

Ethical Issues in the Use of Cost Effectiveness Analysis for the Prioritization of Health Care Resources¹

Dan W. Brock
Brown University

Resources to improve health are and always have been scarce, in the sense that health must compete with other desirable social goals like education and personal security for resources.² It is not possible to provide all the resources to health, including health care and health care research, that might provide some positive health benefits without great and unacceptable sacrifices in other important social goods. This should go without saying, and in other areas of social expenditures resource scarcity is not denied, but in health care many people mistakenly persist in denying this fact. It follows from resource scarcity that some form of health care rationing is unavoidable, where by rationing I mean some means of allocating health care resources that denies to some persons some potentially beneficial health care. That rationing may take many forms. In the United States much rationing is by ability to pay, but in both public programs like the Oregon Medicaid program and in many private managed care plans more systematic efforts to prioritize health care resources have been carried out.

To many health policy analysts it is an unquestioned, and so generally undefended, assumption that in the face of limited health care resources, those resources should be allocated so as to maximize the health benefits they produce, measured by either the aggregate health status or disease burden of a population. Cost effectiveness analysis (CEA) that compares the aggregate health benefits secured from a given resource expenditure devoted to alternative health interventions is the standard analytic tool for determining how to maximize the health benefits from limited resources. Natural, even self-evident, as this maximization standard may appear to many health policy analysts and economists, it assumes a utilitarian or consequentialist moral standard, and more specifically standard of distributive justice, and the utilitarian account of distributive justice is widely and I believe correctly taken to be utilitarianism's most problematic feature.

Cost effectiveness analysis comparing alternative health interventions in the quality-adjusted life years (QALYs) produced from a given level of resources constitutes a quantitative method for prioritizing different interventions to improve health. There are many unresolved technical and methodological issues in QALYs and CEA, none of which will be my concern here. My concern will be instead with the ethical issues in the construction and use of CEAs for the prioritization of health care resources. The specific issues that I shall briefly discuss below all constitute potential ethical criticisms of CEA as a normative standard, specifically criticisms concerning justice or equity, and so one might hope concerns for justice or equity could be integrated into these quantitative methodologies. There are at least two reasons, however, for caution, at least in the near term, about the possibility of integrating some of these ethical concerns into cost effectiveness models and analyses. First, although a great deal of work in economics and health policy has gone into the development and validation of measures of health status and the burdens of disease, as well as of cost effectiveness methodologies, very little work has been done on how to integrate concerns of ethics and equity into cost effectiveness measures, although I shall mention one means of doing so later. The theoretical and methodological work

necessary to do so remains largely undone. Second, each of the issues of ethics and equity that I take up below remain controversial. Since no clear consensus exists about how each should be treated, there is in turn no consensus about what qualifications or constraints they might justify placing on the cost effectiveness goal of maximizing health.

This second difficulty is not likely to be only a near term limitation, awaiting further work on the ethical issues that I will identify. Instead, these issues represent deep divisions in normative ethical theory and in the ethical beliefs of ordinary people; I believe they are likely a permanent fact of ethical life. As I understand and shall present these ethical issues, in most cases there is not a single plausible answer to them. Even from within the standpoint of a particular ethical theory or ethical view, these issues' complexity means that different answers may be appropriate for a particular issue in the different contexts in which CEAs are used. Thus, what is necessary at this point is work developing more clearly and precisely the nature of the issues at stake, the alternative plausible positions on them together with the arguments for and against those positions. Until much more of this work is done, we will not know how deep the conflicts go and the degree to which any can be resolved.

Norman Daniels and James Sabin have recently argued that because ethical theories and theories of justice are indeterminate and/or in conflict on some of these issues, we must turn to fair procedures to arrive at practical solutions to them for health policy (Daniels and Sabin, 1997, pp. 303-350.). As practical policy matters that need resolution now they are no doubt correct, and a single quantitative measure or model of equity and justice for health care resource prioritization is certainly not possible now, if it will ever be. But that is not to deny that much important work remains to be done on the substantive issues of equity in health care, and that work should inform the deliberations of those taking part in the fair procedures that we will need to reach practical resolutions and compromises on these issues in real time. What then are some of the main issues of equity raised by cost effectiveness approaches to resource allocation of health care?

First issue: How should states of health and disability be evaluated?

Any CEA in health care requires some summary measure of the health benefits of interventions designed to improve the health status and reduce the burden of disease of a given population. Early summary measures of the health status of populations and of the benefits of health interventions often assessed only a single variable, such as life expectancy or infant mortality. The usefulness of life expectancy or infant mortality rates is clearly very limited, however, since they give us information about only one of the aims of health interventions, extending life or preventing premature loss of life, and they provide only limited information about that aim. They give us no information about another, at least as important, aim of health interventions, to improve or protect the quality of life by treating or preventing suffering and disability.

Multi-attribute measures like the Sickness Impact Profile (Bergner, Bobbitt, Carter and Gibson, 1981, 787-805) and the SP 36 (Ware and Sherbourne, 1992, pp.473-483) provide measures of different aspects of overall health related quality of life (HRQL) on which a particular population can be mapped, and an intervention assessed for its impact on these different components of health, or HRQL. Since these measures do not assign different relative value or importance to the different aspects or attributes of HRQL, they do not provide a single overall summary measure of HRQL. Thus, if one of two populations or health interventions scores higher in some respect(s) but lower in others, no conclusion can be drawn about whether the overall HRQL of one population, or from one intervention, is better than the other. Much quantitative based resource prioritization requires a methodology that combines in a single measure the two

broad kinds of benefits produced by health interventions--extension of length of life and improvements in the quality of life (Brock, 1992).

Typical summary measures of the benefits over time of health interventions that combine and assign relative value to these two kinds of benefits include QALYs and Disability-Adjusted Life Years (DALYs). Measures like QALYs and DALYs require a measure of the health status of individuals and in turn populations at different points in time, such as the Health Utilities Index (HUI) (Torrance, et.al., 1992) and the Quality of Well-Being Scale (QWB) (Kaplan and Anderson, 1988, PP. 203-235), so as to be able to measure the health benefits in terms of changes in HRQL and length of life produced by different health interventions. The construction of any measure like the HUI or QWB requires a two step process: first, different states of disability or conditions limiting HRQL are described; second, different relative values or utilities are assigned to those different conditions.

