



American Journal of EPIDEMIOLOGY

Volume 139

Number 8

April 15, 1994

Copyright © 1994 by The Johns Hopkins University

School of Hygiene and Public Health

Sponsored by the Society for Epidemiologic Research

REVIEWS AND COMMENTARY

Invited Commentary: Ecologic Studies—Biases, Misconceptions, and Counterexamples

Sander Greenland¹ and James Robins²

Many authors have pointed out that relative-risk estimates derived from ecologic data are vulnerable to biases not found in estimates derived from individual-level data. Nevertheless, biases in ecologic studies still are often dealt with in the same manner as biases in other observational studies, and so are not given adequate treatment. This commentary reviews and illustrates some of the more recent findings about bias in ecologic estimates. Special attention is given to problems of ecologic confounder control when individual risks follow a nonlinear model, and to misconceptions about ecologic bias that have appeared in the literature. *Am J Epidemiol* 1994;139:747–60.

bias (epidemiology); biometry; epidemiologic methods; models, statistical; risk; statistics

Editor's note: Two additional invited commentaries on this topic follow on pp. 761 and 765, and a response by Greenland and Robins appears on p. 769.

Although the problem of ecologic bias was described in the social science literature over 40 years ago (1), specialized epidemiologic discussions have become common only in the last decade (2–6). These discussions may have been inspired by the recognition that, under special assumptions, eco-

logic studies can supply estimates of individual-level relative risks (7).

As various discussants have pointed out (3–6, 8, 9), ecologic relative-risk estimates can be subject to biases not present in estimates from individual-level observational studies of the same populations (case-control and cohort studies). Unlike an individual-level study, an ecologic study does not link individual outcome events to individual exposure or covariate histories, nor does it link individual exposure and covariate histories to one another. It is these linkage failures that are the source of the special biases of ecologic studies (1, 2, 6). There are already a number of epidemiologic reviews that illustrate problems arising from this linkage failure (2, 4, 6, 9), and we do not want to replicate them here. Instead, we will focus on more recent findings, not

Received for publication April 24, 1992, and in final form December 3, 1993.

¹ Department of Epidemiology, UCLA School of Public Health, Los Angeles, CA.

² Departments of Epidemiology and Biostatistics, Harvard School of Public Health, Boston, MA.

Correspondence to Dr. Sander Greenland, 1720 Tuna Canyon Rd., Topanga, CA 90290.

illustrated in these reviews, of differences in the way certain methodological problems affect ecologic and individual-level results (3, 5, 8, 9). Our examples illustrate this point for effects of confounder misspecification, model-form misspecification, and confounder measurement error.

Similar to individual-level studies, ecologic studies can provide valid tests of the no-effect hypothesis under broader conditions than those required for valid estimation from the study (10). Nevertheless, our examples illustrate that this result may be of less utility than the corresponding individual-level result: Under departures from the null, a test in either type of study may indicate no association, or even one in the wrong direction, because of uncontrolled nonlinear effects; but, even when such nonlinearities may be easily detected and accounted for in individual-level data, they are often undetectable and impossible to control in the ecologic summary of the same data. These problems can be compounded by the paradoxical effects of measurement error and cross-level bias in ecologic studies (8).

Our presentation is not intended as a criticism of any particular study or as a general condemnation of ecologic studies. We do, however, suggest that the problems we discuss need to be considered when one critically evaluates ecologic results.

CONDITIONS FOR NO CONFOUNDING IN ECOLOGIC AND INDIVIDUAL-LEVEL STUDIES

Traditionally, an ecologic study of disease examines the relation between outcome means (such as disease rates) and exposure means across various populations. In most but not all examples, the populations are defined geographically, and so we will refer to the populations as regions. The algebraic view provided by Piantadosi et al. (4) and the geometric view provided by Walter (6) succinctly summarize earlier results, which focused on linear regression. It can be seen from either view that, under the simple linear-regression model, ecologic bias can-

not occur unless the exposure-specific disease rates vary across region. This point suggests that ecologic bias can be thought of as arising from cross-regional variation in risk-factor distributions, and this suggestion in turn has led to a view of ecologic bias as a problem analogous to the problem of confounding in individual-level studies (4, 6, 10). The following examples illustrate that, in applying this analogy, one must take care to distinguish ecologic and individual levels of confounding.

