Fenwick, E. and O'Brien, B.J. and Briggs, A. (2004) Cost-effectiveness acceptability curves - facts, fallacies and frequently asked questions. *Health Economics* 13(5):pp. 405-415.

Cost-effectiveness acceptability curves – facts, fallacies and frequently asked questions

Elisabeth Fenwick[a,]*, Bernie J. O'Brien[b,c] and Andrew Briggs[d]
[a] *Centre for Evaluation of Medicines, St Joseph's Hospital, Hamilton, Canada*
[b] *Formerly of Centre for Evaluation of Medicines, St Joseph's Hospital, Hamilton, Canada*
[c] *Formerly of Department of Clinical Epidemiology and Biostatistics, McMaster University, Hamilton, Canada*
[d] *Health Economics Research Centre, University of Oxford, Oxford, UK*

Summary

Cost-effectiveness acceptability curves (CEACs) have been widely adopted as a method to quantify and graphically represent uncertainty in economic evaluation studies of health-care technologies. However, there remain some common fallacies regarding the nature and shape of CEACs that largely result from the 'textbook' illustration of the CEAC. This 'textbook' CEAC shows a smooth curve starting at probability 0, with an asymptote to 1 for higher money values of the health outcome (λ). But this familiar 'ogive' shape which makes the 'textbook' CEAC look like a cumulative distribution function is just one special case of the CEAC. The reality is that the CEAC can take many shapes and turns because it is a graphic transformation from the cost-effectiveness plane, where the joint density of incremental costs and effects may 'straddle' quadrants with attendant discontinuities and asymptotes. In fact CEACs: (i) do not have to cut the y-axis at 0; (ii) do not have to asymptote to 1; (iii) are not always monotonically increasing in λ; and (iv) do not represent cumulative distribution functions (cdfs). Within this paper we present a 'gallery' of CEACs in order to identify the fallacies and illustrate the facts surrounding the CEAC. The aim of the paper is to serve as a reference tool to accompany the increased use of CEACs within major medical journals. Copyright

Introduction

The technique of representing uncertainty in cost-effectiveness analysis through the use of cost-effectiveness acceptability curves (CEACs) has been widely adopted as a method to quantify and graphically represent uncertainty in economic evaluation studies of health-care technologies. CEACs can now be found in major medical journals such as BMJ [1–5], New England Journal of Medicine [6,7], Lancet [8], Circulation [9], Chest [10] as well as health economics and health technology assessment journals [11–17]. Given the widespread dissemination of this technique it is important to ensure that both analysts and potential users of the information understand the nature and interpretation of these curves. This paper examines a variety of cost-effectiveness scenarios, presented on the cost-effectiveness plane, and presents a 'gallery' of CEACs with associated explanations. The paper aims to identify the facts, dispel the fallacies and thus address the frequently asked questions about CEACs. The goal of the paper is to confront the confusion surrounding the interpretation of the curves and provide a useful reference tool to accompany the

*Correspondence to: Department of Economics and Related Studies, University of York, York YO10 5DD, UK.
E-mail: ealf100@york.ac.uk

Received 2 July 2003
Accepted 2 March 2004

dispersion of CEACs within the economic evaluation literature.

First, we provide an introduction to CEACs, including a recap of the various methods for their construction. Next, we list of the common fallacies and present the facts about CEACs. Various cost-effectiveness scenarios are then examined and these facts illustrated through a 'gallery' of CEACs. Finally, some additional issues regarding the interpretation and use of CEACs are discussed and simple rules to aid interpretation are developed.

