ELSEVIER

Journal of Health Economics 18 (1999) 681–708

JOURNAL OF HEALTH ECONOMICS

www.elsevier.nl/locate/econbase

Life-cycle preferences over consumption and health: when is cost-effectiveness analysis equivalent to cost–benefit analysis?

Han Bleichrodt [a,*], John Quiggin [b]

[a] *iMTA, Erasmus University, PO Box 1738, 3000 DR Rotterdam, Netherlands*
[b] *Department of Economics, James Cook University, Townsville, Australia*

Received 8 October 1998; received in revised form 25 March 1999

Abstract

This paper studies life-cycle preferences over consumption and health status. We show that cost-effectiveness analysis is consistent with cost–benefit analysis if the lifetime utility function is additive over time, multiplicative in the utility of consumption and the utility of health status, and if the utility of consumption is constant over time. We derive the conditions under which the lifetime utility function takes this form, both under expected utility theory and under rank-dependent utility theory, which is currently the most important nonexpected utility theory. If cost-effectiveness analysis is consistent with cost–benefit analysis, it is possible to derive tractable expressions for the willingness to pay for quality-adjusted life-years (QALYs). The willingness to pay for QALYs depends on wealth, remaining life expectancy, health status, and the possibilities for intertemporal substitution of consumption. © 1999 Elsevier Science B.V. All rights reserved.

JEL classification: I10; D61

Keywords: Economic evaluation; Value of life; Health; Utility theory; Willingness to pay

* Corresponding author. Tel.: +31-10-408-8566; fax: +31-10-408-9094; E-mail: bleichrodt@bmg.eur.nl

PII: S0167-6296(99)00014-4

1. Introduction

The (non-)equivalence of cost-effectiveness analysis and cost–benefit analysis has received much attention in the literature on the economic evaluation of health care (Phelps and Mushlin, 1991; Johannesson, 1995a; Donaldson, 1998). This attention emanates from a concern about the theoretical properties of cost-effectiveness analysis. Notwithstanding the fact that it is currently the most common tool in the economic evaluation of health care, cost-effectiveness analysis, unlike cost–benefit analysis, has no foundation in economic welfare theory.

The most widely used outcome measure in cost-effectiveness analysis are quality-adjusted life-years (QALYs). There is a well-established literature describing the conditions under which QALY-based decision making is consistent with preferences over lifetime health profiles (Pliskin et al., 1980; Bleichrodt, 1995; Bleichrodt and Quiggin, 1997; Bleichrodt et al., 1997; Miyamoto et al., 1998; Ried, 1998). Much less is known about the consistency between QALY-based decision making [1] and individual preferences when lifetime utility depends not only on health status, but also on consumption. In this paper, we derive a set of conditions that is both necessary and sufficient for the consistency of QALY-based decision making with life-cycle preferences over consumption and health status. Cost–benefit analysis is always consistent with life-cycle preferences over consumption and health status. This follows because cost–benefit analysis imposes no assumptions on the lifetime utility function over consumption and health status. Consequently, to derive the conditions under which cost-effectiveness analysis is consistent with life-cycle preferences over consumption and health simultaneously answers the question under which conditions cost-effectiveness analysis is equivalent to cost–benefit analysis.

Two recent papers by Garber and Phelps (1997) and Meltzer (1997) also analyzed the allocation of health resources within a life-cycle framework, where both health status and consumption are arguments of the utility function. These papers focus on the question when the cost side of cost-effectiveness analysis is consistent with life-cycle preferences over consumption and health status. Their main concern is whether future consumption should be regarded as a cost of life-saving medical interventions. Although Garber and Phelps assume that the health argument of the utility function may be represented in terms of QALYs, neither they nor Meltzer consider the conditions under which an individual concerned about both health status and general consumption would seek to maximise QALYs. This paper complements the papers by Garber and Phelps and by Meltzer by focusing on the question when the outcome side of cost-effective-

[1] Throughout the text, we interchangeably use the terms cost-effectiveness analysis and QALY-based decision making. Our central result, derived in Section 2, can be applied to other outcome measures as well. We briefly comment on other outcome measures in Section 5.

ness analysis, i.e., QALY maximisation, is consistent with life-cycle preferences. Throughout the paper, we assume that costs are measured in a consistent way and we restrict attention to the valuation of outcomes.

The paper is structured as follows. In Section 2, we derive the central result of this paper, that QALY maximization is consistent with life-cycle preferences only if utility is additive over time, if the one-period utility functions can be multiplicatively decomposed in a utility function over consumption and a utility function over health status, and if the utility of consumption is constant over time. Section 3 provides an axiomatic analysis of the conditions under which the utility function is additive over time and the one-period utility functions are multiplicative in the utility of consumption and the utility of health status. An axiomatic analysis of this model is provided both under expected utility theory and under rank-dependent utility theory, currently the most important nonexpected utility model. Section 3 is technical, but we have tried to increase its accessibility by displaying each step in the derivation separately, by making ample use of examples to illustrate the conditions that are successively introduced, and by deferring all technical assumptions and proofs to Appendix A. In Section 4, we derive constancy of consumption. We show that under the utility function of Section 3, consumption is constant if an individual's rate of time preference is equal to the interest rate. Section 4 further examines the implications for the valuation of longevity and the willingness to pay for a QALY gained in case cost-effectiveness analysis is consistent with cost–benefit analysis, i.e., in case the model of Section 3 holds with constant consumption. Johannesson (1995a) has argued that cost-effectiveness can only be a useful tool if information is available on the willingness to pay for a QALY gained. To date, no such information exists. We derive expressions for the willingness to pay for a QALY gained both when life duration is certain and when life duration is risky. Section 5 concludes.

2. The central result

Our aim is to show under which conditions QALY-based decision making is consistent with life-cycle preferences over consumption and health status. The general idea of our argument follows from the uniqueness properties of the additive utility function and a special property of health. Let $(h_1,\dots,h_T)$ denote a sequence of health states, where h_t stands for health status in time period t. Similarly, $((c_1,h_1),\dots,(c_T,h_T))$ denotes a sequence of pairs of consumption level and health status, where (c_t,h_t) denotes consumption level and health status in period t.

The first step in our derivation is to realize that according to the QALY model the utility of the sequence $(h_1,\dots,h_T)$ is equal to $\sum_{t=1}^{T} q(h_t)$, where $q(h_t)$ is the quality weight or utility of health status in period t. Two things should be noted about the QALY utility function. First, it is additive and, second, the one-period utility functions are identical. It is well known from the literature on additive representations (e.g., Wakker, 1989) that if the one-period utility functions are additive then they are unique up to positive linear transformations. That is, if $q(h_t)$ and $q'(h_t)$ are both additive one-period utility functions then

$$q(h_t) = \alpha + \beta q'(h_t) \tag{1}$$

where α can be any real number and β is a strictly positive real number. Note that because the $q(h_t)$ are unique up to positive linear transformations, we are free to choose their scale and location. Throughout, we set the utility of death equal to zero.

Suppose now, that an individual is not only concerned about the sequence of health states, but also cares about consumption. That is, we now study preferences over the sequence $((c_1,h_1),\dots,(c_T,h_T))$. We assume that an individual always prefers more consumption to less as long as health status is better than death. Consequently, utility is strictly increasing in consumption. If the utility function over $((c_1,h_1),\dots,(c_T,h_T))$ is consistent with the QALY utility function then we know by the argument of the preceding paragraph that it must be a positive linear transformation thereof. Hence, it must have the following form:

$$U((c_1,h_1),\dots,(c_T,h_T)) = \sum_{t=1}^{T} (v(c)q(h_t) + w_t(c)) \tag{2}$$

where v is a utility function over consumption that must be strictly positive ($v(c)$ corresponds to β in Eq. (1)) and the w_t are period-specific utility functions over consumptions that are real-valued (the w_t correspond to α in Eq. (1)).

We next employ a special characteristic of the life-cycle model for consumption and health, namely, the fact that an individual will derive no more utility from consumption once he has died. Bequest motives may lead individuals to consume amounts less than their full wealth. However, it is reasonable to suppose, as is commonly done in economic analyses of bequest motives, that any utility of bequests is additively separable from $U((c_1,h_1),\dots,(c_T,h_T))$. Therefore, the existence and form of the utility of bequests has no implications for the form of $U((c_1,h_1),\dots,(c_T,h_T))$. By our scaling convention, q(death) is equal to zero and it follows from Eq. (2) that the $w_t(c)$ must be equal to zero to capture the property that the individual derives no more utility from consumption after he has died. Consequently, life-cycle preferences over consumption and health status are consistent with QALY maximization only if:

$$U((c_1,h_1),\dots,(c_T,h_T)) = \sum_{t=1}^{T} v(c)q(h_t) \tag{3}$$

Eq. (3) has three interesting properties. First, the utility of consumption is constant over the individual's life-cycle. Because we assume that utility is strictly increasing in consumption, it follows that consumption itself is constant over the individual's life-cycle. Second, the utility of consumption is positive. If there exists a subsistence level of consumption beyond which additional life-years are negatively valued, then consumption has to be above this subsistence level. This is not a major restriction. Rosen (1988) has shown that if consumption falls below the subsistence level, an individual will convexify his preferences by randomizing between death and survival at a consumption level which is higher than the subsistence level. Hence, Eq. (3) captures all cases of economic interest.

