is an independent population risk factor for death (if due in part to subjects at increased risk preferentially leaving work), and (2) exposure history is an independent population risk factor for subsequent employment history, one should undertake the following formal approach.

(1) Construct measured graphs based on covariate and vital status information recorded every Δt years with Δt chosen short enough (e.g. 6 months in occupational mortality studies) that recorded covariate (e.g. exposure and employment) history is a suitable approximation to true covariate history.

(2) Determine one's finest causal graph A (e.g. causal graph 1 in our study).

(3) Determine the finest causal graph B (e.g. causal graph 2) coarser than causal graph A that one is willing to assume is a randomized graph. This requires empirically untestable assumptions about the randomness of nature's exposure and/or covariate assignments. (One could perform a sensitivity analysis by analyzing the data under various assumptions as to the finest randomized graph.)

(4) Estimate any G-causal parameters of randomized graph B (or of any coarser randomized graph) that are of substantive interest by employing statistical models to estimate the conditional probabilities required by the G-computation algorithm (see Section 4). Perform a sensitivity analysis by varying the form of these models.

(5) If any null hypothesis of substantive interest can be expressed as the G-null hypothesis of a randomized graph perform a "distribution-free" G-null test. (See Example 3.3 and Section 4, particularly the last paragraph. Example 3.3 is the only published example that we know of in which a causal null hypothesis of interest can be validly tested using *cohort* data but *not* using *case–control* data.) In general, to increase the power of the G-null test, one would increase the number of controls selected per case and either (a) group the measured exposure levels into three or four categories and record data at less frequent intervals (say, $\Delta t = 1$ or 2 years), provided one feels asured that important bias is not thereby introduced (see Example 2.2) or (b) test to determine whether it is valid to utilize an m-modified G-null test such as that described in Example 2.2.

The above approach is formal and does not preclude informal criticism and speculation. For example, if one's interest is in the direct effect of exposure controlling for employment history, employment history is believed to be an intermediate variable on the causal pathway from exposure to disease, and causal graph 1 is not believed to be a randomized graph because of the healthy worker survivor effect, one can still speculate about the magnitude of the non-identifiable G-causal parameters of interest, i.e. those of causal graph 1. (See Caveat 3.1.)

It is common for current risk factor status to be a predictor of future (or current) exposure and to be influenced by past exposure in nonoccupational studies in which there is "treatment by indication." For example, since physicians frequently withdraw women from exogenous estrogens at the time they develop hypertension or angina, it follows that, in a study of the overall effect of exogenous estrogens on cardiac mortality, the cardiac risk factors angina and hypertension would determine subsequent exposure and might be, in part, determined by past exposure to estrogens. In such studies, the analytic approach outlined in steps 1–5 above should be employed.

The approach developed in this paper is shown in Ref. [4] to be equivalent to an approach in which an observational study is identified with a hypothetical double blind randomized trial in which data on each subject's assigned treatment protocol have been erased from the data file. In such a trial, causal inferences can be made by comparing mortality as a function of treatment protocol, since, in a double blind randomized trial missing data on treatment protocol, the association of mortality with treatment protocol can still be estimated.

Finally we note that one may be willing to assume that the causal graph of Fig. 2. is *randomized with respect to survival*, and yet be unwilling to assume the graph is randomized, where we define a causal graph to be *randomized with respect to survival* if it satisfies the conditions of Definition 1 when the phrase "covariate and survival histories" is replaced by the phrase "survival histories." The effect of the above assumptions on our ability to make causal inferences from the data contained on measured graph 2 is considered in Ref. [12].

Acknowledgements—Research for this paper was funded in part by grants from the following organizations: National Institute of Environmental Health Sciences No. 5 R23 ES03405. American Lung Association No. LA 3/22/85, National Institute for Occupational Safety and Health No. 5 R03 OH01867, National Institute of Environmental Health Sciences No. 5 P30 ES00002, the Alfred P. Sloan Foundation and Mobil Corporation, and the American Heart Association. The author wishes to thank Sander Greenland, David Freedman, Alice Whittemore, and several referees for valuable suggestions.