Definition of the CEACS

CEACs were originally introduced to represent the uncertainty concerning the cost-effectiveness of a health-care intervention in the context of decisions involving two interventions, as an alternative to confidence intervals around ICERs [18]. The CEAC is derived from the joint density of incremental costs (ΔC) and incremental effects (ΔE) for the intervention of interest, and represents the proportion of the density where the intervention is cost-effective for a range of values of λ. Parametric estimation is possible by assuming a parametric functional form for the joint density ($\Delta C, \Delta \mathrm{E}$) (for example, the joint normal distribution assumed by Van Hout [18]). The CEAC is then either determined directly by integrating the joint density [18–20] or estimated indirectly via parametric bootstrapping of the distribution [18]. Alternatively, the joint density can be generated by non-parametric bootstrapping [21,22], when patient level data are available for re-sampling, or by Monte Carlo simulation [23]. Where bootstrapping or Monte Carlo simulation is employed, the CEAC is determined as the proportion of the ($\Delta C, \Delta E$) points where the intervention is cost-effective [21]. Where the joint density ($\Delta C, \Delta E$) is plotted on the incremental cost-effectiveness plane [24], this proportion is easily identifiable as the proportion falling to the southeast of a ray through the origin with slope equal to λ.

CEACS – the facts and fallacies

Fallacies about CEACs

The 'textbook' illustration of the CEAC shows a smooth curve starting at probability zero with an asymptote to 1, as we consider higher money values for a health outcome (λ). However, this familiar 'ogive' shape represents just one of the possible shapes that CEACs can take. This typical example, which results from the scenario where the entire joint density ($\Delta C, \Delta E$) is contained within the northeast quadrant (where the intervention is both more costly and more effective), is largely responsible for the common fallacies regarding the shape and nature of CEACs:

(i) CEACs always cut the y-axis at 0;
(ii) CEACs always asymptote to 1;
(iii) CEACs are monotonically increasing in λ and
(iv) CEACs represent cumulative distribution functions.

Facts about CEACs

The reality is that the CEAC can take many shapes and turns because it is a graphic transformation from the cost-effectiveness plane, where the joint density of incremental costs and effects ($\Delta C, \Delta E$) may 'straddle' quadrants with attendant discontinuities and asymptotes. Where there is a non-negligible probability that the joint density extends beyond the NE quadrant the CEAC will not have the familiar shape:

(i) *Do CEACs always cut the y-axis at* 0?

In fact, by the definition of the x-axis, the point where the CEAC intersects the y-axis represents the position where the decision-maker is unwilling to pay anything for health gains ($\lambda = 0$). Generating this point involves determining the proportion of the joint density that involves cost-savings. If there is any evidence that the intervention is less costly, with any of the joint density ($\Delta C, \Delta E$) falling into the SE or SW quadrants, the CEAC will not intersect the y-axis at 0.

(ii) *Do CEACs always asymptote to* 1?

In fact, the value to which the CEAC asymptotes represents the position where the decision-maker is willing to pay an infinite amount for additional units of health gain (e.g. QALYs). Generating this point involves determining the proportion of joint density that involves health gains. If there is any evidence that the intervention is less effective, with any of the joint density ($\Delta C, \Delta E$) falling into the SW or NW quadrants, the CEAC will not asymptote to 1.

(iii) *Do CEACs increase monotonically in λ?*

In fact, the nature of the CEAC involves including the joint density $(\Delta C, \Delta E)$ that falls in the NE quadrant (more costly, more effective) as cost-effective, whilst simultaneously excluding the joint density $(\Delta C, \Delta E)$ that falls in the SW quadrant (less costly, less effective) as no longer cost-effective, as λ (the value of the willingness to trade between cost and effect) increases. Hence, the shape of the CEAC is entirely dependent upon the location of the joint density $(\Delta C, \Delta E)$ within the CE plane.

(iv) *Do CEACs represent a cumulative density function*?

Given facts (i) – (iii) CEACs cannot constitute cumulative density functions (see Appendix A).

Cost-effectiveness scenarios

Within this section we illustrate the facts through examination of a series of cost-effectiveness scenarios (Figures 1, 4 and 5) and identification of the associated CEACs (Figures 2, 3 and 6). For the purposes of illustration we assume a joint normal distribution for $(\Delta C, \Delta E)$, with a correlation between cost and effect of ± 0.5, and show the uncertainty on the cost-effectiveness plane through ellipses of equal probability covering 5, 50 and 95% of the integrated joint density. In keeping with the accepted standards for economic evaluation [25] each scenario represents a subgroup with a homogeneous treatment effect. The rules presented do not depend upon the assumption of joint normality, and are applicable regardless of the method used to determine the joint distribution $(\Delta C, \Delta E)$. It should be noted that the underlying descriptions relate to any density occupying the quadrants discussed although the particular level of the probabilities are specific to the examples provided.