Most importantly, in Eq. (3), the utility of health status is multiplied by the utility of consumption. Consequently, a given gain in quality of life will be more appreciated at higher levels of consumption. This implies that in the allocation of health care resources, larger welfare gains can be obtained by devoting resources to those individuals who have a high level of general consumption. A comparable result was derived by Pratt and Zeckhauser (1996) for the valuation of risk reductions. This result is ethically troubling, since an important motive for the use of cost-effectiveness analysis is the desire to avoid the adverse distributional implications commonly seen to arise from applications of cost–benefit analysis. However, the need for a multiplicative utility structure shows that, if cost-effectiveness analysis is to be placed within the realm of econonomic welfare theory, such implications cannot be escaped.

In Section 3, we present an axiomatic analysis of the model

$$U\big((c_1,h_1),\dots,(c_T,h_T)\big)=\sum_{t=1}^{T}\upsilon(c_t)\,q(h_t) \tag{4}$$

This model is slightly more general than Eq. (3) because consumption is allowed to vary over time. If consumption varies over the life-cycle, an optimal allocation of health resources will assign more weight to health status in high consumption years. Conversely, other things equal, an individual with preferences that can be described by Eq. (4) will choose higher levels of consumption in those years in which his health status is high. The selection of a constant consumption level has to follow from an optimization problem which requires specification of the budget constraint. We present such an optimization problem in Section 4.

3. Characterizations of Eq. (4)

We focus on preferences under risk. Preferences under certainty follow by restricting attention to *riskless lotteries*: lotteries that yield an outcome with probability one, i.e., with certainty. Life-cycle preferences are first analyzed under expected utility theory, which is still the dominant theory in health utility measurement. However, it is by now widely accepted that people do not behave according to expected utility theory (Camerer, 1995). Several nonexpected utility theories have been developed among which rank-dependent utility theory (Quiggin, 1982; Yaari, 1987; Quiggin and Wakker, 1994) is currently the most influential theory. We also derive Eq. (4) under rank-dependent utility theory.

We start with the case where utility is not discounted. The modeling of discounting is similar under expected utility and under rank-dependent utility and we leave it to the end of the axiomatic analysis. We characterize two discounting models: constant rate discounting and a general discounting model which is consistent with the most important alternatives for constant rate discounted utility that have been proposed in the literature on intertemporal preferences (Loewenstein and Prelec, 1992; Harvey, 1994).

3.1. Expected utility theory

3.1.1. Some notation

In this subsection, we introduce the main concepts used in our axiomatic analysis. To improve accessibility, we have moved all technical assumptions regarding these concepts to Appendix A.

We assume that there are T points in time. We express this by saying that there is a set $S=\{1,\dots,T\}$ of time points. The set of all outcomes, i.e., sequences $[(c_1,h_1),\dots,(c_T,h_T)]$ is denoted by X. For ease of notation, we sometimes refer to pairs (c_t,h_t) as y_t. The consumption levels c_t are in each period elements of a set C^+, which consists of all attainable consumption levels. The plus sign serves as a reminder that we only consider consumption levels that are above the subsistence level. The health states h_t are elements of a set H, which consists of all attainable health states. We assume that H consists only of those health states that are at least as good as death. Under expected utility, the axiomatic analysis also holds if health states worse than death are included. Under rank-dependent utility, however, the generalization to health states worse than death is more arduous.

We assume that the sets C^+ and H are equal in each time period. This assumption is made for convenience and does not restrict our analysis. The analysis can straightforwardly be extended to cover cases where C^+ and H vary over time.

3.1.2. Preference conditions and representation theorem

Let P be the set of all *lotteries* over X. A typical element of P is $[p_1,x^{(1)};\dots;p_m,x^{(m)}]$ which gives outcome $x^{(1)}=[(c_1^1,h_1^1),\dots,(c_T^1,h_T^1)]$ with probability p_1, outcome $x^{(2)}=[(c_1^2,h_1^2),\dots,(c_T^2,h_T^2)]$ with probability p_2, etc. and m is any natural number. In medical decision making, lotteries can be interpreted as treatments, the outcomes of which are risky. We assume that the set P contains all riskless lotteries, which are of the type $[1,x]$.

An individual is assumed to have preferences over P. We denote the individual's preference relation over P by $\succcurlyeq$, which stands for ''at least as preferred as''. We write $x \succ y$ if it is true that $x \succcurlyeq y$ but not $y \succcurlyeq x$. That is, $x \succ y$ means that x is *strictly preferred to y*. We write $x \sim y$ if both $x \succcurlyeq y$ and $y \succcurlyeq x$ are true. That is, $x \sim y$ means that x is *indifferent to y*.

A preference relation over outcomes can be derived from the preference relation over lotteries by restricting attention to riskless lotteries. We assume that the preference relation over outcomes satisfies *monotonicity with respect to consumption*: higher consumption levels are strictly preferred to lower levels as long as health status is unequal to death. We also assume a sort of *monotonicity with respect to health status*: if health state h^1 is preferred to health state h^2 for a given level of consumption $c \in C^+$, then h^1 is preferred to h^2 for all consumption levels.

A function V is said to *represent* the preference relation $\succcurlyeq$, if for any two lotteries P_1 and P_2 that belong to the set P, the individual considers P_1 at least as preferred as P_2 if and only if the value of V at P_1 is at least as great as the value of V at P_2. This definition is expressed mathematically as: $P_1 \succcurlyeq P_2$ if and only if $V(P_1) \geq V(P_2)$.

We assume that the expected utility axioms (von Neumann and Morgenstern, 1953; Jensen, 1967) hold. Then, the preference relation $\succcurlyeq$ is represented by the following function:

$$\mathrm{EU}\left(p_1, x^{(1)}; \ldots; p_m, x^{(m)}\right) = \sum_{i=1}^{m} p_i U\left(x^{(i)}\right) \tag{5}$$

where U is a utility function over outcomes. The utility function in Eq. (5) is still entirely general. Our aim is to give a set of conditions that ensure that the utility function can be written as $\sum_{t=1}^{T} v(c_t) q(h_t)$. We derive this representation by imposing four conditions.

The first condition, marginality, ensures that U is a weighted additive function $U((c_1,h_1),\ldots,(c_T,h_T)) = \sum_{t=1}^{T} U_t(c_t,h_t)$ of one-period utility functions U_t.

Definition 1: The preference relation over P satisfies *marginality* if for all P_1, $P_2 \in P$ with equal marginal probability distributions over $y_1, \ldots, y_T$: $P_1 \sim P_2$.

An example may clarify the effect of marginality. Consider a simple model in which there are only two time periods. Let y' be the outcome (20,000, good health), i.e., a consumption level of US\$20,000 and good health, and let y'' be the outcome (5000, bad health). Consider the two lotteries P_1 and P_2 displayed in Table 1. P_1 yields a probability $1/2$ of obtaining y'' in both periods and a probability $1/2$ of obtaining y' in both periods. Marginality implies that the individual is indifferent between P_1 and P_2. This follows, because both P_1 and P_2 yield in each time period a probability $1/2$ of y'' and a probability $1/2$ of y'. Hence, P_1 and P_2 have equal marginal probability distributions over y_1 and y_2.

The above example shows that marginality excludes all complementarity between time periods. It might well be that the individual prefers P_1 to P_2 because he dislikes variation in his consumption and health status levels. Such an

Table 1
An example of marginality

Lottery	Outcome with probability $1/2$	Outcome with probability $1/2$
P_1	(y'', y'')	(y', y')
P_2	(y'', y')	(y', y'')

aversion to variation is not permitted by marginality. Alternatively, the individual might prefer P_2 to P_1 because he wants to avoid the possibility that he obtains the bad outcome y' in both periods. Such a preference for ‘‘catastrophe avoidance’’ is also excluded by marginality.

The next step in our derivation is to make the one-period utility functions identical. That is, we have to impose a condition which allows the representation of preferences by the utility function $U((c_1,h_1),\dots,(c_T,h_T)) = \sum_{t=1}^{T} U(c_t,h_t)$. Before we introduce the condition that achieves this end, symmetry, we define a new concept. A *permutation function* $\pi(t)$ is a function that specifies a rearrangement of the time periods. For example, if $\pi(t) = s$ then the point t is moved to point s in the rearrangement of the time periods.

Definition 2: The preference relation over P satisfies *symmetry* if for all $t \in S$, for all $y \in Y$ with h unequal to death, and for all permutation functions π it is true that $(y_1,\dots,y_T) \sim (y_{\pi(1)},\dots,y_{\pi(T)})$.

Let y' and y'' be as in the previous example. Then, by symmetry, the individual is indifferent between the outcome x' which yields y' in the first period and y'' in the second period and the outcome x'' which yields y'' in the first period and y' in the second. This follows because x'' can be obtained from x' by applying the permutation $\pi(1) = 2$ and $\pi(2) = 1$.

