



Practice of Epidemiology

Observed and Expected Mortality in Cohort Studies

David B. Richardson*, Alexander P. Keil, Stephen R. Cole, and Richard F. MacLehose

* Correspondence to Dr. David B. Richardson, Department of Epidemiology, School of Public Health, University of North Carolina at Chapel Hill, Chapel Hill, NC 27599 (e-mail: david.richardson@unc.edu).

Initially submitted March 3, 2016; accepted for publication July 27, 2016.

Epidemiologists often compare the observed number of deaths in a cohort with the expected number of deaths, obtained by multiplying person-time accrued in the cohort by mortality rates for a reference population (ideally, a reference that represents the mortality rate in the cohort in the absence of exposure). However, if exposure is hazardous (or salutary), this calculation will not consistently estimate the number of deaths expected in the absence of exposure because exposure will have affected the distribution of person-time observed in the study cohort. While problems with interpretation of this standard calculation of expected counts were discussed more than 2 decades ago, these discussions had little impact on epidemiologic practice. The logic of counterfactuals may help clarify this topic as we revisit these issues. In this paper, we describe a simple way to consistently estimate the expected number of deaths in such settings, and we illustrate the approach using data from a cohort study of mortality among underground miners.

cohort studies; mortality; standardized mortality ratio; statistics

Abbreviations: CI, confidence interval; CMR, causal mortality ratio; CRR, causal rate ratio; SMR, standardized mortality ratio.

In cohort mortality studies, epidemiologists often compare the observed number of deaths in the cohort with the expected number of deaths, obtained by multiplying person-time accrued in the cohort by the mortality rate in a reference population, ideally representing the rate of death in the absence of exposure. When the expected number is computed in this manner, while employing standardization to account for covariates, the observed:expected ratio is called a standardized mortality ratio (SMR). Standardization is a commonly used tool in observational research for estimation of causal contrasts, by which we mean comparisons of events in a single population under 2 different exposure conditions (1).

The results of SMR calculations allow for an appealing presentation: The observed and expected numbers of deaths are reported alongside each other. This presentation dates back to classical papers by Case et al. (2–4), who described the approach as a comparison between the mortality observed in the study cohort and a standard (“expected”) mortality based on the male population of England and Wales. Case and Lea wrote, “‘Expected’ in this context means the number that would have occurred had the mortality that affects the

general population, defined as the male population of England and Wales, acted with the same severity in the series studied” (2, p. 64). This interpretation has persisted in much of the more contemporary epidemiologic literature (5–8). Checkoway et al., for example, write that the expected number “can be thought of as the number of cases that would have occurred in the cohort had the cohort experienced the same stratum-specific rates as the reference population during the specified time interval of the study” (8, p. 144).

Such descriptions suggest an intuitive interpretation of the expected numbers of deaths routinely reported in cohort mortality studies. The contrast between the observed number of deaths and the expected value is often interpreted as the effect of exposure among the exposed (9), if in the absence of exposure the cohort members would have experienced the reference mortality rates.

Unfortunately, such an interpretation of the standard SMR calculation is generally inadvisable. The product of the person-time and (ideal) reference rates used to compute the expected number of deaths may differ substantially from the true expected number of deaths in the study cohort in the absence of exposure. This is because hazardous (or salutary) exposure

affects not only the number of deaths but also the person-time distribution observed in the study cohort, undermining the interpretability of this widely used quantity. These issues were raised in the statistical literature more than 2 decades ago by Hartz et al. (10), Berry (11), and Keiding and Vaeth (12) and were noted by Rothman et al. (13) but have had little impact on epidemiologic practice as it relates to interpretation of the standard SMR calculation. One reason may have been the lack of clarity concerning the estimand of interest in such calculations, whether expected counts, mortality risks, or relative death rates. The logic of counterfactuals may help as we revisit these issues. In this paper, we describe an approach with which to calculate the expected number of deaths in the absence of exposure, discuss conditions for a causal interpretation of such estimates, and illustrate the approach using data from a cohort mortality study of underground miners.

METHODS

Consider a study in which deaths have been ascertained without loss to follow-up for a closed occupational cohort of n men. Define study entry as time 0 and a potential duration of study follow-up of T years. Define D as the time of death (possibly occurring after time T , in which case D is unobserved). We use the subscript i to denote the values of variables for cohort member i .