The determination of a person's or group's different health related conditions in terms of the various areas of function on the HUI or QWB both before and after a particular health intervention is an empirical question, which should be answered by appeal to relevant data regarding the burden of a particular disease and the reduction in that burden that a particular health intervention can be expected to produce. Needless to say, often the relevant data are highly imperfect, but that is a problem to be addressed largely by generating better data, not by ethical analysis.

The second step of assigning different relative values or utilities to the different areas and levels of function described by a measure like the HUI is typically done by soliciting people's preferences for life with the various functional limitations. This raises the fundamental question of whose preferences should be used to determine the relative value of life with different limitations in function and how they should be obtained. The developers of the DALY used the preferences of expert health professionals, in part for the practical reason that they are more knowledgeable about the nature of different health statuses, but the degree to which various conditions reduce overall HRQL is not a matter to be settled by professional expertise. Moreover, health professionals may have systematic biases that skew their value judgments about quality of life from those of ordinary persons. Other measures like the HUI and QWB use the value judgments of a random group of ordinary citizens to evaluate different states of disability or limitations in function. The utilities so determined for different functional attributes and their levels in the HUI are shown in Figure 1.

A central issue concerning whose evaluations of different states of disability or functional limitation should be used arises from the typical responses of individuals to becoming disabled: adaptation, that is improving one's functional performance through learning and skills development, coping, that is altering one's expectations for performance so as to reduce the self-perceived gap between them and one's actual performance; and adjustment, that is altering one's life plans to give greater importance to activities in which performance is not diminished by disability (Murray, 1996). The result is that the disabled who have gone through these processes often report less distress and limitation of opportunity and a higher quality of life with their disability than the non disabled in evaluating the same condition. Moreover, those valuations will assign less value to extending the lives of persons with disabilities. If the evaluations of disability states by the non disabled are used for ranking different states of health and disability, then disabilities will be ranked as more serious health needs, but these rankings are open to the charge that they are distorted by the ignorance of the evaluators of what it is like to live with the conditions in question. If the evaluations of the disabled themselves are used, however, the rankings are open to the charge that they reflect a different distortion by unjustifiably

underestimating the burden of the disability because of the process of adaptation, coping, and adjustment that the disabled person has undergone. Moreover, they will assign less value to prevention or rehabilitation for disability because of the results of this process. The problem here is to determine an appropriate evaluative standpoint for ranking the importance of different disabilities which avoids these potential distortions (Brock, 1995, pp. 159-184).

Since the preferences for different states of disability or HRQL used to determine their relative values should be informed preferences, it is natural to think that the preferences of those who actually experience the disabilities should be used. Because they should have a more informed understanding of what it is actually like to live with the particular disability in question, we can hope to avoid uninformed evaluations. But this is to miss the deeper nature of the problem caused by adaptation, coping, and adjustment to disabilities.

Fundamental to understanding the difficulty posed by adaptation, coping, and adjustment to disabilities for preference evaluation of HRQL with various disabilities is that neither the nondisabled nor the disabled need have made any mistake in their different evaluations of quality of life with that disability. They arrive at different evaluations of the quality of life with that disability because they use different evaluative standpoints as a result of the disabled person's adaptation, coping, and adjustment. Disabled persons who have undergone this process can look back and see that before they became disabled they too would have evaluated the quality of life with that disability as nondisabled people now do. But this provides no basis for concluding that their pre-disability evaluation of the quality of life with that disability was mistaken, and so in turn no basis for discounting or discarding it because mistaken. The problem that I call the perspectives problem is that the nondisabled and the disabled evaluate the quality of life with the disability from two different evaluative perspectives, neither of which is mistaken. It might seem tempting to use the non-disabled's preferences for assessing the importance of prevention or rehabilitation programs, but the disabled's preferences for assessing the importance of life-sustaining treatments for the disabled, but this ignores the necessity of a single unified perspective in order to compare the relative benefits from, and prioritize, the full range of different health interventions.

Moreover, what weight to give to the results of coping with one's condition may depend on the causes of that condition, for example disease or injury that are no one's fault as opposed to unjust social conditions. Most measures of HRQL include some measure of subjective satisfaction or distress, a factor that is importantly influenced by people's expectations. In a society which has long practiced systematic discrimination against women, for example, women may not be dissatisfied with their unjustly disadvantaged state, including the health differences that result from that discrimination. The fact that victims are sufficiently oppressed that they accept an injustice as natural and cope with it by reducing their expectations and adjusting their life plans should not make its effects less serious, as measures of HRQL with a subjective satisfaction or distress component would imply.

When measures like the HUI or QWB are applied across different economic, ethnic, cultural, and social groups, the meaningful states of health and disability and their importance in different groups may vary greatly; for example, in a setting in which most labor is manual labor limitations in physical functioning will have greater importance than it does in a setting in which most individuals are engaged in non physical, knowledge-based occupations, where certain cognitive disabilities are of greater importance. Different evaluations of health conditions and disabilities as seem to be necessary for groups with significantly different relative needs for different functional abilities, but then cross-group comparisons of health and disability, and of the relative value of health interventions, in those different groups will not be possible. The health

program benefits will have been measured on two different and apparently incommensurable valuational scales. These differences will be magnified when summary measures of population health are employed for international comparisons across very disparate countries.

Some of this variability of perspective may be avoided by a focus on the evaluation of disability instead of handicap, as these are traditionally distinguished, such as in the 1980 International Classification of Impairments, Disabilities and Handicaps (ICIDH) (International Classification of Impairments, Disabilities and Handicaps, 1980). The ICIDH understands disabilities as “any restriction or lack (resulting from an impairment) of ability to perform an activity in the manner or within the range considered normal for a human being,” whereas handicap is “a disadvantage for a given individual, resulting from an impairment or disability, that limits or prevents the fulfillment of a role that is normal (depending on age, sex, and social and cultural factors) for that individual.” There will be greater variability between individuals, groups, and cultures in the relative importance of handicaps than of disabilities since handicaps take account of differences in individuals’ roles and social conditions that disabilities do not. But it is problematic whether these differences should be ignored in prioritizing health resources for individuals, groups, and societies, that is, whether disabilities or handicaps are the correct focus for evaluation.

Second issue: Do all QALYs count equally?

QALYs standardly assume that an additional year of life has the same value regardless of the age of the person who receives it, assuming that the different life years are of comparable quality. A year of life extension for an infant, a forty-year-old, and an eighty-year-old all have the same value in QALYs produced, and in turn in a cost effectiveness analysis using QALYs, assuming no difference in the quality of the year of life extension. This is compatible, of course, with using age-based quality adjustments for interventions affecting groups of different age patients to reflect differences in the average quality of life of those different groups; for example, if average quality of life in a group of patients of average age 85 is less than that of patients of average age 25, a year of life extension for the 25 year old would have greater value in QALYs than would a year of life extension for the 85 year old.