Joint density occupying a single quadrant

Ellipses A–D (Figure 1) represent scenarios where there is a negligible probability that the joint density $(\Delta C, \Delta E)$ extends beyond a single quadrant, the resulting CEACs are illustrated in Figure 2.

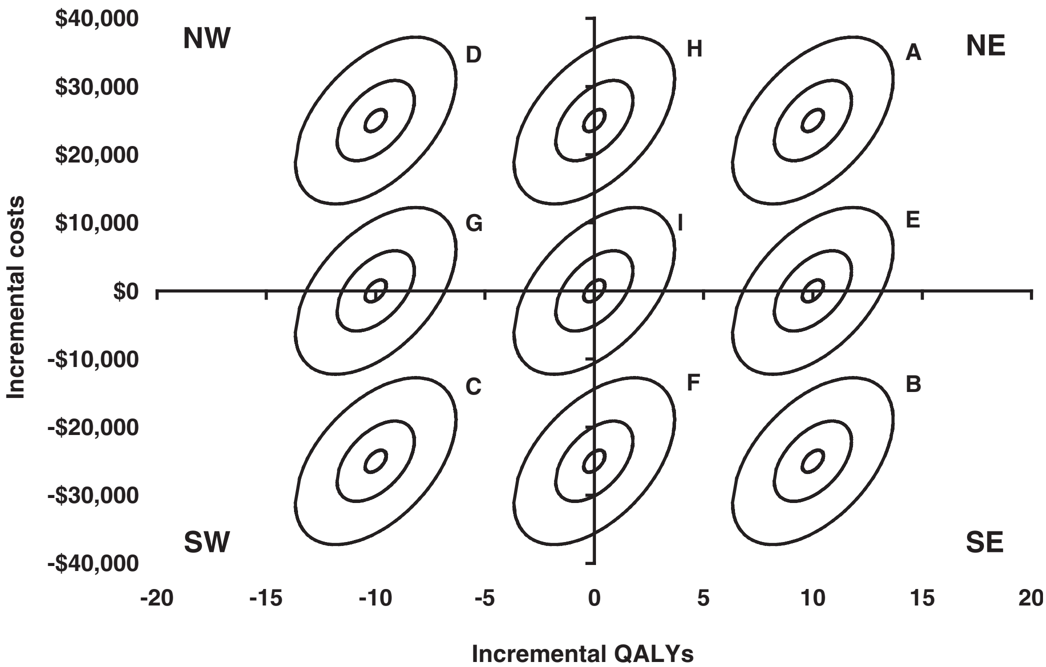

Figure 1. Nine cost-effectiveness scenarios illustrated on the incremental cost-effectiveness plane. Each ellipse represents 5, 50 and 95% of the joint distribution of cost and effect. Each example involves a correlation of 0.5