Expected number of deaths in a homogenous cohort with a constant hazard rate

Suppose we denote the constant hazard rate as h^a , where $a = 1$ indicates the occupational cohort of interest and $a = 0$ indicates a reference mortality rate (e.g., from a region or nation). We use the superscripts 1 and 0 to denote exposed and (reference) unexposed persons, respectively, throughout the paper. While a constant hazard may not be broadly applicable, we use this simple situation to demonstrate the potential problem with SMRs. Furthermore, a constant hazard rate might be assumed if, for example, the cohort of men were aged 30 years at study entry in 1960 and $T = 5$ years (i.e., assuming a constant hazard rate for males in a 5-year interval of age and calendar time).

Let Y_i^a denote an indicator of death for subject i given exposure status a ; the indicator equals 0 for a survivor ($D_i^a > T$) and 1 for a decedent ($D_i^a < T$). The expected value of Y_i^a for a person with the hazard rate h^a over the span of T years is

$$1 - \exp(-h^a T).$$

Letting O denote the total number of deaths in a cohort of size n given occupational exposure ($a = 1$), and assuming independence, the expected value of O is

$$\sum_{i=1}^n [1 - \exp(-h^1 T)],$$

while the expected number of deaths in a cohort of size n in the absence of occupational exposure ($a = 0$), denoted E , is

$$\sum_{i=1}^n [1 - \exp(-h^0 T)].$$

Letting P^1 denote the person-time in the study cohort and noting that $h^1 = (O/P^1)$, the total expected person-time in the study cohort, given occupational exposure, is

$$\frac{\sum_{i=1}^n 1 - \exp(-h^1 T)}{h^1};$$

similarly, the expected person-time in the study cohort in the absence of occupational exposure, denoted P^0 , is

$$\frac{\sum_{i=1}^n 1 - \exp(-h^0 T)}{h^0}.$$

Using these quantities, we can define several contrasts of potential interest. For example, suppose we are interested in the comparison between the observed number of deaths in the cohort over the study period and the expected number of deaths in the cohort in the absence of exposure. We may describe this using the causal mortality ratio (CMR), which we define as the ratio of the number of observed deaths in the cohort to the number that would be expected in the same group of individuals if the cohort members had experienced the reference mortality rate over the span T ,

$$\text{CMR} = O/E.$$

Alternatively, we may be interested in the comparison between the observed rate of death in the cohort over the study period and the expected rate of death in the cohort in the absence of exposure. We may describe this causal rate ratio (CRR) as the ratio of the death rate over follow-up (observed deaths divided by person-time) in the study cohort to the rate that would be expected in the same group of individuals if the cohort members had experienced the reference mortality rate,

$$\text{CRR} = h^1/h^0 = (O/P^1)/(E/P^0).$$

Similar to the case with the SMR, we can estimate the CMR or CRR in an exposed population using an estimate of the mortality rate h^0 from a reference population. A causal interpretation of the CMR or CRR estimate may be made under the conditions laid out by Hernán (1), a central assumption being that the reference rate represents the mortality rate that would have observed in the cohort in the absence of exposure to the occupational hazard.

Under the classical approach to calculation of the expected number of events and the SMR, the expected number of deaths is calculated as the product of the mortality rate in the (unexposed) reference population and the person-time accrued in the (exposed) study cohort (14). The total expected number of deaths based on this approach, denoted \tilde{E} , is

$$h^0 P^1 = \frac{h^0}{h^1} \sum_{i=1}^n [1 - \exp(-h^1 T)].$$

The mixed nature of this expression (i.e., the product of the hazard rate in the reference population and the person-time accrued in the exposed occupational cohort) leads to problems. If the occupational cohort actually had experienced the same hazard rate as the reference population, the person-time in the study cohort would have changed as well. As a result, \tilde{E} may not closely approximate E . These quantities will be equal only if $h^1 = h^0$. If $h^1 > h^0$, then $\tilde{E} < E$ and the SMR, O/\tilde{E} , will be greater than the CMR, O/E . This agrees with the observation by Hartz et al. (10), as well as Keiding and Vaeth (12), that the expected number of deaths as calculated under the standard method is biased downward when the study population mortality rate, h^1 , is larger than that of the standard population, h^0 .

However, the ratio O/\tilde{E} does yield the CRR,

$$\frac{O}{\tilde{E}} = \frac{\sum_{i=1}^n [1 - \exp(-h^1 T)]}{\frac{h^0}{h^1} \sum_{i=1}^n [1 - \exp(-h^1 T)]} = \frac{h^1}{h^0};$$

and under the condition of constant hazards, the SMR does estimate the CRR. This conforms to the observation by Keiding and Vaeth (12) that the SMR is useful for statistical inference regarding relative death rates.

