

Supported by sanofi-aventis



Christopher McCabe PhD
 Professor of Health Economics, Institute of Health Sciences, University of Leeds

What is cost–utility analysis?

- Analytic techniques used for **economic evaluation** in healthcare (cost–benefit analysis, cost-effectiveness analysis and cost–consequences analysis) are designed to compare alternative courses of action, in terms of costs and outcomes. The choice of technique depends on the decision they intend to influence.
- **Quality-adjusted life-years (QALYs)** measure health as a combination of the duration of life and the health-related quality of life.
- The primary outcome of a **cost–utility analysis** is the cost per QALY, or **incremental cost-effectiveness ratio (ICER)**, which is calculated as the difference in the expected cost of two interventions, divided by the difference in the expected QALYs produced by the two interventions.
- The results of cost–utility analysis are compared with a **threshold ICER**; interventions with an ICER below this threshold are funded, whereas those with an ICER above the threshold tend not to be.
- Cost–utility analysis was developed to help decision-makers **compare** the value of alternative interventions that have very **different health benefits**, and it facilitates these comparisons without recourse to placing monetary values on different health states. Cost–utility analysis specifies what value is attached to specific health states, and thus increasingly facilitates the **transparency** of **resource allocation** processes.
- The challenges associated with the utilisation of cost–utility analysis for resource allocation are primarily related to the adequacy of QALYs for capturing the **value that society attaches** to healthcare interventions. Cost–utility analysis is, therefore, not a sufficient basis for resource allocation decisions; however, it is a useful technique and it performs a necessary function better than previous methods.

For further titles in the series, visit:
www.whatisseries.co.uk

What is cost–utility analysis?

Economic evaluation in healthcare

Economic evaluation is a family of analytic techniques, which are designed to compare two or more alternative courses of action in terms of costs and outcomes. Economic evaluation allows decision-makers to consider the relative and, potentially, the absolute value of alternative uses of available resources. Broadly speaking, there are three types of economic evaluation: cost–benefit analysis, cost–effectiveness analysis and cost–consequences analysis. The difference between these techniques is the way in which outcomes are measured.

- Cost–benefit analysis attaches a monetary value to outcomes.
- Cost–effectiveness analysis uses a natural unidimensional index of outcome; for example, the glycated haemoglobin (HbA1c) measure of blood glucose or the occurrence of a myocardial infarction.
- Cost–consequences analysis reports a profile of outcomes for each alternative but does not combine the outcomes into a single unit of effect.

Outcomes for economic evaluation

The different types of economic evaluation are defined by the approach that is used to measure the outcomes or effects of the intervention. The choice of technique in economic evaluation should reflect the decision that the economic evaluation is designed to inform. Thus, if an evaluation is designed to inform the choice between two alternative diabetic medications for the same patient group, a cost–effectiveness analysis using HbA1c as the outcome measure may be appropriate. When a decision is between alternatives that have both health and non-health effects, such as an environmental control regulation, a cost–consequences analysis may be appropriate. When it is possible to attach a monetary value to all the effects of the interventions being compared, a cost–benefit analysis may be the most appropriate technique.

Quality-adjusted life-years

Quality-adjusted life-years (QALYs)

measure health as a combination of the duration of life and the health-related quality of life (HRQoL). HRQoL is measured on a preference scale anchored at 1 (perfect or best imaginable health) and 0 (a quality of life [QoL] as bad as being dead). Figure 1 illustrates how QALYs can capture improvements in both life expectancy and QoL. It is this facility that allows decision-makers, in principle, to use QALYs to **compare the value of interventions** across the full range of healthcare activities.

The Q in QALYs is obtained by asking individuals to trade off improvements in their health status against either life expectancy (time trade-off) or risk of death (standard gamble). The results of these exercises are called ‘**utilities**’. The utility of a health state has been shown to vary systematically on the basis of who is asked, how the question is asked, and how the health state is described.

