



PROGRAM ON THE GLOBAL DEMOGRAPHY OF AGING

Working Paper Series

Valuing Lives Equally and Welfare Economics

David Canning
October 2007

Program on the Global Demography of Aging
Harvard School of Public Health
PGDA Working Paper No. 27

Abstract

Welfare economics, in the form of cost-benefit and cost-effectiveness analysis, is at present internally inconsistent and ethically unappealing. We address these issues by proposing two ethical axioms: society prefers Pareto improvements and society values lives lived at a "standard" level of health and income equally. We show that there exists a unique social preference ordering satisfying these axioms. Welfare economics is reconstructed to produce rankings consistent with this social preference ordering. The result is that we should always measure willingness to pay in life years, not money units. A standardized life year becomes an interpersonally comparable unit of value.

JEL Codes: D61, D71, I10

I would like to thank Dean Jamison, Michael Moore and Abhijit Banerjee as well as participants in the joint Harvard/MIT/Boston University Health Economics Seminar for useful discussion on these ideas.

dcanning@hsph.harvard.edu

<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 AG024409-01

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

Walden, David Henry Thoreau

1. Introduction

We take as given the axiom that society prefers Pareto improvements. If a project makes everyone better off we should accept it. This ensures efficiency but still leaves open the question of distribution – if a project makes some people better off and some worse off how should we evaluate it? We need a second axiom to assess distributional effects. Our second axiom is that there exists a standard or "reference" endowment such that, if everyone had this endowment, society would value additional life years positively and equally for each person (though it may choose to apply a social rate of time preference to discount future life years). If everyone has the reference endowment, including personal abilities, health status, income, and access to non-traded goods, society is indifferent as to who gets an extra life year. Society need not value life years equally when people have different endowments.

We show that these axioms generate a unique social preference ordering. This social preference ordering is shown to have all of the desirable properties we would like in a coherent social ranking, provided people value life. 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. Our social preference ordering can be represented by a 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 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.

We use our social preference ordering to construct a new approach to welfare economics and project appraisal. If society is at the reference point, so that everyone has the reference point endowment, we get standard cost effectiveness for health interventions as currently practiced. The welfare gain of a project is the sum of life years gained. Away from the reference point the analysis changes however; life years gained must be quality adjusted, where the adjustment

factor is the number of healthy life years at the reference endowment that would give the same utility gain as a life year in the current state. Our "quality adjusted" life years are therefore adjusted for the full utility flow from the life being extended, not just adjusted for health related disability. The adjustment depends on the individual's preferences for a year of life in her current state versus life lived at the reference endowment.

Cost benefit analysis undergoes a more substantial change under our social preference ordering. When we are at the reference point, the net benefit is the sum of the willingness of people to pay for the project in life years. Away from the reference endowment the life years being paid must be quality adjusted to life years lived at the reference point based on individual preferences. This "life-metric" approach tends to give more weight to the preferences to the poor than in 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.

Perhaps surprisingly, our two axioms completely determine society's attitudes about inequality and tradeoffs between efficiency and equity. The fact that the social preference ordering satisfying the two axioms is unique is a very strong result. If our two axioms are accepted any two states can be socially ranked given information on individual preferences. For example, we can rank social states that have different levels of average income and income inequality. We show how social aversion to income inequality depends on individual preferences over money and life. We may also be averse to inequality in life spans. Our social welfare function implies a trade off between gains in the average lifespan against the variance in life span, with the weight given to avoiding variance in lifetimes being the social rate of time preference.

2. Relationship to existing cost benefit and cost effectiveness analysis.

Welfare economics is in a very unsatisfactory state. The criteria being used face a number of serious difficulties, both in terms of their ethical underpinnings and internal consistency. The predominant methodology in use 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 positive¹. If the total willingness to pay summed over individuals who gain 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. This 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 (Boardman (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 rule fails to be transitive. Such inconsistencies in the social decision rule seem undesirable.

¹ 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.

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 the objective is to maximize total healthy life years produced with a fixed budget. This assumes society values lives, or more precisely discounted, healthy life years, equally across people and wants to maximize the total number produced, independent of who gets them. While appealing, valuing lives, or life years, equally across people has some ethical difficulties of its own. In particular, it appears to rule out exchanges where one person gets extra health care, and lifespan, by paying another to forego care. This exchange may reduce total life produced but make both better off, if the buyer prefers the lifespan and the seller prefers the money received. Always valuing lives equally is inconsistent with the Pareto principle that society should prefer allocations in which everyone is better off (Weinstein and Manning (1997)).

In addition, cost effectiveness analysis is inconsistent with cost benefit analysis except under extreme assumptions on the nature of individual preferences (Johannesson (1995), Dolan and Edlin (2002)). The use of these two inconsistent criteria makes it difficult to allocate resources coherently between the health sector and other sectors of the economy.

The approach taken in this paper is to derive project evaluation criteria from axioms that represent ethical principles. The Pareto principle is a natural ethical criterion for social choice. Our second axiom comes from the principle that lives of different people should in some sense be valued equally by society, the principle that underlies cost-effectiveness analysis. However we only assume this principle holds at one allocation on goods, when every one has the same "reference" endowment. This is clearly much weaker than valuing lives equally in all circumstances. Ranking healthy lives equally at all times has been defended (e.g. Culyer

(1989))². However it means that essentially we cannot judge between qualities of life. If we think of one person's life as better than another in some sense it seems reasonable that we should be more interested in saving the "better" life (Broome (1985), Broome (2002)). More formally, it can be shown that always ranking lives equally implies that our social welfare function must simply be the sum of healthy life years lived, and that it is not affected in any way by the quality of lives lived (Hasmanand and Østerdal (2004)), which violates the Pareto principle. Valuing lives equally is only compatible with the Pareto principle if we limit severely the circumstances in which the rule is applied.

