PERSPECTIVE

Epidemiology for ophthalmologists: an introduction to concepts, study designs, and interpreting findings

H Kuper, C Gilbert

Br J Ophthalmol 2005;89:378–384. doi: 10.1136/bjo.2003.038430

This review provides an overview of the types of information epidemiological research can provide and how these data can be used. The aim is to provide the readers with basic epidemiological skills to allow them to read critically scientific articles and to gain proficiency in communicating about epidemiological research. All examples in the review are drawn from the ophthalmic literature. The first part of the review is relatively conceptual and focuses on epidemiological theory, including case definition, measures of the burden of disease, sampling and the interpretation of results. In the second part different study designs are described—specifically, cross sectional surveys, cohort studies, case-control studies, and randomised controlled trials, and the strengths and limitations of each highlighted.

The intention of this paper is to give an overview of the types of information epidemiological research can provide, and how these data can be used. Broadly, epidemiological research entails assessing the burden and aetiology of disease, and the efficacy of preventive measures or treatments in populations. This information has a variety of uses (table 1). Epidemiological research often relates closely to other areas of research so that findings generated in one sphere are explored, tested, or evaluated in another (fig 1). Epidemiology has been placed unashamedly at the centre of this interconnecting web of endeavour, as the information has such a key part to play in public health and service provision—or should do.

The first part of this review focuses on epidemiological research issues, measures of the burden of disease, and causation. In the second part we discuss different study designs, highlighting the strengths and limitations of each, using examples from the ophthalmic literature.

EPIDEMIOLOGICAL CONCEPTS

Epidemiological studies address questions such as: “How common is the disease?” “Who is most affected?” “What is the underlying cause?” and “What can be done to prevent or treat it?” These questions are answered by conducting studies on a sample of people. The purpose of any study must be clearly articulated, as all aspects of the design, methodology, and analysis flow from this (fig 2).

Case definition

Any epidemiological investigation starts with the definition of the disease or condition of interest, the question of “When is a case a case?” In terms of blindness, a person is defined as blind by the World Health Organization if they have a presenting visual acuity (VA) of $<3/60$ in the better eye, or a central visual field of $<10^\circ$. If the VA is $<6/60$ – $3/60$ the person has severe visual impairment (SVI), and if $<6/18$ – $6/60$ they are visually impaired (VI). With these definitions, globally in the year 2000 there were about 50 million blind people, 25 million with SVI and 125 with VI.¹ Changing the definition of blindness will change the estimated burden of disease. If blindness were VA $<6/60$ there would be 75 million blind people globally, rather than 50 million. We must, therefore, be clear about what constitutes a case and what does not.

Populations and sampling

Measuring the condition of interest in an entire population is not usually feasible. For ease and speed the condition of interest is measured only in a sample taken from the target population, and the findings are extrapolated back to this population.² “Population” can mean the whole population of a country or region, or (more often) the population defined by particular characteristics (for example, diabetic women).

Measures of the burden of disease

The most commonly used measures of disease burden are prevalence and incidence.^{3,4} Prevalence is the proportion of the population with the condition of interest. (A common mistake is to use the term “prevalence rate” rather than “prevalence.” Since prevalence is a proportion, not a rate, the term “prevalence rate” is incorrect and should be avoided.)

Prevalence estimates are used to calculate the expected number of people with disease in the population. In a recent national survey of blindness in Bangladesh, a sample of 11 624 individuals aged ≥ 30 were examined, of whom 162 had a VA of $<3/60$ in the better eye.⁵ The prevalence of blindness was therefore 162/11 624, or 1.39%, giving a national (age standardised) estimate of 650 000 blind people in those aged ≥ 30 .