Expected number of deaths when the hazard rate is not constant

The assumption of a constant hazard may be overly restrictive in many settings. Here, we allow it to vary with baseline characteristics and over the course of time in the study. Let \mathbf{W} denote a vector of baseline characteristics such as race, sex, age at entry, and calendar year of entry. Denote the mortality rate at time t in the occupational cohort by $h^1(t|\mathbf{W})$, and denote the reference hazard rate function (e.g., race, sex, age, and calendar period-specific death rates for a region or nation) by $h^0(t|\mathbf{W})$. Furthermore, suppose we allow that potential follow-up time to vary between cohort members, as occurs when there is staggered entry into the study and the administrative end of study follow-up is a single calendar date. Therefore, we again define study entry as time 0 and now denote potential study follow-up as T_i years for person i .

Suppose that follow-up time has been grouped into discrete time intervals, where $L(u)$ is the duration of follow-up over the u th time period. Let $S^a(u|\mathbf{W}_i)$ denote the probability of surviving through time u or 1 minus the probability of being deceased by time u —that is,

$$S^a(u|\mathbf{W}_i) = 1 - \sum_{v=0}^u h^a(v|\mathbf{W}_i) S^a(v-1|\mathbf{W}_i) L(v),$$

where we define $S^a(-1|\mathbf{W}_i) = 1$ and assume that the rate is suitably small for this approximation; namely, our estimands are conditional on survival until entrance into the study.

The expected value of Y_i^a for a person with the hazard rate $h^a(t|\mathbf{W}_i)$ over the span of T_i years is

$$\sum_{u=0}^{T_i} h^a(u|\mathbf{W}_i) S^a(u-1|\mathbf{W}_i) L(u),$$

which simplifies to $1 - \exp(-h^a T)$ if h^a and T are constant (see Web Appendix 1, available at <http://aje.oxfordjournals.org/>). Letting O denote the total number of deaths in a cohort of size n given occupational exposure ($a = 1$), the expected value of O is

$$\sum_{i=1}^n \sum_{u=0}^{T_i} h^1(u|\mathbf{W}_i) S^1(u-1|\mathbf{W}_i) L(u).$$

Letting E denote the overall number of deaths in a cohort of size n in the absence of occupational exposure ($a = 0$), the expected value of E is

$$\sum_{i=1}^n \sum_{u=0}^{T_i} h^0(u|\mathbf{W}_i) S^0(u-1|\mathbf{W}_i) L(u).$$

An estimator for person-time, P^a , in the cohort of interest ($a = 0$ or $a = 1$) is

$$\sum_{i=1}^n \sum_{u=0}^{T_i} S^a(u|\mathbf{W}_i) L(u).$$

This is an estimator of the person-time in the cohort ($a = 0$ or $a = 1$) because the average time contributed to the study by a cohort member, i , is the area under the survival curve, which is estimated by $\sum_{u=0}^{T_i} S^a(u|\mathbf{W}_i) L(u)$.

Using these quantities, we can define 2 contrasts of potential interest: $CMR = O/E$ and $CRR = (O/P^1)/(E/P^0)$.

Under the classical approach to calculation of expected events and the SMR, the expected number of deaths is calculated as the product of the hazard rate in the (unexposed or $a = 0$) reference population and the person-time accrued in the (exposed or $a = 1$) study cohort (14). The contribution to the expected number of deaths made by person i , denoted \tilde{Y}_i^0 , is

$$\sum_{u=0}^{T_i} h^0(u|\mathbf{W}_i) S^1(u-1|\mathbf{W}_i) L(u),$$

and the total expected number of deaths based on this approach, denoted \tilde{E} , is

$$\sum_{i=1}^n \sum_{u=0}^{T_i} h^0(u|\mathbf{W}_i) S^1(u-1|\mathbf{W}_i) L(u).$$

\tilde{Y}_i^0 may not closely approximate Y_i^0 , and \tilde{E} may not closely approximate E . These quantities will be equal only if the sharp null hypothesis is true: $h^0(u|\mathbf{W}_i) = h^1(u|\mathbf{W}_i)$. Therefore, except at the null, $SMR = O/\tilde{E}$ will not equal $CMR = O/E$, since the survival curve in the calculation of \tilde{E} does not represent the survival curve that would have

been obtained if, counter to the fact, the study population had experienced the reference hazard rate.

