CLINICAL EPIDEMIOLOGY FOR THE MRCGP EXAMINATION

"Today's therapy, instigated solely as a result of clinical experience, becomes tomorrow's bad joke."
Clinical Epidemiology. A Basic Science for Clinical Medicine, D.L.Sackett, R.B. Haynes, G.H. Guyatt, P. Tugwell
2nd. Edition. 1991. Little, Brown

The concept of "evidence based medicine" (EBM) has received much attention in recent years. It was first advocated by Archie Cochrane in "Effectiveness and Efficiency" and has been taken on David Sackett (now of Oxford University) and his - colleagues from McMaster and Ottawa Universities 1i1. The development of EBM is inextricably linked with the landmark publication of "Effective Care in Pregnancy and Childbirth" by Ian Chalmers and colleagues in 1989 [2]. It involves the application of clinical trial evidence to everyday care as a means of closing the gap between research and everyday practice. Four steps are involved:

1. Accurate identification of the question to be investigated.
2. A search of the literature to select relevant articles.
3. An evaluation of the evidence in the literature selected.
4. Implementation of the findings in clinical practice.

The fears of some clinicians that these developments threaten the concept of the individual doctor - patient relationship are an understandable emotional reaction to the change threatening current practice. In fact the obverse is true. In order to deliver evidence based medicine to individual patients greater clinical skills are required. Diagnoses must be ever more accurate, communication skills need to be honed to a fine art to achieve a jointly agreed and understood management plan between doctor and patient, and new skills learnt to master the scientific basis of clinical practice.

It must be recognized even by evidence-based enthusiasts that there are limits to this approach. As knowledge about specific effective interventions becomes clearer, the difficulties of applying this knowledge and judgment to individual patients who may will have multiple pathologies or risk factors means that increased professional expertise will be demanded of doctors. As treatment improves the stakes involved in delivering optimal clinical care increase. Combining multiple interventions into clinical strategies on an evidence base is problematical. Two interventions can be combined in two different ways. Five interventions results in a possible one hundred and twenty combinations. The risk of "cookbook medicine" taking over is not credible but neither is continuing with practice based solely on opinion and clinical experience.

Continuing Medical Education for the 21st. Century.

There is increasing emphasis on effectiveness and efficiency from patients and professional leaders. The challenge is to achieve the best care for individuals and the population in the face of increasing health care costs, demographic change and the pattern of disease (notably the ageing population and the increase in chronic health problems), biomedical advances and communication of knowledge previously only vested intra-professionally. These forces lead inevitably to changing requirements in medical education; the inclusion of clinical epidemiology in the MRCGP syllabus is a response to this obligation.

Attitudes, behaviour, critical appraisal skills, consultation skills flexible thinking, and access to data about and implementation of effective care are now as important as the mere possession of biomedical information (Fig. 1.).

Fig. 1.

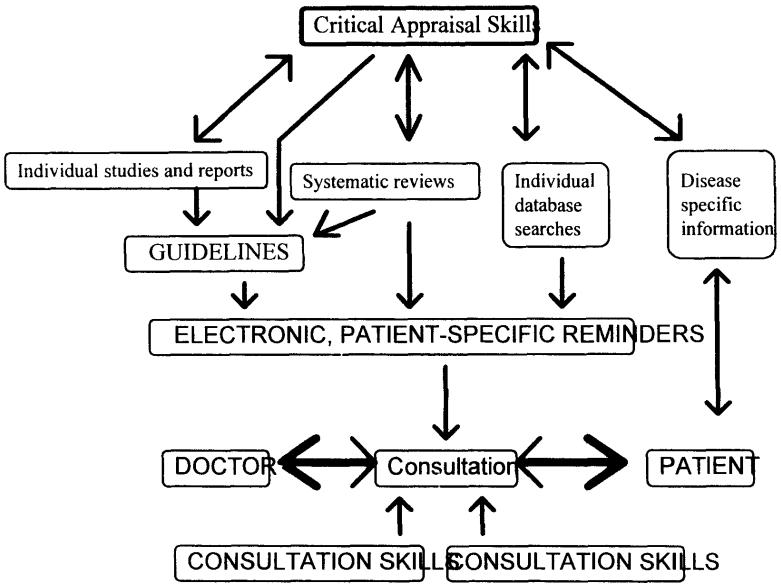

Chance, Bias and Confounding Variables.

