Controlling the healthy worker survivor effect: an example of arsenic exposure and respiratory cancer

H M Arrighi, I Hertz-Picciotto

Abstract

Objective—This investigation sought to examine whether methods proposed to control the healthy worker survivor effect would influence the shape or magnitude of the dose-response curve for respiratory cancer induced by arsenic.

Methods—Results from an unadjusted analysis are compared with results obtained by applying four different methods for control of the healthy worker survivor effect to data on arsenic exposure and respiratory cancer. The four methods are: exposure lag, adjustment for work status, cohort restriction, and the G null test.

Results—Cohort restriction gave erratic results depending upon the minimum years of follow up used. Exposure lag substantially increased the rate ratios and a non-linear shape (decreasing slope) compared with an unlagged analysis. Adjusting for work status (currently employed *v* retired or otherwise not employed) yielded slightly higher rate ratios than an unadjusted analysis, with an overall shape similar to the baseline analysis. Results from the G null test procedure of Robins (1986), although not directly comparable with the baseline analysis, did show an adverse effect of exposure that seemed to reach a maximum when exposure was lagged between 10 and 20 years.

Conclusions—All results confirm an adverse effect of arsenic exposure on respiratory cancer. In these data, it seems that the healthy worker survivor effect was not strong enough to mask the strong effect of arsenic exposure on respiratory cancer. Nevertheless, several methods show a stronger association between arsenic exposure and respiratory cancer after adjustment for the healthy worker survivor effect, suggesting that for weaker causal associations, studies not controlling for this source of bias will have low power to detect results. Although the G methods are theoretically the most unbiased, further work elucidating the validity of the assumptions underlying lagging, adjustment for work status, and the G methods are needed before clear recommendations can be made.

(*Occup Environ Med* 1996;53:455-462)

Keywords: healthy worker effect; bias; arsenic; respiratory tract neoplasm

Department of
Epidemiology, School
of Public Health, The
University of North
Carolina at Chapel
Hill, Chapel Hill, NC,
USA
H M Arrighi
I Hertz-Picciotto

Currently with:
Amgen, Thousand
Oaks, CA, USA
H M Arrighi

Correspondence to:
Dr H M Arrighi, Amgen,
M/S 17-2-C, 1840
DeHavilland Drive,
Thousand Oaks, CA 91320-
1789, USA.

Accepted 21 February 1996

Risk estimates from investigations in highly exposed populations are used to extrapolate risk at lower levels and to establish exposure standards. For many chemical and physical agents, occupational cohorts are the most highly exposed. However, occupational studies are subject to bias from the healthy worker survivor effect, by which workers who are least healthy are most likely to leave work.[1] This differential may attenuate the dose-response relation and has the potential to alter the shape of the dose-response curve.[2,3] The present investigation sought to examine whether methods proposed to control the healthy worker survivor effect would influence the shape or size of the dose-response for respiratory cancer induced by arsenic.

A common belief is that the healthy worker effect is small for cancer, especially lung cancer. Although smaller than for cardiovascular mortality, particularly in the period shortly after hire,[4] a healthy worker effect for cancer has been found in numerous occupational studies that show mortality rates lower than expected for selected cancers.[5,6,7] Robins[8,9] empirically shows the bias of the healthy worker survivor effect for mortality from lung cancer in an occupational cohort exposed to arsenic. Reinforcing these empirical findings are: (*a*) the recognition that the healthy worker survivor effect may operate after the workforce is selected if there are factors related to both the end of employment and mortality, and (*b*) the realisation that such factors are plausible even with cancers.[1,8]

Although use of internal comparisons can successfully control the effect of healthier people being hired into the workforce, the healthy worker survivor effect continues to operate even when internal control populations are used.[8] We reviewed in a previous paper how the healthy worker survivor effect has been historically conceptualised.[1] At least four methods have been proposed to control for the healthy worker survivor effect in cohort studies of occupational mortality.

(1) The analysis is restricted to long term survivors where it is assumed that the healthy worker survivor effect is minimal.[4,10,11] Workers with shorter periods of follow up are excluded as they are assumed to have a different propensity for mortality.

(2) The exposure is lagged so that recent exposures are ignored. The rationale for this approach is that recent exposures are incurred only among the healthiest survivors.[12]

(3) A covariate is introduced to control for current employment status (active *v* inactive

employment).[13] This method assumes that the healthy worker survivor effect is a case of traditional confounding: (a) work status is related to exposure in that those with higher exposure may leave work more readily and while off work, incur no further exposure; (b) being off work is an independent risk factor for the outcome.