In the case of a constant hazard rate, we noted that the classical SMR was useful for statistical inference regarding relative death rates, and in fact yielded the CRR. However, this is not the case generally if hazard rates are not constant.

$$\text{SMR} = \frac{O}{\bar{E}} = \frac{\sum_{i=1}^n \sum_{u=0}^{T_i} h^1(u|W_i) S^1(u-1|W_i) L(u)}{\sum_{i=1}^n \sum_{u=0}^{T_i} h^0(u|W_i) S^1(u-1|W_i) L(u)}$$

need not equal $\text{CRR} = (O/P^1)/(E/P^0)$:

$$\frac{\sum_{i=1}^n \sum_{u=0}^{T_i} h^1(u|W_i) S^1(u-1|W_i) L(u) / \sum_{i=1}^n \sum_{u=0}^{T_i} S^1(u-1|W_i) L(u)}{\sum_{i=1}^n \sum_{u=0}^{T_i} h^0(u|W_i) S^0(u-1|W_i) L(u) / \sum_{i=1}^n \sum_{u=0}^{T_i} S^0(u-1|W_i) L(u)}$$

These quantities will be equal if the sharp null hypothesis is true: $h^0(u|W_i) = h^1(u|W_i)$; otherwise, the SMR may not approximate the CRR.

Cause-specific mortality and competing risks

Suppose we are interested in cause-specific mortality rather than all-cause mortality. The same approach is taken for calculation of the expected number of deaths; however, we now allow that the probability of survival depends upon 2 categories of cause of death: A and B, where B denotes death due to all causes other than A. Allowing $h_A^0(t|W)$ to denote the discrete-time hazard rate of outcome A in the absence of exposure and $h_B^0(t|W)$ to denote the hazard of outcome B in the absence of exposure, the expected number of deaths due to cause A, denoted E_A , is calculated as

$$\sum_{i=1}^n \sum_{u=0}^{T_i} h_A^0(u|W_i) S^0(u-1|W_i) L(u),$$

where $S^0(u|W_i)$ is defined as the overall probability of survival up to time u , given as

$$S^0(u|W_i) = 1 - \left\{ \sum_{v=0}^u h_A^0(v, W_i) S^0(v-1, W_i) L(v) + \sum_{v=0}^u h_B^0(v, W_i) S^0(v-1, W_i) L(v) \right\}.$$

The CMR for the specific cause of death A is the ratio of the observed number of deaths due to cause A, denoted O_A , to the number of deaths due to cause A that would have been expected in the same group of individuals if the cohort had experienced the reference mortality rates ($h_A^0(t|W)$ and $h_B^0(t|W)$), denoted E_A —that is,

$$O_A/E_A.$$

Under the classical SMR approach, the total expected number of deaths due to cause A, denoted \bar{E}_A , can be calculated as

$$\sum_{i=1}^n \sum_{u=0}^{T_i} h_A^0(u|W_i) S^1(u-1|W_i) L(u),$$

and, again, the nature of the expression leads to problems. E_A need not approximate \bar{E}_A because $S^1(u|W_i)$ need not equal $S^0(u|W_i)$. These quantities will be equal if the sharp null hypothesis is true. Therefore, except at the null, the SMR for cause A, O_A/\bar{E}_A , need not equal the CMR for cause A, O_A/E_A .

Calculation

The CMR and CRR require the same information as that needed for SMR calculations. These quantities can be readily calculated using a simple program written for the SAS statistical package (SAS Institute, Inc., Cary, North Carolina) (15) (see Web Appendix 2).

Example: mortality among uranium miners

We illustrate these methods using data for a cohort of white men who entered follow-up of the Colorado Plateau miners' study in 1960 and were followed through 2005 for ascertainment of deaths; less than 1% of this cohort was lost to follow-up (16). The expected numbers of deaths due to all causes and to lung cancer were calculated using life-table methods, applying reference mortality rates for US white males as a function of attained age and calendar period. These race- and sex-specific national death rates for 5-year age and calendar-period intervals were based on the reference rate files compiled for the Life Table Analysis System of the National Institute for Occupational Safety and Health (17). We calculated expected numbers of deaths due to all causes and to lung cancer and calculated the CMR for each outcome; 95% confidence intervals were calculated using Byar's method (18), as implemented in the OpenEpi calculator (19). We also calculated the expected numbers of deaths due to all causes and to lung cancer applying the classical SMR method, the SMR for each outcome, and approximate 95% confidence intervals (again using Byar's method). Finally, to illustrate these calculations in a cohort essentially followed to extinction, we repeated these calculations in a subcohort of 587 white men who were aged 50 years or older at study entry (i.e., on January 1, 1960) and calculated the all-cause mortality CRR for this subcohort.