Example 1

Suppose we wish to test the hypothesis that environmental radon levels affect lung-cancer risk. In an ecologic study, each region has associated with it two key distributions relevant to this question: the regional distribution of radon levels and the regional distribution of background lung-cancer risks ("background risks" here means the risks that would be present if no environmental radon were present). Suppose that radon distributions have somehow been randomly assigned *across* regions and the number of regions is large. Apart from statistical variation, we may then deduce that there will be no ecologic association of radon distributions with lung-cancer rates across regions if there are no radon effects on either lung-cancer risk or the regional distributions of background risk factors. This conclusion holds even if within each region radon levels are associated with background risk, so that an individual-level study in any region would find a spurious association if control for background risk was not complete.

To make this example more concrete, suppose that radon has no effect, that age, sex, and smoking are the only important determinants of background risk, and that across regions no feature of radon distributions is associated with any feature of the joint age-sex-smoking distributions. The latter condition would guarantee absence of an ecologic association of radon and lung cancer. This would be true even if, within each region, older smokers tended to preferentially

live in low-radon dwellings; in this case, age and smoking effects would make radon appear protective at the individual (within-region) level, even though there would be no association at the ecologic level.

Example 2

Consider the converse of example 1. Suppose that, within each region, radon exposure has somehow been randomly assigned across individuals, and there are many individuals in each region. Apart from statistical variation, we may then deduce that there will be no individual-level association of radon levels with lung-cancer rates within regions if there are no radon effects. This conclusion holds even if, across regions, radon distributions are associated with the distributions of background risk across regions, so that an ecologic study would find a spurious association if control for the distribution of background risk was not complete.

To make this example more concrete, again suppose that radon has no effect, that age, sex, and smoking are the only important covariates, and that radon levels *within* regions are unassociated with age, sex, or smoking. The latter condition would guarantee absence of an individual-level association of radon and lung cancer. This would be true even if older smokers tended to preferentially live in regions with low radon levels; in this case, age and smoking effects on regional rates would make radon appear protective at the ecologic (cross-regional) level, even though there would be no association at the individual level.

The preceding examples illustrate that conditions that guarantee no confounding in an ecologic study are logically independent of the conditions that guarantee no confounding in an individual-level study of the same population. In practice, epidemiologists are (rightfully) loathe to assume either such condition in observational studies. When randomization is not a tenable hypothesis, attention turns to identification, measurement, and control of potential confounders—the determinants of background risk. Such control crucially hinges

on adequate measurement of the potential confounders.

EFFECTS OF NONLINEARITY AND NONADDITIVITY

In individual-level studies, a confounder measured without error can be fully controlled by using stratification or by using an approximately correct regression model, and the degree of confounding produced by a confounder can be estimated by comparing the estimates obtained before and after adjustment for a confounder. Nevertheless, a crude measurement or surrogate for a confounder may be inadequate to achieve full control. For example, adjustment for a smoking indicator (yes/no) may be inadequate to remove all confounding by smoking, even if the indicator is measured without error. To ensure full control, a more detailed summary of smoking history would be needed, and thus we would say that an analysis based on the smoking indicator alone would suffer from confounder misspecification. Similar comments would apply for any other variable with a complex, time-varying history, such as nutrient intake.

The situation is no different in principle for ecologic studies. Here, however, each confounder is much more complex than an individual history, and hence misspecification is even more difficult to avoid. For example, if in an individual-level study the main potential confounder is the smoking history of an individual, in the corresponding ecologic study the analogous potential confounder will be the distribution of *all* smoking histories across *all* individuals within each region. In typical ecologic studies, the summaries available for each region may be grossly inadequate to control confounding by the summarized covariate. This is especially true if nonlinearity or nonadditivity of effects ("effect modification") is present at the individual level, as one would expect in studies of epithelial cancers (such as lung cancer).

We will give several highly simplified and extreme examples that we hope will

make clear how bias due to ecologic confounder misspecification can arise from nonlinearity and nonadditivity of the underlying individual-level regression. These examples show that ecologic covariate summaries can be inadequate to detect and control confounding by a covariate with nonlinear and nonadditive effects, and that ecologic product terms can be inadequate to ameliorate this problem. In these examples, all measured variables are discrete and vary according to simple systematic rules. (Parallel examples can be given in which all measured variables are continuous and have random components (9), but such examples

require integral calculus for complete analysis.) For concreteness, we will suppose that our objective is to estimate the effect of radon levels in dwellings (measured in picocuries/liter (pCi/liter)) on lung-cancer rates. The most prominent confounders are age, smoking, and sex. For simplicity, we will eliminate age and sex from the examples by restricting our hypothetical study to males in a narrow age range. The examples illustrate the anomalies that can arise when the available ecologic summary of a confounder is inadequate for control of the confounder, and the confounder's effect is nonlinear.