Symmetry renders irrelevant the point in time at which a particular outcome occurs. If symmetry holds, the individual’s preferences are unaffected when a permutation of the time periods turns a decreasing sequence, i.e., a sequence in which outcomes become worse over time, like x'' in the above example, into an increasing sequence, i.e., a sequence in which outcomes improve over time, like x' in the above example. Symmetry is at odds with the common assumption that people have positive time preference, i.e., prefer good outcomes to occur sooner rather than later. If people have positive time preference then they always prefer decreasing sequences to increasing sequences. Later in this section, we show how symmetry can be replaced by conditions that are compatible with positive time preference.

Marginality and symmetry ensure that the utility function over sequences of consumption and health status is additively decomposable over time and that the one-period utility functions are equal for each time period. The final two conditions, standard gamble invariance and the zero condition, establish that the one-period utility function can be multiplicatively decomposed, i.e., it can be written as $U(c_t,h_t) = v(c_t)q(h_t)$. Standard gamble invariance implies that $U(c_t,h_t)$ can be written as $v(c_t)q(h_t) + w_t(c_t)$. The zero condition then allows to set $w_t(c_t) = 0$ for all t and for all consumption levels.

Definition 3: The preference relation over P satisfies *standard gamble invariance* if for all c, $c' \in C^+$: $(c,h) \succcurlyeq [p,(c,h');\ 1-p,\ (c,h'')]$ if and only if $(c',h) \succcurlyeq [p,(c',h');\ 1-p,\ (c',h'')]$ with all lotteries elements of P.

To illustrate, let h, h', and h'' be three health states such that h' is better than h which is better than h''. For example, h can be a mild form of asthma, h' no asthma, and h'' a severe form of asthma. Suppose that annual consumption is held fixed at a given level, say US$20,000, and that at this consumption level the individual considers having a mild form of asthma for certain at least as good as undergoing a risky treatment which yields a probability p of no asthma and a probability $1-p$ of a severe form of asthma. Then standard gamble invariance says that the individual should still consider a mild form of asthma for certain at least as good as undergoing the risky treatment if consumption is held fixed at another level, say US$5,000. This example illustrates the effect of standard gamble invariance. Standard gamble invariance enables consideration of preferences over health states irrespective of the level at which consumption is held fixed.

Standard gamble invariance is typically invoked in health utility measurement. In assessing the utility of a health state by the standard gamble (Torrance, 1986), it is commonly assumed that life-years can be held fixed and that the value at which life-years are held fixed does not affect preferences. This idea is similar to standard gamble invariance as we use it here. The only difference is that we propose to hold consumption fixed instead of life-years in the assessment of preferences over health states.

Because standard gamble invariance allows consideration of preferences over health status with consumption held constant, it also allows the definition of a separate utility function over health status which does not depend on consumption. In Appendix A, we prove that imposing standard gamble invariance on top of marginality and symmetry implies that the one-period utility functions $U(c_t,h_t)$ can be written as $w_t(c_t)+\upsilon(c_t)q(h_t)$ with $w_t(c_t)$ real and $\upsilon(c_t)$ positive. This model differs from Eq. (4) by the terms $w_t(c_t)$. To complete our characterization, we impose a condition that implies that $w_t(c_t)=0$ for all t and for all c_t. The condition that has this effect, the zero condition, is a condition that is naturally satisfied in the medical context. In words, the zero condition says that a person derives no more utility from consumption once he has died. Formally, the zero condition is defined as:

Definition 4: The preference relation satisfies the *zero condition* if for all consumption levels c, $c' \in C^+$: $(c, \text{death}) \sim (c', \text{death})$.

Theorem 1 summarizes the derivation of Eq. (4). A formal proof of Theorem 1 is given in Appendix A.

Theorem 1: Under expected utility theory, the following two statements are equivalent: (i) Life-cycle preferences are consistent with the maximization of $U[(c_1 h_1),\dots,(c_T,h_T)]=\sum_{t=1}^{T}\upsilon(c_t)q(h_t)$. (ii) The preference relation $\succcurlyeq$ over P satisfies marginality, symmetry, standard gamble invariance, and the zero condition.

3.1.3. An assessment of the conditions of Theorem 1

Before moving on to the characterization of Eq. (4) under rank-dependent utility theory, we briefly comment on the descriptive validity of the conditions described in Section 3.1.2. The conditions are new in the context of life-cycle preferences over consumption and health. Therefore, they have not been subjected to empirical tests and we can only speculate about their empirical contents.

We believe that marginality and symmetry are the most restrictive conditions of the characterization. The zero condition is unobjectionable in the medical context. Standard gamble invariance is implied by the stronger condition of utility independence. Utility independence is commonly assumed both in medical decision analysis (Torrance et al., 1982; Torrance et al., 1995; Torrance et al., 1996) and in general decision analysis (Keeney and Raiffa, 1976) and it is believed to be a reasonable condition in most decision contexts. Moreover, there is some evidence that preferences over health status and duration satisfy utility independence (Miyamoto and Eraker, 1988; Bleichrodt and Johannesson, 1997).

Marginality excludes all complementarity between time periods. This is a strong restriction. In the example presented in Section 3.1.2, marginality implies that people are indifferent with respect to variations in consumption and in health status. However, empirical evidence has shown that people have a tendency to overweigh their status quo or endowment and are averse to changes therein (Samuelson and Zeckhauser, 1988: Kahneman et al., 1990). Such a tendency is incompatible with marginality. In Section 3.2, we show that under rank-dependent utility theory marginality is no longer imposed and is replaced by an alternative condition. This may enhance the descriptive validity of Eq. (4).

Finally, symmetry excludes positive time preference. Positive time preference is commonly assumed in economic analyses and is confirmed in empirical studies. The existence of positive time preference calls for the replacement of symmetry. This is the topic of Section 3.3.

3.2. Rank-dependent utility theory

3.2.1. Some notation

The main difference between rank-dependent utility theory and expected utility theory is that rank-dependent utility theory does not assume linearity of the utility of a lottery in probability. Under rank-dependent utility theory, preferences over lotteries are represented by the functional

$$\mathrm{RDU}\left(p_1, x^{(1)}; \dots; p_m, x^{(m)}\right) = \sum_{i=1}^{m} \pi_i U\left(x^{(i)}\right) \tag{6}$$

where the π_i are decision weights that depend on, but are in general not equal to the probabilities. If the decision weights are equal to the probabilities, i.e., $\pi_i = p_i$ for all i, then rank-dependent utility theory is equal to expected utility theory. That is, rank-dependent utility theory includes expected utility theory as a special case.

In comparison with expected utility theory, there is one change in the mathematical framework. Remember that under expected utility we studied preferences over the set of lotteries P. Under rank-dependent utility theory, we consider preferences over a subset of P, the set of rank-ordered lotteries. A lottery $(p_1, x^{(1)}; \ldots; p_m, x^{(m)})$ is said to be a *rank-ordered lottery* if its outcomes are ranked in decreasing order of preference, i.e., $x^{(1)} \succcurlyeq x^{(2)} \succcurlyeq \ldots \succcurlyeq x^{(m)}$. For example, the lottery [p, full health; $1-p$, death] is rank-ordered, because full health is a better health state than death. However, the lottery [p, death; $1-p$, full health] is not rank-ordered. We denote the set of rank-ordered lotteries by $P\downarrow$.

The set of rank ordered lotteries contains all riskless lotteries. This follows because each outcome is at least as preferred as itself: $x \succcurlyeq x$ for all $x \in X$. Consequently, we can define preferences over outcomes by restricting attention to the subset of riskless lotteries.

We introduce one new notation. Let A be a subset $\{s, \ldots, t\}$ of the set of time points $S = \{1, \ldots, T\}$. Obviously, $1 \leq s \leq t \leq T$. By $a_A x$, we denote the outcome $(x_1, \ldots, x_{s-1}, a_s, \ldots, a_t, x_{t+1}, \ldots, x_T)$, i.e., the outcome that is obtained when the elements $x_s, \ldots, x_t$ of the sequence $x = (x_1, \ldots, x_T)$ are replaced by the corresponding elements $a_s, \ldots, a_t$ of the sequence $(a_1, \ldots, a_T)$. For example, let T be 40, let S be $\{25, \ldots, 30\}$, let x be the constant sequence yielding a consumption level of US\$20,000 and good health in each period and let a be the constant sequence yielding a consumption level of US\$5000 and bad health in each period. Then, $a_A x$ is the outcome yielding a consumption level of US\$20,000 and good health from periods 1 to 24, a consumption level of US\$5000 and bad health from periods 25 to 30, and a consumption level of US\$20,000 and good health from periods 31 to 40.

In case the set A consistst of just one point in time, say point t, then we write $a_t x$ instead of $a_A x$.

3.2.2. Preference conditions and representation theorem

We use the same steps as in Section 3.1.2 to characterize the representation $\sum_{t=1}^{T} v(c_t) q(h_t)$. Hence, the first step is to find a condition that allows utility to be written as the additive sum of the one-period utility functions, i.e., $U((c_1, h_1), \ldots, (c_T, h_T)) = \sum_{t=1}^{T} U_t(c_t, h_t)$. Under expected utility theory, marginality served this purpose. Marginality is no longer available under rank-dependent utility theory, because it requires that the evaluation function for lotteries is linear in probabily. As we have explained, this is not necessarily true under rank-dependent utility theory. Therefore, marginality cannot be invoked and we have to find an alternative condition which achieves the additive decomposition under rank-dependent utility theory. The following condition, generalized utility independence, has this effect.