RESULTS

Among the 3,254 white male Colorado Plateau miners alive in 1960 and subsequently followed through 2005, the observed number of deaths was 2,428 (Table 1). The expected number of deaths due to all causes was 2,167.3, yielding a CMR of 1.1 (95% confidence interval (CI): 1.1, 1.2). A standard SMR calculation yields an SMR of 1.4 (95% CI: 1.4, 1.5) based on 1,693.2 expected deaths. Note that the expected number of deaths under the SMR calculation is substantially less than the number of deaths expected in the cohort over the study period based on the reference

Table 1. Distribution of 3,254 White Male Miners in the Colorado Plateau Miners Study, by Mortality Outcome, 1960–2005

Cause of Death	Observed No. of Deaths (O)	Classical SMR Calculation		CMR Calculation	
		Expected No. of Deaths (\bar{E})	SMR (O/\bar{E})	Expected No. of Deaths (E)	CMR (O/E)
All causes	2,428	1,693.2	1.4	2,167.3	1.1
Lung cancer	549	148.9	3.7	186.0	3.0
All other causes	1,879	1,544.3	1.2	1,981.3	1.0

Abbreviations: CMR, causal mortality ratio; SMR, standardized mortality ratio.

mortality rates. This is because the standard SMR uses the observed person-time to calculate the expected number of deaths, while the CMR uses the person-time that would have been expected had the study population actually had the hazard rates in the reference population. These lower reference rates lead to an increase in person-time for the CMR calculations and thus an increased number of expected deaths. There were 549 deaths due to lung cancer observed among the 3,254 white male miners. The expected number of deaths due to lung cancer was 186.0, yielding a CMR of 3.0 (95% CI: 2.7, 3.2). A standard SMR calculation yields an SMR of 3.7 (95% CI: 3.4, 4.0) based on 148.9 expected lung cancer deaths; again, the expected number derived by the SMR calculation is less than the number of lung cancer deaths expected in the cohort over the study period based on the reference mortality rates.

In the subcohort of 587 white male miners who entered follow-up at age ≥ 50 years in 1960 and were followed through 2005, 585 of the men were deceased by the end of study follow-up (Table 2). The expected number of deaths due to all causes was 579.0, yielding a CMR of 1.0 (95% CI: 0.9, 1.1). A standard SMR calculation yields an SMR of 1.3 (95% CI: 1.2, 1.4) based on 443.0 expected deaths; again, the expected number of deaths under the SMR calculation is substantially less than the number of deaths expected in the cohort over the study period based on the reference mortality rates.

There were 97 deaths due to lung cancer observed in the subcohort of 587 white male miners. The expected number of lung cancer deaths was 34.1, yielding a CMR of 2.8 (95% CI: 2.3, 3.5). Given a CMR for all causes near 1 and an elevated CMR for death due to lung cancer (Table 2), it is useful to note that the CMR for competing events is 0.9 (95% CI: 0.8, 1.0) based on 544.9 expected deaths due to all causes other than lung cancer (compared with 488 observed). A

standard SMR calculation, by contrast, yields an SMR of 1.2 (95% CI: 1.1, 1.3) based on 415.9 expected deaths due to all causes other than lung cancer.

While the CMR for this subcohort is 1.0, the effect of the exposure is reflected by evidence of life-shortening. The observed number of person-years in this subcohort is 9,542, while 11,203 person-years would have been expected if this subcohort had experienced mortality rates comparable to those of the reference population. Using the observed and expected numbers of deaths and person-years in this subcohort, we calculate the CRR. We compare the observed all-cause death rate (61.3 per 1,000 person-years) with the rate expected based on the numbers of deaths and person-years that would be expected if this subcohort had experienced mortality rates comparable to those of the reference population (51.7 per 1,000 person-years). The ratio of these rates (CRR = 1.2) differs slightly from the SMR (SMR = 1.3), reflecting a difference between a comparison of observed and expected mortality rates in the study population and the classical SMR.

DISCUSSION

If exposure affects mortality rates, the expected number of deaths calculated in the standard fashion for SMR analyses does not correspond to the number of people who would be expected to be deceased in the cohort in the absence of exposure. The calculations for the CMR and CRR described in our paper do not require the assumption that exposure does not affect the distribution of person-time, and they may provide more reliable measures of effect in many instances.