Example 3

Suppose that our study data are limited to regional values of mean radon, mean smoking (in packs per day), and lung-cancer rates among males aged 70–74 years, for 41 regions indexed by $r = 0, \dots, 40$. Suppose also that

a) Radon levels are homogeneous within each region (so that the mean level is every resident's level) but vary across regions according to the rule

$$\text{regional radon level} = x_r = 0.1 + 0.3\sqrt{r},$$

ranging from 0.1 in region 0 to 2.0 in region 40.

b) Let u be a standard Gaussian (normal) random variable. The percents of men who smoke 0, 1, and 2 packs (0, 20, or 40 cigarettes) per day are

$$p_{0r} = 53 - 0.2r + 1.5u \text{ (expectation ranging from 45 to 53) and}$$

$$p_{1r} = 34 + 0.4r - u \text{ (expectation ranging from 34 to 50), so that}$$

$$p_{2r} = 100 - p_{0r} - p_{1r} = 13 - 0.2r - 0.5u \text{ (expectation ranging from 5 to 13).}$$

Here, smokers are fewer but tend to smoke more in low-numbered regions; these trends counterbalance each other, so that the expected mean cigarettes per day $E(20p_{1r} + 40p_{2r})/100$ shows no trend across region.

c) The lung-cancer rate per 10^5 person-years among individuals at radon level x smoking s cigarettes a day is

$$R(x, s) = 40(1 + 0.2x)\exp(0.1s), \quad (1)$$

so that radon has a linear "no-threshold" dose-response relation to lung cancer at each level of smoking, but also has a multiplicative interaction with smoking. Note that the rate ratio for 5 pCi/liter radon versus none is $1 + 0.2(5) = 2.0$, and the smoking rate ratio for one pack (20 cigarettes) per day versus none is $\exp[(0.1)20] = 7.4$.

Conditions (a) and (b) imply that radon and smoking levels will in expectation be nearly uncorrelated *across* regions. Conditions (a)–(c) imply that the lung-cancer rate in region r

will be

$$\begin{aligned} R_r &= [p_{0r}R(x_r, 0) + p_{1r}R(x_r, 20) + p_{2r}R(x_r, 40)]/100 \\ &= 0.40(1 + 0.2x_r)(p_{0r} + p_{1r}e^2 + p_{2r}e^4). \end{aligned} \quad (2)$$

A set of ecologic data generated under assumptions (a)–(c) is given in figures 1–3. Even though the lung-cancer rates R_r show the strong upward relation to smoking one would expect from model 1, and the ecologic correlation between radon and smoking is only 0.01, there is a significant negative ecologic association of radon with lung cancer rates. For these data, an ordinary ecologic linear regression of the regional rates R_r on radon level x_r and mean smoking level s_r yields the fitted line

$$\hat{R}_r = -1.5 - 39.8x_r + 40.8s_r,$$

which in turn yields a rate-ratio estimate for 5 pCi/liter of radon versus none (evaluated at the mean cross-regional smoking level, 12 cigarettes per day) of

$$\frac{-1.5 - 39.8(5) + 40.8(12)}{-1.5 + 40.8(12)} = 0.59.$$

Similarly, an ecologic log-linear regression yields

$$\hat{R}_r = \exp(5.07 - 0.093x_r + 0.094s_r),$$

which in turn yields rate-ratio estimates of $\exp[(-0.093)5] = 0.63$ for 5 pCi/liter radon versus none and $\exp[(0.094)20] = 6.6$ for one pack a day of cigarettes versus none. More generally, under assumptions (a)–(c), the expected radon rate-ratio estimates from the ecologic linear




FIGURE 1. Lung cancer rates plotted against regional radon levels, with least-squares line, example 3.




FIGURE 2. Lung cancer rates plotted against regional mean cigarette consumption, with least-squares line, example 3.

and log-linear regressions are 0.59 and 0.62, even though the expected smoking rate-ratio estimates from the same regressions are nearly unbiased, mean smoking has almost no expected ecologic correlation with radon, and smoking has no association with radon within regions (because radon is constant within regions).

Because radon and mean smoking are nearly uncorrelated in example 3, the results of the linear and log-linear regressions are unchanged if mean smoking is dropped from the regression. Thus, in both ecologic regressions, smoking does not appear to be a confounder. However, both regressions yield an inverse association of radon and lung cancer, despite the fact that radon is a positive risk factor in the underlying model used to generate the data, and the estimate of smoking effect is nearly unbiased. This occurs for two reasons: 1) the strong ecologic association of the radon and smoking distributions—an association that is not adequately summarized by the association of radon and mean smoking level; and 2) the nonlinear smoking effect. In essence, ecologic control of smoking, based as it is only on the smoking mean, is ineffective at removing the confounding by smoking. (Note that, in the example, there are no other sources of confounding: radon and smoking were assumed to be the only determinants of lung cancer within age and sex groups.)