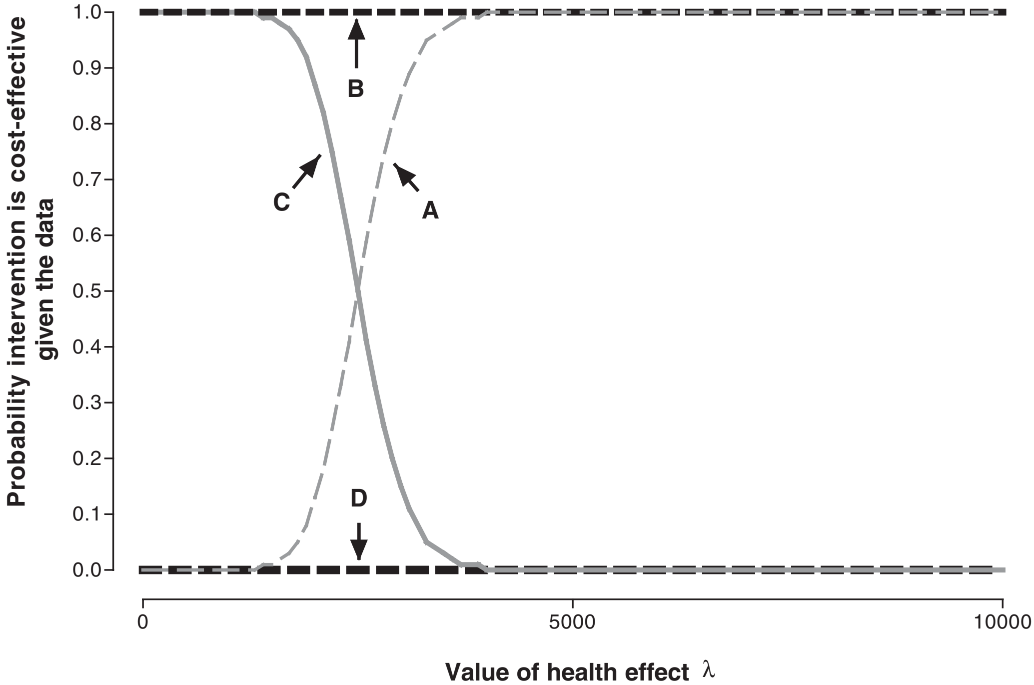

Figure 2. CEACs where there is a negligible probability that the joint density extends beyond a single quadrant of the CE plane: scenarios A–D

(A) The traditional 'textbook' case – the ellipse is contained within the NE quadrant (more costly, more effective). The CEAC cuts the y-axis at 0 because none of the density involves cost-savings. The CEAC asymptotes to 1 because the entire density involves health gains. Hence, the CEAC is an increasing function of λ.

(B) The case of dominance – the ellipse is contained within the SE quadrant (less costly, more effective). The CEAC cuts the y-axis at 1 because the entire density involves cost-savings. The CEAC asymptotes to 1 because the entire density involves health gains. Hence, the CEAC forms a horizontal line at 1.

(C) The inverse of the traditional case – the ellipse is contained within the SW quadrant (less costly, less effective). The CEAC cuts the y-axis at 1 because the entire density involves cost savings. The CEAC asymptotes to 0 because none of the density involves health gains. Hence, the CEAC is a decreasing function of λ – this scenario and the resulting CEAC is the inverse of the traditional case.[a]

(D) The dominated case – the ellipse is contained within the NW quadrant (more costly and less effective). The CEAC cuts the y-axis at 0 because none of the density involves cost-savings. The CEAC asymptotes to 0 because none of the joint density involves health gains. Hence, the CEAC forms a horizontal line at 0 – this scenario and the resulting CEAC is also the inverse of the dominance case.

Joint density occupying 2 quadrants

Ellipses E–H (Figure 1) represent scenarios where 95% of the joint density $(\Delta C, \Delta E)$ occupies two quadrants, the resulting CEACs are illustrated in Figure 3.

(E) The ellipse falls within the NE and SE quadrants (more or less costly, more effective). The CEAC does not cut the y-axis at 0 because some (here 50%) of the joint density $(\Delta C, \Delta E)$ involves cost-savings. The CEAC asymptotes to 1 because the entire density

