THE FRAMEWORK FOR COST-EFFECTIVENESS

The Rationale

In all societies, resources are scarce in the sense that there will never be a sufficient number to produce all the goods and services that people value. This is not just a problem of the public sector, but is common to all types of activities regardless of whether they are public or private. For example, time can be considered the ultimate scarce resource from an individual perspective. There is never enough time to do all the things that we would like to do, so we have to choose between the myriad of competing ways of spending time. It is the same with any resource - land, capital, labour etc. A resource could be used in many ways that would produce benefits for people and choices between alternative uses.

A form of analysis known as economic appraisal has been developed to determine if a proposed use of a scarce resource is the "best". It has been developed largely to consider the best use of public sector resources where "best" is defined in two ways (Evans & Hurley 1995). The first is linked to the concept of allocative efficiency. Under this definition, "best" implies using the resources to produce the mix of goods and services which society would value most highly. This requires a comparison of a wide range of different possible uses of those resources - would it be better to subsidize the treatment of tuberculosis, to increase the distribution of impregnated mosquito nets, or to construct a tertiary care facility in an isolated community, for example.

A narrower, but complementary version of "best" is linked to the concept of productive efficiency. It relates to the fact that there are always alternative ways of meeting a specific goal (i.e. reducing the burden imposed by malaria) and provides a way to determine which of the alternative strategies would achieve this goal more effectively for a given expenditure of resources. Another way of saying the same thing is that this approach identifies the strategy which uses the fewest resources (the least cost method) to produce a benefit of a given size (i.e. reduce the number of deaths due to malaria by 10%).

Both types of questions about which is the best use of resources are important and economic appraisal techniques have been developed for both questions. All of them require an explicit or implicit comparison between alternatives. In the health sector cost-benefit analysis (CBA), cost-effectiveness analysis (CEA), and cost-utility analysis (CUA) have been used most frequently.

Defining technical terms

All forms of economic appraisal require calculating the resources used by an intervention (inputs) and the benefits (outputs) that are produced. CBA focuses on allocative efficiency - it examines whether the proposed use of resources is better than all possible alternatives. It does this by valuing the outputs and inputs in money terms. The costs, at least in theory, are the benefits which would have been earned, and which now must be foregone, if the inputs had been used in the best alternative way. If the benefits of the proposed alternative exceed the costs (i.e. they exceed the benefits that could have been obtained by using the resources elsewhere), the proposed alternative is allocatively efficient and the investment should be made. There is a long history of attempts to apply CBA to health interventions but there are major problems - not the least of which is the difficulty of measuring health improvements - the benefit of an intervention - in money terms (Haddix et al. 1996; Gold et al. 1996; Drummond et al. 1997). Regardless of which method of valuing benefits is used (the human capital or the willingness-to-pay approach), there is a fundamental bias favouring interventions which improve

the health status of the rich over that of the poor. The ethical implications of this form of appraisal explain the observation that CEA and CUA have been used more commonly than CBA in the health sector and the rest of this chapter will focus on those techniques.

In CEA, costs are measured in the same way as in CBA but benefits are measured in physical units designed to capture the improvement in health status. Different types of outcome indicators are used to reflect the health improvement. Final outcome indicators represent the actual improvement in health (i.e. the number of lives or years of lives saved). Intermediate outcome indicators reflect a stage in the process of improvement which is believed to be directly proportional to the final outcome (i.e. the number of children fully immunized reflects the success of vaccination strategies on the grounds that the resultant improvement in population health is directly proportional to the number of children immunized).

CEA requires explicit comparison of alternatives. The cost-per-year of life saved by two programs, or the cost-perchild fully immunized of two strategies, are compared. CEA using final outcome indicators can explore both allocative and production efficiency. For example, it can be used to compare a program to introduce insecticide impregnated mosquito nets with a program to extend childhood immunization coverage, in terms of the cost-peryear of life saved (allocative efficiency). On the other hand, CEA using intermediate outcome indicators can be used only to compare alternative ways of reaching a specified more limited goal - it explores production efficiency only. For example, the cost-effectiveness of mass treating children for helminth infections without first seeking to learn their individual infection status could be compared with the cost-effectiveness of a program where children were first screened and only those found to be infected are treated. The alternatives could be compared in terms of the cost per infected child treated, for example. In principle, the alternative with the lower cost per unit of outcome is the more efficient, although this can at times be misleading. This is discussed in more detail later in this chapter and additional information can be found in some of the references (i.e. Drummond et al. 1997).

