Applications of Non-Expected Utility

Han Bleichrodt, Erasmus University of Rotterdam, The Netherlands
Ulrich Schmidt, University of Kiel, Germany

1 Introduction

Since the work of von Neumann and Morgenstern (1944) expected utility (EU) has been the dominant framework for analyzing decision situations under risk and uncertainty. Starting with the well-known paradoxes of Allais (1953) and Ellsberg (1961), however, a large body of experimental evidence has been gathered which indicates that individuals tend to violate the assumptions underlying EU systematically. This empirical evidence has motivated researchers to develop alternative theories of choice under risk and uncertainty able to accommodate the observed patterns of behavior. These models are usually referred to as non-expected utility. Nowadays the rank-dependent models, in particular prospect theory, have become the most prominent alternative and, accordingly, these models will also be the main focus of our paper.

If the decisions of subjects are not in line with EU, applied models which rest on it may make wrong predictions. Therefore, applications of non-expected utility models may lead to a better accommodation of real world data. In general, applications of non-expected utility can be regarded as part of behavioral economics, a research stream which integrates psychological concepts into economis analysis and has received increasing attention in recent years. Non-expected utility models can in principle be applied to every economic setting involving risk. Due to this fact, it is impossible to cover all fields of applications in the present article. We have decided to focus on three fields, insurance economics, auctions, and health economics. Health economics is treated more extensively than the two other fields because it has become recently a very important research topic and to our knowledge no review of applications non-expected utility in the health domain exists.

The article is organised as follows. Section 2 gives notation and basic concepts. Section 3 describes expected utility and Section 4 describes important non-expected utility models. Section 5 to 7 are devoted to a discussion of applications of non-expected utility. Section 5 discusses applications in insurance, Section 6 to auctions, and Section 7 survey applications of non-expected utility in the health domain.

2. Notation and Basic Concepts

Let X denote a set of *outcomes*, which can be quantitative, for example money amounts or life durations, but also qualitative, e.g. states of health. The set of all probability measures or *prospects* over X will be denoted by P. A prospect $p \in P$ assigns a nonnegative probability p_i to outcome $x_i \in X$ and we have $p(X) = 1$. The set P includes all *riskless prospects*, i.e. prospects that assign probability one to one of the outcomes. The probability measure which assigns probability one to outcome x is denoted by δ_x. For convenience we restrict attention to probability measures with finite support, i.e. for all $p \in P$ there exists a finite $W \subset X$ with $p(W) = 1$.

Models of decision making analyze the preference of a decision maker between prospects which will be formalized by the binary relation $\succcurlyeq \subset P \times P$. For $p, q \in P$, $p \succcurlyeq q$ means that p is at least as good as q (weak preference). The strict preference relation $\succ$ and indifference relation $\sim$ are defined as usual. By restricting attention to riskless prospects the preference relation $\succcurlyeq$ defines a preference relation over outcomes, i.e. for all outcomes $x,y \in X$, $x \succcurlyeq y$ iff $\delta_x \succcurlyeq \delta_y$. A real-valued function V on P is called utility function if it *represents* $\succcurlyeq$ on P, i.e.

(1) $\quad p \succcurlyeq q \Leftrightarrow V(p) \geq V(q)$ for all $p, q \in P$.

We will denote prospects giving outcome x_i with probability p_i, $i = 1,\ldots,n$, as $(p_1, x_1;\ldots;p_n, x_n)$. Within this notation we implicitly assume that outcomes are rank-ordered from best to worst, i.e. $x_1 \succcurlyeq \ldots \succcurlyeq x_n$. *Binary prospects*, i.e. prospects that yield just two outcomes x_1 and x_2 with positive probabilities p_1 and $p_2 = 1-p_1$ will be denoted $(p_1, x_1; x_2)$ for short.

3. Expected Utility

Expected utility (EU) holds if the utility function representing preference can be restricted to be of the following form:

(2) $$V(p) = \sum_{i=1}^{n} u(x_i)p_i ,$$

i.e. the utility of a prospect equals the expected value of the utility of the single outcomes. The central condition in EU is the well known independence axiom: for all $\lambda \in [0, 1]$ and for all $r \in P$, $p \succcurlyeq q \Rightarrow \lambda p + (1-\lambda)r \succcurlyeq \lambda q + (1-\lambda)r$. The independence axiom has intuitive appeal and is accepted as a principle of rational choice by most authors. However, it is often violated in empirical studies. This empirical evidence has motivated the development of non-expectd utility models which usually rely on weakened variants of the independence axiom.

A decision maker is defined to be *risk-averse* if she dislikes mean-preserving spreads in risk. A mean-preserving spread results from increasing one outcome and decreasing a worse outcome without affecting the expected value of a prospect. Consequently, risk aversion holds if $(p_1, x_1;...; p_i, x_i; ...; p_j, x_j; ...; p_n, x_n) \succcurlyeq (p_1, x_1;...; p_i, x_i + \varepsilon/p_i; ...; p_j, x_j - \varepsilon/p_j;...; p_n, x_n)$ for all positive ε. It follows that risk aversion in EU is equivalent to a concave utility function u.

Two common graphical representations, the two-outcome diagram and the triangle diagram, may help to clarify some properties of EU. The two-outcome diagram (see Panel A of Figure 1) restricts attention to binary prospects with outcomes x_1 and x_2, which occur with probabilities p_1 and $1-p_1$. Setting the total differential of a constant utility $V = p_1u(x_1) + (1-p_1)u(x_2)$ equal to zero yields the slope of indifference curves

(3) $$\frac{dx_1}{dx_2} = \frac{1 - p_1}{p_1}\frac{u'(x_2)}{u'(x_1)}.$$

Indifference curves have a negative slope and are convex if risk aversion is assumed. Moreover, their slope equals the negative probability ratio along the 45°-axis which is also called certainty line since we have $x_1 = x_2$, i.e. the individual is in a riskless position along this line.

In the triangle diagram (Panel B of Figure 1) there are three fixed outcomes, $x_1 > x_2 > x_3$ with varying probabilites. Taking $p_2 = 1 - p_1 - p_3$, a fixed utility level is given by $V = p_1u(x_1) + (1 - p_1 - p_3)u(x_2) + p_3u(x_3)$. Solving for the best outcome, we get the equation for an indifference curve:

(4) $$p_1 = \frac{V - u(x_2)}{u(x_1) - u(x_2)} + p_3\frac{u(x_2) - u(x_3)}{u(x_1) - u(x_2)}.$$

It follows, that indifference curves are upwards sloping, parallel lines since the slope is independent of the utility level V. Note that a higher degree of risk aversion leads to steeper indifference curves. Consequently, parallel indifference curves mean that the degree of risk aversion is constant for all prospects consisting of these three outcomes.