We show that our two axioms, the Pareto principle and valuing lives equally at the reference endowment, generate a unique social preference ordering. This 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 satisfies the Pareto principle, non-dictatorship, and the independence of irrelevant alternatives, which Arrow (1950) has proposed as desirable properties for a social choice rule.

The one weakness in our approach is incompleteness; 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. Since Arrow (1950) shows that no social preference ordering, based only on individual preferences, can

² While the life years of people at different incomes are ranked equally in cost effectiveness analysis, the life years of those with disabilities are given a lower weighting. It seems inconsistent to give lower weights to the disabled and not the poor when the results of poverty may make lives as unpleasant as being disabled.

satisfy his desirable properties if completeness is included, incompleteness appears to be a necessary weakness in our theory. While our social preference ordering is incomplete, there appears to be a wide range of circumstances in which it can be usefully applied.

Our second axiom appears to be non-welfarist in the sense that it does not depend on individual preferences; society likes people to have longer life spans independent of what they themselves want. However, limiting preferences to those where additional lifespan is desirable means that we do not have a conflict between individual preferences and social preferences. In particular, we avoid the result that non-welfarist social preferences necessarily violate the Pareto principle (Kaplow and Shavell (2001)).

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, Donaldson 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.

Somanathan (2006) has advocated measuring willingness to pay in life units, rather than in money units, when undertaking cost benefit analysis. However, we only use this measure when the allocation of endowments matches the reference point. When the current allocation of endowments is far away from the reference point the current life years people are willing to pay have to be adjusted to life years at the reference point, as is done for cost effectiveness analysis.

We have to adjust each life year people are willing to pay for its quality, converting it to a volume of standardized (reference point) life years.

The shift from a money-metric to a life-year metric for cost benefit analysis can be seen simply as an issue of choice of a numeraire, or unit of measurement. However, in cost benefit analysis the choice of numeraire has real effects (Berlage and Renard (1985)) because the numeraire becomes an interpersonally comparable unit of value. While any good could in principle be chosen to fulfill this role, on ethical grounds we give a privileged place to life-years.

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)). We may therefore recognize a moral claim to healthy life without recognizing claims to other goods in the same way. Rawls (1971) uses an index of primary goods to measure wellbeing for the purposes of distributive justice. This gives an objective measure of wellbeing independent of people preferences; we use healthy life years lived at the reference income level as such a measure. 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; 1999)).

Our approach gives a single method that can be applied to both the health sector and more general project appraisal. The lack of Pareto efficiency under standard cost effectiveness analysis often leads economists to argue that we should move towards cost benefit analysis in the analysis of health issues (Fuchs and Zeckhauser (1987)). Our resolution maintains the Pareto principle but leads to fairly minor changes in cost effectiveness analysis (life years gained must be adjusted for all factors that impact their quality, not just disability); instead it is cost benefit analysis that is required to undergo a major revision.

3. 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 non-traded goods and $l \in [0, \infty)$ is lifespan. We assume that the consumer has a complete, reflexive, transitive, and continuous preference relationship 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 given by

$$xp = \sum_j x_j p_j \leq \sum_j e_j p_j = m$$

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. The continuity of the utility function assumed in A3 could be derived from a continuous preference ordering. Note that the utility function U can be changed by any positive monotonic transformation and still represent the same 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 which makes life worth living and puts a bound on how low \underline{x} can be. The 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 traded goods obtained through trade. This is given by:

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

We can now define the indirect utility function:

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

Now let

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

Proposition 1

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

Proof in Appendix.

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

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

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

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)).

Define a reference point R of endowments other than lifespan by $(p^r, m^r, z^r, .)$. We define the direct life metric utility function as $\phi_R(x, z, l)$ implicitly by

$$v(p^r, m^r, z^r, \phi_R(x, z, l)) = U(x, z, l)$$

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

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

$$v(p^r, m^r, z^r, \psi_R(p, m, z, l)) = v(p, m, z, l)$$

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

$$\psi_R(p, m, z, l) = \phi_R(x(p, m, z, l), z, l)$$

Further

$$v(p^r, m^r, z^r, \psi_R(p^r, m^r, z^r, l)) = v(p^r, m^r, z^r, l)$$

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

$$\psi_R(p^r, m^r, z^r, l) = l$$

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

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

4. Social Preferences

We now consider a society with n people. 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:

$$A \succ B \Leftrightarrow A \succeq B \text{ and } B \not\succeq A$$

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

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) $(p, z_i, m_i, l_i) \sim (x_i(p, z_i, m_i, l_i), z_i, l_i)$

(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 resource allocation 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 $R = (p^r, m^r, z^r, .)$. This gives each agent the same price vector, the same money endowment and the same endowment of non-traded goods. When the endowment of a variable is the same for every agent we do not subscript the variable with the person.

Axiom 2

There exists a reference point $R = (p^r, m^r, z^r, .)$ at which lives are valued equally:

$$(p^r, m^r, z^r, l_i) \succ (p^r, m^r, z^r, l'_i) \Leftrightarrow \sum_i l_i > \sum_i l'_i, \text{ for } l_i \geq 0, l'_i \geq 0$$

This implies if everyone has the reference endowment we can derive the social preference as the sum of lifespans lived. Society is indifferent as to who gets an extra life year. Note that in this axiom we do not apply the principle that social preferences should depend on individual preferences. A utilitarian might argue that we should allocate extra units of life to those who would most enjoy it. Here we take the view that, given everyone has the same reference endowment, society does not wish to differentiate between people when it comes to allocating extra life years. If there is a social rate of time preference we may wish to discount future life years relative to current life years – we analyze the implications of discounting future life years in section 6.

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 $\sum_i \psi_i(p, z_i, m_i, l_i)$ over Γ and $\sum_i \phi_i(x_i, z_i, l_i)$ 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

$$R = (p^r, m^r, z^r, \cdot).$$

Proof in Appendix.