In the World Bank Study, World Development Report 1993: Investing in Health, (World Bank, 1993), the alternative DALY measure was developed to measure the burden of disease in reducing life expectancy and quality of life. Probably the most important ethical difference between QALYs and DALYs is that DALYs assign different value to a year of life extension of the same quality, depending on the age at which an individual receives it; specifically, life extension for individuals during their adult productive work years is assigned greater value than a similar period of life extension for infants and young children or the elderly. The principal justification offered for this feature of DALYs was the different social roles that individuals typically occupy at different ages and the typical emotional, physical, and financial dependence of the very young and the elderly on individuals in their productive work years (Murray, 1994).

This justification of age-based differences in the value of life extension implicitly adopts an ethically problematic social perspective on the value of health care interventions that extend life, or maintain or restore function, that is, an evaluation of the benefits to others of extending an individual's life, or maintaining or restoring his or her function, in addition to the benefit to that individual of doing so. This social perspective is in conflict with the usual focus in clinical decision making and treatment only on the benefits to the individuals who receive the health care interventions in question. Typical practice in health policy and public health contexts is more ambiguous on this point, since there benefits to others besides the direct recipient of the intervention are sometimes given substantial weight in the evaluation and justification of health

programs; for example, treatment programs for substance abuse are argued to merit high priority because of their benefits in reductions in lost work days and in harmful effects on family members of the substance abuser. This social perspective is ethically problematic because it gives weight to differences between individuals in their social and economic value to others; in so doing, it discriminates against persons with fewer dependencies and social ties, which arguably is not ethically relevant in health care resource allocation. The social perspective justifying the DALY measure is therefore ethically problematic, in a way the alternative QALY measure is not, if the value of health benefits for individuals should focus on the value to the individuals treated of the health benefits, not on the social value for others of treating those persons. The ethical difficulty here is briefly explored and further in the section below on what costs and benefits should count in a CEA.

Giving different value to life extension at different ages, however, might be justified ethically if done for different reasons. For example, Norman Daniels has argued that because everyone can expect to pass through the different stages of the life span, giving different value to a year of life extension at different stages in the life span need not unjustly discriminate against individuals in the way giving different weight to life extension for members of different racial, ethnic, or gender groups would unjustly discriminate (Daniels, 1988). Each individual can expect to pass through all the life stages in which life extension is given different value, but is a member of only one race, ethnic group, and gender. Thus, all persons are treated the same at comparable stages of their lives regarding the value of extending their lives, and so the use of DALYs would not constitute unjust age discrimination comparable to gender, ethnic or racial discrimination.

Moreover, individuals, and in turn their society, might choose to give lesser weight to a year of life extension beyond the normal life span than to a year of life extension before one has reached the normal life span based on a conception of what equality of opportunity requires, or on what Alan Williams calls the "fair innings argument" (Williams, 1997, pp. 117-132). People's plans of life and central long term projects will typically be constructed to fit within the normal life span, and so the completion of these central projects will typically require reaching, but not living beyond, the normal life span (Daniels, 1988; Brock 1989).

Third issue: What costs and benefits should count in cost effectiveness analyses of health programs?

It is widely agreed that cost effectiveness analyses in health should reflect the direct health benefits for individuals of their medical treatment, such as improving renal function or reducing joint swelling, and of public health programs, such as reducing the incidence of infectious diseases through vaccination programs. The direct costs of medical treatment and public health programs, such as the costs of health care professionals' time and of medical equipment and supplies, should also be reflected. But medical and public health interventions typically also have indirect benefits and costs. For example, some disease and illness principally affects adults during their working years, thereby incurring significant economic costs in lost work days associated with the disease or illness, whereas other disease and illness principally affects either young children, such as some infectious diseases, or the elderly, such as Alzheimer's' dementia, who in each case are not typically employed and so do not incur lost wages or lost work time from illness. Should an indirect economic burden of disease of this sort be given weight in a cost effectiveness analysis used to prioritize between different health interventions?

From an economic perspective, as well as from a broad utilitarian moral perspective, indirect benefits and costs are real benefits and costs of disease and of efforts to treat or prevent it, even if not direct health benefits and direct treatment costs; they should be reflected in the overall cost effectiveness accounting of how to use scarce health resources so as to produce the

maximum aggregate benefit. A possible moral argument for ignoring these indirect costs and benefits in health resource prioritization is grounded in a conception of the moral equality of persons. Giving priority to the treatment of one group of patients over another because treating the first group would produce indirect benefits for others (for example, other family members who were dependent on these patients) or would reduce indirect economic costs to others (for example, the employers of these patients who incur less lost work time) could be argued to fail to treat each group of patients with the equal moral concern and respect that all people deserve; in particular, doing so would fail to give equal moral concern and weight to each person's health care needs. Instead, giving lower priority to the second group of patients simply because they are not a means to the indirect benefits or cost savings produced by treating the first group of patients gives the second group of patients and their health care needs lower priority simply because they are not a means to these indirect benefits or cost savings to others. It would violate the Kantian moral injunction against treating people solely as means for the benefit of others.

In public policy we often use a notion of "separate spheres," which in this case could be used to argue that the purpose of health care and of public health is health and the reduction of disease, and so only these goals and effects should guide health care and public health programs (Kamm, 1993; Walzer, 1983). There are obvious practical grounds for the separate spheres view from the difficulty of fully determining and calculating indirect benefits and costs. But the Kantian moral argument could serve as a principled moral basis for ignoring indirect benefits and costs in a cost effectiveness analysis to be used to prioritize health resources and interventions that serve different individuals or groups.

Fourth issue: Should discount rates be applied to health care benefits?

It is both standard and recommended practice in cost effectiveness analyses, within health care and elsewhere, to assume a time preference by applying a discount rate to both the benefits and costs of different programs under evaluation, although the reasons for doing so and the proper rate of discount are controversial (Gold, 1996, Ch.7). It is important to separate clearly the ethical issue about whether health benefits should be discounted from other economic considerations for discounting, as well as to be clear why the issue is important for health policy. It is not ethically controversial that a discount rate should be applied to economic costs and economic benefits; a dollar received today is worth more than a dollar received 10 years from now because we have its use for those ten years, and there is a similar economic advantage in delaying the incurring of economic costs. The ethical issue is whether a discount rate should be applied directly to changes in life extension and well-being or health. Is an improvement in well-being, such as a specific period of life extension, a reduction in suffering, or an improvement in function, extending, say, for one year of substantially less value if it occurs twenty years from now than if it occurs next year?