Incidence differs from prevalence as it relates to new, rather than existing, cases. Broadly, there

Abbreviations: ARMD, age related macular degeneration; CIR, cumulative incidence ratio; IRR, incidence rate ratio; RCT, randomised controlled trials; SVI, severe visual impairment; VA, visual acuity; VI, visual impairment

See end of article for authors' affiliations

Correspondence to: Hannah Kuper, Department of Infectious and Tropical Diseases, London School of Hygiene and Tropical Medicine, Keppel Street, London WC1E 7HT, UK; hannah.kuper@lshtm.ac.uk

Accepted for publication 1 October 2004

Figure 1 Overview of research.

are two measures of incidence: *cumulative incidence* and *incidence rate*. Cumulative incidence is the number of new cases that occur in a population which is disease free at baseline (that is, excluding prevalent cases) over a specified period of time. Approximately seven million people become blind every year—that is, there are seven million incident cases of blindness per year.¹ The disease free population at risk is 6000 million (global population) minus the number of existing cases of blindness (50 million), giving 5950 million. The global cumulative incidence of blindness in the year 2000 was therefore approximately seven million divided by 5950 million, or 0.1%.

The incidence rate uses *person time* at risk as the denominator, rather than the population at risk. During the follow up of a group of people who are disease free at baseline, each person may (a) develop the disease of interest, (b) be lost to follow up, (c) develop a competing disease or die so that they are no longer at risk of becoming a case, or (d) remain disease free. To calculate person time the clock starts ticking at enrolment and stops when one of these events occurs or follow up ends. The total person time at risk is the sum of the person time at risk for each subject in the study. Ten person years of follow up can be produced by one person followed for 10 years, 10 people each followed for 1 year, or 20 people each followed for 6 months. The person time of follow up therefore depends on the number of subjects and the amount of time they are each followed. To calculate the incidence rate we divide the number of incident cases that occur during follow up by the total person time at risk. Incidence rate is expressed as cases per unit of person time (for example, four per 100 000 person years) (box 1).

Prevalence and incidence are distinct, but related, concepts. If the prevalence is relatively low (<10%) then prevalence approximates the incidence multiplied by the average duration of the disease.^{3 4}

Box 1 The calculation of person time

× Event
○ Censored (ie, either lost to follow up, or competing risk, or died)

UNEXPOSED GROUP

Years: 0 1 2 3 4 5

Person A: —————○ 5 person years

Person B: ————○ 1 person year

Person C: —————× 3 person years

Person D: —————○ 3 person years

Person E: —————× 4 person years

Summary: 2 events in 16 person years

Incidence rate = 2/16 per person year = 0.125 per person year = 125 per 1000 person years of follow up

EXPOSED GROUP

Years: 0 1 2 3 4 5

Person A: —————× 3 person years

Person B: —————○ 3 person years

Person C: ————× 2 person years

Person D: —————× 4 person years

Person E: —————× 4 person years

Summary: 4 events in 16 person years

Incidence rate = 4/16 per person year = 0.25 per person year = 250 per 1000 person years of follow up

Incidence rate ratio = IR(exposed)/IR(unexposed) = 0.25/0.125 = 2

Measures of associations

To explore questions about the aetiology of disease or efficacy of treatment we compare the prevalence or incidence in different groups defined by the exposure of interest. This allows us to estimate the association between exposure and disease. This is usually done by calculating the ratio of the incidence (or prevalence) in one group compared to the other, although the absolute difference in the incidence (or prevalence) between two groups can also be used.

The ratio of prevalence in two groups is the *prevalence ratio*; for cumulative incidences this is the *cumulative incidence ratio* (CIR) and for incidence rates the *incidence rate ratio* (IRR).

Table 1 Types and use of epidemiological studies

Study design	Type of information collected	How this information can be used
Descriptive and analytical studies	● Burden of disease (prevalence and incidence) ● Distribution of disease, with ● identification of high risk groups ● Risk factors for disease (aetiology)	Policy Advocacy Priority setting Planning and resource allocation Target setting Evaluation Impact assessment
Intervention studies	● Effectiveness of preventive measures ● Effectiveness of treatments	Health promotion Preventive measures Therapeutic interventions
Meta-analysis and systematic reviews	● Summary of evidence of effectiveness of preventive measures ● Summary of evidence of effectiveness of treatments	Policy Guidelines Planning

Figure 2 Designing epidemiological studies.

Relative risk refers to either CIR or IRR. Another relative measure is the *odds ratio*, where the odds of exposure are compared between those with and those without disease to estimate the exposure disease association. For relative measures, a value of 1.0 indicates no association between exposure and disease. A value >1.0 shows that there may be a positive association between exposure and disease (exposure *may* cause disease), and a value <1.0 indicates an inverse association between exposure and disease (exposure *may* protect from disease).