In any study of the effect of a medical intervention on the natural history of a disease one would expect to see results clearly stated. They may show that the intervention has been successful and reduced deaths or disability. We need to know whether it is likely that these results have occurred by chance, whether the results could be biased by the design of the study or by the inclusion or exclusion of some patients, and also to consider whether another factor (termed a confounding variable), independent of the intervention, is producing an erroneous conclusion.

Chance and bias are straightforward issues. An example of a confounding variable (Fig. 2) is the study which shows an association of coffee drinking with an increased rate of carcinoma of the bronchus. In fact it is smoking that is the intervention responsible for the lung cancer and coffee drinkers were more likely to be smokers.

Fig. 2.

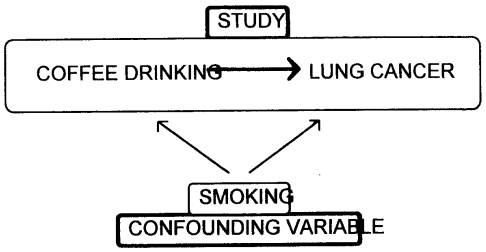

The solution to avoid this confounding variable would be to have two groups of coffee drinkers - smokers and non-smokers. The rates of lung cancer could then be determined in both groups.

A confounding variable may therefore be considered as a particular type of bias, and several biases may occur (and therefore need to be taken account of) in a single study. If we were to run a study to determine whether regular exercise lowers the risk of coronary heart disease we might do so by offering aerobic classes to employees of a large company and then measuring the number of coronary events in the groups who did and did not volunteer for the classes. The events would be determined by regular check ups including a careful history, an electrocardiogram and a review of GP and hospital records.

The results of this (hypothetical) study show that the exercise group had fewer cardiac events. However, the review of records also showed that the exercise group smoked less. Selection bias could also operate if the exercise group were at lower risk before the programme began - did they have less hypertension, lower blood cholesterol and more favourable family histories? Measurement bias may occurred because those participants who knew they had

had coronary events could be more likely to attend for their study check up and report their problem. Finally the lower cigarette consumption in the exercise group would be a confounding bias.

TYPES OF STUDY

It is helpful to separate out studies that are observational (and are therefore hypothesis-forming) from those that are analytical (and therefore hypothesis testing).

Observational - Hypothesis Forming:

(i) Case reports and case studies.
Many important advances in medical knowledge begin with simple descriptions of a small series of cases presenting clinically to astute doctors. e.g. five cases of male homosexuals in San Francisco with pneumocystis carnii as the cause of their pneumonia, subsequently shown to have HIV infection and AIDS. There is usually no attempt to determine whether causal association in the study - the purpose of the report is to raise awareness. Proof will only be provided by more extensive investigation.

(ii) Cross Sectional Studies.
Also termed prevalence studies (see below), this method involves a survey of a given population and attempts to correlate between personal factors and disease states. It cannot measure cause and effect, nor can it determine changes between exposure and disease. Again, it may lend weight to a more rigorous investigation being required.

(iii) Correlational Studies.
Sometimes termed ecological or geographic studies, these look at the number of cases in a given population at any given time (the prevalence) or the number of new cases occurring in a given time (the incidence), and compare the prevalence or incidence with another population. Limited information as to causation can be obtained but useful inferences can sometime be forthcoming e.g. migrant studies of Japanese from - their home country to the United States and their rate in successive generations of acquiring the pattern of ischaemic heart disease of Americans.

Analytical - Hypothesis Testing:

In order to determine whether a possible factor really is involved in a disease, or a particular intervention really does improve the treatment of that disease we need a different type of analysis.