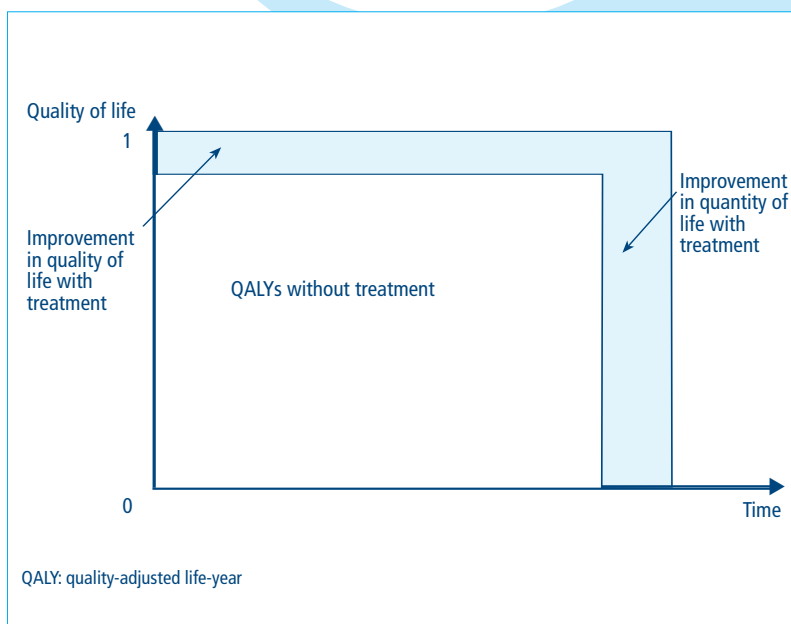


Figure 1. Quality-adjusted life-years combine quality and quantity of life

The utility of a health state is combined with the time spent in the health state to calculate the total QALYs lived in that health state. It is assumed that the utility of a health state is not affected by previous or subsequent health states, nor by the amount of time spent in that health state. This allows the QALYs to be summed over a profile of health states to estimate the total QALYs over an extended period of time – up to a lifetime.*

QALYs are normally calculated using utilities from ‘off-the-shelf’ health status instruments. The most widely used of these are the EQ-5D, the Health Utilities Index Mark 3 and the SF-6D.† The utilities for these instruments are based upon surveys of the general population. There is evidence of variation in the sensitivity of these instruments to the impact of different health problems on an individual’s QoL. Increasingly, condition-specific health status measures are being developed to address this issue; however, there are concerns about the comparability of these utilities across diseases. In response to this, decision-makers are increasingly specifying a preferred generic health status instrument for so-called ‘reference case’ analyses, while accepting supplementary analyses that use condition-specific measures.

Incremental cost-effectiveness ratio

The primary outcome of a **cost-utility analysis** is the **incremental cost-effectiveness ratio (ICER)**, otherwise known as the **cost per QALY**. This is calculated as the difference in the expected cost of two interventions, divided by the difference in the expected QALYs produced by the two interventions.

The early literature on cost-utility analysis assumed that the results of these analyses would be used to construct a cost-utility league table: a ranked list of ICERs, with the most efficient intervention at the top and the

least efficient intervention at the bottom. These could be used by decision-makers to identify which treatments to fund, by starting with the most efficient intervention (the one with the lowest cost per QALY) and moving down the league table until the entire available budget was exhausted. However, this would require information on the incremental cost-effectiveness of all healthcare interventions provided by the health service. This level of information is not, and is never likely to be, available. The results of cost-utility analyses are, therefore, now compared with a **threshold ICER**. Interventions with an ICER below the threshold are normally funded, whereas interventions with an ICER above the threshold tend not to be. Interventions with a high ICER may be funded on the basis of other considerations such as the severity of the condition and the availability of alternative treatments. The threshold ICER is often referred to as the willingness to pay for health gain.

Willingness to pay for health gain and cost-effectiveness threshold

In a centrally funded healthcare system, like the NHS in the UK, the willingness to pay for health is implied by the budget allocated to the health service by parliament. Once the budget has been determined, the function of a prioritisation process, such as that developed by the National Institute for Health and Clinical Excellence, is to promote the efficient use of that budget. To this end, it is important that new technologies that are introduced to the system are at least as efficient as the technologies that are displaced from the system in order to pay for them. In this type of system the ICER threshold is, therefore, the estimate of the ICER of the least efficient intervention currently provided. The use of cost-utility analysis does not require that we know the ICER for every intervention available; only an estimate of the ICER for the least efficient interventions, which will have to be displaced to fund the implementation of new interventions. Recent work has shown how these can be estimated from programme

* Note the QALY model makes a number of other assumptions. See publications listed under ‘Further reading’ for more details.

† See *What is quality of life?* (details under ‘Further reading’) for a more detailed discussion of these instruments.

budget data. There is increasing interest in the feasibility of modifying the threshold to reflect differences in the value society attaches to health gain depending on the characteristics of the individuals who receive it.* A detailed discussion of these issues is outside the remit of this paper.†

Advantages of cost-utility analysis

Cost-utility analysis was developed to address the problem of conventional cost-effectiveness analysis, which did not allow decision-makers to compare the value of interventions for different health problems. The need for such a comparison persists and, as healthcare costs continue to increase, it is likely to become more rather than less pressing. Unlike cost-benefit analysis, the conventional technique for evaluating different uses of public funding, cost-utility analysis **facilitates** these **comparisons** without recourse to placing monetary values on different health states and indeed life itself; with all the technical and ethical challenges associated with this.

Cost-utility analysis can capture the value of improvements in morbidity and mortality. The utilities can now be obtained from standardised and validated health status instruments, making the evidence required to inform cost-utility analysis relatively straightforward and cheap to acquire – certainly when compared with the cost of acquiring evidence on clinical effectiveness, and indeed the cost of many of the treatments being reviewed.

Cost-utility analysis makes it clear what value is attached to specific health states, and this **transparency** then allows discussion among the stakeholders involved in decision-making (for example, patients, doctors, budget holders) about the accuracy and robustness of these utilities. Cost-utility analysis thus increasingly facilitates the transparency of resource allocation processes.