The proof is based on the fact that the Pareto principle implies that if everyone is indifferent between two social states then society must be indifferent. To see this, suppose we have two states between which everyone is indifferent, but society strictly prefers one to the other.

Consider convex combinations of the two states that are close to the state society thinks is worse.

By strict concavity all agents are better off at these convex combinations, so by the Pareto principle they must be socially preferred to both the end points. But now we have that one end point is socially strictly worse than the other but it has points arbitrarily close that are strictly socially better. This violates continuity of the social preferences.

This means that society will be indifferent when we shift any allocation to an allocation at 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 through axiom 2. By transitivity of the social preferences, the unique ordering of the reference point allocations must be the same as the social ranking of the two original allocations.

It may seem odd that we can reconcile the Pareto Principle and valuing lives equally. Assume that we are at the reference point in terms of endowments. Take the case where we have to choose between adding one year of life to person A's lifespan and adding one and a half years to person B's lifespan. Axiom 2, valuing lives equally, says we must give the life to person 2. However, suppose person A is willing to pay \$100,000 dollars for his extra life year but B is only

willing to pay \$50,000 for her year and a half of life. We can give person A a year of life and compensate person B with \$75,000 say, making both better off. This Pareto dominates just giving extra life to B, is ranked higher on our social welfare function, and is socially preferred. These two decisions are not in conflict. Note that axiom 2 only applies when we are redistributing life among people holding everything else constant; when compensation is also being paid in the form of other goods axiom 2 is silent. In particular, axiom 2 does not imply society must prefer a year and a half of extra life to B over a year of extra life to A plus a transfer of \$75,000 to B.

If no compensation is actually paid, a traditional cost benefit analysis would still rank giving a year of life to A higher than a year and a half of life to B, on the grounds that A has a higher willingness to pay. Giving a year of life to A maximizes social welfare in terms of total welfare produced measured in money units. Our social welfare function favors giving a year and a half of life to B because this maximizes welfare in a life metric. When we face distribution questions, the life metric and money metric approaches give different answers. However, when faced with an actual Pareto improvement each person is better off. So the sum of individual utilities increases using either metric, and the Pareto improvement is always socially preferred.

Nothing in the proof hinges on the fact that the reference point is the same for each person. We can set different reference points for different people if we choose. There is a strong ethical argument that the reference point should be the same for each person. However, a plausible alternative is to take the reference point to be the status quo at the start of the welfare analysis. Using the initial endowments as the reference point leads to a substantial simplification in implementation since we value current live spans equally and there is no need to adjust these to “standardized” life years lived at some other reference point. The difficulty with taking the

status quo as the reference point is that this can only be done once, after which the reference point is fixed. As time passes decisions must be made based on this historical distribution of endowments which is problematic if it lacks a strong ethical basis. In response to a change in the allocation of resources, changing the reference point to keep it at the current allocation is unsatisfactory since it is equivalent to choosing an entirely new social preference ordering, which not surprisingly can cause a reversal of social preferences over pairs of choices.

5. A New Approach Cost Effectiveness and Cost Benefit Analysis

We now sketch how cost effectiveness and cost benefit analysis can be carried out so that projects are ranked in a manner consistent with the social ordering constructed in section 4. In fact, both cost effectiveness and cost benefit are now conceptually very easy. We have a social welfare function that represents our ordering. This is the sum of individual utilities measured in our "life metric." These life metric utilities can be constructed through revealed preference; for any proposed allocation for a consumer, what lifespan at the reference point would make the consumer indifferent? The project with the higher sum of life metric utilities should be ranked higher and is socially preferred. We now show what this means in practice and how it differs from standard approaches.

We first consider cost effectiveness. Take a policy that redistributes lifespans. For example society has to choose between two medical interventions that give different life year gains to different groups of people. Let h_i be the net effect (either positive or negative) on the life span of person i from moving from intervention 1 to intervention 2. Suppose initially society is at the reference point $(p^R, m_i^R, z_i^R, l_i^R)$. That is, at the resource allocation that holds before the

policy is implemented, society is indifferent between who gets additional lifespan. The gain in social welfare from this policy, measured in the life metric at the reference point, is given by

$$\sum_i \psi_i(p^R, m_i^R, z_i^R, l_i^R + h_i) - \sum_i \psi_i(p^R, m_i^R, z_i^R, l_i^R) = \sum_i h_i$$

It follows that the change in the distribution of lifespans is socially preferred if and only if the total number of life years lived increases. If we have two ways of spending a fixed sum of money on health care, the only effect of which is to change lifespans, leaving the other elements of the resource allocation constant, we should prefer the policy that generates more total life years.

At the reference point we get our standard from of cost effectiveness. However when the initial position is not the reference point the effect of change in lifespans is harder to analyze.

The gain in welfare from a redistribution of lifespans is now

$$\sum_i \psi_i(p, m_i, z_i, l_i + h_i) - \sum_i \psi_i(p, m_i, z_i, l_i)$$

In general, this is not the life years gained by the policy but rather the sum of the lifespan increases at the reference point, which would give the same gain in utility.

In general we can think of agents as consuming a vector of goods. These goods can be indexed by time to represent consumption at different ages. To simplify matters for exposition we take a model where we have only one tradable good and one non-tradable good. This “good” can be interpreted as a fixed rate of consumption that occurs over the entire lifespan. We assume that the non-tradable is a measure of health status. For example it could be a scalar measure of disability. Further let us assume that the price of the tradable good is one. In addition the consumer has an endowment of the tradable goods (whatever lifespan is) of m . It follows that we have $v(1, m, z, l) = U(m, z, l)$ and we can take the indirect utility function to be the direct utility function with the quantity of the tradable good consumed being m . We can write the indirect

