

An Introduction to Veterinary Epidemiology *

Mark Stevenson †

EpiCentre, IVABS ‡

Massey University, Palmerston North, New Zealand

October 24, 2007

Contents

1 Introduction	6
1.1 Host, agent, and environment	7
1.2 Individual, place, and time	7
Individual	7
Place	8
Time	8
1.3 Causation	10
2 Measures of health	15
2.1 Prevalence	15
2.2 Incidence	16
Incidence risk	16
Incidence rate	17
The relationship between prevalence and incidence	18
2.3 Other measures of health	19
Attack rates	19
Secondary attack rates	20
Mortality	20
Case fatality	21

*Notes for Veterinary Biometrics and Epidemiology, taught as course 227.407 within the BVSc program at Massey University. Contributions from Dirk Pfeiffer, Cord Heuer, Nigel Perkins, and John Morton are gratefully acknowledged.

†M.Stevenson@massey.ac.nz

‡URL: <http://epicentre.massey.ac.nz>

Proportional mortality	21
2.4 Adjusted measures of health	21
Stratum-specific rates	21
Direct adjustment	22
Indirect adjustment	23
3 Study design	27
3.1 Descriptive studies	27
Case reports	27
Cases series	28
Descriptive studies based on rates	28
3.2 Analytical studies	28
Ecological studies	29
Cross-sectional studies	29
Cohort studies	30
Case-control studies	31
Hybrid study designs	32
3.3 Experimental studies	34
Randomised clinical trials	34
Community trials	35
3.4 Comparison of major the major study designs	35
4 Measures of association	37
4.1 Measures of strength	38
Risk ratio	38
Incidence rate ratio	38
Odds ratio	38
4.2 Measures of effect in the exposed population	39
Attributable risk (rate)	39
Attributable fraction	39
4.3 Measures of effect in the total population	40
Population attributable risk (rate)	40
Population attributable fraction	40
4.4 Using the appropriate measure of effect	41

5	Confounding and interaction	43
5.1	Confounding	43
5.2	Interaction	44
5.3	Identifying the presence of confounding and interaction	44
5.4	Methods for dealing with confounding	47
	Restriction	47
	Randomisation	47
	Stratification	47
	Matching	47
	Multivariate methods	48
5.5	A worked example	48
6	Diagnostic tests	51
6.1	Screening versus diagnosis	51
6.2	Sensitivity and specificity	51
6.3	Accuracy and precision	52
	Accuracy	52
	Precision	52
6.4	Test evaluation	53
	Sensitivity	53
	Specificity	54
	Positive predictive value	54
	Negative predictive value	55
6.5	Prevalence estimation	56
6.6	Diagnostic strategies	56
	Parallel interpretation	57
	Serial interpretation	57
6.7	Screening and confirmatory testing	57
6.8	Likelihood ratios	58
7	Sampling	62
7.1	Probability sampling methods	62
	Simple random sampling	62
	Systematic random sampling	63
	Stratified random sampling	63
	Cluster sampling	64
7.2	Non-probability sampling methods	65
7.3	Sources of error and how to reduce error	65
	Random error	66

	Bias	67
7.4	Sampling techniques	67
	Methods of randomisation	67
	Replacement	68
	Selection proportional to size	68
7.5	Sample size	69
	Simple and systematic random sampling	69
	Sampling to detect disease	70
8	Outbreak investigation	72
8.1	Verify the outbreak	72
	What is the illness?	72
	Is there a true excess of disease?	72
8.2	Investigating an outbreak	72
	Establish a case definition	72
	Enhance surveillance	73
	Describe outbreak according to individual, place and time	73
	Develop hypotheses about the nature of exposure	74
	Conduct analytical studies	74
8.3	Implement disease control interventions	74
9	Appraising the literature	76
9.1	Description of the evidence	76
9.2	Internal validity	77
	Non-causal explanations	77
	Positive features of causation	77
9.3	External validity	78
	Can the results be applied to the eligible population?	78
	Can the results be applied to the source population?	78
	Can the results be applied to other relevant populations?	79
9.4	Comparison of the results with other evidence	79
	Consistency	79
	Specificity	79
	Plausibility	79
	Coherency	80

10 Exercise: outbreak investigation	81
10.1 The problem	81
10.2 Diagnosis	81
10.3 Measures of disease frequency	81
10.4 Investigation	82
10.5 Measures of association	82
10.6 Recommendations	83
10.7 Clinical trial	83
10.8 Financial impact	83
11 Review questions	85
11.1 Host, agent, environment	85
11.2 Measures of health	85
11.3 Measures of association	86
11.4 Study design	87
11.5 Diagnostic tests	87
12 Resources	88

1 Introduction

By the end of this unit you should be able to:

- Compare and contrast clinical approaches and epidemiological approaches to disease management.
- Describe the factors that influence the presence of disease in individuals.
- Describe the factors that influence the presence of disease in populations.
- Explain what is meant by a point source and a common source epidemic.
- Explain what is meant by the term causation.

Epidemiology is the study of diseases in populations. Epidemiologists attempt to characterise those individuals in a population with high levels of disease and those with low levels. They then ask questions that help them discover what the high rate group is doing that the low rate group is not or *vice versa*. This allows the factors influencing the risk of disease to be identified. Once identified, measures can be applied to reduce exposure to these factors — reducing the overall burden of disease in the population. This allows disease to be controlled even if the precise pathogenic mechanism (or the aetiologic agent) is not known.

It is useful to distinguish epidemiological from clinical approaches to disease management. The **clinical approach** is focussed on individual animals and is aimed at diagnosing a disease and then treating it. It involves physical examination and generation of a list of differential diagnoses. Further examinations, laboratory tests and possibly response to treatment are then used to narrow the list of differential diagnoses to a single diagnosis. In an ideal world this will always be the correct diagnosis. The success of this approach depends on two conditions:

- That the true diagnosis is on the list of differential diagnoses; and
- Clinical signs arise from a single disease process.

Research in health professionals has shown that the final diagnosis is nearly always drawn from the initial differential list. If the disease is not on the initial list of differentials then it tends not to become the final diagnosis. Diseases may be omitted from the list because the clinician is not familiar with them (exotic or unusual diseases) or because the disease is ‘new’ and has never been identified before. The single cause idea is true in some diseases (e.g. parvo virus causes a characteristic clinical syndrome in dogs) however in many cases there are multiple causative factors interacting in a complex web that may or may not produce disease.

The **epidemiological approach** to disease management is conceptually different in that there is no dependency on being able to precisely define the aetiological agent. It is based on observing differences and similarities between diseased and non-diseased animals in order to try and understand what factors may be increasing or reducing the risk of disease.

In practice, clinicians unwittingly use a combination of clinical and epidemiological approaches in their day-to-day work. If the problem is relatively clear-cut then an epidemiological approach plays a very minor role. If the condition is new or more complex then the epidemiological approach is preferred since it will provide a better understanding of what makes individuals susceptible to disease and — once these factors are known — the measures required to control the disease become better defined.

1.1 Host, agent, and environment

Whether or not disease occurs in an **individual** depends on an interplay of three factors:

- The host;
- The agent; and
- The environment

The host is the animal (or human) that may contract a disease. Age, genetic makeup, level of exposure, and state of health all influence a host's susceptibility to developing disease. The agent is the factor that causes the disease (bacteria, virus, parasite, fungus, chemical poison, nutritional deficiency etc) — one or more agents may be involved. The environment includes surroundings and conditions either within the host or external to it, that cause or allow disease transmission to occur. The environment may weaken the host and increase its susceptibility to disease or provide conditions that favour the survival of the agent.

1.2 Individual, place, and time

The level of disease in a **population** depends on an interplay of three factors:

- Individual factors: what types of individuals tend to develop disease and who tends to be spared?
- Spatial factors: where is the disease especially common or rare, and what is different about these places?
- Temporal factors: how does disease frequency change over time, and what other factors are associated with these changes?

Individual

Individuals can be grouped or distinguished on a number of characteristics: age, sex, breed, coat colour and so on. An important component of epidemiological research is aimed at determining the influence of individual characteristics on the risk of disease. Figure 1 shows how mortality rate for drowning varied among children and young adults in the USA during 1999. The rate was highest in those aged 1 – 4 years: an age when children are mobile and curious about everything around them, even though they do not understand the hazards of deep water or how to survive if they fall in. What conclusions do we draw from this? Mortality as a result of drowning is highest in children aged 1 – 4 years: preventive measures should be targeted at this age group.




Figure 1: Mortality from drowning by age: USA, 1999. Reproduced from: Hoyert DL, Arias E, Smith BL, Murphy SL, Kochanek KD (2001) Deaths: final data for 1999. National Vital Statistics Reports volume 49, number 8. Hyattsville MD: National Center for Health Statistics.

Place

The spatial pattern of disease is typically a consequence of environmental factors. Environmental factors include aspects of climate (temperature, humidity, rainfall) as well as aspects of animal management (management of animals in a certain area of a country may result in high rates of disease that may not be seen in other areas). Geographic Information Systems and easy access to spatial data (e.g. satellite images) have facilitated the ability to conduct spatial epidemiological analyses in recent years. Figure 2 shows the geographical distribution of BSE incidence risk in British cattle from July 1992 to June 1993. These maps show a higher density of disease in the south of the country, compared with the north.

Time

Temporal patterns of disease in populations are presented graphically using epidemic curves. An epidemic curve consists of a bar chart showing time on the horizontal axis and the number of new cases on the vertical axis, as shown in Figure 3. The shape of an epidemic curve can provide important information about the nature of the disease under investigation. An **epidemic** occurs when there is a rapid increase in the level of disease in a population. An epidemic is usually heralded by an exponential rise in the number of cases in time and a subsequent decline as susceptible animals are exhausted. Epidemics may arise from the introduction of a novel pathogen (or strain) to a previously unexposed (naïve) population or as a result of the re-growth of susceptible numbers some time after a previous epidemic due to the same infectious agent. Epidemics may be described as being either common source or propagated.

In a **common source epidemic**, subjects are exposed to a common noxious influence. If the group is exposed over a relatively short period then disease cases will emerge over one incubation period. This is classified as a common point source epidemic. The epidemic of leukaemia cases in Hiroshima following the atomic bomb blast would be a good example of a point source

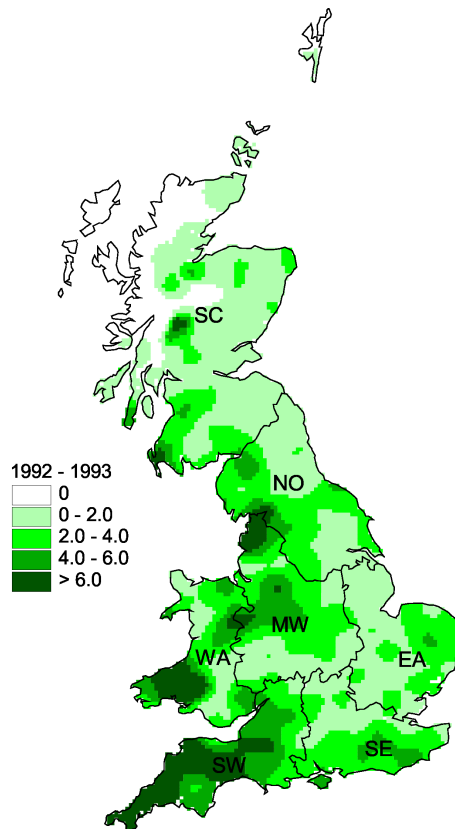


Figure 2: Incidence risk of BSE across Great Britain (expressed as confirmed BSE cases per 100 adult cattle per square kilometre), July 1992 – June 1993. Reproduced from Stevenson et al. (2000).

epidemic. The shape of this curve rises rapidly and contains a definite peak at the top, followed by a gradual decline. Exposure can also occur over a longer period of time, either intermittently or continuously. This creates either an intermittent common source epidemic or a continuous common source epidemic. The shape of this curve rises rapidly (associated with the introduction of the agent). The down slope of the curve may be very sharp if the common source is removed or gradual if the outbreak is allowed to exhaust itself.

A **propagated epidemic** occurs when a case of disease serves as a source of infection for subsequent cases and those subsequent cases, in turn, serve as sources for later cases. In theory, the epidemic curve of a propagated epidemic has a successive series of peaks reflecting increasing numbers of cases in each generation. The epidemic usually wanes after a few generations, either because the number of susceptibles falls below a critical level, or because intervention measures become effective.

Sometimes epidemic curves can show characteristics of being both common source and propagated. Figure 4 shows the epidemic curve for foot-and-mouth disease in the county of Cumbria (Great Britain) in 2001. This epidemic started as a common (point) source, then took on the characteristics of a propagative epidemic over time.

Endemic describes levels of disease which do not exhibit wide fluctuations over time. Epidemic curves for endemic disease might show evidence of seasonal variation (as in the case of monthly reports of human leptospirosis cases in the USA, shown on the left in Figure 5). If data are




Figure 3: Epidemic curves. The plot on the left is typical of a propagated epidemic. The curve on the right is typical of a common source epidemic.

graphed over extended periods, long-term trends might be evident (as in the reported wildlife and dog rabies cases in the USA from 1946 to 1965, shown on the right in Figure 5).

1.3 Causation

In a medical context a cause is an event, condition, or characteristic without which disease cannot occur (Rothman 1976). Causes have the following characteristics: (1) they must precede the effect, (2) they can be either host or environmental factors (e.g. characteristics, conditions, actions of individuals, events, natural, social or economic phenomena), and (3) they can be either positive (the presence of an exposure) or negative (the absence of exposure, such as vaccination).

It is easiest to conceptualise causation by regarding causal factors as the pieces of a pie. Disease occurs when we have assembled enough causal factors to to produce a full pie. For some diseases (especially infectious conditions) it may be that exposure to the infectious agent will cause disease: in this situation there is only one piece to the pie. For other diseases there may be many reasons why some exposed individuals don't develop the disease yet others do: in this situation, the pie is made up of many pieces. The following descriptors are used when talking about causation:

- Component causes are conditions that are causally related to the presence of disease (the pieces of the pie). Factors such as high cholesterol, smoking, lack of exercise, genetics, and the presence of concurrent diseases are all component causes of coronary heart disease in humans.
- Sufficient causes are the set of conditions without any one of which disease would not have occurred (the whole pie). Sufficient causes are not usually a single factor but several. Accumulation of a set of sufficient causes is synonymous with occurrence (although not necessarily diagnosis) of disease.




Figure 4: Weekly hazard of foot-and-mouth disease infection for cattle holdings (solid line) and ‘other’ holdings (dashed line) in Cumbria (Great Britain) in 2001. Reproduced from Wilesmith et al. (2003).

- A necessary cause is one that must be present for the disease to occur (the most important piece of the pie). If chicken salad has been identified as sufficient causes of salmonellosis in a foodborne disease outbreak, *Salmonella* spp. would be a necessary cause of diarrhoea.

Causes operate in different ways by:

- Predisposing individuals to disease (e.g. age, sex, previous illness).
- Enabling disease to occur (e.g. low income, poor nutrition, bad housing, inadequate medical care).
- Precipitating disease (e.g. exposure to a specific infectious agent).
- Reinforcement (e.g. repeated strenuous activity may aggravate an established disease or state).
- Interaction, where the effect of two or more causes acting together is often greater than would be expected on the basis of summing the individual effects (e.g. the risk of lung cancer in subjects that smoke and who were exposed to asbestos is greater than the additive effect of the two factors considered together).

Epidemiological studies measure the relative contribution of risk factors to disease occurrence. This allows us to set priorities for disease control and opens up the possibility of preventing disease by targetting those risk factors that are most influential in determining disease outcome.

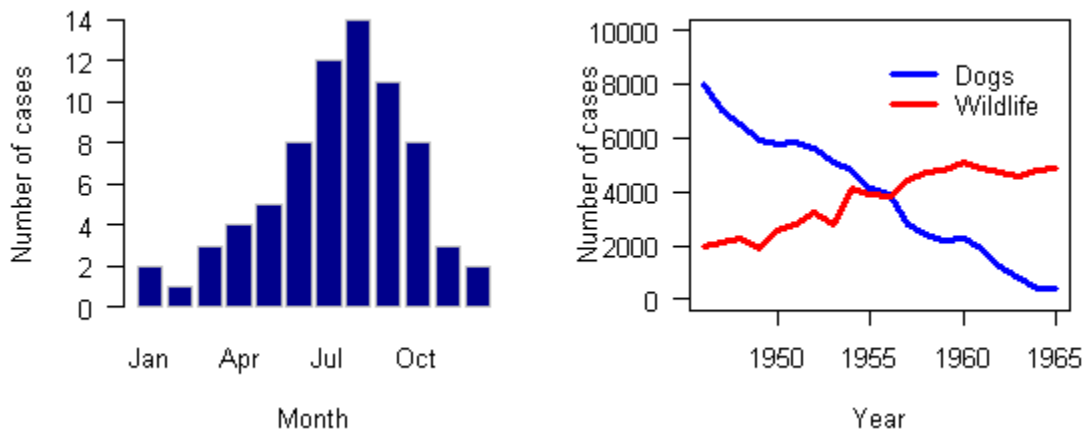


Figure 5: Temporal trends. The plot on the left shows monthly reports of human leptospirosis from 1980 - 1995. The plot on the right shows the annual number of wildlife and dog rabies cases in the USA from 1946 to 1965.

Epidemiological research usually starts by identifying **associations** between risk factors and the presence of disease. Over time, the list of associative factors can be reduced (or expanded) to a set of **causative** factors. Koch (1884) was the first to provide a framework for identifying causes of infectious disease. He specified that the following criteria (known as Koch's postulates) had to be met before an agent could be considered as the cause of a disease:

- The agent has to be present in every case of the disease.
- The agent has to be isolated and grown in pure culture.
- The agent has to cause disease when inoculated into a susceptible animal and the agent must then be able to be recovered from that animal and identified.

In the late nineteenth century Koch's postulates brought a degree of order and discipline to the study of infectious diseases, although the key assumption of 'one-agent-one-disease' was highly restrictive since it failed to take account of diseases with multiple aetiologic factors, multiple effects of single causes, carrier states, and non-agent factors (such as age and sex). Based on John Stuart Mill's rules of inductive reasoning from 1856, Evan developed a unified concept of causation which is now the generally accepted means for identifying cause-effect relationships in modern epidemiology. Evan's unified concept of causation includes the following criteria:

- The proportion of individuals with disease should be higher in those exposed to the putative cause than in those not exposed.
- Exposure to the putative cause should be more common in cases than in those without the disease.
- The number of new cases should be higher in those exposed to the putative cause than in those not exposed, as shown in prospective studies.
- Temporally, the disease should follow exposure to the putative cause.

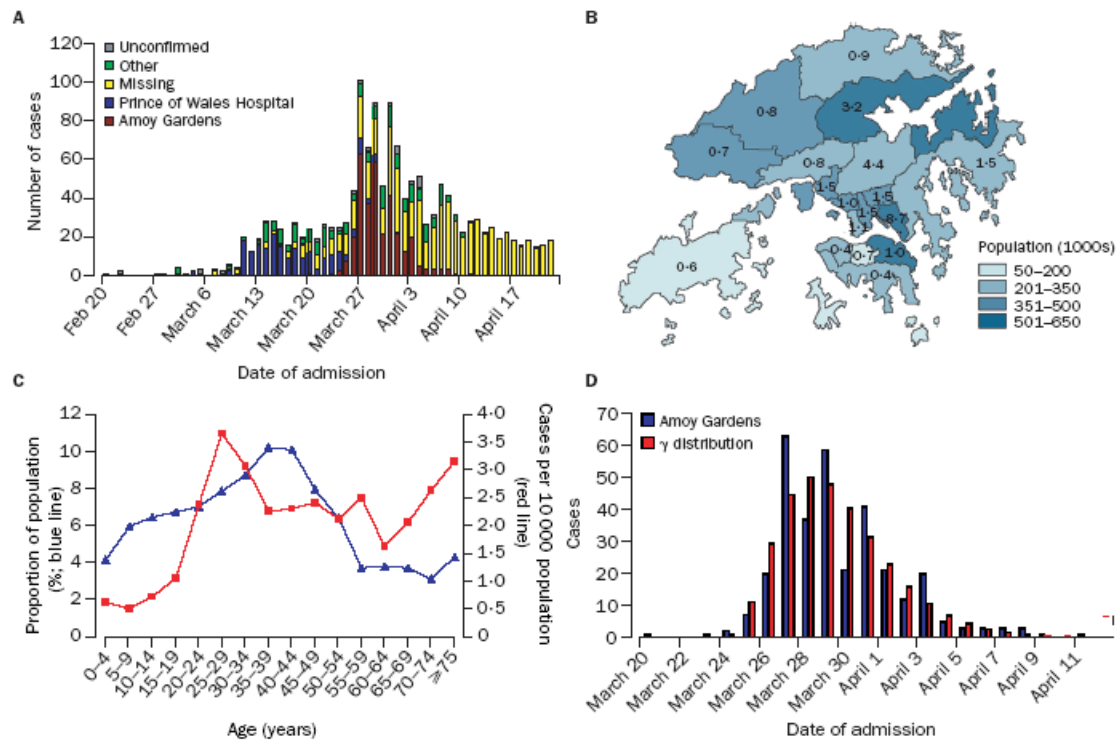


Figure 6: Descriptive epidemiology of Severe Acute Respiratory Syndrome in Hong Kong, February to April, 2003. A: Temporal pattern of SARS epidemic in Hong Kong by cluster of infection. B: Spatial distribution of population of Hong Kong and district-specific incidence (per 10 000 population) over course of epidemic to date. C: Age distribution of residents of Hong Kong and age-specific incidence (per 10 000 population) over course of epidemic to date. D: Detail of temporal pattern for Amoy Gardens cluster, according to day of admission, and fitted gamma distribution. Reproduced from Donnelly et al. (2004).

- There should be a measurable biologic spectrum of host responses.
- The disease should be reproducible experimentally.
- Preventing or modifying the host response should decrease or eliminate the expression of disease.
- Elimination of the putative cause should result in lower incidence of disease.

Bradford Hill (1965) elaborated on Evans criteria as part of work that identified smoking as a cause of lung cancer. Hill's criteria are as follows:

- Strength of association. Strong associations are more likely to be causal because they are unlikely to be due entirely to bias and confounding. Weak associations do not eliminate causation.
- Consistency. Causation is supported if the cause-effect relationship has been identified by a number of different researchers.



Figure 7: The diagram above shows a disease that has three sufficient causal complexes, each having five component causes. A is a necessary cause since it appears as a member of each sufficient cause. B, C, and F are not necessary causes since they fail to appear in all three sufficient causal sets.