Future benefits are appropriately discounted when they are more uncertain than proximate benefits. Proximate benefits, such as restoration of an individual's function, also are of more value than distant benefits if they make possible a longer period of, and thus larger, benefit by occurring sooner. But neither of these considerations require the use of a discount rate--they will be taken account of in the measurement of expected benefits of alternative interventions. The ethical question is whether an improvement in an individual's well-being is of lesser value if it occurs in the distant future than if it occurs in the immediate future, simply and only because it occurs later in time. This is a controversial issue in the literature on social discounting and my own view is that no adequate ethical justification has been offered for applying a discount rate directly to changes in health and well-being, though I cannot pursue the justifications offered by proponents of discounting here. The avoidance of paradoxes that arise if a discount rate is applied to costs and

the same discount rate is not applied to benefits, has influenced many economists to support use of the same discount rate for costs and benefits (Keeler and Cretin, 1983, pp. 300-306), but I believe these are properly dealt with not through discounting, but rather through directly addressing the ethical issues they raise, usually about equity between different generations.

The policy importance of this issue is relatively straightforward in the prioritization of health care interventions. Many health care programs take significantly different lengths of time to produce their benefits. Applying a discount rate to those benefits leads to an unwarranted priority to programs producing benefits more rapidly. It results in a program that produces benefits in health and well-being say twenty years into the future being given lower priority than an alternative health care program that produces substantially less overall improvement in health and well-being, but produces that improvement much sooner. Many public health and preventive interventions, for example, vaccination programs and changes in unhealthy behavior, reap their health benefits years into the future. If those benefits are unjustifiably discounted, they will be given lower priority than alternative programs that produce fewer aggregate benefits. The result is a health policy that produces fewer overall health benefits over time than could have been produced with the same resources.

Fifth issue: What life expectancies should be used for calculating the benefits of life saving interventions?

In calculating QALYs it is standard practice to take account of differences in the average ages and in turn life expectancies of patients served by different health care programs; for example, a treatment for a life-threatening childhood disease would produce more QALYs than a comparable treatment for a life-threatening disease affecting primarily the elderly. Similarly, accurate estimates of the expected QALYs from different interventions would adjust for differences in the average life expectancies of patients caused by diseases other than those treated by the interventions; for example, an intervention that improved the quality of life of patients with cystic fibrosis, who have a much lower than average life expectancy as a result of their disease, would produce fewer QALYs than an intervention with a comparable improvement in lifetime quality of life for patients with average life expectancies undiminished by disease. In fact, this latter case raises difficult issues about discrimination against people with disabilities. But there are other differences in the life expectancies of different groups that an accurate estimate of QALYs produced by health interventions serving those groups would seemingly have to reflect; for example, there are significant differences in the life expectancies between different genders, racial and ethnic groups, and socio-economic groups within most countries. Internationally, these differences in life expectancies between different countries are often much larger. Should these differences affect calculations of the QALYs gained by health care and public health interventions that extend life or improve quality of life? An accurate estimation of the additional life years actually produced by those interventions should not ignore differences in life expectancies that the health care interventions will not affect, but the result will be that it is less valuable to save the life of a poor person in an underdeveloped country than a rich person in a developed country.

The differences in life expectancies between different racial, ethnic, and socio-economic status groups within a single country, as well as the very large differences between life expectancies in economically developed and poor countries, are often principally the result of unjust conditions and deprivations suffered by those with lower life expectancies. It would seem only to compound those injustices to give less value to interventions that save lives or improves quality of life for groups with lower life expectancies caused by the unjust conditions and deprivations from which they suffer. Differences in life expectancies between the genders, on the other hand, are believed to rest in significant part on biological differences, not on unjust social

conditions. Whether the biologically based component of gender differences in life expectancies should be reflected in measures like QALYs or DALYs is more controversial. For example, on the one hand, the lower life expectancy of men does not result from any independent injustice, but, on the other hand, it is explicit public policy and required by law in the United States to ignore this gender-based difference in most calculations of pension benefits and annuity costs so as to avoid gender discrimination. The developers of the DALY explicitly chose to use a single uniform measure of life expectancy (except for the biologically based gender difference), specifically that observed in Japan which has the highest national life expectancy, to measure gains from life saving interventions. They justified their choice in explicitly ethical terms as conforming to a principle of "treating like events as like," although the reasoning was not pursued in any detail (Murray, 1994, 7). How this issue is treated can have a substantial impact on the priorities that result from the cost effectiveness analysis, especially at the international level where country differences tend often to be greater than group differences within specific countries.

Each of the preceding five ethical issues can be considered issues in the construction of a cost-effectiveness analysis in health care. The other issues I want to briefly note can be considered issues in the use of cost effectiveness analysis in health resource prioritization. They are each issues of distributive justice or equity raised by the fact that a cost effectiveness analysis is insensitive to the distribution of health benefits and of the costs of producing them. Yet people's beliefs about equity and justice directly affect the relative priority they assign to different health interventions. One standard response to this point is that a CEA can only be an aid to policy making in general, and health resource prioritization in particular, and that policy makers must take account of considerations of equity in final policy decisions and choices. But as with the ethical issues in the construction of CEAs, much work remains to be done to clarify and assess alternative positions on these issues of equity so the policy choices on them can at least be better informed, even if they remain controversial. Here, there is only space to state three of the main equity issues in the use of CEAs and some of the principal ethical considerations supporting different positions on them (Daniels, 1993, pp. 224-233). After doing that, I shall mention an alternative quantitative methodology that, unlike CEA, incorporates considerations of equity within the quantitative analysis.

Sixth Issue: What priority should be given to the sickest or worst off?