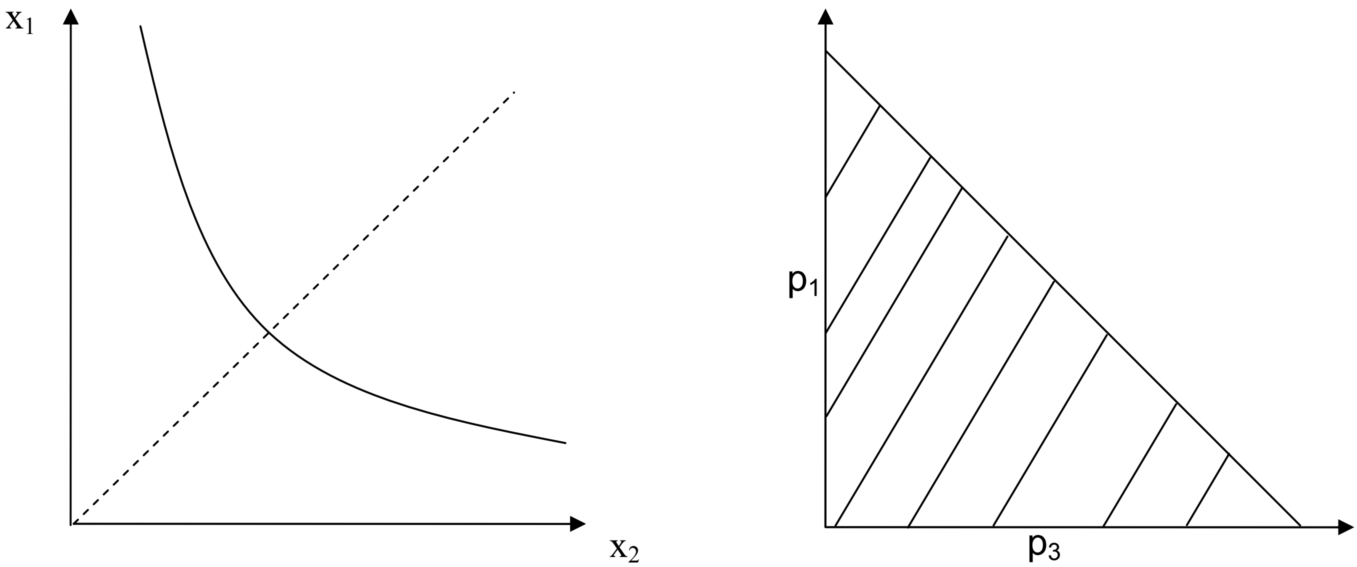

A: The Two-outcome Diagram B: The Triangle Diagram

Figure 1: Graphical Representations of EU

The descriptive validity of the independence axiom of EU was first questionend by Allais (1954). A well-known experimental design is the so-called common-ratio effect: there is a choice between p = (1, \$3000) and q = (0.8, \$4000; 0.2, \$0) and a second choice between p* = (0.25, \$3000; 0.75, \$0) and q* = (0.2, \$4000; 0.8, \$0). It turns out that many people choose p in the first problem and q* in the second one which violates EU. If we normalize utility by u(4000) = 1 and u(0) = 0, choosing p implies u(3000) > 0.8 while choosing q* implies 0.25u(3000) < 0.2 which is obviously a contradiction.

4. Non-expected utility

The experimental evidence against the independence axiom has motivated the development of various alternative models, for overviews see Starmer (2000), Schmidt (2004), and Sugden (2004). Two classes of models can be distinguished, utility theories with the betweenness property and rank-dependent models. These classes are disjoint in the sense that only the EU belongs to both. Utility theories with the betweenness property generalize EU by implying that indifference curves in the triangle diagram are also linear but not necessarily parallel. By this generalization the Allais paradox can be explained if indifference curves become steeper for higher utility level. This means that subjects are more risk averse when choosing between "good" prospects than for the choice between "bad" prospects, which is also called fanning out hypothesis (Machina 1982). Formally, betweenness is defined by $p \succ (\sim)\ q \Rightarrow p \succ (\sim)\ \lambda p + (1-\lambda)q \succ (\sim)\ q$ for all $1 > \lambda > 0$.

In the following we will focus on the family of rank-dependent models since these are currently the most popular in applications. One central model within this family is prospect theory (Tversky and Kahneman 1992, Wakker and Tversky 1993).[1] Prospect theory (PT) has three important differences compared to EU. First, it assumes that decision makers to not evaluate outcomes as final wealth levels but rather as deviations from a status quo, i.e. as gains and losses relative to a reference point. Second, decision makers are *loss averse*, which means that a given loss has a greater impact on the desirability of a prospect than a gain of equal size. Third, people do not evaluate probabilities linearly as in EU, but transform probabilities. Compared with EU, probabilities are replaced by decision weights π_i in all rank-dependent models. These decision weights are constructed by transforming probabilities through a weighting function w. In prospect theory, probability weighting can be different for gains and losses. Altogether, for a prospect consisting of k gains and n-k losses, we have in PT the following representation of preferences (recall that outcomes are rank-ordered from best to worst):

(5) $$V(p) = \sum_{i=1}^{k} v(x_i)\pi_i^+ + \sum_{i=k+1}^{n} v(x_i)\pi_i^- .$$

In this equation, the outcomes x_i are gains and losses relative to a reference point and not final wealth positions as in EU. The decision weights are defined as follows:

(6) $$\pi_i^+ = w^+(\sum_{j=1}^{i} p_j) - w^+(\sum_{j=1}^{i-1} p_j) \text{ and}$$

(7) $$\pi_i^- = w^-(\sum_{j=i}^{n} p_j) - w^-(\sum_{j=i+1}^{n} p_j),$$

with both weighting functions strictly increasing and satisfying $w^+(0) = w^-(0) = 0$ and $w^+(1) = w^-(1) = 1$. Note that in the domain of gains decumulative probabilities are transformed whereas in the domain of losses cumulative probabilites are transformed.

The value function v plays the same role as the utility function u in EU and is strictly increasing. The hypothesis of loss aversion can be captured by assuming that v is steeper in the domain of losses than in the domain of gains. A second important hypothesis is diminishing sensitivity according to which marginal utility is decreasing as one moves away from the reference point. Consequently, the value function is concave for gains and convex for losses. This leads to the reflection effect, according to which people are often risk averse for gains and risk seeking for losses.

---

[1] Prospect theory is sometimes referred to as cumulative prospect theory to distinguaish it from the original version of prospect theory proposed by Kahneman and Tversky 1979.

Empirical evidence confirms concave utility for gains and convex utility for losses (Tversky and Kahneman 1992, Abdellaoui 2000, Abdellaoui et al. forthcoming). There is also a lot of evidence supporting loss aversion, both in the laboratory and in field studies (Camerer 2000, Schmidt and Traub 2002, Pennings and Smidts 2003, Abdellaoui et al. forthcoming). Empirical evidence on probability weighting indicates that w has an inverse-S shaped form, indicating that people are sensitive to changes in probability around 0 (the impossibility effect) and 1 (the certainty effect) and much less so for intermediate probabilities (Tversky and Kahneman 1992, Tversky and Fox 1995, Wu and Gonzalez 1996, Gonzalez and Wu 1999, Abdellaoui 2000, Bleichrodt and Pinto 2000).