- Specificity. A single exposure generally causes a single disease. This is a hold-over from the concepts of causation that were developed for infectious diseases, though there are many exceptions (e.g. smoking is associated with lung cancer as well as many other diseases). When present, specificity does provide evidence of causality, but its absence does not preclude causation.
- Temporality. A cause must always precede the effect, though sometimes this can be difficult to establish due to long induction periods and long latent (sub-clinical) periods.
- Dose-response relationship. As the level of exposure is increased, the rate of disease should also increase (though non-linear effects are common).
- Plausibility and coherence. Does a causal interpretation fit with known facts of natural history and biology of disease, including distribution in time and space and laboratory experiments (that is, does the causative relationship make ‘biological sense’?).
- Experimental evidence. Investigator-initiated interventions that modify exposure through prevention, treatment, or removal should result in less disease.
- Analogy. Has a similar relationship been observed with another exposure and/or disease? (e.g. BSE and scrapie/transmissible mink encephalopathy).

Hill’s intention was to provide a set of guidelines that could be used to determine if associations are causal, providing the following cautionary statement: ‘none of my viewpoints can bring indisputable evidence for or against the cause and effect hypothesis and none can be regarded as *sine qua non*¹.’

¹ *sine qua non*: an essential condition or element

2 Measures of health

By the end of this unit you should be able to:

- Differentiate between ratios, proportions and rates.
- Describe the terms incidence and prevalence, and use them appropriately.
- Describe the difference between risk and rate as applied to measures of incidence.

A fundamental task in epidemiological research is to quantify the occurrence of disease. This can be done by counting the number of affected individuals however, to compare levels of disease among groups of individuals, time frames and locations, we need to consider counts of cases in context of the size of the population from which those cases arose.

A **proportion** is a fraction in which the numerator is included in the denominator. Say we have a herd of 100 cattle and over a 12-month period we identify 58 diseased animals. The proportion of diseased animals is $58 \div 100 = 0.58 = 58\%$.

A **ratio** defines the relative size of two quantities expressed by dividing one (numerator) by the other (denominator). The odds of disease in our herd of 100 cattle is 58:42 or 1.4 to 1.

The term morbidity is used to refer to the extent of disease or disease frequency within a defined population. Two important measures of morbidity are **prevalence** and **incidence**. As good epidemiologists we must take care to use these terms correctly.

2.1 Prevalence

Strictly speaking, prevalence refers to the number of cases of a given disease or attribute that exists in a population at a specified time. Prevalence risk is the proportion of a population that has a specific disease or attribute at a specified point in time. Many authors use the term 'prevalence' when they really mean prevalence risk, and these notes will follow this convention.

$$\text{Prevalence} = \frac{\text{Number of existing cases}}{\text{Size of population}} \quad (1)$$

Prevalence can be interpreted as the probability of an individual from a population having a disease at a specified point in time. Two types of prevalence are reported in the epidemiological literature: (1) **point prevalence** equals the number of disease cases in a population at a single point in time (a snapshot), (2) **period prevalence** equals the point prevalence at the beginning of a study period plus the number of new cases that occurred during the remainder of the study period.

In 1944 the cities of Newburgh and Kingston, New York agreed to participate in a study of the effects of water fluoridation for prevention of tooth decay in children (Ast and Schlesinger 1956). In 1944 the water in both cities had low fluoride concentrations. In 1945, Newburgh began adding fluoride to its water — increasing the concentration ten-fold while Kingston left its supply unchanged. To assess the effect of water fluoridation on dental health, a survey was conducted among school children in both cities during the 1954 – 1955 school year. One measure of dental decay in children 6 – 9 years of age was whether at least one of a child's 12 deciduous cuspids or first or second deciduous molars was missing or had clinical or X-ray evidence of tooth decay.

Of the 216 first-grade children examined in Kingston, 192 had evidence of tooth decay. Of the 184 first-grade children examined in Newburgh 116 had evidence of tooth decay. Assuming complete survey coverage, there were 192 prevalent cases of tooth decay among first-grade children in Kingston at the time of the study. The prevalence of tooth decay was $192 \div 216 = 89$ cases per 100 children in Kingston and $116 \div 184 = 63$ cases per 100 children in Newburgh.

Reference: Ast DB, Schlesinger ER (1956). The conclusion of a ten-year study of water fluoridation. American Journal of Public Health, 46: 265-271.

2.2 Incidence

Incidence measures how frequently initially susceptible individuals become disease cases as they are observed over time. An incident case occurs when an individual changes from being susceptible to being diseased. The count of incident cases is the number of such events that occur in a defined population during a specified time period. There are two ways to express incidence: **incidence risk** and **incidence rate**.

Incidence risk

Incidence risk (also known as cumulative incidence) is the proportion of initially susceptible individuals in a population who become new cases during a defined follow-up period.

$$\text{Incidence risk} = \frac{\text{Number of incident cases}}{\text{Number of individuals initially at risk}} \quad (2)$$

The follow-up period may be arbitrarily fixed (e.g. the 5-year incidence risk of arthritis) or it may vary among individuals (e.g. the lifetime incidence risk of arthritis). In an investigation of a localised epidemic the follow-up period may be simply defined as the duration of the epidemic.

- Individuals have to be disease-free at the beginning of the observation period to be included in the numerator or denominator of this calculation.
- The time period to which the risk applies must be specified.
- Incidence risk is usually reported as the number of cases of disease per 100 head of population over a specified follow-up period.

Last year a herd of 121 cattle were tested for tuberculosis using the tuberculin test and all tested negative. This year the same 121 cattle were tested and 25 tested positive.

The incidence risk would then be 21 cases per 100 cattle for the 12-month follow-up period. We can also say that the probability of an animal becoming positive to the tuberculin test for the 12-month period was 21%.

Populations at risk can be either closed or open. A closed population has no additions during the follow-up period and no or few losses to follow-up. An open population is where individuals

are recruited (e.g. as births or purchases) and leave (e.g. as sales or deaths) throughout the follow-up period. Incidence risk can be measured directly when the population is closed and all subjects are followed for the entire study period. When the population is open incidence risk cannot usually be measured directly, but can be estimated by making one of the following adjustments to the number of individuals initially at risk (denominator):

- Number at risk = population size at the mid-point of the study period.
- Number at risk = $[N_{start} + \frac{1}{2}N_{new}] - [\frac{1}{2}N_{lost}]$
- Number at risk = $[N_{start} + \frac{1}{2}N_{new}] - [\frac{1}{2}(N_{lost} + N_{cases})]$. This approach assumes that only one case of disease is considered per individual.

Incidence rate

Incidence rate (also known as incidence density) is the number of new cases of disease that occur per unit of individual time at risk, during a defined follow-up period. The denominator of incidence rate is measured in units of animal (or person) time.

$$\text{Incidence rate} = \frac{\text{Number of incident cases}}{\text{Amount of at-risk experience}} \quad (3)$$

Because the denominator is expressed in units of animal- or person-time at risk those individuals that are withdrawn or are lost to follow-up are easily accounted-for. Consider a study of clinical mastitis in five cows over a 12-month period, as shown in Table 1.

Table 1: Hypothetical mastitis data

ID	Details	Events	Days at risk
1	Calve 01 Aug, mastitis 15 Aug, mastitis 15 Sep, mastitis 15 Oct, sold 15 Nov	3	106
2	Calve 01 Aug, mastitis 15 Nov, dry off 15 May,	1	365
3	Purchased 01 Dec, mastitis 01 Jan, Dry off 15 May	1	243
4	Calve 01 Aug, Sold 16 Nov	0	107
5	Calve 01 Oct, Died 05 Oct	0	4
Total		5	825

On the basis of the data presented in Table 1 the incidence rate of clinical mastitis for the 12-month period is 5 cases per 825 cow-days at risk (equivalent to 2.2 cases of clinical mastitis per cow-year at risk). Incidence rate:

- Accounts for individuals that enter and leave the population throughout the follow-up period (i.e. open populations).
- Can account for multiple disease events in the same individual (e.g. cow 1 in Table 1).

- Is usually reported as the number of cases of disease per animal time at risk (e.g. cases of disease per 100 cow-years at risk).

To calculate incidence rate correctly, it is necessary to record detailed information for each individual under study. When this is not possible, at-risk experience can be estimated using the same methods to estimate the size of an open population quoted earlier:

- At-risk experience = population size at the mid-point of the study period \times length of study period.
- At-risk experience = $\{[N_{start} + \frac{1}{2}N_{new}] - [\frac{1}{2}N_{lost}]\} \times$ length of study period.
- At-risk experience = $\{[N_{start} + \frac{1}{2}N_{new}] - [\frac{1}{2}(N_{lost} + N_{cases})]\} \times$ length of study period. This approach assumes that only one case of disease is considered per individual.

Gardner et al (1999) studied on-the-job back sprains and strains among 31,076 material handlers employed by a large retail merchandising chain. Payroll data for a 21-month period during 1994 – 1995 were linked with job injury claims. A total of 767 qualifying back injuries occurred during 54,845,247 working hours, yielding an incidence rate of 1.40 back injuries per 100,000 worker-hours.

Reference: Gardner LI, Landsittel DP, Nelson NA (1999). Risk factors for back injury in 31,076 retail merchandise store workers. *American Journal of Epidemiology*, 150: 825 - 833.

The distinction between incidence risk and incidence rate may be confusing. A useful analogy is to think of these measures in terms of driving a car. Incidence risk is analogous to the distance travelled by the car during a specified period of time (e.g. 120 km in two hours). Incidence rate is equivalent to the average speed of the car throughout the journey (e.g. 60 km/hour). Incidence risk is a measure of the proportion of a group who develop disease over a particular time and is therefore a function of both the underlying incidence rate (the speed the car travels) and the length of the follow-up period (the total travel time).

The relationship between prevalence and incidence

Table 2 compares the main features of the three measures of disease frequency. Figure ?? provides a worked example for calculating the various measures of disease frequency. The example is based on a herd of 10 animals which are all disease-free at the beginning of the observation period and followed for a 12-month period. Disease status is assessed at monthly intervals.

Providing incidence rate is constant throughout a follow-up period, incidence risk can be estimated from incidence rate:

- Closed population: incidence risk = incidence rate \times length of study period.
- Open population (when the study period is short): incidence risk \sim incidence rate \times length of study period.
- Open population: incidence risk = $1 - \exp(-\text{incidence rate} \times \text{length of study period})$.

Table 2: A comparison of the main features of prevalence, incidence risk, and incidence rate.

Item	Prevalence	Incidence risk	Incidence rate
Numerator	All cases counted on a single occasion	New cases occurring during a specified follow-up period	New cases occurring during a specified follow-up period
Denominator	All individuals examined - cases and non-cases	All susceptible individuals present at the start of the study	Sum of time periods during which all individuals could have developed disease
Time	Single point or period	Defined period	Measured for each individual from beginning of study until disease event
Study	Cross-sectional	Prospective cohort study	Prospective cohort study
Interpretation	Probability of having disease at a point in time	Risk of developing disease over a specified period	How quickly new cases develop over a specified period

There is a relationship between incidence and prevalence, which is only valid for closed populations. Imagine a closed population of cattle, all of whom survive for the duration of a study of 5 years. The number of new cases of disease (e.g. lameness requiring amputation of the claw) and the total number of cases at the end of the study period (the 5-year period prevalence) will be identical. If the duration of the disease is less than lifelong (e.g. an episode of lameness) then the period prevalence will be less than the 5-year incidence risk. Providing incidence rate is constant, prevalence can be estimated from incidence rate as follows:

- $\text{Prevalence} = (\text{incidence rate} \times \text{duration of disease}) \div (\text{incidence rate} \times \text{duration of disease} + 1)$.

Re-arranging we can derive an expression for disease duration:

- $\text{Duration of disease} = (\text{prevalence}) \div (\text{incidence rate} \times 1 - \text{prevalence})$.

In a herd of dairy cows the incidence rate of lameness is estimated to be 0.006 cases per cow-day at risk. The mean duration of disease is 7 days. The estimated prevalence of disease is $(0.006 \times 7) \div (0.006 \times 7 + 1) = 0.041$, that is 4.1 cases per 100 cows.

2.3 Other measures of health

Attack rates

Attack rates are usually used in outbreak situations where the period of risk is limited and all cases arising from exposure are likely to occur within the risk period. Attack rate is defined as the number of cases divided by the number of individuals exposed. ‘Attack risk’ would be a better way to describe this parameter.

Animal	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Diseased?	Months at risk
A					Disease								yes	4
B													no	12
C								Withdrawn					no	7
D		Disease											yes	1
E													no	12
F						Disease							yes	5
G											Disease		yes	10
H													no	12
I													no	12
J						Withdrawn							no	5
Total													4	80

Number of disease events: 4 Number present at start: 10
 Number of withdrawals: 2
 Number present at end: 8

Prevalence in June: 33% (3 cases in 9 animals)
 Prevalence in December: 50% (4 cases in 8 animals)

Incidence risk (accounting for withdrawals): 44% (4 cases in 9 animals)
 Incidence risk (approximate): 40% (4 cases in 10 animals)

Incidence rate (exact): 4 cases per 80 animal-months at risk
 Incidence rate (approximate): 4 cases per 84 animal-months at risk

Figure 8: Calculation of measures of disease frequency.

Secondary attack rates

Secondary attack rates are used to describe infectiousness. The assumption is that there is spread of an agent within an aggregation of individuals (e.g. a herd or a family) and that not all cases are a result of a common-source exposure. Secondary attack rates are the number of cases at the end of the study period less the number of initial (primary) cases divided by the size of the population that were initially at risk.

Mortality

Mortality risk (or rate) is an example of incidence where death is the outcome of interest. Cause-specific mortality risk is the incidence risk of fatal cases of a particular disease in the population at risk of death from that disease. The denominator includes both prevalent cases of the disease (that is, the individuals that haven't died yet) as well as individuals who are at risk of developing the disease.

Case fatality

Case fatality risk (or rate) refers to the incidence of death among individuals who develop the disease. Case fatality risk reflects the prognosis of disease among cases, while mortality reflects the burden of deaths from the disease in the population as a whole.

Proportional mortality

As its name implies, proportional mortality is simply the proportion of all deaths that are due to a particular cause for a specified population and time period:

$$\text{Proportional mortality} = \frac{\text{Number of deaths from the disease}}{\text{Number of deaths from all causes}} \quad (4)$$

2.4 Adjusted measures of health

Adjusted rates are used when we want to compare the level of disease in different populations. In human epidemiology it is common to adjust populations on the basis of age because the occurrence of many health conditions is age-dependent. Adjustment in veterinary epidemiology is not often used but factors such as age, breed, and production type (e.g. beef-dairy) would be rational choices as adjustment variables.

The age adjustment process removes differences in the age composition of two or more populations to allow comparisons between these populations to be made, independent of age structure. For example, a county's age-adjusted death rate is the weighted average of the age-specific death rates observed in that county, with the weights derived from the age distribution in an external population standard. Different standard populations have different age distributions and the choice will affect the resulting age-adjusted rate. If the age-adjusted rates for different counties are calculated with the same weights (that is, using the same population standard), the effect of any differences in the county's age distributions is removed.

There are two methods for adjusting disease rates: **direct adjustment** and **indirect adjustment**. Before adjustment, it is useful to consider stratum-specific rates of disease.

Stratum-specific rates

Stratum-specific rates report incidence for different subgroups of the population (e.g. age) rather than for the entire population. Stratum-specific rates are recommended for comparing subgroups between or within populations when rates of disease are strongly stratum dependent. For example, the incidence risk of lameness in a herd of dairy cows (i.e. a population) might be 4.0 cases per 100 cows. Stratifying disease incidence by age yields the data in Table 3.

The data in Table 3 shows a characteristic pattern. Incidence of disease is relatively high in young animals (i.e. 2 years of age), low in middle age (i.e. 3 and 4 – 8 years of age), then increases as animals get older (producing a 'bathtub' shaped distribution). If we were certain that this pattern was consistent across herds it would be valid to adjust these disease frequency estimates so one herd can be compared with another. If there was no consistent relationship

Table 3: Stratum-specific incidence risk of lameness in a herd of dairy cows.

Age group	Incidence risk (cases per 100 cows)
2 years	4.5
3 years	3.0
4 – 8 years	2.5
> 8 years	6.0
Total	4.0

between age and incidence of disease we should use stratum-specific rates to compare herds, in preference to adjustment (described below). The other thing to remember when you are comparing the frequency of disease across different populations is to ensure that the numerators and denominators (i.e. the disease events and the population at risk) are defined consistently over time and place. Look for: (1) consistency in definition of the disease event; (2) consistency of surveillance intensity within populations over time; and (3) consistency of surveillance intensity between populations over time. When comparing age-specific rates where the age categories are relatively large, it is important to consider the possibility of residual confounding by age.

Disease frequency estimates based on small numbers of events can fluctuate widely from one time period to the next for reasons other than a true change in the underlying frequency of occurrence of the event. Calculation of incidence is not recommended when there are fewer than five events in the numerator, because the calculated risk (or rate) will be unstable and have wide confidence intervals. Small counts should be included, where possible, even if the rates are not reported, so that the counts can be combined into larger totals (for example, three or five year averages) which would be more stable.

Directly and indirectly adjusted rates are recommended when making comparisons in the rates of age-related health events between different populations or for comparing trends in a given population over time. Directly and indirectly adjusted rates should be used only for the purpose of comparison. Because an adjusted rate is based on an external standard population, it does not reflect the absolute frequency of the event in a population.

Direct adjustment

With direct adjustment the observed stratum-specific rates are known and an estimated population distribution is used as the basis for adjustment. A standard population structure is typically used: if we were stratifying by sex we might say that in a standard population 50% of the total population would be allocated to the male strata and 50% to the female strata. The choice of the standard population for direct adjustment is not crucial; however, where possible it is desirable to select a standard that is demographically sensible. The directly adjusted disease count for the i^{th} strata is then:

$$\text{Directly adjusted count}_i = \text{STD } P_i \times \text{OBS } R_i \tag{5}$$

Where:

STD P_i : the size of the standard population in the i^{th} strata

OBS R_i : the observed rate in the i^{th} strata

Consider a study of leptospirosis seroprevalence in Scottish dogs, the details of which are shown in Table 4.

Table 4: Seroprevalence of leptospirosis in urban dogs, stratified by city.

City	Positive	Sampled	Seroprevalence
Edinburgh	61	260	23%
Glasgow	69	251	27%
Total	130	511	25%

The crude data suggests that Glasgow has a slightly higher seroprevalence of leptospirosis amongst its dog population. However, what about the composition of the two populations that were studied? Male dogs are known to have a higher incidence rate for leptospirosis because of their sexual behaviour, and it might be that more male dogs were sampled in Glasgow. Sex-specific prevalence estimates (Table 5) confirm the role of population structure.

Table 5: Seroprevalence of leptospirosis in urban dogs, stratified by city and sex.

City	Positive		Sampled		Seroprevalence		
	Male	Female	Male	Female	Male	Female	Total
Edinburgh	15	46	48	212	31%	22%	23%
Glasgow	53	16	180	71	29%	22%	27%
Total	68	62	228	223	30%	22%	25%

The confounding effect of sex can be removed by producing gender-adjusted prevalence estimates (Table 6). Direct adjustment involves adjusting the crude values to produce estimates which would be expected if the potentially confounding characteristics were similarly distributed in the two study populations.

Direct adjustment involves specifying the frequency of each level of a potential confounder (for example, sex) to produce a ‘standard population.’ In this example, we use a standard population comprised of 250 males and 250 females. The values for each study group are then weighted by the frequency of each level of the confounder.

The directly adjusted prevalence estimates are similar which suggests the difference between the cities is due to the different sex structures of the two populations.

Indirect adjustment

With indirect adjustment the stratum-specific rates are unknown and a known population distribution is used as the basis for adjustment. Indirect adjustment provides an estimate of the

Table 6: Directly adjusted seroprevalence of leptospirosis in urban dogs, stratified by city.

City	Positive		Sampled		Seroprevalence
	Male	Female	Male	Female	
Edinburgh	$0.31 \times 250 = 77$	$0.22 \times 250 = 55$	250	250	$(77 + 55) / 500 = 26\%$
Glasgow	$0.29 \times 250 = 72$	$0.22 \times 250 = 55$	250	250	$(72 + 55) / 500 = 25\%$
Total	$77 + 72 = 149$	$55 + 55 = 110$	500	250	$(149 + 110) / 1000 = 25\%$

expected number of cases, given the stratum-specific population size. Once we have this information it is usual to divide the observed number of disease cases by the expected number to yield a standardised morbidity/mortality ratio (SMR). The indirectly adjusted count for the i^{th} strata is:

$$\text{Indirectly adjusted count}_i = \text{STD } R_i \times \text{OBS } P_i \quad (6)$$

Where:

STD R_i : the standard rate in the i^{th} strata of the population

OBS P_i : the observed population size in the i^{th} strata

Suppose we have a study area divided into N contiguous regions labelled $i = 1, 2, 3, \dots, N$. Let $O_i = O_1, O_2, O_3, \dots, O_N$ denote the observed number of cases of disease in each region. Let $n_i = n_1, n_2, n_3, \dots, n_N$ denote the population size in each region. The expected number of disease cases per region (that is, the indirectly adjusted disease count) E_i is given by:

$$E_i = n_i \left(\frac{\sum_{i=1}^N O_i}{\sum_{i=1}^N n_i} \right) \quad (7)$$

The standardised mortality ratio (SMR) and the standard error of the SMR for the i^{th} region are given as:

$$\text{SMR}_i = \frac{O_i}{E_i} \quad (8)$$

$$\text{SE SMR}_i = \frac{\sqrt{n_i}}{E_i} \quad (9)$$

To visualise spatial variation in disease risk, it is common to plot the SMR for each region i in the form of a choropleth map (a map where areas are coloured according to the value of the outcome of interest). There are some problems with this technique however: (1) when observed disease counts in individual regions equal zero, the SMR won't allow one to differentiate disease

risk among individual regions, and (2) when the expected disease count approaches zero, SMR estimates become unstable. An alternative is to determine the Standardised Mortality Difference (SMD) as $(O_i - E_i)$. The expected value of the SMD for the i^{th} region is zero.

We know that the prevalence of a given disease throughout a country is 0.01% (0.0001). If we are presented with a region with 20,000 animals the expected number of cases of disease in this region will be $0.0001 \times 20,000 = 2$.

