



**PROGRAM ON THE GLOBAL
DEMOGRAPHY OF AGING**

Working Paper Series

**Axiomatic Foundations of Cost
Effectiveness Analysis**

David Canning

December 2009

PGDA Working Paper No. 51

<http://www.hsph.harvard.edu/pgda/working.htm>

The views expressed in this paper are those of the author(s) and not necessarily those of the Harvard Initiative for Global Health. The Program on the Global Demography of Aging receives funding from the National Institute on Aging, Grant No. 1 P30 AG024409-06.

Axiomatic Foundations of Cost Effectiveness Analysis

David Canning

**Harvard School of Public Health
December 2009**

Abstract

Cost-effectiveness analysis, which ranks projects by quality adjusted life years gained per dollar spent, is widely used in the evaluation of health interventions. We show that cost effectiveness analysis can be derived from two axioms: society prefers Pareto improvements and society values discounted life years, lived in perfect health, equally for each person. These axioms generate a unique social preference ordering, allowing us to find the cost effectiveness threshold to which health projects should be funded, and to extend cost effectiveness analysis to give a consistent method of project evaluation across all sectors of the economy.

I would like to thank Abhijit Banerjee, David Bloom, Dean Jamison and Michael Moore as well as participants in the joint Harvard/MIT/Boston University Health Economics Seminar and the 2nd Biannual Conference of the American Society of Health Economics for useful discussion on these ideas.

dcanning@hsph.harvard.edu

Key words: cost-benefit, cost-utility, social choice, welfare economics

This paper is an updated version of "Valuing lives equally and welfare economics" PGDA working paper No. 27, 2007.

"the cost of a thing is the amount of what I will
call life which is required to be exchanged for it"

Walden, Henry David Thoreau

1. Introduction

In the health sector there is unease about the idea that we should value lives in terms of people's willingness, and ability, to pay. The ethical difficulties involved in cost-benefit analysis have led to the use of cost-effectiveness analysis (Weinstein and Stasson (1977)) in which projects are evaluated in terms of the life years gained per dollar spent. The use of this criterion assumes society values lives, or more precisely discounted, quality adjusted, life years, equally across people and wants to maximize the total number of life years produced, independently of who gets them.

Detailed standards for cost effectiveness analysis have been produced by the World Health Organization (Edejer et al. (2003)) and are widely used as a criterion for priority setting among health interventions (e.g. Jamison et al. (2006)). Medical journals frequently publish articles that estimate the cost effectiveness of particular treatments and procedures, and it could be argued it is the most common method of economic project appraisal¹. In Britain, the National Institute for Health and Clinical Excellence explicitly uses cost effectiveness as a criterion in selecting treatments for use in the National Health Service, and recommends procedures that meet a cost effectiveness threshold (Devlin and Parkin (2004)). In the United States, cost effectiveness is not used as a criterion for Medicare provision but it has been argued that its use could improve the quality of services and reduce costs (Neumann et al. (2005)).

¹ A search of the ISI Web of Knowledge database of academic articles produced 16,678 articles with 'cost effectiveness' or 'cost effective' in the title and 3,632 articles with 'cost benefit' in the title (searching independently of word ordering).

In cost effectiveness analysis we value the discounted quality adjusted life years of each person equally. This contrasts with cost benefit analysis where projects are evaluated by summing each person's willingness to pay in money units. Cost effectiveness analysis is inconsistent with cost benefit analysis except under extreme assumptions on the nature of individual preferences (Bleichrodt and Quiggin (1999), Johannesson (1995), Dolan and Edlin (2002)) or assuming that the distributions of income and health are already socially optimal (Hansen et al. (2004)).

The fact that cost effectiveness and cost benefit analysis are usually inconsistent makes it difficult to allocate resources coherently between the health sector and other sectors of the economy. In particular, cost effectiveness analysis can allocate a fixed health budget between competing health interventions but has no way of deciding what the size of the health budget should be. Setting the health budget is equivalent to society setting a money value on a quality adjusted life year, but the willingness to pay for a life year differs across people, and there is at present no clear way to amalgamate these valuations.

One response to these problems would be to abandon cost effectiveness analysis and use cost benefit analysis for health projects just as we do for other projects in other sectors of the economy (Fuchs and Zeckhauser (1987)). This leads to higher levels of health care for those with a greater willingness to pay but dismisses the strong ethical appeal of the argument that life is different from other commodities. Life, particularly healthy life, can be argued to have a special moral importance on the grounds that it is a prerequisite for the opportunity to carry out other activities (Daniels (2008)). Life and health also have a privileged role in the capabilities approach to evaluating wellbeing on the grounds that they are essential to a having a reasonable

opportunity set and freedom of choice (Sen (1985), Sen (1999)). We may therefore recognize a moral claim to healthy life without recognizing claims to other goods in the same way.

Cost effectiveness analysis usually begins by assuming that we wish to maximize discounted quality adjusted life years gained with a fixed health budget. The weakness of this approach is that it leaves open the question of how health care should be financed and how to judge interventions in other sectors, such as transport, which have both economic and health outcomes. Our approach is to find ethical axioms that apply to all social decisions. We choose these axioms so that cost effectiveness as currently employed is the correct decision rule in the health sector and investigate the implication of these axioms for other social choices.

Our first axiom is the Pareto principle; a project that makes no one worse off and at least one person better off should be socially preferred. This axiom guarantees efficiency. We also require a second axiom on society's preferences on distribution. We assume that if everyone maintained their current endowment of other goods, but had perfect health, society would value additional discounted life years equally across individuals. If every person has perfect health, and their current endowment of other goods, the only thing that varies between possible states is the life span of each person. In comparing across these states with perfect health status but different life spans for different individuals we assume society always prefers a state with a higher sum of discounted life years to a state with a lower sum.

We show that, taken together, the Pareto principle, and the assumption that society values lives in perfect health equally, generate a unique social preference ordering over states and this social preference ordering implies the use of cost effectiveness analysis in the health sector. A fixed sum allocated to health should be spent to maximize the sum of discounted, quality adjusted life years gained, independently of who gets them, which implies implementing projects

in order of their cost effectiveness. The quality adjustment factor for each individual is the rate they are willing to trade life years in their current health state for life years lived in perfect health.

Our approach extends the work of Garber and Phelps (1997) and Meltzer (1997) who analyze how cost effectiveness analysis can be derived from the optimal decision making of an individual. The Pareto efficiency axiom implies that the health interventions used by an individual maximize their utility given their budget constraint. Our second axiom extends their approach by allowing for comparisons across people.

While our axiomatic approach can be used to justify cost effectiveness analysis it can also be used to extend its reach. We show that the unique social preference ordering generated by our two axioms can be represented by a utilitarian social welfare function. We simply ask each person how many life years, lived at full health and their current endowment of other goods, would be required to make them indifferent to any proposed social state. This gives a measure of the utility of the state for each person in terms of equivalent quality adjusted life years. Social welfare at the proposed social state is the just the sum of these discounted quality adjusted life years across people.

Any public choice problem can then be addressed by finding the policy that maximizes this social welfare function. Usually cost effectiveness analysis provides a ranking of health projects but leaves open where the threshold for funding should be placed. We use our social welfare function to derive the optimal threshold (in terms of dollars per quality adjusted life year gained) for projects that should be funded, which is equivalent to deciding how much should be spent on the health sector. We show that the optimal threshold depends on the method of funding – we find results for lump sum transfers and a proportional income tax. The funding method

matters because our social welfare function values everything in quality adjusted life years. The "cost" of the money used to finance the health system is the quality adjusted life years those who pay for the system would be willing to give up to retain the money. Since this varies across people, and the method of financing affects who pays, the financing method affects the "cost" of the funds and where the threshold should be set.

An implication of our axioms is that society cares about poverty as well as the average level of income. Redistributing money from the rich to the poor will improve social welfare if the willingness to pay for a life year is higher for rich people than poor people (which is another way of saying that the value of money, measured in life years, is higher for the poor than the rich).

The social preference ordering our framework generates can also be used to evaluate projects in other sectors of the economy. Using our cost effectiveness axioms the net benefit of a project is the sum of the willingness of people to pay for the project in discounted quality adjusted life years. This "life-metric" approach to measuring willingness to pay tends to give more weight to the preferences to the poor than standard "money metric" cost benefit analysis; a rich man may often be willing to pay more money than a poor man for a project, while not being willing to give up more life years. In terms of cost benefit analysis all we do is move from the usual measurement of willingness to pay in money units, to willingness to pay measured in quality adjusted life years. However, in welfare analysis the choice of numeraire has real effects (Berlge and Renard (1985), Brekke (1997)) because the numeraire becomes an interpersonally comparable unit of value.

There are three reasons why the results in this paper may appear surprising at first sight. The first is Arrow's impossibility theorem which appears to rule out a consistent social

preference ordering of the type we construct. Arrow (1950) puts forward four desirable properties for a social choice mechanism, and shows that it is impossible for a mechanism to possess all four properties. Our approach avoids Arrow's impossibility theorem by adhering to only two of his four desirable properties. Our social ordering obeys the Pareto principle and non-dictatorship. However it is incomplete; we cannot rank all preferences. In particular we limit our analysis to people who want longer, rather than shorter, life spans. We also violate the independence of irrelevant alternatives axiom. When comparing two arbitrary health states we require the existence of an "irrelevant" third state, perfect health, to measure them against.

The second is that it has been shown that always valuing lives equally is inconsistent with the Pareto principle (Weinstein and Manning (1997), Hasmanand and Østerdal (2004)), and thus our two axioms appear inconsistent. However, we do not value all lives equally. We assume that if people had perfect health and their current endowment of other goods society would value their discounted life years equally. We assume society puts equal value only on life years lived in this state and may put a higher, or lower, value on life years lived with a different health status, or a different endowment of other goods.

The third issue is that our second axiom appears to be non-welfarist in the sense that it does not depend on individual preferences; society prefers people to have longer life spans independently of what they themselves want. Kaplow and Shavell (2001) show that any non-welfarist social preferences necessarily violate the Pareto principle for some individual preferences. However, we limit admissible individual preferences to those where additional lifespan is desirable for every person. This means that we do not have a conflict between individual preferences and social preferences.

