

and the estimated rate for females is

$$\frac{5\,420}{2\,306\,300 \times 8} = 294 \text{ per } 10^6 \text{ person-years.}$$

6.7 The follow-up of each subject can be represented by a line on a three-dimensional Lexis diagram with axes: age, period, and time since treatment. Age and period were divided into five-year bands and time since treatment into three-year bands. Observed deaths and person-years can be assigned to cells in the resulting three-dimensional table. Multiplication of person-years by national rates gives the expected number of deaths for each cell. Table 6.6 is formed by adding this table over age and period.

7 Competing risks and selection



7.1 Censoring in follow-up studies

Up to this point we have lumped all the different reasons for censoring together. In this chapter we look at this practice more carefully and make a distinction between censoring due to practical difficulties in maintaining follow-up (such as migration, refusal to participate further and so on), and censoring due to competing causes of failure.

The first class of events causes removal of a subject from observation, but after censoring the subject is still at risk of failure – a subject does not cease to run the risk of a myocardial infarction simply because he or she has ceased to participate in a follow-up study. Such observations are censored in the sense that this later experience is removed from our view. The second class of censoring events also causes removal of a subject from observation, but this time the subject is no longer at risk from the failure of interest. This is obviously true when a subject dies from a competing cause, but onset of a non-fatal competing disease can also remove a subject from the risk under study. For example, in a study of myocardial infarction in previously healthy subjects, a subject who suffers the onset of lung cancer would be considered as no longer at risk — although patients with lung cancer suffer myocardial infarctions quite frequently, the aetiology is so different as to be regarded as a different type of event.

7.2 Competing causes

The termination of follow-up by a competing cause is not due to imperfection of any one study, but is intrinsic to all imaginable studies. The binary model which underlies the measurement of disease frequency by rates and risks assumes only one type of failure. To allow for more than one type, the model must be extended. Fig. 7.1 illustrates a model with two causes of failure over a single study period of fixed duration. There are now three possible outcomes, labelled F1 and F2 for the two types of failure and S for survival. The probabilities of F1 and F2 are referred to as π_1 and π_2 , so the probability of survival is $1 - \pi_1 - \pi_2$. In incidence studies, π_1 and π_2 represent *cause-specific* failure probabilities or risks.

It is easy to use likelihood to estimate the parameters π_1 and π_2 . If N




Fig. 7.1. Two causes of failure

subjects are studied and we observe D_1 failures of the first type and D_2 failures of the second type, the likelihood is

$$(\pi_1)^{D_1} (\pi_2)^{D_2} (1 - \pi_1 - \pi_2)^{N - D_1 - D_2},$$

and the log likelihood is

$$D_1 \log(\pi_1) + D_2 \log(\pi_2) + (N - D_1 - D_2) \log(1 - \pi_1 - \pi_2).$$

This takes its maximum value when $\pi_1 = D_1/N$ and $\pi_2 = D_2/N$ so that the most likely values correspond with the intuitive measures — the proportions of subjects failing due to each cause.

Exercise 7.1. In a 5-year follow-up study of 1000 subjects, 27 suffered myocardial infarctions during the study period while 8 suffered strokes. (If any subject suffered both events, only the first was counted.) Estimate the cause-specific risks for these conditions. If myocardial infarctions and strokes are grouped together as ‘cardiovascular events’, what is the estimated risk of a cardiovascular event?

Fig. 7.2 illustrates the extension of this model to describe observation of a subject through several consecutive bands. Superscripts denote band and subscripts continue to indicate the type of failure. As in the case of a single cause, the π parameters are defined as conditional probabilities. For example, π_1^3 represents the probability of failure F1 during the third band, conditional upon survival through all preceding bands. The log likelihood behaves as if the time bands form separate studies involving different groups of subjects, so for each band the cause-specific failure probabilities are estimated by the proportion of those subjects at risk during the band, failing from the specified cause.

Exercise 7.2. The conditional probabilities of F1 and F2 remain constant at 0.1 and 0.2 respectively over three bands. List the 7 possible outcomes and calculate their probabilities.




Fig. 7.2. Consecutive time bands

Table 7.1. Log likelihood contributions for a subject during one click

Outcome	Log likelihood
F1	$\log \lambda_1 + \log h$
F2	$\log \lambda_2 + \log h$
S	$-(\lambda_1 + \lambda_2)h$