If the actual number of cases of disease in this region is 5, then the standardised mortality (morbidity) ratio is $5 \div 2 = 2.5$. That is, there were 2.5 times more cases of disease in this region, compared with the number of cases expected.

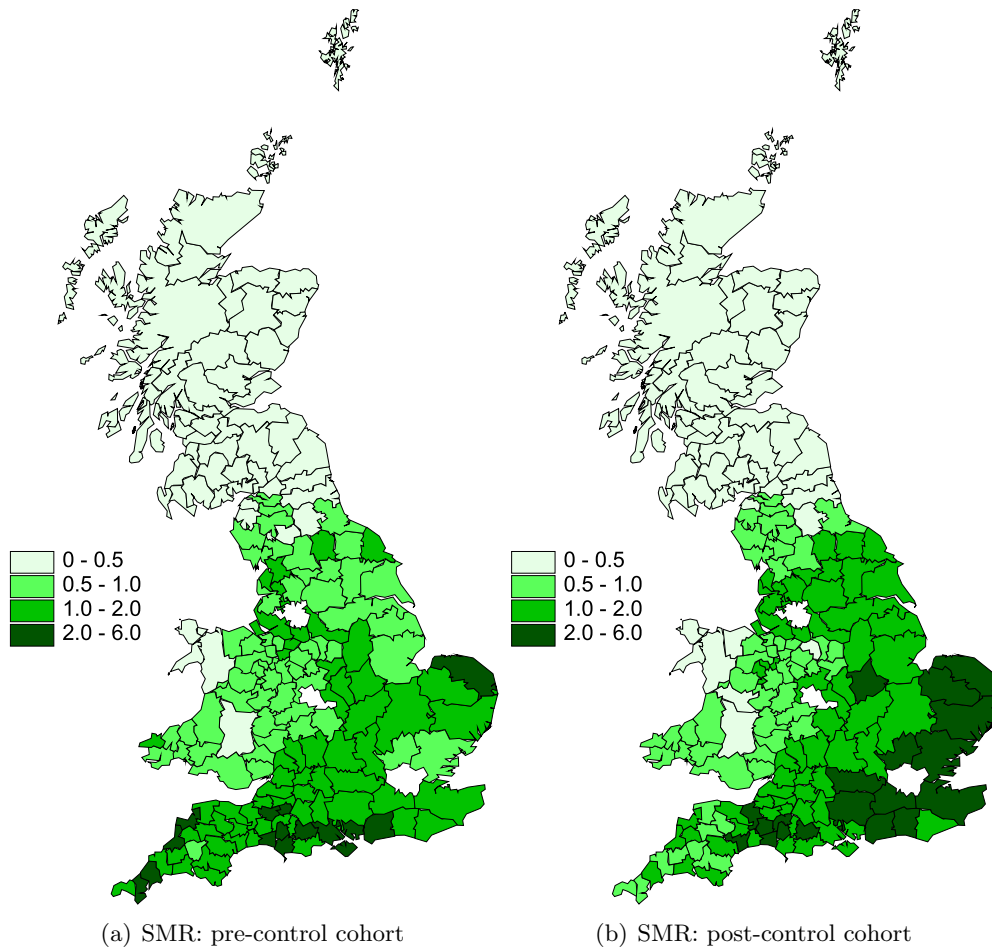


Figure 9: An example of the use of indirect standardisation used to describe the change in spatial distribution of disease risk over time. Choropleth maps of area-level standardised mortality ratios (SMRs) for bovine spongiform encephalopathy in British cattle 1986 – 1997, for (a) cattle born before the 18 July 1988 ban on feeding meat and bone meal to ruminants, and (b) cattle born between 18 July 1988 and 30 June 1997. The above maps show a shift in area-level risk over time (even though the incidence of BSE reduced markedly from 1988 to 1997). Reproduced from Stevenson et al. (2005).

3 Study design

By the end of this unit you should be able to:

- Describe the difference between descriptive and analytical epidemiological studies (giving examples of each).
- Describe the major features of the following study designs: case reports, case series, ecological studies, cross-sectional studies, cohort studies, case-control studies, clinical trials, and randomised clinical trials.
- Describe the strengths and weaknesses of cross-sectional studies, cohort studies, case-control studies, and clinical trials.

A study generally begins with a research question. Once the research question has been specified the next step is to choose a study design. A study design is a plan for selecting study subjects and for obtaining data about them. Figure 10 shows the major types of epidemiological study designs. There are three main study types: (1) descriptive studies, (2) analytical studies, and (3) experimental studies.

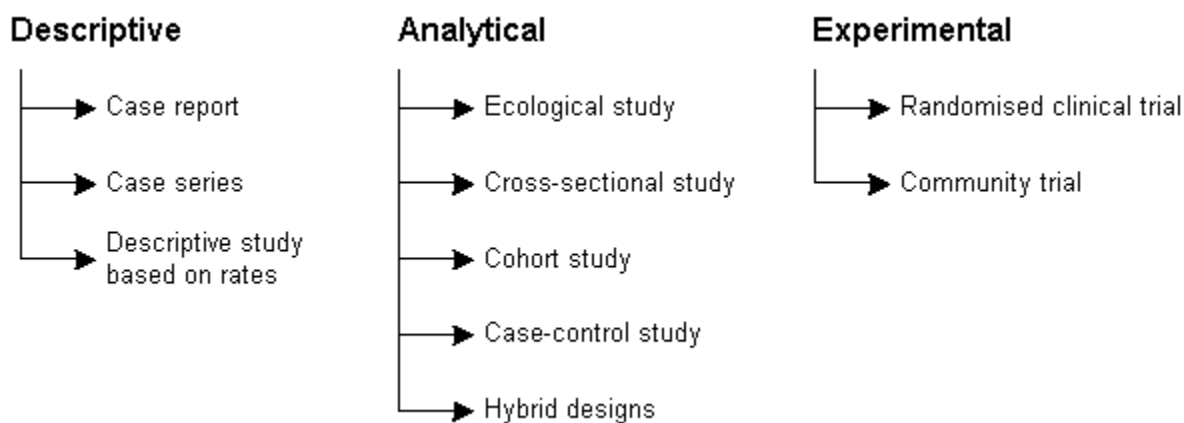


Figure 10: Tree diagram outlining relationships between the major types of epidemiologic study designs.

Descriptive studies are those undertaken without a specific hypothesis. They are often the earliest studies done on a new disease in order to characterise it, quantify its frequency, and determine how it varies in relation to individual, place and time. Analytical studies are undertaken to identify and test hypotheses about the association between an exposure of interest and a particular outcome. Experimental studies are also designed to test hypotheses between specific exposures and outcomes — the major difference is that in experimental studies the investigator has direct control over the study conditions.

3.1 Descriptive studies

The hallmark of a descriptive study is that it is undertaken without a specific hypothesis.

Case reports

A case report describes some ‘newsworthy’ clinical occurrence, such as an unusual combination of clinical signs, experience with a novel treatment, or a sequence of events that may suggest

previously unsuspected causal relationships. Case reports are generally reported as a clinical narrative.

Trivier et al (2001) reported the occurrence of fatal aplastic anaemia in an 88 year-old man who had taken clopidogrel, a relatively new drug on the market that inhibits platelet aggregation. The authors speculated that his fatal illness may have been caused by clopidogrel and wished to alert other clinicians to a possible adverse effect of the drug.

Reference: Trivier JM, Caron J, Mahieu M, Cambier N, Rose C (2001). Fatal aplastic anaemia associated with clopidogrel. *Lancet*, 357: 446.

Cases series

Whereas a case report shows that something can happen once, a case series shows that it can happen repeatedly. A case series identifies common features among multiple cases and describes patterns of variability among them.

After bovine spongiform encephalopathy (BSE) appeared in British cattle in 1987, there was concern that the disease might spread to humans. A special surveillance unit was set up to study Creutzfeld-Jacob disease (CJD), a rare and fatal progressive dementia that shares clinical and pathological features of BSE. In 1996 investigators at the unit described ten cases that met the criteria for CJD but had all occurred at unusually young ages, showed distinctive symptoms and, on pathological examination, had extensive prion protein plaques throughout the brain similar to BSE.

Reference: Will RG, Ironside JW, Zeidler M, Cousens SN, Estibeiro K, Alperovitch A et al (1996). A new variant of Creutzfeld-Jacob disease in the UK. *Lancet*, 347: 921 - 925.

Descriptive studies based on rates

Descriptive studies based on rates quantify the burden of disease on a population using incidence, prevalence, mortality or other measures of disease frequency. Most use data from existing sources (such as birth and death certificates, disease registries or surveillance systems). Descriptive studies can be a rich source of hypotheses that lead later to analytic studies.

Schwarz et al (1994) conducted a descriptive epidemiological study of injuries in a predominantly African-American part of Philadelphia. An injury surveillance system was set up in a hospital emergency centre. Denominator information came from US census data. These authors found a high incidence of intentional interpersonal injury in this area of the city.

Reference: Schwarz DF, Grisso JA, Miles CG, Holmes JH, Wishner AR, Sutton RL (1994). A longitudinal study of injury morbidity in an African-American population. *Journal of the American Medical Association*, 271: 755 - 760.

3.2 Analytical studies

Analytical studies are undertaken to test a hypothesis. In epidemiology the hypothesis typically concerns whether a certain **exposure** causes a certain **outcome** — e.g. does cigarette smoking cause lung cancer? The term exposure is used to refer to any trait, behaviour, environmental factor or other characteristic as a possible cause of disease. Synonyms for exposure are: potential risk factor, putative cause, independent variable, and predictor. The term outcome generally refers to the occurrence of disease. Synonyms for outcome are: effect, end-point, and dependent variable.

The hypothesis in an analytic study is whether an exposure actually causes an outcome (not merely whether the two are associated). Each of Hill's criteria for causation are usually required to be met to support a case for causality, but probably the most important is that exposure must precede the outcome in time.

Ecological studies

In an ecological study the unit of analysis is a group of individuals (such as counties, states, cities, or census tracts). Summary measures of exposure and summary measures of outcome are compared and inference is made at the individual level. Ecological studies are relatively quick and inexpensive to perform and can provide clues to possible associations between exposures and outcomes of interest. A major disadvantage of ecological studies is that of ecological fallacy: the assumption that an observed relationship in aggregated data will hold at the individual level.

Yang et al (1998) conducted an ecological study examining the association between chlorinated drinking water and cancer mortality among 28 municipalities in Taiwan. The investigators found a positive association between the use of chlorinated drinking water and mortality from rectal, lung, bladder, and kidney cancer.

Reference: Yang CY, Chiu HF, Cheng MF, Tsai SS (1998). Chlorination of drinking water and cancer in Taiwan. *Environmental Research*, 78: 1 - 6.

Cross-sectional studies

In a cross-sectional study a random sample of individuals from a population is taken at a point in time. Individuals included in the sample are examined for the presence of disease and their status with regard to the presence or absence of specified risk factors. Cross sectional studies commonly involve surveys to collect data. Surveys range from simple one-page questionnaires addressing a single variable, to highly complex, multiple page designs. There is a whole sub-field of epidemiology associated with design, implementation and analysis of questionnaires and surveys.




Figure 11: Schematic diagram of a cross-sectional study.

Advantages: Cross-sectional studies are relatively quick to conduct and their cost is moderate, compared with other study designs.

Disadvantages: Cross-sectional studies cannot provide information on the incidence of disease in a population — only an estimate of prevalence. Difficult to investigate cause and effect relationships.

Anderson et al (1998) studied 4,063 children aged 8 to 16 years who had participated in the National Health and Nutrition Examination Survey to assess the relationship between television watching and body-mass index. At a single examination, each child was asked a series of questions about their usual amount of television viewing. Height, weight and a series of other body measurements were taken at the same time.

Boys and girls who reported watching four or more hours of television per day had significantly greater body mass indexes than boys and girls who reported watching fewer than two hours of television per day.

Reference: Anderson RE, Crespo CJ, Bartlett SJ, Cheskin LJ, Pratt M (1998). Relationship of physical activity and television watching with body weight and level of fatness among children. Results from the Third National Health and Nutrition Examination Survey. *Journal of the American Medical Association*, 279: 938 - 942.

Cohort studies

A cohort study involves comparing disease incidence over time between groups (cohorts) that are found to differ on their exposure to a factor of interest. Cohort studies can be distinguished as either **prospective** or **retrospective** (Figure 12).

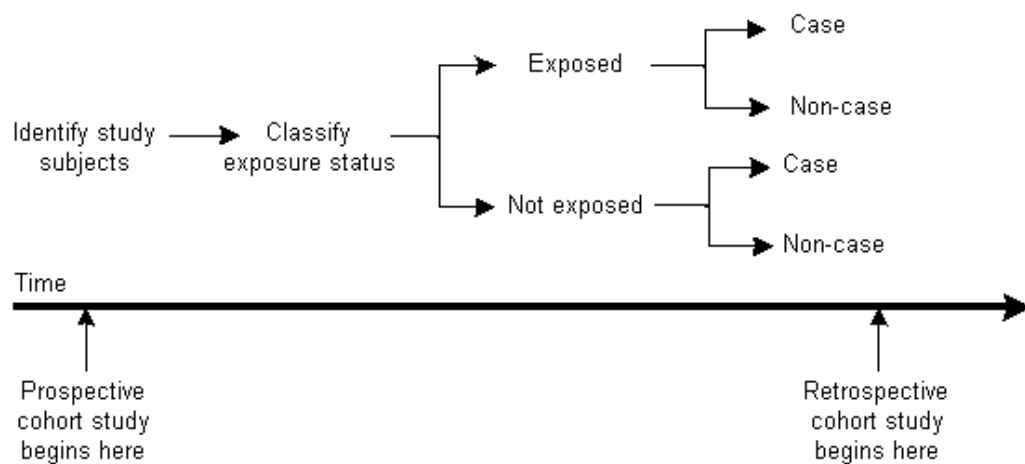


Figure 12: Schematic diagram of a prospective and retrospective cohort study.

A prospective cohort study begins with the selection of two groups of non-diseased animals, one exposed to a factor postulated to cause a disease and the other unexposed. The groups are followed over time and their change in disease status is recorded during the study period. A retrospective cohort study starts when all of the disease cases have been identified. The history of each study participant is carefully evaluated for evidence of exposure to the agent under investigation.

Advantages: Because subjects are monitored over time for disease occurrence, cohort studies provide estimates of the absolute incidence of disease in exposed and non-exposed individuals. By design, exposure status is recorded before disease has been identified. In most cases, this provides unambiguous information about whether exposure preceded disease. Cohort studies are well-suited for studying rare exposures. This is because the relative number of exposed and non-exposed persons in the study need not necessarily reflect true exposure prevalence in the population at large.

Disadvantages: Prospective cohort studies require a long follow-up period. In the case of rare diseases large groups are necessary. Losses to follow-up can become an important problem. Often quite expensive to run.

To assess the possible carcinogenic effects of radio-frequency signals emitted by cellular telephones, Johansen et al (2001) conducted a retrospective cohort study in Denmark. Two companies that operate cellular telephone networks provided names and addresses for all 522,914 of their clients for the period 1982 to 1995. The investigators matched these records to the Danish Central Population Register. After cleaning the data 420,095 cellular telephone subscribers remained and formed the exposed cohort. All other Danish citizens during the study years became the unexposed cohort. The list of exposed and unexposed individuals were then matched with the national cancer registry. The resulting data allowed calculation of cancer incidence rates.

Overall, 3,391 cancers had occurred among cellular telephone subscribers, compared with 3,825 cases expected based on age, gender, and calendar-year distribution of their person time at risk.

Reference: Johansen C, Boise J, McLaughlin J, Olsen J (2001). Cellular telephones and cancer — a nationwide cohort study in Denmark. *Journal of the National Cancer Institute*, 93: 203 - 237.

Case-control studies

In a case-control study a group of cases and non-cases ('controls') are selected and we compare the frequency of exposure factors in the cases with that of the controls. Cases are those study subjects who have developed the outcome of interest whereas controls are those who have not developed the outcome of interest at the time of selection. The key thing is that the set of controls represent a set of individuals whose exposure to the factor of interest reflects the exposure in the population from which the cases were drawn. In most situations the individual is the unit of interest, this design applies equally as well to aggregates of individuals (such as litters, pens, and herds). Figure 13 is a diagrammatic representation of the case-control design.




Figure 13: Schematic diagram of a case-control study.

The key issue when designing a case-control study is to ensure that cases and controls are similar in every way except for the exposure factors hypothesised to be associated with the disease of

interest. Controls should be drawn from the same general population as the cases — this is necessary to protect against the possible distortions from effect modifiers (termed confounders). There are three approaches that might be used to ensure that cases and controls are similar:

- **Restricted sampling.** If breed is a likely confounder you might only select one breed in the study (the dominant breed in the source population).
- **Matching.** Each case is matched with a control that has identical (or at least similar) values of the confounding variable (e.g. age and sex). This method provides direct control over known confounders and under certain conditions the efficiency of the analysis is improved. Disadvantages: (1) recruitment of suitable controls can be difficult (when it is difficult to find a suitable match); (2) you cannot quantify the effect of the matching variable on the risk of disease; (3) analysis of the data must take into account the effect of matching; (4) it is possible to overmatch, which decreases the efficiency of the study (and sometimes introduces bias).
- **Analytical control.** Multivariable regression techniques may be applied to remove the effect of known confounders.

Advantages: Case-control studies are an efficient method for studying rare diseases. Because subjects have experienced the outcome of interest at the start of the study, case-control studies are quick to run and are considerably cheaper than other study types.

Disadvantages: Case-control studies cannot provide information on the disease incidence in the studied population. The study is reliant on the quality of past records or recollection of study participants. It can also be very difficult to ensure an unbiased selection of the control group and, as a result, the representativeness of the sample selection process is difficult to guarantee.

Muscat et al (2000) sought to test the hypothesis that cellular telephone use affects the risk of brain cancer. From 1994 to 1998 at five academic medical centres in the USA they recruited 469 cases aged 18 to 80 years with newly diagnosed cancer originating in the brain. Controls (n = 422) were inpatients without brain cancer at those hospitals, excluding those with leukaemia or lymphoma. Controls were sampled to match the cases on age, sex, race and month of admission. Each case and control was then interviewed about any past subscription to a cellular telephone service. Overall 14.1% of cases and 18.0% of controls reported ever having had a subscription for a cellular telephone service. After adjusting for age, sex, race, education, study centre, and month and year of interview, the risk of developing brain cancer in a cellular telephone user was estimated to be 0.85 (95% CI 0.6 – 1.2) times as great as in a non-user. Reference: Muscat JE, Malkin MG, Thompson S, Shore RE, Stellman SD, McRee D et al. (2000). Handheld cellular telephone use and risk of brain cancer. *Journal of the American Medical Association*, 284: 3001 - 3007.

Hybrid study designs

A **nested case-control study** is similar to a cohort study with the key difference that a sample of non-cases are selected for analysis (rather than the entire cohort, as in the case of a cohort study). Figure 14 shows a diagram of a nested case-control design.

Advantages: Nested case-control studies are useful when it is either too costly or not feasible to perform additional analyses on an entire cohort (e.g. if collection of specimens and laboratory analysis of specimens is expensive). Compared with standard case-control studies, nested studies: 1) can utilise exposure and confounder data originally collected before the onset of the disease, thus reducing potential recall bias and temporal ambiguity, and 2) include cases and controls

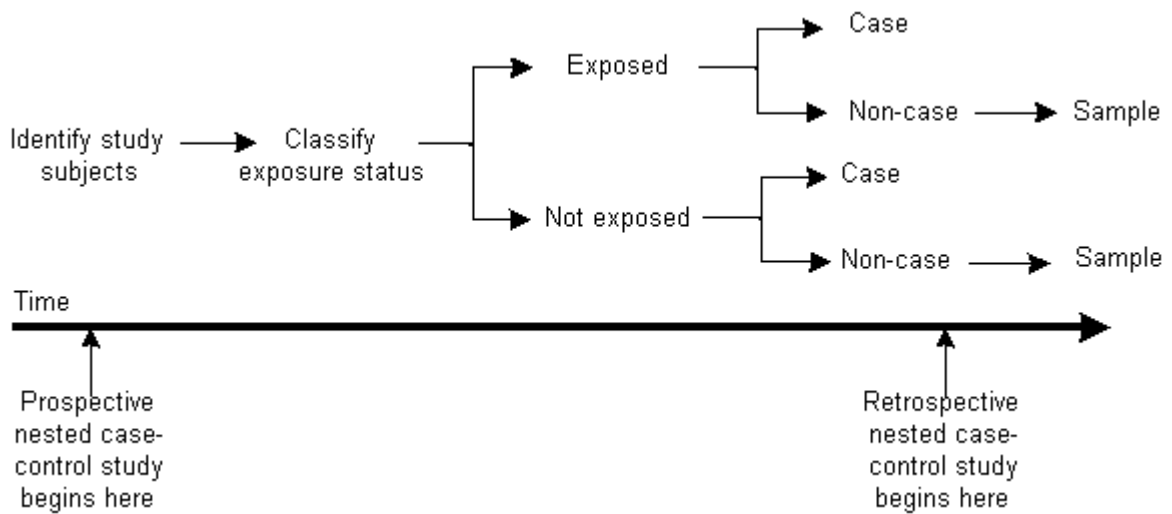


Figure 14: Schematic diagram of a nested case-control study.

drawn from the same cohort, decreasing the likelihood of selection bias. The nested case-control study is thus considered a strong observational study, comparable to its parent cohort study in the likelihood of an unbiased association between an exposure and an outcome.

Disadvantages: A concern, usually minor, is that the remaining non-diseased persons from whom the controls are selected when it is decided to do the nested study, may not be fully representative of the original cohort due to death or losses to follow-up.

To determine if *Helicobacter pylori* infection was associated with the development of gastric cancer, Parsonnet et al (1991) identified a cohort of 128,992 persons who had been followed since the mid-1960s. Of the original cohort, 189 patients developed gastric cancer. The investigators carried out a nested case-control study by selecting all of the 189 gastric cancer patients as cases and another 189 cancer-free individuals from the same cohort as controls. *H. pylori* infection status was determined using serum obtained at the beginning of the follow-up period. All total of 84% of the confirmed gastric cancer cases had been infected previously with *H. pylori*, while only 61% of the controls had been infected. This indicated a positive association between *H. pylori* infection and gastric cancer risk.

Reference: Parsonnet J, Friedman GD, Vandersteen DP, Chang Y, Vogelman JH, Orentreich N, Sibley RK (1991). *Helicobacter pylori* infection and the risk of gastric-carcinoma. *New England Journal of Medicine*, 325(16): 1127 - 1131.