An objection to the results in this paper is that the second axiom, that society values discounted life years lived in perfect health, with the current endowment of other goods, equally for each person is somewhat arbitrary. We could value life years lived at some other reference point, with imperfect health or a different endowment of other goods, equally, giving different results. We could take some other commodity, such as money, as our standard and argue that society values additional units of money equally for each person.

We consider the second axiom to be a statement of the cost effectiveness principle on a limited domain. At the current endowment of goods, but perfect health for every person, we would value additional discounted life years equally for each person. We think that someone who advocates cost effectiveness analysis as currently practiced would subscribe to this axiom. If they also accept the Pareto principle that society prefers Pareto improvements, we show that they then have a particular social preference ordering over all states of the economy. This includes how to quality adjust life years when people are not in perfect health, how to set the cost threshold that cost effective interventions need to achieve to be financed, how to rank projects in other sectors of the economy, and the welfare effects of redistributing income. Our approach makes clear the ethical axioms that are required to underpin cost effectiveness analysis and the implications of these axioms for other public choices.

We leave open the issue of whether the axioms we use are defensible on ethical grounds. People who think that health care should be allocated to maximize aggregate health, independently of willingness to pay, who care about poverty reduction, and think Pareto improvements are desirable may be attracted to the axioms we propose and using a life year metric for social welfare. Our approach shows how to combine these views into a consistent method of social choice.

However our second axiom incorporates views on distribution and there may be a dispute about what constitutes a desirable distribution. People who think that money is equally valuable to each person and who consequently see no benefit in income redistribution, and who think that health care should be distributed to those most willing to pay for it, will be attracted by cost-benefit analysis and using a money metric in social decisions.

While the use of cost benefit analysis rather than the generalized cost effectiveness analysis proposed here is at heart an issue of social preferences on distribution, generalized cost effectiveness, based on our axiomatic approach, also has some technical advantages. We show it generates a consistent method of social choice over the whole economy, so that if state A is strictly socially preferred to state B then B is not preferred to A. In addition the rule is transitive so if state A is socially preferred to state B, and B is socially preferred to C, then A is socially preferred to C. Project evaluation using cost benefit analysis does not provide a consistent ordering, exhibiting preference reversals, so that moving from A to B may be desirable on a cost-benefit criterion while once we are at B moving back to A may also be desirable. Ranking states using a cost-benefit decision rule also lacks transitivity (Blackorby and Donaldson (1985)) .

In section 2 we consider a stylized model of the economy and individual preferences. In section 3 we assume that society prefers Pareto improvements and that society values additional discounted life years lived at in perfect health, and current endowments of other goods, equally for each person. These assumptions are shown to imply the existence of a unique, well behaved, social preference ordering for our economy. In section 4 we show how cost effectiveness analysis follows from our social preference ordering and in section 5 we examine how the threshold for cost effective projects should depend on the method used to finance the health

system. In section 6 we generalize cost effectiveness analysis to rank projects in all sectors of the economy. Section 7 concludes.

2. Individual Preferences

We assume that there exist 3 types of commodities: traded goods, non-traded goods, and life span. We also assume that there is no uncertainty. An allocation for a consumer is a bundle (x, z, l) where $x \in R^G$ is the vector of consumption of the G traded goods, $z \in R^H$ is the vector of consumption of H non-traded goods and $l \in [0, \infty)$ is lifespan. We can think of health status as one of the non traded goods.

We assume that the consumer has a complete, reflexive, transitive, and continuous preference ordering over the space of consumption bundles. There exists a continuous utility function U that represents these preferences, though any positive monotonic transformation of U also represents the same preferences.

Suppose the consumer is endowed with the bundle of goods (e, z, l) . The agent has a budget constraint for traded goods given by

$$xp = \sum_j x_j p_j \leq \sum_j e_j p_j = m \quad (1)$$

There are prices p_j for each traded good at which trades may be made. We assume that all prices are positive. We denote the “money” value of endowment at the price vector p by m . Let F denote the set of feasible consumption bundles. Let F_x denote the feasible set of traded goods. We make the following assumptions:

A1. F is a bounded, closed, and convex set in $R^G \times R^H \times R^+$.

A2. The set F_x is bounded below in the sense that there exists some $\underline{x} \in R^G$ with the property that

$x \in F_x \Rightarrow x \geq \underline{x}$. l is bounded below by zero.

A3. The utility function U is continuous.

A4. The utility function U is strictly concave.

A5. The agent's utility function U is strictly increasing in at least one component of x . The function U is strictly increasing in l .

Assumptions A3-A5 are much stronger than is required for our results. For example, the continuity of the utility function assumed in A3 could be derived from a continuous preference ordering. Assumption 4 assures that, given the budget constraint, the consumption bundle of traded goods chosen is unique. Assumption 5 means that there is always a valuable tradable good and ensures that the agent's budget constraint is binding. Assumption 5 also implies that holding all else equal the agent strictly prefers a longer life span. The assumption that agents strictly prefer more life to less is implicitly an assumption that the vector of goods being consumed is above some minimal level that makes life worth living and puts a bound on how low \underline{x} can be. We will have one non-tradable that we regard as health which we shall think of as a good. The other non-tradables may be goods or bads.

Let us assume that that the agent faces (p, m, z, l) where p is the price vector, m is his endowment of money (or the money value of an endowment of goods at the prices p), z is an endowment of non traded goods and l is his lifespan. Consider the agent's optimal consumption of goods obtained through trade. This is given by:

$$x(p, m, z, l) = \arg \max_x \{U(x, z, l) \mid x \in F, px \leq m\} \quad (2)$$

Note that changing an agent's life span may change their optimal consumption bundle. We can now define the indirect utility function:

$$v(p, m, z, l) = \max_x \{U(x, z, l) \mid x \in F, px \leq m\} \quad (3)$$

Now let

$$S = \{(p, m, z, l) : px \leq m \Rightarrow (x, z, l) \in F\} \quad (4)$$

This is the set of endowments and trading constraints that limit consumption to be in the consumption set F .

Proposition 1

$v(p, m, z, l)$ is continuous and strictly increasing in l on the set S .

Proof in Appendix.

We now fix a reference endowment in terms of prices, money, and non traded goods other than lifespan. Let his reference point be (p^r, m^r, z^r, \cdot) . For the present we consider an arbitrary reference point; in the next section the reference point chosen will have ethical significance.

A6. We limit the admissible allocations such that for all $(x, z, l) \in F$ and $(p, m, z, l) \in S$ we have, for some \bar{l} ,

$$v(p^r, m^r, z^r, 0) \leq U(x, z, l) \leq v(p^r, m^r, z^r, \bar{l}) \quad (5)$$

$$v(p^r, m^r, z^r, 0) \leq v(p, m, z, l) \leq v(p^r, m^r, z^r, \bar{l})$$

This assumption limits the range of allocations we can consider. The first inequality says that all allocations under consideration are at least as good as never being born. This rules out some allocations that are so bad that the agent would rather not exist. The second inequality rules out allocations that are better than an unbounded lifespan at the reference point.

We now examine the existence of a life metric utility function. The issues raised are similar to those for a money metric utility function (examined by Weymark (1985)). We define the direct life metric utility function over consumption bundles $\phi(x, z, l)$ implicitly by

$$v(p^r, m^r, z^r, \phi(x, z, l)) = U(x, z, l) \quad (6)$$

This is the lifespan lived at the reference point endowment that would give the same utility to the agent as the allocation (x, z, l) . We can define life metric indirect utility function by

$\psi(p, m, z, l)$ implicitly by

$$v(p^r, m^r, z^r, \psi(p, m, z, l)) = v(p, m, z, l) \quad (7)$$

This is the lifespan lived at the reference point endowment that would give the same utility to the agent as the endowment (p, z, m, l) . It is immediate that

$$\psi(p, m, z, l) = \phi(x(p, m, z, l), z, l) \quad (8)$$

Further

$$v(p^r, m^r, z^r, \psi(p^r, m^r, z^r, l)) = v(p^r, m^r, z^r, l) \quad (9)$$

and hence, since v is strictly increasing in l we have for all l ,

$$\psi(p^r, m^r, z^r, l) = l \quad (10)$$

Proposition 2 $\psi(p, m, z, l)$ exists, is unique, and is continuous over S.

$\phi(x, z, l)$ exists, is unique, and is continuous over F.

Proof in Appendix

Our approach to constructing life metric utility replicates the approach used by Hammond (1994) to construct money metric utility. The only difference between the two approaches is in the range of allocations covered by the metric. The money metric cannot measure utility in states that are preferable to an infinite quantity of money or are worse than having no money. We cannot measure utility in the life metric in states that are preferable to any bounded lifespan or are worse than never being born.

3. Social Preferences

We now consider a society with n people. To make matters simple we think of one cohort all being born at time zero. The social planner evaluates their welfare at time zero based on their planned lifetime consumption. We have no uncertainty. The feasible set of consumption bundles and preferences of each person i are assumed to obey the model set out in section 2. Each person i has a utility function U_i and associated indirect utility function v_i .

We wish to construct social preference orderings over resource endowments and allocations of goods. We use the symbol \succeq for a weak social preference (as least as good as). Given these social preferences we can define strict social preference \succ and social indifference over states A and B by:

$$A \succ B \Leftrightarrow A \succeq B \text{ and } B \not\succeq A \quad (11)$$

$$A \sqsupseteq B \Leftrightarrow A \succeq B \text{ and } B \succeq A$$

The set $(p, z_i, m_i, l_i) \in \Gamma$ where $\Gamma = S_1 \times S_2 \times \dots \times S_n$ contains the admissible resource allocations for the society. We also have social preferences over consumption bundles $(x_i, m_i, l_i) \in \Omega = F_1 \times F_2 \times \dots \times F_n$. We assume society can also choose between a consumption bundle and a resource allocation. By considering the consumption bundles individuals choose given their endowment, $x_i(p, z_i, m_i, l_i)$, we have that a resource allocation generates a unique consumption bundle (by strict concavity of the utility function).