Chance and sample size

An estimate measured in a sample should reflect the level of certainty that we have in this estimate. We can measure the

prevalence of myopia in a sample of 20 children randomly selected from a population of 200 children. However, if we randomly select a new sample of 20 children the estimated prevalence may differ. This sampling variability is the result of chance and is an inevitable consequence of measuring disease in a selected sample rather than in the whole target population. We express the degree of confidence that we have in our estimate by using confidence intervals which provide an upper and lower range for the estimate. In the national survey of blindness in Bangladesh the 95% confidence interval for the estimate of the number of blind ranged from 524 000 to 725 800 blind people.³ The interpretation of 95% confidence intervals for the sample estimate is that the “true” value for the population lies within this range 95% of the time.⁶ Confidence intervals depend on the size of the study sample: as the sample size increases we are more confident that the findings in our sample reflect the “truth” and so the estimates will have narrower confidence intervals. If the study is too small, the 95% confidence intervals of the estimates will be wide, giving imprecise measures of disease burden, or failing to show a difference between two groups even if one does exist. Although they are distinct concepts, when the 95% confidence interval does not include the null value this usually reflects a statistically significant result.

INTERPRETING RESULTS

We do not conduct studies to make inferences about a specific sample, but to find out general “truths” about whether an exposure is associated with a disease, or an intervention prevents a disease or delays its progression. To begin to make general inferences from the results of individual studies we must be confident that the results were not the result of chance, bias, or confounding (fig 3).

Chance

Variations due to chance are an inevitable consequence of sampling, but the effects can be minimised by having a study that is sufficiently large. The term “power” is often used and this refers to the likelihood of the study resulting in a statistically significant finding, given that there really is a “true” difference between our comparison groups.⁶ If a study does not have a statistically significant result one is left wondering whether this is because the sample was not big enough to detect it or there really is no difference.

Confounding

All associations are potentially influenced by the effects of confounding, which can be thought of as alternative explanations for an association.^{3–4} We can best describe this in stages.

- (a) An exposure appears to be associated with a disease.

Exposure → Disease

- (b) But the exposure is associated with a third factor, the confounder, which is an independent cause of the disease (a variable on the pathway from exposure to disease is *not* a confounder).

- (c) The apparent association between the exposure and disease, may be wholly, or partly, the result of the confounder.

A study was undertaken in Australia to assess whether ultraviolet radiation (exposure) was associated with cataract (disease).⁷ The authors reported a statistically significant positive correlation and argued for the avoidance of sunlight as a preventive measure. However, less wealthy people were more likely to be exposed to sunlight, and poverty is independently related to the prevalence of cataract. The association between sunlight and cataract may, therefore, have been confounded by poverty.

If everyone in the study had the same status with respect to the confounder (for example, all came from the same socioeconomic group) then the factor could no longer exert a confounding influence. This principle is used to control for confounding. In the design, confounding can be controlled by restricting the sample to people with certain characteristics (for example, the same socioeconomic status) or through randomisation (discussed later). In the analysis, multivariate regression or stratified analyses are used as a means of keeping confounders constant, and produce estimates that are “adjusted” for confounder(s). For instance, in an age adjusted analysis we control for age to reduce its confounding effect on the exposure disease association.

Figure 3 Schematic representation of how chance, bias, and confounding can lead to the finding in the sample (result of the study) being different from the “truth” in the target population.

Bias

Study findings may be influenced by bias, which is the deviation of results from the truth.⁶ There are two main types of bias: *selection bias* and *information bias* (also called *measurement bias*).^{3,4} Selection bias is the error caused by systematic differences in characteristics between those who take part in a study and those who do not.⁶ Selection bias could arise in a study to find the association between physical disability and age related macular degeneration (ARMD) if people who are both physically disabled and have ARMD are preferentially included in the study (for example, are more likely to be found at home). This could produce a spurious association between physical disability and ARMD. Information bias is caused by inaccuracy in the measurement of exposure or disease that results in different quality of information between comparison groups.⁶ If we interview patients with disease in a hospital but the healthy comparison subjects in their homes we may generate a spurious association between an exposure and disease because the patients in the hospital were nervous or the interviewer examined them more carefully. The degree to which bias may deviate the results from the “truth” cannot be measured, but can be minimised by good study design and paying attention to the quality of data collection.