A limitation of CEA as defined above is that the final outcome indicators that are the most commonly used reflect improvements in the duration of life. Health interventions can improve quality of life as well as, or instead of, duration and it has traditionally been difficult to incorporate the improvements in both quantity and quality into one outcome indicator. Recently, QALYs, DALYs and a number of variations of the theme (ie: Healthy Year Equivalents) have been developed to do this. Outcome indicators are described more fully in another chapter in this volume. However, CEA based on an outcome indicator such as a QALY, which combines the impact of an intervention on both mortality and morbidity (or disability) has become known as cost-utility analysis (CUA). These outcome indicators involve a weighting system in which years of death prevented are added to years of life in which quality has been improved using a weighting system. Years of life saved are valued more highly (usually given a weight of one) while years of health improvement are given a weight of between zero and one. The greater the improvement in quality of life, the closer is the weighting of a year of health improvement to one. The weights used to derive QALYs were called utilities - hence the term cost-utility analysis. CUA can be regarded as a special form of CEA so the term CEA is used as a generic term to describe all forms of cost-effectiveness comparisons, including CUA, in the remainder of this chapter. More details of the differences between QALYs and DALYs can be found in the references at the end of the chapter (i.e. Murray 1996).

Levels of Analysis

CEA can be used at national, sub-national and facility levels. At the national level, some governments have defined a basic package of services to which they would like to ensure that all people have access. CEA has been used as one of the determinants of which interventions to include and which to exclude on the grounds that those which improve public health standards more efficiently should be given preference (World Bank 1993). Under decentralization, where responsibility for planning and/or financing services has been transferred to lower levels (i.e. districts), health planners can decide if schistosomiasis control is better undertaken in terms of mass treatment regardless of infection status, or whether it would be more efficient to test individuals for the presence of disease before treatment. At the facility level, hospitals have asked whether it is more cost-effective to undertake specific types of surgery on an outpatient or inpatient basis. Recently, CEA has also been used to help set global research priorities in health (Ad Hoc Committee on Health Research Relating to Future Intervention Options 1996).

Uses and Misuses of CEA

CEA is a way of explicitly identifying and measuring the costs expected to be incurred undertaking an intervention and the benefits that are expected to accrue. This process reveals the areas where good evidence of effectiveness is available and makes it less easy to make decisions on personal hunches or biases. However, accurate data on costs and/or effectiveness often are not available, and estimates have to be made using "best guesses". In such cases, sensitivity analysis can help to determine how robust the results are to changes in critical assumptions.

It should, however, be remembered that CEA is only one input, although an important one to decision making. The calculations do not take equity into account explicitly, for example, nor do they include other factors that a society or government might consider important determinants of social welfare. There are other problems in interpreting the results that will be discussed later, but overall, the technique is a valuable input to the development of evidencebased health policy.

THEORETICAL CONSTRUCTS

Introduction

This chapter does not attempt to be a "how to" manual. Many such manuals exist and a step-by-step guide to undertaking a CEA would require a volume rather than a chapter. One of the best guides is Phillips et al. (1993) but even this volume does not discuss in detail all the theoretical constructs and controversies which surround CEA. These issues require an even longer document. They are discussed well in Gold et al. (1996) and only the most important principles are covered here.

Steps in CEA analysis

We have divided CEA into five main steps, although it would be equally appropriate to further subdivide them.

Step 1: Define Alternatives

All CEA requires a comparison of alternatives so it is important to specify clearly those under consideration. Most commonly, this involves a straight comparison of alternatives (i.e. mass treatment versus selective treatment for helminth infections), but it is also possible to compare a proposed intervention with the option of doing nothing. The analyst should clarify the nature of the interventions and the types of people who are expected to benefit from them, the current status quo for approaching the disease/condition, and the exact manner in which the intervention will be organized (Drummond et al. 1997)

Step 2: Identify and Measure Costs

CEA is most appropriate for decisions about how best to use public resources (Evans & Hurley 1995). Accordingly, the viewpoint guiding the overall decision about cost-effectiveness should be that of society as a whole and all costs, regardless of who pays them, should be included. For budgeting purposes; however, governments might wish to know how much an intervention will cost them over time and they would also estimate costs from the narrower perspective of the government. And in some cases, it might even be appropriate to estimate costs from the perspective of only one part of government (i.e. in countries where drugs are funded by a separate part of government to those which fund the rest of the health sector). In this case, the costing exercise would show the resource implications to the section of government responsible for pharmaceuticals. But these narrower viewpoints should be done in addition to, rather than in place of, the cost-effectiveness analysis from the perspective of society as a whole.