(4) The healthy worker survivor effect is assumed to operate simultaneously as a confounding and an intermediate variable. Under this paradigm, the analysis considers differences in occupational exposure at discrete points in time while simultaneously controlling for past exposure and employment status.[8] This analysis can be implemented with the G null test procedure. These four methods have not previously been compared by application to real data.

This paper directly compares these four methods with data from an occupational cohort exposed to arsenic.

These data provide a setting where exposure has a strong dose-response relation for mortality from respiratory cancer. Also, the dose-response curve has been found to have a non-linear shape in which arsenic exposure seems to exert a greater effect at lower levels of exposure than would be expected if potency was extrapolated linearly from higher levels of exposure.[23] The impact of these methods on both the magnitude and shape of the dose-response relation is evaluated and results are reviewed to highlight the strengths and weaknesses of each method.

Methods and materials

DATABASE

The occupational cohort exposed to arsenic used in this analysis has been previously described.[14][15] Briefly, this cohort consists of all white male workers employed for one year or more between 1940 and 1964 at a copper smelter in Tacoma, Washington. The primary occupational exposure was arsenic. Measures of exposure were primarily based on periodic samples of air arsenic, specifically, arsenic trioxide.[15] These measurements were confined to departments where exposure to arsenic was considered to be a problem. Follow up was complete to the end of 1976, encompassing 2802 workers, 73 000 person-years of observation, and 1061 deaths. Respiratory cancer accounted for 104 deaths and was strongly associated with arsenic exposure.[14][15]

STATISTICAL METHODS

Four methods to control the healthy worker survivor effect were applied to these data: (a) restricting the cohort to those surviving a minimum period, (b) lagging exposure, (c) inclusion of active v inactive work status as a covariate, (d) the G null test procedure, which considers the healthy worker survivor effect to be simultaneously a confounding and intermediate variable. Each of these techniques was first applied as an independent method to control the healthy worker survivor effect. Subsequently, lagging was used simultaneously with adjustment for work status and with the G null test procedure.

The restriction method of Fox and Collier[10] limits the analysis to survivors of a given period, under the assumption that among long term survivors, the healthy worker survivor effect is reduced or non-existent.[4] Initially, this method is similar to crude stratification of workers on duration of follow up, but ignores the strata with shorter follow up. Three minimum survival periods were evaluated in these analyses: 10, 15, and 20 years.

The method of lagging exposure[12] is identical in its implementation to the lagging used in adjusting for the latency between the time of exposure and manifestation of cancer mortality, although the motivation differs. To adjust for a long latency, exposures occurring shortly before the event are not relevant. To control the healthy worker survivor effect, recent exposures are ignored because exposures nearest the event could only have been acquired by those who survived the job, namely the healthiest employees. Three lagging periods were used: 10, 15, and 20 years.

A third method includes a covariate for current employment status at the plants of interest. The healthy worker survivor effect is considered to be a confounder in the traditional sense and is represented by employment status. Because one year survival rates for lung cancer are low, employment status was extended one year beyond the actual date of the end of employment to encompass those deaths that are related to disease incurred during employment but occur shortly after the end of employment.[16]

A fourth analytical strategy[8] takes into account the possibility that the healthy worker survivor effect operates simultaneously as a confounding and as an intermediate variable. Because both exposure and work status vary with time, and because each can influence the other, a different paradigm is needed. Essentially, this procedure examines differences in exposure at each point in time while simultaneously controlling for past exposure and employment. To implement the analysis, a nested, matched case-control strategy was used. Cases of respiratory cancer were matched to controls by age, date of hire, past exposure, and past employment. Exposure differences between cases and their matched controls were evaluated by the G null test procedure.[8] This method considers periods off work to be an analytically and conceptually distinct category.

The modified G null test is more powerful than the unmodified test but requires a more restrictive assumption on how long past exposures and employment status predict current exposures and employment status. For example, a modified G null test with m set to 2 assumes that the current job task (exposure concentration) will predict employment status or exposure only during the next two years. Exposure was lagged at 10, 15, and 20 years in combination with the G null test with and without a two year predictive period.