The development of PT was influenced by the rank-dependent utility (RDU) model of Quiggin (1982) which differs from EU only by probability weighting. Formally, RDU is given by

(8) $$V(p) = \sum_{i=1}^{k} u(x_i)\pi_i \ .$$

The construction of the decision weights is identical to that in Eq. 6 with a weighting function w. An interesting special case of RDU is the dual theory (DT) of Yaari (1987) which is given by (8) with the restriction $u(x_i) = x_i$. Although utility is linear, due to probability weighting we may also have risk averion in DT. More precisely, a decision maker in DT exhibits risk aversion if the weighting function is convex. This is because a convex weighting function, compared to untransformed probability, underweighs the probabilities of the best outcomes and overweighs the probabilities of the worst outcomes. In RDU risk aversion can be produced either by a convex weighting function or a concave utility function or both (Chateauneuf and Cohen 1994).

The utility of prospects in a two-outcome diagram for RDU is given by

(9) $$V = w(p_1)u(x_1) + (1-w(p_1)u(x_2) \text{ if } x_1 \succcurlyeq x_2$$

$$V = w(1-p_1)u(x_2) + \big(1-w(1-p_1)\big)u(x_1) \text{ if } x_2 \succ x_1$$

Calculating the slope of indifference curves yields

(10) $$\frac{dx_1}{dx_2} = \frac{1-w(p_1)}{w(p_1)}\frac{u'(x_2)}{u'(x_1)} \text{ if } x_1 \succcurlyeq x_2$$

$$\frac{dx_1}{dx_2} = \frac{w(1-p_1)}{1-w(1-p_1)}\frac{u'(x_2)}{u'(x_1)} \text{ if } x_2 \succ x_1$$

In case of a convex weighting function we have $1 - w(p_1) > w(1 - p_1)$ and $w(p_1) < 1 - w(1 - p_1)$. Consequently, risk aversion implies that indifference curves have a kink at the 45°-line. This kink is an important difference between EU and RDU (as well as PT and DT) and can be characterized by the concepts of first-order and second-order risk aversion. Consider a random variable ε with $E(\varepsilon) = 0$ and variance σ_ε^2. From Pratt (1964) it is known that the risk premium RP for avoiding tε in the case of EU with differentiable utility function can for a sufficiently small t be approximated by $RP \cong -(t^2/2)\, \sigma_\varepsilon^2 u''(x)/u'(x)$. The risk premium is, thus, proportional to t^2 and approaches zero faster than t which means that for small risks no risk premium will be demanded. This behavior has been termed second-order risk aversion by Segal and Spivak (1990). Second-order risk aversion does not only hold for EU with differentiable utility function but for all non-expected utility models which are smooth in the sense of Frèchet-differentiability. If we have a kink along the 45°-axis, however, Segal and Spivak (1990) have shown that RP is proportional to t which is called first-order risk aversion and yields $dRP/dt\big|_{t=0+} \neq 0$. First-order risk aversion implies that a decision maker will demand a risk premium also for infinitesimal small risks.

5. Applications of non-expected utility in insurance economics

Insurance Economics is a straightforward field for applying non-expected utility. In an important article Machina (1995) has shown that all classical results in insurance economics derived under EU carry over to non-expected utility as long as the representing utility functional exhibits second-order risk aversion. In the case of first-order risk aversion, however, some differences occur. We will show this by considering classical results by Mossin (1968), Arrow (1971), and Borch (1960).

Consider an individual with wealth $y > 0$ which is subject to a random loss $\tilde{L}$. If the individual insures the loss, she will receive an indemnity I(L) for paying a premium P(I(L)). In the case of coinsurance the indemnity is given by $I(L) = \alpha L$ where α can be chosen by the insured and is between zero and unity. The premium is usually given by $P(I(L)) = (1 + \lambda)\alpha E(\tilde{L})$ where $\lambda \geq 0$ is a loading factor for profits and fixed costs of the insurer. A well-known theorem by Mossin (1968) now states that the insured will choose full coverage ($\alpha = 1$) if and only if insurance is fair ($\lambda = 0$). For $\lambda > 0$ partial coverage ($\alpha < 1$) will be chosen. Mossin´s theorem can be easily explained in the case where $\tilde{L}$ has only two possibe realizations, $L > 0$ with probability p_1 and $L = 0$ with probability $1 - p_1$. Final wealth x is now given by $x_1 = y - L + \alpha L - (1 + \lambda)\alpha p_1 L$ and $x_2 = y - (1 + \lambda)\alpha p_1 L$, where $(1 + \lambda)\alpha p_1 L$ is the premium for a coverage of α. Taking differentials with respect to α yields $dx_1 = (1 - (1 + \lambda)p_1)Ld\alpha$ and $dx_2 = -(1 + \lambda)p_1 Ld\alpha$. Consequently, the slope of the budget line in a two outcome diagram is

$$(11) \quad \frac{dx_1}{dx_2} = -\frac{1 - (1 + \lambda)p_1}{(1 + \lambda)p_1} .$$

If full coverage is demanded, the individual is at a position on his certainty line. We know from Eq. 3 that the slope of indifference curves along the certainty line equals $-(1 - p_1)/p_1$. Full coverage is, thus, only optimal if $\lambda = 0$ since only then the slope of the budget line equals the slope of indifference curves along the certainty line.

The demand for coinsurance with non-expected utility prefernces was analyzed, among others by Doherty and Eeckhoudt (1995), Schmidt (1996), Schlesinger (1997), and Segal and Spivak (1990). Recall from Eq. 10 that the slope of indifference curves at the certainty line for RDU equals $-w(1 - p_1)/(1 - w(1 - p))$ for the case $x_2 > x_1$ which is the only relevant case if overinsurance is ruled out. Therefore, full coverage is optimal as long as the indifference cure is flatter than the budget line, i.e.

$$(12) \quad -\frac{1 - (1 + \lambda)p_1}{(1 + \lambda)p_1} \geq \frac{-w(1 - p_1)}{1 - w(1 - p_1)} ,$$

which yields $1 + \lambda \leq (1 - w(1 - p_1))/p_1$. In the case of risk aversion w is convex and, therefore, $(1 - w(1 - p_1))/p_1 > 1$. Consequently, due to the kink of indifference curves in the case of first-order risk aversion full coverage is also optimal for strictly positive loading factors. This means that Mossin´s (1968) theorem carries over only partly to non-expected utility preferences. Note that in the case of DT indifference curves are linear. Therefore, the individual either demands full coverage in case that inequality (12) holds or no coverage at all.