It is a commonplace that most theories of distributive justice require some special concern for those who are worst off or most disadvantaged; for example, it is often said that the justice of a society can be measured by how it treats its least well off members. In the context of health care allocation and the prioritization of health interventions, the worst off with regard to need for the good being distributed might reasonably be thought to be the sickest patients. In many cases, the sickest will be given priority by a CEA comparing treating them as opposed to less sick patients; the sickest have greater possible improvements in HRQL because they begin from a lower HRQL, and so, for example, in comparing fully effective treatments those for the sickest will produce the greater benefits. But in other cases giving priority to the sickest will require a sacrifice in aggregate health benefits. An abstract example makes the point most concisely. Suppose Group A patients have a very serious disease that leaves them with a health utility level of .25 as measured by the HUI, and this would be raised only to .45 with the best available treatment because no treatment is very effective for their disease; for example, patients with severe chronic obstructive pulmonary disease or with severe chronic schizophrenia that is largely resistant to standard pharmacological treatments. A similar number of Group B patients have a health utility level of .60 because they have a considerably less serious disease, but since treatment for their disease is more effective, although no more costly, it would raise their health utility level to .90; for

example, patients with asthma, or with milder forms of pulmonary disease or schizophrenia that both leave them less disabled without treatment and are more responsive to treatment. Should we give priority to treating Group B because doing so would produce a 50% greater aggregate health benefit at the same cost, as the CEA standard implies, or to treating Group A who are the sickest? In some empirical studies, both ordinary people and health professionals prefer to sacrifice some aggregate health benefits in order to treat the sickest patients, although the degree of sacrifice they are prepared to make is variable and not statistically reliable (Nord, 1993, pp. 227-238).

One difficulty raised by this issue is determining what weight to give to this particular aspect of equity--concern for the worst off. Virtually no one would prefer to treat the sickest, no matter how costly their treatment and how small the benefit to them of doing so, and no matter how beneficial and inexpensive treatment for the less sick might be. However, there seems no objective, principled basis for determining how much priority to give the sickest, that is, how much aggregate health benefits should be sacrificed in order to treat or give priority to the sickest. Instead, the most one can say for most people, and indeed for many theories of distributive justice as well, is that concern both for maximizing overall benefits with scarce health resources and for helping the worst off or sickest are parts of their overall moral views, but there is a large range of indeterminacy regarding the proper trade off between these two concerns when they are in conflict.

One issue in understanding this concern for the worst off important for health care priorities is whether it should focus on who is worst off at a point in time or instead over an extended period of time, such as a lifetime. When choosing between patients to receive a scarce resource, such as in organ transplantation, it is often plausible to focus on lifetime well being, since otherwise we may give priority to the patient who is worst off at the time the distributive choice is made, but whose lifetime level of well being is far higher than the other patient. Frances Kamm has defended a notion of need in this context according to which the neediest patient is the patient whose life will have gone worst if he or she does not get the scarce resource, such as an organ transplant (Kamm, 1993, ch. 8). However, some justifications for giving priority to the worst off may support focusing on the sickest here and now.

What are the ethical justifications for giving priority to the worst off? I can mention only two possibilities here. One is that we must give priority to the worst off in order to avoid increasing the already unjustified disadvantage or inequality they suffer relative to those better off. But it is worth noting that a concern for the worst off is not always the same as a concern to produce equality in outcomes. In the example above of Groups A and B, equality could be achieved by what Derek Parfit has called "leveling down", that is by bringing B's health utility level down to that of A's instead raising A's level up to that of B (Parfit,). If equity here is equivalent to equality in outcomes, then if it were not possible to raise A's level above .40 with treatment, equity would seem to support not treating Group B and letting their condition deteriorate until it reached the lower level of Group A. The fact that no one would defend doing this suggests that this aspect of our notion of equity or justice is best captured by the idea of giving priority to improving the condition of the worst off, rather than by a simple concern for equality in outcomes. A different justification for giving priority to treating the sickest, offered by some participants in Nord's research, is that it would be subjectively more important to the sickest to obtain treatment, even if the health benefits they receive from treatment are less than those that would go to the less sick; this justification might support focusing on who is worst off at the point in time at which the decision about who to treat is made, not whose lifetime well-being will be lowest (Nord, 1993, 227-238).

One further issue concerning the priority to the worst off should be mentioned. In the context of health resource prioritization in health policy it seems natural to understand the worst off as the sickest. But this may not always be correct. At the most fundamental ethical level in our general theories of equity and distributive justice, our concern should be for those who are overall or all things considered worst off, and they will not always be the sickest. It could be argued that giving priority to the worst off in health resource prioritization sometimes requires giving priority to those with the lowest levels of overall well-being, even at some cost to aggregate health benefits produced and at the cost of not treating sicker persons whose overall well-being is much higher. A preference for health interventions that raise the level of well-being of those who are worst off in overall well-being, instead of giving priority to the sickest, might be justified in order not to increase the unjustified disadvantage suffered by those with the lowest overall level of well-being. If, instead, the priority to the worst off in health resource prioritization should focus only on health states and so on the sickest, a justification of this narrowed focus is needed.

Seventh Issue: When should small benefits to a large number of persons receive priority over large benefits to a small number of persons?

Cost effectiveness and utilitarian standards require minimizing the aggregate burden of disease and to maximizing the aggregate health of a population without regard to the resulting distribution of disease and health, or who gets what benefits. The issue about priority to the worst off focuses on who gets the benefits. A different issue concerns what benefits different individuals get. Some would argue that health benefits are often qualitatively different and so cannot all be compared on a single scale like the HUI, or in turn by a single measure like QALYs, but that is not the issue of concern now. In its most general form the issue about aggregation concerns what ethical limits there are, if any, on aggregating together different size benefits for different persons in comparing and prioritizing different health interventions; CEA accepts no such limits. There are many forms in which this issue can arise which cannot be pursued here (Kamm, 1993, Part II), but the version that has received the most attention, and which Daniels has called the aggregation problem, is when, if ever, large benefits to a few individuals should take priority over greater aggregate benefits to a different and much larger group of individuals, each one of whom receives only a small benefit. This issue arises when a very serious disease or condition for those affected that is also very costly to prevent or treat is compared with a much more prevalent disease or condition that both has a very small impact on each individual affected and is very inexpensive to treat or prevent in any one individual. Applying cost effectiveness or utilitarian standards, preventing or treating the very prevalent but low impact disease or condition at a given cost will receive higher priority when doing so produces greater aggregate benefits than using the same funds to treat or prevent the disease or condition that has a very great impact on each individual affected. The example that received considerable attention in the United States arose in the Oregon Medicaid priority setting process where capping teeth for exposed pulp was ranked just above an appendectomy for acute appendicitis, a potentially life-threatening condition. Because an appendectomy is approximately 150 times as expensive as capping a tooth for exposed pulp, the aggregate benefit of capping a tooth for 150 patients was judged to be greater than the benefit of an appendectomy for one patient. Since Medicaid coverage decisions were to be made according to the list of treatment/condition pairs ranked in terms of their relative cost effectiveness, it could have turned out, depending on the overall level of resources available to the Medicaid program, that tooth capping would have been covered but appendectomies not covered.