Both forms of the G null test were used; for

Figure 1 Cohort restriction.
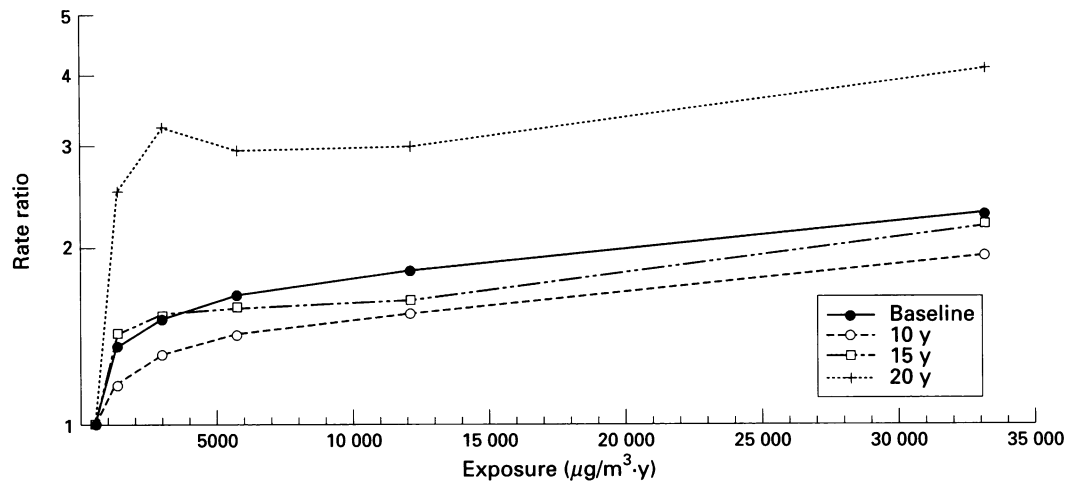
the modified G null test it was assumed that for two years past exposures and employment status were related to current exposures and employment status. A test of significance (χ^2 with one degree of freedom) was calculated by the G null test procedure with software provided by Drs James Robins and Donald Blevins. An odds ratio (OR) is provided to assist in the interpretation of the results but cannot be interpreted as the effect of exposure during an extended time interval relative to no or low exposure during that interval. In further G null test analyses, exposure was lagged at 10, 15, and 20 years. A test for trend provides evidence for the direction of an effect but not its magnitude. For notation purposes a "+" trend test corresponds to a positive coefficient indicating an adverse effect of arsenic exposure on mortality from respiratory cancer.

For the unadjusted analysis and the first three methods, mortality during follow up was modelled by Poisson regression, with the unit of observation being one person-year. The log of the mortality rate was modelled as a linear function of covariates, according to the formula:

$$\ln\,[\lambda(t)] \; = \; \alpha \; = \; \Sigma\beta_i x_i + \Sigma y_j x_j\ (t)$$

where $\lambda(t)$ represents the mortality rate per person-year; α represents the baseline rate in the population; β_i represents the coefficient for each of the corresponding time constant covariates; x_i and y_j, represent the coefficient for each corresponding time dependent covariate, $x_j(t)$.

All Poisson regression models were adjusted for current age and birth cohort. Some authors[17] have argued that occupational cohort mortality analyses need to control for the time since hire or time since start of follow up. Models adjusted for current age and birth cohort yielded essentially similar results as those models adjusted for age and time since hire and those models adjusted for age and time since start of follow up.[18] Before applying these methods, a baseline model was fitted, unadjusted for the healthy worker survivor effect—that is, this model did not restrict the cohort, lag exposure, or include covariates for work status. The SAS system of software was used to perform the analyses.[19] Person-years were calculated with an adaptation of the per-

son-time program described by Pearce and Checkoway.[20]

EXPOSURE METRICS

For all analyses except the G null test, exposure was modelled as cumulative exposure. With the same exposure cut off points as Enterline *et al*[15] this resulted in some sparse cells and hence instability in the multivariate analyses. For this reason, the two highest categories were collapsed so that exposure was categorised in $\mu g/m^3$.years with the following cut off points: < 750, < 2000, < 4000, < 8000,< 20000, and $\geqslant$ 20000. To allow maximum sensitivity to non-linearities, each exposure category was a separate covariate in the model, with the lowest category as the control. Thus, the dose-response shape is not imposed by the model. The median arsenic concentration in each category was used in all graphs and models.

A cumulative exposure measure is incompatible with the underlying theory of the G null test; for these analyses, exposure was categorised directly from the periodic industrial hygiene samples with the following cut off points for air arsenic concentration: < 200, < 400, < 1000, and $\geqslant$ 1000 $\mu g/m^3$. Exposure concentrations were estimated at six month intervals beginning with the date of hire. Periods off work were distinguished from any periods on work but with no exposure.

Results

The 104 respiratory cancer deaths were distributed as follows: 11 occurred among men aged less than 50 years, 23 among men aged 50 to 59 years, 46 among men aged 60 to 69 years, and 24 among men aged 70 or more years.