It is natural to think of social preferences as being over the consumption bundles that people actually consume. However, it is useful to also think of social preferences over endowments. If we undertake a policy to change someone's lifespan or access to a non-traded good, this changes their endowment and consumption of these goods. However, such policies can also affect the individual's optimal consumption bundle of traded goods and, in principle, we want to take into account these induced changes in our analysis.

Definition: A social preference ordering \succeq over (Γ, Ω) is well behaved if it is :

(i) reflexive

(ii) transitive

(iii) continuous

(iv) complete

and

(v) $(p, z_i, m_i, l_i) \sqsupseteq (x_i(p, z_i, m_i, l_i), z_i, l_i)$

(vi) $(p, z_i, m_i, l_i) \succeq (p', z'_i, m'_i, l'_i)$ if and only if $(x_i(p, z_i, m_i, l_i), z_i, l_i) \succeq (x_i(p', z'_i, m'_i, l'_i), z'_i, l'_i)$

Conditions (i)-(iv) are standard. Conditions (v) and (vi) imply that a resource allocation can be identified with the consumption bundle it generates after agents trade. An endowment A is preferred to an endowment B if and only if the consumption bundle associated with A is preferred to the consumption bundle associated with B.

These assumptions imply that we can consider social preferences over resource allocations as equivalent to the social preferences over the consumption bundles chosen by consumers with these endowments.

Definition. A consumption bundle (x_i, z_i, l_i) over n people is weakly Pareto superior to (x'_i, z'_i, l'_i) if and only if for each person i, $U_i(x_i, z_i, l_i) \geq U_i(x'_i, z'_i, l'_i)$ and for at least one person k, $U_k(x_k, z_k, l_k) > U_k(x'_k, z'_k, l'_k)$

Axiom 1

If (x_i, z_i, l_i) is weakly Pareto superior to (x'_i, z'_i, l'_i) then it is strictly socially preferred.

We now assume at some reference point lives are equally valuable. Let the reference point be $R = (p^r, m_i^r, z_i^r, \dots)$. Each agent faces the same price vector, but different agents may have different endowments of money and non-traded goods at the reference point. There may be an ethical argument for treating people symmetrically and giving each person an identical reference point. However, current approaches to cost effectiveness analysis require different reference points for different people. We will undertake our theoretical analysis for a general reference point that may vary cross individuals, and investigate the choice of reference point that generates

cost effectiveness analysis in the next section. At the reference point society treats each person's life as equally valuable and social welfare depends only on life spans.

Axiom 2

Let $\delta \geq 0$ be the rate of social time preference. There exists a reference point $R = (p^r, m_i^r, z_i^r, \cdot)$

at which lives are valued equally. For any $l_i \geq 0, l_i' \geq 0$

$$(p^r, m_i^r, z_i^r, l_i) \succ (p^r, m_i^r, z_i^r, l_i') \Leftrightarrow \sum_i \int_0^{l_i} e^{-\delta t} dt > \sum_i \int_0^{l_i'} e^{-\delta t} dt$$

$$(p^r, m_i^r, z_i^r, l_i) \square (p^r, m_i^r, z_i^r, l_i') \Leftrightarrow \sum_i \int_0^{l_i} e^{-\delta t} dt = \sum_i \int_0^{l_i'} e^{-\delta t} dt$$

This implies if everyone has their reference endowment, and there is no discounting, we can derive the social welfare as the sum of discounted lifespans lived. It is immediate that for $\delta = 0$ we have

$$(p^r, m_i^r, z_i^r, l_i) \succ (p^r, m_i^r, z_i^r, l_i') \Leftrightarrow \sum_i l_i > \sum_i l_i', \text{ for } l_i \geq 0, l_i' \geq 0$$

$$(p^r, m_i^r, z_i^r, l_i) \square (p^r, m_i^r, z_i^r, l_i') \Leftrightarrow \sum_i l_i = \sum_i l_i', \text{ for } l_i \geq 0, l_i' \geq 0$$

At the reference point the only variation between people that society considers is differences in life span, and society is indifferent as to who gets an extra discounted life year at the reference point. Axiom 2 is clearly weaker than assuming that additional life years have equal social value in all circumstances; the more common approach to justifying cost-effectiveness analysis.

An important point is that our axioms do not conflict in any way. Note that Axiom 2 only applies at the reference point where the only variation between allocations is in life spans. If we

have a Pareto improvement, where all life spans either stay the same or rise, Axiom 2 implies the Pareto improvement is preferred so there is no conflict. In cases where some people have increases in life spans, and others have reductions, we can apply Axiom 2 to find which allocation society prefers but in these cases the Pareto principle does not provide a ranking, and again there is no conflict. When one allocation is not at the reference point, or the changes under consideration involve changes in both life span and other goods, the Pareto principle may provide a ranking but Axiom 2 does not apply.

We now show that our axioms generate a unique social preference ordering. If the two axioms are accepted all social states (where preferences obey the assumptions set out in section 2) can be ranked.

Proposition 3. There exists a unique well-behaved social preference ordering on (Γ, Ω) that satisfies Axioms 1 and 2. This social preference ordering can be represented by the social

welfare functions $\sum_i \int_0^{\psi_i(p, z_i, m_i, l_i)} e^{-\delta t} dt$ over Γ and $\sum_i \int_0^{\phi_i(x_i, z_i, l_i)} e^{-\delta t} dt$ over Ω where ψ_i and ϕ_i are the

life metric indirect and direct utility functions respectively of person i given the reference point

$$R = (p^r, m_i^r, z_i^r, \cdot).$$

Proof in Appendix.

Proposition 3 says that to get social welfare we need to find the "life metric" utility of each person as set out in section 2, measured in life years, and add these up over people, after discounting the life years if appropriate. The proof is based on the fact that the Pareto principle, and continuity of the social preference ordering, implies that if everyone is indifferent between

two social states then society must be indifferent. This means that society will be indifferent when we shift from any allocation to the reference point but with life spans adjusted to keep each individual just as well off as before. We can then compare any two allocations by shifting them to the reference point and valuing the different implied distributions of life spans using axiom 2. By transitivity of the social preferences, the ordering of the reference point allocations must be the same as the social ranking of the two original allocations.

We give results both for direct and indirect utility functions. The use of direct utility functions is appropriate when changes in lifespan and health can occur without changing the consumption of any other good. However, in many cases we can think of life spans and health as endowments and when these change and individual will re-optimize, changing their consumption of other goods. For example suppose we have three goods, a consumption good x , a health state z , and lifespan l . Suppose we take x to represent the flow of consumption good in each period of life while z is the flow of health. With a constant personal discount rate δ_i we might write lifetime direct utility as

$$U_i(x_i, z_i, l_i) = \int_0^{l_i} e^{-\delta_i t} u_i(x_i, z_i) dt \quad (12)$$

where $u_i(x_i, z_i)$ is the flow of utility per period, which we assume for simplicity is independent of age. Note that the individual's rate of time preference δ_i need not equal the social rate of time preference δ . We could use this model to analyze the effect of changes in health and life span holding the flow of consumption steady each period.

However we might imagine a case where the individual has the utility function as set out in (12) but maximizes this subject to the budget constraint that the net present value of consumption is less than an initial stock of wealth m_i . That is

$$\max_{x_i(t)} U_i(x_i(t), z_i, l_i) = \int_0^{l_i} e^{-\delta_i t} u_i(x_i(t), z_i) dt \quad s.t. \quad \int_0^{l_i} e^{-r t} x_i(t) dt \leq m_i \quad (13)$$

We can solve this optimal control problem using a Hamiltonian. In the social case where the rate of interest equals the rate of time preference, that is $\delta_i = r$, the optimal path of consumption is steady over time and we can solve (13) for per period consumption to give the indirect utility function

$$v_i(m_i, z_i, l_i) = \int_0^{l_i} e^{-\delta_i t} u_i\left(\frac{m_i r}{1 - e^{-r l_i}}, z_i\right) dt \quad (14)$$

With a fixed sum to allocate over his lifetime an increase in life span has two effects. It increases life years lived, but it forces the agent to spread their fixed wealth out over a longer life and consume less in each period. We could imagine more complex cases where a change in health affects worker productivity, and wage rates, and hence labor supply decisions and income. The use of the indirect utility function allows these indirect effects through changes in behavior to be taken into account when we evaluate the welfare effects of a change in health or lifespan.

This unique social preference ordering has most of the desirable properties we would like in a coherent social ranking. Our social preferences generate a continuous, reflexive and transitive partial order, overcoming the reversal problems in standard cost benefit analysis. It is easy to show that it satisfies the Pareto principle and non-dictatorship which Arrow (1950) has proposed as desirable properties for a social choice rule.

The two weaknesses in our social welfare function from a theoretical standpoint are incompleteness and dependence on “irrelevant alternatives”. There are a set of social states and individual preferences which our social preference ordering cannot rank. We need to assume that individuals always prefer more life to less, ruling out cases where people prefer a shorter

lifespan, or are indifferent as to their lifespan. In addition, we cannot socially rank states that are so bad that individuals would strictly prefer never to have been born. Nor can we rank states that are so good that they are preferred to living forever with the reference allocation. However, our social preference ordering is complete except on these states. Secondly the ranking of two social states depends on the existence of the reference point as well as the two states being considered. Removing this “irrelevant alternative” from the choice set would destroy our ability to compare states.

Since Arrow (1950) shows that no social preference ordering based only on individual preferences can satisfy his desirable properties if completeness and independence of irrelevant alternatives are included, some weakening of these assumptions appears to be required. While our social preference ordering is incomplete in general, there appears to be a wide range of circumstances in which it can be usefully applied, and it is complete over the restricted set of preferences and endowments assumed in section 2. Cost effectiveness analysis usually measures health status relative to a reference point of full health. When we compare two health states we require the third “irrelevant alternative” of full health to exist for our ranking to occur.