7.3 Cause-specific rates

The same argument can be extended to rates by dividing the time scale into clicks. Fig. 7.1 now represents the possible outcomes for one subject during a single click. The conditional failure probabilities are

$$\pi_1 = \lambda_1 h, \quad \pi_2 = \lambda_2 h,$$

where h is the duration of a click and λ_1 and λ_2 are cause-specific rates — conditional probabilities per unit time. Because the probabilities of failure are very small, we can make the approximation

$$\log(1 - \pi_1 - \pi_2) \approx -\pi_1 - \pi_2 = -(\lambda_1 + \lambda_2)h,$$

and the contributions to the log likelihood of a single subject during a single click are then those shown in Table 7.1. The total log likelihood is obtained by summing such terms over subjects and over clicks. There are D_1 clicks which result in failure of type F1 and these contribute a total of

$$D_1 \log(\lambda_1) + D_1 \log(h)$$

to the log likelihood. Since the second term does not depend upon parameters it can be ignored. Similarly the D_2 failures of type F2 contribute $D_2 \log(\lambda_2)$. Because every subject, regardless of eventual outcome, survives all the clicks save the last, the sum of all of these log likelihood contributions over both subjects and clicks is

$$\sum -(\lambda_1 + \lambda_2)h = -(\lambda_1 + \lambda_2)Y,$$

where Y is the total person-time of observation of the cohort. The grand total of all these contributions to the log likelihood is

$$D_1 \log(\lambda_1) + D_2 \log(\lambda_2) - (\lambda_1 + \lambda_2)Y.$$

A minor rearrangement of this expression leads to

$$D_1 \log(\lambda_1) - \lambda_1 Y + D_2 \log(\lambda_2) - \lambda_2 Y$$

so that the log likelihood is the sum of two parts, both Poisson in form, the first referring to F1 and the second to F2. The fact that the log likelihood falls into two distinct parts, one for each cause, justifies the standard practice of analyzing each cause separately, allowing for competing causes only in that they curtail further observation. The argument is easily generalized to allow for more than two causes.

7.4 Interpreting cause-specific rates

There has been some controversy as to whether the practice of estimating cause-specific rates in this way requires us to assume *independence* of causes – an assumption which might often not be justified. In fact, the split of the log likelihood into a sum of separate parts, one for each cause-specific rate, does not arise as a result of any assumption of independence of causes, but out of the way cause-specific rate parameters are defined. The rate for cause 1 is defined as the probability per unit time of failure due to cause 1, conditional upon the subject having previously survived *all* causes of failure. This quantity is not truly specific to one cause. Influences which directly influence one cause can, because of this, have an indirect affect on rates for another cause. The term *cause-specific* is misleading. For example, it is likely that myocardial infarction and stroke compete for the same high risk subgroup of the population: those with advanced atherosclerosis. A preventive measure which reduced the incidence rate of myocardial infarction without reducing the prevalence of atherosclerosis would result in an *increase* in the rate of stroke, since more of the atherosclerotic group would survive to be at risk from stroke.

It is a common practice to apply the formula

$$\log(\text{Cumulative survival probability}) = -\text{Cumulative rate.}$$




Fig. 7.3. Elimination of cause F2

to the cumulative *cause-specific* rate to calculate a *cause-specific* survival probability, interpreted as the probability of survival which would be observed if all other causes of failure were eliminated. However, this interpretation *does* depend on the assumption that the different causes of failure are independent. This is illustrated in Fig. 7.3. If the causes are independent, subjects who would have failed due to F2 have exactly the same conditional probabilities of failure due to F1 as those who would not. Under these circumstances, elimination of cause F2 will have no effect on the subsequent rate for F1, and the exponential function of minus the cumulative cause-specific rate for F1 can be interpreted as a survival probability when cause F2 is eliminated. More generally we might expect elimination of other causes to have an effect on the rate for the remaining one and the cumulative cause-specific rate will then have no such interpretation. Since the independence of different causes is usually untestable, it is best to avoid such interpretations and to leave estimates of cumulative cause-specific failure rates, calculated by the modified life table or Aalen-Nelson method, without converting them to cumulative probabilities of survival. Conversely, if the actuarial life table and Kaplan-Meier methods of Chapter 4 are applied to cause-specific failure probabilities, the resulting 'survival probability' should be transformed to a cumulative rate by taking minus its logarithm.

7.5 Selection bias

We now turn to the other reasons for censoring in follow-up studies. The statistical theory is exactly the same as for competing causes – we simply relabel the two causes as failure and loss to follow-up (Fig. 7.4). However, the question of dependence between failure and censoring takes on a new significance, because censoring arises as a result of the imperfection of




Fig. 7.4. Loss to follow-up

real studies, rather than unavoidable biological realities. We would like to estimate what would have happened in an ideal study in which no censoring occurred, but in order to do this we need to assume that censoring and failure are independent. More precisely, this means that those lost to follow-up due to censoring must have the same probabilities of failure in later bands as those remaining under observation. If this is the case, censoring is said to be *non-informative*. If not, the study results will be subject to *selection bias*.