This result, and other less extreme examples like it, was highly counter-intuitive and unacceptable to most people, whose intuitive rankings of the relative importance or priority of health interventions are based on a one-to-one comparison of one tooth capped as opposed to one

appendectomy performed. In the face of these results Oregon made a fundamental change in its prioritization methodology, abandoning the cost effectiveness standard in favor of what was essentially a relative benefit standard applied to different treatment/condition pairs that did not take account of differences in their costs. This was not a minor problem requiring tinkering at the margins of the CEA standard, but a fundamental challenge to it and so required a fundamental revision in it.

Yet it is by no means clear that no such aggregation can be ethically justified. The very case that precipitated Oregon's Medicaid revision was a 12 year old boy in need of a bone marrow transplant as the only effective chance to save his life. Oregon denied coverage under its Medicaid program on the grounds that it could do greater good by using its limited resources to improve prenatal care for pregnant women, in this case giving higher priority to small benefits to many over a potentially much larger benefit to a few. Moreover, many public policy choices appear to give higher priority to small benefits to many over even life saving benefits to few; for example, governments in the United States support public parks used by tens or hundreds of thousands of persons, while reducing funding for public hospitals resulting in quite predictable loss of life.

The cost effectiveness or utilitarian standard that permits unlimited aggregation of benefits might be defended by distinguishing between the clinical context in which physicians treat individual patients and the public health and health policy context in which health resource allocation decisions are made that will affect different groups in the population. In the clinical context, physicians forced to prioritize between individual patients typically will first treat the patient who will suffer the more serious consequences without treatment, or who will benefit the most from treatment, even if doing so will prevent her treating a larger number of less seriously ill patients. But from a public health or health policy perspective, it could be argued that the potential overall or aggregate effects of alternative interventions on population health is the appropriate perspective. However, the Oregon example makes clear that even when allocating public resources for interventions to improve the health of a population it is ethically controversial whether always giving priority to producing the maximum aggregate benefits, even when that is done by giving small benefits to many at the cost of forgoing large benefits to a few, is justified.

Just as with the problem of what priority to give the worst off, part of the complexity of the aggregation problem is that for most people some, but not all cases of aggregation are ethically acceptable and equitable. The theoretical problem then is to develop a principled account of when, and for what reasons, different forms of aggregation satisfy requirements of equity and when they do not (Kamm, 1993). There is no consensus on this issue either among ordinary persons or within the literature of health policy or ethics and political philosophy. As with the problem about priority to the worst off, the complexities of this issue have received relatively little attention in bioethics and moral and political philosophy, and there is much difficult but important work to be done.

Eighth Issue: The conflict between fair chances and best outcomes.

The third ethical issue in the use of CEA for health resource utilization that I will mention here has been characterized as the conflict between fair chances and best outcomes (Daniels, 1993, pp. 224-233). The conflict is most pressing when the health intervention is life saving and not all those whose lives are threatened can be saved, but it arises as well when threats are only to individuals' health and well-being. In the context of health care, this issue first received attention in organ transplantation where there is a scarcity of life saving organs such as hearts and lungs resulting in thousands of deaths each year of patients on waiting lists for an organ for transplant;

an abstract example from transplantation can illustrate the issue most clearly and succinctly (Brock, 1988).

Suppose two patients are each in need of a heart transplant to prevent imminent death, but there is only one heart available for transplant. Patient A has a life expectancy with a transplant of ten years and patient B has a life expectancy with a transplant of nine years (of course, precise estimates of this sort are not possible, but the point is that there is a small difference in the expected benefits to be gained depending on which patient gets the scarce organ), with no difference in their expected quality of life. Maximizing health benefits or QALYs, as a CEA standard requires, favors giving the organ to patient A, but patient B might argue that it is unfair to give her no chance to receive the scarce heart. Just as much as A, she needs the heart transplant for life itself and will lose everything, that is her life, if she does not receive it. It is unfair, B might argue, to give the organ to A because the quite small increment in expected benefits from doing so is too small to justly determine who lives and who dies. Instead, she argues, each of them should receive a fair chance of getting the organ and having their health needs met; in this case, that might be done by giving each an equal chance of receiving the transplant through some form of random selection between them, or by a weighted lottery that gives the patient who would benefit more some greater likelihood of being selected to receive the organ, but still gives the patient who would benefit less some significant chance of getting it instead (Broome, 1984, pp. 38-55; Kamm, 1993, Part III; Brock, 1988).

Most prioritization and rationing choices arise not from physical scarcity of the needed health resource, as in organ transplantation, but from economic scarcity, limits in the money society devotes to health care. Will this issue of equity arise in health resource prioritization and allocation choices forced by economic scarcity? Two considerations will often mitigate the force of the ethical conflict between fair chances and best outcomes there. First, allocation of resources in health care is typically not an all or nothing choice, as in the case of selecting recipients for scarce organs, but is usually a matter of the relative priority for funding to be given to different health programs or interventions. That one health program A promises a small gain in aggregate health benefits over a competing program B need not entail that A is fully funded and B receives no funding, but only that A should receive higher priority for, or a higher level of, funding than B. Persons with the disease or condition that A treats will have a somewhat higher probability of being successfully treated than will those who have the disease or condition that B treats; in the case of prevention, those at risk of A will have a somewhat higher probability of successful prevention than will those at risk of B. When there is significant resource scarcity this will involve some sacrifice in aggregate health benefits that might have been produced by always preferring the more cost effective alternative. But doing so means that Individuals who are served by B have no complaint that the small difference in expected benefits between programs A and B unfairly them from having their health needs met at all. Instead, the small difference in expected benefits between programs A and B need only result in a comparably small difference in the resources devoted to A and B; it is not obvious that this is unfair to those patients served by B, whose needs are somewhat less well served than patients in program A because of B's lower priority and level of funding.

The second consideration that may mitigate some the conflict between fair chances and best outcomes in health resource prioritization forced by economic scarcity is that often, probably usually, the diseases and health problems to be treated or prevented are not life threatening, but instead only impact on individuals' quality of life, and often for only a limited period of time. In these cases, the difference in health benefits between individuals who receive a needed health intervention that is given a higher priority and individuals who do not receive a needed health

intervention because their condition is given lower priority, is much less, making the unfairness arguably less compelling.