utility function of person i as $U_i(m_i, z_i, l_i)$. Let us take as the reference point the allocation $(m^R, z^R, .)$ which is the same for all consumers. It is normal in cost effectiveness analysis to take as the reference point the state in which a person is in good health; we do this as well, though it is not necessary. As should be clear from our discussion of the reference point, setting the reference health standard at a lower level may affect social preferences. Let $l_i^* = \phi_i(m_i, z_i, l_i)$ be the life span at the reference point which gives consumer i the same level of utility as (m_i, z_i, l_i) ; this is just the life metric utility level. This also has the property that

$$U_i(m_R, z_R, l_i^*) = U_i(m_i, z_i, l_i)$$

It is useful to assume the utility function is differentiable so we can examine marginal rates of substitution and the effects of “small” changes in endowments. Differentiating with respect to l_i we have

$$\left. \frac{\partial U_i}{\partial l_i^*} \right|_R \frac{\partial l_i^*}{\partial l_i} = \frac{\partial U_i}{\partial l_i}$$

where we add the subscript R in $\left. \frac{\partial U_i}{\partial l_i^*} \right|_R$ to emphasize it is calculated at (m_R, z_R, l_i^*) rather than the point (m_i, z_i, l_i) .

Hence

$$\frac{\partial l_i^*}{\partial l_i} = \frac{\frac{\partial U_i}{\partial l_i}}{\left. \frac{\partial U_i}{\partial l_i^*} \right|_R}$$

Adding lifespan to a consumer increases his life metric utility. This gain depends on the gain in utility from the lifespan increase divided by the gain in utility from an increase in lifespan at the

reference income and current utility equivalent lifespan. For small changes in lifespans h_i we have that social welfare increases by

$$\sum_i \frac{\partial \psi_i(p, m_i, z_i, l_i)}{\partial l_i} h_i = \sum_i \frac{\partial l_i^*}{\partial l_i} h_i = \sum_i \frac{\frac{\partial U_i}{\partial l_i}}{\frac{\partial U_i}{\partial l_i^*} \Big|_R} h_i$$

It follows that when giving a social ranking to life span increases, the life span increases of different people should be weighted differently, where the weight is number of life years at the reference point that give the same gain in utility as a life year at the current endowment (that is, the marginal rate of substitution between life years at the reference point and life years at the current endowment).

To make matters more concrete, suppose person i has the utility function

$U_i(m_i, z_i, l_i) = l_i(z_i m_i)^{\alpha_i}$, with $0 < \alpha < 1$. In order to satisfy our axioms we require the money and non-traded good endowment of the agent always to be positive so that life is valuable. The reference allocation is (m_R, z_R, \cdot) which we assume and has the property that $(z_R m_R) > 0$, 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: $U(m_i, z_i, l_i) = U(m_R, z_R, l_i^*)$. This implies

that $l_i^* = l_i \left(\frac{z_i m_i}{z_R m_R} \right)^{\alpha_i}$ and social welfare is just the sum of these life metric utilities.

The value of a marginal increase in lifespan to person i in social welfare is

$$\frac{\partial l_i^*}{\partial l_i} = \frac{\frac{\partial U_i}{\partial l_i}}{\frac{\partial U_i}{\partial l_i^*} \Big|_R} = \left(\frac{z_i m_i}{z_R m_R} \right)^{\alpha_i}$$

It follows that we weight gains to life span for richer, and healthier, people more heavily in our social welfare function. Even though we value lives equally at the reference point, the Pareto principle makes us value lives unequally elsewhere.

Weighting life span changes equally across people in all states of the world violates Pareto efficiency. It appears if we want to maintain Pareto efficiency we have to weight people differently sometimes. Weighting different people differently, in terms of the social value of lifespan increases, raises many of the same objections used against cost benefit analysis as a method for allocating health care. However our outcome, while also unequal, is somewhat different. Standard cost benefit analysis assesses benefits by the sum of the individuals' willingness to pay for them. Given our utility function, the willingness to pay, in money units, for a marginal increase in lifespan by individual i is:

$$\frac{\frac{\partial U_i}{\partial l_i}}{\frac{\partial U_i}{\partial m_i}} = \frac{m_i z_i}{\alpha_i l_i}$$

If variations in income are much larger than variations in lifespans this implies a much larger degree of disparity in weights on life span increases in standard cost-benefit than is used in our social ordering. Note also what happens when α_i is very small for one person.

We now turn to reconstructing cost benefit analysis to make it consistent with our social preference ordering. The increase in social welfare from a change in the allocation of the non-tradable z_i to an allocation z'_i is :

$$\sum_i \psi_i(p, m_i, z'_i, l_i) - \sum_i \psi_i(p, m_i, z_i, l_i)$$

We can write the indirect utility function as $v(1, m, z, l) = U(m, z, l)$. Let us take as a reference point $(1, m^r, z^r, l^r)$. Now setting $l_i^* = \psi_i(m_i, z_i, l_i)$ we have

$$U_i(m^r, z^r, l_i^*) = U(m_i, z_i, l_i)$$

Differentiating with respect to z_i gives

$$\left. \frac{\partial U_i}{\partial l_i^*} \right|_R \frac{\partial l_i^*}{\partial z_i} = \frac{\partial U_i}{\partial z_i}$$

Hence

$$\frac{\partial l_i^*}{\partial z_i} = \frac{\frac{\partial U_i}{\partial z_i}}{\left. \frac{\partial U_i}{\partial l_i^*} \right|_R} = \frac{\frac{\partial U_i}{\partial z_i}}{\frac{\partial U_i}{\partial m_i}} \frac{\frac{\partial U_i}{\partial m_i}}{\left. \frac{\partial U_i}{\partial l_i^*} \right|_R} = \frac{\frac{\partial U_i}{\partial z_i}}{\frac{\partial U_i}{\partial m_i}} \frac{\frac{\partial U_i}{\partial l_i}}{\left. \frac{\partial U_i}{\partial l_i^*} \right|_R}$$

It follows that

$$\frac{\partial l_i^*}{\partial z_i} = \frac{\frac{\partial l_i^*}{\partial l_i}}{\frac{\partial m_i}{\partial l_i}} \frac{\partial m_i}{\partial z_i}$$