In the previous literature, a number of reasons have been given for not calculating expected numbers of deaths based upon expected survival functions. Keiding and Vaeth

Table 2. Distribution of 587 White Male Miners in the Colorado Plateau Miners Study Who Were Aged ≥ 50 Years at Study Entry in 1960 and Were Followed Until 2005, by Mortality Outcome

Cause of Death	Observed No. of Deaths (O)	Classical SMR Calculation		CMR Calculation	
		Expected No. of Deaths (\bar{E})	SMR (O/\bar{E})	Expected No. of Deaths (E)	CMR (O/E)
All causes	585	443.0	1.3	579.0	1.0
Lung cancer	97	27.1	3.6	34.1	2.8
All other causes	488	415.9	1.2	544.9	0.9

Abbreviations: CMR, causal mortality ratio; SMR, standardized mortality ratio.

correctly noted that the approach requires one to know the period of study (12). However, in the context of occupational cohort mortality studies, which is one setting where SMR calculations are commonly performed, the administrative end of study follow-up is almost always explicitly defined. This permits definition of the expected number of deaths over an explicitly defined interval of study follow-up. We agree with Keiding and Vaeth, however, that the CMR should not be used if the end of the study period cannot be explicitly defined.

It has been noted that when examining long-term mortality in a cohort, a comparison of mortality risks may not be the preferred summary measure of association (20). The CMR is based upon estimates of cumulative numbers of events at the end of follow-up; consequently, the expected mortality due to all causes is a function of the length of the study period, and in the long run, it goes to unity (because, in the long run, we all die). The standard SMR does not have this feature (as illustrated in Table 2). In the current paper, we commenced with the assumption that the expected number of deaths was a quantity of primary interest. When this is the case, reporting “expected” deaths as calculated under the method used in classical SMR calculations may erroneously suggest that a deleterious exposure is responsible for a larger absolute excess of mortality than could logically be expected. In contrast, the CMR can be interpreted as a causal parameter because it represents a contrast of the same group of individuals under 2 different scenarios. Indeed, the expected number of deaths in the CMR calculation is a special case of the g -formula (21) used for calculation of the cumulative incidence under no exposure. The fact that the CMR converges to unity in the long run does not suggest that the proposed CMR is without utility in long-term cohort studies; to the contrary, it accurately represents the fact that long-term comparisons of the observed and expected numbers of deaths due to all causes are not particularly informative. Further, the CMR need not be defined solely at the end of follow-up but can be defined as a function of time since the beginning of follow-up (22).

Some authors have suggested that in long-term mortality studies, attention should be focused on estimation of mortality rate ratios rather than mortality risk ratios (20). Given constant hazard rates, the SMR estimates the CRR, and Keiding and Vaeth suggested that its associated test statistic is optimal for statistical inference for relative death rates (12). However, we have shown that under more general conditions, the SMR will not estimate the ratio of observed and expected numbers of deaths (the CMR), nor will it provide a consistent estimate of the ratio of rates (the CRR), if the null condition is not true. Interestingly, if the proportional hazards assumption holds, the SMR will estimate this constant hazard ratio, which can be interpreted as the hazard ratio for any individual in the study population. However, this does not necessarily equal the CRR, which is a marginal, population-average quantity. Moreover, the SMR is standardized with respect to attained age and calendar time, and the distribution of person-time with respect to these characteristics is affected by a hazardous (or salubrious) exposure. Robins and Morgenstern (23) pointed out

how complications could arise in interpretation of epidemiologic measures in such settings; and Rothman et al. (13) noted that when exposure affects the distribution used to construct the standardization weights, comparison of standardized incidences will not properly reflect the net exposure effect.

The CRR presents a simple counterfactual comparison between the observed rate in the study cohort and the rate that would have been observed in the same population if it had experienced the reference hazard rate, so that the groups are comparable with respect to baseline covariates (such as age at entry into the study). The cumulative risk, and hence the CMR, may be useful in scenarios in which we wish to estimate net impacts of exposure on mortality, such as for harm reduction or to formalize the tradeoffs of policy choices. The CRR may be useful in other scenarios, and we also provide an approach to estimation of that quantity. As illustrated in our example, such methods readily allow for a summary rate ratio comparison standardized to baseline covariates, interpretable as the ratio of the marginal rate in the study cohort to the marginal rate that would have been expected in the same population over the study period if the cohort members had experienced the reference rate; and the CRR need not equal the classical SMR.