(i) Case Control Studies.
Case control studies take a sample of patients with the disease - the cases - and match these cases with a sample of the population who do not have the disease - the controls. The controls need to be as similar as possible to the cases (except in respect of the risk exposure) to reduce bias. The case control study then looks backwards in time and tries to determine the frequency of exposure to the identified risk factor in both cases and controls.

Fig 3.

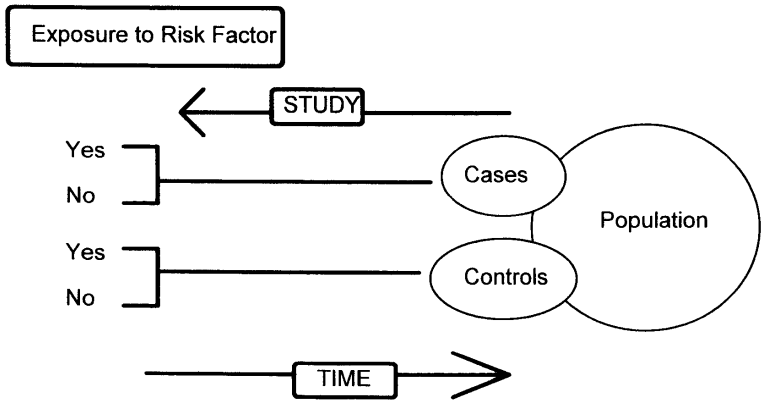

The results are usually presented in a table:-

		DISEASE	
		Cases	**Controls**
	Yes	**a**	**b**
EXPOSURE	**No**	**c**	**d**

The odds ratio of the exposure resulting in the disease can then be calculated from: $\frac{ad}{bc}$

An odds ratio of I would show no association, a value below this a protective effect of exposure and numbers in excess of one a possible association, though a causal relationship would require further consideration in almost all circumstances. It should be noted that an odds ratio from a case control study is not a measure of the risk in the general population - as an inherent part of their design case control studies cannot provide incidence data.

Hypothetical results from a case control study designed to see whether lung cancer is linked to smoking might produce the following table:

Table 1.

EXPOSURE		DISEASE Cases	Controls
(lung cancer)	Yes	56	230
(smoking)	No	7	246

The odds ratio would therefore be $\frac{56 \times 246}{7 \times 230} = \frac{13776}{1610} = 8.6.$

This would be a large enough odds ratio to indicate the possibility that there was a true association between exposure and disease.

An example of the usefulness of the case control design was published in 1994 [3]. Several case series had previously shown that in patients with low back pain a magnetic resonance imaging scan (MRI) had demonstrated lumbosacral disc abnormalities in the majority of patients. However, when a control group was also studied a similar incidence of disc abnormalities was found in the control group. This results an odds ratio approximating to 1 and therefore doubts being expressed concerning the hypothesis that the abnormalities seen on MRI scanning in cases of low back pain are related to the cause of the pain.

(ii) Cohort Studies.

In a cohort study a sample of the population who have the potential to develop a disease are assembled. This sample is then classified into characteristics (possible risk factors) that might be related to outcome. Observation over time then takes place with collection of data to see which members of the cohort experience the outcome being measured. Cohort studies are sometimes called longitudinal or incidence studies.

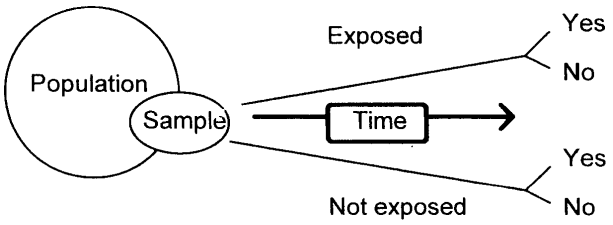

Fig. 4.

Sometimes cohort studies are performed where the sample is selected historically. A good example of this is the UK birth cohorts where all babies born in a single week in 1948, 1958 and 1970 have been followed throughout their lives. The sample is available for follow up in the present but the cohort is assembled in the past. A concurrent or prospective cohort study assembles the cohort in the present and is then destined to follow the cohort forward with follow up at a designated point or points in the future.