Causation

If the effects of confounding, chance, and bias are ruled out then causation becomes a possible explanation for the exposure disease association. Bradford-Hill outlined guidelines for causality, arguing that while we cannot prove causality, there is a stronger case if certain criteria are fulfilled.⁸ Firstly, the exposure must precede the disease (*temporality*). The case for causation is strengthened if the association between the exposure and disease is large (*strength of evidence*), consistent across different studies (*coherence*), in line with current biological knowledge (*experimental evidence*) and generally credible (*plausibility*). Showing that increasing the level of the exposure further increases the risk of disease also supports the case for causality (*dose-response relation*). Bradford-Hill also argued that a role for causation was strengthened if a specific exposure leads to one specific disease (*specificity*), but since some exposures (for example, smoking) can lead to many diseases this criterion is weak and will not be discussed here. Furthermore, the criterion of making comparisons with other disease-exposure associations (*analogy*) is too vague to be useful. For a full discussion of causation and causal inference see Rothman and Greenland³ or Rothman. For a more philosophical discussion of the question of causation we refer you to MacMahon and Trichopoulos.⁹

STUDY DESIGNS

The study designs used in epidemiology can be grouped as descriptive or analytical studies (cross sectional, cohort, and case-control studies), where investigators observe the events as they unfold naturally, and intervention studies (randomised controlled trial), where investigators manipulate exposure(s) to risk factors/treatments to assess their impact on disease.^{6,10,11}

Cross sectional survey

Cross sectional surveys allow us to measure the prevalence of disease, so that the burden of disease can be estimated.^{3,4} To conduct a cross sectional survey, we carefully sample the desired number of study participants and examine and/or interview them to determine whether they have the disease and exposure(s) of interest. These data can be used to calculate the overall prevalence or to compare the prevalence in the exposed and unexposed groups and thereby explore aetiological questions.

Bourne and colleagues conducted a cross sectional survey in Bangkok to estimate the burden of glaucoma.¹² They sampled 701 people aged >50 years and examined them for glaucoma. The estimated prevalence of glaucoma was 3.8%, but was higher in women (6.0%) than men (3.2%), giving a prevalence ratio of 1.86 (95% CI: 0.9 to 4.0). Since the 95% confidence interval included the null value, the “truth” may be that there is no sex difference in glaucoma prevalence.

With cross sectional surveys we must emphasise the adage that “association does not equal causation.” One reason for this is that exposure and disease are measured at the same time and so we cannot be sure that the exposure preceded the disease (the “temporality” criterion cannot be tested). A survey could show that people with myopia read more for pleasure than those with normal vision, but this may be because myopia is caused by reading or because people with uncorrected myopia have poor distance vision and reading becomes a more important leisure activity. Furthermore, as a general principle, people with prevalent disease often have a benign and long lasting form of disease that is neither fatal nor readily treated. Prevalent cases may therefore not be representative of all cases that have occurred, and so surveys may identify only risk factors for prevalent, rather than incident, disease.

Cohort study

Cohort studies allow us to measure predictors of disease incidence.^{3,4} To conduct a cohort study, a group of people free from the disease of interest are recruited and characterised as “exposed” or “unexposed” with respect to the risk factor under investigation. The participants are followed over time and the number of incident cases of disease determined. The incidence is calculated for the exposed and unexposed groups, and the incidence ratio (CIR or IRR) is estimated. Cohort studies can be closed, where people are enrolled only at the beginning of follow up (after which the study is closed), or open, where enrolment occurs over time.