Costs are traditionally divided into direct and indirect costs. Direct costs are the resources used up because of an intervention - drugs, hospitalisation, transport, food etc. Indirect costs are the reductions in productive potential because of an intervention. Direct costs can be incurred by different players, and should be identified and measured separately before an attempt is made to combine them. For example, some costs might be incurred by the government, some by households, and some by religious or non-government organizations. It is important not to double count costs when government subsidizes some of them. There is dispute in the literature about the extent to which indirect costs (for example, the value of time used in seeking care) should be included in the analysis. Here, we recommend that indirect costs directly related to utilization of the intervention under discussion should be included, though there are complications in undertaking this in practice. For a more complete discussion of the options, see Gold et al. (1997).

Direct costs are further divided into capital and recurrent costs. Capital items are those which provide a flow of benefits for more than one time period, usually defined as a year. The economic costs of such items include not only the rate at which the equipment deteriorates, but the potential benefits forgone by investing resources in the capital goods rather than elsewhere. Capital costs can be expressed as a single lump sum but often it is more convenient to express them as an annual value over the life of the capital item. There are well-established formulae for deriving the annual equivalent value of a capital item based on the replacement cost, the expected lifetime of the good, and the discount rate (see Phillips et al. 1993; Drummond et al. 1997).

Recurrent (or operating) costs are the value of the resources used each year in providing the intervention. They include labour costs as well as the cost of items used up during the year including drugs, fuel and other transport costs, etc. Methods of itemizing and valuing these costs are provided in Creese & Parker (1993) and Phillips et al. (1993).

Some costs, both capital and recurrent, are not unique to one intervention. For example, a malaria control program might share offices and administrative staff with other programs aimed at the control of other communicable diseases. In undertaking a CEA of malaria control strategies, some of the time of the administrative staff and some of the office space should be charged to the malaria interventions. There are a number of methods of allocating shared (or overhead) costs of this nature which are well described in Drummond et al. (1997).

Related to the concept of shared costs is the fact that the unit costs of expanding or contracting a service - known as marginal costs - need not be the same as the average costs of the current intervention. The average cost of a day in hospital includes a component of shared costs, such as the costs of buildings, equipment and administrative staff. These costs are often simply divided by the number of patient days to determine the shared cost per bed. Shared costs are then added to any specific costs related to the bed - drugs, investigations, laundry etc. - to obtain the costs of each bed. If the hospital was already operating at full capacity, expansion would require the construction of a new facility and the costs per additional bed might be considerably higher than the currently observed average cost. On the other hand, expansion would be considerably cheaper than the current average cost where the hospital was operating below full capacity as no additional shared costs would be required per additional bed. Care should be taken to determine if marginal costs are likely to differ to average costs, and to make the appropriate adjustments if necessary.

Most interventions are implemented over a number of years and costs are incurred on a yearly basis. People generally prefer to receive $1000 today than in one year, and to pay $1000 in one year than today, partly because of the potential earnings available if the money was invested for the intervening year. Clearly, $1000 does not equal $1000 but value is also dependent of when costs are incurred. The further into the future the $1000 cost must be paid, the better it is from the perspective of the payer - or the less it is worth. Discounting is the method used to translate future costs into their present values. Formulae and tables can be found in any text (i.e. Drummond et al. 1997). Discounting of costs is routinely undertaken and is not controversial. The rate at which future costs should be discounted is more controversial with the general consensus being that somewhere between 2.5% and 5% is appropriate when real (inflation free) rather than nominal (allowing for inflation) costs are used (Gold et al. 1996). Gold et al. recommend using both 3% and 5% in a sensitivity analysis.

Step 3: Identify and Measure Outcomes

It is necessary to distinguish effectiveness from efficacy. Efficacy describes the percentage improvement in outcome under ideal conditions while effectiveness is the percentage improvement in reality, taking into account less than perfect compliance and coverage. Results from randomised clinical trials reflect efficacy rather than effectiveness because they apply the interventions to a narrow spectrum of homogenous patients and efforts are taken to ensure that providers and patients comply with recommended protocols. Often the only available evidence relates to efficacy, particularly for new tools. In CEA, these estimates should be modified to allow for probable reduced levels of compliance and coverage when a tool or strategy is introduced in field conditions.