This approach to producing a well behaved social preference ordering, and social welfare function, by giving up completeness, has been examined by Chichilnisky and Heal (1983) and can be contrasted with the approach which maintains completeness but assumes the social planner has direct information in the form of a cardinal, interpersonally comparable, measure of each individual’s utility. Sen (1977) and Blackorby et al. (1984) discuss the link between information available to the planner and the type of social preferences that can be derived. Our rankings depend only on individuals’ preference orderings, though our axioms allow us to

generate from these preferences a cardinal, interpersonally comparable, utility measure for individuals with preferences from a restricted domain.

We can construct “life metric” utility by asking people what life span would be required, lived at the reference endowment, to make them indifferent between this and the state under consideration. Note that all that is required for this is for people to give rankings between states; we do not require any direct information on the intensity of their preferences. Our social preference ordering can, however, be represented by a cardinal utilitarian social welfare function made up of the sum of these individual “life metric” utilities.

There is a common objection that valuing lives equally must violate the Pareto principle – we show this is not the case. The argument is that if one person is willing to pay more for a life year than another we should “value” life more highly for the first and give them the life year, while potentially compensating the second to ensure both are better off. When compensation is actually paid we have a Pareto improvement; our first axiom results in the Pareto improvement being socially preferred, even when the sum total of life years declines. However, when compensation is not paid we face a purely distributional question; which person does society think deserves the extra life year more? Traditional cost-benefit analysis favors giving the life year to the person who is willing to pay more. On the other hand, we assume that in this case (at the reference point) society values the claims of each person to an extra life year equally, independently of their willingness to pay for life in money units.

Our two axioms imply that we wish to maximize a utilitarian social welfare function that is the sum of people’s individual utilities. The only unusual aspect of this utilitarian approach is that utility must be measured in life years. Valuing lives equally in our formulation does not make maximizing life years lived a social goal; rather, it makes life years, lived at the standard

level of income and health, a measuring rod for utility and social welfare. It is more usual for economists to use money as a measuring rod. The choice of a measuring rod for utility has no effect on Pareto efficiency; if everyone is better off their utility goes up whatever the metric, but it does have significant consequences when we add over gains and losses to decide distributional issues.

While we have a unique social ordering, the ordering depends on the choice of reference point. Choosing a different reference point will generate a new, unique, social preference ordering. The effect of changing the reference point is familiar from the difference between compensating and equivalent variation in money metric welfare economics. Moving from an *ex ante* to an *ex post* reference point can change the ranking of social states.²

The choice of reference point appears arbitrary but it is really an ethical decision about distribution. Any reference point could be assumed but this choice, together with the Pareto principle, will generate a social preference ordering and will have implications for how society chooses between possible allocations. In this paper we focus on one particular reference point, the endowment that generates cost effectiveness analysis as currently practiced. Our axioms, together with this particular reference point, can therefore be used as a foundation for cost effectiveness analysis. In addition, since the axioms generate a social preference ordering we can extend cost effectiveness analysis to answer questions not usually addressed within its framework.

4. Cost Effectiveness

² We could reconstruct an axiomatic basis for cost benefit using the same approach as set out here with money metric rather than life metric utility. We can then derive compensating or equivalent variation as the correct welfare measure depending on if we assume society values additional units of money equally for each person at the *ex ante*, or *ex post*, allocation.

Suppose we have a fixed sum of money to spend on health care, how should this money be allocated? We assume a fixed sum K has been allocated to health and that providing a lifespan l_i at health state z_i to person i has cost $C(l_i, z_i)$. We assume only one non-traded good z_i for each person and we identify this good as health status. we assume health policy affects only lifespan and health status and not endowments of other goods. We assume all costs are at time zero. Let I be the initial state. Given a reference point $(p^r, z_i^r, m_i^r, .)$ we can measure the utility of each person in the initial state I , with health and lifespan $z_i(I), l_i(I)$ before the health care funds K have been spent, by $l_i^*(I)$ where

$$(p, z_i(I), m_i, l_i(I)) \square_i (p^r, z_i^r, m_i^r, l_i^*(I)) \quad (15)$$

Given the budget constraint we maximize social welfare given the budget constraint:

$$\text{Max}_{l_i, z_i} \sum_i \int_0^{l_i^*} e^{-\delta t} dt \quad \text{s.t.} \quad C(l_1, l_2, \dots, l_n, z_1, z_2, \dots, z_n) \leq K \quad (16)$$

For a choice between two mutually exclusive policies A and B that both cost K we simply ask each person how many years lived at the reference point would be required make them indifferent to the health and lifespan associated with each policy.

$$(p, z_i(A), m_i, l_i(A)) \square_i (p^r, z_i^r, m_i^r, l_i^*(A)), \quad (p, z_i(B), m_i, l_i(B)) \square_i (p^r, z_i^r, m_i^r, l_i^*(B)) \quad (17)$$

Given a description of the health state z_i and life span l_i each person chooses a lifespan lived at the reference point that would leave them indifferent to the effects of the policy under consideration.

A key issue is how we choose the reference point in order to derive cost effectiveness as the appropriate social decision rule. In cost effectiveness analysis the reference health state is usually defined as full health, or perfect health. We take this to be z^p . However the prices, and

income, to be received in the reference state are usually left undefined. We assume that these are defined implicitly at their initial levels. Hence we take as our reference point :

$$p^r = p, m_i^r = m_i^0, z_i^r = z_i^P \quad (18)$$

Our reference point is perfect health and income for each individual at the initial level (without the intervention). We assume prices are fixed throughout the analysis we conduct. It would in principle be possible to define reference point prices as initial prices and take into account in our analysis any induced effect on prices from the health intervention.

Taking this reference point we may regard $l_i^*(A)$ as a measure of quality adjusted life years (QALYS) equivalent to policy A using the time trade-off method. It is the number of life years, lived in perfect health with the current endowment of other goods, that are equivalent to the individual to living $l_i(A)$ years in health state $z_i(A)$ as would be obtained with policy A, holding other endowments equal. We then add up these reference point life years after discounting them at the social rate of time preference. Policy A is strictly preferred to policy B if and only if

$$\sum_i \int_0^{l_i^*(A)} e^{-\delta t} dt > \sum_i \int_0^{l_i^*(B)} e^{-\delta t} dt \quad (19)$$

The inequality in (19) can be written as

$$\sum_i \int_{l_i^*(I)}^{l_i^*(A)} e^{-\delta t} dt > \sum_i \int_{l_i^*(I)}^{l_i^*(B)} e^{-\delta t} dt \quad (20)$$

We take the gain in quality adjusted life years for each person, discount these at the social rate of time preference, and add up over people. The socially preferred policy is the one that gives the biggest gain in discounted quality adjusted life years. The incremental cost

effectiveness ratio of a policy that produces life span and health (l_i, z_i) at a cost $C(l_i, z_i)$ can be written as

$$\frac{C(l_i, z_i)}{\sum_i \int_{l_i^*(t)}^{l_i^*(l_i, z_i)} e^{-\delta t} dt} \quad (21)$$

This is the cost per discounted quality adjusted life year gained by the policy.

If we can envision small changes in policy that result in differences in lifespan and health states across people the Lagrangian for the problem (16) can be written as

$$L(l_i, z_i, \lambda) = \sum_i \int_0^{l_i^*} e^{-\delta t} dt + \lambda(K - C(l_1, l_2, \dots, l_n, z_1, z_2, \dots, z_n)) \quad (22)$$

The social welfare function is continuous since by proposition 2 each indirect utility function is continuous. It what follows we assume that all our utility functions, indirect utility functions, and the social welfare function, are twice differentiable so that we can explore marginal conditions. Crouzeix (1983) shows how differentiability of the indirect utility functions (and hence our social welfare function) can be inferred from assumptions on the differentiability of the underlying utility functions. We also assume the cost function is twice differentiable. Our optimization problem (22) has the first order conditions for interventions that affect life years lived

$$e^{-\delta l_i^*} \frac{dl_i^*}{dl_i} - \lambda \frac{dC}{dl_i} = 0 \quad \text{for all } i \quad (23)$$

This can be rewritten as

$$e^{-\delta l_i^*} \frac{dl_i^*}{dl_i} \frac{dl_i}{dC} = \lambda \quad \text{for all } i \quad (24)$$

A more familiar way of writing this is as

$$\frac{1}{e^{-\delta l_i^*}} \frac{dl_i}{dl_i^*} \frac{dC}{dl_i} = \frac{1}{\lambda} \quad \text{for all } i \quad (25)$$

The gain cost per life year gained by individual i is $\frac{dC}{dl_i}$. The quality adjustment factor is $\frac{\partial l_i}{\partial l_i^*}$, the number of life years lived in the current state that equal a life year at the reference point for person i . Future quality adjusted life years are discounted at the social discount rate δ . This says that the marginal cost of an extra life year for each person, suitably adjusted by the quality of the life year relative to a year of perfect health and discounted if it is a future gain, should be the same for each person. The cost of a discounted quality adjusted life year for each person when social welfare is maximized should equal a common threshold value $1/\lambda$.

Similarly if we consider interventions that improve health status we have

$$e^{-\delta l_i^*} \frac{dl_i^*}{dz_i} \frac{dz_i}{dC} = \lambda \quad \text{for all } i \quad (26)$$

Equation (26) indicates that to evaluate improvements in health status we convert these into equivalent (in the sense of having the same utility) changes in life years lived at the reference point, and discount these future, reference-point, life years.

Combining equations (24) and (26) we have

$$\frac{dl_i^*}{dl_i} \frac{dl_i}{dC} = \frac{dl_i^*}{dz_i} \frac{dz_i}{dC} \quad \text{for all } i \quad (27)$$

Note that the social rate of time preference does not appear in equation (27). Equation (27) is exactly the marginal condition we can derive for maximizing an individual's utility subject to fixed health budget for that individual. It follows that when judging if we should spend an extra dollar on a extending a person's life, or improving the same person's health status and quality of life, the comparison depends only on their own preferences and not on social preferences.

To make matters more concrete, we illustrate quality adjustment using an example. For simplicity, suppose we have only one good in addition to health status and lifespan. Suppose each person i has a utility function as set out in equation (12) with consumption of the “traded” good per unit time fixed independently of health policy. We can think of a flow of both consumption and health status that is constant over the person’s lifetime. The individual's lifetime utility is the sum of the discounted flow of period utility.