An alternative to coinsurance is deductible insurance. For deductible insurance the indemnity is given by

$$(13) \quad I(L) = \begin{cases} L - d & \text{if } L \geq d \\ 0 & \text{if } L < d. \end{cases}$$

Arrow (1971) has shown that for a given loading factor deductible insurance is the most preferred form of insurance contract for every insured who is a risk averse expected utility maximizer. Moreover, it can be shown that the optimal deductible d equals zero if and only if the loading factor equals zero.

Deductible insurance with non-expected utility has been analyzed by Karni (1992), Schlee (1995), and Schlesinger (1997). Karni (1992) has shown that the optimality of deductible insurance carries over to all non-expected utility models which satisfy second-order risk aversion. Schlesinger (1997) has generalized this result by showing that final wealth under every possible insurance contract is a mean-preseving spread in risk of final wealth under deductible insurance. Consequently, also every risk averse non-expected utility maximizer will prefer deductible insurance. In contrast to EU, however, the optimal deductible under first-order risk aversion may equal zero even for strictly positive loading factors.

Let us finally analyze efficient risk sharing. Consider n individuals who have to share state-dependent outcomes. For simplicity we assume that there exist only two possible outcomes x_1 and x_2 with $x_1 > x_2$. Borch (1960) has shown that efficient risk sharing under EU can be characterized by equal marginal rates of substitution, i.e. for any two individuals i and j it must be true that

$$(14)\qquad -\frac{1-p_1}{p_1}\frac{u_i'(x_2^i)}{u_i'(x_1^i)} = -\frac{1-p_1}{p_1}\frac{u_j'(x_2^j)}{u_j'(x_1^j)},$$

where p_1 is the probability of final wealth level x_1. An important conclusion from this equation is that an efficient risk sharing agreement leaves each individual with some residual wealth uncertainty. This can be explained as follows: suppose i is in a riskless position. This means that her marginal rate of substitution equals $-(1 - p_1)/p_1$. According to Eq. 14, the marginal rate of substitution for all other individuals has to equal also $-(1 - p_1)/p_1$, i.e. they are also in a riskless position. But this is impossible if there is aggregate risk.

Schmidt (1996, 1999a) has shown that this result does not carry over to first-order risk aversion, i.e. it is possible that some individuals are in a riskless position. Recall from Eq. 10 that the marginal rate of substitution at a riskless position is, in contrast to EU, not identical for all individuals but determined the probability weighting function. If one individual has a rather convex weighting function while the weighting function of the other individual is nearly linear indifference curves may have identical slope at the certainty line of the first individual. In the case of DT indifference curves are linear and their slope is solely determined by the weighting function. It turns out that efficient risk sharing here assigns all risk to the least risk averse individual while all others enjoy a riskless position.

For further applications of non-expected utility to insurance economics the reader is referred to Konrad and Skaperdas (1993), Wang et al. (1997), and Schmidt (1999b).

6. Applications of non-expected utility in auctions

Due to the increasing importance of auctions in the real world, the literature on auctions has grown rather rapidly in recent years. Since the analysis of many auction designs like combinatorial auctions is rather complex even for risk neutral bidders, applications of non-expected utility are rare in this context. In the present chapter we will focus on auctions of a single object and stick to the independent private values framework. In this framework each bidder i has a private valuation v_i of the auctioned object Z which is not known by the other bidders. Formally, we have $[Z - v_i] \sim \delta_0$ where $[Z - v_i]$ denotes receiving the object for paying v_i. All valuations are drawn from the same distribution over an interval $[v^+, v^-]$.

	first-price	second-price
open bids	descending bid auction	ascending bid auction
sealed bids	first-price sealed-bid auction	second-price sealed-bid auction

Table 1: The standard auction formats

The basic literature (see Engelbrecht-Wiggans 1980 for a review) assumes risk neutrality of bidders and analyzes four auction formats stated in Table 1. In the ascending bid auction open bidding prevails until no bidder is willing to raise the last bid. It is obvious that the optimal maximal bid of bidder i is v_i since on the one hand it does not make sense to bid more than v_i and on the other hand it is always possible to make a gain if the highest bid among the other bidders is below v_i. In the second-price sealed bid auction each bidder submits secretly a bid to the auctioneer and the highest bidder wins the auction and has to pay the second highest bid. Note that the own bid does not determine the price one has to pay (because this is determined by the second highest bid) but only whether one will buy the object for a given price or not. Since bidder i is willing to buy the object for all prices which do not exceed v_i the optimal bid is v_i. Consequently, the the two second-price auctions are demand revealing and the revenue of the auctioneer is the second highest valuation in both cases as long as bid increments are infinitesimal small in the ascending bid auction.

The first-price sealed bid auction equals the second-price sealed-bid auction but the highest bidder has to pay his own bid and not only the second highest bid. In this auction there is no dominant bidding strategy as the optimal bid is a trade-off between winning probability and profit. If valuations are distributed uniformly, optimal bids in the unique Nash equilibrium are given by $v_i(n - 1)/n$ where n is the number of participating bidders. This is also the Nash equilibrium of the descending bid auction which runs as follows: the auctioneer starts anouncing a prohibitively high price for the object and then continously decreases the price until one bidder accepts to buy the object for the current price. It is obvious that both first-price auctions are strategically equivalent because in both cases the bidder determines his bid without knowing any bid of her competitors. Moreover, the bid is also the price in both cases. Since also the two second-price auctions are strategically equivalent it remains to compare first-price with second-price auctions. Note that $v_i(n - 1)/n$ is precisely the expected value of the second highest valuation if v_i is the highest valuation. Consequently, the expected price for the bidder and, therefore, the expected revenue of the auctioneer is identical in all four standard auctions. This well-known revenue equivalence theorem first established by Vickrey (1961) is also valid if valuations are not uniformly distributed. However, if bidders are not risk neutral but risk averse expected utility maximizers, optimal bids in the first-price auctios exceed those in second-price auctions (Milgrom and Weber 1982). This can be explained as follows: the optimal strategy in second-price auctions is independent of risk attitude. In first-price auctions, however, the optimal bid is determined in a trade-off between winning probability and potential profit. Risk averse bidders are willing to solve this trade-off at a higher bid which increases the winning probability but decreases potential profit.