We have focused on comparisons of observed and expected numbers of deaths that are summarized in terms of a ratio measure. However, the quantities O and E also may be subtracted to yield a difference measure that can be interpreted as a causal parameter. Importantly, E_A need not equal \hat{E}_A , and consequently this quantity need not equal the difference between observed and “expected” values derived under the classical approach to the SMR, except when $h^1 = h^0$.

We have illustrated these issues using empirical data. For example, we calculated the expected mortality among 587 men who entered the Colorado Plateau miners’ study at age ≥ 50 years and were followed for the subsequent 46 years (Table 2). While nearly all of these ≥ 50 -year-olds would be expected to be deceased by the end of the study period in the absence of exposure (based on US mortality rates), a standard SMR calculation yields an estimate of the expected number of deaths in the cohort by the end of study follow-up that implies (somewhat absurdly) that approximately 25% of the cohort is expected to be alive at the end of follow-up (when the minimum age would be 96 years).

In addition, we illustrated the difference in these approaches for a specific category of cause of death, namely lung cancer mortality among underground miners. Given competing risks, it is useful to report both the event of interest and the competing event. As illustrated in Table 2, where the CMR for all-cause mortality approaches unity, an elevated lung cancer CMR implies a CMR below unity for non-lung cancer mortality. When exposure is harmful (causing lost years of life), we illustrate that the expected number of deaths in the absence of exposure (i.e., death due to the specific cause of interest and death due to the competing event) will tend to be greater under the CMR calculation than under the SMR. This reflects the fact that if exposure affects

the mortality rate (e.g., increasing lung cancer mortality rates), it will also tend to affect the distribution of person-time (e.g., leading to fewer years of observation in the presence of exposure than would be observed in the absence of exposure). Of course, if the hazard rate or the length of the study period is exceedingly small, an exposure's effect on the distribution of person-time will be small, and the SMR and CMR will be similar in magnitude.

The notion of counterfactuals has offered epidemiologists a useful framework for thinking about causation in observational studies. While certain assumptions must be made for the CMR to estimate the causal effect of an exposure (namely, exchangeability within strata of age, sex, and calendar year), the potential-outcomes framework provides a formal structure for discussing what such observed and expected contrasts might look like. The comparison of observed and expected events derived under traditional SMR calculations does not easily fit within this potential-outcomes framework; however, using the approach described in this paper, we illustrate how such a calculation can be performed to obtain a consistent estimate of the mortality ratio. While we have focused here on the CMR as applied to analyses of mortality, the CMR could be usefully applied for other outcomes, such as analyses of cancer incidence, where reference rates might be based on registry data.

In the current paper, we do not address issues of confounding bias that may arise due to noncomparability of the study cohort and the reference population (known in the causal inference literature as the assumption of conditional exchangeability) (1, 24). Rather, we assume appropriate selection of reference rates to focus our discussion on a problem in interpretation of the "expected" number of events even in settings where appropriate reference rates are available. Nonetheless, by clarifying how to calculate a CMR, which has a well-defined interpretation under specific conditions, we set a foundation for strengthening interpretation of analyses that involve comparison of mortality in a study cohort to expectations derived from calculations involving external reference mortality rates. Our prior work has suggested one approach to address such concerns using a negative control outcome (25). It may be useful to further explore how the negative control outcome approach could be applied in calculations of expected counts as derived here for the CMR calculation.

In summary, in this paper we illustrate how to calculate expected counts in a defined population over a specified period of follow-up without requiring the assumption that exposure does not affect the distribution of person-time. This strong assumption is required for interpreting the SMR as a contrast between observed and expected mortality. No such assumption is required for the CMR. Using the CMR allows us to sharpen our interpretation of results obtained from cohort mortality analyses by directly comparing observed event counts with expected event counts in the absence of exposure. The CMR and CRR require the same information as SMR calculations, and thus they may be straightforward to implement in the future as additional quantities estimated in statistical packages routinely used for computing the SMR.

ACKNOWLEDGMENTS

Author affiliations: Department of Epidemiology, School of Public Health, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina (David B. Richardson, Alexander P. Keil, Stephen R. Cole); and Division of Epidemiology and Community Health, School of Public Health, University of Minnesota, Minneapolis, Minnesota (Richard F. MacLehose).

A.P.K. was supported by the National Institute of Environmental Health Sciences (grant T32ES007018-30).

D.B.R. was supported by the National Institute for Occupational Safety and Health (grant R03 OH010946).