We take the reference endowments to be $(x_i^0, z^p, .)$ current income and perfect health, and assume that $u_i(x_i^0, z^p) > 0$, so agents prefer being alive to being dead at the reference point.

Utility measured in life years at the reference point is given by l_i^* in the implicit function:

$U(x_i, z_i, l_i) = U(x_i, z^p, l_i^*)$ or for our specific utility function given in (12):

$$\int_0^{l_i^*} e^{-\delta t} u_i(x_i^0, z^p) dt = \int_0^{l_i} e^{-\delta t} u_i(x_i, z_i) dt \quad (28)$$

Given that the consumption of other goods remains unchanged, the value of a marginal increase in lifespan to person i in life metric utility is

$$\frac{\partial l_i^*}{\partial l_i} = \frac{\frac{\partial U_i}{\partial l_i}}{\frac{\partial U_i}{\partial l_i} \Big|_R} = e^{-\delta_i(l_i - l_i^*)} \frac{u_i(x_i^0, z_i)}{u_i(x_i^0, z^p)} \quad (29)$$

Note that for a person in full health we have $z_i = z^p$ and $l_i^* = l_i$. For such a person a life year and a quality adjusted life year are the same. For a person in less than full health we will have $u_i(m_i, z_i) < u_i(m_i, z^p)$ and $l_i^* < l_i$. For this person an extra life year will be worth less than a quality adjusted life year because it has a lower health state and because it comes at the end of

their lifespan (whereas a marginal quality adjusted life year would come sooner, at the end of a shorter life lived at full health).

Similarly the value of an improvement in health status in life metric utility is given by

$$\frac{dl_i^*}{dz_i} = e^{\delta_i l_i^*} \frac{(1 - e^{-\delta_i l_i^*})}{\delta_i} \frac{\frac{du_i}{dz_i}}{u_i(x_i^0, z^p)} \quad (30)$$

Notice that the QALY compares extra life in the current health state to life at full health, in both cases with income fixed at the individual's initial level. This means that when we quality adjust life years we adjust for health status but not for income. If we were to use the same reference point for each person, with fixed levels of both health and income, we would have to quality adjust life years in both dimensions.

Different methods of quality adjustment have been proposed for QALYs. We need to find the value of a life year lived in each health state in terms of a life year lived in perfect health. It has been proposed that we use the patients own preferences for quality adjustment, or a measure of societies preferences as a whole, or expert opinion. In our formulation quality adjustment is individual specific and reflects person being affects own preferences in trading off life years at the current health state versus life years at full health.

Our approach also makes clear that a full reference point, over all goods, needs to be defined for quality adjustment. While measures of a QALY take perfect health as a reference point they should be precise in what will happen to other goods. One approach is measure welfare relative to a year of life with perfect health but the consumption of all other goods held at their existing levels. An alterative is to think of the measuring rod as a year of life with perfect health but the endowments of all other goods held at their existing levels. In this case we allow the agent to optimize their labor supply and consumption plans in light of their health status.

This issue is reflected in the different results found by Garber and Phelps (1997), who fix the consumption of other goods, and Meltzer (1997) who fixes the endowment of other goods, but allows consumption to change subject to a budget constraint, when analyzing the cost effectiveness of health interventions for an individual. In this section we have assumed other consumption is held fixed when we undertake a health intervention. In the next section we consider health interventions combining with financing mechanisms that also change people's incomes and we have to take account of induced changes in consumption.

5. Financing Health Care

We now turn to the issue of how much of the government budget should be devoted to health care; an alternative way of thinking about this issue is how the threshold for cost effective interventions should be set. For simplicity we focus only on interventions that extend life span but do not affect the health state. We assume people are endowed with money at the beginning on life and the government finances health care using some of this money. We use the model set out in equation (13) and assume each person has a corresponding indirect utility function

$$v_i(p, m_i, z_i, l_i)^3.$$

We first assume that health care is financed through lump sum transfers. Each person is required to give a lump sum k_i , whose amount may be specific to the individual, to the government at the start of life for health care. Life metric utility $l_i^*(p, z_i, m_i, l_i)$ is given by the solution to $(p, z_i, m_i, l_i) \square_i (p, z^p, m_i^0, l_i^*)$. We assume the prices (in this case just the interest rate) are

³ We do not need to assume that the individual's rate of time preference equals the rate of interest though the resulting indirect utility function may then be more complex than set out in equation (14).

fixed throughout the analysis. We take the reference point health to be perfect health. We fix the reference point for income as the initial level of income prior to paying any tax.

Formally, using our social welfare function the social planner has the problem:

$$\text{Max}_{l_i, k_i} \sum_i \int_0^{l_i^*(p, z_i, m_i - k_i, l_i)} e^{-\delta t} dt \quad \text{s.t.} \quad C(l_1, l_2, \dots, l_n) \leq \sum_i k_i \quad (31)$$

the Lagrangian for this problem is

$$L(l_i, k_i, \lambda) = \sum_i \int_0^{l_i^*(p, z_i, m_i - k_i, l_i)} e^{-\delta t} dt + \lambda (\sum_i k_i - C(l_1, l_2, \dots, l_n)) \quad (32)$$

and the first order conditions are

$$e^{-\delta l_i^*} \frac{dl_i^*}{dl_i} \frac{dl_i}{dC} = \lambda \quad \text{for all } i, \quad e^{-\delta l_i^*} \frac{dl_i^*}{dm_i} = \lambda \quad \text{for all } i \quad (33)$$

The first condition is that the money raised for health care should be spent cost effectively; we equalize the marginal cost of a discounted quality adjusted life year across people. This is equivalent to maximizing the discounted, quality adjusted, life years gained with the money raised. The second condition is that the value of money measured in discounted quality adjusted life years is the same across people. It is useful write these conditions as

$$\frac{1}{e^{-\delta l_i^*}} \frac{dC}{dl_i^*} = \frac{1}{\lambda} \quad \text{for all } i, \quad \frac{1}{e^{-\delta l_i^*}} \frac{dm_i}{dl_i^*} = \frac{1}{\lambda} \quad \text{for all } i \quad (34)$$

In this formulation we have that the marginal cost of a discounted quality adjusted life year should be the same for each person, set at a cost effectiveness threshold $1/\lambda$. The second condition is now that each person's willingness to pay for a discounted quality adjusted life year equals the same threshold. Note that these conditions together imply

$$\frac{dC}{dl_i^*} = \frac{dm_i}{dl_i^*} \quad \text{for all } i \quad (35)$$

which is the usual efficiency condition that the cost of a quality adjusted life year for each person equals that person's willingness to pay for it. This efficiency condition means that each person would rather have the health care than be given the money that would be saved if they were prepared to forgo the care.

Garber and Phelps (1997) argue that for Pareto efficiency the marginal cost of health (a quality adjusted life year) should equal an individual's willingness to pay. Cost effectiveness analysis equalizes the marginal cost of health across all individuals, which appears to violate this condition, leading them to argue that it is incompatible with Pareto efficiency. Equation (34) shows this conflict can be resolved; at the social optimum the cost of a discounted, quality adjusted, life year is the same across people, and equal to each person's willingness to pay. Lump sum transfers allow us to resolve the conflict by redistributing money from those with a low value of money to those with a high value of money, measured by their willingness to pay quality adjusted life years for money.

A competitive market would achieve efficiency by allocating health care to those with the highest willingness to pay for it. However, this would not be socially desirable based on our second axiom. We can instead achieve efficiency by redistributing income (through the lump sum transfers) to make the willingness to pay for money, in terms of health forgone, equal across households. The social welfare function that follows from our axioms has significant implications for society's views on the optimal distribution of income. In cost benefit analysis efficiency dictates health care should be given to those most willing to pay money for it, while money transfers are used to address distributional concerns. In cost effectiveness analysis we

address distributional issues by allocating health care equitably, money should then be given to those most willing to pay health for it to ensure efficiency⁴.

With lump sum transfers we can achieve full Pareto efficiency in the allocation ;we keep transferring money until the marginal willingness to pay for money, in quality adjusted life years is the same for each person. However, it may be that due to information constraints, incentive effects, or political forces that resist redistribution, individual specific lump sum transfers are not feasible. In many countries health care is financed out of general taxation. We model this as a proportional income tax. The United Kingdom finances health care in this way and in addition, through the National Institute for Clinical Excellence, uses a cost effectiveness threshold for health interventions (Devlin and Parkin (2004)). We can ask the question, how should this threshold, and the implied tax rate, be set? This question cannot be addressed through standard cost effectiveness analysis but our axiomatic approach allows us to investigate the issue.

Again we consider only interventions to extend life spans. With health care financed only by a proportional tax rate τ we have the social problem

$$\text{Max}_{l_i, \tau} \sum_i \int_0^{l_i^*(p, z_i, m_i(1-\tau), l_i)} e^{-\delta t} dt \quad \text{s.t.} \quad C(l_1, l_2, \dots, l_n) \leq \sum_i m_i \tau \quad (36)$$

Forming the Lagrangian as before, the first order conditions for a maximum can be written as

$$\frac{1}{e^{-\delta l_i^*}} \frac{dC}{dl_i^*} = \frac{1}{\lambda}, \quad \frac{\sum_i m_i}{\sum_i m_i / v_i} = \frac{1}{\lambda} \quad \text{where } v_i = \frac{1}{e^{-\delta l_i^*}} \frac{dm_i}{dl_i^*} \quad (37)$$

Again we have that the taxes raised should be spent according to the cost effectiveness rule. The marginal cost of a discounted quality adjusted life year should be the same for each person, set at

⁴ Some economists may think that achieving efficiency by allocating health care to those most willing to pay money for is natural, while achieving efficiency by allocating money to those most willing to pay health for it is bizarre. The two are symmetrical from an efficiency perspective (as can be seen by re-labeling the goods); the advocacy of one approach over the other is evidence of the advocate's views on distribution.

cost effectiveness threshold $1/\lambda$. The second condition is that the cost effectiveness threshold for funding should be a weighted harmonic mean of each individual's willingness to pay for a discounted quality adjusted life year, where the weights are income levels. The rationale for the weighting by income is that most of the burden of the proportional tax falls on the rich. The harmonic mean comes from the fact that we compute the value of the money losses due to the tax in quality adjusted life years units, which is the inverse of the willingness to pay money for an increase in quality adjusted life years.