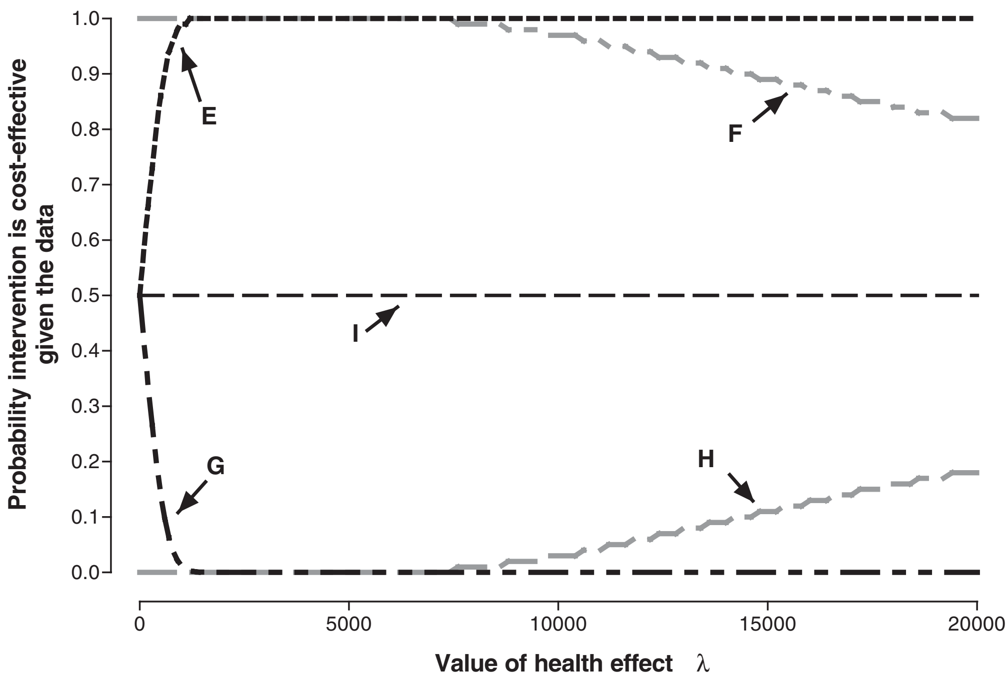

Figure 3. CEACs when the joint density lies in two quadrants of the CE plane, and a special case of all four quadrants: scenarios E–I

involves health gains. Hence, the CEAC is an increasing function of λ.

(F) The ellipse falls within the SE and SW quadrants (less costly, more or less effective). The CEAC cuts the y-axis at 1 because the entire density involves cost-savings. The CEAC asymptotes to a value less than 1 because not all of the joint density involves health gains (here only 50%). Hence, the CEAC is a decreasing function of λ.

(G) The ellipse falls within the SW and NW quadrants (more or less costly, less effective). The CEAC does not cut the y-axis at 0 because some of the joint density involves cost-savings (here 50%). The CEAC asymptotes to 0 because none of the density involves health gains. Hence, the CEAC is a decreasing function of λ.

(H) The ellipse falls within the NE and NW quadrants (more costly, more or less effective). The CEAC cuts the y-axis at 0 because none of the density involves cost-savings. The CEAC asymptotes to a value less than 1 because not all of the density involves health gains (here only 50%). Hence, the CEAC is an increasing function of λ.

Ellipses occupying all quadrants

Ellipses I–M (Figures 1, 4 and 5) represent scenarios where substantial portions of the joint density $(\Delta C, \Delta E)$ cover all four quadrants of the CE plane, the resulting CEACs are illustrated in Figures 3 and 6.

(I) Scenario I (Figure 1) illustrates a special case where the ellipse is centred upon the origin (the intervention is expected to be exactly as costly and effective as the comparator) with equal portions of the density falling into each of the quadrants. The CEAC does not cut the y-axis at 0 because some of the density involves cost-savings. The CEAC does not asymptote to 1 because not all of the density involves health gains. Here the CEAC cuts the y-axis at 0.5 because half of the joint density lies in the SW and SE quadrants, and asymptotes to 0.5 as only half of the joint density lies in the NE and SE quadrants. For this special case, an identical amount of the density is included as cost-effective (NE quadrant) as is excluded as no longer cost-effective (SW quadrant) as λ increases. As a