In a **case-crossover study** a set of cases (subjects) is identified and a period of time before the onset of disease is selected (termed the case window) wherein the exposure to the risk factor of interest is evaluated. For each subject a second, non-overlapping time window (the control window) of the same length as the case window is selected, during which the subject did not experience the disease. This design is suitable for studying the transient effects of exposures that can vary over time and precipitate acute events (e.g. epilepsy episodes, asthma attacks). The design is efficient in that each case acts as its own control.

A study was conducted to determine if sleep disturbance was a risk factor for injury in children (Valent et al 2001). A set of cases were identified and each child was asked if their sleep was disturbed in the 24 hours before the injury occurred (the case window) and in the 24 hours before that (the control window). Among 181 boys, 40 had less than 10 hours sleep on both days; 111 had less than 10 hours on neither day; 21 had less than 10 hours only on the day before the injury; and 9 had less than 10 hours sleep on the penultimate day before the injury. The odds ratio for injury, comparing days without and with 10 hours or more sleep was 2.33 (95% CI 1.02 – 5.79).

Reference: Valent F, Brusaferrro S, Barbone F (2001) A case-crossover study of sleep and childhood injury. *Pediatrics*, 107: E23.

A **panel study** combines the features of cross-sectional and a prospective cohort designs. It can be viewed as a series of cross-sectional studies conducted on the same subjects (the panel) at successive time intervals (sometimes referred to as waves). This design allows investigators to relate changes in one variable to changes in other variables over time.

A **repeated survey** is a series of cross-sectional studies performed over time on the same study population, but each is sampled independently. Whereas panel studies follow the same individuals from survey to survey, repeated surveys follow the same study population (which may differ in composition from one survey to the next). Repeated surveys are useful for identifying overall trends in health status over time.

3.3 Experimental studies

Randomised clinical trials

The randomised clinical trial is the epidemiologic design that most closely resembles a laboratory experiment. The major objective is to test the possible effect of a therapeutic or preventive intervention. The design's key feature is that a formal chance mechanism is used to assign participants to either the treatment or control group. Subjects are then followed over time to measure one or more outcomes, such as the occurrence of disease. All things being equal, results from randomised trials offer a more solid basis for inference of cause and effect than results obtained from any other study design.

Advantages: Randomisation generally provides excellent control over confounding, even by factors that may be hard to measure or that may be unknown to the investigator.

Disadvantages: For many exposures it may not be ethical or feasible to conduct a clinical trial (e.g. exposure to pollution). Expensive. Impractical if long periods of follow-up required.

Bacterial vaginosis affects an estimated 800,000 pregnant women each year in the USA and has been found to be associated with premature birth and other pregnancy complications. To determine whether treatment with antibiotics could reduce the incidence of adverse pregnancy outcomes, Carey et al (2000) screened 29,625 pregnant women to identify 1953 who had bacterial vaginosis, met certain other eligibility criteria, and consented to participate. Women were randomly assigned to receive either: (1) two 2 gram doses of metronidazole, or (2) two doses of a similar-appearing placebo.

Bacterial vaginosis resolved in 78% of women in the treatment group, but in only 37% of women in the placebo group. Pre-term labour, postpartum infections in the mother or infant, and admission to the neonatal intensive care unit were equally common in both groups.

Reference: Carey JC, Klebanoff MA, Hauth JC, Hillier SL, Thom EA, Ernest JM et al (2000). Metronidazole to prevent preterm delivery in pregnant women with asymptomatic bacterial vaginosis. *New England Journal of Medicine*, 342: 534 - 540.

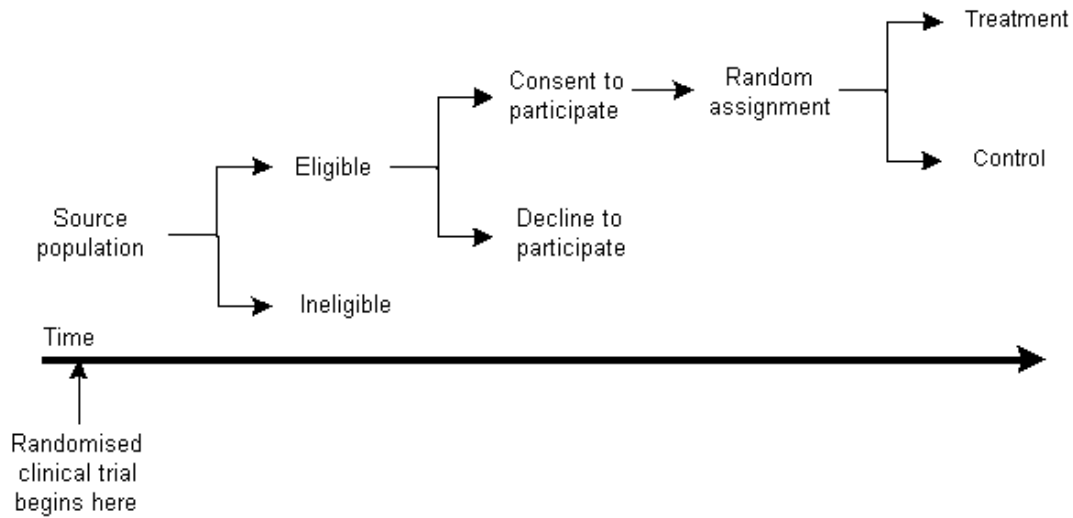


Figure 15: Schematic diagram of a randomised clinical trial.

Community trials

Instead of randomly assigning individuals to treatment or control groups, community trials assign interventions to entire groups of individuals. In the simplest situation one group (community) receives the treatment and another serves as a control.

3.4 Comparison of major the major study designs

Cohort studies involve enumeration of the denominator of the disease measure (individual time at risk) while case-control studies only sample from the denominator. Cohort studies therefore provide an estimate of incidence and risk whereas case-control studies can only estimate ratios. Prospective cohort studies provide the best evidence for the presence of cause-effect relationships, because any putative cause has to be present before disease occurs. Since these study designs are based on observation within a largely uncontrolled environment it is possible that there are still other unmeasured factors which produce cause-effect relationships that might be identified. The prospective cohort study is inefficient for studying rare diseases, which is a particular strength of the case-control study. A carefully designed cross-sectional study is more likely to be representative of the population than a case-control study.

Table 7: Comparison of the features of the cohort, case-control and cross-sectional study designs.

Criteria	Cohort	Case-control	Cross-sectional
Sampling	Separate samples of exposed and non-exposed individuals	Separate samples of diseased and non-diseased individuals	Random sample of study group
Time	Usually prospective (but may be retrospective)	Usually retrospective	Single point
Causality	Causality through evidence of temporality	Preliminary causal hypothesis	Association between disease and risk factor
Risk	Incidence risk, incidence rate	None	Prevalence
Measure of association	Risk ratio, odds ratio	Odds ratio	Risk ratio, odds ratio

4 Measures of association

By the end of this unit you should be able to:

- Given disease count data, construct a 2×2 table and explain how to calculate the following measures of association: relative risk, odds ratio, attributable rate, and attributable fraction.
- Interpret the following measures of association: relative risk, odds ratio, attributable rate, and attributable fraction.
- Describe those situations where relative risk is not a valid measure of association between exposure and outcome.

Risk is the probability that an event will happen. A characteristic or factor that influences whether or not an event occurs, is called a risk factor.

- Worn tyres are a risk factor for motor vehicle accidents.
- High blood pressure is a risk factor for coronary heart disease.
- Vaccination is a protective risk factor in that it usually reduces the risk of disease.

If we identify those risk factors that are causally associated with an increased likelihood of disease and those causally associated with a decreased likelihood of disease, then we are in a good position to make recommendations about health management. Much of epidemiological research is concerned with estimating and quantifying risk factors for disease.

Associations between putative risk factors (exposures) and an outcome (usually a disease) can be investigated using analytical observational studies. Consider a study where subjects are disease free at the start of the study and all are monitored for disease occurrence for a specified time period. If both exposure and outcome are binary variables (yes or no), the results can be presented as a 2×2 table.

	Diseased	Non-diseased	Total
Exposed	a	b	$a + b$
Non-exposed	c	d	$c + d$
Total	$a + c$	$b + d$	$a + b + c + d = n$

Based on data presented in this ‘standard’ format, various measures of association can be calculated. These fall into three main categories: (1) measures of strength, (2) measures of effect, and (3) measures of total effect. To calculate these parameters, it helps to first work out some summary parameters:

Incidence risk in the exposed population: $R_E = a/(a + b)$

Incidence risk in the non-exposed population: $R_O = c/(c + d)$

Incidence risk in the total population: $R_{total} = (a + c)/(a + b + c + d)$

Odds of disease in the exposed population: $O_E = a/b$

Odds of disease in the non-exposed population: $O_O = c/d$

Observed associations are not always causal and/or may be estimated with bias. The interpretation of the measures of association described below assumes that relationships are causal and have been estimated without bias.

4.1 Measures of strength

Risk ratio

Where incidence risk has been measured, the incidence risk ratio is defined as the ratio of the risk of disease (i.e. the incidence risk) in the exposed group to the risk of disease in the unexposed group:

$$RR = \frac{R_E}{R_O} \quad (10)$$

The incidence risk ratio provides an estimate of how many times more likely exposed individuals are to experience disease, compared with non-exposed individuals. If the incidence risk ratio equals 1, then the risk of disease in the exposed and non-exposed groups are equal. If the incidence risk ratio is greater than 1, then exposure increases the risk of disease with greater departures from 1 indicative of a stronger effect. If the incidence risk ratio is less than 1, exposure reduces the risk of disease and exposure is said to be protective. The incidence risk ratio cannot be estimated in case-control studies, as these studies do not allow calculation of risks. Odds ratios are used instead — see below.

Risk ratios range between 0 and infinity.

Incidence rate ratio

In a study where incidence rate has been measured (rather than incidence risk) the incidence rate ratio (also known as the rate ratio) can be calculated. This is the ratio of the incidence rate in the exposed group to that in the non-exposed group. The incidence rate ratio is interpreted in the same way as the risk ratio.

The term relative risk is used as a synonym for both incidence risk ratio and incidence rate ratio.

Odds ratio

The odds ratio is the odds of disease, given exposure. The odds ratio (OR) is an estimate of incidence risk ratio and is interpreted in the same way. The odds ratio is calculated as:

$$OR = \frac{O_E}{O_O} = \frac{ad}{bc} \quad (11)$$

When the number of cases of disease is low relative to the number of non-cases (i.e. the disease is rare), then the odds ratio will approximate the incidence risk ratio. If the incidence of disease is relatively low in both exposed and non-exposed individuals, then a will be small relative to b and c will be small relative to d . As a result:

$$RR = \frac{a/(a+b)}{c/(c+d)} \simeq \frac{a/b}{c/d} = \frac{ad}{bc} = OR \quad (12)$$

4.2 Measures of effect in the exposed population

Attributable risk (rate)

Attributable risk (or rate) is defined as the increase or decrease in the risk (or rate) of disease in the exposed group that is attributable to exposure. Attributable risk (unlike incidence risk ratio) measures the absolute quantity of the outcome measure that is associated with the exposure. Using the notation defined above, attributable risk (AR) is calculated as:

$$AR = R_E - R_O \quad (13)$$

In a clinical setting attributable risk may also be referred to as attributable risk reduction (ARR) or attributable risk increase (ARI) depending on whether the event risk is decreased or increased in the exposure positive (treatment) group. Another useful way of expressing attributable risk in a clinical setting is in terms of the number needed to treat (NNT). The NNT is the number of subjects who would have to be given the exposure (i.e. treatment) to prevent a negative outcome from occurring. NNT equals the inverse of the attributable risk.

A prospective cohort study was conducted to evaluate the effect of administering oxygen to patients with renal impairment prior to general anaesthesia. The incidence risk of death in oxygen treated patients was 3.5 cases per 100. The incidence risk of death in patients not receiving oxygen was 6.7 cases per 100. The attributable risk was $3.5 - 6.7 = -3.2$ cases per 100. In other words, oxygen treatment prevented death in 3.2% of patients. The NNT for these data was -31.3. This means that around 31 patients would need to be treated with oxygen to prevent one death. NNT gives a good intuitive feel for the treatment benefit and is often useful when communicating the results of such studies to patients.

Attributable fraction

Attributable fraction (also known as the attributable proportion in exposed subjects) is the proportion of disease in the exposed group that is due to exposure. Using the notation defined above, attributable fraction (AF) is calculated as:

$$AF = \frac{(R_E - R_O)}{R_E} = \frac{(RR - 1)}{RR} \quad (14)$$

For case-control studies, attributable fraction can be approximated if the incidence of disease is low:

$$AF_{est} = \frac{(O_E - O_O)}{O_E} = \frac{(OR - 1)}{OR} \quad (15)$$

In vaccine trials, vaccine efficacy is defined as the proportion of disease prevented by the vaccine in vaccinated individuals (equivalent to the proportion of disease in unvaccinated individuals due to not being vaccinated), which is the attributable fraction. A case-control study investigating the effect of oral vaccination on the presence or absence of rabies in foxes was conducted. The following results were obtained:

The odds of rabies in the unvaccinated group was 2.3 times the odds of rabies in the vaccinated group ($OR = 2.30$). Fifty six percent of rabies cases in unvaccinated foxes was due to not being vaccinated ($AF_{est} = 0.56$).

Table 8: Oral vaccination and the risk of rabies in wild foxes.

	Rabies +	Rabies -	Total
Vaccination -	18	30	48
Vaccination +	12	46	58
Total	30	76	106

4.3 Measures of effect in the total population

Population attributable risk (rate)

Population attributable risk (or rate) is the increase or decrease in risk (or rate) of disease in the population that is attributable to exposure. Using the notation defined above, population attributable risk (PAR) is calculated as:

$$PAR = R_{total} - R_O \quad (16)$$

Population attributable fraction

Population attributable fraction (also known as the aetiologic fraction) is the proportion of disease in the population that is due to the exposure. Using the notation defined above, the population attributable fraction (PAF) is calculated as:

$$PAF = \frac{(R_{total} - R_O)}{R_{total}} \quad (17)$$

Methods are available to estimate PAF using data from case-control studies. A cross sectional study investigating the relationship between dry cat food (DCF) and feline urologic syndrome (FUS) was conducted. The following results were obtained:

Table 9: Use of dry cat food and the risk of FUS in cats.

	FUS +	FUS -	Total
DCF +	13	2163	2176
DCF -	5	3349	3354
Total	18	5512	5530

The incidence risk of FUS in the DCF+ group was 5.97 cases per 1000. The incidence risk of FUS in the DCF group was 1.49 cases per 1000. The incidence risk of FUS in DCF exposed cats was 4.01 times greater than the incidence risk of FUS in DCF cats (RR = 4.0).

The incidence risk of FUS in DCF+ cats that may be attributed to DCF is 4.5 per 1000 (AR = 0.0045). In DCF+ cats 75% of FUS is attributable to DCF (AF = 0.75).

The incidence risk of FUS in the cat population that may be attributed to DCF is 1.8 per 1000. That is, we would expect the risk of FUS to decrease by 1.8 cases per 1000 if DCF were not fed

(PAR = 0.0018). Fifty-four percent of FUS cases in the cat population are attributable to DCF (PAF = 0.54).

4.4 Using the appropriate measure of effect

Table 10 outlines which measures of effect are appropriate for each of the three major study designs (case-control, cohort and cross-sectional studies).

Table 10: Epidemiologic measures of association for independent proportions in 2×2 tables.

Parameter	Case-control	Cohort	Cross-sectional
Measures of strength:			
Incidence risk ratio	No	Yes	Yes (prevalence RR)
Incidence rate ratio	No	Yes	No
Odds ratio	Yes	Yes	Yes (prevalence OR)
Measures of effect:			
Attributable risk	No	Yes	Yes
Attributable fraction	No	Yes	Yes
Attributable fraction (est)	Yes	Yes	Yes
Measures of effect in population (total effect):			
Population attributable risk	No	Yes ^a	Yes
Population attributable fraction	No	Yes ^a	Yes
Population attributable fraction (est)	Yes	Yes	Yes

^a If an estimate of the prevalence of exposure or disease incidence for the population is available from another source.

Members of the public often have a poor understanding of relative and absolute risk. A case in point was a recent news item describing the results of a study of risk factors for leukaemia in children (Draper et al. 2005). Children who lived within 200 metres of high voltage lines at birth had a 70% higher incidence risk of leukaemia compared with those that lived 600 metres or more away. While the facts were correctly reported, the interpretation of the scientific evidence was misguided. If the incidence risk of leukaemia in the general population is around 1 in 20,000 a 70% increase elevates this to around 2 cases per 20,000 (a very minor increase in absolute terms).

Pylon study sparks child health fears

Large-scale study finds increased incidence of leukaemia in children who live near high-voltage lines.

KELLY ANDREW and
NIKKI MACDONALD

FEARS that overhead power lines could increase the risk of childhood leukaemia have been heightened by a British study.

In one of the biggest studies of its kind, Oxford University researchers investigated more than 25,000 children with cancer who were born between 1962 and 1995, including 4,900 with leukaemia (a cancer of the blood), compared with a control group of healthy children.

Youngsters who lived within 200 metres of high-voltage lines at birth had a 70 per cent higher incidence of leukaemia than those 600 metres or more away. Those living between 200 metres and 600 metres from the lines were 23 per cent more likely to develop leukaemia.



Unfazed: Rosemary Greenfield says her children Jamie, 11, and Belinda, 9, have lived near a pylon all their lives and are healthy.

ions and child cancer, rather than resolving it. During the past 20 years, sci-

and Belinda, 9, for 12 years. She was not rattled by the findings of the British study. "I was a

Figure 16: Newspaper headline warning of the risk of leukaemia associated with living close to high-voltage electricity lines. Source: The Dominion Post (Wellington, New Zealand) Saturday 4 June 2005.

5 Confounding and interaction

Measures of association evaluate the relationship between exposures and outcomes. In many situations other factors may have an important influence on the relationship that is observed. If these other factors explain (or at least partially explain) the relationship between outcome and exposure then confounding is said to be present. If instead the other factors modify the relationship between exposure and the outcome, then interaction is said to be present.

5.1 Confounding

The general rule with confounding is that the association between an exposure and a given outcome is strengthened, weakened, or eliminated by a third variable or group of variables (confounders). The following criteria are used to determine whether or not a variable confounds the relationship between an exposure and an outcome:

1. The confounding variable is causally associated with the outcome; and
2. The confounding variable is noncausally associated with the exposure; and
3. The confounding variable and the exposure variable are on two separate causal pathways to the outcome; and
4. The strength of the association between the exposure and outcome changes when you account for the confounder.

An example of confounding is seen in the association between coffee consumption (the exposure) and coronary heart disease (the outcome). To assess whether or not smoking confounds the association between coffee consumption and coronary heart disease we ask the following questions:

1. Is smoking causally associated with coronary heart disease?
2. Is smoking noncausally associated with coffee drinking?
3. Is the link between smoking and heart disease and the link between coffee drinking and heart disease on two separate causal pathways?
4. Does the strength of the association between coffee drinking and heart disease change when you account for the presence of smoking?

If the answer to each of these questions is yes we conclude that smoking confounds the relationship between coffee consumption and coronary heart disease.

Figure 17 shows the relationship between three factors: exposure, confounder, and outcome. Think of the arrows that connect each of the three factors as water pipes. The direction of the arrow (which represents the association between one variable and another) can be thought of in terms of a flow of water and the strength of the association can be thought of in terms of the water pressure travelling through the pipe. Our interest is to determine the existence and amount of water pressure in the pipe linking exposure and outcome. When there is no confounding, water arising from the exposure arrives at the outcome directly (at a given pressure). When confounding is present, water arising from exposure arrives at the outcome from two sources: directly from the exposure and via the confounder. In this case, the presence of the confounder ‘dilutes’ the strength of association between the exposure and outcome (producing a drop in water pressure in the pipe from exposure to outcome).




Figure 17: A schematic diagram of the relationship between an exposure, outcome, and confounder. The confounder is causally associated with the outcome of interest and either causally or noncausally associated with the exposure. In the above diagram a unidirectional arrow indicates the association is causal; a bidirectional arrow indicates a noncausal association.

5.2 Interaction

In addition to confounding, extraneous factors can influence the effect of an exposure on an outcome. This biological phenomenon is known as interaction (also known as effect modification). Interaction refers to a difference in the effect of one exposure according to the level of another. Like confounding, interaction is due to the influence of an extraneous factor.

An example of interaction occurred in a cohort study of elderly people conducted by Fransen et al. (2002). In this study the chance of death or institutionalisation within 2 years was much higher for those who had previously suffered a hip fracture at the start of the study. The excess risk associated with hip fracture was significantly higher for men than women. This is an example of interaction between hip fracture status at study start (yes or no) and sex.

Reference: Fransen M, Woodward M, Norton R, Robinson E, Butler J, Campbell A (2002). Excess mortality or institutionalisation following hip fracture: men are at greater risk than women. *Journal of the American Geriatrics Society*, 50: 685 - 690.

There are three types of interaction:

- **Unilateralism:** where A has no effect in the absence of B, but a considerable effect in the presence of B.
- **Synergism:** where the effect of A is in the same direction, but stronger in the presence of B.
- **Antagonism:** where the effect of A works in the opposite direction when acting in the presence of B, to the direction in which it acts in the absence of B.