In the case of non-expected utility let us first analyze the ascending bid auction. Suppose a bidder $j \neq i$ was bidding slightly less than $v_i - \varepsilon$ for an infinitesimal small ε so bidder i has to choose between bidding v_i or quitting the auction. If he quits the auction, the consequence is obviously δ_0. If he bids v_i, he will win the auction with some probability λ and get $[Z - v_i]$. However, with probability $1 - \lambda$ another bidder bids more and i will also get δ_0. Consequently, a maximal bid of v_i is optimal if $\delta_0 \sim \lambda[Z - v_i] + (1 - \lambda)\delta_0$. Suppose that the auctioned object Z is a lottery which is often the case in the real world due to uncertainty of the precise quality of the object. Then this indifference is obviously only satisfied if betweennes holds (see the definition of betweenness in section 4). Consider now quasiconcave preferences defined by $p \sim q \Rightarrow \lambda p + (1 - \lambda)q \succ p$. Since $\delta_0 \sim \lambda[Z - v_i]$ quasiconcavity implies $\lambda[Z - v_i] + (1 - \lambda)\delta_0 \succ \delta_0$ and, thus, the optimal bid exceeds v_i. Quasiconcavity can be interpreted as preference for randomization. This preference causes bidders to stay in the auction even for prices above v_i since doing so yields a random consequence compared to quitting the auction. This result was first obtained by Karni and Safra 1989a, b). Analogously, the optimal bid is lower than v_i in the case of quasiconvex preferences defined by by $p \sim q \Rightarrow p \succ \lambda p + (1\text{-}\lambda)q$. Karni and Safra (1989a, b) also analyzed the second-price sealed-bid auction. In this auction the maximal bid is determined without knowing that another bidder will continue bidding until $v_i - \varepsilon$. Thus there is a chance that the bidder will get the object for a price much lower than v_i which is not longer the case in the ascending bid auction if another bidder was already bidding $v_i - \varepsilon$. Thus, the optimal bid is determined on a higher indifference curve than in the ascending bid auction. If in contrast to EU the degree of risk aversion may vary for different utility levels, the evaluation of the auctioned lottery and consequently the optimal bid may change. More precisely, Grimm and Schmidt (2000) have shown that the optimal bid in the second-price sealed-bid auction is lower that in the ascending bid auction if the preferences satisfy betweenness and fanning out. Altogether it turns out that the advantage of second price auctions, i.e. the fact that they elicit true valuations, does not carry offer to non-expected utility if the auctioned object is a lottery. In other words, there does not exist an incentive-compatible mechanism to elicit certainty equivalents for non-expected utility models. If, however, a deterministic object is auctioned second price auctions elicit true valuations for all preferences which are consistent with first-order stochastic dominance.

First-price auctions with non-expected utility were analyzed by Weber (1982), Karni (1988), and Grimm and Schmidt (2000). Consider a bidder in a descending bid auction whose valuation is already larger than the actual price b. The choice between accepting the actual price and waiting slightly longer is a choice between a sure gain of $[Z - b]$ and the lottery $\lambda[Z - b - \varepsilon] + (1 - \lambda)\delta_0$ where λ is the probability that no other bidder will accept the price $b - \varepsilon$. The optimal bid is thus determined by $[Z - b] \sim \lambda[Z - b - \varepsilon] + (1 - \lambda)\delta_0$. In the first-price sealed-bid auction bidder i lacks information that he can make a sure profit because there is the possibility that another bidder will place a higher bid than v_i. Consequently, the optimal bid is determined on a lower indifference curve. Since only for EU preferences the degree of risk aversion is equal on different indifference curves we can conclude that optimal bids in the descending bid auction and in the first-price sealed-bid auction are always identical if and only if EU holds. Suppose in contrast that preferences are consistent with the fanning out hypothesis. This means that the degree of risk aversion is higher for the decision in the descending bid auction which leads to a higher bid.

7. Applications of non-expected utility in the health domain

Utility theory is widely applied in medical decision making. The common way to evaluate new medical technologies is through cost-utility analysis in which the benefits of these technologies is expressed in terms of utility. The most popular utility model is the quality-adjusted life-years (QALY) model. The QALY model combines the two dimensions of health, life duration and health status, in a single index number and claims that the utility of a T years in health state Q is equal to

(15) $U(Q,T) = V(Q)*T$,

where V(Q) denotes a weight that reflects the utility (or attractiveness) of health state Q. QALYs play an important role in health policy in many countries. For example, in the UK the National Institute for Clinical Excellence (NICE) requires a cost-utility analysis based on QALYs before a treatment is eligible for inclusion in the NHS. In the Netherlands, the Council for Public Health and Care, the main advisory board of the Dutch government on health policy, recently recommended that only treatments that cost less than €80,000 per QALY gained should be included in the basic insurance package. Treatments costing more can only be insured through supplementary insurance.

Sometimes a more general form of the QALY model is proposed in which the utility for life duration is not linear, as in Eq. 13, but can be curved:

(16) $U(Q,T) = V(Q)*W(T)$,

where W(T) denotes the utility for life duration. In what follows we will refer to Eq. 16 as the *nonlinear QALY model* to distinguish it from Eq. 15.

QALYs have two important advantages, they are tractable, which makes them attractive for practical applications, and they are intuitive, one QALY can be interpreted as one year in good health, which makes them easy to communicate to policy makers. There are, however, also important methodological questions surrounding QALYs. In this chapter we will focus on two of these and, in particular, on the insights that non-expected utility has offered to solve these methodological questions. The first question relates to the validity of the QALY model. The QALY model is a simple model, which as explained above has clear advantages. However, this simplicity may also have a price: the QALY model could be too simple and may misrepresent people's preferences for health. To obtain insight into the descriptive validity of the QALY models, Eqs. 15 and 16, we need behavioral foundations that identify the conditions on which the models depend. These conditions can then be tested in experimental studies. As we will explain, non-expected utility has been very useful in fine-tuning these tests.

A second question on which we will focus is the estimation of the utilities V(Q) and W(T). The measures commonly used to measure V(Q) and W(T) yield systematically different results, which cannot be explained under expected utility. The insights from non-expected utility, most notably prospect theory, can help to reconcile some of these differences as we will show in what follows.

7.1. Non-expected utility and tests of the descriptive validity of QALYs

Pliskin et al. (1980) were the first to give a behavioral foundation for the QALY model. Their model was later simplified by Bleichrodt et al. (1997) and Miyamoto et al. (1998). Bleichrodt et al. (1997) and Miyamoto et al. (1998) showed that the crucial condition of the QALY model is that people be risk neutral with respect to life duration. That is, for a given health quality Q they should be indifferent between a risky treatment that gives life duration T_1 with probability p and T_2 with probability $1-p$ and $p*T_1 + (1-p)*T_2$ for sure. Empirical evidence has generally observed that people do not behave according to this condition but are risk averse with respect to life duration. For example, the median subject in Stiggelbout et al. (1994) was indifferent between 4 years for sure and a risky treatment, giving 10 years with probability ½ and 0 years (death) with probability ½.

The analyses of Pliskin et al. (1980), Bleichrodt et al. (1997), and Miyamoto et al. (1998) relied crucially on the assumption that people behave according to expected utility. Without this assumption their behavioral foundations are no longer true. For example, under rank-dependent utility it is very well possible that people have linear utility and are risk neutral with respect to life duration. Consider, for example, the median preference observed by Stiggelbout et al. (1994). If $w(½) = 0.40$ then this response is consistent with a linear utility for life duration.