* See Martin *et al* (details under 'Further reading') for more information.

† See McCabe *et al* (details under 'Further reading') for a review of the use of the cost-effectiveness threshold.

Disadvantages of cost-utility analysis

With many healthcare interventions, there are significant concerns about the ability of cost-utility analysis to capture all the valued characteristics. It is undoubtedly true that QALYs do not capture differences in the process characteristics of interventions, and there is substantial evidence that patients do attach value to these.

There is also concern that the descriptive instruments and the utilities they generate are insufficiently sensitive to differences in treatments for milder conditions. For chronic conditions, the assumption that the utility of a health state is independent of the time spent in that health state is considered problematic. Similarly, that the preceding and subsequent health states do not affect the utility of a specific health state is a strong assumption in the context of chronic conditions, especially conditions where disability accumulates over time.

In addition, there is evidence that the utility that society attaches to health is not independent of the characteristics of the person experiencing that health. Health gain for individuals in a severe health state may be valued more highly than health gain for individuals with milder health problems. Equally, it has been suggested that society would not want to penalise those with a limited capacity to benefit from healthcare when allocating resources. The current formulation of cost-utility analysis does not reflect such preferences.

What about cost-consequences analysis?

Some have argued that the limitations of cost-utility analysis are such that the best service health economists can provide to decision-makers is to describe the costs of alternative interventions and a profile of all the health and non-health impacts of the interventions – this is referred to as cost-consequences analysis.* Two problems

* See Coast (details under 'Further reading') for more information.

are associated with this approach. First, in arriving at a decision, the decision-maker will have to implicitly weigh all the different impacts and relate these to the costs, and then decide which interventions represent the best value. This process invariably takes place in the 'black box' of the decision-maker's head and will, correspondingly, be lacking in transparency. In addition, the values that will drive this decision will be the decision-maker's values – which may or may not reflect the values that society would wish to be used. Second, the quantity of information presented to a decision-maker by cost-consequences analysis will normally be considerably in excess of the volume that humans are able to reliably process, and in these circumstances it is well established that humans use short cuts to simplify the problem of decision-making. Often these short cuts are not consistent with the objectives that have been set; thus, adopting a cost-consequences approach increases the risk of poor decision-making.

Allowing perfection to be the enemy of the merely useful

There are a number of challenges to the utilisation of cost-utility analysis for resource allocation decisions. These are primarily related to the adequacy of QALYs for capturing the value that society attaches to healthcare interventions. However, it is important to remember that cost-utility analysis was developed in response to the need to help decision-makers compare the value of interventions with very different

health benefits. The relative values of cancer and blood pressure treatments still need to be compared; the measurement of neither progression-free survival nor myocardial infarctions avoided will do this. These health benefits are different and they certainly do not capture value, regardless of how sensitive values are to the factors which QALYs are criticised for being insensitive to.

Cost-utility analysis is clearly not a sufficient basis for resource allocation decisions. It fails to capture a number of factors that are potentially important and captures others with varying degrees of sensitivity; however, cost-utility analysis is useful and it performs a necessary function better than previous methods. This usefulness to decision-makers explains the rapid expansion in the utilisation of cost-utility analysis over the last decade. It would, therefore, be a pity if the numerous and well-understood imperfections were used as an argument to abandon this useful tool.

Further reading

1. Brazier JE, Ratcliffe J, Tsuchiya A, Salomon J. *Measuring and Valuing Health Benefits for Economic Evaluation*. Oxford: Oxford University Press, 2007.
2. Briggs A, Claxton K, Sculpher MJ. *Decision Modelling for Health Economic Evaluation*. Oxford: Oxford University Press, 2006.
3. Coast J. Is economic evaluation in touch with society's values? *BMJ* 2004; **329**: 1233–1236.
4. Drummond MF, Sculpher MJ, Torrance GW, O'Brien BJ, Stoddart GL. *Methods for the Economic Evaluation of Health Care Programmes*. Oxford: Oxford University Press, 2005.
5. Fallowfield L. *What is quality of life?* London: Hayward Medical Communications, 2009.
6. Martin S, Rice N, Smith PC. *Further Evidence on the Link Between Health Care Spending and Health Outcomes in England*. *CHE Research Paper 32*. York: University of York, 2007.
7. McCabe C, Claxton K, Culyer A. The NICE cost-effectiveness threshold: what it is and what that means. *Pharmacoeconomics* 2008; **26**: 733–744.

What is

cost-utility analysis?

This publication, along with the others in the series, is available on the internet at www.whatisseries.co.uk

The data, opinions and statements appearing in the article(s) herein are those of the contributor(s) concerned. Accordingly, the sponsor and publisher, and their respective employees, officers and agents, accept no liability for the consequences of any such inaccurate or misleading data, opinion or statement.

Published by Hayward Medical Communications, a division of Hayward Group Ltd.

Copyright © 2009 Hayward Group Ltd.
All rights reserved.



Supported by *sanofi-aventis*