Bowman and colleagues conducted a closed cohort study in the Gambia to examine the association between trichiasis (exposure) and corneal visual loss (outcome).¹³ At baseline 639 people with trachomatous lid scarring but without corneal visual loss were identified. Some subjects had trichiasis (exposed group) while others did not (unexposed group). After 12 years, 326 of the initial cohort were retraced. Of the 26 people with trichiasis at baseline (exposed group) two had developed corneal visual loss (CI: 2/26 or 7.7% over 12 years), compared to six of the 295 people with trachomatous lid scarring without trichiasis (unexposed group) (CI: 6/295 or 2.0% over 12 years). The CIR comparing the exposed and unexposed groups was 3.78 (95% CI: 0.80 to 17.81), showing that the 12 year risk of developing corneal visual loss was almost four times higher in people with trichiasis at baseline compared to those without (although the confidence interval included the null value).

This example illustrates the problem of loss to follow up, which is often encountered by cohort studies. Only 326 of the initial cohort of 639 were traced, and it was not known what happened to the remaining 313 people in terms corneal visual loss. Those lost to follow up often differ from those traced in

terms of their exposure and disease status. Loss to follow up can therefore lead to selection bias. Another problem is *competing risks*, where cohort members die or develop other diseases so that they can no longer develop the disease of interest and the “true” cumulative incidence is underestimated. The problems of competing risks and loss to follow up can be overcome by using person time analyses and calculating incidence rates, as people only contribute person time while they are eligible of becoming a case. Person time analyses also allow people to switch between exposure status, contributing person time first to one exposure group and then to another. However, calculating incidence rates requires ongoing measurement of exposure and disease status, and this is logistically difficult.

The strengths of the cohort study design include the ability to measure the incidence of disease, the assurance (more or less) that the exposure preceded the disease, and the fact that several disease outcomes can be measured. The main drawback is that since most diseases are rare we either need a large sample or long follow up to accumulate enough cases to have sufficient power to make meaningful inferences. This makes cohort studies expensive and time consuming, which is why there are only a few in the ophthalmic literature.

Case-control studies

Case-control studies are used to study the aetiology of disease. Case-control studies are conducted by recruiting people who have the disease of interest (cases) as well as people without the disease (controls).^{3,4,9} The controls should be selected from the same population that gave rise to the cases (that is, if any of the controls had developed the disease they would have been included as a case) so that they represent the exposure distribution in the source population. In practice, controls are usually selected from the general population (population based case-control study) or from the same hospital as the cases (hospital based case-control study). Cases and controls are interviewed, or past medical records/laboratory files examined, to assess their exposure status. The odds of exposure in cases (number exposed versus unexposed) is compared to the odds of exposure in controls, to assess the exposure-disease association. Case-control cannot be used to estimate the burden of disease, since the ratio of cases to controls is determined by the investigators. The ratio of cases to controls is often one to one, but more than one control can be selected per case to increase the statistical power of the study.

Minassian and colleagues conducted a hospital based case-control study to investigate the association between child-bearing and risk of cataract in young women.¹⁴ Cases were women aged 35–45 with bilateral “senile” cataract attending an eye hospital in central India. Controls were women of the same age with clear lenses attending the hospital with other complaints. Cases and controls were interviewed about their history of pregnancy and childbirth. The number of live births was statistically significantly higher in cases than in controls and a dose-response relation between childbearing and risk of cataract was apparent (table 2).

Case-control studies are relatively quick and cheap to carry out. They can also be used to investigate rare diseases and

Table 2 Odds of many versus few pregnancies in cataract cases compared to controls, hospital based case-control study¹⁴

Children born	Cataract cases	Controls	Odds ratio (95% CI)
1–3*	31	134	Baseline
4–6†	49	114	1.9 (1.1 to 3.1)
7–11†	15	14	4.6 (2.0 to 10.6)

*Unexposed group, †exposed group.

Table 3 Baseline characteristics of the participants in the four treatment arms of the onchocerciasis treatment randomised controlled trial¹⁷

Baseline characteristics*	150 µg/kg yearly (n = 166)	800 µg/kg yearly (n = 172)	150 µg/kg 3 monthly (n = 161)	800 µg/kg 3 monthly (n = 158)
Age (years)	39.6	41.1	40.1	38.4
Height (cm)	172.1	172.8	171.3	172.6
Weight (kg)	62.6	64.1	61.8	62.8
Palpable nodules	4.8	4.9	4.9	5.0
Patients with skin microfilariae (%)	99	95	97	98
Microfilariae	62.9	51.4	49.6	57.4

*Values are mean unless otherwise stated.

multiple exposures. Recall bias is a problem, however, since cases may report exposures differently from controls. This problem is less serious when the exposure can be objectively assessed (for example, height), accurately recalled (for example, age), or verified (for example, treatment received). There is also the potential for selection bias, particularly in the selection of the controls. Case-control studies can be unfairly maligned because of these limitations, but if designed and executed thoughtfully they can provide the same information as cohort studies in a shorter time and at lower cost.