Figure 1 shows the results for restriction analyses. Figures 1 to 4 show the impact of different methods to adjust for the healthy worker survivor effect in graphs of rate ratios for respiratory cancer versus cumulative exposure to arsenic. Compared with the baseline analysis, there is little difference in the magnitude of the estimated rate ratios, with an apparent exception when the cohort is restricted to 20 year survivors. When the cohort is restricted to 10 or 15 year survivors,
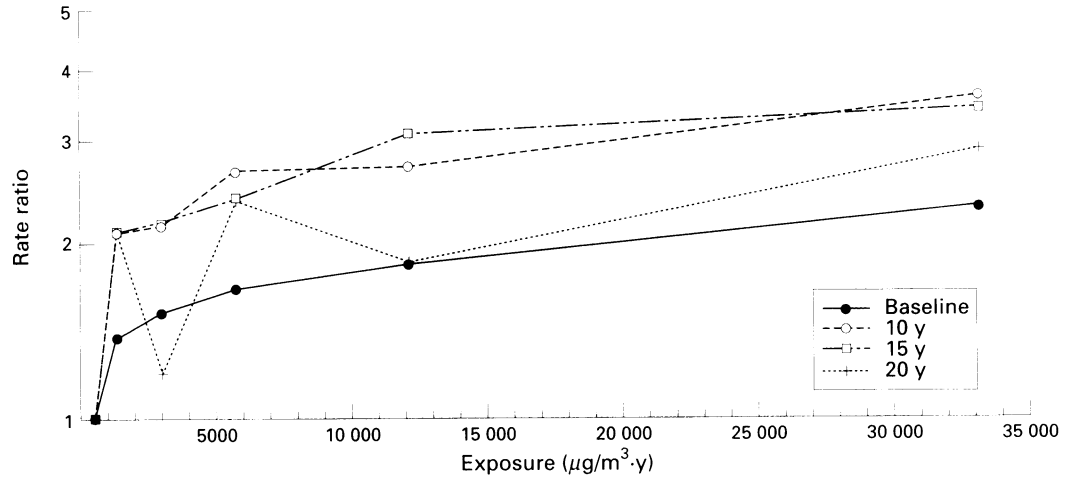
Figure 2   Exposure lagging.

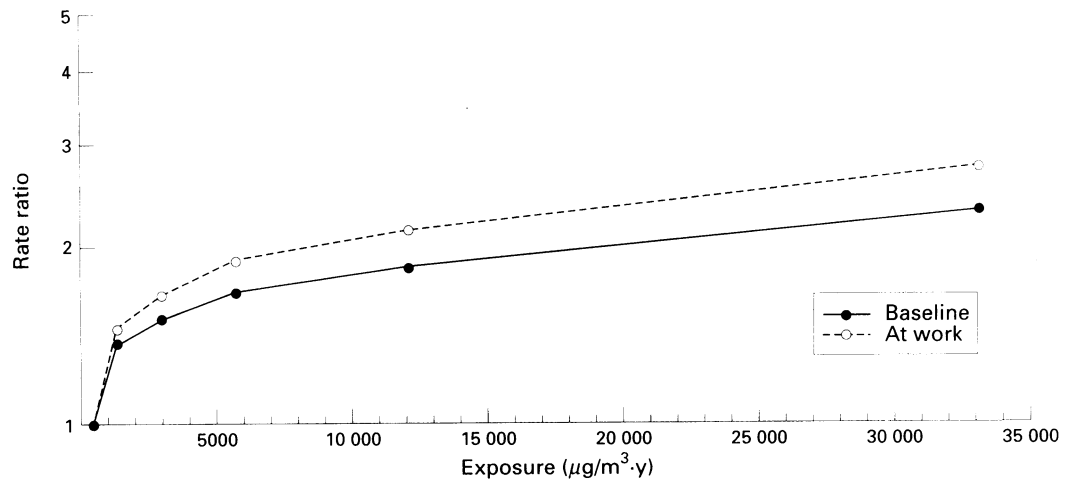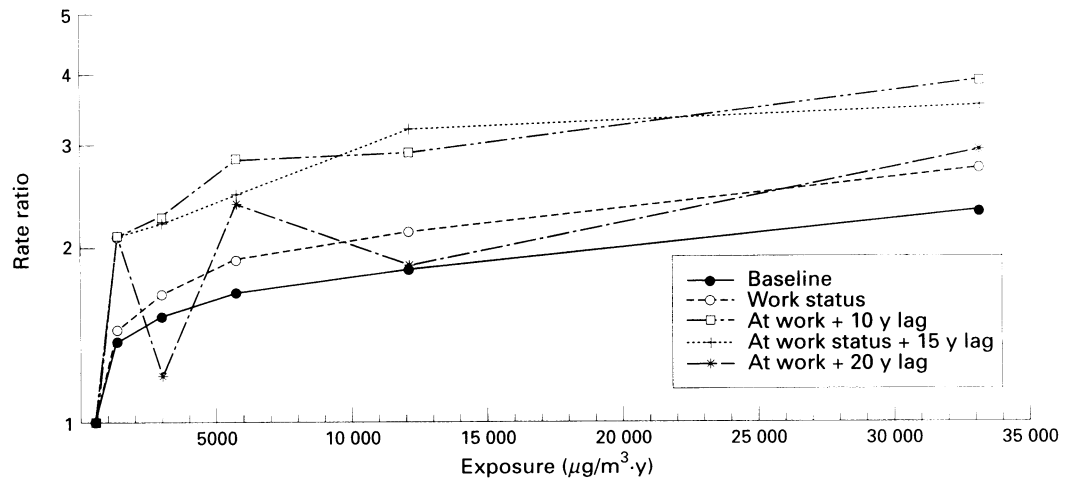*Figure 3   Work status covariate.*