More striking examples occur if radon and mean smoking are correlated across regions, even to a small degree. In such cases, smoking will appear to be a confounder, but it is possible for the smoking-adjusted estimate to be more biased than the unadjusted estimate.

Example 4

Let everything be as in example 3, except let the percent smoking 0 and 1 pack per day be

$$p_{0r} = 53 - 0.25r + 1.5u \text{ (expectations ranging from 43 to 53),}$$

$$p_{1r} = 34 + 0.45r - u \text{ (expectations ranging from 34 to 52).}$$

The expected ecologic correlation of radon and mean smoking will now be about 0.3. If smoking is omitted from the regressions, the expected linear and log-linear ecologic estimates of the rate-ratio for 5 pCi/liter radon versus none are 0.64 and 0.67. If, however, smoking is included in the regression, these expected radon estimates become 0.53 and 0.58, even though the smoking estimates are (as in example 3) nearly unbiased.

The additional bias produced by smoking adjustment in example 4 occurs because the crude estimate is already downwardly biased; since mean smoking is positively correlated with both radon and cancer, adjustment for this variable can only further diminish the radon rate ratio. If mean smoking had been negatively correlated with radon, adjustment for it would have increased the radon rate ratio and thus lessened the bias.

Thus far, our examples have concerned point estimation. Their implications for interval estimation are straightforward, however: Since most confidence intervals are arithmetically or geometrically centered on the point estimate, a bias in the point estimate implies a corresponding bias in the confidence interval. The following example shows that the problems can also afflict ecologic tests of the null hypothesis.




FIGURE 3. Scatterplot of regional radon levels and regional mean cigarette consumption, with least-squares line, example 3.

Example 5

Consider again example 3. If we change this example by setting the radon coefficient in equation 1 to 0 instead of 0.2, the true radon rate-ratio will be 1, but the expected estimates (at 5 pCi/liter) from the linear and log-linear ecologic regressions will be 0.06 and 0.28, whether or not smoking is controlled. This implies that the statistic for testing the null hypothesis will be centered on a value far below 0, and hence will be invalid.

It has been pointed out (11) that use of product ("interaction") terms in ecologic regressions can adjust for bias due to nonadditivity in certain special cases. The above examples show that product terms need not always help, however: If the product of radon and mean smoking is added to the linear and log-linear regressions, the expected radon rate ratios evaluated at the mean cross-regional smoking level would be little changed. Thus, in examples 3–5, the problematic nonlinear effects are not controlled by adding a product term to the ecologic regression.

One may legitimately ask if the same problems would occur if, instead of mean smoking levels, we had some other summary of regional smoking levels, such as percent smoking (i.e., $p_{1r} + p_{2r}$). In examples 3–5, this percent is highly correlated with both radon and lung cancer, so that adjustment for percent smoking would further reduce the estimates. In fact, if percent smoking (instead of mean smoking) is added to the linear and log-linear regressions, the expected radon rate ratios become 0.12 and 0.37 in example 3, 0.21 and 0.36 in example 4, and -0.19 and 0.17 in example 5. (It is a peculiar property of the linear estimate that it can take on negative values.) If the product of radon and percent smoking is also added to these regressions, the ratios become 0.12 and 0.30 in example 3, 0.07 and 0.27 in example 4, and -0.25 and 0.14 in example 5, slightly *worse* than without the product term. Thus, in all three examples, the bias is worse using percent smoking instead of mean smoking.

REQUIREMENTS FOR ECOLOGIC CONTROL OF CONFOUNDING

We have given examples in which smoking could not be adequately controlled by either of the usual ecologic summaries of smoking distributions. What if we had information on both mean and percent smoking, and we entered them both in the ecologic regressions? Here, at last, we would see reasonable ecologic estimates: In example 3, the expected linear and log-linear radon rate ratios would become 2.9 and 3.6, and in example 4, these ratios would become 2.8 and 3.5; these figures are on the correct side of the null and certainly represent dramatic improvements over the figures obtained controlling either mean smoking or percent smoking alone (although they remain biased because the ecologic regressions are misspecified and cannot fully control for the interaction of radon and smoking (cf. references 3, 9)).

In example 5, the ratios would become 1.0 and 1.4 on control of both smoking variables. In fact, the expected linear-regression radon coefficient would be zero, illustrating that a valid ecologic test of the null can be constructed from the linear model if one has sufficiently detailed information on exposure and covariates and all misclassification is nondifferential (10). This result is somewhat analogous to a corresponding result for individual-level studies, although, for individual-level studies, the result holds for both linear and log-linear models.