It is easy to repeat the analysis set out here to find the appropriate cost effectiveness threshold for different methods of financing the health system. In each case we will be concerned with the value of money, measured in quality adjusted life year units, among those who bear the burden of financing health care. Without lump sum transfers based on full knowledge of individual preferences the social planner is in a second best world and cannot achieve full efficiency. Money raised should be spent according to the cost effectiveness rule, but we may restrict the money available to the health system if the funds are coming from the poor, who have a high value (in life years) of money.

In our example here we have assumed that all of the cost of the health intervention is at time zero and falls on initial wealth. We could look at the effect of future costs borne on future wealth. There is a debate in cost effectiveness analysis on the appropriate discounting of future health gains and future money costs, and if these should be discounted at the same rate. This issue does not arise in our framework. Future quality adjusted life years are discounted at the rate δ . For future money costs we ask the people who will bear these costs the number of quality adjusted life years that they would be willing to give up to avoid these costs. These quality adjusted life years are then discounted at the rate δ like any others.

6. Cost Effectiveness Analysis in Other Sectors

The predominant methodology in use for project appraisal outside the health sector is cost-benefit analysis, where a project is desirable if the sum of consumers' willingness to pay (adding over gains and losses) for it is positive⁵. If the total willingness to pay summed over individuals exceeds the cost we could potentially compensate to losers by lump sum money payments from the winners, leading to a Pareto improvement. If it results in an actual Pareto improvement, in which no one loses and some gain, there seems to be a strong ethical case for such a project. However, without the compensation payments the ethical case for the cost-benefit criterion is much weaker; we have winners and losers and need to make a welfare argument that the gains of the winners outweigh the losses of the losers. The results of a cost-benefit calculation can be justified if money is equally valuable to each person, so that money gains are equivalent to welfare gains, but it seems likely that the marginal utility money is lower for the rich than for the poor, making money an unsuitable metric for interpersonal comparisons of welfare (Boardway (1974), Sen (1977)).

In terms of internal consistency, there is the problem that a project may meet the positive net willingness to pay criterion and be rated as socially desirable. However, having decided the project is to be carried out, the willingness to pay to stop the project may exceed the willingness to pay to keep it going. Following the cost benefit rule, society will now decide not to have the project. In addition, even if refined to counter this problem, the cost-benefit rule fails to be transitive. Such inconsistencies in the social decision rule seem undesirable.

⁵ Formally, we can take this to be the compensating variation, the amount of money the agent could give up when the project occurs and be just as well off as before.

We now reconstruct cost benefit analysis to make it consistent with our social preference ordering. Take the case where we have one non-tradable good s in addition to health status z . We also assume that consumption per period equals income per period m_i is fixed and not affected by the provision on the non-traded. This means we can work with the direct utility function rather than indirect utility. If we assume the social discount rate is zero, the gain in social welfare from changing the allocation of the non-traded good from s_i to s'_i is simply:

$$\sum_i l_i^*(m_i, s'_i, z_i, l_i) - \sum_i l_i^*(m_i, s_i, z_i, l_i) \quad (38)$$

Let us take as a reference point $(p, m_i^0, s_i^0, z^p, .)$. We measure utility relative to a reference point with our initial levels of the public good and income, but perfect health. Utility measured in quality adjusted life years is given by the solution to

$$U_i(m_i^0, s_i^0, z^p, l_i^*) = U(m_i, s_i, z_i, l_i) \quad (39)$$

We assume that there is only one traded good and money is converted one for one into this good (we normalize the price to 1). If marginal changes in all variables are possible and our functions are differentiable we have that differentiating with respect to s_i gives

$$\left. \frac{\partial U_i}{\partial l_i^*} \right|_R \frac{\partial l_i^*}{\partial s_i} = \frac{\partial U_i}{\partial s_i} \quad (40)$$

Note that the marginal utility of a quality adjusted life year on the left hand side of equation (40) is measured with variables set at the reference point while the marginal utility of the public good in the right hand side is measured at actual values of the variables. Hence

$$\frac{\partial l_i^*}{\partial s_i} = \frac{\frac{\partial U_i}{\partial s_i}}{\left. \frac{\partial U_i}{\partial l_i^*} \right|_R} = \frac{\frac{\partial U_i}{\partial s_i}}{\frac{\partial U_i}{\partial m_i}} \frac{\frac{\partial U_i}{\partial m_i}}{\frac{\partial U_i}{\partial l_i}} \frac{\partial U_i}{\left. \frac{\partial U_i}{\partial l_i^*} \right|_R} \quad (41)$$

It follows that

$$\frac{\partial l_i^*}{\partial s_i} = \frac{\partial l_i^*}{\partial l_i} \frac{\frac{\partial m_i}{\partial s_i}}{\frac{\partial m_i}{\partial l_i}} \quad (42)$$

These three terms have an intuitive interpretation. $\frac{\partial m_i}{\partial s_i}$ is the willingness to pay in money units

for the non-traded good. $\frac{\partial m_i}{\partial l_i}$ is the willingness to pay in money units for a year of life. $\frac{\partial l_i^*}{\partial l_i}$ is

the rate at which life years convert into "life metric" utility measured at the reference point, our

quality adjustment of life years as in cost effectiveness analysis. We now have a linear

approximation to the gain in social welfare is given by

$$\sum_i l_i^*(x_i, s'_i, z_i, l_i) - \sum_i l_i^*(x_i, s_i, z_i, l_i) \approx \sum_i \frac{\partial l_i^*}{\partial l_i} \frac{\frac{\partial m_i}{\partial s_i}}{\frac{\partial m_i}{\partial l_i}} (s'_i - s_i) \quad (43)$$

If everyone has perfect health we have $\frac{\partial l_i^*}{\partial l_i} = 1$ and the social value of a project at the reference

point is given by consumers' willingness to pay measured in life units (their willingness to pay in

money units divided by their money value of a life year). Away from the reference point of

perfect health, the life years the agent is willing to pay have to be quality adjusted in the same

way as for cost effectiveness analysis in the health sector.

To give a concrete example again let $U(m_i, s_i, z_i, l_i) = l_i(m_i, s_i, z_i)^{\alpha_i}$. This implies individuals do not discount future life years. Take as the reference point the consumption

$(m_i^0, s_i^0 z^p, .)$. Then the willingness to pay money for the non-traded (the usual figure used in cost benefit calculations) is

$$\frac{\partial m_i}{\partial s_i} = \frac{\frac{\partial U_i}{\partial s_i}}{\frac{\partial U_i}{\partial m_i}} = \frac{m_i}{s_i} \quad (44)$$

The willingness to pay for an increase in life span is still

$$\frac{\partial m_i}{\partial l_i} = \frac{\frac{\partial U_i}{\partial l_i}}{\frac{\partial U_i}{\partial m_i}} = \frac{m_i}{\alpha_i l_i} \quad (45)$$

Hence, the willingness to pay for the non-traded good in life units is

$$\frac{\partial l_i}{\partial s_i} = \frac{\frac{\partial U_i}{\partial s_i}}{\frac{\partial U_i}{\partial l_i}} = \frac{\frac{\partial m_i}{\partial s_i}}{\frac{\partial m_i}{\partial l_i}} = \frac{\alpha_i l_i}{s_i} \quad (46)$$

Note that while the rich are willing to pay more money for the non-traded good in this example they are not willing to give up more of their lifespan for it. Those willing to pay the most life years are those with long lives who have little of the non-traded good.

To derive our life metric social welfare we need to convert the willingness to pay in current life years to willingness to pay life years at the reference point. The rate at which current life years convert to life metric utility, given income is fixed at its initial level, is given by

$$\frac{\partial l_i^*}{\partial l_i} = \left(\frac{m_i^0 s_i z_i}{m_i^0 s_i^0 z^p} \right)^{\alpha_i} = \left(\frac{s_i z_i}{s_i^0 z^p} \right)^{\alpha_i} \quad (47)$$

Combining (46) and (47) we get the willingness to pay for the non-traded good in quality adjusted life years in our simple example is

$$\frac{\partial l_i^*}{\partial s_i} = \left(\frac{s_i z_i}{s_i^0 z^p} \right)^{\alpha_i} \frac{\alpha_i l_i}{s_i} \quad (48)$$

This should be compared with equation (44) which is the willingness to pay in money units usually used in cost benefit analysis. In this example while a person's income affects their willingness to pay money for the public good it does not affect their willingness to pay quality adjusted life years. Using cost benefit analysis, and willingness to pay in money units, great weight is given to the preferences of the rich. Those who are willing to pay the most are the rich who have little of the non-traded good. However if we make people pay in life years, not money, in our example the income level has no effect on willingness to pay.

Our example in this section assumes no social discounting; in general the quality adjusted life years each person is willing to pay should be discounted before adding up. Neither have we considered the cost of the project. As with financing health care, the desirability of any project using cost effectiveness analysis depends on who pays for it. The benefits of a project measured by willingness to pay in quality adjusted life years must be compared to the costs measured in the same units. Money costs must be converted to equivalent reductions in quality adjusted life years among those who bear the burden of these costs. Our approach can also be used in cases where there are time costs as well as money costs to a health policy. In addition we can value projects such as water or sanitation, where there are direct use benefits as well as health benefits. Existing approaches to cost effectiveness analysis convert things into either health benefits or money costs, and it is sometimes unclear where things, such as time savings, should be accounted for. In our axiomatic approach all cost and benefits should be converted to discounted quality adjusted life year equivalents before being summed up.

People may have difficulty in expressing their willingness to pay directly in quality adjusted life years. It is however quite feasible to calculate this from information on their

willingness to pay in money units, their willingness to pay for a life year, and their valuation of life years at their current endowment relative to the reference endowment, as set out in equation (42).