5.3 Identifying the presence of confounding and interaction

Confounding and interaction are different phenomena. A variable may manifest itself as either a confounder or interactive factor, as neither, or both. Different strategies are used to identify the presence of confounding and interaction in an exposure and outcome data set. In the first instance, data are stratified to calculate strata-level measures of association. If measures of association differ among strata significantly, then we conclude that interaction is present. If

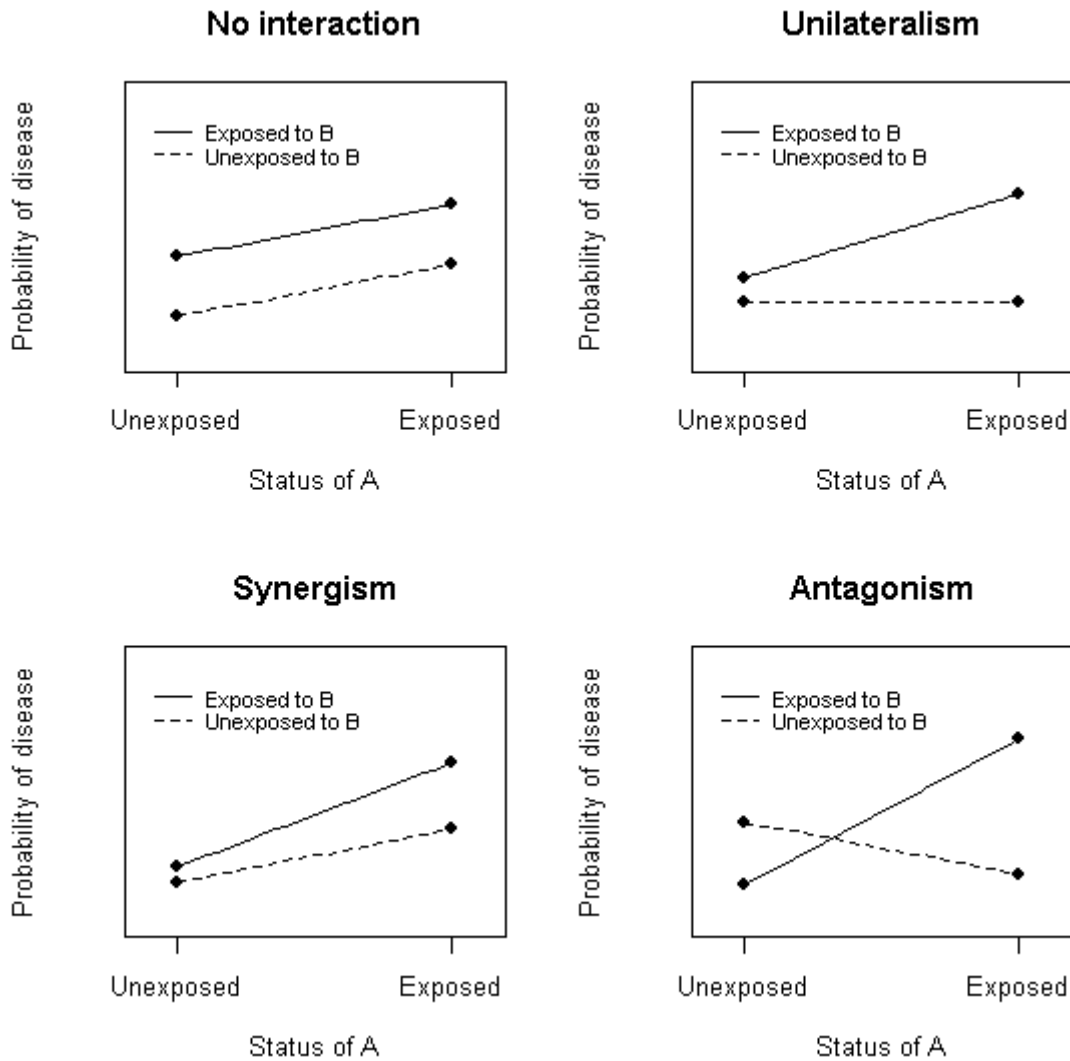


Figure 18: Interaction diagrams showing the four types of interaction that may occur between risk factors A and B.

strata-level measures of association are not significantly different, interaction is said to be absent and the presence of confounding is investigated.

To check for confounding, we compare the crude measure of association with an adjusted measure (e.g. the Mantel-Haenszel adjusted odds ratio). If there are i strata and the total number of subjects in each strata is T_i , the crude odds ratio equals:

$$OR_{\text{crude}} = \frac{\sum_i a_i \sum_i d_i}{\sum_i b_i \sum_i c_i} \quad (18)$$

The Mantel-Haenszel adjusted odds ratio equals:

$$OR_{M-H} = \frac{\sum_i \frac{a_i d_i}{T_i}}{\sum_i \frac{b_i c_i}{T_i}} \quad (19)$$

Risk ratios and rate ratios (and other measures of association) can be adjusted in a similar way. The formula for the adjusted odds ratio is provided here simply to give you an idea of what the adjustment process involves. Formulae for adjusting other measures of association are provided in many standard epidemiological texts — Elwood (1998) provides a very clear description of the approach.

If the adjustment changes the interpretation of the exposure and outcome relationship, we conclude that confounding is present. If the adjustment does not change the interpretation of the exposure-outcome relationship, we conclude that confounding is small or absent. A summary of the approach for distinguishing confounding and interaction in a data set is shown in Figure 19.

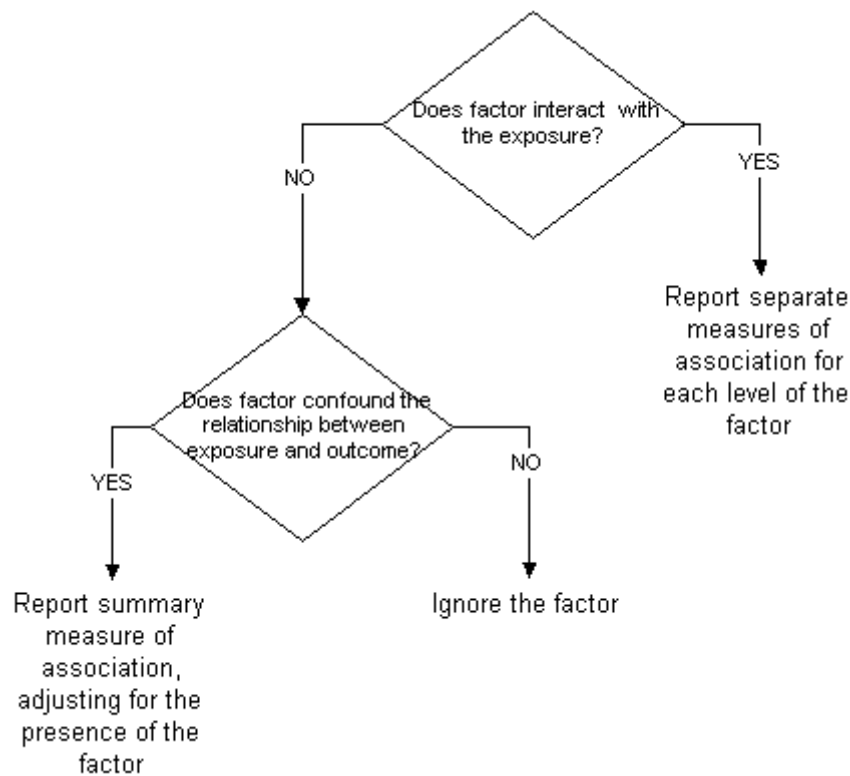


Figure 19: Strategy for distinguishing confounding and interaction.

5.4 Methods for dealing with confounding

Restriction

Think about the association between the number of children a woman has had and the risk of being diagnosed with breast cancer and the confounding effect of age. To deal with age as a confounder we could only include those subjects that were of a certain age in the study population. We could do this with either a cohort or a case-control design. Restriction is clearly an effective method, as it leaves no possibility of confounding, but obviously the disadvantage is that the study then becomes specific to a particular age group, and we cannot generalise the study beyond that target population.

Randomisation

Randomisation is an option for dealing with confounding in prospective intervention studies assessing the effects of an ethical, practical and acceptable intervention which is thought to be beneficial and not likely to be harmful. Randomisation cannot be applied to retrospective studies.

The principle of randomisation is that from a pool of study participants, subjects are randomly assigned to exposure and non-exposure groups. The definition of random is such that each subject in the study has the same chance of being allocated to a particular group, and that the chance of one individual being allocated to one group is not influenced by the allocation of any other member to the group. The advantage of randomisation is that, given large sample sizes, it is likely to produce groups which are similar even in respect to variables which have not been anticipated, designed, or measured.

Stratification

The best way to adjust for a single confounder is to examine exposure-outcome relationships within levels of the confounder. Within each of the confounder levels, there will be no confounding because exposed and non-exposed subjects will all have the same level of the confounder. If the size of the exposure-outcome association is the same at all levels (or strata) of the confounder, then statistical methods can be used to combine the stratum-specific estimates of effect to give an estimate of effect that is adjusted for the confounder. Note that the term effect modification is used to describe the situation where the exposure-outcome relationship varies according to the level of the confounder.

Matching

Matching each exposed subject to an unexposed subject with the same level of a confounder will reduce selection bias. For example, in the hypothetical smoking/heart disease study, smokers and non-smokers could be matched according to sex. When a male smoker is recruited into the study, he is matched with a male non-smoker. When a female smoker is recruited she is matched to a female non-smoker. This will obviously lead to identical percentages of men and women among smokers and non-smokers.

Matching of exposed and non-exposed subjects is only possible in studies where subjects are recruited on the basis of their exposure status. It is not possible in case control studies, where subjects are recruited according to their outcome status (presence or absence of outcome). Thus, in a case-control study of smoking and coronary heart disease, people with heart disease could be matched by sex to people without heart disease. This would result in cases and controls having the same percentage of male and females but would not lead to an even distribution of sex among smokers and non-smokers. The latter condition is the important one for control of confounding by sex. Matching is an excellent design strategy for control of confounders in cohort studies. However, it is inappropriate for this purpose in case-control studies.

The purpose of matching in case-control studies is to improve the statistical power of the study. If matching is done to improve power in a case control study, then the data analysis should take this matching into account. Otherwise, bias can be introduced into the study.

Multivariate methods

Whereas stratification is an excellent method for controlling a single confounder, multivariate methods (statistical modeling) is required if there are multiple confounders. One disadvantage of modeling as a means to control confounding is that the investigator is distanced from the mechanics of the data analysis: stratification permits a much better ‘feel’ for the data and should always precede modeling.

5.5 A worked example

Siscovick et al. (1984) conducted a case control study to evaluate the relationship between primary cardiac arrest and habitual vigorous exercise. They reported the data shown in Table 11:

Table 11: Results of a case control study evaluating the relationship between habitual exercise and primary cardiac arrest (Siscovick et al. 1984).

Non-smokers	Diseased	Non-diseased	Total
Exercise -	36	24	60
Exercise +	32	70	102
Total	68	94	162

Smokers	Diseased	Non-diseased	Total
Exercise -	40	17	57
Exercise +	25	22	47
Total	65	39	104

We would like to know: (1) if there is an interaction between smoking and vigorous habitual exercise on the risk of primary cardiac arrest, and (2) if smoking confounds the association




Figure 20: Interaction plot showing the nature of the interaction between smoking and habitual vigorous exercise on the risk of primary cardiac arrest.

between vigorous habitual exercise and primary cardiac arrest. First, we check for evidence of interaction:

The odds of primary cardiac arrest for non-smokers that did not undertake habitual vigorous exercise was 3.28 (95% CI 1.69 – 6.38) times that of those who did exercise habitually. The odds of primary cardiac arrest for smokers that did not undertake habitual vigorous exercise was 2.07 (95% CI 0.92 – 4.64) times that of those who did exercise habitually. There appears to be a synergistic interaction between smoking, exercise, and risk of primary cardiac arrest (Figure 20).

Although we have evidence to suggest the presence of interaction, we need to test the hypothesis that the strata-level odds ratios are the same using the chi-squared test of homogeneity. The test of homogeneity test statistic is compared with a chi-squared distribution with $n - 1$ degrees of freedom (where n is the number of strata). A test of homogeneity of the stratified odds ratios produces a χ^2 test statistic of 1.03. Since there are two strata the comparison has 1 degree of freedom and the associated P-value is 0.31. We accept the null hypothesis and conclude that the stratum-specific odds ratios are the same (that is, there is no significant interaction). We now use the four criteria outlined above to assess whether smoking is a confounder in the relationship between habitual vigorous exercise and primary cardiac arrest.

Is smoking causally associated with primary cardiac arrest?

	Diseased	Non-diseased	Total
Smoking +	65	39	104
Smoking -	68	194	162
Total	133	133	266

The odds of primary cardiac arrest for smokers was 4.75 (95% CI 2.93 – 7.71) times that of non smokers. We have evidence that smoking is associated with cardiac arrest. A review of the relevant literature would also support the notion that this association is causal.

Is smoking noncausally associated with habitual exercise?

	Exercise+	Exercise -	Total
Smoking +	47	57	104
Smoking -	102	60	162
Total	149	117	266

The odds of being a habitual exercise for smokers was 0.49 (95% CI 0.29 – 0.80) times that of non smokers. It is reasonable to conclude that being a smoker is noncausally associated with (lack of) habitual exercise.

Is the link between habitual exercise and cardiac arrest and smoking and cardiac arrest on two separate causal pathways? Lack of habitual exercise and smoking increase the risk of cardiac arrest by two independent physiological mechanisms. It is reasonable to assume that they are on two separate causal pathways.

Does the strength of the association between habitual exercise and primary cardiac arrest change when you account for the presence of smoking? We compare the crude odds ratio with the Mantel-Haenszel adjusted odds ratio. The odds of primary cardiac arrest in those who undertook habitual vigorous exercise was 2.99 (95% CI 1.81 – 4.95) times greater than those who did not. We apply the Mantel-Haenszel procedure to produce an adjusted odds ratio. After adjusting for smoking status, the odds of primary cardiac arrest for those that did not undertake habitual vigorous exercise was 2.72 (95% CI 1.46 – 5.06). The ratio of the crude odds ratio to the adjusted odds ratio is $2.99 \div 2.72 = 1.10$. We conclude that smoking confounds the association between habitual vigorous exercise and risk of primary cardiac arrest (using a relative difference of greater than 10% to 15% between the crude and adjusted odds ratio as an indicator of the presence of confounding).

6 Diagnostic tests

By the end of this unit you should be able to:

- Explain what is meant by the terms sensitivity and specificity, as applied to diagnostic tests.
- Given testing results presented in a 2×2 table, calculate and interpret test sensitivity and specificity.
- Given testing results presented in a 2×2 table, calculate and interpret test positive and negative predictive value.

A test may be defined as any process or device designed to detect (or quantify) a sign, substance, tissue change, or body response in an animal. Tests included:

- Routine examination of an animal or premises.
- Questions posed during history taking.
- Clinical signs.
- Laboratory findings — haematology, serology, biochemistry, histopathology.
- Post mortem findings.

If tests are to be used in a decision-making context, the selection of an appropriate test should be based on its ability to alter your assessment of the probability that a disease does or does not exist.

6.1 Screening versus diagnosis

In clinical practice, tests tend to be used in two ways:

Screening tests are those applied to apparently healthy members of a population to detect seroprevalence of certain diseases, the presence or disease agents, or subclinical disease. Usually, those animals that return a positive to such tests are subject to further in-depth diagnostic work-up, but in other cases (such as national disease control programs) the initial test result is taken as the state of nature.

Diagnostic tests are used to confirm or classify disease status, provide a guide to selection of treatment, or provide an aid to prognosis. In this setting, all animals are ‘abnormal’ and the challenge is to make a correct diagnosis.

6.2 Sensitivity and specificity

Analytic sensitivity of an assay for detecting a given chemical compound refers to the lowest concentration the test can detect. Analytic specificity refers to the capacity of the test to react to only one chemical compound.

Epidemiologic sensitivity and specificity depend on analytic sensitivity and specificity, but are entirely different concepts. Epidemiologic sensitivity answers the question: ‘Of all individuals that actually had disease X, what proportion tested positive?’ Epidemiologic specificity answers the question: ‘Of all individuals that were free of disease X, what proportion tested negative?’ Figure 21 presents this concept diagrammatically.




Figure 21: Test results measured on a continuous scale, showing the distribution of results that might be obtained for healthy and diseased individuals. The cut-off value for the test is shown by the vertical dashed line: those individuals with a result less than the cut-off value are diagnosed as non-diseased, those individuals with a result greater than the cut-off value are diagnosed as diseased. Using this diagnostic test, disease-positive individuals with a test result in the area marked 'A' will be false negatives. Disease-negative individuals with a test result in the area marked 'B' will be false positives.

6.3 Accuracy and precision

The accuracy of a test relates to its ability to give a true measure of the substance being measured. To be accurate, a test need not always be close to the true value, but if repeat tests are run, the average of the results should be close to the true value. An accurate test will not over- or under-estimate the true value. Results from tests can be 'corrected' if the degree of inaccuracy can be measured and the test results adjusted accordingly.

The precision of a test relates to how consistent the results of the test are. If a test always gives the same value for a sample (regardless of whether or not it is the correct value), it is said to be precise.

Accuracy

Assessment of test accuracy involves running the test on samples with a known quantity of substance present. These can be field samples for which the quantity of substance present has been determined by another, accepted reference procedure. Alternatively, the accuracy of a test can be determined by testing samples to which a known quantity of a substance has been added. The presence of background levels of substance in the original sample and the representativeness of these 'spiked' samples make this approach less desirable for evaluating tests designed for routine field use.

Precision

Variability among test results might be due to variability among results obtained from running the same sample within the same laboratory (repeatability) or variability between laboratories

(reproducibility). Regardless of what is being measured, evaluation of test precision involves testing the same sample multiple times within and/or among laboratories.

6.4 Test evaluation

The two key requirements of a diagnostic test are: (1) the test will detect diseased animals correctly, and (2) the test will detect non-diseased animals correctly. To work out how well a diagnostic test performs, we need to compare it with a ‘gold standard.’ A gold standard is a test or procedure that is absolutely accurate. It diagnoses all diseased animals that are tested and misdiagnoses none.

Histopathological and microbiological examination of the small intestine is generally regarded as the gold standard test for Johne’s disease in cattle. Histopathological examination of the brain stem is the gold standard test for bovine spongiform encephalopathy.

Once samples are tested using the gold standard and the test to be evaluated, a 2×2 table can be constructed, allowing test performance to be quantified. The usual format is shown in Table 12

Table 12: Diagnostic test data presented in a 2×2 table format.

	Diseased	Non-diseased	Total
Test positive	a	b	$a + b$
Test negative	c	d	$c + d$
Total	$a + c$	$b + d$	$a + b + c + d$

Sensitivity

The sensitivity of a test is defined as the proportion of subjects with disease that test positive [$p(T^+|D^+)$]. A sensitive test will rarely misclassify animals with the disease. Sensitivity is a measure of accuracy for predicting events.

$$\text{Sensitivity} = \frac{a}{(a + c)} \tag{20}$$

Sensitivity is:

- The conditional probability of a positive test, given the presence of disease.
- The likelihood of a positive test in a diseased animal.
- The proportion of animals with disease that have a positive test for the disease.
- The true positive rate (relative to all animals with disease).

Specificity

The specificity of a test is defined as the proportion of subjects without disease that test negative [$p(T^-|D^-)$]. A highly specific test will rarely misclassify animals that are not diseased.

$$\text{Specificity} = \frac{d}{(b + d)} \quad (21)$$

Specificity is:

- The conditional probability of a negative test, given the absence of disease.
- The likelihood of a negative test in an animal without disease.
- The proportion of animals without the disease that have a negative test for the disease.
- The true negative rate (relative to all animals without disease).

Sensitivity and specificity are inversely related and in the case of test results measured on a continuous scale they can be varied by changing the cut-off value. In doing so, an increase in sensitivity will often result in a decrease in specificity, and vice versa. The optimum cut-off level depends on the diagnostic strategy. If the primary objective is to find diseased animals (that is, to minimise the number of false negatives and accept a limited number of false positives) a test with a high sensitivity and good specificity is required. If the objective is to make sure that every test positive is ‘truly’ diseased (minimise the number of false positives and accept a limited number of false negatives) the diagnostic test should have a high specificity and good sensitivity.

Positive predictive value

The positive predictive value is the proportion of subjects with positive test results which have the disease.

$$\text{Positive predictive value} = \frac{a}{(a + b)} \quad (22)$$

Positive predictive value is:

- The predictive value of a positive test.
- The post-test probability of disease following a positive test.
- The posterior probability of disease following a positive test.

Negative predictive value

The negative predictive value is the proportion of subjects with negative test results which do not have the disease.

$$\text{Negative predictive values} = \frac{d}{(c + d)} \quad (23)$$

Negative predictive value is:

- The predictive value of a negative test.
- The post-test probability of no disease following a negative test.
- The posterior probability of no disease following a negative test.

Predictive values quantify the probability that a test result for a particular animal correctly identifies the condition of interest. Estimation of predictive values requires knowledge of sensitivity, specificity and the prevalence of the disease in the population. It is important to remember that predictive values are used for interpretation at the individual animal level and cannot be used to compare tests. The effect of prevalence on predictive values is considerable. Given a prevalence of disease in a population of around 30% and we are using a test with 0.95 sensitivity and 0.90 specificity, the predictive value of a positive test would be 0.80 and the predictive value of a negative test would be 0.98. If the prevalence of disease is only 3% and the test characteristics remain the same, the predictive value of a positive and negative test will be 0.23 and 0.99, respectively.

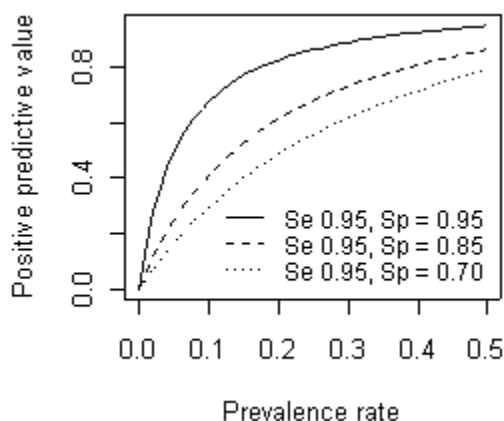


Figure 22: Relationship between prevalence and positive predictive value for tests of different sensitivities and specificities.

Remember the following general rules about diagnostic tests:

- Sensitivity and specificity are properties of a test and don't change with prevalence.
- If the prevalence increases, positive predictive value increases and negative predictive value decreases.

- If the prevalence decreases, positive predictive value decreases and negative predictive value increases.
- The more sensitive a test, the better its negative predictive value.
- The more specific a test, the better its positive predictive value.

6.5 Prevalence estimation

The estimate of disease prevalence determined on the basis of an imperfect test is called the apparent prevalence. Apparent prevalence is the proportion of all animals that give a positive test result. It can be more than, less than, or equal to the actual proportion of diseased animals, the true prevalence. If sensitivity and specificity of a test are known, true prevalence can be calculated using the Rogan and Gladen (1978) formula:

$$p(D^+) = \frac{AP - (1 - Sp)}{1 - [(1 - Sp) + (1 - Se)]} = \frac{AP + Sp - 1}{Se + Sp - 1} \quad (24)$$

Where:

AP: apparent prevalence

Se: sensitivity (0 - 1)

Sp: specificity (0 - 1)

Individual cow somatic cell counts (ICSCC) are used as a screening test for subclinical mastitis in dairy cattle. This test has a sensitivity of 0.90 and a specificity of 0.80. The apparent prevalence of mastitis in this herd using the screening test is 23 cases per 100 cows. True prevalence $p(D^+)$ may be calculated as follows:

$$AP = 0.23$$

$$Se = 0.90$$

$$Sp = 0.80$$

$$p(D^+) = (0.23 + 0.80 - 1) / (0.90 + 0.80 - 1)$$

$$p(D^+) = 0.03 / 0.70$$

$$p(D^+) = 0.04$$

The true prevalence of mastitis in this herd is 4 cases per 100 cows.