As mentioned before in Section 3, evidence abounds that people violate expected utility. These violations of expected utility cast doubt on the validity of previous tests of the QALY model. Several authors have derived tests of the QALY model that are robust to violations of expected utility. Bleichrodt and Quiggin (1997) derived a test of the QALY model that is valid under a large class of non-expected utility models including rank-dependent utility. Recall that $(p,(Q_1,T_1); (Q_2,T_2))$ denotes the risky prospect that gives (Q_1,T_1) with probability p and (Q_2,T_2) with probability $1-p$. As before, we assume that all prospects are *rank-ordered,* i.e. $(Q_1,T_1) \succcurlyeq (Q_2,T_2)$. The condition Bleichrodt and Quiggin (1997) imposed, *constant marginal* utility, says that for all Q and for all ε small enough that the prospects involved are still rank-ordered and the life durations do not exceed the maximum possible life duration, $(p,(Q,T_1);(Q,T_2)) \sim (p,(Q,T_3);(Q,T_4))$ iff $(p,(Q,T_1+\varepsilon);(Q,T_2)) \sim (p,(Q,T_3+\varepsilon);(Q,T_4))$ and $(p,(Q,T_1);(Q,T_2)) \sim (p,(Q,T_3);(Q,T_4))$ iff $(p,(Q,T_1);(Q,T_2+\varepsilon)) \sim (p,(Q,T_3);(Q,T_4+\varepsilon))$. Constant marginal utility was tested and rejected by Bleichrodt and Pinto (2005). Bleichrodt and Miyamoto (2003) extended the analysis of Bleichrodt and Quiggin (1997) to prospect theory where outcomes can be both gains and losses. Miyamoto (1999) proposed another condition, constant proportional coverage, that allows to test the QALY model under rank-dependent utility. *Constant proportional coverage* holds if for all Q and for all $T_1>T_2>T_3$ and $T'_1> T'_2> T'_3$ whenever $(Q,T_2) \sim (p,(Q,T_1);(Q,T_3))$, $(Q,T'_2) \sim (p',(Q,T'_1);(Q,T'_3))$, and $(T_2 - T_3)/(T_1 - T_3) = (T'_2 - T'_3)/(T'_1 - T'_3)$ then $p = p'$. Doctor et al. (2004) showed that the condition is also valid under prospect theory if a plausible assumption about the location of the reference point is made. They tested constant proportional coverage and obtained support for it.

Miyamoto and Eraker (1988) were the first to test the nonlinear QALY model under a general utility theory and obtained support for it. Bleichrodt and Pinto (2005) considered an even more general utility theory and also obtained support for the nonlinear QALY model.

In summary, the insights from non-expected utility have helped to perform more robust tests of the QALY model. It is more plausible that people have linear utility for life duration under non-expected utility models because under these models risk attitude is captured not only by the utility function, as in expected utility, but also in other functions and parameters. For example, in rank-dependent utility part of people's risk attitude is captured by the probability weighting and in prospect theory also loss aversion plays an important part in the explanation of attitudes towards risk. The available on QALYs under non-expected utility is still limited but indicates that the QALY model may have been too easily dismissed. The QALY model may not describe people's preferences for health better than commonly thought. This is of course an important finding given the dominant role that the QALY model plays in practical health economics research.

7.2. Non-expected utility and the measurement of health state utilities

Non-expected utility theory also has implications for the main methods to determine V(Q) and W(T). One of the most widely used methods to measure the health state utilities V(Q) is the *standard gamble*. The standard gamble is, for example, used in the SF-6D a very popular valuation method in applied health economics (Brazier et al. 2002). In the standard gamble method people face a choice between an impaired health state Q for T years for sure and a risky treatment option that gives full health with probability p and death with probability 1−p.[2] The purpose of the standard gamble is to determine the probability p that leads to indifference between these two options. Under expected utility and the nonlinear QALY model it then follows that V(Q) = p.

It is well known that the standard gamble gives systematically higher utilities than other methods to determine health state utilities that do not involve risk. This obviously raises the question as to which method should be preferred. The traditional view was that the standard gamble should be considered the preferred method as it is based on expected utility. The idea was that medical decision analysis is a prescriptive exercise and that expected utility is a prescriptive theory and, hence, health utility measurement should be based on expected utility. The problem with this point of view is that the measurement of utility is a descriptive exercise and that if people do not behave according to expected utility then utilities that are derived under expected utility will be biased. Using biased utilities in cost-utility analysis runs the risk of resulting in the wrong recommendations for health policy.

Conclusive evidence of such biases was observed by Llewellyn-Thomas et al. (1982). They used two different ways to ask the standard gamble. The first way was as described above with full health and death as endpoints in the risky treatment option. We will refer to this format as the *direct method*. In the *chained method* they first established indifference between (Q,T) for sure and a risky treatment (q,(full health,T); (Q′,T)) where Q′ is a worse health state than Q, and then they established indifference between (Q′,T) for sure and a risky treatment (r,(full health,T); death). Under expected utility the first indifference in the chained method entails

(16) $V(Q) = q + (1 - q)*V(Q')$.

The second indifference implies V(Q′) = r and, hence, the chained method gives V(Q) = q + (1 − q)*r. Except for random error, we should therefore observe that p = q + (1 − q)*r when expected utility holds. However, Llewellyn-Thomas et al. observed that the chained method led to systematically higher utilities V(Q) (see also Rutten-van Mölken et al. 1995, Bleichrodt 2001), a clear violation of expected utility casting doubt on the validity of standard gamble measurements.

[2] In principle it is not necessary to use full health and death as the outcomes of the risky treatment but this is the way the standard gamble is commonly asked.

To determine W(T) two methods are commonly used. In the *probability equivalence (PE) method* people are asked for the probability p that makes them indifferent between T_2 years in some health state Q for sure and a risky treatment that gives probability p of T_1 years in health state Q and probability 1−p of T_3 years in health state Q. In the *certainty equivalence (CE) method* people are asked for the number of years S_2 in some health state Q for sure and a risky treatment that gives probability q of T_1 years in health state Q and probability 1−q of T_3 years in health state Q. Of course, if we use the same T_1 and T_3 in the PE and the CE and if we substitute the response from the PE in the CE (i.e. p = q)[3] then we should observe that $S_2 = T_2$ under expected utility. In fact, what is typically observed is that $S_2 > T_2$[4], which leads to a more concave utility for life duration under the PE method than under the CE method and which, obviously violates expected utility (Bleichrodt et al. 2001).

Several authors have tried to solve the above inconsistencies by using non-expected utility. Wakker and Stiggelbout (1995) explored the impact of correcting the standard gamble for probability weighting as in rank-dependent utility. They showed that if people have an inverse S-shaped probability weighting function then the resulting utilities are generally pushed downwards leading to more consistency with other utility measurement methods. Their conjectures were confirmed by two empirical studies. Bleichrodt et al. (1999) observed that correcting the standard gamble for inverse-S shaped probability weighting leads to utilities that are more consistent with people's preferences than using the uncorrected standard gamble utilities. Bleichrodt and Pinto (2000) developed a new parameter-free method to examine the shape of the probability weighting function and tested their method in a medical setting. They found that the probability weighting function was indeed inverse S-shaped. On the other hand, Bleichrodt (2001) showed that correcting the standard gamble for inverse S-shaped probability weighting did not resolve the difference between the direct version of the standard gamble and the chained version of the standard gamble that was first observed by Llewellyn-Thomas et al. (1982) but, instead, exacerbated this difference. Correcting for probability weighting cannot resolve the difference between the PE and the CE either even though the difference tends to be mitigated (Bleichrodt et al. 2001). Finally, Stalmeier and Bezembinder (1999) observed that correction for probability weighting reduced the difference between risky and riskless utilities for life duration but was not sufficient to resolve the difference.