Definition 5: The preference relation $\succcurlyeq$ on $P\downarrow$ satisfies *generalized utility independence* if for all subsets A of S: $[p, a_A z; 1-p, b_A y] \succcurlyeq [p, c_A z; 1-p, d_A y]$

if and only if $[p,a_A w;\ 1-p,b_A v] \succcurlyeq [p,c_A w;\ 1-p,d_A v]$ with all lotteries elements of $P\downarrow$.

An example may clarify the effect of generalized utility independence. For ease of notation, we assume that health status is held constant, say at full health, and we will suppress it from our notation. Let there be two time periods and let $A=\{1\}$. Consider the four lotteries displayed in Table 2. Here, P_1 is the lottery yielding US\$20,000 in both periods with probability p and US\$15,000 in both periods with probability $1-p$. It is easily verified that P_1 corresponds to $[p,a_A z;\ 1-p,b_A y]$, P_2 to $[p,c_A z;\ 1-p,d_A y]$, P_3 to $[p,a_A w;\ 1-p,b_A v]$, and P_4 to $[p,c_A w;\ 1-p,d_A v]$ if $a=z=$ US\$20,000, $c=$ US\$30,000, $b=y=w=$ US\$15,000, and $d=v=$ US\$10,000. Hence, if P_1 is at least as good as P_2, for example, because the individual dislikes variation in consumption over time, then generalized utility independence implies that P_3 should also be at least as preferred as P_4.

Observe that the probability distribution over what happens in the second period is identical for P_1 and P_2 (a probability of $1/2$ of 20,000 and a probability of $1/2$ of 15,000) and also for P_3 and P_4 (a probability of $1/2$ of 15,000 and a probability of $1/2$ of 10,000). Hence, generalized utility independence implies that if two lotteries have identical probability distributions over the outcomes in a given period, then the individual will ignore the outcomes occurring in this period. Individual preferences are only affected by periods in which lotteries have different probability distributions over the outcomes.

Generalized utility independence is a strengthening of utility independence, which says that the individual will ignore periods in which outcomes are common and certain. That is, utility independence says that the individual will ignore those periods for which riskless lotteries coincide. Utility independence can be derived from generalized utility independence by setting $z=y$ and $v=w$ in Definitions. Generalized utility independence extends utility independence by saying that preferences are unaffected by those periods for which lotteries in general, i.e., not only riskless lotteries, coincide.

As remarked before, it is widely believed that preferences approximately satisfy utility independence both in medical decisions and in other decision contexts. Because generalized utility independence extends utility independence in a natural

Table 2
An example of generalized utility independence

Lottery	Outcome with probability p	Outcome with probability $1-p$
P_1	(20,000, 20,000)	(15,000, 15,000)
P_2	(30,000, 20,000)	(10,000, 15,000)
P_3	(20,000, 15,000)	(15,000, 10,000)
P_4	(30,000, 15,000)	(10,000, 10,000)

way, it may be expected that generalized utility independence describes life-cycle preferences over consumption and health reasonably well.

The other three conditions used under expected utility theory in characterizing Eq. (4) can still be used under rank-dependent utility theory. Theorem 2 summarizes the derivation. A formal proof of Theorem 2 is given in Appendix A.

Theorem 2: Under rank-dependent utility theory, the following two statements are equivalent: (i) Life-cycle preferences are consistent with the maximization of $U[(c_1,h_1),\ldots,(c_T,h_T)] = \sum_{t=1}^{T} v(c_t)q(h_t)$. (ii) The preference relation $\succcurlyeq$ on $P\downarrow$ satisfies generalized utility independence. Further $\succcurlyeq$ satisfies symmetry, standard gamble invariance, and the zero-condition.

3.3. Discounting

3.3.1. The general discounting model

The characterizations presented in Theorems 1 and 2 imply that people give equal weights to different time periods, i.e., they are *timing neutral*. In economic evaluations of health care, it is more common to assume that people have positive time preference and discount future time periods. Empirical evidence is also supportive of the existence of positive time preference for health (Olsen, 1993a; Cairns, 1994; Chapman, 1996).

In this section, we characterize two models that incorporate time preference. In one model, we impose no restrictions on the discount weights. Therefore, this model is consistent with most discounting models that have been proposed in the literature on intertemporal preferences (Loewenstein and Prelec, 1992; Harvey, 1994). The other model is the constant rate discounting model, which is commonly applied in cost-effectiveness analysis and also underlies the analyses by Garber and Phelps (1997) and Meltzer (1997).

Timing neutrality is a consequence of imposing symmetry. Therefore, symmetry has to be dropped to allow for time preference. However, symmetry also served to select identical one-period utility functions in Theorems 1 and 2. If symmetry is dropped, the one-period utility functions are no longer necessarily identical. Hence, another condition has to be imposed which ensures that the one-period utility functions can be chosen identical. The following condition, trade-off consistency, serves this end.

Definition 6: The preference relation $\succcurlyeq$ on X satisfies *trade-off consistency* if for all $s, t \in S$: if $(a_s x \preccurlyeq b_s y)$ and $(c_s x \succcurlyeq d_s y)$ and $(a_t v \succcurlyeq b_t w)$ then $(c_t v \succcurlyeq d_t w)$, with all outcomes elements of X.

Trade-off consistency can be explained in terms of strength of preference. Suppose that a is strictly preferred to b, and that c is strictly preferred to d. Then, the preference $(a_s x \preccurlyeq b_s y)$ implies that to obtain the strictly preferred outcome a

instead of b in period s is not sufficient to outweigh getting x instead of y in all periods other than s. The preference $(c_s x \succcurlyeq d_s y)$, however, implies that to obtain the strictly preferred outcome c instead of d in period s is sufficient to outweigh getting x instead of y in the other periods. These two preferences, therefore, indicate that in period s the strength of preference of c over d must be at least as great as the strength of preference of a over b. Trade-off consistency claims that if the strength of preference of c over d is at least as great as the strength of preference of a over b in period s, then there does not exist another period t in which the strength of preference of c over d is smaller than the strength of preference of a over b. That is, if we observe that obtaining a instead of b in period t is sufficient to outweigh getting v instead of w in all other periods then it must be that obtaining c instead of d is also sufficient.

Trade-off consistency thus ensures that utility differences are ordered similarly in different periods. This implies that the one-period utility functions are cardinally equivalent and can be chosen identical. Trade-off consistency does not imply that each period gets the same weight. Under expected utility theory, trade-off consistency implies in combination with marginality that the utility function can be represented by $\sum_{t=1}^{T} \lambda_t U(c_t, h_t)$. Under rank-dependent utility theory, this representation follows from trade-off consistency and generalized utility independence. Additionally imposing standard gamble variance and the zero condition gives the general discounting model $U[(c_1,h_1),\ldots,(c_T,h_T)] = \sum_{t=1}^{T} \lambda_t v(c_t) q(h_t)$.

3.3.2. The constant rate discounted utility model

The constant rate discounted utility model, $U[(c_1,h_1),\ldots,(c_T,h_T)] = \sum_{t=1}^{T} \beta^{t-1} v(c_t) q(h_t)$ can be obtained by imposing one additional condition, stationarity, on the general multiplicative discounting model of Section 3.3.1. It is easily verified that in the constant rate discounted utility model, the ratio between the weights assigned to utility in period t and to utility in period s is equal to β^{t-s}. The ratio between the weights assigned to utility in period $t+e$ and to utility in period $s+e$ is also equal to β^{t-s}. This implies that in the constant rate discounted utility model only the difference in timing between the outcomes, i.e., $t-s$, affects preferences, but not the positions in time at which the outcomes occur.

Stationarity formalizes the idea that preferences depend only on the difference in timing and not on the exact timing of the outcomes. The definition of stationarity is as follows.

Definition 7: The individual preference relation $\succcurlyeq$ on X satisfies *stationarity* if there exists a common outcome $q \in Y$ with health status unequal to death, such that for all x_t, $y_t \in Y$, with health status unequal to death: $(x_1,\ldots,x_{T-1},q) \succcurlyeq (y_1,\ldots,y_{T-1},q)$ if and only if $(q,x_1,\ldots,x_{T-1}) \succcurlyeq (q,y_1,\ldots,y_{T-1})$.

In words, stationarity says that preferences over outcomes are unaffected if we move the common outcome from the last to the first period and delay all other outcomes with one period. Note that the differences in timing between the x_t are unaffected by this permutation of outcomes in time.

Theorem 3 summarizes Sections 3.3.1 and 3.3.2. A formal proof is given in Appendix A.