These three terms are familiar. $\frac{\partial m_i}{\partial z_i}$ 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, as in cost effectiveness

analysis. We now have

$$\sum_i \psi_i(p, m_i, z'_i, l_i) - \sum_i \psi_i(p, m_i, z_i, l_i) \approx \sum_i \frac{\partial l_i^*}{\partial l_i} \frac{\frac{\partial z_i}{\partial m_i}}{\frac{\partial l_i}{\partial m_i}} (z'_i - z_i)$$

At the reference point 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 the life years

the agent is willing to pay have to be adjusted to "life metric" utility, of equivalent life years measured at the reference point, as is done in cost effectiveness analysis

To give a concrete example again let $U(m_i, z_i, l_i) = l_i(m_i z_i)^{\alpha_i}$ and take as the reference point the consumption (m_R, z_R, \cdot) . Then the willingness to pay for the non-traded (the usual figure used in cost benefit calculations) is

$$\frac{\frac{\partial U_i}{\partial z_i}}{\frac{\partial U_i}{\partial m_i}} = \frac{m_i}{z_i}$$

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

$$\frac{\frac{\partial U_i}{\partial l_i}}{\frac{\partial U_i}{\partial m_i}} = \frac{m_i}{\alpha_i l_i}$$

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

$$\frac{\frac{\partial l_i}{\partial z_i}}{\frac{\partial l_i}{\partial m_i}} = \frac{\frac{\partial U_i}{\partial z_i}}{\frac{\partial U_i}{\partial m_i}} = \frac{\frac{\partial m_i}{\partial z_i}}{\frac{\partial m_i}{\partial l_i}} = \frac{\alpha_i l_i}{z_i}$$

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 is

$$\frac{\partial l_i^*}{\partial l_i} = \left(\frac{m_i z_i}{m_R z_R} \right)^{\alpha_i}$$

Hence

$$\frac{\partial l_i^*}{\partial z_i} = \frac{\partial l_i^*}{\partial l_i} \frac{\partial z_i}{\partial m_i} = \left(\frac{m_i z_i}{m_R z_R} \right)^{\alpha_i} \frac{\alpha l_i}{z_i}$$

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, the income level has no effect on willingness to pay. While the rich and poor may be willing to pay the same in terms of life years in our social welfare function, we weight the life years of the rich more heavily since they are higher quality life years. Note that the weight we use to convert a person's life year to a life year at the reference point depends on that person's own preferences of how much their current life year is worth relative to a reference point life year.

One of the standard arguments for cost benefit analysis is that it generates potential Pareto improvements. If we allocate non-traded goods using willingness to pay criteria, those who gain from the goods could potentially compensate the losers to get a Pareto improvement. This argument however does not imply that society prefers the allocation; total money metric losses may be smaller than total money metric gains, but we may weight the losses more heavily. It is important to note that if actual compensation is paid, and an actual Pareto improvement results, our new cost benefit criteria will always rank the Pareto improvement higher than the previous state.

We can generate something very similar to current cost benefit analysis if we assume money is equally valuable to each person and construct the money metric social welfare function, taking the current allocation as the reference point. We cannot shift the reference point as is usually done in existing cost benefit analysis when the situation changes, but this problem

will only become apparent over time as we move away from the original endowment. 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 uniform reference endowment are equally valuable to each person?

6. Inequality Aversion, the Reference Point and Discounting

In our framework there will generally be social aversion to inequality in income. This will occur if the marginal utility of income declines faster than the marginal utility of life extension. In this case as a person grows richer additional units of income are equivalent to ever smaller extensions of life span. Since the social welfare is measured in life years, additional income for the rich is less socially valuable than additional income for the poor.

To illustrate this suppose we have the utility function $U(m_i, z_i, l_i) = l_i \log(z_i) \frac{m_i^{1-\rho}}{1-\rho}$

which is the same for each person i . This implies a constant coefficient of relative risk aversion, ρ over income. The reference allocation is (m_R, z_R, \cdot) which we assume and has the property

$$\log(z_R) \frac{m_R^{1-\rho}}{1-\rho} > 0. \text{ Social welfare is then } \sum_i l_i \frac{\log(z_i) m_i^{1-\rho}}{\log(z_R) m_R^{1-\rho}}$$

Suppose everyone has the same life span l and endowment of non-traded goods, z .

Further assume for expositional purposes³ that we have a continuum of agents normalized to size one so that social welfare can be written as

³ Moving from a finite set of agents to an infinite set may cause problems with our approach since total social welfare may become infinite and comparing between infinities is non-trivial. However, the purpose of this section is expositional and we think of the integral as an approximation to social welfare per capita when the number of agents is large.

$$k \int m^{1-\rho} f(m) dm \quad \text{where} \quad k = l \frac{\log(z)}{\log(z_R)} \frac{1}{m_R^{1-\rho}}$$

and $f(m)$ is the distribution of income. Suppose that income, m , has mean μ and a lognormal distribution with the standard deviation of $\log(m)$ being σ . Then it is easy to show (using the moment generating function of the normal distribution) that social welfare is a strictly increasing transformation of

$$\log \mu - \rho \frac{\sigma^2}{2}$$

Society prefers higher average income levels but is worse off as income inequality increases. Our simple utility function gives a tradeoff between average income and inequality in income that depends on the coefficient of relative risk aversion, a measure of how fast the marginal utility of income is declining. How much society is willing to trade off average income for reduced inequality will in general depend on individual preferences and how fast the marginal utility of money declines relative to the marginal utility of lifespan gains for each person.

While individual preferences determine society's attitude to income inequality we have assumed that at the reference point the lives of each individual are weighted equally. This means that we care about the total life years lived at the reference point but are indifferent as to their distribution. Society may want to discount future life years for individuals more than current life years. It is easy to reformulate axiom 2 to incorporate this.