Figure 4   Exposure lagging and work status covariate.

the point estimates for most exposure categories are slightly lower than for an unrestricted analysis.

Restricting the cohort to 20 year survivors reduces the number of cases by 16% and the number of person-years by 53%. Although the rate ratios among 20 year survivors are large, these are highly unstable due to the presence of only three cases in the lowest (control) exposure category. The overall shape of the dose-response curve is not altered by the restriction analyses.

Figure 2 shows the results when exposure is lagged. A lagged exposure generally produces higher rate ratios at all levels of exposure compared with the baseline (unlagged) analysis. The greatest effect estimates are found when exposure is lagged by 10 or 15 years. Regardless of lag, the shape of all the dose-response curves is similar. The non-linear, concave downward shape of the unlagged analysis is enhanced when exposure is lagged.

Figure 3 shows the dose-response curves for the models with and without the employment

Table 1   G null test procedure sample sizes

	Basic	Modified m = 2
Cases	36	77
Discordant pairs	142	241

Table 3   − 2 Log likelihoods for baseline, exposure lagging, and work status models

Model	− 2 Log likelihood	df
Baseline	3834·39	8
Exposure lag:		
10 y	3828·40	8
15 y	3827·05	8
20 y	3827·09	8
Work status	3832·18	9
Work status and exposure lag:		
10 y	3826·84	9
15 y	3825·95	9
20 y	3826·46	9

status covariate. Adjustment for work status did not alter the shape of the dose-response curve. However, the model that controls for work status shows slightly higher rate ratios at all levels of exposure. The rate ratios from the baseline analysis are 1·4, 1·5, 1·7, 1·8, and 2·3, whereas the rate ratios from the analysis with the work status covariates are 1·4, 1·7, 1·9, 2·1, and 2·7. In these models, current employment was a protective factor for mortality from respiratory cancer. In models adjusted for current age, birth cohort, and exposure, current employment had a rate ratio of 0·70 (95% confidence interval (95% CI) 0·44 to 1·1).

Figure 4 shows results when exposure is lagged by 15 years and the work status covariate is included. The point estimates for exposure are slightly greater at all exposures when the exposure is lagged and the work status covariate is included and compared with the corresponding analyses that only include a lagged exposure or the work status covariate. Similar results were found with lags of five and 10 years (not shown).

Table 1 presents the number of cases contributing to an unlagged G null analysis and the number of case-control divergences. There are 104 cases of respiratory cancer with 4006 matched controls—that is, meeting the matching criteria for date of birth and date of hire but not necessarily similar exposures. For the basic G null test, 142 pairs match on exposure and then diverge and hence contribute to the analysis; these pairs are from only 36 cases. The modified G null test (m = 2) improves statistical efficiency, as 77 cases contribute to the analysis with 241 divergent pairs.

Table 2 provides details from the G null test procedure: the associated χ^2 statistics (all with one degree of freedom), and the OR for both the unmodified G null test and the modified G null test. Both the overall tests (of the three highest point exposure categories v the lowest) and the tests for trend were significant in both the modified and unmodified analyses regardless of the lag period used. A 10 year lag gave the largest χ^2 statistics in both the modified and unmodified G null analyses, although the ORs were largest for comparing the highest three point exposures with the lowest (although only marginally so for the unmodi-

fied G null analysis) with a lag of 20 years.

Table 3 shows the − 2 log likelihoods from the Poisson regression models (figs 2–4). Log likelihoods for the models based on cohort restriction are not comparable with any of the other models because this method reduces the number of observations. Furthermore, they are not comparable with each other because the number of observations vary by the degree of restriction. Models that lag the exposure provide a substantial improvement in fit when contrasted with the baseline model. Although the number of variables in the model is the same as in the baseline model, the difference in the − 2 log likelihood statistics is similar regardless of the duration of the exposure lag. When the employment status variable is added to either the baseline model or to those with exposure lags, there is no significant difference in fit. Assessment of fit does not allow for the possible bias in estimating the exposure effects.