A well known example of selection bias due to censoring arises in clinical trials when patients become so ill that their clinicians are unable ethically to maintain them on a randomized double blind protocol. The randomization code is then broken and the clinician is free to modify the treatment as necessary. If observation of such patients is regarded as censored at this point, the analysis is seriously biased, because these patients have a worse prognosis than those remaining in the trial. It will almost always be preferable to continue the follow-up of these patients and to analyze the data according to the initial treatment assigned. This is known as analysis by *intention to treat*.

Similar considerations apply when there is *late entry* to follow-up studies. Ideally, subjects should be recruited at the starting point for the failure process under study. This is usually the case in clinical epidemiology, where patients are recruited into the study at diagnosis, the natural starting point for a prognostic study. In many epidemiological studies, however, subjects are recruited some time after the natural starting point (see Fig. 7.5) This is known as *late entry*; it can introduce further selection bias if the new subjects have different subsequent probabilities of failure from the survivors they join. For example, clinical follow-up studies are frequently carried out in cohorts initially recruited from patients under treatment in a group




Fig. 7.5. Selection due to loss and late entry

of participating hospitals when the study starts. These cohorts are then extended by addition of new patients as they are diagnosed. When carrying out analyses of survival times from diagnosis, the initial members of such a cohort are late entries, because their diagnosis preceded their entry into the study. This introduces possibilities for selection bias, because the initial cohort could include patients diagnosed elsewhere, but would omit patients diagnosed in the participating hospitals and referred elsewhere for treatment.

In epidemiological studies of the causes of disease, late entry is almost universal and we must be careful that it does not introduce bias. An example of bias arising this way is the *healthy worker* effect, so called because of the widespread empirical finding that occupationally recruited cohorts have lower mortality than general population rates would suggest. This arises partly because of selective recruitment into occupations but mainly because early retirement and job changes in response to ill health act to prevent entry of 'unhealthy' workers into the cohort. The ideal study would recruit subjects on entry to the occupation, but in practice the subjects actually recruited are those in employment on a particular date. Follow-up then starts on that date. Factors such as early retirement, and job changes in response to ill health, can operate in the period between joining the occupation and recruitment to prevent entry of unhealthy workers into the cohort.

The phenomenon of late entry is closely related to the distinction sometimes drawn between 'closed' cohorts (in which only failures and censoring can occur) and 'dynamic' cohorts, which can be refreshed by new entrants. This distinction implies that being closed or dynamic is a property of the cohort, but this is incorrect. The occurrence or non-occurrence of late entry is not a property of the cohort, but depends on the time scale chosen for

analysis. If survival is analyzed by time in study there are no late entries, but in an analysis of the same study by age, or by time since entering an occupation, there will be late entries.

Solutions to the exercises

7.1 The estimated 5-year risk of myocardial infarction is 27/1000 while that for stroke is 8/1000. The risk of a cardiovascular event is 35/1000.

7.2 The outcomes and their probabilities are listed below.

Outcome	Probability
Band 1	
F1	0.1
F2	0.2
Band 2	
F1	$0.7 \times 0.1 = 0.07$
F2	$0.7 \times 0.2 = 0.14$
Band 3	
F1	$0.7 \times 0.7 \times 0.1 = 0.049$
F2	$0.7 \times 0.7 \times 0.2 = 0.098$
S	$0.7 \times 0.7 \times 0.7 = 0.343$

8 The Gaussian probability model

Until now we have been concerned only with the binary probability model. In this model there are two possible outcomes and the total probability of 1 is shared between them. It is an appropriate model when studying the occurrence of events, but not when studying a response for which there are many possible outcomes, such as blood pressure. For this the *Gaussian* or *normal* probability model is most commonly used.

In the Gaussian model the total probability of 1 is shared between many values. This is illustrated in the left panel of Fig. 8.1. When measurements are recorded to a fixed number of decimal places, there is a finite number of possible outcomes but, in principle, such measurements have infinitely many possible outcomes, so the probability attached to any one is effectively zero. For this reason it is the probability *density* per unit value which is specified by the model, not the probability of a given value. This is illustrated in the right panel of the figure. If π is the probability shared between values in a very narrow range, width h units, the probability density is π/h .

8.1 The standard Gaussian distribution

The standard Gaussian distribution has probability density centred at 0. The probability density at any value z (positive or negative) is given by

$$0.3989 \exp \left[-\frac{1}{2}(z)^2 \right].$$

A graph of this probability density for different values of z is shown in Fig. 8.2. There is very little probability outside the range ± 3 .

Tables of the standard Gaussian distribution are widely available, and these readily allow calculation of the probability associated with specified ranges of z . For our purposes it is necessary only to record that the probability corresponding to the range $(-1.645, +1.645)$ is 0.90 and that for the range $(-1.960, +1.960)$ is 0.95.

If the probability model for z is a standard Gaussian distribution then the probability model for $(z)^2$ is called the *chi-squared* distribution on one degree of freedom. Tables of chi-squared distributions can be used to find