In CEA, effectiveness is measured in terms of physical units. As stated earlier, two types of outcome indicators are typically described. Intermediate outcome indicators reflect an aspect of the physical production of the intervention that is believed to be directly proportional to the final improvement in health status. Examples include: the number of children fully immunized (for

comparison of ways of delivering vaccinations) or the number of cases of schistosomiasis treated (for comparison of strategies to deliver praziquantel, an anthelminthic). Intermediate outcome indicators are useful for evaluating productive efficiency (i.e. for a comparison of interventions aimed at the same disease or condition). They are rarely useful in the consideration of allocative efficiency or in the comparison across different types of interventions - it makes no sense to compare the cost per fully immunized child of a vaccination program with the cost per infected child treated in an anti-schistosomiasis intervention.

Final outcome indicators, most commonly the number of lives saved or the number of years of life saved, are used in CEA aimed at reaching conclusions about allocative efficiency. But such indicators do not allow comparison of interventions that extend life with those that improve the quality of life. This is why indicators such as QALYs and DALYs were developed.

An area of considerable controversy relates to the concept of indirect costs discussed in the previous section. Interventions which improve health status can lead to indirect benefits which is the converse of indirect costs (i.e. to an improvement in the productive potential of individuals). Some practitioners of CEA subtract indirect benefits from the costs of the intervention and call the numerator of the cost-effectiveness ratio "net costs" (i.e. Haddix et al. 1996). We argue that this amounts to double counting the benefits. An indicator such as a QALY reflects the overall improvement in health status from an intervention that is valued by society for many reasons, one of which is the increase in productive potential. It does not make sense to subtract out the economic component of the benefit in the numerator when it is already implicitly in the denominator.

On the other hand, some interventions not only incur direct costs but they can lead to savings in direct costs (i.e. direct benefits). For example, a program to distribute impregnated mosquito nets will, hopefully, prevent cases of malaria. A proportion of those cases will have used resources being treated and these savings are usually subtracted from the direct costs of the intervention in the numerator. It is necessary to draw boundaries around this exercise. They are usually drawn to include direct benefits relating only to the disease or condition under question. Accordingly, the use of unrelated health services later in life as a result of an intervention that saves lives now, is not counted as a direct cost of the intervention. The costs of saving the life are included and the reductions in the use of services related to the condition under question are subtracted into the net costs of the numerator.

Finally, an additional area of controversy concerns whether future benefits measured in physical units should be discounted to their present values, and if so, whether they should be discounted at the same rate as costs. In a cost- benefit analysis where benefits are measured in money terms, there is no controversy and money benefits are discounted at the same rate as money costs on the grounds that society and individuals prefer money benefits sooner rather than later. However, there is no consensus about whether outcomes in terms of units such as DALYs and QALYs should be discounted. This is an area of continuing debate and there are good arguments why the discount rate for health benefits need not be identical to the discount rate for costs (Gold et al. 1996). However, current "best practice" is to discount costs and outcomes at the same rate - of somewhere between 3% and 5%.

Step 4: Interpreting Results

The results of CEA are presented in the form of a cost-effectiveness ratio showing the costs incurred per unit of "success" or benefit. In principle, the lower the (discounted) cost per unit of

(discounted) benefit, the more efficient the program. This is, however, only strictly true where one intervention is both less costly and more effective than the alternatives. In many cases, an intervention will be more costly but will produce more benefits than the alternative. If the interventions are mutually exclusive - only one can be done but not both - an incremental analysis must be performed. An incremental cost-effectiveness ratio compares the extra benefit with the extra cost that would be involved by switching from the intervention with the lowest cost-effectiveness ratio to the intervention with the next lowest ratio. In fact, it is commonly argued that incremental analysis is the gold standard because all interventions must be compared to what is currently done. [This is a slightly different question to that posed in the previous section dealing with marginal analysis, which concerned the extra costs (and effectiveness) emerging from an expansion or contraction of a single program.]