Discussion

The baseline analysis showed an exposure effect that rose monotonically but non-linearly with increasing levels of exposure. Methods intended to control for the healthy worker survivor effect tended to enhance the exposure effect, with the notable exception being the restriction analyses. One motivation for this investigation was to find whether the healthy worker survivor effect was responsible for the non-linear dose-response relation which has been found in numerous studies of arsenic and respiratory cancer.[3] The shape of the dose-response curves is strikingly similar across all the methods examined. We conclude that the healthy worker survivor effect does not explain the non-linear shape found in the unadjusted analysis.

Although all methods show an effect of exposure to arsenic on respiratory mortality, an indication of a weak effect of exposure is found in the effect measures at the lowest exposure cate-

Table 2   G null test results

	Unmodified G null test procedure								Modified G null test							
	No lag		10 year lag		15 year lag		20 year lag		No lag		10 year lag		15 year lag		20 year lag	
	χ^2	OR	χ^2	OR	χ^2	OR	χ^2	OR	χ^2	OR	χ^2	OR	χ^2	OR	χ^2	OR
Test for trend	11·4	+	12·2	+	9·79	+	8·47	+	11·9	+	13·3	+	7·02	+	9·57	+
Overall	9·25	2·49	9·34	2·58	7·01	2·42	5·27	2·60	5·32	1·59	8·92	2·09	6·56	2·01	5·76	2·36

χ^2 has 1 degree of freedom; + = adverse effect of arsenic exposure on lung cancer mortality; overall = test of exposure concentrations 2, 3, 4 v 1, with 1 as the lowest and 4 as the highest.

gories. A comparison of the second lowest with the lowest exposure category shows an adverse effect of exposure; however, these differences were not significant in any model. Thus, weak effects of exposure remain difficult to detect, regardless of method to control the healthy worker survivor effect.

Empirically, cohort restriction has the potential to result in lower rate ratios. By itself, this finding could be due to higher rate ratios in short term employees who are excluded by the restriction, or a bias introduced by the method—for example, through reduced variability in exposure. The method of cohort construction of Enterline et al[15] already required a one year survival at employment and hence could have attenuated the impact of this method if such short term workers had lower mortality than the rest of the cohort (note that they would have contributed data to the lowest cumulative exposure groups). However, the assumption of the restriction method is that the healthy worker survivor effect is eliminated or reduced when the cohort is sufficiently restricted.[8 10 11] Evidence is lacking on this assumption. Indeed, health differences may continue to be associated with absenteeism among long term survivors.

Cohort restriction has practical limitations. The available sample becomes less generalisable as it is reduced by a non-random mechanism. As the period of restriction increases, the precision of the estimates decreases due to a reduction in the number of cases and person-years in the analysis. This decreased precision was particularly striking when the cohort was restricted to 20 year survivors and the dose-response curve became erratic.

Lagged analyses resulted in larger risk estimates than the baseline analysis, with the largest exposure effects found at lags of 10 and 15 years. These results suggest that the healthy worker survivor effect was successfully controlled and arsenic exposure has a latency period of 10 to 15 years, or bias was introduced because of a violation of assumptions of this method. The main assumptions of lagging are that (a) the mean latency between exposure and the measured disease outcome is longer than the time the healthy worker survivor effect operates,[12] (b) time off work is equivalent to time on work at zero exposure, and (c) exposure levels are not related to health status, including susceptibility to the outcome under study. Note that if the latency is very short and the effect of exposure is reversible, lagging could cause an effect to be missed. For example, if the risk resulting from an exposure was increased in the first five years after exposure and then dropped to normal levels, analyses in which exposures were lagged by five years or more would miss the effect and would not be the method of choice regardless of considerations about the healthy worker survivor effect. In our analyses, support was strongest for a 10 or 15 year latency, although the model fit was similar regardless of the duration of the latency period. The assumption that time off work and time on work at zero exposure are equivalent could be a problem if people who leave work

either have more health problems or receive poorer health care.

Unlike the lagging method, adjustment for work status considers off job and on job experiences to be different. This distinction is incorporated even more strictly into the G null test. The methods that lag exposure or restrict the cohort do not make this separation. The importance of employment status as a potential confounder has long been recognised.[6 10 21] Our results support the contention that work status does confound the association between exposure and cancer in that all point estimates increased after adjustment for work status.