The authors gratefully acknowledge comments on an earlier draft of this paper provided by Dr. Jessie K. Edwards.

Conflict of interest: none declared.

REFERENCES

- Hernán MA. A definition of causal effect for epidemiological research. *J Epidemiol Community Health*. 2004;58(4):265–271.
- Case RA, Lea AJ. Mustard gas poisoning, chronic bronchitis, and lung cancer; an investigation into the possibility that poisoning by mustard gas in the 1914–18 war might be a factor in the production of neoplasia. *Br J Prev Soc Med*. 1955;9(2):62–72.
- Case RA, Hosker ME, McDonald DB, et al. Tumours of the urinary bladder in workmen engaged in the manufacture and use of certain dyestuff intermediates in the British chemical industry. I. The role of aniline, benzidine, alpha-naphthylamine, and beta-naphthylamine. *Br J Ind Med*. 1954;11(2):75–104.
- Case RA. The expected frequency of bladder tumour in works populations. *Br J Ind Med*. 1953;10(2):114–120.
- Doll R. Mortality from lung cancer in asbestos workers. *Br J Ind Med*. 1955;12(2):81–86.
- Monson RR. Observations on the healthy worker effect. *J Occup Med*. 1986;28(6):425–433.
- Darby SC, Muirhead CR, Doll R, et al. Mortality among United Kingdom servicemen who served abroad in the 1950s and 1960s. *Br J Ind Med*. 1990;47(12):793–804.
- Checkoway H, Pearce NE, Kriebel D. *Research Methods in Occupational Epidemiology*. 2nd ed. New York, NY: Oxford University Press; 2004.
- Miettinen OS, Cook EF. Confounding: essence and detection. *Am J Epidemiol*. 1981;114(4):593–603.
- Hartz AJ, Giefer EE, Hoffmann RG. A comparison of two methods for calculating expected mortality. *Stat Med*. 1983;2(3):381–386.
- Berry G. The analysis of mortality by the subject-years method. *Biometrics*. 1983;39(1):173–184.
- Keiding N, Vaeth M. Calculating expected mortality. *Stat Med*. 1986;5(4):327–334.
- Rothman KJ, Greenland S, Lash TL. *Modern Epidemiology*. 3rd ed. Philadelphia, PA: Lippincott Williams & Wilkins; 2007.
- Monson RR. Analysis of relative survival and proportional mortality. *Comput Biomed Res*. 1974;7(4):325–332.
- SAS Institute Inc. *SAS OnlineDoc*® 9.2. Cary, NC: SAS Institute Inc.; 2007. <http://support.sas.com/documentation/92/>. Accessed May 18, 2016.

16. Schubauer-Berigan MK, Daniels RD, Pinkerton LE. Radon exposure and mortality among white and American Indian uranium miners: an update of the Colorado Plateau cohort. *Am J Epidemiol*. 2009;169(6):718–730.
17. Robinson CF, Schnorr TM, Cassinelli RT 2nd, et al. Tenth revision US mortality rates for use with the NIOSH Life Table Analysis System. *J Occup Environ Med*. 2006;48(7):662–667.
18. Breslow NE, Day NE. *Statistical Methods in Cancer Research. Vol 2. The Design and Analysis of Cohort Studies*. (IARC Scientific Publication no. 82). Lyon, France: International Agency for Research on Cancer; 1987.
19. Dean AG, Sullivan KM, Soe MM. *OpenEpi: Open Source Epidemiologic Statistics for Public Health*. www.OpenEpi.com. Updated April 6, 2013. Accessed May 17, 2016.
20. Steenland K, Beaumont J, Schulte PA, et al. Re: A comparison of two methods for calculating expected mortality [letter]. *Stat Med*. 1985;4:105–106.
21. Robins J. A new approach to casual inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Math Model*. 1986;7(9-12):1393–1512.
22. Cole SR, Richardson DB, Chu H, et al. Analysis of occupational asbestos exposure and lung cancer mortality using the g formula. *Am J Epidemiol*. 2013;177(9):989–996.
23. Robins J, Morgenstern H. The foundations of confounding in epidemiology. *Comput Math Applic*. 1987;14(9-12):869–916.
24. Arrighi HM, Hertz-Picciotto I. The evolving concept of the healthy worker survivor effect. *Epidemiology*. 1994;5(2):189–196.
25. Richardson DB, Keil AP, Tchetgen Tchetgen E, et al. Negative control outcomes and the analysis of standardized mortality ratios. *Epidemiology*. 2015;26(5):727–732.