Suppose everyone is born at the same time and we carry out our evaluation over social states at the beginning of life. Let an agent's life span be l years. The discounted life span, with the discount rate δ , can be written

$$\int_0^l e^{-\delta t} dt = \frac{(1 - e^{-\delta l})}{\delta}$$

For $\delta = 0$ the discounted life span is just l , the actual life span (by L'hospital's rule the discounted life span tends to l as $\delta \rightarrow 0$).

Axiom 2b

Given the fixed discount rate $\delta > 0$ and any $l_i \geq 0, l'_i \geq 0$

$$(p^r, m^r, z^r, l_i) \succ (p^r, m^r, z^r, l'_i) \Leftrightarrow \sum_i \frac{1}{\delta} (1 - e^{-\delta l_i}) > \sum_i \frac{1}{\delta} (1 - e^{-\delta l'_i}),$$

We can now write the social welfare function as

$$\sum_i \frac{1}{\delta} (1 - e^{-\delta \psi_i(p, z_i, m_i, l_i)})$$

Everything in section 2 goes through with this social welfare function instead of simply adding up the simple "life metric" utility gains. It should be clear that in this formulation the choice of a social discount rate is an arbitrary social preference and is not linked to any discounting by individuals.

Our individual preferences do not assume that the same goods are being consumed over time, or the time separability of the utility function, so that the concept of individual time preference need not well defined (Becker and Mulligan (1997)). However, if we impose restrictions that allow for the concept of individual time preference we find there is no necessary connection between individual time preferences and the social rate of time preference used to discount future life years. In cost benefit analysis, it is common to argue that to generate efficiency the social rate of time presence should be related to the market rate of interest or individuals' rate of time preference. In our framework this is not the case. Any social rate of time preference can be assumed while maintaining the Pareto principle.

While the discounting of life years does not affect efficiency, it is related to social attitudes to inequality. In our framework there is generally a social aversion to inequality in income because of declining willingness to pay for higher consumption by giving up life years as consumption rises.

If there is no discounting of future life years, however, society is indifferent as to the distribution of life years across people with similar (reference point) endowments. With discounting, society prefers a more equal distribution of life years within cohorts; an extra life year for someone with a short life span comes before, and is more highly valued, than an extra life year for someone with a high lifespan. We show that if life spans are normally distributed, the social rate of time preference is equivalent to a coefficient of social inequality aversion on life spans.

Suppose every individual has the reference allocation of traded and non traded goods. Again assuming a continuum of agents of mass one we can write social welfare as

$$\int \frac{1}{\delta} (1 - e^{-\delta l}) g(l) dl$$

where $g(l)$ is the distribution of life spans. We assume the distribution of life spans is normal with mean μ_l and standard deviation σ_l . Of course this must be regarded as an approximation since life spans cannot be negative. In addition, while the distribution of adult life spans is unimodal (though skewed) there is also a peak in mortality and life spans between birth and age one which is not captured in a normal distribution.

However, taking the normal distribution as an approximation we then have (using the moment generating function again) that social welfare is a strictly positive transformation of

$$u_l - \delta \frac{\sigma_l^2}{2}$$

Society prefers higher average life spans but dislikes variation in lifespan, the weight on inequality in the social welfare function increasing with the rate of social time preference. This analysis matches that of Edwards (2007) who argues that if people discount the future they will be risk-averse when offered gambles over life span. Our approach differs in that it reflects the effect of social time preference on society's attitudes to a distribution of life spans that are known with certainty. Note that in simple examples society's preferences over income inequality depend on individuals' preferences and their willingness to trade off income for life span, while society's preferences over inequality in lifespan depend on the social rate of time preference. In more complex cases involving endowments that are not at the reference point, both individual preferences and the social rate of time preference will affect social attitudes to inequality.

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 with a standard income level, all (discounted) life years gained are weighted equally. For people whose income or health is not at the reference point, life years gained are 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. Our results suggest that quality adjustment should be extended to all factors that affect the quality of life.

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 our new 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, cost-benefit and cost-effectiveness analysis are coherent. All measurements are now in quality-adjusted life years, and so we compare easily across sectors.

The real argument is of course over the two axioms. It may well be possible to devise other sets of axioms that have the same internal consistency properties. Even if our axioms are accepted we have left open the choice of the reference point and the rate of social discounting of future life years. If these change, all our internal consistency results remain, but our social preference ordering, and ranking of projects change.

This shows the importance of the reference point. Once we value lives equally at one reference point and set the social rate of time preference we completely determine the social preference ordering. Shifting the reference point at which lives are valued equally will usually produce a completely new, and different, social preference ordering. Insisting that life years are equally valuable under two different distributions of income usually leads to a violation of the Pareto principle. Future research should examine how different choice of reference point affects social preferences.

We consider the welfare of a single cohort, of a fixed size, under certainty. This sidesteps difficult issues associated with intergenerational distribution, endogenous population numbers, and uncertainty. With cohorts at different birth years the social rate of time preference affects intergenerational distribution. It is unclear if this should be the same as the rate of time preference we use for choices between members of the same cohort. Endogenous population

numbers raise the issue of how we value potential lives relative to the actual lives of those who have been born. Our approach appears to do little to resolve this question.

While we do not allow for uncertainty, given the Pareto principle, assuming that individual and social preferences obey the axioms of expected utility theory implies a utilitarian social welfare function (Harsanyi (1955), Fishburn (1984)). Our approach also implies a utilitarian social welfare function, which gives some hope that widening the scope of the analysis to incorporate uncertainty may not lead to fundamental problems. However, we leave these difficult 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 .

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.

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 $\sum_i \psi_i(p, z_i, m_i, l_i)$ over Γ and $\sum_i \phi_i(x_i, z_i, l_i)$ 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 $R = (p^r, m^r, z^r, \cdot)$.