It appears, then, that one answer to the bias illustrated in examples 3–5 is to control for more than just a single summary of the covariate distribution (11). More specifically, suppose a covariate has K possible levels. Then, one may summarize regional covariate distributions with a set of K variables p_{1r}, \dots, p_{Kr} , representing the proportion of each region at each covariate level (i.e., the marginal distribution of the covariate). If there are no exposure effects, regional effects, or other risk factors, the rate at level k of the covariate will be a constant β_k across regions, and so the rate

in region r will be

$$R_r = \sum_k \beta_k p_{kr}, \quad (3)$$

which is a linear regression model. Thus, a valid test for exposure effect may be obtained by testing whether an exposure term significantly improves the fit when added to model 3. An equivalent test will be obtained if the set of proportions (the p_{kr}) are replaced by a full-rank linear transform of the set. In examples 3–5, the mean and proportion smoking and the constant term are such a transform, and there are no confounders besides smoking; thus, one obtains a valid test for radon effect if one controls for mean and proportion smoking in a linear model with an intercept term.

Unfortunately, equation 3 is less useful than it might at first appear. In any realistic example, most covariates (such as smoking) will have many levels, but only one or perhaps two ecologic summaries will be available. Furthermore, nonlinearity and nonadditivity of effects among multiple covariates can distort ecologic estimates in ways that cannot be controlled simply by using summaries of the marginal distributions of each covariate (9, 11). In such cases, the within-region *joint* distribution of the covariates (e.g., the joint smoking-age-sex distribution) may be needed to fully control ecologic confounding (11). Data on joint covariate distributions are rarely if ever available in ecologic studies. In contrast, the analogous individual-level study would directly observe the joint covariate distribution.

Despite these limitations, one may considerably improve ecologic results if multiple summary statistics are available on each covariate. There is, however, one further caution that should be recognized when adding such summaries to an ecologic regression: If, as is usually the case, the outcome variable has been standardized, then the covariates need to be standardized using the same standard distribution as used for the outcome. Otherwise, the addition of the covariates to the regression may not help, and may even worsen bias. Furthermore, the exposure must be standardized using the

same standard distribution if bias is to be avoided. In general, if some but not all the variables in a regression have been standardized, or if different variables have been standardized to different distributions, severe bias may result (12). Standardization bias may be especially common in ecologic analyses, because disease rates are usually age-standardized, whereas ecologic exposures and covariates (such as radon, smoking, income, etc.) are usually not (e.g., as in reference 13).

MEASUREMENT ERROR

Apart from basic demographic variables, most covariates used in ecologic regressions represent rather crude measurements. One relatively recent finding is that measurement error can have profoundly different consequences for ecologic and individual-level studies; for example, nondifferential misclassification of a dichotomous exposure will usually produce bias away from the null in ecologic studies, exactly the opposite of its usual effect in individual-level studies (8). As another example, nondifferential misclassification of a dichotomous confounder may have little effect on ecologic control of the confounder (which, however, may have been poor even without misclassification) (14).

For variables with multiple levels, the effects of measurement error on ecologic estimates are less predictable, in part because these effects can interact with confounding and nonadditivity to produce various anomalies.

Example 6

Consider again examples 3–5. Suppose that, in each case, 10 percent of all smokers misreported themselves as nonsmokers, and an additional 10 percent of two-pack smokers misreported themselves as one-pack smokers, regardless of radon or disease status. The observed mean smoking levels would now be $18p_{1r} + 34p_{2r}$ cigarettes per day. With mean smoking in the regressions, the expected radon rate-ratio estimates from

the linear and log-linear regressions would now be 0.54 and 0.59 in example 3, 0.49 and 0.55 in example 4, and 0.03 and 0.27 in example 5. In each case, the expected estimates after misclassification are slightly lower and thus more biased than before, and the additional bias produced by the misclassification is away from the null, even though the misclassification is nondifferential. Note also that in example 4, in which control for mean smoking does nothing if smoking is not misclassified, control for mean smoking after misclassification produces a more biased estimate than no control. If both mean smoking and proportion smoking are entered in the regression, the slight downward bias in the radon estimate produced by smoking misclassification is toward the null, and hence toward the true radon rate ratio.

The potential consequences of misclassification can easily be explored via sensitivity analysis of the ecologic regression (8, 15); such analyses illustrate that the linear rate-ratio estimate can be extraordinarily sensitive to even small degrees of misclassification.