Theorem 3: If we replace symmetry by trade-off consistency in Theorems 1 and 2 then the general multiplicative discounting model $U[(c_1,h_1),\ldots,(c_T,h_T)] = \sum_{t=1}^{T} \lambda_t \upsilon(c_t) q(h_t)$ represents life-cycle preferences over consumption and health status. If stationarity is imposed in addition then the constant rate discounted utility model $U[(c_1,h_1),\ldots,(c_T,h_T)] = \sum_{t=1}^{T} \beta^{t-1} \upsilon(c_t) q(h_t)$ is representing.

3.3.3. An assessment of the conditions

To conclude the axiomatic analysis, we comment on the empirical contents of the conditions used to characterize the two discounted utility models. There exist no tests of trade-off consistency in the medical context. The main effect of trade-off consistency is to impose an additive representation. In an additive representation, preferences over outcomes in one period are independent of (common) outcomes in all other periods. Therefore, trade-off consistency is most likely to hold in decision contexts where complementarity between periods does not affect preferences over outcomes.

Studies that have tested stationarity yield negative results (Cairns and van der Pol, 1997; Bleichrodt and Johannesson, 1998). These studies find that people not only pay attention to differences in timing between outcomes, i.e., to their relative position in time, but also to the moment at which the outcomes occur, i.e., their absolute position in time. The general pattern that emerges from the literature on intertemporal preferences for health is that people are more timing averse, in the sense that their implied rate of time preference is higher, for delays that occur in the near future than for delays that occur in the more distant future. That is, the difference between year s and year t is discounted more than the difference between year $s+e$ and year $t+e$. This finding suggests that other discounting models than the constant rate discounted utility model may be more appropriate. If trade-off consistency holds then these models can be derived as special cases of the general discounted utility model.

4. Constant consumption, the willingness to pay for QALYs, and the valuation of longevity

In this section, we derive when consumption is constant over time. As shown in Section 2, constant consumption combined with Eq. (4) implies that cost-effectiveness analysis is consistent with life-cycle preferences over consumption and health status, or, which is equivalent, that cost-effectiveness analysis is consistent with cost–benefit analysis.

If consumption is constant and Eq. (4) holds, tractable expressions can be derived for the willingness to pay for QALYs and for the valuation of longevity. We extend the models proposed by Rosen (1988) by including health status in the analysis. We start with the deterministic case where the individual knows his life duration with certainty. The deterministic case illustrates the essential ideas while keeping the analysis relatively straightforward. We then turn to the more realistic stochastic case.

4.1. Certainty

4.1.1. The optimization problem

Consider an individual whose preferences can be described by Eq. (4) with constant rate discounting:

$$U=\sum_{t=1}^{T}\left[v(c_t)q(h_t)\frac{1}{(1+a)^{t-1}}\right] \tag{7}$$

where a is the individual's constant rate of time preference. We assume that v, the utility function over consumption, is strictly increasing and concave. The first derivative of v with respect to c is denoted by v_c.

The individual's wealth consists of initial wealth W and a fixed annual labor income w. The individual allocates his wealth between consumption and medical expenditures, which we denote by m. In each period, medical expenditures are a function of the sequence of quality of life levels $(q(h_1),\dots,q(h_T))$ and duration T: $m_t=g((q(h_1),\dots,q(h_T)),T)$. We assume that, for all t, the first derivative of medical expenditures with respect to $q(h_t)$, denoted by g_{h_t} is positive. The first derivative of medical expenditures with respect to duration, g_T, is also positive. The individual faces a pure capital market at which he can borrow and invest at interest rate r. The individual cannot die in debt and has no heirs. Under these assumptions, the individual's budget constraint becomes:

$$W+w\sum_{t=1}^{T}\frac{1}{(1+r)^{t-1}}=\sum_{t=1}^{T}[c_t+m_t]\frac{1}{(1+r)^{t-1}} \tag{8}$$

The Lagrangian expression for this problem is:

$$L=\sum_{t=1}^{T}v(c_t)q(h_t)\frac{1}{(1+a)^{t-1}}+\lambda\left\{W+w\sum_{t=1}^{T}\frac{1}{(1+r)^{t-1}}\right.$$
$$\left.-\sum_{t=1}^{T}\left[(c_t+m_t)\frac{1}{(1+r)^{t-1}}\right]\right\} \tag{9}$$

and the first order conditions are:

$$v_c(c_t)q(h_t)\frac{1}{(1+a)^{t-1}} = \lambda\frac{1}{(1+r)^{t-1}} \text{ for all } t\in[1,T] \tag{10a}$$

$$v(c_t)\frac{1}{(1+a)^{t-1}} = \lambda g_{h_t}\frac{1}{(1+r)^{t-1}} \text{ for all } t\in[1,T] \tag{10b}$$

$$v(c_T)q(h_T)\frac{1}{(1+a)^{T-1}} = \lambda[\{c_T+m_T-w\}\frac{1}{(1+r)^{T-1}} + \sum_{t=1}^{T}\frac{1}{(1+r)^{t-1}}g_T \tag{10c}$$

Under the assumption that the rate of time preference equals the interest rate $(a=r)$, it follows from Eqs. (10a) and (10b) that that the individual will choose c_t and h_t such that they are constant throughout $[1,T]$. Constancy of health status in turn implies constancy of medical expenditures. This analysis answers the question when an individual whose preferences can be described by Eq. (4) will select a constant consumption profile: he will do so when his rate of time preference equals the interest rate. Hence, given Eq. (4), cost-effectiveness analysis is consistent with cost–benefit analysis if the individual's rate of time preference is equal to the interest rate. This observation has important implications for the discussion whether the discount rate for health benefits can be different from the interest rate (the discount rate for costs) (Keeler and Cretin, 1983; Olsen, 1993b). Our analysis shows that if the purpose is to achieve equivalence between cost-effectiveness analysis and cost–benefit analysis then health benefits should be discounted at the interest rate, that is, health benefits should be discounted at the same rate as costs.

4.1.2. The willingness to pay for a QALY gained

We now derive the willingness to pay for a QALY gained if cost-effectiveness analysis is consistent with cost–benefit analysis. That is, we assume Eq. (4) and equality of the rate of time preference and the interest rate. We consider the case where a QALY is gained through an increase in quality of life with duration constant. The case where quality of life is held constant and duration changes is discussed in Section 4.1.3.

Because equality of the rate of time preference and the interest rate implies that $q(h_t)$ is constant for all t, g_{h_t} is constant for all t. Denote the constant value of $q(h_t)$ by q and the constant value of g_{h_t} by g_h. Rewriting Eq. (10b) gives:

$$\frac{v(c)}{g_h} = \lambda \tag{11}$$

which we substitute in Eq. (10a) to give:

$$g_h = \frac{v(c)}{v_c(c)q} = \frac{c}{q\varepsilon} \tag{12}$$

where $\varepsilon = v_c c/v$ is the elasticity of the utility function for consumption. The parameter ε reflects the possibilities for intertemporal substitution of consumption. The higher the ε, the better are the possibilities for intertemporal substitution. As ε tends to unity, $v(c)$ becomes more linear in c and the individual is less concerned about the distribution of consumption over time than in the value of aggregate consumption. In the limiting case, where ε is equal to one, the individual is not interested in the distribution of consumption over time, but only cares about total consumption.

The term g_h indicates how medical expenditures change as a result of a change in quality of life. Hence, g_h can be interpreted as the marginal cost of an additional unit of quality of life. A (discounted) QALY is gained if quality of life increases with $z = 1/(\sum_{t=1}^{T}(1/(1+r)^{t-1}))$ units. This expression displays that the higher the interest rate the greater the future gain in quality of life has to be in order to gain a QALY. Substitution of z in Eq. (12) defines the willingness to pay for a QALY gained. The willingness to pay for a QALY gained depends on four factors. It is increasing in consumption and the interest rate and decreasing in quality of life (ceteris paribus, individuals in worse health are willing to pay more for improvements in quality of life), and the possibilities for intertemporal substitution.

In providing a numerical illustration, we will avoid the complications associated with discounting by focusing on immediate improvements in health. Empirical estimates of ε are in the range 0.20–0.40 (Thaler and Rosen, 1975; Rosen, 1988). We assume that ε is equal to 0.25. Consider a person whose annual consumption is US$20,000 and whose quality of life initially is equal to 0.8. It follows that this individual will be willing to pay about US$100,000 per QALY for immediate improvements in health.

4.1.3. The willingness to pay for longevity

We now derive the willingness to pay for longevity. Substituting Eq. (11) in Eq. (10c) and rearranging gives:

$$g_T = \frac{c\dfrac{1-\varepsilon}{\varepsilon} + (w-m)}{A} \tag{13}$$

where

$$A = (1+r)^{T-1}\sum_{t=1}^{T}\frac{1}{(1+r)^{t-1}} = \sum_{t=1}^{T-1}(1+r)^{t-1}$$

The term g_T reflects the responsiveness of medical expenditures to changes in longevity. That is, g_T indicates the marginal cost of additional units of life or the willingness to pay for longevity. Eq. (13) displays that the willingness to pay for longevity increases in consumption and decreases in the interest rate (the higher the interest rate the less valued are future life years) and in ε. The negative sign of the elasticity ε follows from the fact that, ceteris paribus, additional longevity does not change the individual's total consumption throughout his life, but it changes the distribution of consumption. Because the individual's concern about the distribution of consumption decreases with ε, his willingness to pay for longevity decreases with ε.