Cohort studies also usually present their main results in the form of a table:

	OUTCOME	
	Yes	**No**
Exposed	a	b
Not exposed	c	d

The simplest analysis consists of attributable risk (sometimes called absolute risk or risk difference) and relative risk (sometimes called risk ratio).

Attributable risk answers the question "What is the incidence of disease attributable to exposure" and is simply a - c.

Relative risk answers the question "How many times are exposed persons more likely to develop the disease, relative to non-exposed persons?" i.e. the incidence in the exposed divided by the incidence in the non-exposed.

This is expressed as $\frac{a}{a+b}$ divided by $\frac{a}{c+d}$

As an example let us consider the development of deep vein thromboses (DVT) in oral contraceptive users. Hypothetical results might look something like Table 2.

Table 2.

	OUTCOME (DVT)	
Exposed (on oral contraceptive)	Yes	No
	41	9996
Not exposed (not on o.c.)	7	10009

These results would give an attributable risk of 34 and a relative risk of 6 - significantly large enough numbers to indicate the possibility of a real association between exposure and outcome. However, the possibility of biases very often arises in studies and the risk is greater in designs that are other than randomized trials. In this case, are women at higher risk of DVTs given an oral contraceptive? Is it possible that women on oral contraceptives are more likely to themselves report symptoms of a DVT, whereas women not on the treatment will ignore them? Are doctors more likely to make the diagnosis when their own suspicions have been raised by their patient's current medication? These are real possibilities and a well designed study will provide evidence to restrict or refute influences that may skew the result.

A good example of a cohort is the Framingham study [41 which was started in 1949 when a sample of 5209 men and women aged between 30 and 59 were selected as a representative sample from about 10,000 persons of that age living in Framingham, near

Boston, USA. The study was set up to identify factors associated with coronary heart disease and 5127 of the cohort were free of the disease when first examined. As is well known, the risk factors that have been identified that are associated with the development of coronary heart disease are elevated blood pressure, hypercholesterolaemia, cigarette smoking, diabetes mellitus and left ventricular hypertrophy. Since the sample is representative of the population and its size is related to the true population, real incidence figures are available from cohort studies. This is one of their major advantages.

(iii) Randomized controlled trials.

Randomized controlled trials (RCTs) are often referred to as "the gold standard" when evidence based medicine is discussed. This is because their design restricts the biases that may influence the results of case control and cohort studies. They are undoubtedly the standard of excellence for assessing the effects of treatment.

Fig. 5.

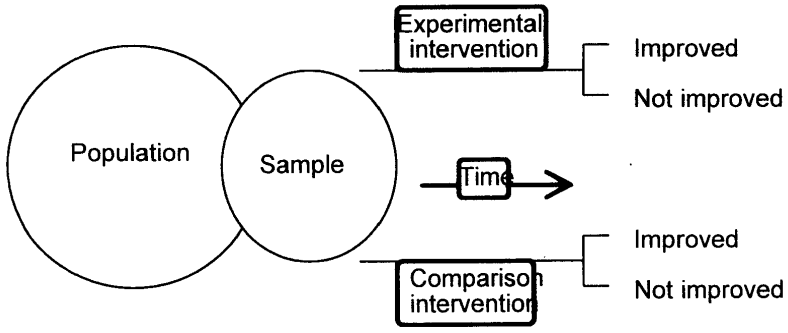

The design of randomized controlled trials is familiar. The patients to be studied are selected by defined criteria from a larger number of patients with the condition under investigation. Those who then agree to participate in the , study are then randomized (by a system analogous to tossing a coin) into two groups of comparable prognosis. Randomization produces two groups which differ only by chance - the purpose is not to produce equal groups, though in large trials the groups that emerge are balanced. Two comparable interventions are then applied to the groups and the outcomes measured. Ideally patients, their attending physicians and the study investigators should all be unaware of which patients received which treatments - a process known as blinding. Both randomisation and blinding are used to avoid bias, with errors in the results obtained therefore restricted to chance. Where small improvements in outcomes are expected from the intervention under investigation, large numbers of patients are required for the trial.