CROSS-LEVEL BIAS

One important breakdown in the analogy between ecologic-bias and individual-level confounding occurs because in etiologic research, the target of inference for both ecologic and individual-level studies is the same: Both study effects at the individual level. For an individual-level study, these target effects are at the same level as the units of analysis. But, for an ecologic study, these target effects are at a finer level than the units of analysis. As a result, an ecologic study can be unbiased for ecologic effects (in particular, ecologic confounding may be absent) and yet still be biased for individual-level effects. Many of the classic social science examples of ecologic bias are of this form.

Example 7

A study of 19th-century suicide patterns in Prussian provinces, cited by Morgenstern

(2), found that suicide rates increased monotonically with the percent of province that was Protestant (as opposed to Catholic). This result might have been ecologically unconfounded, in the sense that percent Protestant might have been unassociated with any other ecologic risk factor, so that the observed association represented a true effect of percent Protestant on provincial suicide rates. It does *not*, however, follow that being Protestant was a risk factor for suicide in individuals. In fact, the opposite could have been true: The increased suicide rate in predominantly Protestant provinces could have been due to an increase in suicides among Catholics when the latter were subjected to the pressures of minority status in predominantly Protestant provinces. In this case, the variable "percent Protestant" would have been responsible for the excess suicide rate via an individual-level interaction with the individual-level religion variable, and being Protestant would actually have been protective against suicide among individuals.

Note that the phenomenon just described would have revealed its presence or absence if any province free of the true individual-level risk factor in the interaction (Catholicism) had been observed: If Catholics were indeed the source of the excess suicides in the predominantly Protestant provinces, in the absence of ecologic confounding a completely Protestant province should have exhibited a lower suicide rate than the predominantly Protestant provinces. Conversely, if Protestantism was the individual-level risk factor responsible for the observed ecologic trend, a completely Protestant province should have had the highest rate of all.

The bias phenomenon in example 7 may be ascribed to the fact that an ecologic (regional) variable (percent Protestant) had effects on individual risk, in addition to effects of the corresponding individual-level variable (religion) that the ecologic variable summarized. In both social and infectious-disease epidemiology, as well as in community intervention studies, ecologic effects

may be of direct interest; a classic example is the phenomenon of “herd immunity,” in which the overall prevalence of immunity in a region, as well as individual immune status, determines the risk of individuals within the region (16). For noncontagious diseases, however, ecologic effects may not be of direct interest, and may in fact obscure the individual-level effect of interest.

Example 8

Suppose we obtained data on mean airborne asbestos levels, mean smoking levels, and lung-cancer rates at a large number of industrial facilities that used asbestos-containing materials, and suppose we wished to use these data to estimate the excess lung-cancer caseload due to this industrial asbestos exposure. (Here, the ecologic “regions” are the facilities.) A regression of the rates on asbestos and smoking means could produce a severely attenuated asbestos-lung cancer association under the following scenario: Suppose that within each facility there was a wide range of exposure levels according to specific job assignment. Employers at facilities with extensive asbestos use had become aware of their potential liability and the potential interaction of asbestos and smoking at a fairly early time, and had selectively reassigned smokers to low-exposure jobs; in contrast, no selective reassignment was done at facilities with low exposure. The employer actions at high-exposure facilities would have prevented a large number of cases that would have been produced by asbestos-smoking interactions, and so would have greatly reduced the lung-cancer rates at the high-exposure facilities from what they would have been otherwise. The asbestos-lung cancer dose-response from the ecologic regression would be reduced (and in theory could even be reversed) by such a phenomenon. This would be true even if asbestos followed the same linear “no-threshold” model that radon followed in example 3.

The attenuated dose-response in example 8 reflects a “true” ecologic effect of the ecologic asbestos variable on the region-

factory-specific disease rates: High asbestos levels led to employer intervention, which led to lower lung-cancer rates than would have occurred without the intervention. But this ecologic dose-response understates the effect of high individual-level asbestos exposure on individual risks. This downward bias of the ecologic estimate for the individual-level effect is *not*, however, a form of ecologic confounding: The bias would occur with full magnitude even if smoking distributions were identical in all the study facilities, so that the asbestos and smoking marginal distributions were ecologically unassociated.

On the other hand, the bias in example 8 does depend on the asbestos-smoking association and interaction on the individual level. As in example 7, this bias could have revealed its presence if some homogeneously exposed facilities were observed and those facilities were identified as such. In such facilities, no internal association of asbestos and smoking would be possible, and so some sort of ecologic association of asbestos and smoking would be necessary for smoking to distort the ecologic association of asbestos and lung cancer among those facilities.