Proof. We first address existence. Consider the social welfare function defined by

$\sum_i \psi_i(p, z_i, m_i, l_i)$ on Γ and $\sum_i \phi_i(x_i, z_i, l_i)$ on Ω . 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 lifespans 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 then 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. Our social welfare function generates a well behaved social preference ordering satisfying the two axioms. Let us denote this social preference ordering as \succeq while \succ and \sim denote the induced strict preference and indifference relations. If this ordering is not unique there exists a second, different, well behaved social preference satisfying axioms one and two. Let us denote this ordering as \succeq' while \succ' and \approx denote the induced strict preference and indifference relations.

If the two orderings were the same we would have

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

Since they are different we can find two allocations such that

$$(p, z_i, m_i, l_i) \succeq (p', z'_i, m'_i, l'_i) \text{ while } (p, z_i, m_i, l_i) \prec (p', z'_i, m'_i, l'_i)$$

or

$$(p, z_i, m_i, l_i) \prec (p', z'_i, m'_i, l'_i) \text{ while } (p, z_i, m_i, l_i) \succeq (p', z'_i, m'_i, l'_i)$$

We begin with the case where $(p, z_i, m_i, l_i) \succeq (p', z'_i, m'_i, l'_i)$ while $(p, z_i, m_i, l_i) \triangleleft (p', z'_i, m'_i, l'_i)$

By construction of the preface ordering \succeq we have

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

and it follows by axiom 3 that we must have

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

where, by construction

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

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

Converting each expression to an allocation, let

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

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

and

$$(x_R(l_i^*), z^r, l_i^*) = (x(p^r, m^r, z^r, \psi_R(p, z_i, m_i, l_i)), z^r, \psi_R(p, z_i, m_i, l_i)) \text{ where } l_i^* = \psi_R(p, z_i, m_i, l_i)$$

$$(x_R(l_i'^*), z^r, l_i'^*) = (x(p^r, m^r, z^r, \psi_R(p, z'_i, m'_i, l'_i)), z^r, \psi_R(p, z'_i, m'_i, l'_i)) \text{ where } l_i'^* = \psi_R(p, z'_i, m'_i, l'_i)$$

Hence we have, by conditions v and vi

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

$$(x_R(l_i^*), z^r, l_i^*) \succeq (x_R(l_i'^*), z^r, l_i'^*)$$

while for all i

$$(x_i, z_i, l_i) \sim_i (x_R(l_i^*), z^r, l_i^*)$$

$$(x'_i, z'_i, l'_i) \sim_i (x_R(l_i'^*), z^r, l_i'^*)$$

Hence we have one allocation preferred to another. However when we move both allocations to the reference point with different life spans, in such a way as to keep every agent indifferent, the social ranking is reversed. Now consider the strict convex combination allocation for

$$a(\lambda) = \lambda(x'_i, z'_i, l'_i) + (1 - \lambda)(x_R(l_i^*), z^r, l_i^*)$$

By A4 we have that $a(\lambda)$ is Pareto superior to both (x'_i, z'_i, l'_i) and $(x_R(l_i^*), z^r, l_i^*)$

Hence $a(\lambda) \succ (x'_i, z'_i, l'_i)$ and $a(\lambda) \succ (x_R(l_i^*), z^r, l_i^*)$ for $0 < \lambda < 1$.

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

$$a(0) = (x_R(l_i^*), z^r, l_i^*) \succeq (x'_i, z'_i, l'_i)$$

Similarly taking limits as $\lambda \rightarrow 1$

$$a(1) = (x'_i, z'_i, l'_i) \succeq (x_R(l_i^*), z^r, l_i^*)$$

It then follows that

$$(x_R(l_i^*), z^r, l_i^*) \approx (x'_i, z'_i, l'_i)$$

Similarly we have that

$$(x_i, z_i, l_i) \approx (x_R(l_i^*), z^r, l_i^*)$$

And hence

$$(x'_i, z'_i, l'_i) \approx (x_R(l_i^*), z^r, l_i^*) \leq (x_R(l_i^*), z^r, l_i^*) \approx (x_i, z_i, l_i)$$

by transitivity we have

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

a contradiction with

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

Now we turn to the case where $(p, z_i, m_i, l_i) \prec (p', z'_i, m'_i, l'_i)$ while $(p, z_i, m_i, l_i) \succeq (p', z'_i, m'_i, l'_i)$.

The argument is very similar to the previous case. By construction of the preface ordering \succeq we have

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

and it follows by axiom 3 that we must have

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

where by construction

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

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

Converting each expression to an allocation let

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

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

and

$$(x_R(l_i^*), z^r, l_i^*) = (x(p^r, m_i^r, z_i^r, \psi_R(p, z_i, m_i, l_i)), z_i^r, \psi_R(p, z_i, m_i, l_i)) \text{ where } l_i^* = \psi_R(p, z_i, m_i, l_i)$$

$$(x_R(l_i'^*), z^r, l_i'^*) = (x(p^r, m_i^r, z_i^r, \psi_R(p, z'_i, m'_i, l'_i)), z_i^r, \psi_R(p, z'_i, m'_i, l'_i)) \text{ where } l_i'^* = \psi_R(p, z'_i, m'_i, l'_i)$$

Hence we have by conditions v and vi

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

$$(x_R(l_i^*), z^r, l_i^*) \triangleleft (x_R(l_i'^*), z^r, l_i'^*)$$

while for all i

$$(x_i, z_i, l_i) \sim_i (x_R(l_i^*), z^r, l_i^*)$$

$$(x'_i, z'_i, l'_i) \sim_i (x_R(l_i'^*), z^r, l_i'^*)$$

Hence we have one allocation preferred to another. However when we move both allocations to the reference point with different life spans, in such a way as to keep every agent indifferent, the social ranking is reversed. Now consider the strict convex combination allocation for

$$a(\lambda) = \lambda(x'_i, z'_i, l'_i) + (1 - \lambda)(x_R(l_i^*), z^r, l_i^*)$$

By A4 we have that $a(\lambda)$ is Pareto superior to both (x'_i, z'_i, l'_i) and $(x_R(l_i^*), z^r, l_i^*)$

Hence $a(\lambda) \triangleright (x'_i, z'_i, l'_i)$ and $a(\lambda) \triangleright (x_R(l_i^*), z^r, l_i^*)$ for $0 < \lambda < 1$.