Our extension of cost effectiveness for project appraisal outside the health sector has some parallels in the literature. Equation (42) can also be regarded as a form of weighted cost benefit analysis where we weight each person's willingness to pay in money terms to take account of distributional concerns (Harberger (1978), Brent (1984)). Brekke (1997) investigates the effects of changing the numeraire in cost benefit analysis. Somanathan (2006) has advocated measuring willingness to pay in life units, rather than in money units. Our use of discounted quality adjusted life years as the numeraire is very similar to these approaches.

The cost benefit analysis we have set out has several advantages over the standard approach. We do not get preference reversals and the rule is transitive. However, the question of whether we prefer existing cost benefit or our new approach is essentially an ethical one. Does society think additional money is equally valuable to each person at the current endowment, or does it think additional life years at the reference endowment would be equally valuable to each person?

Our social welfare function measures everything in units of quality adjusted life years and simply add these up over people. Becker et al. (2005) use a money metric approach to convert life span into money units to give find "full income" improvements over the last century. Fleurbaey and Gaulier (2009) construct a social welfare function by adding up "full" incomes combining life span with income in money units, allowing for social inequality aversion. The difference between these approaches and ours is in the choice of a different

numeraire, money rather than life years, for measuring individual utility before adding up over people. This choice has substantial consequences for rankings.

7. Conclusion

Overall our approach gives quite similar project appraisal for the health sector as is used in current cost effectiveness analysis. At the reference point, which we can take to be a life lived in good health, all (discounted) life years gained are weighted equally. For people whose health is not at the reference point, each life year gained is weighted by their own judgment of how many life years at the reference point would be equivalent. This type of weighting is currently carried out using quality adjusted life years to adjust for different health states.

For non-health projects however, our project appraisal is quite different. Instead of using willingness to pay in money terms as a metric we use willingness to pay in life years; how much life would a person be willing to give up for the project? These life years are then adjusted for quality, as in cost effectiveness analysis, before being summed to give the total, quality adjusted, life year value of the project. Our approach has the advantage of internal consistency. All projects can be ranked, and we avoid the reversals of ranking that occur in standard cost-benefit analysis. In addition our generalized cost-effectiveness analysis is a coherent method of project evaluation. All measurements are now in quality-adjusted life years, and so we compare easily across sectors.

For our axioms to generate cost effectiveness analysis we need to impose a particular reference point. All welfare is measured relative to a life year lived in perfect health with the endowment of other goods the same as at present. If the reference point changes, all our internal consistency results remain, but our social preference ordering, and ranking of projects change.

Shifting the reference point at which lives are valued equally will usually produce a completely new, and different, social preference ordering. This is manifested by a different method of quality adjusting life years when we carry out in cost effectiveness analysis. For example we could argue that the correct reference point was perfect health and a fixed “adequate” level of income. In this case we would quality adjust life years not only for their health status relative to perfect health but their income relative to the “adequate” level. Our approach shows these arguments about how to quality adjust life years are equivalent to arguments about the correct social preference ordering.

We consider the welfare of a single cohort, of a fixed size, under certainty. This sidesteps difficult issues associated with intergenerational distribution, and uncertainty. With many generations we have to address issues of discounting both across and within generations. Introducing uncertainty, and assuming that social preferences satisfy the axioms of expected utility theory, places additional restrictions on the allowable social preference orderings (Harsanyi (1955), Fishburn (1984)). We leave these issues for future research.

Appendix

Proposition 1 $v(p, m, z, l)$ is continuous and strictly increasing in l on the set S .

Proof. Continuity is straightforward. Note that the budget set for traded goods does not depend on the lifespan l . Let $l_n \rightarrow l$ and denote the optimal feasible consumption bundle at l by $x^*(l)$ so that $v(p, m, z, l) = U(x^*(l), z, l)$. Fix $\varepsilon > 0$.

Suppose for infinitely many n we have $v(p, m, z, l_n) < v(p, m, z, l) - \varepsilon$.

Since the consumption set of traded goods is closed and bounded compact we can choose a subsequence l_{n_k} such that $x^*(l_{n_k})$ converges. Then since U is continuous we have

$$\lim_{n_k \rightarrow \infty} v(p, m, z, l_{n_k}) = \lim_{n_k \rightarrow \infty} U(x^*(l_{n_k}), z, l_{n_k}) \geq \lim_{n_k \rightarrow \infty} U(x^*(l), z, l_{n_k}) = U(x^*(l), z, l) = v(p, m, z, l)$$

This contradicts every point in the infinite sequence $v(p, m, z, l_n)$ being at least ε below $v(p, m, z, l)$.

Now suppose that for infinitely many n we have $v(p, m, z, l_n) > v(p, m, z, l) + \varepsilon$.

Again by compactness we can construct a converging subsequence $x^*(l_{n_k})$ converging to x' say

. Hence

$$\lim_{n_k \rightarrow \infty} v(p, m, z, l_{n_k}) = \lim_{n_k \rightarrow \infty} U(x^*(l_{n_k}), z, l_{n_k}) = U(x', z, l) \leq U(x^*(l), z, l) = v(p, m, z, l)$$

which contradicts every point in the infinite sequence $v(p, m, z, l_n)$ being at least ε above $v(p, m, z, l)$.

It follows that for any $\varepsilon > 0$ we have $|v(p, m, z, l_n) - v(p, m, z, l)| \leq \varepsilon$ for all but finitely many n and it follows that $v(p, m, z, l)$ is continuous in l .

To see that indirect utility is strictly increasing in l , note that when lifespan increases the agent can feasibly consume the same set of communities as before, with a higher lifespan. Since utility is strictly increasing in l utility at this feasible bundle is strictly higher than before. Optimal consumption must give at least as high a utility as this feasible consumption, and hence indirect utility function is strictly increasing in l . \square

Proposition 2 $\psi_R(p, m, z, l)$ exists, is unique and continuous.

$\phi_R(x, z, l)$ exists, is unique and continuous over F.

proof. Given (p, m, z, l) , then by assumption 6

$$v(p^r, m^r, z^r, 0) \leq v(p, m, z, l) \leq v(p^r, m^r, z^r, \bar{l})$$

Now consider the indirect utility function $v(p^r, m^r, z^r, l)$ as a function of l alone. This function is continuous, and strictly increasing by proposition 1. Hence by implicit value theorem for continuous functions (Jittorntrum (1978)) there exists a unique l^* such that

$$v(p^r, m^r, z^r, l^*) = v(p, m, z, l) \text{ and } l^* = \psi_R(p, m, z, l) \text{ is continuous over } (p, m, z, l) \in S.$$

The proof for $\phi_R(x, z, l)$ is similar. \square

Proposition 3. There exists a unique well-behaved social preference ordering on (Γ, Ω) that satisfies Axioms 1 and 2. This social preference ordering can be represented by the social

welfare functions $\sum_i \int_0^{\psi_i(p, z_i, m_i, l_i)} e^{-\delta t} dt$ over Γ and $\sum_i \int_0^{\phi_i(x_i, z_i, l_i)} e^{-\delta t} dt$ over Ω where ψ_i and ϕ_i are the

life metric indirect and direct utility functions respectively of person i given the reference point

$$R = (p^r, m_i^r, z_i^r, \cdot).$$

Proof. We first address existence. Consider the social welfare function $\sum_i \int_0^{\psi_i(p, z_i, m_i, l_i)} e^{-\delta t} dt$ over Γ

and $\sum_i \int_0^{\phi_i(x_i, z_i, l_i)} e^{-\delta t} dt$ over Ω where ψ_i defined by. This generates social preferences over states by

taking weak preference if and only if the social welfare function gives a value that is at least as

high as the alternative. By proposition 2 every resource allocation in Γ and consumption bundle

in Ω can be ranked by this function so the preference ordering is complete on (Γ, Ω) . It is easy

to see it is a reflexive and transitive social preference ordering since the ordering of the real

numbers is reflexive and transitive. Proposition 2 also ensures that this social welfare function,

and the associated social preferences, are continuous. Hence these social preferences satisfy

conditions (i)-(iv). Now consider preferences over comparisons of resource endowments with consumption bundle. Conditions (v) and (vi) are satisfied immediately by the definitions of the direct and indirect life metric utility functions. Hence this social preference ordering is well behaved.

This social ordering also satisfies the Pareto principle. Suppose (p, z_i, m_i, l_i) is weakly Pareto superior to (p', z'_i, m'_i, l'_i) . Then for every consumer i we have

$\psi_i(p, z_i, m_i, l_i) \geq \psi_i(p', z'_i, m'_i, l'_i)$ and for some consumer k we have

$\psi_k(p, z_k, m_k, l_k) > \psi_k(p', z'_k, m'_k, l'_k)$ and hence $\sum_i \psi_i(p, z_i, m_i, l_i) > \sum_i \psi_i(p', z'_i, m'_i, l'_i)$ so that weak

Pareto improvements are ranked higher on our social order. Similar arguments apply to comparisons of consumption bundles, and our social preferences satisfy axiom 1.

The social ordering also satisfies condition (vi), we value lives equally at the reference point. To see this, consider two allocations that have different life spans at the reference point:

(p^r, m_i^r, z_i^r, h_i) , $(p^r, m_i^r, z_i^r, h_i')$. We the have that

$$\sum_i \psi_i(p^r, m_i^r, z_i^r, h_i) > \sum_i \psi_i(p^r, m_i^r, z_i^r, h_i') \Leftrightarrow \sum_i h_i > \sum_i h_i'$$

Hence an allocation of lifespans at the reference point ranks higher than on our social welfare function if and only if the total years of life gained is larger. It follows that the ranking generated by the social welfare function $\sum_i \psi_i(p, z_i, m_i, l_i)$ satisfies axiom 2.

We now turn to uniqueness. We first prove that our axioms imply that if each individual is indifferent between two endowment vectors then society must be indifferent between them.

Consider two endowment vectors (p, z_i, m_i, l_i) and (p', z'_i, m'_i, l'_i) and assume that each individual is indifferent between their allocations in the two vectors.