Randomised controlled trials

A randomised controlled trial (RCT) is an intervention study that forms a special subset of cohort studies.^{3 4 15} RCTs are typically used to assess the benefit of a new drug or treatment, but they are also useful for evaluating the impact of a preventive measure (for example, health education). In an RCT, people are selected and randomised to receive either the intervention (treatment under investigation) or the control (placebo or standard treatment). The purpose of randomisation is to make the intervention and control groups as similar as possible with respect to important confounders (both known and unknown) so that the only difference between them is that one group receives the intervention while the other does not. The two groups are monitored over time for the defined outcomes, allowing the relative risk to be calculated.

People will not always comply with the assigned intervention; they may be lost to follow up, refuse treatment, or else privately use different treatments. However, data analysis should only take account of the treatment group to which they were assigned (intent to treat analyses) to avoid breaking the randomisation, as otherwise the difference between the groups could be the result of baseline differences in risk factors rather than the result of the intervention. Participants in an RCT often will not know whether they are receiving the intervention or the control (the study is blinded or masked) and ideally the investigator will also be unaware of their treatment status (the study is double blinded/masked). Blinding helps to reduce information bias.¹⁶

Gardon and colleagues conducted an RCT in Cameroon to assess the effectiveness of new treatment regimens for onchocerciasis.¹⁷ They randomised 657 people with onchocerciasis to receive 150 µg/kg ivermectin yearly (standard practice group), 150 µg/kg every 3 months, 400 µg/kg then 800 µg/kg yearly, or 400 µg/kg then 800 µg/kg every 3 months. The randomisation resulted in a similar profile of baseline characteristics in the four groups, so that the only difference was the treatment assigned (table 3). Both the participants and investigators were blinded to the treatment group status. The primary outcome was the vital status of female worms which was assessed after 3 years of follow up. Significantly more female worms had died in the three monthly treated groups than in the standard practice group (OR = 1.84, 95% CI: 1.23 to 2.75 for 150 µg/kg, and

OR = 2.17, 95% CI: 1.42 to 3.31 for 800 µg/kg) and 3 monthly treatment was also superior to standard practice in terms of reducing the number of female worms, itching, skin lesions, and transmission of onchocerciasis. These results are reliable because the effects of confounding, selection bias, and information bias have been minimised.

RCTs are often held up as the gold standard of study designs. They are, however, expensive, take a long time to generate results, and can only be used to answer certain types of questions. There are also ethical issues concerning the use of RCTs, as it is immoral to give a treatment that is known to be worse, or to withhold a treatment that is better, than standard practice or placebo. In designing RCTs there must therefore be a state of equipoise where the potential benefits of a new treatment are equally and oppositely outweighed by the potential harm. A committee should be established to ensure that equipoise is maintained throughout the trial.

Box 2 Suggestions for further reading

Epidemiology methods

- Hennekens CH, Buring JE, Mayrent SL. *Epidemiology in medicine*. Boston: Little, Brown, 1987. Easy to read introduction to epidemiological methods and techniques. Now slightly out of date.
- Last JM, ed. *A dictionary of epidemiology*. New York: Oxford University Press, 2001.⁶ Useful as a reference for epidemiological terms, but not appropriate as an introductory text.
- MacMahon B, Trichopoulos D. *Epidemiology: principles and methods*. 2nd ed. Boston: Little, Brown and Company, 1996.⁹ Stylishly written introduction to epidemiological methods, providing useful practical tips for conducting epidemiological investigations.
- Rothman KJ. *Epidemiology: an introduction*. New York: Oxford University Press, 2002.⁴ Up to date comprehensive introductory text incorporating more sophisticated concepts.
- Rothman KJ, Greenland S. *Modern epidemiology*. 2nd ed. Philadelphia: Lippincott-Raven, 1998.³ Provides broad overview of epidemiological methods as well as specialist topics in epidemiology. Especially suitable for the intermediate epidemiologist.