These two considerations may mitigate, but they do not fully avoid, the conflict between fair chances and best outcomes in prioritization decisions about health interventions forced by economic scarcity. When a more cost effective health program is developed for one population instead of a different less cost effective health program for a different population, individuals who would have been served by the second program will have a complaint that they did not have a fair chance to have their needs served only because of a small gain in the benefits that are produced by the first program. The fair chances versus best outcome conflict will arise in prioritizing health interventions in health policy; how this conflict can be equitably resolved is complex, controversial, and unclear.

Ninth Issue: Does use of CEA to set health care priorities unjustly discriminate against the disabled?

In several contexts using CEA to set health care priorities will result in assigning lower priority to both life extending and quality of life improving treatment for disabled than nondisabled persons with the same health care needs (Brock, 1995, pp. 159-184). Here are five such contexts. First, since already disabled persons have a lower HRQL from their disability than nondisabled persons, treatment that extends their life for a given number of years produces fewer QALYs than the same treatment that extends the life of a nondisabled person for the same number of years. Second, if two groups of patients with the same HRQL have the same need for a life sustaining or quality of life improving treatment, but one will be restored to normal function and the other will be left with a resultant disability, more QALYs will be produced by treating the first group. Third, persons with disabilities often have a lower life expectancy because of their disability than otherwise similar nondisabled persons. As a result, treatments that prevent loss of life or produce lifetime improvements in quality of life will produce fewer QALYs when given to disabled than to nondisabled persons with the same health care needs. Fourth, disabilities often act as comorbid conditions making a treatment less beneficial in QALYs produced for disabled than for nondisabled persons with the same health care needs. Fifth, the presence of a disability can make treatment of disabled persons more difficult and so more costly than for nondisabled persons with the same health care needs; the result is a lower cost effectiveness ratio for treating the disabled persons.

In each of the five cases above, disabled persons have the same medical and health care need as nondisabled persons, and so the same claim to treatment on the basis of their needs. But treating the disabled will produce less benefit, that is fewer QALYs, because of their disability than treating the nondisabled. Thus, their disability is the reason for their receiving lower priority for treatment. This at least arguably fails to give equal moral concern to disabled persons' health care needs and is unjust discrimination against them on grounds of their disability. Indeed, United States Health and Human Services Secretary Louis Sullivan denied Oregon's initial request for a waiver of federal regulations for its proposed revisions to its Medicaid plan on the grounds that Oregon's method of prioritization of services was in violation of the Americans with Disabilities Act (ADA).³ Sullivan cited some of the five kinds of cases I noted above in support of that position, and Oregon in turn made essentially ad hoc revisions in its ranking to avoid the putative violation of the ADA.

Disabled persons charge that in cases like the first I cited above concerning life saving treatment, the implication of use of CEA to prioritize health care is that saving their lives, and so their lives themselves, have less value than nondisabled persons' lives. They quite plausibly find that implication of CEA threatening and unjust. There are means of avoiding these problems

about discrimination against persons with disabilities, but they involve abandoning fundamental features of CEAs. For example, one response to the first case cited above would be to give equal value to a year of life extension, whatever the quality of that life, so long as it is acceptable to the person whose life it is (Kamm, 1993, Part I). But that has problematic implications too since, for example, a small percentage of persons in surveys say they would want their lives sustained even if they were in a persistent vegetative state. I cannot pursue the issues further here, but I believe the problem of whether CEA unjustly discriminates against the disabled is a deep and unresolved difficulty for use of CEA and QALYs to prioritize health care.

The sixth, seventh, and eighth issues above all raise possible criticisms of the maximization standard embodied in CEA; in each case, the claim is that equity requires attention to the distribution of health benefits and costs to distinct individuals. Steadfast utilitarians or consequentialists will reject the criticisms and hold fast to the maximization standard. But most people will accept some departure from the maximization standard of CEA; there are two broad strategies for how to do so. The first and probably most common is to propose CEA as an aid to policy makers who must make prioritization and allocation choices in health care, but then to remind those policy makers that they must take account of these considerations of equity as well in their decision making; this may be, but usually is not, accompanied by some guidance about alternative substantive positions, and reasons in support of them, on the equity issues. Moreover, some use of CEA in health policy and health program evaluation does not raise these last three issues of equity; for example, CEA of alternative treatments that each have uniform but different benefits for a group of patients with a particular medical condition. And outside of a CEA, either QALYs or DALYs can be used for evaluating alternative interventions, or for monitoring changes over time in health status or the burdens of disease, in a given group or population.

The second strategy for responding to concerns about equity seeks to develop a quantitative tool that measures the specific weight people give to different equity concerns in comparing interventions that raise issues of distributive justice because they serve different individuals or benefit individuals differently. The most prominent and promising example is Erik Nord's "person trade-off" approach which explicitly asks people how many outcomes of one kind they consider equivalent in social value to X outcomes of another kind, where the outcomes are for different groups of individuals (Nord, Working Paper 38). For example, people can be asked, as in our earlier example, to compare treatment A for very severely ill patients who are at .25 on the HUI without treatment and who can be raised only to .45 with treatment, with treatment B of less severely ill patients who are at .60 and can be raised to .90 with treatment; filled out detailed examples, of course, will make the comparisons more understandable. Respondents are then asked how many patients treated with A would be equivalent in social value to treating 100 patients with B. Answers to questions of this form will tell us in quantitative terms how much importance people give to treating the sickest when doing so conflicts with maximizing aggregate health benefits.

The person trade-off approach is designed to permit people to incorporate concerns for equity or distributive justice into their judgments about the social value of alternative health programs. There has been relatively little exploration and use of this methodology in health care evaluation in comparison with the mass of methodological work on and studies of aggregate QALYs and CEAs, in part because many health policy analysts and health economists assume, often with little or no argument, that the social value of health programs is the sum of the individual utilities produced by the program. As I noted in the introduction to the paper, the early stage we are now at in the development and use of the person trade-off approach is a reason for

caution at the present time about using it to settle issues of equity in health resource prioritization. While the utilitarian assumption in CEA is rejected in most philosophical work on distributive justice, as well as in the preferences ordinary people express for different health outcomes and programs, I also noted in the introduction a second more important reason for caution about bringing considerations of equity into health policy decision making through a quantitative methodology like the person trade-off methodology--the issues of distributive justice that must be addressed by equitable health resource prioritization represent deep and long-standing divisions in moral and political philosophy about which there is not now, and may never be, anything approaching consensus. There is a strong case to be made, though I cannot pursue it here, that important value conflicts about justice of this sort should be addressed in public, democratic political processes, or in fair, participatory and accountable procedures within private institutions like managed care organizations (Daniels and Sabin, 1997, pp. 303-350). The person trade-off method can be a useful aid to those deliberative decision making processes in providing more structure and precision to different people's views about equity in health care resource prioritization and trade-offs, but it is not a substitute for that deliberation. Despite these briefly noted reservations, I do emphasize that for purposes of resource prioritization and allocation, the person trade-off approach is the proper perspective, in comparison with CEA, because it correctly reflects that the choices are typically about how health benefits and costs are distributed to different individuals.