Converting each endowment to an allocation, let

$$(x_i, z_i, l_i) = (x(p, z_i, m_i, l_i), z_i, l_i)$$

$$(x'_i, z'_i, l'_i) = (x(p', z'_i, m'_i, l'_i), z'_i, l'_i)$$

Now consider the strict convex combination allocation for

$$a(\lambda) = \lambda(x'_i, z'_i, l'_i) + (1 - \lambda)(x_i, z_i, l_i)$$

By A4 preferences are strictly convex so we have that $a(\lambda)$ is Pareto superior to both

$$(x'_i, z'_i, l'_i) \text{ and } (x_i, z_i, l_i) \text{ for } \lambda \in (0, 1).$$

Hence by axiom 1 $a(\lambda) \succ (x'_i, z'_i, l'_i) \square (x_i, z_i, l_i)$ for $0 < \lambda < 1$.

Now by continuity of the social preference ordering taking limits as $\lambda \rightarrow 0$ we have

$$a(0) = (x'_i, z'_i, l'_i) \succeq (x_i, z_i, l_i)$$

Similarly taking limits as $\lambda \rightarrow 1$

$$a(1) = (x_i, z_i, l_i) \succeq (x'_i, z'_i, l'_i)$$

It then follows that

$$(x'_i, z'_i, l'_i) \square (x_i, z_i, l_i)$$

Now by condition (v) for a well behaved social preference ordering we have

$$(p', z'_i, m'_i, l'_i) \square (p, z_i, m_i, l_i)$$

Hence if every person is indifferent between two allocations society must be indifferent between them.

Now consider two arbitrary (agents need not be indifferent between them) endowments

$$(p', z'_i, m'_i, l'_i), (p, z_i, m_i, l_i)$$

Suppose $\sum_i \psi_i(p, z_i, m_i, l_i) \geq \sum_i \psi_i(p', z'_i, m'_i, l'_i)$ so based on the social preference ordering

constructed in the first part of the proof (p, z_i, m_i, l_i) is weakly preferred to (p', z'_i, m'_i, l'_i) .

Now suppose that there exists an alternative, well behaved, social preference ordering satisfying axioms 1 and 2 with $(p', z'_i, m'_i, l'_i) \succ (p, z_i, m_i, l_i)$

By definition

$$(p^r, m^r, z^r, \psi_i(p, m_i, z_i, l_i)) \sqsupseteq_i (p, z_i, m_i, l_i)$$

and hence

$$(p^r, m^r, z^r, \psi_i(p, m_i, z_i, l_i)) \sqcup (p, z_i, m_i, l_i)$$

since we have shown that if all individuals are indifferent between two endowment vectors any well behaved social preference ordering satisfying axiom 1 must be indifferent between them.

Similarly

$$(p^r, m^r, z^r, \psi_i(p', m'_i, z'_i, l'_i)) \sqsupseteq_i (p', z'_i, m'_i, l'_i)$$

and

$$(p^r, m^r, z^r, \psi_i(p', m'_i, z'_i, l'_i)) \sqcup (p', z'_i, m'_i, l'_i)$$

By assumption

$$\sum_i \psi_i(p, z_i, m_i, l_i) \geq \sum_i \psi_i(p', z'_i, m'_i, l'_i)$$

hence by axiom 2

$$(p^r, m^r, z^r, \psi_i(p, m_i, z_i, l_i)) \succeq (p^r, m^r, z^r, \psi_i(p', m'_i, z'_i, l'_i))$$

It follows that we have

$$(p, z_i, m_i, l_i) \sqsupseteq (p^r, m^r, z^r, \psi_i(p, m_i, z_i, l_i)) \succeq (p^r, m^r, z^r, \psi_i(p', m'_i, z'_i, l'_i)) \sqsupseteq (p', z'_i, m'_i, l'_i)$$

By transitivity of the social preference ordering, this implies that

$$(p, z_i, m_i, l_i) \succeq (p', z'_i, m'_i, l'_i)$$

A contradiction.

Hence the set of weakly preferred points to a particular allocation is the same for any well behaved social preference ordering satisfying axioms 1 and 2. Similarly, by considering a pair of allocations such that $\sum_i \psi_i(p, z_i, m_i, l_i) > \sum_i \psi_i(p', z'_i, m'_i, l'_i)$ we can use the same argument to show that $(p, z_i, m_i, l_i) \succ (p', z'_i, m'_i, l'_i)$ in well behaved social preference orderings satisfying axioms 1 and 2. It follows that the set of strictly preferred points to a particular allocation is the same for any well behaved social preference ordering satisfying axioms 1 and 2. It follows that there is only one well behaved social preference ordering satisfying our two axioms. \square

References

Arrow, K. A difficulty in the concept of social welfare. *The Journal of Political Economy* 1950; 58; 328–46.

Becker, G.S., T.J. Philipson and R. Soares. The quantity and quality of life and the evolution of world inequality. *American Economic Review* 2005; 95; 277-91.

Berlage, L. and R. Renard. The discount rate in cost-benefit analysis and the choice of a numeraire. *Oxford Economic Papers* 1985; 37; 691-99.

Blackorby, C. and D. Donaldson. Consumers' surpluses and consistent cost-benefit tests. *Social Choice and Welfare* 1985; 1; 251-62.

Blackorby, C., D. Donaldson and J.A. Weymark. Social choice with interpersonal utility comparisons: A diagrammatic introduction. *International Economic Review* 1984; 25; 327-56.

Bleichrodt, H. and J. Quiggin. Life-cycle preferences over consumption and health: When is cost-effectiveness analysis equivalent to cost-benefit analysis? *Journal of Health Economics* 1999; 18; 681-708.

Boardway, R.W. The welfare foundations of cost-benefit analysis. *The Economic Journal* 1974; 84; 926-39.

Brekke, K.A. The numeraire matters in cost-benefit analysis. *Journal of Public Economics* 1997; 64; 117-23.

Brent, R.J. Use of distributional weights in cost-benefit analysis: A survey of schools. *Public Finance Review* 1984; 12; 213-30.

Chichilnisky, G. and G. Heal. Necessary and sufficient conditions for a resolution of the social choice paradox. *Journal of Economic Theory* 1983; 31; 68-87.

Crouzeix, J.P. Duality between direct and indirect utility functions : Differentiability properties. *Journal of Mathematical Economics* 1983; 12; 149-65.

Daniels, N. *Just health: Meeting health needs fairly*: Cambridge University Press; Cambridge; 2008.

Devlin, N. and D. Parkin. Does NICE have a cost-effectiveness threshold and what other factors influence its decisions? A binary choice analysis. *Health Economics* 2004; 13; 437-52.

Dolan, P. and R. Edlin. Is it really possible to build a bridge between cost-benefit analysis and cost-effectiveness analysis? *Journal of Health Economics* 2002; 21; 827-43.

- Edejer, T.T.-T., R. Baltussen, T. Adam, R. Hutubessy, A. Acharya, D.B. Evans and C.J.L. Murray eds. *Making choices in health: Who guide to cost-effectiveness analysis*: World Health Organization; Geneva; 2003.
- Fishburn, P. On harsanyi's utilitarian cardinal welfare theorem. *Theory and Decision* 1984; 17; 21–28.
- Fleurbaey, M. and G. Gaulier. International comparisons of living standards by equivalent incomes. *Scandinavian Journal of Economics* 2009; 111; 597-624.
- Fuchs, V.R. and R. Zeckhauser. Valuing health--a "priceless" commodity. *The American Economic Review* 1987; 77; 63-268.
- Garber, A.M. and C.E. Phelps. Economic foundations of cost-effectiveness analysis. *Journal of Health Economics* 1997; 16; 1–31.
- Hammond, P.J. Money metric measures of individual and social welfare allowing for environmental externalities. In *Models and measurement of welfare and inequality*, ed. W. Eichhorn: Springer-Verlag; Berlin and Heidelberg; 1994. 694–724.
- Hansen, B.O., J.L. Hougaard, H. Keiding and L.P. Osterdal. On the possibility of a bridge between cba and cea: Comments on a paper by dolan and edlin. *Journal of Health Economics* 2004; 23; 887-98.
- Harberger, A.C. On the use of distributional weights in social cost-benefit analysis. *Journal of Political Economy* 1978; 86; S87-120.
- Harsanyi, J.C. Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *The Journal of Political Economy* 1955; 63; 309-21.
- Hasmanand, A. and L.P. Østerdal. Equal value of life and the pareto principle. *Economics and Philosophy* 2004; 20; 19-33.
- Jamison, D.T., J.G. Breman, A.R. Measham, G. Alleyne, M. Claeson, D.B. Evans, P. Jha, A. Mills and P. Musgrove eds. *Disease control priorities in developing countries*: Oxford University Press/The World Bank; New York NY; 2006.
- Jittorntrum, K. An implicit function theorem. *Journal of Optimization Theory and Applications* 1978; 25; 575-78.
- Johannesson, M. The relationship between cost-effectiveness analysis and cost-benefit analysis. *Social Science and Medicine* 1995; 41; 483-9.
- Kaplow, L. and S. Shavell. Any non-welfarist method of policy assessment violates the pareto principle. *Journal of Political Economy* 2001; 109; 281-86.

Meltzer, D. Accounting for future costs in medical cost-effectiveness analysis. *Journal of Health Economics* 1997; 16; 33-64.

Neumann, P.J., A.B. Rosen and M.C. Weinstein. Medicare and cost-effectiveness analysis. *New England Journal of Medicine* 2005; 353; 1516-22.

Sen, A.K. *Commodities and capabilities*: Oxford University Press; Oxford; 1985.

Sen, A.K. *Development as freedom*: Knopf; New York; 1999.

Sen, A.K. On weights and measures: Informational constraints in social welfare analysis. *Econometrica* 1977; 45; 1539-72.

Somanathan, E. Valuing lives equally: Distributional weights for welfare analysis. *Economics Letters* 2006; 90; 122-25.

Weinstein, M.C. and W.I. Manning. Theoretical issues in cost-effectiveness analysis. *Journal of Health Economics* 1997; 16; 121–28.

Weinstein, M.C. and W.B. Stasson. Foundations of cost-effectiveness analysis for health and medical practices. *New England Journal of Medicine* 1977; 296; 716-21.

Weymark, J.A. Money-metric utility functions. *International Economic Review* 1985; 26; 219-32.