REFERENCES

1. Gilbert ES. Some confounding factors in the study of mortality and occupational exposures. **Am J Epidemiol** 1982; 116: 177–189.
2. Robins JM. A new theory of causality in observational survival studies—Application to the healthy worker effect. **Biometrics** 1985; 41: 311.
3. Robins JM. A statistical method to control for the healthy worker effect in intracohort comparisons. **Am J Epidemiol** 1984; 120: 465.
4. Robins JM. A new approach to causal inference in mortality studies with a sustained exposure period—Application to control of the healthy worker survivor effect. **Math Model** 1986; 7: 1393–1512.
5. Rubin DB. Bayesian inference for causal effects: The role of randomization. **Ann Stat** 1978; 6: 34–58.
6. Mantel N. Chi-square tests with one degree of freedom: Extensions of the Mantel–Haenszel procedure. **J Am Stat Assoc** 1963; 58: 690–700.
7. Robins JM. Errata for "A new approach to causal inference in mortality studies with a sustained exposure period—applications to control of the healthy worker survivor effect." **Comput Math Appl** 1987; to appear.
8. Thomas DC. General relative risk models for survival time and matched case-control analysis. **Biometrics** 1981; 37: 673–686.
9. Whittemore AS, McMillan A. Lung cancer mortality among U.S. uranium miners: A reappraisal. **J Natl Cancer Inst** 1983; 71(3): 489–499.
10. Lee AM, Fraumeni JF. Arsenic and respiratory cancer in man: An occupational study. **J Natl Cancer Inst** 1969; 42: 1045–1052.
11. Breslow NE. Contribution to discussion of a paper by D. R. Cox. **J R Stat Soc** 1972; B34: 216–217.
12. Robins JM. Addendum to "A new approach to causal inference in mortality studies with a sustained exposure period—Application to control of the healthy worker survivor effect." **Comput Math Appl** 1987; To appear.

APPENDIX

The Monte-Carlo G-Computation Algorithm

To estimate, say, $S(t_{81}|G)$, for a particular generalized treatment, G, of randomized graph 2, we can use the following Monte-Carlo algorithm. Note first that any single path of intra- and internodal lines that lie on the highlighted subgraph G and connects the left circumferences of the t_1 node to a node at t_{81} can be uniquely identified by one of 2^{80} possible employment history paths $L(t_{80})$. We choose at random one of these 2^{80} employment history paths, say, $L_i(t_{80})$, by flipping a coin 80 times with success probabilities $P_1, \ldots, P_{80}$ chosen in such a way that P_{80} depends on the outcomes of the preceding 79 tosses $H_1, \ldots, H_{79}$ [where $H_j = 1(0)$ if toss j was a success (failure)] as follows:

$$P_1 = p[1(t_1) = 1], P_2 = p[1(t_2) = 1 | L_i(t_1) = H_1, E(t_1)], \ldots,$$

and

$$P_{80} = p[1(t_{80}) = 1 | L_i(t_{79}) = (H_1, \ldots, H_{79}), E(t_{79})],$$

where $L_i(t_s)$ is the randomly selected employment history through t_s, $1(t_s) = 1$ if out of work at t_s, and exposure history, $E(t_s)$, is uniquely determined by G and $L_i(t_s)$. Note P_s is the intranodal conditional probability entered in Step 1 of the G-computation algorithm associated with the right circumference point defined by $1(t_s) = 1$, $L_i(t_{s-1})$, $E(t_{s-1})$. Then, $L_i(t_{80}) = (H_1, \ldots, H_{80})$. We next compute

$$S[L_i(t_{80})] \equiv \prod_{s=1}^{80} p[D > t_{s+1} | L_i(t_s) = (H_1, \ldots, H_s), E(t_s), D > t_s].$$

$S[L_i(t_{80})]$ is the product of conditional internodal survival probabilities entered in Step 1 of the G-computation algorithm that lie on the unique path of intra- and internodal lines determined by G and $L_i(t_{80})$. Then,

$$\sum_{i=1}^{N} S[L_i(t_{80})]/N$$

converges in probability to $p(D > t_{81}|G)$ as $N \to \infty$ where N is the number of random paths chosen as above. In our experience, the choice $N = 200$ often gives Monte-Carlo standard errors of 0.005 or less. In practice, of course, we will only have estimates of the intra- and internodal conditional probabilities needed to estimate $S(t_{81}|G)$. A PASCAL program implementing this algorithm is available from the author.