6.6 Diagnostic strategies

Clinicians commonly perform multiple tests to increase their confidence that a patient has a particular diagnosis. When multiple tests are performed and all are positive, the interpretation is straightforward: the probability of disease being present is relatively high. It is far more likely however, that some of the tests return a positive result and others will be negative. We can deal with this problem by interpreting test results in **parallel** or **series**.

Parallel interpretation

Parallel interpretation means that when multiple tests are run an individual is declared positive if at least one of the multiple tests returns a positive result. Interpreting test results in parallel increases the sensitivity and therefore the negative predictive value for a given disease prevalence. However, specificity and positive predictive value are lowered. As a consequence, if a large number of tests are performed and interpreted in this way then virtually every individual will be considered positive.

Serial interpretation

Series interpretation means that when multiple tests are run an individual is declared positive if all tests return a positive result. Series interpretation maximises specificity and positive predictive value which means that more confidence can be attributed to positive results. It reduces sensitivity and negative predictive value, and therefore it becomes more likely that diseased animals are being missed.

6.7 Screening and confirmatory testing

With a screening and confirmatory test strategy (as often used in disease control schemes) a test is applied to every animal in the population to screen the population for positives. Ideally, this test should be easy to apply and low in cost. It also should be a highly sensitive test so that it misses only a small number of diseased or infected animals. Its specificity should still be reasonable, so that the number of false positives subjected to the confirmatory test remains economically justifiable.

Individuals that return a negative result to the screening test are considered definitive negatives and not submitted to any further examination. Any animal positive to the screening test is subjected to a confirmatory test. The confirmatory test can require more technical expertise and more sophisticated equipment, and be more expensive, because it is only applied to a reduced number of samples. But it has to be highly specific, and any positive reaction to the confirmatory test is considered a definitive positive.

The same principles apply to disease control and eradication schemes. We firstly apply a test to detect disease: individuals identified as positive are removed from the population. To efficiently identify positives we need a highly sensitive test. During this early phase of a program the apparent prevalence will be higher than the true prevalence, as a consequence of test specificity being less than 1.00. As the program continues, test positive animals are identified and culled. The population prevalence of disease declines. As prevalence declines, the positive predictive value of testing declines which increases the gap between apparent and true prevalence. The proportion of false positives will then increase. At this stage a highly specific test is required. In some cases it may become necessary to use a number of tests interpreted in series to increase specificity.

SPINs and SNOUTs. SPecific tests are needed to rule a diagnosis IN, and highly SeNsitive tests are needed to rule them OUT. When a disease is rare however (with a prevalence of less than 0.01) the specificity of a test is rarely high enough to give adequate positive predictive value. Only the sensitivity is useful in the rare disease case. To remember this:

Thinking about,
SPIN and SNOUT
In cases where
Disease is rare.
Dont use SPIN,
But keep SNOUT in.

Positive and negative predictive values are more useful to the clinician than sensitivity and specificity. Although predictive values vary with prevalence, a common mistake is to assume they are fixed. This can lead to serious errors. You need to know the prevalence of disease to derive a valid estimates of positive and negative predictive value.

6.8 Likelihood ratios

Diagnostic testing is often undertaken to help us decide whether or not an individual is diseased. Because diagnostic tests are imperfect (that is, false positives and false negatives occur) we should move away from the ‘test positive = disease positive’, ‘test negative = disease negative’ paradigm and think about testing as a process that provides us with a probability estimate of the presence of disease in the tested individual. Likelihood ratios offer a means for doing this.

The likelihood ratio for a positive test tells us how likely we are to find a positive test result in a diseased individual compared with a non-diseased individual. The likelihood ratio for a positive test is estimated on the basis of dividing the probability of a particular test result in the presence of disease (sensitivity) by the probability of the test result in the absence of disease (1 - specificity). The likelihood ratio for a negative test equals (1 - sensitivity) divided by the specificity. Thus:

$$LR^+ = \frac{Se}{1 - Sp} \quad (25)$$

$$LR^- = \frac{1 - Se}{Sp} \quad (26)$$

Where:

Se: sensitivity (0 - 1)

Sp: specificity (0 - 1)

Likelihood ratios (LR) can be calculated using single cut-off values, so that one obtains only one pair of likelihood ratios, one for a positive (LR+) and another for a negative test result (LR-). More powerful information can be extracted from the diagnostic test by using multilevel likelihood ratios. In this case ranges of test results will have associated likelihood ratio values.

Likelihood ratios provide a quantitative measure of the diagnostic information contained in a particular test result. If we consider the expectation of the likelihood that an animal has a certain condition (= pre-test odds of disease) the likelihood ratio of the test multiplied by the pre-test odds gives us a revised estimate of the odds of disease (= post-test odds). This result can be re-expressed as a probability to make it more interpretable. The relationship between

odds and probability is as follows:

$$\text{Odds of event} = \frac{\text{Probability of event}}{1 - \text{Probability of event}} \quad (27)$$

$$\text{Probability of event} = \frac{\text{Odds of event}}{1 + \text{Odds of event}} \quad (28)$$

Individual cow somatic cell counts (ICSCC) are used as a screening test for sub-clinical mastitis in dairy herds. A client has a herd of dairy cows where the (true) prevalence of subclinical mastitis is estimated to be around 5%. Your herd testing authority provides you with the likelihood ratios for categories of ICSCC values in Table 13.

Table 13: Individual cow somatic cell count ranges and their (hypothetical) positive likelihood ratios.

ICSCC (cells/mL)	< 100	100 – 200	200 – 300	300 – 400	> 400
LR (+)	0.14	0.37	2.50	14.50	40.80

You are called to examine an individual cow from this herd and find that she has an ICSCC of 320,000 cells/mL. What is the probability that this cow really has mastitis? The posterior probability of mastitis in this cow is determined as follows:

1. The pre-test probability of mastitis: $50 \div 1000 = 0.05$.
2. The pre-test odds of mastitis: $0.05 \div (1 - 0.05) = 0.053$.
3. The post-test odds of mastitis given a positive test result: $\text{pre-test odds} \times \text{LR}(+) = 0.053 \times 14.5 = 0.76$.
4. The post-test probability of mastitis given a positive test result: $0.76 / (1 + 0.76) = 0.43$.

The post-test probability of a cow with a ICSCC of 320,000 cells/mL being mastitic is around 43%.

Post-test probabilities can be quickly determined in practice by using a nomogram, as shown in Figure 23. On the left hand side of the nomogram we mark the pre-test probability that the individual being examined has disease. We next identify the point defining the likelihood ratio of a positive test result along the middle scale. Finally, we draw a straight line from the pre-test probability estimate through the likelihood ratio value to the corresponding post-test probability value on the right-hand side of the chart.

A nice feature of this approach to evaluating test information is that sequential testing can be easily handled. If we are using serial interpretation, the post-test probability of disease from the first test becomes the pre-test probability for the second test.

To continue the mastitis example described above lets imagine that we examine our cow and as part of that examination we test milk from each quarter using a rapid mastitis test (RMT). We are told that the sensitivity and specificity of the RMT is 0.70 and 0.80, respectively. Our cow returns a positive result to the RMT. What now is this cow's probability of being mastitic?




Figure 23: Nomogram for post-test probability calculations using likelihood ratios of a positive test result.

The likelihood ratio of a positive RMT is $(0.70 / 1 - 0.80) = 3.5$. If the pre-test probability of disease is 43% we can use a nomogram to estimate the posterior probability of disease, given a positive test, as 72%. We are now 72% certain that this cow has mastitis.

The advantage of the likelihood ratio method of test interpretation is that we can better appreciate the value (i.e. the increase in post-test probability) provided by each diagnostic test that is applied (in the above example, ICSCC provided more information compared with the RMT). If the cost of each test applied is known the cost per unit increase in post-test probability can be determined, enabling us to be more objective in our use of diagnostic resources.




Figure 24: Diagram showing how the estimated probability of disease changes after applying a series of diagnostic tests. In our example of the cow with mastitis, we had a prior belief that the probability of the cow being mastitic was 5%. After considering the ICSCC result this probability increased to 43%. After applying a rapid mastitis test and getting a positive result, the probability of the cow having mastitis increased to 72%.

7 Sampling

By the end of this unit you should be able to:

- Explain the key features of simple random sampling, systematic random sampling, stratified random sampling, and cluster sampling.
- Describe the advantages of disadvantages of simple random sampling, systematic random sampling, stratified random sampling, and cluster sampling.
- Describe ways to reduce error when making inferences from sampled data.

Epidemiologists frequently examine populations to:

- Detect the presence of a disease;
- Demonstrate that a disease is not present within a population; and
- Establish the level of occurrence of a disease within a population.

To produce accurate estimates of disease we must be able to measure populations effectively. The exact level of disease within a population will be obtained if every individual within the population is examined (and if there was no measurement error). This technique is a census. However, in many situations a census is impossible and/or excessively expensive. Usually an accurate estimate can be obtained by examining some of the animals (a sample) from the population.

7.1 Probability sampling methods

A probability sample is one in which every element in the population has a known non-zero probability of being included in the sample.

Simple random sampling

Simple random sampling occurs when each subject in the population has an equal chance of being chosen.




Figure 25: Simple random sampling. If a sample of five cows was required, five random numbers between 1 and 10 would be generated and cows selected on the basis of the generated random numbers.

Systematic random sampling

With systematic random sampling, the selection of sampling units occurs at a predefined equal interval (known as the sampling interval). This process is frequently used when the total number of sampling units is unknown at the time of sampling (e.g. in a study where patients that enter an emergency department of a hospital on a given day are to be sampled — at the start of the study day we do not know the total number of patients seen by the end of the day).

Suppose we are studying inpatient medical records on an ongoing basis for a detailed audit. The total number of records in the population is not likely to be known in advance of the sampling since the records are to be sampled on an ongoing basis (and so it would not be possible to use simple random sampling). However, it would be possible to guess the approximate number of records that would be available per time period and to select a sample of one in every k records as they become available.

We require a total of 300 records over a 12-month period to complete the study. If there are, on average, ten new discharge records available per day then total number of records available per year is estimated to be $10 \times 365 = 3650$. To obtain the required number of records per year in the sample, the sampling interval k should be the largest integer in the quotient $3650 \div 300$. Since the value of the quotient is 12.17, the sampling interval k would be 12. Thus, we would take a sample of 1 from every 12 records.

One way to implement this procedure is to identify each record as it is created with a consecutive number. At the beginning of the study a random number between 1 and 12 is chosen as the starting point. Then, that record and every twelfth record beyond it is sampled. If the random number chosen is 4, then the records in the sample would be 4, 16, 28, 40, 52, and so on.

Stratified random sampling

Stratified sampling occurs when the sampling frame is divided into groups (strata) and a random selection within each stratum are selected. Stratified sampling is frequently undertaken to ensure that there is adequate representation of all groups in the population in the final sample. The simplest form is proportional stratified random sampling, where the number sampled within each stratum is proportional to the total number within the stratum.

Suppose that you wish to determine the prevalence of disease in the pig population of a region. Previous surveys have indicated that 70% of the regions pigs are located in very large, intensive specialised pig farms, 20% of pigs are found within smaller farming units (frequently as a secondary enterprise on large dairy farms), and 10% of pigs are kept singly within small plots around towns (by people whose major occupation is not farming). With proportional stratification, a sample would be selected at random from within each stratum such that the aggregated sample would consist of 70% pigs obtained from the large intensive farms, 20% pigs obtained from the smaller pig farms, and 10% pigs obtained from small plots near towns.

In some situations obtaining a sample from a particular stratum is more difficult or costly than for other strata. In the example described it may be more costly to sample from the pigs held in small plots around towns. This may be due to an incomplete register of smallholdings, difficulties in contacting pig owners and arranging suitable times to visit and perhaps extra travel requirements. In this situation, a technique known as non-proportional sampling may be adopted.

An advantage of stratified sampling is that the precision of parameter estimates is improved. If the population can be divided into logical strata whereby the variation within each stratum is small compared with the variation between strata a more precise estimate will be obtained.

We wish to determine average total lactation milk volume (total litres) produced by dairy cows in a region. The region contains two breeds of cattle. One breed (Friesian) is characterised by production of large volumes of milk with low concentrations of milk solids. The other breed (Jersey) is characterised by production of small volumes of milk with high concentrations of milk solids. By dividing the population into breed strata and sampling within each stratum, the average lactation milk volume production of each breed can be estimated with accuracy. The mean milk production for cows within the region can also be estimated by calculation of a weighted mean based upon each stratum mean and the stratum size.

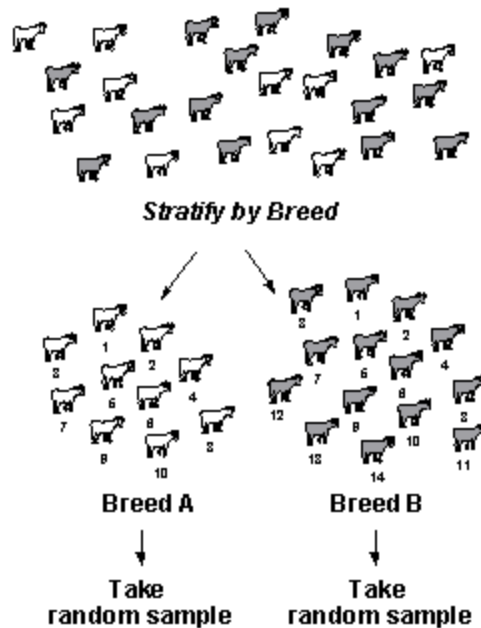


Figure 26: Stratified random sampling. A group of animals are stratified by breed and a random sample within each breed taken.

Cluster sampling

Cluster sampling occurs when the sampling frame is divided into logical aggregations (clusters) and a random selection of clusters is performed. The individual sampling units (known as primary sampling units) within the selected clusters are then examined. Clustering may occur in space or time. For example, a litter of piglets is a cluster formed within a sow, a herd of dairy cows is a cluster within a farm, and a fleet of fishing boats is a cluster formed within space (that is, a port or harbour).

Although cluster sampling has a number of advantages (including the advantage of being economical) it has the disadvantage that the standard errors of estimates are often high compared with those obtained from samples of the same number of listing units chosen by other sampling designs. The reason for this is that listing units within the same cluster are tend to be more homogenous than those listing units from different clusters. There are two types of cluster sampling:

- One stage cluster sampling occurs when clusters are selected by simple random sampling and then, once selected, all of the listing units within the cluster are examined.

- Two stage cluster sampling occurs when clusters are selected by simple random sampling and then, once selected, a random sample of listing units within each cluster are selected for examination. Estimation of population characteristics is straightforward in this situation when each cluster has the same number of listing units. Estimation of population characteristics is not straightforward when each cluster contains different numbers of listing units (in this case, you will need to consult a statistician).

The number of clusters to sample and the number of listing units within each cluster to sample will depend upon the relative variation of the factor of interest between clusters, compared with within clusters, and the relative cost of sampling clusters compared with the cost of sampling individual listing units.

- When the between-cluster variation is large relative to the within-cluster variation, you will have to sample many more clusters to get a precise estimate.
- When the between-cluster variation is small relative to the within-cluster variation, you will have to sample many more individual listing units within each cluster to get a precise estimate.

7.2 Non-probability sampling methods

Non-probability sampling occurs when the probability of selection of an individual within a population is not known and some groups within the population are more or less likely than other groups to be selected. Non-probability sampling methods include:

- Convenience sampling: where the most accessible or amenable sampling units are selected;
- Purposive sampling: where the most desired sampling units are selected; and
- Haphazard sampling: where sampling units are selected using no particular scheme or method. Inherent in this type of sampling is the problem that subconscious forces may influence the person selecting the units in an attempt to 'balance' the sample. For example, a young animal may be preferred for the next selection immediately after an older animal has been selected.

Non-probability sampling will produce biased population estimates, and the extent of that bias cannot be quantified.

7.3 Sources of error and how to reduce error

When you derive an estimate from a sample you want it to be precise and accurate. A precise estimate has confidence intervals that are small. An accurate has confidence intervals that are centred on the true population value. There are two types of error that can exist within a sample estimate: random errors and bias. The difference between random error and bias may be explained using the following diagram:




Figure 27: The distribution of bullets fired at the target on the left show little evidence of random error and bias. The distribution of the bullets fired at the centre target show a high degree of random error and a low degree of bias. The distribution of the bullets fired at the target on the right show a low degree of random error and a high degree of bias.

Random error

Random error is caused by chance. A random selection of individuals taken to make up a sample will differ slightly from each other. These differences will result in sample estimates that differ slightly from each other and also from the target population. Random error is the inherent error that arises from using a sample to make a measurement of a population. The influence of random error may be reduced by:

1. Increasing the size of the sample taken. Using the central limit theorem it can be demonstrated that a fourfold increase in sample size will result in a halving of the confidence interval.
2. Modifying the sample selection procedure to ensure that only the target group is sampled. For example, you may be interested in the performance of only one particular breed of dairy cow. You can design the study to ensure that you sample animals only from farms that contain this breed of cow. Stratified sampling is a technique that reduces sample variance by dividing the population into individual strata. Each stratum contains individuals that are similar, and so the variance within strata is less than the variation between strata. You would typically obtain samples from individual strata that have less variation than similar-sized samples obtained from the whole (unstratified) population.
3. Using an appropriate scale of measurement. Ratio estimators may result in a reduction in confidence intervals in some situations. Suppose, for example, that you wish to determine whether farmed lambs have reached the correct weight for sale. You could take a sample of lambs and estimate the average weight of the sample and from that an associated confidence interval. If the weight of lambs in the population is quite variable and you do not select a large sample it is likely that the associated confidence interval will be wide (and will include the target value). An alternative is to dichotomously classify each lambs weight within the sample with respect to the target weight (i.e. describe it as either above or below target weight). You can then calculate an estimate of the proportion of lambs that have obtained target weight (along with associated confidence intervals). You are more likely to produce narrow confidence intervals for this ratio estimate and are thus able to make a more confident decision regarding the sale of the lambs.

Bias

Bias is caused by systematic error, a systematic error being one that is inherent to the technique being used that results in a predictable and repeatable error for each observation. Bias may present itself in two ways:

1. Non-observational errors are due to inappropriate sample selection. These errors may arise from failure to include an important group of individuals within the sampling frame (resulting in their exclusion from selection), or as a result of missing data. In some situations data may be missing from a particular group of individuals within the sample.
2. Observational errors are due to inappropriate measurements. These may be attributable to false responses (i.e. participants make untrue statements) or to measurement errors.

7.4 Sampling techniques

Random sampling means that each unit of interest within the population has the same probability of selection into the sample as every other unit. The probability of selection of individual units must not differ. This is irrespective of accessibility, ease of collection or other differences that may exist between individuals. There are several important considerations to take into account before collecting a random sample:

- The target population must be defined and identified.
- A study group (sometimes called a study population) that is representative of the target population must be identified. The study group should not differ in composition from the target population.
- A sampling frame is produced. The sampling frame is a means of identifying every unit of interest (sampling unit) within the study group.
- Sampling units are selected from the sampling frame using a random (probabilistic) approach such that each sampling unit within the sampling frame has an equal chance of being selected.

Methods of randomisation

There are two principal techniques for random sampling, physical randomisation and the use of random numbers. Physical randomisation is a process where sampling units are selected using physical systems that contain random elements. These include the selection of numbered marbles from a bag, the use of a die, or the toss of a coin.

Random numbers are a sequence of numbers comprising individual digits with an equal chance that any number from 0 to 9 will be present. Tables of random numbers can be used for sample selection. Some computer programs can generate random numbers. These programs use algorithms to produce the sequence of numbers. The sequence of numbers that is generated depends upon the value chosen as the starting value for the algorithm (the seed value). Whilst

there is an equal probability that any digit from 0 to 9 will be present in a position chosen at random from the sequence, the actual digit present at each point of the sequence is determined by the seed value. In other words, the exact sequence of random numbers can be reproduced if the process is repeated using the same seed value. Computer-generated random numbers are frequently called pseudo-random numbers for this reason.

Replacement

Samples may be taken in one of two ways: sampling with replacement or sampling without replacement. In sampling with replacement, each selected unit is examined and recorded and then returned to the sampling frame. These units may then be selected into the sample again.

In sampling without replacement, each selected unit is examined and recorded and then withdrawn from the sampling frame. These units are excluded from selection into the sample again. Intuitively, sampling without replacement is the most logical — it is better to have different information from new animals as opposed to having copies of information obtained from the repeated sampling of a single animal. However, there are statistical reasons why sampling with replacement may be employed in certain circumstances. These reasons relate to the mathematics of the estimation process. In sampling with replacement the probability of selection of a unit remains the same from the first selection through to the last selection. The distribution of results within the final sample is described by the binomial distribution. In sampling without replacement, the probability of selection of the next unit changes each time a selection is made. This is due to a reduction in size of the denominator as each unit is drawn. The distribution of results is described by the (more complex) hypergeometric distribution.

The difference between the two sampling procedures is not important when samples are drawn from large populations. Often, the binomial distribution is used to approximate the hypergeometric distribution when analysing the results of samples drawn without replacement from large populations.

Selection proportional to size

Selection proportional to size is a commonly used technique for selecting clusters in one- and two-stage cluster sampling. Selection proportional to size is carried out by creating a cumulative list of cluster populations and selecting a systematic sample from a random starting point.

For example, suppose you need to take a random sample of three herds from a list of 10 herds shown Table 14. We first divide the total population (6700) by the number of herds to be selected (3) to obtain the sampling interval ($6700 \div 3 = 2233$). We next choose a random number between 1 and 2233. Suppose our chosen number is 1814. This should be fitted in position in the list to identify the first herd in the sample. Since 1814 lies between 1601 and 1900, the first selected cluster is herd 4. Next, add the sampling interval to the initial random number: $1814 + 2233 = 4047$. The next cluster to be selected is herd number 6. Add the sampling interval again: $4047 + 2233 = 6280$ and herd 10 is chosen.

This procedure leads to clusters being chosen with the selection probability being proportional to cluster size. This technique is desirable if, in addition, a constant numbers of subjects are selected within each chosen cluster. Then, overall, each subject in the population will have

Table 14: A cumulative list of herd sizes.