Conclusion

I have distinguished above nine distinct issues about equity and justice that arise in the construction and use of cost effectiveness analysis to minimize the burdens of disease and to maximize health outcomes. In each case the concern for equity is in my view valid and warrants some constraints on a goal of unqualified maximization of health outcomes. There has not been space here to pursue at all fully any of these eight issues regarding equity and justice--each is complex, controversial, and important. In each case, my point has been that there are important ethical and value choices to be made in constructing and using the measures; the choices are not merely technical, empirical, or economic, but moral and value choices as well. Each requires explicit attention by health policy makers using CEA. In a few cases I have indicated my own view about how the potential conflict between equity and utilitarian maximization might be resolved, but in other cases I have simply summarized briefly some arguments for giving the particular concern about equity some weight when it conflicts with maximization of utility. For some of these issues, the literature and research is at a relatively early stage and one cannot be confident about how the issues should be resolved or even about the range of plausible positions and supporting reasons on them. However, this is not grounds for ignoring the issues, but instead for getting to work on them and for ensuring that they receive explicit attention and deliberation in decisions about health resource prioritization and allocation.

BIBLIOGRAPHY

- Bergner, M., Bobbitt, R.A., Carter, W.B., and Gibson, B.S.: 1981 'The sickness impact profile: Development and final revision of a health status measure,' Medical Care 19, 787-805.
- Brock, D.W.: 1988 'Ethical issues in recipient selection for organ transplantation,' in Organ Substitution Technology: Ethical, Legal, and Public Policy Issues (ed.) D. Mathieu. Boulder and London: Westview Press.
- Brock, D.W.: 1989, 'Justice, Health Care, and the Elderly,' Philosophy & Public Affairs 18, 3 297-312
- Brock, D.W.: 1992 'Quality of Life Measures in Health Care and Medical Ethics', in A. Sen and M. Nussbaum (eds.), (The Quality of Life., Oxford: Oxford University Press.
- Brock, D.W.: 1995, 'Justice and ADA: does prioritizing and rationing health care discriminate against the disabled?' Social Theory and Policy 12 159-84.
- Broome, J.: 1984, 'Selecting People Randomly,' Ethics 95, 38-55.
- Daniels, N.: 1988, Am I My Parents Keeper? An Essay on Justice Between the Young and the Old, New York: Oxford University Press, .
- Daniels, N.: 1993 'Rationing Fairly: Programmatic Considerations,' Bioethics 7/2-3, 224-233.
- Daniels, N. and Sabin, J.: 1997, 'Limits to health care: fair procedures, democratic deliberation, and the legitimacy problem for insurers,' Philosophy and Public Affairs 26,4 303-50
- Gold, M.R. et. al.: 1996, Cost-Effectiveness in Health and Medicine New York: Oxford University Press, Ch. 7.
- International Classification of Impairments, Disabilities and Handicaps (Geneva: World Health Organization, 1980).
- Kamm, F.M.: 1993, Morality/Mortality. Volume One. Death and Whom to Save From It Oxford: Oxford University Press
- Kaplan, R.M. and Anderson, J.P.: 1988, 'A general health policy model: update and applications,' Health Services Research, June 23, 203-35.
- Keeler, E.B. and Cretin, S.: 1983, 'Discounting of life-saving and other nonmonetary effects,' Management Science 29, 300-306.
- Murray, C.J.L.: 1994 'Quantifying the burden of disease: the technical basis for disability-adjusted life years,' . in Global Comparative Assessments in the Health Sector: Disease Burden, Expenditures and Intervention Packages, eds. C.J.L. Murray and A.D. Lopez (Geneva: World Health Organization,)
- Murray, C.J.M.:1996, 'Rethinking DALYs, in The Global Burden of Disease: A Comprehensive Assessment of Mortality and Disability From Disease, Injuries, and Risk Factors in 1990 and Projected to 2020. Geneva: World Health Organization.
- Nord, E.: 'The Person Trade-off Approach to Valuing Health Care Programs,' Working Paper 38, National Centre for Health Program Evaluation (Fairfield Hospital, Fairfield Victoria, Australia) 7.
- Nord, E.: 1993 'The trade-off between severity of illness and treatment effect in cost-value analysis of health care,' Health Policy 24, 227-38.
- Parfit, D.: "Equality or priority," The Lindley Lecture. Copyright: Department of Philosophy, University of Kansas.
- Torrance, G.W. et.al., 1992 'Multi attribute preference functions for a comprehensive health status classification system.' Working Paper No. 92-18. Hamilton, Ontario: McMaster University, Center for Health Economics and Policy Analysis,

- Walzer, M. 1983, Spheres of Justice New York: Basic Books, .
- Ware, J.E. and Sherbourne, D.C.: 1992, "The MOS 36-item short form health survey," Medical Care 30, 473-83.
- Williams, A. 1997, 'Intergenerational Equity: An Exploration of the "Fair Innings" Argument,' Health Economics 6,2 117-32.
- World Bank, 1993 World Development Report 1993: Investing in Health, Oxford: Oxford University Press.

NOTES

¹ This paper draws heavily on my "Considerations of Equity in Relation to Prioritization and Allocation of Health Care Resources," in Ethics, Equity and Health For All, eds Z. Bankowski, J.H. Bryant and J. Gallagher (Geneva: CIOMS, 1997) and "Ethical Issues in the Development of Summary Measures of Population Health States" in Summarizing Population Health: Directions for the Development and Application of Population Metrics (Washington DC: National Academy Press, 1998).

2

Interventions that would improve health should be understood broadly, and in particular extend substantially beyond health care. It is widely agreed that other factors such as improved sanitation and economic conditions have contributed more to the health gains of the past century than has health care. However, in this paper I shall largely confine myself to health care interventions

³Unpublished letter from Secretary of Health and Human Services Louis Sullivan to Oregon Governor Barbara Roberts, August 3, 1992.