Finally the RCT design needs to consider whether the objective is to find out whether offering treatment helps in normal clinical practice or whether the treatment is efficacious under ideal circumstances. For example, if we were investigating a new antibiotic used for pneumonia we could design an RCT where the outcomes could either be clearance of the causative organism from the sputum or the length of stay in hospital. The first study might show that the organism was cleared faster than placebo or an alternative antibiotic and (all other aspects being equal) would produce a valid assessment of the drug's efficacy. However, for this result to be generalisable to everyday practice we would want to know that patients got better quicker and were able to leave hospital earlier. Conducting a trial with such an outcome would, however, potentially lead to the introduction of other variables (e.g. concurrent or intercurrent illnesses, variations in administration procedures and policies on discharge from hospital) which could bias the result. RCTs therefore often try and strike a balance between validity and generalisabilty. They may often only answer one or other question - and the subsequent debate fills up the correspondence columns of medical journals.

An additional problem is that RCTs are, by definition, measuring treatment being provided in an experimental (and therefore artificial) setting. Transferring a valid result from the RCT carries the risk of sub-optimal results due to the different setting and conflicting pressures of everyday clinical care.

Presenting the results of an RCT would produce table 3. The results of an RCT are often only presented as a relative risk reduction (RRR) e.g. "magicillin reduces the length of stay in hospital in patients with pneumonia by 45%." Whilst the RRR answers the question "How much

better is the active treatment than the comparison intervention?", it does not take into account the incidence of the disease in the population. If we are to assess the value of magicillin to society we need the absolute risk reduction (ARR) which answers the question "How many fewer patients will get the outcome I am measuring if I use active treatment instead of the comparison intervention".

Table 3.

OUTCOME	Yes	No
Comparison intervention	a	b
Experimental intervention	c	d

Absolute risk reduction is therefore the comparison intervention patients with the outcome out of the total of the comparison patients minus the experimental patients with the outcome out of the total patients on experimental treatment.

i.e. $\frac{a}{a+b} - \frac{c}{c+d}$

Relative risk reduction is ARR in a ratio to the outcomes measured in the comparison group.

i.e.. $\frac{\frac{a}{a+b} - \frac{c}{c+d}}{\frac{a}{a+b}}$

These complicated formulae become clearer if we consider real data from the recent 4S study [5].

Table 4.

OUTCOME (death)	Yes	No	Total
Comparison intervention (placebo)	256	1967	2223
Experimental intervention (simvastatin)	182	2039	2221

The ARR is (256/2223) - (182/2221) = 0.115 - 0.082 = 0.033.
The RRR is 0.033/0.115 = 0.29 or expressed as a percentage 29%.

Treating patients with established coronary heart disease (CHD) with simvastatin for a mean duration of 5.4 years in the 4S study therefore reduced all cause mortality by 29%. All in all, a pretty impressive result - even when the particular circumstances of a RCT and the patients excluded from the study are taken into account. However, in order to assess the benefits when the study is applied to the population we need to consider the incidence of deaths from coronary from heart disease. The ARR takes this into account but the figure of 0.033 is difficult to interpret. The figure contains more useful information than the crude risk reduction but the decimal form is unfamiliar to clinicians. What does 0.033 mean in practice?

This difficulty is solved by dividing the ARR into 1 i.e. by taking its reciprocal. This turns out to be the number of patients we need to treat with the experimental intervention to prevent one outcome. 1/0.033 = 30. We therefore now know the number needed to treat - we need to treat 30 patients with coronary heart disease for 5.4 years with simvastatin to prevent one death - a much more accessible and meaningful statement than "the absolute risk reduction is 0.033".