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

$$a(0) = (x_R(l_i^*), z^r, l_i^*) \succeq (x'_i, z'_i, l'_i)$$

Similarly taking limits as $\lambda \rightarrow 1$

$$a(1) = (x'_i, z'_i, l'_i) \succeq (x_R(l_i^*), z^r, l_i^*)$$

It then follows that

$$(x_R(l_i^*), z^r, l_i^*) \approx (x'_i, z'_i, l'_i)$$

Similarly we have that

$$(x_i, z_i, l_i) \approx (x_R(l_i^*), z^r, l_i^*)$$

And hence

$$(x'_i, z'_i, l'_i) \approx (x_R(l_i^*), z^r, l_i^*) \triangleleft (x_R(l_i^*), z^r, l_i^*) \approx (x_i, z_i, l_i)$$

Hence since \triangleright is well behaved we have

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

a contradiction with

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

It follows that the well behaved social preference ordering satisfying axioms 1 and 2 is unique.

The fact that the social preference ordering can be represented by the social welfare function $\sum_i \psi_i(p, z_i, m_i, l_i)$ on Γ and $\sum_i \phi_i(x_i, z_i, l_i)$ on Ω comes directly from the constructive proof of the existence of the social preference ordering given above.

References

- Arrow, K. (1950). "A Difficulty in the Concept of Social Welfare." *The Journal of Political Economy* 58: 328–346.
- Becker, G. S. and C. B. Mulligan (1997). "The Endogenous Determination of Time Preference." *The Quarterly Journal of Economics* 112(3): 729-758.
- Berlage, L. and R. Renard (1985). "The Discount Rate in Cost-Benefit Analysis and the Choice of a Numeraire." *Oxford Economic Papers* 37(4): 691-699.
- Blackorby, C., D. Donaldson, et al. (1984). "Social Choice with Interpersonal Utility Comparisons: A Diagrammatic Introduction." *International Economic Review* 25(2): 327-356.
- Boardway, R. W. (1974). "The Welfare Foundations of Cost-Benefit Analysis." *The Economic Journal* 84: 926-939.
- Broome, J. (1985). "The Economic Value of Life." *Economica* 52: 281-294.
- Broome, J. (2002). Measuring the burden of disease by aggregating wellbeing. *Summary Measures of Population Health: Concepts, Ethics, Measurement and Applications*. C. J. L. Murray, J. A. Salomon, M. C. D. and A. D. Lopez. Geneva, World Health Organization: 91–113.
- Chichilnisky, G. and G. Heal (1983). "Necessary and sufficient conditions for a resolution of the social choice paradox." *Journal of Economic Theory* 31(1): 68-87.
- Culyer, A. J. (1989). "The normative economics of health care finance and provision." *Oxford Review of Economic Policy* 5: 34–58.
- Daniels, N. (2008). *Just Health: Meeting Health Needs Fairly*. Cambridge, Cambridge University Press.
- Dolan, P. and R. Edlin (2002). "Is it really possible to build a bridge between cost-benefit analysis and cost-effectiveness analysis?" *Journal of Health Economics* 21(5): 827-843.
- Edwards, R. D. (2007). *The Cost of Uncertain Life Span*, Department of Economics, Queens College, City University of New York.
- Fishburn, P. (1984). "On Harsanyi's utilitarian cardinal welfare theorem." *Theory and Decision* 17: 21–28.
- Fuchs, V. R. and R. Zeckhauser (1987). "Valuing Health--A "Priceless" Commodity." *The American Economic Review* 77: 63-268.
- Hammond, P. J. (1994). Money Metric Measures of Individual and Social Welfare Allowing for Environmental Externalities. *Models and Measurement of Welfare and Inequality*. W. Eichhorn. Berlin and Heidelberg, Springer-Verlag: 694–724.
- Harsanyi, J. C. (1955). "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility." *The Journal of Political Economy* 63(4): 309-321.

- Hasmanand, A. and L. P. Østerdal (2004). "Equal Value of Life and the Pareto Principle." *Economics and Philosophy* 20: 19-33.
- Jittorntrum, K. (1978). "An implicit function theorem." *Journal of Optimization Theory and Applications* 25(4): 575-578.
- Johannesson, M. (1995). "The relationship between cost-effectiveness analysis and cost-benefit analysis." *Social Science and Medicine* 41(4): 483-9.
- Kaplow, L. and S. Shavell (2001). "Any Non-Welfarist Method of Policy Assessment Violates the Pareto Principle." *Journal of Political Economy* 109: 281-286.
- Rawls, J. (1971). *A Theory of Justice*. Cambridge, Harvard University Press.
- Sen, A. K. (1977). "On Weights and Measures: Informational Constraints in Social Welfare Analysis." *Econometrica* 45(7): 1539-1572.
- Sen, A. K. (1985). *Commodities and Capabilities*. Oxford, Oxford University Press.
- Sen, A. K. (1999). *Development As Freedom*. New York, Knopf.
- Somanathan, E. (2006). "Valuing lives equally: Distributional weights for welfare analysis." *Economics Letters* 90(1): 122-125.
- Weinstein, M. C. and W. I. Manning (1997). "Theoretical issues in cost-effectiveness analysis." *Journal of Health Economics* 16: 121-128.
- Weinstein, M. C. and W. B. Stasson (1977). "Foundations of cost-effectiveness analysis for health and medical practices." *New England Journal of Medicine* 296(13): 716-21.
- Weymark, J. A. (1985). "Money-Metric Utility Functions." *International Economic Review* 26(1): 219-232.