We refer to an incorrect extrapolation from unbiased estimates of ecologic effects to unobserved individual-level effects as “cross-level bias” (although the latter term is sometimes used to refer to any ecologic bias). As the examples suggest, it is possible to detect cross-level bias if one can identify and observe homogeneously exposed regions. In the extreme, ecologic bias in comparisons limited to homogeneously exposed regions (as in examples 3–6) can be viewed as purely an issue of confounding, albeit with special complexities of measurement and control of confounders. If exposure has no effect on any individual, then there will be no individual-level or ecologic effects, and so cross-level bias cannot occur. Thus, cross-level bias will not affect the validity of an ecologic test of the null hypothesis, although it still must be considered in interpreting a significant result.

More generally, one may show that cross-level bias will not occur if the individual-level effects of all variables (including unmeasured background factors) follow a multiple linear-regression model with no regional effects and no interactions (17). Nevertheless, given the usual inappropriateness of the multiple-linear model and the absence of homogeneously exposed regions, the possibility of cross-level bias adds a dimension to ecologic bias beyond that of simple confounding.

MISCONCEPTIONS IN ECOLOGIC REGRESSION

Before closing, we would like to point out a few mathematical errors that have appeared in the literature on ecologic bias.

Most of the literature on ecologic biases has been based on multiple-linear regression models, and much of the recent controversy concerning ecologic studies revolves around whether the conclusions based on such models are trustworthy (5, 11, 18–22). In individual-level studies, it is generally recognized that linear-regression models are not well suited to modeling rates and do not conform well to most cancer and cardiovascular data sets. One invalid rationale for their use in ecologic studies is based on the notion that a linear model should approximate the true model because of Taylor's theorem (18). The chief flaw in this rationale is that adequacy of a linear Taylor approximation to a regression depends on nonlinear and nonadditive effects being small over the range of the regressor variables. Such a condition is not likely to be satisfied in most ecologic regressions, such as those involving age and cancer. Other mathematical flaws in this type of rationale have been detailed elsewhere (23).

More relevant to our present discussion is that inadequacy of a multiple-linear approximation to the ecologic regression is *not* the source of the biases we discuss; rather, it is nonlinearities in the individual-level regression that can lead to ecologic bias, *even*

if the ecologic regression appears to be approximately linear, as in examples 3–5. In typical applications, large nonlinearities will be present in the individual-level regression; for example, age and smoking have enormous nonlinear and nonadditive effects on lung cancer (24).

Another, related, misconception is that important departures from linearity in the individual-level model will be detected by a test of fit of the ecologic linear model. But again, the ecologic relations may appear linear even when the individual-level model generating the data is highly nonlinear. Examples 3–5 and figures 1 and 2 illustrate this point: The R^2 for each of the linear models that include smoking exceed 0.9, despite the fact that the individual-level model is far from a multiple-linear form. In contrast, a test of fit of the linear model for individual-level data following (say) model 1 would almost certainly reject the linear model if the sample size was at all adequate.

Another misconception is that use of a large number of regions will somehow insure a random cross-regional relation between exposure and the covariates (18). In reality, large numbers do not insure randomness of exposure and covariate distributions in observational studies, whether ecologic or individual-level, nor do they insure that such biases as might occur will cancel each other out. Furthermore, uncorrelatedness of particular distributional summaries (such as means) does not guarantee that such uncorrelatedness holds for all important aspects of the distributions; example 3 is an illustration of this point.

Finally, it is sometimes assumed that for ecologic bias to occur, region itself must be a confounder on the individual level after other factors are controlled. Examples 3–5 provide counterexamples: In each case, there is no term for region in the individual-level rate model (model 1), and so region has no effect at the individual level. Hence, an individual-level study could ignore region in the analysis without introducing any bias.

CONCLUSION

Covariate control in ecologic studies requires attention to details not ordinarily of concern in individual-level studies. When, as is usually the case, important nonlinearity or nonadditivity can be expected among exposure and covariate effects, it may be necessary to obtain and control for multiple summaries of joint covariate distributions in order to insure that control is adequate. Unfortunately, for many potential confounders (such as diet and life-style), such summaries are unlikely to be available, and those marginal summaries that are available may be too crude to provide effective control. For covariates for which summaries are available, one must be alert to avoid improper application of control methods (e.g., improper combinations of standardization and regression). If important nonlinearity or nonadditivity is present and one cannot observe homogeneously exposed regions, one may need to consider the possibility of cross-level bias. One will also have to consider the potentially paradoxical effects of misclassification on ecologic regression results (8). Finally, although not discussed here, ecologic studies of open populations will also need to consider potential biases due to migration (25).