Numbers needed to treat (NNTs) are now starting to be quoted in trials in the mainstream peer-reviewed medical journals. The clinical effectiveness industry is also busy calculating NNTs for current interventions. Some of these are presented in table 5 (N.B. Refer to the original studies for full details - this data is accurate and very interesting as a crude comparison between interventions, but the full picture from the original papers is required to obtain the nuances of e.g. trial design, withdrawals, exclusions, blinding and other potential biases.)

Table 5. Source: [6].

INTERVENTION	OUTCOME	NNT
Streptokinase + aspirin v. placebo (ISIS 2)	prevent 1 death at 5 weeks	20
tPA v. streptokinase (GUSTO trial)	save 1 life with tPA usage	100
Simvastatin v. placebo in CHD	prevent 1 major event in 5y	15
Treating hypertension in the over-60s	prevent 1 major event in 5y	18
Aspirin v. placebo in healthy adults	prevent MI or death in 1y	500

Now the clinical effectiveness picture begins to make a little more sense. We can advocate streptokinase with aspirin in myocardial infarction, treating hypertension in the over 60s and using simvastatin in coronary heart disease, whilst being very cautious at first glance about primary prevention of CHD with aspirin and about the overall benefits of tPA over streptokinase. We need to know more about the particular studies to determine their generalisability and whether there are some special subgroups of patients where the benefits might be greater or less than the population in general, but the numbers needed to treat allows some useful comparisons between the proportional benefits of different medical treatments and their overall contribution to healthcare.

Even so there are caveats to be added. We have not considered the side effects of our interventions. How many patients with hypertension will develop impotence, gout or diabetes as a result of our treatment? How serious a risk is there of rhabdomyelosis or hepatitis with simvastatin? How great is the risk of causing a haemorrhagic stroke or serious anaphylactic reaction with streptokinase?

Further development is therefore likely towards a combined index which will result in accessible compilation of data that will incorporate both the benefits and the risks of interventions, together with an indication of the likely improvement in the quality as well as the quantity of life. Still there will be difficulties in applying this data to individual patients with multiple pathologies and risk factors. But it is easy to envisage not very far in the future expert guidance software on the GP's desktop that will calculate the odds of different interventions based on an biological data for that patient - patient and clinician then discussing and compiling a management plan based on evidence rather than clinical experience and opinion. Clinicians therefore need to understand how to access and assess information on effective interventions - individual studies, meta-analyses and systematic reviews - and to be effective communicators of this new information to their patients.

TESTS OF SIGNIFICANCE

Statistics is for many clinicians a concept even more detestable than management. This is due to our own value systems, mathematical ineptitude and the fact that mathematics and statistics are almost always taught by highly competent and qualified mathematicians. Unfortunately this means that not only do they speak another language from their students but they also find it a frustrating experience in trying to instruct what are to them very simple concepts. Disillusion quickly sets in upon both parties, confusion and bewilderment are not far behind and another biological scientist thinks understanding statistics is an impossibility.

Two principles stand out when it comes to statistics.
1. Since most of the really important evidence-based medicine is based on randomized clinical trials, only knowledge of what probability and confidence intervals are and what they mean is required.
2. For those who wish to learn a little more there is an understandable introduction to statistics written by a psychologist in terms that non-mathematicians can understand. Most medics usually only come to this advanced stage of development after the passage of some time and the internal kindling of a spark of interest by a chance event, rather than being driven by the external forces of needing to pass an examination. "Simple Statistics" [7] is a truly wonderful book and deserves to be regarded as a classic.

Probability.

Trials are analyzed on the basis that there is no difference between the treatments. This is termed the "null hypothesis". The probability that the observed differences could have occurred by chance is tested and the familiar p value is obtained. By convention, if a result is obtained which could only have occurred by chance once in twenty times this is judged to

be "significant". Once in twenty is the same as five times in a hundred and this is expressed as p = 0.05.

For example, in a randomized controlled trial there is found to be fewer deaths with treatment A than with treatment B. We need to know whether this result could have occurred by chance. Our statistician with the computer software tells us that the p value is 0.001. This means that there is only a 1 in 1000 chance of that result occurring by chance and there is a significant difference between our treatments.