Ophthalmic epidemiology

- Johnson GJ, Weale R, Minassian DC, *et al*, eds. *The epidemiology of eye disease*. 2nd ed. New York: Oxford University Press, 2003. Summary of epidemiological methods as they relate to eye diseases, as well as the epidemiology of specific eye diseases.

We hope that this overview has shown you that there is no ideal study design. It is more a case of “horses for courses” as different designs are appropriate in different situations.

CONCLUSIONS

In this review we gave a whirlwind account of epidemiology as it relates to ophthalmologists. We hope that this will increase your ability to read papers critically, and your effectiveness when communicating about epidemiology, whether orally or in writing. In one review we cannot hope to do more than touch on a number of important issues and for those interested in exploring these topics more deeply we provide a list of suggestions for further reading (box 2).

Authors' affiliations

H Kuper, C Gilbert, International Centre for Eye Health, Clinical Research Unit, Department of Infectious and Tropical Diseases, London School of Hygiene and Tropical Medicine, Keppel Street, London WC1E 7HT, UK

REFERENCES

- 1 **Dandona L**, Foster A. Patterns of blindness. In: Tasman W, Jaeger EA, eds. *Duane's clinical ophthalmology*. Philadelphia: Lippincott-Raven, 2001.
- 2 **Levy PS**, Lemeshow S. In: *Sampling of populations: methods and applications*. 3rd ed. New York: John Wiley, 1999.
- 3 **Rothman K**, Greenland S. In: *Modern epidemiology*. 2nd ed. Philadelphia: Lippincott-Raven, 1998.
- 4 **Rothman K**. *Epidemiology: an introduction*. New York: Oxford University Press, 2002.
- 5 **Dineen BP**, Bourne RR, Ali SM, *et al*. Prevalence and causes of blindness and visual impairment in Bangladeshi adults: results of the National Blindness and Low Vision Survey of Bangladesh. *Br J Ophthalmol* 2003;**87**:820–8.
- 6 **Last JM**, ed. *A dictionary of epidemiology*. New York: Oxford University Press, 2001.
- 7 **Hollings F**, Moran D. Cataract—the ultraviolet risk factor. *Lancet* 1981;**2**:1249–50.
- 8 **Hill AB**. The environment and disease: association or causation? *Proc R Soc Med* 1965;**58**:295–300.
- 9 **MacMahon B**, Trichopoulos D. In: *Epidemiology principles and methods*. 2nd ed. New York: Little, Brown and Company, 1996.
- 10 **Egger M**, Smith GD, Phillips AN. Meta-analysis: principles and procedures. *BMJ* 1997;**315**:1533–7.
- 11 **Egger M**, Davey Smith G, Altman DG, eds. In: *Systematic reviews in health care: meta-analysis in context*. 2nd ed. London: BMJ, 2001.
- 12 **Bourne RR**, Sukdom P, Foster PJ, *et al*. Prevalence of glaucoma in Thailand: a population based survey in Rom Klao District, Bangkok. *Br J Ophthalmol* 2003;**87**:1069–74.
- 13 **Bowman RJ**, Jatta B, Cham B, *et al*. Natural history of trachomatous scarring in the Gambia: results of a 12-year longitudinal follow-up. *Ophthalmology* 2001;**108**:2219–24.
- 14 **Minassian DC**, Mehra V, Reidy A. Childbearing and risk of cataract in young women: an epidemiological study in central India. *Br J Ophthalmol* 2002;**86**:548–50.
- 15 **Piantadosi P**. *Clinical trials a methodologic perspective*. New York: John Wiley, 1997.
- 16 **Moher D**, Schulz KF, Altman DG. The CONSORT statement: revised recommendations for improving the quality of reports of parallel-group randomised trials. *Lancet* 2001;**357**:1191–4.
- 17 **Gardon J**, Boussinesq M, Kamgno J, *et al*. Effects of standard and high doses of ivermectin on adult worms of *Onchocerca volvulus*: a randomised controlled trial. *Lancet* 2002;**360**:203–10.