To assess production efficiency, a comparison of the alternative ways of addressing the problem will suffice but consideration of allocative efficiency requires a comparison of the proposed intervention with all possible alternative ways of using those resources. It is not possible to make the myriad of calculations required in the context of a single study and analysts and policy makers have resorted to league tables, which show the cost-effectiveness of a wide variety of alternative interventions. The cost-effectiveness of the proposed intervention (or the incremental cost-effectiveness) can be compared with the cost-effectiveness of other types of interventions that are either currently undertaken or could be undertaken. If the proposed intervention is more efficient than alternatives currently undertaken, the implication is that population health status could be improved by switching resources from the less efficient to the more efficient options.

The problems with using league tables are well known (Gold et al. 1997) and often arise because different methods have been used in the studies that are included. Another problem is that incremental analyses are not easily transferable between settings because they are dependent on the baseline interventions currently being undertaken. League tables showing the cost-effectiveness of a set of interventions reported in terms of cost per life saved exist for the USA (Tengs et al. 1995; Tengs 1996), and another set can be constructed in terms of QALYs for interventions in developed countries (Drummond et al. 1997). The only large study comparing the costeffectiveness of interventions across both developed and developing countries measures outcomes in terms of DALYs (Jamison et al. 1993).

Although incremental analyses clearly show if a proposed change from the present situation would be an efficient use of scarce public resources, it does not explicitly question whether what is presently done is efficient. It can be shown that an intervention which improves efficiency from the standpoint of current interventions, would not be chosen if the policy makers had the ability to re-evaluate their entire portfolio. In some counties, health sector reforms do allow a broader examination of which portfolio of interventions a government would wish to subsidize and in such cases, CEA based on average costs and outcomes is a good guide to the mix of interventions which would achieve the largest improvement in population health status. But if it is not possible to reconsider what is currently done, incremental analysis is the gold standard.

Step 5: Account for Uncertainties

Because many of the critical parameters in CEA are not known with certainty, or may take a range of values, sensitivity analysis can be used to determine the extent to which the results are robust to changes in critical parameter values. Computer programs increasingly allow analysts to undertake multivariable sensitivity analysis showing the range of values of a cost-

effectiveness ratio when many parameters vary at the same time, although traditionally studies have considered changes in each critical parameter separately.

PRACTICAL CONSIDERATIONS

Undertaking CEA is not costless. It requires technical skills, financial resources, and might require field studies to collect data. Care needs to be taken to determine if a field study is required as it is not possible even for the wealthiest countries to conduct all the studies necessary to provide a complete database on the cost-effectiveness of all available interventions, and to keep it updated over time.

The results of studies undertaken in other settings are one way of reducing the cost. Often effectiveness (or efficacy) data will be available from field or clinical trials and there may be no physiological or epidemiological reason why efficacy should vary across sites. Adjustments might need to be made for local variations in compliance and coverage, or for a different group of patients likely to avail themselves of the intervention. Costs will generally differ across sites, but estimates of local costs can often be made from studies published in other areas if those studies provide details of the physical use of resources as well as their money values or prices. The analyst can determine if similar levels of physical resource use will be required in the local setting and can use local prices to obtain costs. For this reason, we strongly recommend that all studies should report costs broken down by the quantities of resources used and the unit prices. A variation of this theme is that the relative cost-effectiveness of alternatives may have been compared in another setting, and the analyst might feel there is no reason why the relative numbers should vary across settings.

At other times, there may be no data on effectiveness or efficacy, or there may be reasons for thinking that physiological or epidemiological factors differ across sites. For example, the impact of impregnated mosquito nets might depend on the type of vector (mosquito) and its biting patterns, as well as the seasonal patterns of transmission. In such cases, field trials might be required. It is difficult to collect good economic data as part of a field trial of efficacy or even effectiveness, as the trial itself can change behaviours dramatically. For example, a trial of the impact of ivermectin on lymphatic filariasis transmission requires patients to have their infection status measured at different intervals. If they are sick, they would be treated by the research team as a matter of ethics, precluding the ability to determine "normal" treatment practices and their costs. A substantial literature exists on how to collect cost data from drug trials in the North so as to minimize bias. Even though it does not specifically relate to the south, some of the lessons are important for all studies that seek to measure costs as part of a trial (Coyle et al. 1998).