The usefulness of p values is limited on some occasions. A result of 0.049 is by convention significant (since it is less than 0.05), whereas one of 0.051 is by convention not significant. Clearly that is nonsense. The second problem is that the magnitude of the differences between treatments is not explained by probability. No statistical test can definitely prove anything. All statistics can do is quantify the likelihood that the observed result is a real effect rather than due to chance. Clinical significance should always be considered as well as statistical significance.

Confidence Intervals.

The confidence interval (CI) around a result observed in a sample of patients in a study indicates the range of values within which it is fairly certain (usually 95% certain) that the result of the same intervention applied to the true population would lie.

For example, we have seen that the results of the 4S study show that we need to treat 30 patients with established ischaemic heart disease with simvastatin for 5.4 years to save one life. If we apply confidence intervals to the data we get 95% CI of 18 - 80. In other words, if we use simvastatin in the same way as the 4S researchers did, in the population as a whole we will save one life for somewhere between every 18 and - every 80 people treated.

ADVANTAGES AND DISADVANTAGES OF DIFFERENT TYPES OF STUDY

We need to look for studies that produce the strongest evidence in order to provide valid answers to clinical questions. This means reducing the biases which in turn means that a well designed randomized clinical trial will always be the preferred type of study. Enthusiasts of evidence based medicine will often only consider in their systematic reviews evidence from RCTs and reject results from other types of study. Experience has shown that many interventions adopted on the basis of evidence other than a well done RCT has subsequently been shown to be harmful when that RCT is done [8].

However, the non-experimental case control and cohort studies clearly do have a place. They are often the only methods that are applicable to determine adverse effects - it would be unethical to conduct a randomised controlled trial in which the investigators were to expose the active group of participants to something that was likely to do them harm. For example, imagine in the 1950s discovering for the first time that there was evidence from a case series and geographical data that smoking seemed to be associated with lung cancer. Would it be ethical to take 200 medical students and get half of them to smoke 20 cigarettes a day for thirty years and get the other half to be non-smokers? A much more sensible approach would be to construct a cohort or case control study and reduce the possibility of an erroneous result by limiting the potential for bias. Case control studies are also particularly useful to analyse rare disorders.

However, a clear hierarchy of evidence exists with RCTs providing the strongest evidence, next come cohort studies and then case control studies. Figure 6 illustrates the influence of bias in studies of the effectiveness of beast screening.

Fig.6. Breast cancer mortality in studies of breast cancer screening. Women aged 50 and over (55 in Malmo, 45 in UK). Sources [9, 10].

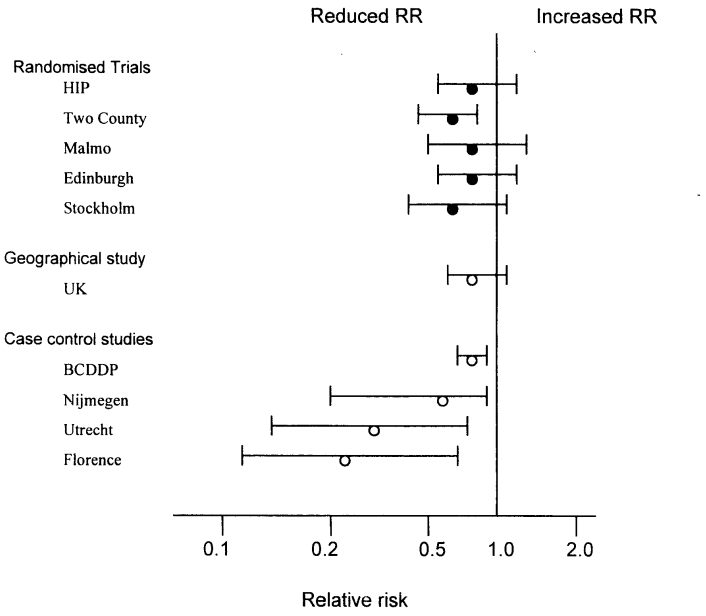

Figure 6 shows that all the studies have a relative risk of less than 1, i.e. screening produces a protective effect - a reduction in mortality in the screened women. The horizontal lines and bars indicate the 95% confidence limits. It will be noticed immediately that the geographical and case control studies show greater benefits than the randomized trials. Three biases operate to produce this effect:

- lead time bias - screening advances the date of the diagnosis and hence the survival time, although the date of death is not altered by the earlier detection.
- length time bias - the preferential detection of slowly growing tumours.
- selection bias - tendency for people who volunteer for screening to be atypical of
- the population from which they come.

Selection bias is removed by randomisation whilst the others remain. These details are not in themselves important but they illustrate why caution is required when considering evidence other than from randomised trials.

If a study is well designed and conducted the results can be considered to be valid. However, other factors such as the subjects being equivalent to those in one's own practice, the result being clear statistically (and in addition having clinical significance) and the setting of the study (e.g. in a health care system similar to the NHS) will influence whether the results are generalisable.

SCREENING FOR DISEASE

There is an inherent attraction in being able to detect a disease at an early stage that will lead to a greater proportion of the detected cases being successfully treated. With cervical and breast screening already well established and advocates for prostate and colonic cancer becoming ever more vocal, there is an increasing need to look closely at current practice and judge it and future programmes against explicit criteria of benefit for individuals and for the population.

"Wilson's Criteria". (Wilson and Jungen. World Health Organisation, 1968)

1. Condition should be common and important / serious
2. Natural history of disease understood with latent period in which disease can be
3. detected
4. Successful treatment by an agreed method available when detection occurs

5. Screening test should be safe, acceptable to patients, screening to be continuous, on a group agreed to be high risk
6. Test should be cheap (or at least cost-effective)
7. Screening programme delivered via an agreed policy

		DISEASE	
		POSITIVE	NEGATIVE
TEST	POSITIVE	a	b
	NEGATIVE	c	d

A successful screening test will have few false negatives (sensitivity = a/a+c) and few false positives (specificity = d/b+d). Judging screening programmes is further complicated by lead time bias, length time bias, and compliance bias. Not surprisingly the Chief Medical Officer has accepted the first recommendation of the new (1996) National Screening Committee - no new local screening programmes are to introduced in the UK until rigorous evaluations are made available by this expert group.

References.

1. D.L.Sackett, R.B. Haynes, G.H. Guyatt, P. Tugwell. Clinical Epidemiology. A Basic Science for Clinical Medicine. 2nd. Edition. 1991. Little, Brown.

2. Chalmers, M. Edkin and M.J.N.C. Keirse (eds). Effective Care in Pregnabcy and Childbirth (Vols 1&2). Oxford. Oxford University Press. 1989.

3. M.C.Jensen et al. Magnetic resonance imaging of the lumbar spine in people without back pain. New Engl J Med 1994; 331: 69-73.

4. TR Dawber. The Framingham Study. The Epidemiology of Atherosclerotic Disease. Cambridge, MA. Harvard University Press, 1980.

5. T.R.Pedersen et al. Randomised trial of cholesterol lowering in 4444 patients with coronary heart disease: the Scandinavian Simvastatin Survival Study (4S). Lancet 1994; 344: 1383-89.

6. Bandolier. No 17. 1994. Oxford.

7. Frances Clegg. Simple Statistics. 1982. Cambridge.

8. The Cardiac Arrythmia Suppression Trial (CAST) Investigators. Special Report. Preliminary Report: Effect of ecainide and flecainide on mortality in a randomised trial of arrythmia suppression after myocardial infarction. N Endl J Med 1989; 321: 406.

9. Breast Screening Advisory Committee. Breast Cancer Screening: Evidence and Experience since the Forrest Report. 1991. NHSBSP.

10. NE Day. Screening for Breast Cancer. British Medical Bulletin. 1991; 47: 400-15.

Neal Maskrey. February 1997