However, as work status is also an intermediate variable in the sense that it is influenced by past exposure, is related to disease, and itself affects future exposure, there is the potential that this approach may actually introduce confounding. For this reason, the conceptual contribution of Robins et al is of great importance. In particular, they showed that controlling for work status as a pure confounder can introduce bias if it is also an intermediate variable.[9 22] In fact, traditional thinking regarding confounders is now coming under closer scrutiny by others. Weinberg[23] presents several scenarios in which variables that are correlates of exposure should not be controlled as confounders, and confirms earlier work by Robins[8 9] that shows that (a) such adjustment introduces bias, and (b) even if not intermediate on a causal pathway, it may be incorrect to adjust for correlates of exposure if they have no independent causal relation to outcome.

The causal paradigm underlying the G method of Robins is based on the concept of a "counterfactual", defined as the outcome an exposed person would have experienced, had (counter to the fact) she or he not been exposed. As each person, at each point in time, can only receive one level of exposure, the only valid comparison is among those whose past exposure experience is identical, or as close to identical as possible. Hence the G methods examine those points in time when exposure diverges for workers whose past exposures and times off work are matched. Because past exposure predicts exposure at work, the history of employment status and past exposure influence the G methods. Standard methods of analysis are, at least theoretically, biased even with adjustment for employment status[8 9]; alternative methods[8 22] are necessary to provide an unbiased estimate.

Results from the analyses of the G null test procedure also showed an adverse effect of exposure. This procedure provides no direct information on the size of effect of long term exposure nor the shape of the dose-response curve, although several extensions of this test do estimate causal parameters.[22 24] (The ORs are summary ORs from a stratified analysis of odds of exposure at a single point in time conditional on no previous differences in exposures and work history, comparing cases with controls.) All of the G methods require that exposures at different times never be combined.[22] Hence, traditional measures of occupational exposure (cumulative, peak, intensity of expo-

sure, or years employed) are incompatible with this technique. For the unmodified G null test, once a case and its matched control differ in exposure history, subsequent exposures are not considered as these are influenced by differential survivorship. With the modified G null test, subsequent exposures can be considered if cases and their matched controls converge again; a single case-control pair may contribute information to more than one exposure contrast.

Although the cohort included over 100 respiratory cancer cases and 2800 employees, few of the G null analyses have more than 40 totally matched pairs. The G null test requires an exact match of previous exposure and employment between cases and controls. Employees with similar but not identical employment and exposure histories will not match. A case may match with multiple controls and contribute more than once to each of the test statistics, even in the unmodified test procedure. Thus, a test consisting of 76 matched pairs may have information contributed from fewer than 76 different cases with some cases contributing a disproportionate number of matches. Despite matching multiple controls per case—for example, 50 or more in our analyses—the requirement of identical exposure histories may produce an insufficient number of risk sets to test for an effect, a limitation originally noted by Robins.[8] Use of a subset of the cases may also limit the ability to generalise from any findings.

Partly in response to the limitations of the G null test, a more generalised and flexible method, the G estimation procedure, has been developed by Robins and colleagues[22] who used structural nested failure time models, a variant of the accelerated failure time model. G estimation provides: (*a*) a readily interpretable and meaningful measure of the exposure effect and (*b*) a flexible approach to matching of past exposure and employment. In particular, one obtains a point estimate and CI for the overall change in survival time, while simultaneously adjusting for past exposure and off work history along with other time dependent and fixed covariates. Additionally, structural nested failure time models allow past exposure to be of several forms including the traditional measures of exposure (cumulative, peak, or intensity). These models allow one to relax the requirement of identical exposure histories by modelling current exposure as a function of past exposure; thus, an expectation is that more cases will have suitable controls resulting in a statistically more efficient and powerful model relative to the G null test procedure. To date, application of these models to occupational cohorts has not been attempted. Further analyses of the cohort exposed to arsenic with structural nested failure time models and G estimation are planned.

Conclusion

Overall the shape of the dose-response curves was similar in the unadjusted Poisson regression and with the methods of cohort restriction, exposure lagging, and inclusion of the age specific work status covariates. At lags of 0, 10, 15, and 20 years, the G null test procedure detected an adverse effect of exposure.

Cohort restriction seems inadequate, with empirical evidence seeming to support theoretical deficiencies. It assumes that the healthy worker survivor effect ends after a fixed period of restriction. Both lagging of exposure and adjustment for work status empirically yielded a greater effect than unadjusted analyses, which could be interpreted to suggest confounding by the healthy worker survivor effect and some success in controlling such confounding from each method. These two methods, which seemed to enhance the differences between exposed and unexposed people, may provide more sensitive analyses of effects that are weak and therefore difficult to detect in epidemiological studies. Only lagging of exposure substantially improved the fit of the models.