[3] Or for that matter substitute the response from the CE in the PE (i.e. $T_2 = S_2$).
[4] For money outcomes this was already observed by

The missing element in the corrections by Wakker and Stiggelbout (1995) was that they did not account for loss aversion, the other main deviation from expected utility modeled by prospect theory. Bleichrodt et al. (2001) derived new formulas that allow the measurement of health state utilities under prospect theory. The crucial step in their analysis was the location of the reference point. Bleichrodt et al. (2001) conjectured that in the standard gamble and in the PE method people take the sure outcome as their reference point, because this outcome is given. In the CE method, however, the sure outcome has to be determined and it is therefore unlikely to serve as the reference point. People will adopt a reference point in the CE but it will not be the sure outcome. Their data confirmed these hypotheses. Bleichrodt et al. (2001) tested whether their formulas could explain the discrepancy between the PE and the CE methods and their results were very encouraging: the corrections made the discrepancy vanish. Further support for the hypotheses in Bleichrodt et al. (2001) comes from several studies in the literature, some of which recorded people's though processes in responding to PE and CE questions (Stalmeier and Bezembinder 1999, Morrison 2000, Robinson et al. 2001, van Osch et al. 2004,2006). Bleichrodt et al. (forthcoming) extended the analysis of Bleichrodt et al. (2001) to other utility measurement procedures and found that prospect theory performed clearly better than expected utility and rank-dependent utility and solved many inconsistencies that were observed under expected utility.

Oliver (2003) observed, however, that the corrections of Bleichrodt et al. (2001) could not entirely explain the differences between the direct version of the standard gamble and the chained version of the standard gamble. He found that the best performing model was a version of prospect theory in which there was loss aversion but no probability weighting. Stalmeier (2002) found evidence, however, that the difference between the direct version of the standard gamble and the chained version of the standard gamble is produced by a more elementary violation of rationality than probability weighting or loss aversion. He observed that people tended to give the same probability in the direct version and to both questions in the chained version, i.e. $p = q = r$. That is, people anchor on their response and do not adjust it for differences in health quality. Such basic violations of rationality are hard to accommodate by any formal theory of decision under risk.

In summary, the insights from prospect theory seem to have improved the measurement of the utility in the health domain. Inconsistencies are significantly reduced when corrections for probability weighting and loss aversion are applied. An important implication of these findings is that the standard gamble as it is commonly used will result in utilities that are far too high. As was shown by Bleichrodt et al. (2001), under the common findings in the literature (inverse-S shaped probability weighting and losses loom larger than gains) uncorrected standard gamble utilities are clearly biased upwards. In particular, there are serious reasons to suspect that the widely used SF-6D method produces biased utilities and care should be taken in applying this algorithm.

References:

Abdellaoui, M. 2000. Parameter-free elicitation of utilities and probability weighting functions. *Management Science* **46** 1497-1512.

Abdellaoui, M., H. Bleichrodt, C. Paraschiv. forthcoming. Measuring loss aversion under prospect theory: A parameter-free approach. *Management Science.*

Allais, M. 1953. Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'ecole américaine. *Econometrica* **21** 503-546.

Arrow, K. J. 1971. *Essays in the theory of risk-bearing*. North Holland, Amsterdam.

Bleichrodt, H. 2001. Probability weighting in choice under risk: An empirical test. *Journal of Risk and Uncertainty* **23** 185-198.

Bleichrodt, H., J. M. Abellan, J. L. Pinto, I. Mendez. forthcoming. Resolving inconsistencies in utility measurement under risk: Tests of generalizations of expected utility. *Management Science*

Bleichrodt, H., J. Miyamoto. 2003. A characterization of quality-adjusted life-years under cumulative prospect theory. *Mathematics of Operations Research* **28** 181-193.

Bleichrodt, H., J. L. Pinto. 2000. A parameter-free elicitation of the probability weighting function in medical decision analysis. *Management Science* **46** 1485-1496.

Bleichrodt, H., J. L. Pinto. 2005. The validity of QALYs under non-expected utility. *Economic Journal* **115** 533-550.

Bleichrodt, H., J. L. Pinto, P. P. Wakker. 2001. Using descriptive findings of prospect theory to improve the prescriptive use of expected utility. *Management Science* **47** 1498-1514.

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

Bleichrodt, H., J. van Rijn, M. Johannesson. 1999. Probability weighting and utility curvature in QALY based decision making. *Journal of Mathematical Psychology* **43** 238-260.

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

Borch, K. 1960. The safety loading of reinsurance premiums. *Skandinavisk Aktuarietskrift* 153-184.

Brazier, J., J. Roberts, M. Deverill. 2002. The estimation of a preference-based measure of health from the SF-36. *Journal of Health Economics* **21** 271-292.

Camerer, C. F. 2000. Prospect theory in the wild: Evidence from the field. D. Kahneman and A. Tversky, eds. *Choices, values and frames*. Cambridge University Press, New York, 288-300.

Chateauneuf, A., M. Cohen. 1994. Risk seeking with diminishing marginal utility in a non-expected utility model. *Journal of Risk and Uncertainty* **9** 77-91.

Doctor, J. N., H. Bleichrodt, J. Miyamoto, N. R. Temkin, S. Dikmen. 2004. A new and more robust test of QALYs. *Journal of Health Economics* **23** 353-367.

Doherty, N. A., L. Eeckhoudt. 1995. Optimal insurance without expected utility: The dual theory and the linearity of insurance contracts. *Journal of Risk and Uncertainty* **10** 157-179.

Ellsberg, D. 1961. Risk, ambiguity and the Savage axioms. *Quarterly Journal of Economics* **75** 643-669.

Engelbrecht-Wiggans, R. 1980. Auctions and bidding models: A survey. *Management Science* **26** 119-142.

Gonzalez, R., G. Wu. 1999. On the form of the probability weighting function. *Cognitive Psychology* **38** 129-166.

Grimm, V., U. Schmidt. 2000. Equilibrium bidding without the independence axiom: A graphical analysis. *Theory and Decision* **49** 361-374.

Hershey, J. C., P. J. H. Schoemaker. 1985. Probability versus certainty equivalence methods in utility measurement: Are they equivalent? *Management Science* **31** 1213-1231.

Kahneman, D., A. Tversky. 1979. Prospect theory: An analysis of decision under risk. *Econometrica* **47** 263-291.