Finally, the willingness to pay for longevity also increases in the surplus of the annual wage rate over the annual medical expenditures. During the additional time that the individual lives he is able to create more wealth. The higher his wage rate the more wealth he creates and the greater his possibilities for increasing consumption. The term $(w-m)$ does not appear in Section 4.1.2 where we considered QALY gains through increases in quality of life, because there duration was held constant and, hence, the individual's wealth was unaffected.

4.2. Risk

4.2.1. The optimization problem

We now turn to the more realistic case where the individual is uncertain about his life duration. Let f_t be the probability of living for t time periods and let F_t be the cumulative probability of surviving until time period t at most. That is, $F_t = \sum_{s=1}^{t-1} f_s$. Then, the probability of being alive at the beginning of time period t, denoted by S_t, is equal to $1-F_t$. The period-specific death rate ρ_t is defined as the probability of dying during time period t given that one has survived up to time period t. The period-specific death rate is a conditional probability defined as f_t/S_t. Hence, S_t and ρ_t are related as follows:

$$S_t = \prod_{s=1}^{t-1} (1-\rho_s) \tag{14}$$

We assume that the individual maximizes expected utility. For each t, if he lives exactly t periods, then his utility is defined by Eq. (4) with constant rate discounting and $T=t$. Therefore, the individual's expected utility is equal to:

$$\mathrm{EU} = \sum_{t=1}^{\infty} S_t v(c_t) q(h_t) \frac{1}{(1+a)^{t-1}} \tag{15}$$

We assume that the individual participates with a cohort of identical individuals in an actuarially fair annuity system (Yaari, 1965; Rosen, 1988). Each individual hands over his wealth to an insurance company in exchange for a contract that ensures him his optimal consumption and medical expenditure bundles until death.

Obviously, individual choices are restricted by the total wealth available. The budget constraint facing each individual is based on overall life expectancy. If initial wealth is positive then those individuals who die early effectively subsidize the insurance pool and the claims of the individuals who live longer than expected are financed out of these ''subsidies''. Under these assumptions, the individual's budget constraint becomes:

$$W + w\sum_{t=1}^{\infty} S_t \frac{1}{(1+r)^{t-1}} = \sum_{t=1}^{\infty} S_t(c_t + m_t)\frac{1}{(1+r)^{t-1}} \tag{16}$$

where the medical expenditures are in each period a function of the infinite sequences of the quality of life levels $(q(h_1), q(h_2), \dots)$ and the age-specific death rates $(\rho_1, \rho_2, \dots)$: $m_t = g((q(h_1), q(h_2), \dots), (\rho_1, \rho_2, \dots))$. The first derivatives of medical expenditures with respect to $q(h_t)$ and ρ_t are denoted by g_{h_t} and g_{ρ_t}, respectively.

Maximization of expected utility, Eq. (15), subject to the budget constraint, Eq. (15), yields the following first order conditions, in which common terms S_t have been cancelled.

$$\upsilon_c(c_t)q(h_t)\frac{1}{(1+a)^{t-1}} = \lambda\frac{1}{(1+r)^{t-1}} \text{ for all } t \in [1,\infty] \tag{17a}$$

$$\upsilon(c_t)\frac{1}{(1+a)^{t-1}} = \lambda g_{h_t}\frac{1}{(1+r)^{t-1}} \text{ for all } t \in [1,\infty] \tag{17b}$$

$$\begin{aligned} &\sum_{t=1}^{\infty}\frac{S_t}{1-\rho_s}\frac{1}{(1+a)^{t-1}}\upsilon(c_t)q(h_t) \\ &\quad = \lambda\left\{\sum_{t=1}^{\infty}\frac{S_t}{1-\rho_s}\frac{1}{(1+r)^{t-1}}(c_t+m_t-w)\right. \\ &\qquad \left. -g_{\rho_s}\sum_{t=1}^{\infty}S_t\frac{1}{(1+r)^{t-1}}\right\} \text{ for all } s \in [1,\infty] \end{aligned} \tag{17c}$$

The first order conditions in the decision problem under risk, Eqs. (17a) and (17b) are identical with the conditions under certainty, Eqs. (10a) and (10b), because the term S_t occurs both in Eqs. (15) and (16) and cancels out. Hence, the problem under risk becomes similar to that under certainty and it follows immediately that the optimal paths for consumption and quality of life are constant only if the rate of time preference equals the interest rate. Hence, the conclusion derived under certainty, that cost-effectiveness analysis will only be equivalent to cost–benefit analysis if the individual's rate of time preference is equal to the interest rate, is still valid if life duration is uncertain.

4.2.2. The willingness to pay for a QALY gained

Even though our conclusions about the consistency between cost-effectiveness analysis and cost–benefit analysis are unaffected by the introduction of risk, there is a change in the willingness to pay for a QALY gained. Again, we consider first the case where a QALY is gained through an increase in quality of life and duration is constant.

Let the rate of time preference be equal to the interest rate. Denote the constant value of $q(h_t)$ by q and the constant value of g_{h_t} by g_h. It then follows from Eqs. (17a) and (17b) that g_h, the willingness to pay for a QALY gained is still equal to:

$$g_h = \frac{v(c)}{v_c(c)q} = \frac{c}{q\varepsilon} \tag{18}$$

However, under uncertainty quality of life has to increase with

$$\frac{1}{S_t \sum_{t=1}^{\infty} \frac{1}{(1+r)^{t-1}}} \tag{19}$$

to generate an additional QALY. This increase exceeds the increase in the certainty case by the factor S_t. Because future life-years are uncertain, the survival probabilities act as an additional discount factor of future life-years. Alternatively stated, the willingness to pay for a given increase in quality of life will be lower in the uncertainty case because the risk of premature death leads to a lower weight for future periods.

4.2.3. The willingness to pay for longevity

Expected life-duration changes if the period-specific death rates change. Under uncertainty, the willingness to pay for longevity therefore follows from the willingness to pay for changes in the age-specific death rates. We determine the latter expression. Because expected life-duration is negatively related to changes in the period-specific death rates, the sign of the factors that determine the willingness to pay for increases in (expected) longevity is equal to the sign of the factors that determine the willingness to pay for reductions in the period-specific death rates.

Substitution of $\lambda = v(c)/g_h$ and Eq. (18) in Eq. (17c) gives after some rearranging:

$$g_{\rho_s} = -\frac{\left(c\frac{1-\varepsilon}{\varepsilon} + w - m\right)}{1-\rho_s} \tag{20}$$

Obviously, Eq. (20) is negative, because individuals will desire compensation for increases in the age-specific death rates. Eq. (20) shows that the willingness to pay for reductions in the age-specific death rates are positively related with consumption and the surplus of the wage rate over medical expenditures, and negatively with the possibilities for intertemporal substitution. Finally, the willingness to pay for reductions in the period-specific death rates increases with the death rate. Ceteris paribus, the higher the age-specific death rate, the higher the individual's willingness to pay for reductions in the age-specific death rate.

5. Discussion

This paper has derived the conditions under which QALY maximization is consistent with life-cycle preferences over consumption and health. By implication, this answers the question under which conditions cost-effectiveness analysis will give the same results as cost–benefit analysis. We have shown that cost-effectiveness analysis is equivalent to cost–benefit analysis if the lifetime utility function over consumption and health status is additive over time, if the one-period utility functions can be multiplicatively decomposed in a utility function over consumption and a utility function over health status, and if the utility of consumption is constant over time.

We have derived that the utility function has this form under expected utility if the preference relation satisfies marginality, symmetry, standard gamble invariance, and the zero condition and if the individual's rate of time preference is equal to the interest rate. The latter condition has an interesting implication for the debate about the appropriate rate of discount for health benefits: our analysis shows that if the aim is to achieve consistency between the results of cost-effectiveness analysis and individual preferences then health benefits must be discounted at the same rate as health costs.

Expected utility is now widely believed to be descriptively invalid and we therefore have also provided an axiomatic analysis of the above utility function under the most influential nonexpected utility theory, rank-dependent utility theory. Under rank-dependent utility theory, marginality is dropped and is replaced by generalized utility independence.

If cost-effectiveness analysis is consistent with cost–benefit analysis then tractable expressions can be derived for the willingness to pay for a QALY gained through increases in quality of life and for the willingness to pay for longevity. This analysis answers a research question posed by Johannesson (1995a), who argued that it is important to have information on the factors that determine the willingness to pay for a QALY gained if cost-effectiveness analysis is to be a useful tool in societal decisions about the allocation of health care resources. We show that the willingness to pay for a QALY gained is determined by four factors: wealth, life expectancy, health status and the possibilities for intertemporal substitution of consumption. The willingness to pay for a QALY gained increases with wealth and with life-expectancy and decreases with health status and the possibilities for intertemporal substitution of consumption.