However, these two methods both rely on assumptions that may be problematic. Lagging requires that the latency period be longer than the period over which exposure predicts employment history. Both disease latency and the time frame for associations between past exposure and subsequent work history are usually not known. Adjustment for work status assumes that this variable is a pure confounder—that is, that exposure is unrelated to subsequent work status. The underlying theory of the dual nature of the healthy worker survivor effect as a confounding and an intermediate variable needs to be seriously explored. The G null and G estimation approaches are unique in taking into account the relations between exposure, employment status, underlying risk of disease, and subsequent exposure history and employment status. Despite concerns about implementation of the G null technique, these methods represent a significant contribution to epidemiological theory.

1 Arrighi HM, Hertz-Picciotto I. The evolving concept of the healthy worker survivor effect. *Epidemiology* 1993;5:189–96.
2 Hertz-Picciotto I, Holtzman DA. Issues in conducting a cancer risk assessment using epidemiologic data: arsenic as a case study. *Exp Pathol* 1989;37:219–23.
3 Hertz-Picciotto I, Smith AH. Observations on the dose-response curve for arsenic exposure and lung cancer. *Scand J Work Environ Health* 1993;19:217–26.
4 Choi BCK. Definition, sources, magnitude, effect modifiers, and strategies of reduction of the healthy worker effect. *J Occup Med* 1992;34:979–88.
5 Blair A, Stewart P, O'Berg M, Gaffey W, Walrath J, Ward J, et al. Mortality among industrial workers exposed to formaldehyde. *J Natl Cancer Inst* 1986;76:1071–84.
6 Ott MG, Holder BB, Langner RR. Determinants of mortality in an industrial population. *J Occup Med* 1976;18:171–7.
7 Steenland K, Stayner L. The importance of employment status in occupational cohort mortality studies. *Epidemiol* 1991;2:418–21.
8 Robins J. A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Mathematical Modeling* 1986;7:1393–512.

9 Robins J. A graphical approach to the identification and estimation of causal parameters in mortality studies with sustained exposure periods. *J Chron Dis* 1987;40(suppl 2):139S-61S.

10 Fox AJ, Collier PF. Low mortality rates in industrial cohort studies due to selection for work and survival in the industry. *Br J Prev Soc Med* 1976;30:225-30.

11 Fox AJ, Collier PF. Mortality experience of workers exposed to vinyl chloride monomer in the manufacturer of polyvinyl chloride in Great Britain. *Br J Ind Med* 1977; 34:1-10.

12 Gilbert E. Some confounding factors in the study of mortality and occupational exposures. *Am J Epidemiol* 1982; 116:77-188.

13 Gilbert E, Marks S. An analysis of the mortality of workers in a nuclear facility. Rad Res 1979;79:122-48.

14 Enterline PE, Marsh GL. Cancer among workers exposed to arsenic and other substances in a copper smelter. *Am J Epidemiol* 1983;116:895-911.

15 Enterline PE, Henderson VL, Marsh GL. Exposure to arsenic and respiratory cancer: a reanalysis. *Am J Epidemiol* 1987;125:929-38.

16 Wing S, Shy CM, Wood JL, Wolf S, Cragle DL, Frome EL. Mortality among workers at Oak Ridge National Laboratory. *JAMA* 1991;265:1397-402.

17 Flanders WD, Cárdenas VM, Aust H. Time-since-hire in internal comparisons of cumulative exposure in occupa-

tional cohort studies. *Epidemiol* 1993;4:336-41.

18 Arrighi HM, Hertz-Picciotto I. Controlling for time-since-hire in occupational studies using internal comparisons and cumulative exposure. *Epidemiol* 1995;6:415-8.

19 SAS Institute. *SAS/STAT(TM) user's guide, version 6, volume 2.* Cary, NC: SAS Institute 1989.

20 Pearce N, Checkoway H. A simple computer program for generating person-time data in cohort studies using time-related factors. *Am J Epidemiol* 1987;125:1085-91.

21 Howe GR. Methodological issues in cohort studies: point estimators of the rate ratio. *Int J Epidemiol* 1986;15: 257-62.

22 Robins JM, Blevins D, Ritter G, Wufsohn M. G-estimation of the effect of prophylaxis therapy for Pneumocystis carini pneumonia on the survival of AIDS patients. *Epidemiol* 1992;3:19-336.

23 Weinberg C. Toward a clearer definition of confounding. *Am J Epidemiol* 1993;137:1-8.

24 Robins JM. The analysis of randomized and nonrandomized AIDS treatment trial using a new approach to causal inference in longitudinal studies. In Sechrest L, Freeman H, Mulley A (eds). *Health services research methodology: a focus on AIDS. (Conference proceedings).* Washington DC: National Center for Health Services Research and Health Care Technology, Public Health Service, US Department of Health and Human Services, 1989; 113-59.