Karni, E. 1988. On the equivalence between descending bid auctions and first price sealed bid auctions. *Theory and Decision* **25** 211-217.

Karni, E. 1992. Optimal insurance: A nonexpected utility analysis. G. Dionne, eds. *Contributions to insurance economics*. Kluwer, Boston,

Karni, E. 1995. Non-expected utility and the robustness of the classical insurance paradigm: Discussion. *Geneva Papers on Risk and Insurance Theory* **20** 51-56.

Karni, E., Z. Safra. 1989a. Ascending bid auctions with behaviorally consistent bidders. *Annals of Operations Research* **19** 435-446.

Karni, E., Z. Safra. 1989b. Dynamic consistency, revelations in auctions and the structure of preferences. *Review of Economic Studies* **56** 421-434.

Konrad, K. A., S. Skaperdas. 1993. Self-insurance and self-protection: A nonexpected utility analysis. *Geneva Papers on Risk and Insurance Theory* **18** 131-146.

Llewellyn-Thomas, H., H. J. Sutherland, R. Tibshirani, A. Ciampi, J. E. Till, N. F. Boyd. 1982. The measurement of patients' values in medicine. *Medical Decision Making* **2** 449-462.

Machina, M. 1982. 'expected utility' analysis without the independence axiom. *Econometrica* **50** 277-323.

Machina, M. J. 1995. Non-expected utility and the robustness of the classical insurance paradigm. *Geneva Papers on Risk and Insurance Theory* **20** 9-50.

Milgrom, P. R., R. J. Weber. 1982. A theory of auctions and competitive bidding. *Econometrica* **50** 1089-1122.

Miyamoto, J. M. 1999. Quality-adjusted life-years (QALY) utility models under expected utility and rank dependent utility assumptions. *Journal of Mathematical Psychology* **43** 201-237.

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

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

Morrison, G. C. 2000. The endowment effect and expected utility. *Scottish Journal of Political Economy* **47** 183-197.

Mossin, J. 1968. Aspects of rational insurance purchasing. *Journal of Political Economy* **76** 553-568.

Oliver, A. 2003. The internal consistency of the standard gamble: Tests after adjusting for prospect theory. *Journal of Health Economics* **22** 659-674.

Pennings, J. M. E., A. Smidts. 2003. The shape of utility functions and organizational behavior. *Management Science* **49** 1251-1263.

Pliskin, J. S., D. S. Shepard, M. C. Weinstein. 1980. Utility functions for life years and health status. *Operations Research* **28** 206-223.

Pratt, J. W. 1964. Risk aversion in the small and in the large. *Econometrica* **32** 83-98.

Robinson, A., G. Loomes, M. Jones-Lee. 2001. Visual analog scales, standard gambles, and relative risk aversion. *Medical Decision Making* **21** 17-27.

Rutten-van Mölken, M. P., C. H. Bakker, E. K. A. van Doorslaer, S. van der Linden. 1995. Methodological issues of patient utility measurement. Experience from two clinical trials. *Medical Care* **33** 922-937.

Schlee, E. E. 1995. The comparative statics of deductible insurance in expected- and non-expected utility theories. *Geneva Papers on Risk and Insurance Theory* **20** 57-72.

Schlesinger, H. 1997. Insurance demand without the expected-utility paradigm. *Journal of Risk and Insurance* **64** 19-39.

Schmidt, U. 1996. Demand for coinsurance and bilateral risk-sharing with rank-dependent utility. *Risk, Decision and Policy* **1** 217-228.

Schmidt, U. 1999a. Efficient risk-sharing and the dual theory of choice under risk. *Journal of Risk and Insurance* **66** 597-608.

Schmidt, U. 1999b. Moral hazard and first-order risk aversion. *Journal of Economics* **Supplement 8** 167-179.

Schmidt, U. 2004. Alternatives to expected utility: Some formal theories. S. Barbéra, P. J. Hammond and C. Seidl, eds. *Handbook of utility theory, vol. 2*. Kluwer, Dordrecht, the Netherlands, 757-838.

Schmidt, U., S. Traub. 2002. An experimental test of loss aversion. *Journal of Risk and Uncertainty* **25** 233-249.

Segal, U., A. Spivak. 1990. First order versus second order risk aversion. *Journal of Economic Theory* **51** 111-125.

Stalmeier, P. F. M. 2002. Discrepancies between chained and classic utilities induced by anchoring with occasional adjustments. *Medical Decision Making* **22** 53-64.

Stalmeier, P. F. M., T. G. G. Bezembinder. 1999. The discrepancy between risky and riskless utilities: A matter of framing? *Medical Decision Making* **19** 435-447.

Starmer, C. 2000. Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk. *Journal of Economic Literature* **28** 332-382.

Stiggelbout, A. M., G. M. Kiebert, J. Kievit, J. W. H. Leer, G. Stoter, J. C. J. M. de Haes. 1994. Utility assessment in cancer patients: Adjustment of time tradeoff scores for the utility of life years and comparison with standard gamble scores. *Medical Decision Making* **14** 82-90.

Sugden, R. 2004. Alternatives to expected utility. S. Barbéra, P. J. Hammond and C. Seidl, eds. *Handbook of utility theory, vol. 2*. Kluwer, Dordrecht, the Nethelands, 685-755.

Tversky, A., C. Fox. 1995. Weighing risk and uncertainty. *Psychological Review* **102** 269-283.

Tversky, A., D. Kahneman. 1992. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty* **5** 297-323.

van Osch, S., M.C., W. B. van den Hout, A. M. Stiggelbout. 2006. Exploring the reference point in prospect theory: Gambles for length of life. *Medical Decision Making* **26** 338-346.

van Osch, S. M. C., P. P. Wakker, W. B. van den Hout, A. M. Stiggelbout. 2004. Correcting biases in standard gamble and time tradeoff utilities. *Medical Decision Making* **24** 511-517.

Vickrey, W. 1961. Counter speculation, auctions and competitive sealed tenders. *Journal of Finance* **16** 8-37.

von Neumann, J., O. Morgenstern. 1944. *The theory of games and economic behavior*. Princeton University Press, Princeton, NJ.

Wakker, P. P., A. M. Stiggelbout. 1995. Explaining distortions in utility elicitation through the rank-dependent model for risky choices. *Medical Decision Making* **15** 180-186.

Wakker, P. P., A. Tversky. 1993. An axiomatization of cumulative prospect theory. *Journal of Risk and Uncertainty* **7** 147-176.

Wang, S. S., V. R. Young, H. H. Panjer. 1997. Axiomatic characterization of insurance prices. *Insurance: Mathematics and Economics* **21** 173-183.

Weber, R. 1982. The Allais paradox, Dutch auctions, and alpha-utility theory. Working Paper, Nr. 536, Kellogg Graduate School of Management, Northwestern University, Evanston.

Wu, G., R. Gonzalez. 1996. Curvature of the probability weighting function. *Management Science* **42** 1676-1690.