Costing studies have unfortunately been sometimes regarded as technically undemanding, and many examples can be found of unsound analyses. Incorrect assumptions and methods can lead to order of magnitude differences in outcomes and can bias policy decisions in inappropriate directions. Moreover, proponents of a particular course of action sometimes do studies, and their methodologies should be checked carefully. It is, therefore, important that studies clearly specify their assumptions. It would also be valuable for analysts to be required to make the spreadsheets used in their calculations available to a wider audience. If this was expected as part of routine practice, it would act as a form of quality control but would also help analysts trying to transfer the results of studies undertaken in one setting to their local setting.

REFERENCES - COST EFFECTIVENESS

(Ordered by tool number in the table)

Tool 1
Canadian Coordinating Office for Health Technology Assessment), 1997. **Guidelines for Economic Analysis of Pharmaceuticals: Canada.** 2nd. ed., Canadian Coordinating Office for Health Technology Assessment (CCOHTA), Ottawa.

Tool 2
Commonwealth Department of Human Services and Health, 1995. **Guidelines for the Pharmaceutical Industry on Preparation of Submissions to the Pharmaceutical Benefits Advisory Committee, Including Submissions Involving Economic Analyses.** 2nd. ed., Australian Government Publishing Service, Canberra.

Tool 3
Drummond MF, Jefferson TO, on behalf of the British Medical Journal Economic Evaluation Working Party, 1996. **Guidelines for authors and peer reviewers of economic submissions to the BMJ.** British Medical Journal, 313:275-283.

Tool 4
Phillips M, Mills A & C Dye, 1993. **Guidelines for Cost-Effectiveness Analysis of Vector Control.** World Health Organization, Geneva (WHO/CWS/93.4).

Tool 5
Gold RM, Siegel JE, Russell LB & Weinstein MC, 1996. **US Consensus Statement: How to undertake cost-effectiveness analysis in health and medicine.** Cost Effectiveness in Health and Medicine. Oxford University Press, New York.

Tool 6
Haddix AC, Teutsch SM, Shaffer PA, & Dunet DO, 1996. **Prevention Effectiveness: A Guide to Decision Analysis and Economic Evaluation.** Oxford University Press, New York.

Tool 7
Drummond MF, O'Brien BJ, Stoddart GL & Torrrance G, 1997. **Methods for the Economic Evaluation of Health Care Programmes.** 2nd. ed., Oxford University Press. New York.

Tool 8
Tengs TO, Adams M, Pliskin JS, Safran DG, Siegel JE, Weinstein MC & Graham JD, 1995. **"Five-hundred lifesaving interventions and their cost-effectiveness"**, Risk Analysis, 15:369-390.

Tool 9
Jamison DT, Mosley W.H., Measham AR & Bobadilla J.L., 1993. **Disease Control Priorities in Developing Countries.** Oxford University Press, New York.

Commonwealth Department of Human Services and Health (CDHSH), 1994. **Manual of Indicators.** Australian Government Publishing Service, Canberra, Australia.

Evans D, Freund D, Dittus R, Robertson J & Henry D, 1993. **Background document on the use of economic analysis as a basis for inclusion of pharmaceutical products on the Pharmaceutical Benefits Scheme.** Australian Government Publishing Service, Canberra.

Weinstein MC, Siegel JE, Gold RM, Kamlet MS, Russell LB for the Panel on Cost-Effectiveness in Health and Medicine, 1996. **Recommendations of the Panel on Cost-Effectiveness in Health and Medicine**, Journal of the American Medical Association, 276:1253-1258.

Tengs TO, 1996. "Enormous variation in the cost-effectiveness of prevention: implications for public policy" **Current Issues in Public Health,** 2:13-17.

Additional References:

Ad Hoc Committee on Health Research Relating to Future Intervention Options, (1996): **Investing in Health Research and Development**. WHO, Geneva (Document TDR/Gen/96.2)

Coyle D., Davies L. & Drummond M.F., 1998. "Trials and tribulations. Emerging issues in designing economic evaluations alongside clinical trials", *International Journal of Technology Assessment in Health Care,* 14:135-144.

Evans D.B. & Hurley S.F., 1995. "The application of economic evaluation techniques in the health sector: the state of the art", *Journal of International Development*, 7(3):503-524.

Murray C.J.L., 1996. "Rethinking DALYs" in Murray C.J.L. & Lopez A.D. (1996). **The Global Burden of Disease.** Published by Harvard University Press, Boston on behalf of WHO and the World Bank.

World Bank, 1993. **World Development Report: Investing in Health.** Oxford University Press.

December 8, 2008