As with other observational studies, ecologic studies can give useful results if biases such as those discussed here can be ruled out or quantified. Nevertheless, bias evaluation can be especially difficult in ecologic studies of geographic regions because of the many potentially interacting covariates that may differ across regions. When biases cannot be ruled out with available data, further exploration will require individual-level studies.

ACKNOWLEDGMENTS

Dr. Greenland was supported by the UCLA Center for Occupational and Environmental Health. Dr. Robins was supported by grant nos. K04-ES00180, 5-P30-ES00002, R01-A132475,

and R01-ES03405 from the National Institutes of Health.

The authors thank Drs. Hermann Brenner, Irva Hertz-Picciotto, Jennifer Kelsey, and Hal Morgenstern for their helpful comments on the manuscript.

REFERENCES

1. Robinson WS. Ecological correlations and the behavior of individuals. *Am Sociol Rev* 1950; 15:351-7.
2. Morgenstern H. Uses of ecologic analysis in epidemiologic research. *Am J Public Health* 1982; 72:1336-44.
3. Richardson S, Stücker I, Hémon D. Comparison of relative risks obtained in ecological and individual studies: some methodological considerations. *Int J Epidemiol* 1987;16:111-20.
4. Piantadosi S, Byar DP, Green SB. The ecological fallacy. *Am J Epidemiol* 1988;127:893-904.
5. Greenland S, Morgenstern H. Ecological bias, confounding, and effect modification. *Int J Epidemiol* 1989;18:269-74.
6. Walter SD. The ecologic method in the study of environmental health. II. Methodologic issues and feasibility. *Environ Health Perspect* 1991; 94:67-73.
7. Beral V, Chilvers C, Fraser P. On the estimation of relative risk from vital statistical data. *J Epidemiol Community Health* 1979;33:159-62.
8. Brenner H, Savitz DA, Jöckel K-H, et al. Effects of nondifferential exposure misclassification in ecologic studies. *Am J Epidemiol* 1992;135: 85-95.
9. Greenland S. Divergent biases in ecologic and individual-level studies. *Stat Med* 1992;11: 1209-23.
10. Robins JM. When are ecological studies preferable to individual studies? Presented at the 21st Annual Meeting of the Society for Epidemiologic Research, Vancouver, British Columbia, Canada, June 15-17, 1988.
11. Richardson S, Hémon D. Ecological bias and confounding. (Letter). *Int J Epidemiol* 1990;19: 764-6.
12. Rosenbaum PR, Rubin DB. Difficulties with regression analyses of age-adjusted rates. *Biometrics* 1984;40:437-43.
13. Cohen BL. A test of the linear-no threshold theory of radiation carcinogenesis. *Environ Res* 1990;53:193-220.
14. Brenner H, Greenland S, Savitz DA. The effects of nondifferential confounder misclassification in ecologic studies. *Epidemiology* 1992;3:456-9.
15. Greenland S, Brenner H. Correcting for nondifferential misclassification in ecologic analyses. *Appl Stat* 1993;42:117-26.
16. Halloran ME, Haber M, Longini IM Jr, et al. Direct and indirect effects in vaccine efficacy and effectiveness. *Am J Epidemiol* 1991;133: 323-31.

17. Firebaugh G. A rule for inferring individual-level relationships from aggregate data. *Am Sociol Rev* 1978;43:557-72.
18. Cohen BL. Ecological versus case-control studies for testing a linear-no threshold dose-response relationship. *Int J Epidemiol* 1990;19:680-4.
19. Greenland S, Morgenstern H. Neither within-region nor cross-regional independence of exposure and covariates prevents ecological bias. (Letter). *Int J Epidemiol* 1991;20:816-17.
20. Richardson S, Hémon D. Ecological bias and linear dose response relationship. (Letter). *Int J Epidemiol* 1991;20:817-18.
21. Cohen BL. Problems in ecological studies. (Letter). *Int J Epidemiol* 1992;21:422-4.
22. Greenland S, Morgenstern H. From Sander Greenland and Hal Morgenstern. (Letter). *Int J Epidemiol* 1992;21:424-5.
23. White H. Maximum likelihood estimation of misspecified models. *Econometrica* 1982;50:1-9.
24. Doll R, Peto R. Cigarette smoking and bronchial carcinoma: dose and time relationships among regular smokers and lifelong non-smokers. *J Epidemiol Community Health* 1978;32:303-13.
25. Polissar L. The effect of migration on comparison of disease rates in geographic studies in the United States. *Am J Epidemiol* 1980;111:175-82.