Herd	n	Cumulative n
1	1000	1000
2	400	1400
3	200	1600
4	300	1900
5	1200	3100
6	1000	41000
7	1600	5700
8	200	5900
9	350	6250
10	450	6700

an equal probability of being in the sample. This technique is said to be self-weighting and simplifies the analytical procedures required to make inferences about the population from the sample. If other techniques are used, a weighted analysis must be used.

Note that when this technique is used it is possible for the same cluster to be selected twice if the cluster has a population size that is greater than the sampling interval. This is unlikely to happen if the proportion of clusters selected is small, unless one cluster is very much larger than the others. If this occurs, one should select two subsamples of subjects from within the cluster. It is not valid to select another cluster instead, or to repeat the sampling procedure until no clusters are repeated, since either of these two approaches invalidates the required probabilities. If no estimate of cluster population sizes is available it will be impossible to carry out selection proportional to size and clusters must be selected by simple random sampling methods. If this is the case, responses will need to be weighted in any analyses that are undertaken. This requires a count of the total number of sampling units in each selected cluster.

7.5 Sample size

The choice of sample size involves both statistical and non-statistical considerations. Non-statistical considerations include the availability of time, money, and resources. Statistical considerations include the required precision of the estimate, and the variance expected in the data. In descriptive studies we need to specify the desired level of confidence that the estimate obtained from sampling is close to the true population value $(1 - \alpha)$. In analytical studies we may also be interested in the power $(1 - \beta)$ of the study to detect real effects.

Simple and systematic random sampling

The following formulae may be used to derive sample sizes appropriate to estimate population parameters (population total, mean, and proportion) on the basis of a simple random sample.

$$\text{Total: } n \geq \frac{z^2 SD^2}{\epsilon^2} \quad (29)$$

$$\text{Mean: } n \geq \frac{z^2 SD^2}{\epsilon^2} \quad (30)$$

$$\text{Proportion: } n \geq \frac{z^2(1 - P_y)P_y}{\epsilon^2} \quad (31)$$

Where:

z : the reliability coefficient (e.g. $z = 1.96$ for an alpha level of 0.05)

SD: the population standard deviation of the variable of interest

ϵ : the maximum absolute difference between the sample estimate and the unknown population value

P_y : the unknown population proportion

We want to estimate the mean bodyweight of deer on a farm. We anticipate the standard deviation of body weight in farmed deer of this age is around 30 kg. We would like to be 95% certain that our estimate is within 10 kg of the true mean. How many deer should we include in our sample?

$$SD^2 = 30 \times 30 = 900.$$

$$\text{Sample size} = (1.96 \times 1.96 \times 900) / (10 \times 10) = 34.$$

A sample of 34 deer are required.

We want to estimate the sero-prevalence of brucellosis in a population of cattle. The expected prevalence is 15% and we would like to take enough samples to be 95% sure that our estimate is within 20% of the actual prevalence of disease. How many cattle should be included in our sample?

$$P_y = 0.15.$$

$$\text{Relative error} = 0.20.$$

$$\text{Absolute error} = 0.20 \times 0.15 = 0.03$$

$$\text{Sample size} = [1.96 \times 1.96 \times (1 - 0.15) \times 0.15] / (0.03 \times 0.03) = 544.$$

A sample of 544 cattle are required.

Sampling to detect disease

Veterinarians are frequently asked to test groups of animals to confirm the absence of disease. The number of animals that should be tested to provide a specified level of confidence that disease is detected is given by:

$$n = (1 - \alpha^{\frac{1}{D}}) \times (N - \frac{D - 1}{2}) \quad (32)$$

Where:

N : the population size

α : 1 - confidence level (usually $\alpha = 0.05$)

D : the estimated minimum number of diseased animals in the group (that is, population size \times the minimum expected prevalence)

What is the approximate number of animals that should be tested in a herd of 200 to confirm the presence of disease if the expected prevalence is 20%?

$$N = 200$$

$$\alpha = 0.05$$

$$D = 0.20 \times 200 = 40$$

$$n = (1 - 0.05^{1/40}) \times (200 - [40 - 1] / 2)$$

$$n = 0.072 \times 180.5$$

$$n = 13$$

A minimum of 13 animals need to be tested.

8 Outbreak investigation

By the end of this unit you should be able to:

- Describe the steps to take during an outbreak investigation, including description of the outbreak by animal, place and time.
- Explain why it is important to establish a case definition when investigating a disease outbreak.
- List methods you might use to enhance surveillance once an outbreak of disease is identified

An outbreak is a series of disease events clustered in time. During an outbreak the investigator asks the questions:

- What is the problem?
- Can something be done to control it?
- Can future occurrences be prevented?

These notes outline an approach to investigating outbreaks of disease in animal populations. Although the term outbreak implies a sudden (and possibly spectacular) event (e.g. an outbreak of botulism in feedlot cattle), be aware that outbreaks can be of a more insidious nature: some causing subclinical losses in a population of animals over an extended period before being identified, characterised and investigated.

8.1 Verify the outbreak

What is the illness?

Once a suspected outbreak is identified, identifying the specific nature of the illness is an important early step. An attempt should be made to characterise cases (leading towards a formal case definition, see below). Usually it will not be possible to make a definitive diagnosis at this stage. What is required is a working definition of the disease or syndrome: for example ‘ill thrift in recently weaned calves’ or ‘sudden death in grower pigs.’

Is there a true excess of disease?

The first issue to be certain of is whether or not the outbreak is genuinely an unusual event worthy of special attention. The number of cases per unit time should be substantially greater than what is normal for the group of individuals under investigation. It is common to have owners and others concerned about a possible outbreak which is transient increase in the normal level of endemic disease.

8.2 Investigating an outbreak

Establish a case definition

A case definition is the operational definition of a disease for study purposes. A good case definition has two parts: (1) it specifies characteristics shared by all members of the class being

defined, and (2) it specifies what distinguishes them from all outside the class. A case definition ensures that the outcome of interest is consistently defined across space (e.g. among different investigation centres in a large scale outbreak) and over time.

In an outbreak of this severe and often fatal pneumonia in delegates attending the 58th annual meeting of the American Legion, Department of Pennsylvania a case was considered Legionnaires' disease if it met clinical and epidemiologic criteria. The clinical criteria required that a person have onset between 1 July and 18 August 1976, an illness characterised by cough and fever (temperature of 38.9 degrees or higher) or any fever and chest x-ray evidence of pneumonia. To meet the epidemiologic criteria, a patient either had to have attended the American Legion Convention held 21 – 24 July 1976, in Philadelphia, or had to have entered Hotel A between 1 July 1976 and the onset of illness.

Reference: Fraser DW, Tsai TR, Orenstein W, Parkin WE, Beecham HJ, Sharrar RG, Harris J, Mallison GF, Martin SM, McDade JE, Shepard CC, Brachman PS (1977). Legionnaires' disease — description of an epidemic of pneumonia. *New England Journal of Medicine*, 297:1189-1197.

Enhance surveillance

When it is suspected that an outbreak is occurring, enhanced surveillance can be useful to identify additional cases. Enhanced surveillance may involve both heightening awareness to increase passive case reports and implementing targeted surveillance. Techniques include directly contacting field practitioners by telephone, facsimile or email, via health department web pages and email discussion groups. For large outbreaks media releases (print, television, radio) can be extremely effective.

Describe outbreak according to individual, place and time

Collect historical, clinical and productivity data on those individuals that are affected (cases) and those that are not affected (non-cases). It is a mistake to concentrate exclusively on diseased animals. If possible, all cases of diseased animals should be included in the investigation. If there are large numbers of unaffected individuals you may select a representative sample of unaffected individuals for examination (controls). You may consider matching controls with some characteristic of the cases e.g. age and gender.

Plot an epidemic curve by identifying the first case (index case) and then graphing subsequent numbers of cases per day or per week from the index case through to the end of the outbreak. An extremely rapid increase in the number of cases from the index case suggests a common source epidemic (all the diseased animals were exposed to the source at about the same time). If the number of disease animals is increasing over time, this is more indicative of a propagated epidemic which is more typical of contagious disease or prolonged exposure to the agent via vectors or toxins.

Location is often an important risk factor for disease. Draw a sketch map of the area or the layout of the pens and the number of cases within pens. This includes examination of animal movements and recent additions to the herd or flock. The investigator should inspect the drawing for possible interrelationships among cases, and between location of cases and other physical features.

Develop hypotheses about the nature of exposure

At this stage, you will probably have some suspicions about what has caused the outbreak — that is, you will have started to form some hypotheses. Your next job is to test these hypotheses using the various analytical techniques described below.

Conduct analytical studies

Part of the data collection procedure above will have entailed collecting individual-level details such as age, sex, breed, date of parturition, stage of production. Individuals should be categorised according to the presence of each attribute. Attack rate tables divide the cohort of interest into exposed and non-exposed groups. Attack rates are then calculated for each exposure by dividing the number diseased by the group size (Table 15).

Table 15: Attack rate table for an outbreak of food poisoning.

Food	Exposed				Unexposed			
	Ill	Well	Total	AR (%)	Ill	Well	Total	AR (%)
Ham	36	5	41	88	2	11	13	15
Salad	40	4	44	91	9	6	15	60
Prawns	16	15	31	52	10	13	23	43

The exposure which is most likely to have served as a vehicle for an outbreak is that with the greatest difference in attack rate for exposed and unexposed individuals. An alternative is to calculate the risk ratio of disease for each exposure. Essentially this is the attack rate for the exposed individuals divided by the attack rate for unexposed individuals — the exposure with the highest risk ratio being the likely vehicle for the outbreak. It is also useful to calculate the population attributable fraction for each exposure. This will identify the percent of the risk of disease in the exposed group that is due to exposure. The closer this value is to 100% the more likely the exposure accounted for the outbreak.

8.3 Implement disease control interventions

At this stage it may be possible to produce a hypothesis regarding the cause of the outbreak. If further investigation is warranted then other epidemiological studies (case-control, prospective cohort etc) may be designed and implemented. You may also use more complex analytical techniques to analyse data already collected (multivariate techniques).

Dozens become sick from food tainted by salmonella

Three sent to hospital with severe symptoms

SCOTT ROBERTS
STAFF REPORTER

At least 75 people are sick and three are in hospital after eating a Mother's Day buffet tainted with salmonella bacteria last Sunday at the Royal Botanical Gardens in Burlington.

About 300 people attended the brunch. Health officials have only been able to contact about 170 and are expecting the number of ailing people to rise.

"This is a very high attack rate for salmonella," said Dr. Bob Nosal, medical officer of health for the Halton Region. "Right now about 40 per cent of those contacted are ill. We usually don't see numbers that high."

Nosal said this strain of salmonella was potent, causing severe diarrhea and vomiting in many victims. Other symptoms include fever, nausea, headaches and abdominal pain.

Public health investigators are trying to pinpoint how the salmonella poisoning occurred, which specific foods were involved and who was to blame. They obtained 14 food samples from the buffet, which have been sent off to the Ontario Public Health Laboratory in Toronto for testing. Results are expected in a few days.

Health officials are also administering detailed questionnaires to those who attended the brunch in hopes of finding

the cause. They are also asking for stool samples from those who became ill.

Royal Botanical Gardens contracts out its food services.

Halton Region public health officials would not disclose the name of the catering company responsible for the Mother's Day meal. Royal Botanical Gardens officials were not available for comment last night.

The Halton Region Health Department is not allowing the catering company to prepare any more buffets at the Royal Botanical Gardens until further notice.

"My understanding is that there haven't been significant problems or issues with this catering company in the recent past," Nosal said.

Figure 28: Report of an outbreak of Salmonellosis in humans arising from a contaminated buffet lunch. Source: The Globe and Mail (Toronto, Canada) Thursday 19 May 2005.

9 Appraising the literature

By the end of this unit you should be able to:

- Describe, in your own words, the four main areas that should be considered when appraising the scientific literature.
- Explain what is meant by the terms internal and external validity.
- Explain the difference between the eligible population and the study population.

Reading the literature is necessary to keep up to date with new developments and to learn more about a particular area of science that interests us.

Fortunately, there appears to be no shortage of literature available to read, and our ability to source this literature easily has been facilitated by the Internet (either in the form of peer-reviewed articles published on-line by established journals or as pre-print publications published by individuals on their own web pages). Although the Internet allows information to be widely disseminated, the quality of that information varies widely. As a result, as good scientists, we need to be discerning about what we read and (more importantly) what we believe. A systematic method of appraising (or evaluating) the literature helps us to do this. These notes outline a systematic approach to appraising the epidemiological literature, which consists of:

- Describing the evidence;
- Assessing the internal validity of the study;
- Assessing the external validity of the study; and
- Comparing the results with other available evidence.

These notes outline an approach for critically appraising epidemiological studies (i.e. those that investigate the relationship between a set of exposures and a defined outcome). An excellent series of articles providing guidelines for appraising other types of articles appeared in the *British Medical Journal* in 1997. See Greenhalgh (1997a), Greenhalgh (1997b), Greenhalgh (1997c), Greenhalgh (1997d), Greenhalgh (1997e), Greenhalgh (1997f), Greenhalgh (1997g), and Greenhalgh and Taylor (1997). Much of the technical material from these articles has been compiled into a very readable textbook on the subject by the same author (Greenhalgh 2006).

9.1 Description of the evidence

The first step in evaluating a scientific article is to understand exactly what relationship was being evaluated and what hypothesis was being tested. The reader should be able to identify the exposure variable(s) and the outcome variable. It is also necessary to categorise the study in terms of its design (survey, case-control, observational cohort, intervention cohort). Definition of the subjects that were studied in terms of source populations, the eligibility criteria, and the participation rates of the different groups that are being compared.

Having defined the topic of study, it is then useful to summarise the main result — what is the result in terms of the association between exposure and outcome? It should be possible to express the main result in a simple table and obtain from the paper the means to calculate the appropriate measure of association (risk ratio, odds ratio, difference in proportions) and the appropriate test of statistical significance.

9.2 Internal validity

Non-causal explanations

Having described the study the next step is to assess its internal validity — that is, for the subjects who were studied, does the evidence support a causal relationship between the exposure and the outcome? We consider the three possible non-causal mechanisms which could produce the observed results:

- Are the results likely to be affected by **bias**?
- Are the results likely to be affected by **confounding**?
- Are the results likely to be affected by **chance** variation?

It is useful to consider each of these aspects separately. The order of these non-causal explanations is important. If there is severe observation bias, no analytical manipulation of the data will overcome the problem. If there is confounding, then appropriate analysis will (in most cases) overcome the problem. The assessment of chance variation should be made on the main result of the study, after considering issues of bias and confounding.

Positive features of causation

Is there a correct temporal relationship? For a relationship to be causal, the putative exposure must act before the outcome occurs. In a prospective study design where exposed and non-exposed subjects are compared, this requirement is established by ensuring that subjects do not already have the outcome of interest when the study starts. The ability to clarify time relationships is weaker in retrospective studies, and care is required to ensure that possible causal factors did in fact occur before the outcome of interest.

A difficulty in all study designs, but more so in retrospective studies, is that the occurrence in biological terms of the outcome of interest may precede the recognition and documentation of that outcome by a long and variable period of time (e.g. some cancers).

Is the relationship strong? A stronger association, that is a larger risk ratio, is more likely to reflect a causal relationship. As a measured factor gets closer to a biological event on the causal pathway, the risk ratios become larger.

The fact that a relationship is strong does not protect us against certain non-causal relationships, however if the relationship that is observed is due to bias, then the bias must be large and therefore easy to identify. If a strong relationship is due to confounding, either the association of the exposure with the confounder must be very close, or the association of the confounder with the outcome must be very strong.

Is there a dose-response relationship? In some circumstances the demonstration of a smooth dose-response relationship may be a strong argument against an identified relationship arising as a result of bias. In general, we should expect uni-directional dose-effect relationships and evidence that this is not the case should be considered carefully.

Consistency of the association? A causal relationship will be expected to apply across a wide range of subjects. An association identified in one study that is consistent with the same association identified in a different groups of subjects is supportive of causation. The difficulty with consistency is that very large data sets are required to assess the similarity or otherwise of associations in different subgroups of subjects. Even with adequate numbers, the subgroups to be compared need to be defined on *a priori* grounds.

Specificity of association? It has been argued that a specific association between one causal factor and one outcome (i.e. exposure to the defined causal factor results in a specific syndrome), is good evidence for causality.

An argument against the negative health effects of smoking arose from the observation that smoking was shown to be associated with the occurrence of a number of cancers and other serious diseases and therefore demonstrated non-specificity of action, making the hypothesis of a causal link with lung cancer less likely.

Specificity may be useful, if we do not make it an absolute criterion, as one causal agent may in truth produce various outcomes, and one outcome may result from various agents. The concept is often useful in study design: as a check on response bias we may deliberately collect information on factors which we expect to be the same in groups that we are comparing (similar results across groups will indicate a lack of observation bias).

9.3 External validity

If the internal validity of a study is poor, then there is no point in proceeding further — if the results are not valid for the subjects that were studied, its application to other groups of subjects is irrelevant.

Can the results be applied to the eligible population?

The relationship between the **study population** (those that participated in the study) and the **eligible population** (those that met the study inclusion criteria but did not take part) should be well documented. Losses due to non-participation have to be considered carefully as they are likely to be non-random, and the reasons for the losses may be related to the exposure or the outcome.

Can the results be applied to the source population?

The important issue is not whether the subjects studied are ‘typical’, but whether the association between outcome and exposure given by the study participants is likely to apply to other groups. In assessing this applicability, we need to be specific about the factors which are likely to affect the association.

Most clinical trials are done on patients in teaching hospitals. If a new therapy for a particular type of neoplasia is shown to be effective in such a trial, we would readily apply the results to patients in a district hospital who had a similar stage and type of tumour and were of similar age, even though the trial patients cannot be said to be representative of district hospital patients in a general or statistical sense.

Can the results be applied to other relevant populations?

In general, the difficulties of applying results from one groups of subjects to another will be minimal for issues of basic physiology and maximal for effects in which cultural and psychosocial aspects are dominant.

9.4 Comparison of the results with other evidence

For many clinical questions a large amount of evidence is available which comes from different types of studies. In these circumstances it is useful to consider a hierarchy of evidence. Given that studies are adequately performed within the limitations of the design used, the reliability of the information from them can be ranked as follows:

1. Randomised trials.
2. Cohort and case-control studies.
3. Other comparative studies.
4. Case series, descriptive studies, clinical experience.

Randomised clinical trials, if properly performed on adequate numbers of subjects, provide greatest evidence because of the unique advantages in overcoming problems of bias and confounding.

Consistency

This is the most important characteristic used in the judgement that an association is causal. To say that the result is consistent requires that the association has been observed in a number of different studies, each of which individually can be interpreted as showing a causal explanation, and which have enough variation in their methodology and study populations to make it unlikely that the same biases or confounding factors apply in all the studies. Lack of consistency argues against causality.

Specificity

Whether a difference in results between two studies is interpreted as inconsistency or as specificity depends on whether the difference is anticipated by a hypothesis set up before the comparison is made. If not, but a plausible mechanism can be found or if the difference itself found consistently, then the hypothesis may be modified to take into account the specificity which has been shown.

Plausibility

Plausibility refers to the observed association being biologically understandable on the basis of current knowledge concerning its likely mechanisms.

However, any dramatically new observation may be in advance of current biological thinking and its lack of plausibility may reflect deficiencies in biological knowledge rather than error in observation. For example:

- John Snow effectively prevented cholera in London 25 years before the isolation of the cholera bacillus and the general acceptance of the principle that the disease could be spread by water.
- Percival Pott demonstrated the causal relationship between exposure to soot and scrotal cancer some 150 years before the relevant carcinogen was isolated.

Coherency

An association is regarded as coherent if it fits the general features of the distribution of both the exposure and the outcome under assessment; thus if lung cancer is due to smoking, the frequency of lung cancer in different populations and in different time periods should relate to the frequency of smoking in those populations at earlier relevant time periods.

If the exposure variable under study causes only a small proportion of the total disease, the overwhelming influence of other factors may make the overall pattern inconsistent.

10 Exercise: outbreak investigation

This exercise has been adapted from Gardner (1990b).

A veterinarian in a mixed practice has been investigating an ongoing diarrhoea problem in neonatal pigs in a 150-sow breeding/finishing herd. In the 12 months prior to the outbreak, 7% of litters had diarrhoea but over recent weeks the proportion of litters affected has increased to about 40%. As part of the investigation the veterinarian submitted 3 acutely affected pigs to the regional diagnostic laboratory. Of the 3 pigs, 1 was infected with *E. coli* serotype 08 but other pathogenic bacteria and viruses were not isolated from the other 2 pigs. Lesions in all 3 pigs were consistent with an acute enteritis. The veterinarian asks you to assist.

As background to the problem, the veterinarian provides you with a map showing the layout of the sheds, a description of normal management procedures, and recent records for farrowing sows as detailed below:

10.1 The problem

Shed design. The shed has 16 concrete-floored pens (oriented in a single row in a west - east direction. Pen 1 is near the entrance door at the western end of the shed and pens run in numerical sequence to pen 16 which is located near the extraction fans. The pit underneath the sows is flushed at least twice daily. During the study, pen 14 was under repair and was not used.

Management - treatments. Sows are moved into cleaned and disinfected pens in the farrowing shed on about day 110 of gestation. Sows farrow with minimal supervision. On the first day of life, pigs have their needle teeth clipped and are provided with heat lamps. No vaccines are given to sows or baby pigs for control of enteric disease. Sows are fed *ad libitum* during lactation with a high energy ration (15.5 MJ DE/kg). During gestation, they are fed about 2.0 to 2.5 kg of a lower energy ration plus about 0.5 kg/day of recycled manure for control of enteric infections and parvovirus. Piglets in litters with diarrhoea are treated with oral furazolidone and electrolytes are offered *ad libitum* in shallow bowls in each pen.

Records. Records are provided from a recent set of 26 farrowings (April 2002) for you to examine before your visit. Before April 2002 the records of diarrhoea were insufficiently detailed to be of value in the current investigation.

10.2 Diagnosis

How valid are owner-diagnoses of scours-related deaths? How could you improve their validity in the future?

10.3 Measures of disease frequency

Estimate the following rates from the data:

- The scours-specific mortality rate.
- The proportional mortality rate for scours.

- The case fatality rate for scours.
- The proportion of litters affected with scours.
- The preweaning mortality rate.