Even though we have focused on QALY-based decision making, our central result is also valid for other outcome measures. For example if healthy-years equivalents (HYEs) are used the utility function over consumption and health still has to be multiplicative in the utility of consumption and the utility of health status and the utility of consumption still has to be constant over time. That is, standard gamble invariance, the zero condition, and equality between the interest rate and the individual's time preference all have to hold. However, symmetry and marginality respectively generalized utility independence can be dropped. A problem with using HYEs in this context is that the HYE, as intended by Mehrez and Gafni (e.g., Gafni and Birch, 1997), is not a utility and that additional conditions have to be imposed to use HYEs in life-cycle problems involving both consumption and health. In particular, the utility function over years in full health has to be linear (Johannesson, 1995b).

Our analysis is based on the view that economic evaluation should have a foundation in welfare economics. There is a different conception of economic evaluation, which places cost-effectiveness analysis outside the realm of welfare economics (e.g., Culyer, 1989; Williams, 1993; Donaldson, 1998). This extra-welfarist or decision-making approach posits that principles of optimization theory coupled with an exogenously specified objective function and an exogenously specified resource constraint suffice as a foundation for cost-effectiveness analysis. As noted by Johannesson (1995a) and Weinstein and Manning (1997), the decision-making approach provides little guidance if the object of cost-effectiveness analysis is to compare the efficiency of different programs and may well lead to problems of suboptimization. It is not our intention to resolve the controversy about the role of cost-effectiveness analysis. We only observe that there exists a perception of cost-effectiveness analysis which requires a foundation of cost-effectiveness analysis in welfare economics. This perception provides the rationale for this paper.

The aim of our axiomatic analysis is to reveal the conditions under which cost-effectiveness analysis is equivalent to cost–benefit analysis. Let us emphasize that we do not intend to argue that these conditions have descriptive or normative force. In fact, as already indicated in Section 3, we believe that several of the conditions are unlikely to hold. For example, marginality is a strong condition, and empirical evidence indicates that marginality does not hold in medical decision making (Maas and Wakker, 1994). Under rank-dependent utility theory, marginality is replaced by generalized utility independence, which may be more realistic being a strengthening of utility independence. There exists some empirical support for utility independence in the medical context. Symmetry also seems too restrictive given that empirical research (e.g., Cairns, 1994; Chapman, 1996) indicates that people have time preference, i.e., they give different decision weights to different time periods. However, as shown in Section 3, symmetry can be replaced by conditions that allow differential weighting of time points. The most common model that allows for differential weighting of time points is the constant rate discounted utility model. The central assumption in this model is stationarity. Empirical evidence on the validity of stationarity in medical decision making is negative. However, more general discounting models can be a good description (e.g., Cairns and van der Pol, 1997; Bleichrodt and Johannesson, 1998).

Future research should test the validity of the conditions identified in this paper. Difficulties in the empirical estimation of willingness to pay for changes in health status have spurred the use of cost-effectiveness analysis as a tool in the allocation of health care resources. If the viewpoint is accepted that economic evaluation of health care should have a foundation in welfare economic theory and if the conditions identified in this paper do not hold, then the way to advance for methodological research in economic evaluation is to try and solve the empirical problems surrounding cost–benefit analysis instead of resorting to cost-effectiveness analysis.

Acknowledgements

Werner Brouwer and two anonymous referees provided many helpful comments. Han Bleichrodt's research was made possible by a fellowship from the Royal Netherlands Academy of Arts and Sciences and John Quiggin's by a fellowship from the Australian Research Council.

Appendix A. Technical assumptions and proofs

A.1. Structural assumptions

The set X is a *Cartesian product* of the one-period outcome sets Y, which are assumed identical. The one-period outcome set Y is a Cartesian product of C^+ and H. C^+ is a subset of the set of nonnegative real numbers, which is a convex subset of a linear space over IR and hence endowed with the Euclidean topology. We assume that H is a connected topological space and that X and Y are both endowed with the product topology.

The set P consists of all *simple lotteries*: lotteries with finite support. The preference relation $\succcurlyeq$ over P satisfies the von Neumann–Morgenstern axioms (Jensen, 1967). Preferences over X are derived by restricting attention to riskless lotteries. Preferences over Y are derived by restricting attention to constant outcomes: $(c^{(1)},h^{(1)}) \succcurlyeq (c^{(2)},h^{(2)})$ if and only if the sequence that yields the pair $(c^{(1)},h^{(1)})$ in each time period is at least as preferred as the sequence that yields the pair $(c^{(2)},h^{(2)})$ in each time period. Preferences over consumption and over health status are derived from the preference relation over Y by restricting attention to those pairs (c,h) in which one of the attributes is held constant. That is, the preference relation over C^+ is defined as: for all $c^{(1)}$, $c^{(2)} \in C^+$ and for all $h \in H/\{\text{death}\}$, $c^{(1)} \succcurlyeq c^{(2)}$ iff $(c^{(1)},h) \succcurlyeq (c^{(2)},h)$. The preference relation over C^+ is assumed to satisfy *monotonicity*: for all c^1, $c^2 \in C^+$ such that $c^1 > c^2$ and for all $h \in H/\{\text{death}\}$, it is true that $(c^1,h) \succ (c^2,h)$. The preference relation over H

is defined as: for all $h^{(1)}$, $h^{(2)} \in H$ and for all $C \in C^{+}$, $h^{(1)} \succcurlyeq h^{(2)}$ iff $(c,h^{(1)}) \succcurlyeq (c,h^{(2)})$. We assume *preferential independence* of H from C^{+}: if $(c,h^{(1)}) \succcurlyeq (c,h^{(2)})$ for one $c \in C^{+}$ then $(c,h^{(1)}) \succcurlyeq (c,h^{(2)})$ for all $c \in C^{+}$. Loosely speaking, preferential independence can be interpreted as a monotonicity condition for the preference relation over H.

We say that consumption is *essential* if there exist c, $c' \in C^{+}$ and $h \in H$ such that $(c,h) \succ (c',h)$. Similarly, we say that health status is essential if there exist $c \in C^{+}$ and h, $h' \in H$ such that $(c,h) \succ (c,h')$. Recall that $a_t x$ denotes the outcome $x \in X$ with x_t replaced by a: $(x_1,\dots,x_{t-1},a,x_{t+1},\dots,x_T)$. A point in time $t \in S$ is essential if there exist $a_t x$, $b_t x \in X$ such that $a_t x \succ b_t x$. Essentiality of either consumption or health status implies that at least one point in time must be essential. We assume that both consumption and health status are essential (otherwise our problem would become trivial) and that at least two points in time are essential.

A.2. Proof of Theorem 1

By Theorem 4 in Fishburn (1965) marginality implies that $U(x) = \sum_{t=1}^{T} U_t(x_t)$. The proof that symmetry implies that all utility functions can be chosen identical has been given in Bleichrodt and Quiggin (1997).

Fix a $c_t^1 \in C^{+}$ and define $q(h_t)$ as $U(c_t^1,h_t)$. By standard gamble invariance, for all $c_t \in C^{+}$, $U(c_t,h_t)$ is strategically equivalent to $q(h_t)$. Hence, for all $c_t \in C^{+}$, $U(c_t,h_t)$ is a positive linear transform of $q(h_t)$: $U(c_t,h_t) = w_t(c_t) + \upsilon(c_t)q(h_t)$ with $w_t(c_t)$ real and $\upsilon(c_t)$ positive. Positivity of $\upsilon(c_t)$ follows because we used weak preference ($\succcurlyeq$) in the definition of standard gamble invariance. If we had used indifference, $\upsilon(c_t)$ would have been real and preference reversals would have been possible. Denote death by $h=0$ and scale $q(h_t)$ such that $q(0)=0$. By the zero condition, for all c_t^2, $c_t^3 \in C^{+}$: $w_t(c_t^2) + \upsilon(c_t^2)^* 0 = w_t(c_t^3) + \upsilon(c_t^3)^* 0$. Hence, $w_t(c_t)$ is constant. By the uniqueness properties of the von Neumann–Morgenstern utility function, we may subtract a constant to give $U(c_t,h_t) = \upsilon(c_t)q(h_t)$. By monotonicity $\upsilon(c_t)$ is increasing.

A.3. Proof of Theorem 2

Generalized utility independence implies utility independence of all subsets $A \subseteq S$ by setting $y=z$ and $\upsilon = w$. Miyamoto and Wakker (1996) have shown for two attributes that utility independence implies that $U(x)$ is either additive or multiplicative. Their argument can easily be generalized to more than two attributes. We illustrate the case where $U(x)$ is multiplicative. The case where $U(x)$ is additive can de derived in a similar fashion. By utility independence of x_1 from $x_2,\dots,x_T$ and by utility independence of $x_2,\dots,x_T$ from x_1 we have $U(x) = f_1(x_1) f_{2\dots T}(x_2,\dots,x_T)$, with f_1 and $f_{2\dots T}$ utility functions over Y and Y^{T-1}, respectively. Applying utility independence on Y^{T-1} gives $f_{2\dots T}(x_2,\dots,x_T) = f_2(x_2) f_{3\dots T}(x_3,\dots,x_T)$. Repeating this procedure gives the multiplicative utility function.