10.4 Investigation

Outline your approach to investigating this diarrhoea problem (at this stage there is no need to calculate any factor-specific rates). What initial conclusions or hypotheses did you formulate after examining the history and laboratory findings, and temporal and spatial patterns of disease?

10.5 Measures of association

Analyse the records from the 26 April farrowings and calculate some factor-specific rates or relative risks either by hand or by using computer software available for that purpose. For example:

- What was the risk ratio of scours in parity 1 litters, compared with litters from all other parities?
- What was the risk ratio of scours in litters from sick sows, compared with litters from healthy sows?
- What was risk ratio of scours in large litters, compared with small litters?
- What was the risk ratio of scours in litters born in pens 1 – 8, compared with litters born in pens 9 – 16?

Test the statistical significance of the difference between the two rates in each case. How helpful are the data in allowing you to formulate better hypotheses? Could confounding be a problem and how would you deal with it at this stage of the study?

We are interested in testing the hypothesis that the proportion of exposed individuals that are disease positive differs from the proportion of non-exposed individuals that are disease positive. Because this is nominal (count) data, a chi-squared test is the appropriate method to test this hypothesis. This involves three steps:

1. A statement of the null hypothesis: ‘The proportion of exposed individuals that are diseased does not differ from the proportion of non-exposed individuals that are diseased.’
2. Calculation of a chi-squared test statistic. Using the standard notation the formula for the chi-squared test statistic for data presented in a 2×2 table is:

$$\chi_1^2 = \frac{n(ad - bc)^2}{(a + c)(b + d)(a + b)(c + d)} \quad (33)$$

3. We will use an alpha level of 0.05 to test this hypothesis and apply a one-tailed test. Specifying an alpha level of 0.05 means that there is a 5% probability of incorrectly rejecting the null hypothesis (when it is in fact true). The critical value that separates the upper 5% of the χ^2 distribution with 1 degree of freedom from the remaining 95% is 3.841 (from statistical tables). Thus, if our calculated chi-squared test statistic is greater than 3.841 we can reject the null hypothesis and accept the alternative hypothesis, concluding that the proportions diseased among exposed and non-exposed individuals differ.

10.6 Recommendations

What recommendations, if any, would you make to your colleague and to his client based on your findings (without the data from the clinical trial or cohort study)?

10.7 Clinical trial

Design either a clinical trial or a prospective cohort study to test one of your hypotheses in detail.

10.8 Financial impact

Estimate the financial impact of the losses due to diarrhoea in this set of 26 litters. The following data has been provided:

Litter	Pen	Sow	Parity	Farrow	Born	Weaned	Death due to		
							Overlay	Scours	Other
1	9	124	1	03 Apr 02	12	9	1	2	0
2	4	121	1	03 Apr 02	9	6	1	2	0
3	12	76	3	04 Apr 02	8	8	0	0	0
4	13	164	2	05 Apr 02	11	9	0	2	0
5	16	27	6	06 Apr 02	7	7	0	0	0
6	1	18	4	09 Apr 02	10	6	0	4	0
7 ^a	7	3	2	10 Apr 02	14	8	2	2	2
8	3	69	8	10 Apr 02	10	9	1	0	0
9	11	13	5	11 Apr 02	8	8	0	0	0
10	2	101	3	12 Apr 02	12	7	2	1	2
11	8	83	6	14 Apr 02	11	10	1	0	0
12	5	79	2	15 Apr 02	11	11	0	0	0
13	10	62	4	18 Apr 02	9	8	1	0	0
14 ^a	6	74	1	18 Apr 02	10	7	0	3	0
15	4	27	1	19 Apr 02	9	6	0	3	0
16	15	61	7	23 Apr 02	6	5	1	0	0
17	12	52	5	24 Apr 02	12	10	0	0	2
18	3	107	2	26 Apr 02	15	9	4	2	0
19	16	27	3	26 Apr 02	10	9	1	0	0
20	1	159	1	27 Apr 02	6	6	0	0	0
21	13	41	2	28 Apr 02	6	6	0	0	0
22	7	131	4	29 Apr 02	8	6	0	2	0
23	9	83	6	30 Apr 02	7	6	0	0	1
24	2	79	3	30 Apr 02	9	9	0	0	0
25	8	128	5	30 Apr 02	12	10	1	1	0
26	11	169	4	30 Apr 02	11	10	0	0	1
Total					253	205	16	24	8

^a Sow sick at farrowing.

Item	Value	Target
Percent of litters with scours in 12 months before outbreak	7%	< 5%
Preweaning mortality in 12 months before outbreak	11.5%	< 12%
Post weaning mortality	5%	< 3%
Gross margin per pig marketed	\$35.00	-
Treatment costs per litter	\$10.00	-
E. coli vaccine	2 × \$2.50	-
Labour cost to vaccinate one pig	\$0.30	-

11 Review questions

11.1 Host, agent, environment

You are discharging a 2 year-old male domestic shorthair cat who has spent 10 days in your clinic recovering from the complications associated with obstruction of the urinary tract. As the cat's owner is writing out a cheque for \$1500 he asks 'will my cat experience another attack of FUS in the future and what can I do to prevent it?' What advise would you give, from an epidemiological perspective?

Think about three or four health problems or diseases that you or your friends have had. List each of the host, agent, and environmental factors that may have been causative for each disease you have listed.

Can you think of circumstances when exposure to a causal factor does not change disease incidence?

List five or six broad and fundamental influences on health and disease, that is, those influences that change the population patterns of disease.

Reflect on some medical and public health activities which were widely practiced but are now known to be wrong, some dangerously so. Your reflection should include some historical activities say, before the turn of the twentieth century and more recent ones. Also reflect on some current policies and practices that may meet the same fate.

11.2 Measures of health

Imagine you are in a country where no animal demographic data is available. An epidemic of pneumonia is suspected in the cattle population. You are asked to develop a plan to prevent and control the epidemic. Which questions do you need to answer to start a rational control strategy for this disease? Which epidemiological data do you need to answer the questions?

What benefits are there from investigating the changes in disease frequency in a population over time?

Consider the reasons why a variation in disease pattern might be artefact rather than real. Can you group them into three or four categories of explanation? What explanations can you think of for a real change in disease frequency? Can you group these into three or four categories of explanation?

Imagine you are asked to describe the health status of a population of animals to a senior public servant. The person you are talking to has no previous background in animal (or human) health. What kinds of measures would you choose to portray the health of the animal population? Consider not only the specific types of data, but also the qualities of the data you would seek out.

Imagine a population of 10,000 new army recruits. You are interested in studying the incidence and prevalence rate of gunshot wounds on war duty. Assume all gunshot wounds lead to permanent visible damage. You follow the recruits for one year. All of the study population survive, all medical records are available, and all recruits are available to interview and examination. Assume the occurrence of gunshot wounds is spread evenly through the year, and that at the

time of entering the army, no recruits had gunshot wounds. Over the year you determine that 20 recruits had a gunshot wound.

- What is the incidence risk of gunshot wounds? What is the incidence rate of gunshot wounds?
- What is the point prevalence rate of having had a gunshot wound at the beginning, middle, and end of the year?
- What is the period prevalence rate over the year?
- If the incidence rate remains the same over time, what is the prevalence rate of ever being scarred by the end of five years?
- What is the average duration of a gunshot wound, among those scarred, by the end of the first year?
- What is the estimated point prevalence rate over the five-year period?

What might be your denominator for a study defining the incidence rate of:

- Calf mortality.
- Clinical mastitis.
- Bovine spongiform encephalopathy.

11.3 Measures of association

Reflect on the terms ‘risk factor’ and ‘cause of disease.’ What is the difference between these terms?

Consider why the risk ratio might provide a false picture of the effect of a risk factor on disease and hence the strength of association.

Imagine that the incidence of chronic obstructive pulmonary disease (COPD) in horse is compared in two areas of a country: one with polluted air (A) and the other not (B). In the polluted area there were 20 cases of COPD in a population of 100,000. In the other area there were 10 cases in a population of 100,000.

- What is the risk ratio of COPD in area (A)?
- What is the risk ratio of COPD in area (B)?
- Do we know the precision of these estimates of risk ratio?
- What explanations are there for the risk ratio estimate in area (A)?
- What questions will you need to consider before concluding that there is a real association between pollution and COPD?

Imagine that exposure to a dry cat food triples the incidence of a feline urologic syndrome (FUS), that is, the risk ratio is 3. This disease has a baseline incidence of 1 per cent per year in the non-exposed group. Imagine also that the baseline incidence is double in castrated male cats (that is, 2 per cent) and that the risk ratio associated with exposure to dry cat food is the same, three. You follow 100 entire and 100 castrated male cats that are fed dry cat food, and an equivalent number of cats fed moist food. The study lasts for 5 years. Create a 2×2 table to show the data for castrates and entire male cats and calculate the odds ratio of disease in the exposed group in relation to those not exposed. Compare the odds ratio with the risk ratio of 3.

The Ministry of Health has made available a sum of \$100,000 for a health promotion programme to reduce coronary heart disease mortality. We can spend it on encouraging people to stop smoking or encouraging them to do more exercise. Assume the risk ratio associated with both risk factors is 2, that changes in prevalence rate are equally permanent, and that the cardioprotective effect occurs quickly. Which choice will give a better return in lives saved?

- First, make a judgement on which of the two preventive programs you prefer.
- Now consider which is more common: smoking or lack of exercise?
- Calculate the population attributable risk when the prevalence rate of smoking is 20%, 30%, 40% and 50% and the prevalence rate of lack of exercise is 60%, 70%, and 80% (these are realistic prevalence rates in industrialised countries). Has the result altered or substantiated your earlier judgement?

11.4 Study design

Imagine a cohort study which aims to determine the incidence of arthritis in large breeds of dogs. The follow-up period for the study is five years. Describe the advantages and disadvantages of the two approaches for measuring incidence.

Imagine a study of the incidence of congestive heart disease in large breeds of dogs, based on post mortem records collected at a University teaching hospital over a five-year period. Again, consider the advantages and disadvantages of the two approaches for measuring incidence.

Is there a difference between a clinical case series and a population case series?

How might epidemiology study the potential role in disease causation of factors which vary little between individuals within a region or country. For example: fluoride content of water, hardness or softness of water supplies, annual exposure to sunshine?

What is the essential feature that differentiates a cross-sectional study from a cohort study?

Explain what you understand by the term 'error'. What is the difference, if any, between error and bias?

11.5 Diagnostic tests

A client of yours manages a study beef herd which, for the past ten years, has consistently tested negative for tuberculosis. A positive reactor has been found after the latest round of testing. What would you advise?

12 Resources

EpiCentre, Massey University	http://epicentre.massey.ac.nz/
University of Guelph, Department of Pop Medicine	http://www.ovc.uoguelph.ca/PopMed/
Royal Veterinary College, University of London	http://www.rvc.ac.uk/
University of Michigan School of Public Health	http://www.sph.umich.edu/epid/
Epidemiology Monitor	http://www.epimonitor.net/
Association of Teachers of Veterinary Public Health	http://www.cvm.uiuc.edu/atvphpm/
Epidemiology for the uninitiated — BMJ	http://www.bmj.com/epidem/
Centers for Disease Control and Prevention	http://www.cdc.gov/
EXCITE	http://www.cdc.gov/excite/
Epidemiology Supercourse	http://www.pitt.edu/~super1/
VEIN links: Evidence Based Medicine	http://vein.library.usyd.edu.au
Post Graduate Foundation in Veterinary Science	http://www.pgf.edu.au/
EBM Resources	http://www.dartmouth.edu/~biomed/
MAF, New Zealand	http://www.maf.govt.nz
AFFA, Australia	http://www.affa.gov.au
Canadian Food Inspection Agency	http://www.inspection.gc.ca
Health Canada	http://www.hc-sc.gc.ca/
International EpiLab	http://www.dfvf.dk/
The Cochrane Collaboration	http://www.cochrane.org/index0.htm

References

- Altman, D., & Bland, J. (1994a). Statistics Notes: Diagnostic tests 1: sensitivity and specificity. *British Medical Journal*, *308*, 1552.
- Altman, D., & Bland, J. (1994b). Statistics notes: Diagnostic tests 2: predictive values. *British Medical Journal*, *309*, 102.
- Anderson, R., Crespo, C., Bartlett, S., Cheskin, L., & Pratt, M. (1998). Relationship of physical activity and television watching with body weight and level of fatness among children. Results from the Third National Health and Nutrition Examination Survey. *Journal of the American Medical Association*, *279*, 938 - 942.
- Ast, D., & Schlesinger, E. (1956). The conclusion of a ten-year study of water fluoridation. *American Journal of Public Health*, *46*, 265 - 271.
- Carey, J., Klebanoff, M., Hauth, J., Hillier, S., Thom, E., & Ernest, J. (2000). Metronidazole to prevent preterm delivery in pregnant women with asymptomatic bacterial vaginosis. *New England Journal of Medicine*, *342*, 534 - 540.
- Dawson-Saunders, B., & Trapp, R. (1994). *Basic and Clinical Biostatistics*. New York: Appleton and Lange.
- Deeks, J., & Altman, D. (2004). Statistics Notes: Diagnostic tests 4: likelihood ratios. *British Medical Journal*, *329*, 168 - 169.
- Dohoo, I., Martin, S., & Stryhn, H. (2003). *Veterinary Epidemiologic Research*. Charlottetown, Prince Edward Island, Canada: AVC Inc.
- Donnelly, C., Ghani, A., Leung, G., Hedley, A., Fraser, C., Riley, S., et al. (2004). Epidemiological determinants of spread of causal agent of severe acute respiratory syndrome in Hong Kong. *Lancet*, *361*, 1761 - 1766.
- Draper, G., Vincent, T., Kroll, M., & Swanson, J. (2005). Childhood cancer in relation to distance from high voltage power lines in England and Wales: a case-control study. *British Medical Journal*, *330*, 1290.
- Elwood, J. (1988). *Causal Relationships in Medicine*. New York, USA: Oxford University Press.
- Elwood, J. (1998). *Critical Appraisal of Epidemiological Studies and Clinical Trials*. New York, USA: Oxford University Press.
- Farquharson, B. (1990). On-farm trial. In D. Kennedy (Ed.), *Epidemiological Skills in Animal Health. Refresher Course for Veterinarians. Proceedings 143* (p. 207 - 212). Postgraduate Committee in Veterinary Science, University of Sydney, Sydney, Australia.
- Fletcher, R., Fletcher, S., & Wagner, E. (1996). *Clinical Epidemiology*. Baltimore, USA: Williams and Wilkins.
- Fransen, M., Woodward, M., Norton, R., Robinson, E., Butler, J., & Campbell, A. (2002). Excess mortality or institutionalisation following hip fracture: men are at greater risk than women. *Journal of the American Geriatrics Society*, *50*, 685 - 690.

- Fraser, D., Tsai, T., Orenstein, W., Parkin, W., Beecham, H., Sharrar, R., et al. (1977). Legionnaires' disease — description of an epidemic of pneumonia. *New England Journal of Medicine*, *296*, 1189-1197.
- Gardner, I. (1990a). Case study: Investigating neo-natal diarrhoea. In D. Kennedy (Ed.), *Epidemiology at Work. Refresher Course for Veterinarians. Proceedings 144* (p. 109 - 129). Postgraduate Committee in Veterinary Science, University of Sydney, Sydney, Australia.
- Gardner, I. (1990b). Reporting disease outbreaks. In *Epidemiology at Work: Refresher Course for Veterinarians* (p. 29 - 42). Quarantine Station, North Head, NSW: Post Graduate Committee in Veterinary Science, The University of Sydney.
- Gardner, L., Landsittel, D., & Nelson, N. (1999). Risk factors for back injury in 31,076 retail merchandise store workers. *American Journal of Epidemiology*, *150*, 825 - 833.
- Greenhalgh, T. (1997a). How to read a paper: Assessing the methodological quality of published papers. *British Medical Journal*, *315*, 305 - 308.
- Greenhalgh, T. (1997b). How to read a paper: Getting your bearings (deciding what the paper is about). *British Medical Journal*, *315*, 243 - 246.
- Greenhalgh, T. (1997c). How to read a paper: Papers that report diagnostic or screening tests. *British Medical Journal*, *315*, 540 - 543.
- Greenhalgh, T. (1997d). How to read a paper: Papers that summarise other papers (systematic reviews and meta-analyses). *British Medical Journal*, *315*, 672 - 675.
- Greenhalgh, T. (1997e). How to read a paper: Papers that tell you what things cost (economic analyses). *British Medical Journal*, *315*, 596 - 599.
- Greenhalgh, T. (1997f). How to read a paper: Statistics for the nonstatistician. II: 'Significant' relations and their pitfalls. *British Medical Journal*, *315*, 422 - 425.
- Greenhalgh, T. (1997g). How to read a paper: The Medline database. *British Medical Journal*, *315*, 180 - 183.
- Greenhalgh, T. (2006). *How to Read a Paper: The Basics of Evidence-Based Medicine*. London: British Medical Journal Books.
- Greenhalgh, T., & Taylor, R. (1997). How to read a paper: Papers that go beyond numbers (qualitative research). *British Medical Journal*, *315*, 740 - 743.
- Hill, A. (1965). The environment and disease: Association or causation? *Proceedings of the Royal Society of London. Series C, Medicine*, *58*, 295 - 300.
- Hoyert, D., Arias, E., Smith, B., Murphy, S., & Kochanek, K. (1999). *Deaths: final data for 1999. National Vital Statistics Reports Volume 49, Number 8*. Hyattsville MD: National Center for Health Statistics.
- Johansen, C., Boise, J., McLaughlin, J., & Olsen, J. (2001). Cellular telephones and cancer — a nationwide cohort study in Denmark. *Journal of the National Cancer Institute*, *93*, 203 - 237.

- Kelsey, J., Thompson, W., & Evans, A. (1986). *Methods in Observational Epidemiology*. London: Oxford University Press.
- Leung, W.-C. (2002). Measuring chances. *Student British Medical Journal*, 10, 268 - 270.
- Levy, P., & Lemeshow, S. (1999). *Sampling of Populations Methods and Applications*. London: Wiley Series in Probability and Statistics.
- Martin, S., Meek, A., & Willeberg, P. (1987). *Veterinary Epidemiology Principles and Methods*. Ames, Iowa: Iowa State University Press.
- Morris, R. (1990). Disease outbreak! What can you do? In D. Kennedy (Ed.), *Epidemiological Skills in Animal Health. Refresher Course for Veterinarians. Proceedings 143* (p. 321 - 327). Postgraduate Committee in Veterinary Science, University of Sydney, Sydney, Australia.
- Muscat, J., Malkin, M., Thompson, S., Shore, R., Stellman, S., & McRee, D. (2000). Handheld cellular telephone use and risk of brain cancer. *Journal of the American Medical Association*, 284, 3001 - 3007.
- Noordhuizen, J., Frankena, K., Hoofd, C. van der, & Graat, E. (1997). *Application of Quantitative Methods in Veterinary Epidemiology*. Wageningen: Wageningen Pers.
- Oleckno, W. (2002). *Essential Epidemiology Principles and Applications*. Prospect Heights, Illinois: Waveland Press.
- Olsen, J. (2003). What characterises a useful concept of causation in epidemiology? *Journal of Epidemiology and Community Health*, 57, 86 - 88.
- Parsonnet, J., Friedman, G., Vandersteen, D., Chang, Y., Vogelman, J., Orentreich, N., et al. (1991). *Helicobacter pylori* infection and the risk of gastric-carcinoma. *New England Journal of Medicine*, 325(16), 1127 - 1131.
- Petrie, A., & Watson, P. (2005). *Statistics for Veterinary and Animal Science*. London: Blackwell Science.
- Rogan, W., & Gladen, B. (1978). Estimating prevalence from results of a screening test. *American Journal of Epidemiology*, 107, 71 - 76.
- Rothman, K. (1976). Causes. *American Journal of Epidemiology*, 104, 587 - 592.
- Rothman, K., & Greenland, S. (1998). *Modern Epidemiology*. Philadelphia, USA: Lippincott - Raven.
- Schlesselman, J. (1982). *Case-Control Studies Design, Conduct, Analysis*. London: Oxford University Press.
- Schwarz, D., Grisso, J., Miles, C., Holmes, J., Wishner, A., & Sutton, R. (1994). A longitudinal study of injury morbidity in an African-American population. *Journal of the American Medical Association*, 271, 755 - 760.
- Selvin, S. (1996). *Statistical Analysis of Epidemiological Data*. London: Oxford University Press.

- Siscovick, D., Weiss, N., Fletcher, R., & Lasky, T. (1984). The incidence of primary cardiac-arrest during vigorous exercise. *New England Journal Of Medicine*, *311*, 874 - 877.
- Smith, R. (1995). *Veterinary Clinical Epidemiology — A Problem-Oriented Approach*. Boca Raton, Florida: CRC Press.
- Stevenson, M., Morris, R., Lawson, A., Wilesmith, J., Ryan, J., & Jackson, R. (2005). Area-level risks for BSE in British cattle before and after the July 1988 meat and bone meal feed ban. *Preventive Veterinary Medicine*, *69*, 129 - 144.
- Stevenson, M., Wilesmith, J., Ryan, J., Morris, R., Lawson, A., Pfeiffer, D., et al. (2000). Descriptive spatial analysis of the epidemic of bovine spongiform encephalopathy in Great Britain to June 1997. *Veterinary Record*, *147*, 379 - 384.
- Thrusfield, M. (2005). *Veterinary Epidemiology*. London: Blackwell Science.
- Trivier, J., Caron, J., Mahieu, M., Cambier, N., & Rose, C. (2001). Fatal aplastic anaemia associated with clopidogrel. *Lancet*, *357*, 446.
- Valent, F., Brusaferrro, S., & Barbone, F. (2001). A case-crossover study of sleep and childhood injury. *Pediatrics*, *107*, E23.
- Webb, P., Bain, C., & Pirozzo, S. (2005). *Essential Epidemiology*. Cambridge, UK: Cambridge University Press.
- Wilesmith, J., Stevenson, M., King, C., & Morris, R. (2003). Spatio-temporal epidemiology of foot-and-mouth disease in two counties of Great Britain in 2001. *Preventive Veterinary Medicine*, *61*, 157 - 170.
- Will, R., Ironside, J., Zeidler, M., Cousens, S., Estibeiro, K., & Alperovitch, A. (1996). A new variant of Creutzfeld-Jacob disease in the UK. *Lancet*, *347*, 921 - 925.
- Yang, C., Chiu, H., Cheng, M., & Tsai, S. (1998). Chlorination of drinking water and cancer in Taiwan. *Environmental Research*, *78*, 1 - 6.