Now, we use generalized utility independence to distinguish between the multiplicative and the additive utility function. We give a proof by contradiction that generalized utility independence implies that the utility function over life-years must be additive. Suppose the utility function is multiplicative instead and let generalized utility independence hold. Then we have

$$w(p)Av_t(a) + [1 - w(p)]Bv_t(b) \geq w(p)Av_t(c) + [1 - w(p)]Bv_t(d) \tag{A1}$$

if and only if

$$w(p)Cv_t(a) + [1 - w(p)]Dv_t(b) \geq w(p)Cv_t(c) + [1 - w(p)]Dv_t(d) \tag{A2}$$

with $A = v_1(z_1)\dots v_{t-1}(z_{t-1})v_{t+1}(z_{t+1})\dots v_T(z_T)$, $B = v_1(y_1)\dots v_{t-1}(y_{t-1})v_{t+1}(y_{t+1})\dots v_T(y_T)$, $C = v_1(w_1)\dots v_{t-1}(w_{t-1})v_{t+1}(w_{t+1})\dots v_T(w_T)$, and $D = v_1(v_1)\dots v_{t-1}(v_{t-1})v_{t+1}(v_{t+1})\dots v_T(v_T)$.

From Eqs. (A1) and (A2), we derive

$$v_t(a) - v_t(c) \geq \frac{[1 - w(p)]B}{w(p)A}\{v_t(d) - v_t(b)\} \tag{A3}$$

if and only if

$$v_t(a) - v_t(c) \geq \frac{[1 - w(p)]D}{w(p)C}\{v_t(d) - v_t(b)\} \tag{A4}$$

which is clearly not always true. We derive a contradiction by assuming the multiplicative utility function and hence the additive utility function must hold. The rest of the proof is similar to the proof of Theorem 1.

A.4. Proof of Theorem 3

Wakker (1989) has shown that the structural assumptions and trade-off consistency ensure that there exists an additive representation $\sum_{t=1}^{T}\lambda_t U(x_t)$ over X. Because both P and $P\downarrow$ contain all riskless loteries, the preference relations over P and $P\downarrow$ are also representing under certainty. Further, there exist additive representations over P and $P\downarrow$ by marginality and generalized utility independence, respectively. It follows that $U(x) = \sum_{t=1}^{T}\lambda_t U(x_t)$ represents preferences over P and $P\downarrow$. Fishburn (1970) has shown that if stationarity is imposed as well then $U(x) = \sum_{t=1}^{T}\beta^{t-1}U(x_t)$. The rest of the proof is similar to the proof of Theorem 1.

References

Bleichrodt, H., 1995. QALYs and HYEs: under what conditions are they equivalent?. Journal of Health Economics 14, 17–37.

Bleichrodt, H., Johannesson, M., 1997. The validity of QALYs: an experimental test of constant proportional trade-off and utility independence. Medical Decision Making 17, 21–32.

Bleichrodt, H., Johannesson, M., 1998. Time preference for health: a test of stationarity versus decreasing timing aversion. Working paper, Erasmus University, Rotterdam.

Bleichrodt, H., Quiggin, J., 1997. Characterizing QALYs under a general rank dependent utility model. Journal of Risk and Uncertainty 15, 151–165.

Bleichrodt, H., Wakker, P., Johannesson, M., 1997. Characterizing QALYs by risk neutrality. Journal of Risk and Uncertainty 15, 107–114.

Cairns, J.A., 1994. Valuing future benefits. Health Economics 3, 221–229.

Cairns, J.A., van der Pol, M., 1997. Saving future lives: a comparison of three discounting models. Health Economics 6, 341–350.

Chapman, G.B., 1996. Temporal discounting and preferences for sequences of health and money. Journal of Experimental Psychology: Learning, Memory, and Cognition 22, 771–791.

Culyer, A.J., 1989. The normative economics of health care finance and provision. Oxford Review of Economic Policy 5, 34–58.

Donaldson, C., 1998. The (near) equivalence between cost-effectiveness analysis and cost–benefit analysis: fact or fallacy?. Pharmacoeconomics 13, 389–396.

Fishburn, P.C., 1965. Independence in utility theory with whole product sets. Operations Research 13, 28–45.

Gafni, A., Birch, S., 1997. QALYs and HYEs: spotting the differences. Journal of Health Economics 16, 601–608.

Garber, A.M., Phelps, C.E., 1997. Economic foundations of cost-effectiveness analysis. Journal of Health Economics 16, 1–31.

Harvey, C.M., 1994. The reasonableness of non-constant discounting. Journal of Public Economics 53, 31–51.

Jensen, N.E., 1967. An introduction to Bernoullian utility theory: I. Utility functions. Swedish Journal of Economics 69, 163–183.

Johannesson, M., 1995a. The relationship between cost-effectiveness analysis and cost–benefit analysis. Social Science and Medicine 41, 483–489.

Johannesson, M., 1995b. Quality-adjusted life-years versus healthy-years equivalents: a comment. Journal of Health Economics 14, 9–16.

Kahneman, D., Knetsch, J.L., Thaler, R.H., 1990. Experimental test of the endowment effect and the Coase theorem. Journal of Political Economy 98, 1325–1348.

Keeler, E.B., Cretin, S., 1983. Discounting of life-saving and other nonmonetary effects. Management Science 29, 300–306.

Keeney, R.L., Raiffa, H., 1976. Decisions with Multiple Objectives. Wiley, New York.

Loewenstein, G., Prelec, D., 1992. Anomalies in intertemporal choice: evidence and an interpretation. Quarterly Journal of Economics 107, 573–597.

Maas, A., Wakker, P.P., 1994. Additive conjoint measurement for multiattribute utility. Journal of Mathematical Psychology 38, 86–101.

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

Miyamoto, J.M., Eraker, S.A., 1988. A multiplicative model of the utility of survival duration and health quality. Journal of Experimental Psychology: General 117, 3–20.

Miyamoto, J.M., Wakker, P.P., 1996. Multiattribute utility theory without expected utility foundations. Operations Research 44, 313–326.

Miyamoto, J.M., Wakker, P.P., Bleichrodt, H., Peters, H., 1998. The zero-condition: a simplifying assumption in QALY measurement and multiattribute utility. Management Science 44, 839–849.

Olsen, J.A., 1993a. Time preference for health gains: an empirical investigation. Health Economics 2, 257–265.

Olsen, J.A., 1993b. On what basis should health be discounted?. Journal of Health Economics 12, 39–53.

Phelps, C.E., Mushlin, A.I., 1991. On the (near) equivalence between cost-effectiveness and cost–benefit analysis. International Journal of Technology Assessment in Health Care 7, 12–21.
Pliskin, J.S., Shepard, D.S., Weinstein, M.C., 1980. Utility functions for life years and health status. Operations Research 28, 206–224.
Pratt, J.W., Zeckhauser, R.J., 1996. Willingness to pay and the distribution of risk and wealth. Journal of Political Economy 104, 747–763.
Quiggin, J., 1982. A theory of anticipated utility. Journal of Economic Behavior and Organization 3, 323–343.
Quiggin, J., Wakker, P.P., 1994. The axiomatic basis of anticipated utility: a clarification. Journal of Economic Theory 64, 486–499.
Ried, W., 1998. QALYs versus HYEs: What's right and what's wrong. a review of the controversy. Journal of Health Economics 17, 607–625.
Rosen, S., 1988. The value of changes in life expectancy. Journal of Risk and Uncertainty 1, 285–304.
Samuelson, W., Zeckhauser, R., 1988. Status quo bias in decision making. Journal of Risk and Uncertainty 1, 7–59.
Thaler, R., Rosen, S., 1975. The value of saving a life: evidence from the labor market. In: Terleckyj, N. (Ed.), Household Production and Consumption. National Bureau of Economic Research, New York.
Torrance, G.W., 1986. Measurement of health state utilities for economic appraisal: a review. Journal of Health Economics 5, 1–30.
Torrance, G.W., Boyle, M.H., Horwood, S.P., 1982. Application of multi-attribute utility theory to measure social preferences for health states. Operations Research 30, 1043–1049.
Torrance, G.W., Furlong, W.J., Feeny, D.H., Boyle, M.H., 1995. Multi-attribute preference functions: health utilities index. Pharmacoeconomics 7, 503–520.
Torrance, G.W., Feeny, D.H., Furlong, W.J., Barr, R.D., Zhang, Y., Wang, Q., 1996. Multi-attribute utility function for a comprehensive health status classification system: health utilities index mark 2. Medical Care 34, 702–722.
von Neumann, J., Morgenstern, O., 1953. Theory of Games and Economic Behavior. Princeton Univ. Press, Princeton, NJ.
Wakker, P.P., 1989, Additive Representations of Preferences: A New Foundation of Decision Analysis. Kluwer, Dordrecht.
Weinstein, M.C., Manning, W.I., 1997. Theoretical issues in cost-effectiveness analysis. Journal of Health Economics 16, 121–128.
Williams, A., 1993, Cost–benefit analysis: applied welfare economics or general decision aid. In: Williams, A., Giardina, E. (Eds.), Efficiency in the Public Sector. Edward Elgar, London.
Yaari, M.E., 1965. Uncertain lifetime, life insurance and the theory of the consumer. Review of Economic Studies 32, 137–150.
Yaari, M.E., 1987. The dual theory of choice under risk. Econometrica 55, 95–115.