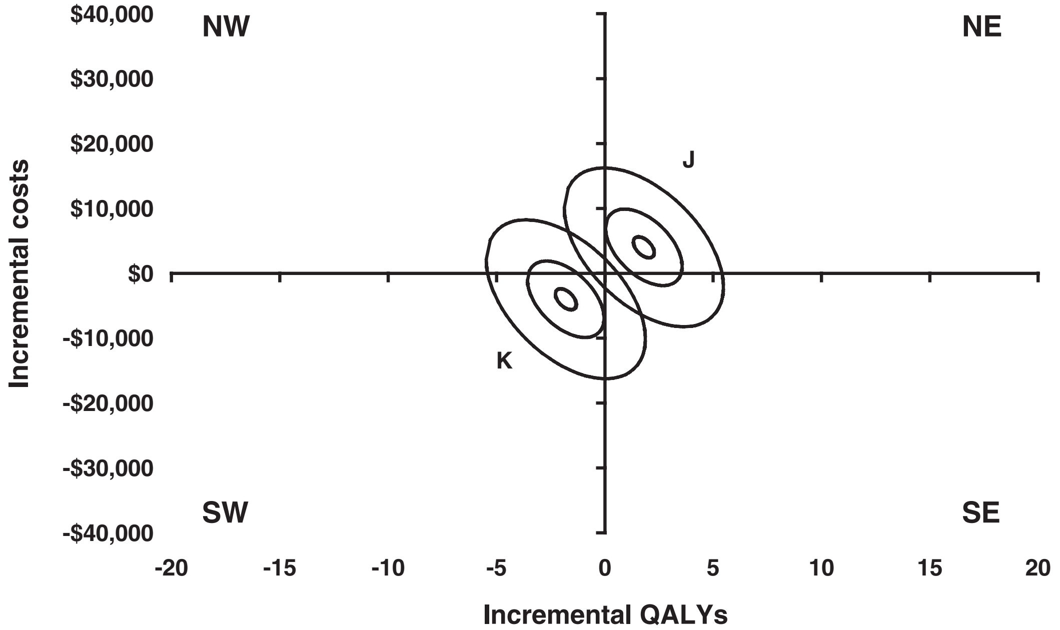

Figure 4. Two cost-effectiveness scenarios where the joint density covers all quadrants of the CE plane. Each example involves a correlation of -0.5

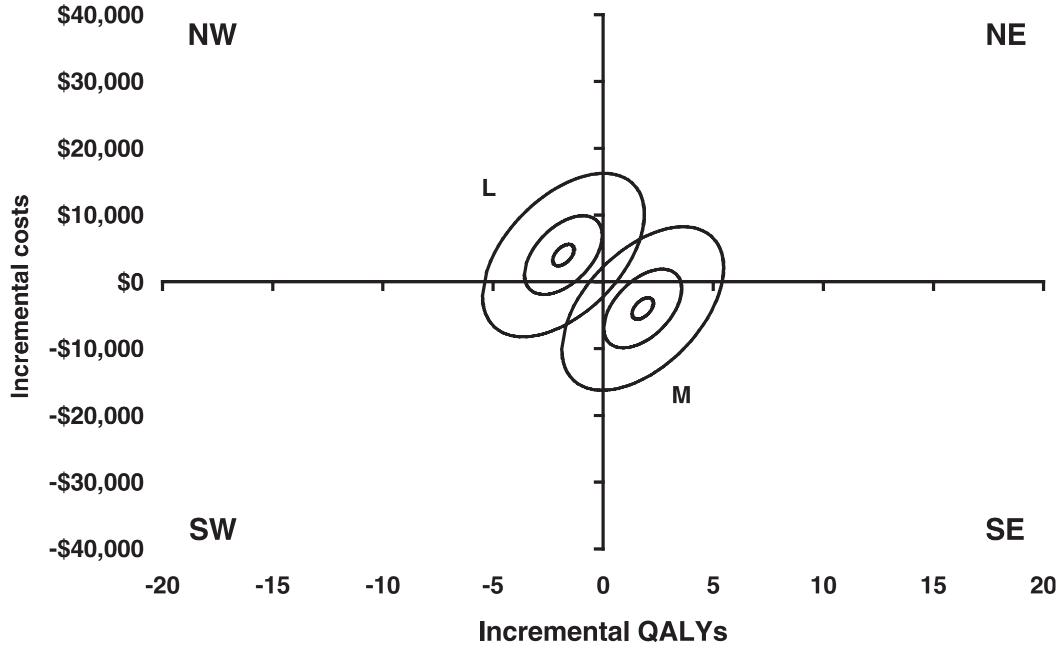

Figure 5. Further two cost-effectiveness scenarios where the joint density covers all quadrants of the CE plane. Each example involves a correlation of $+0.5$

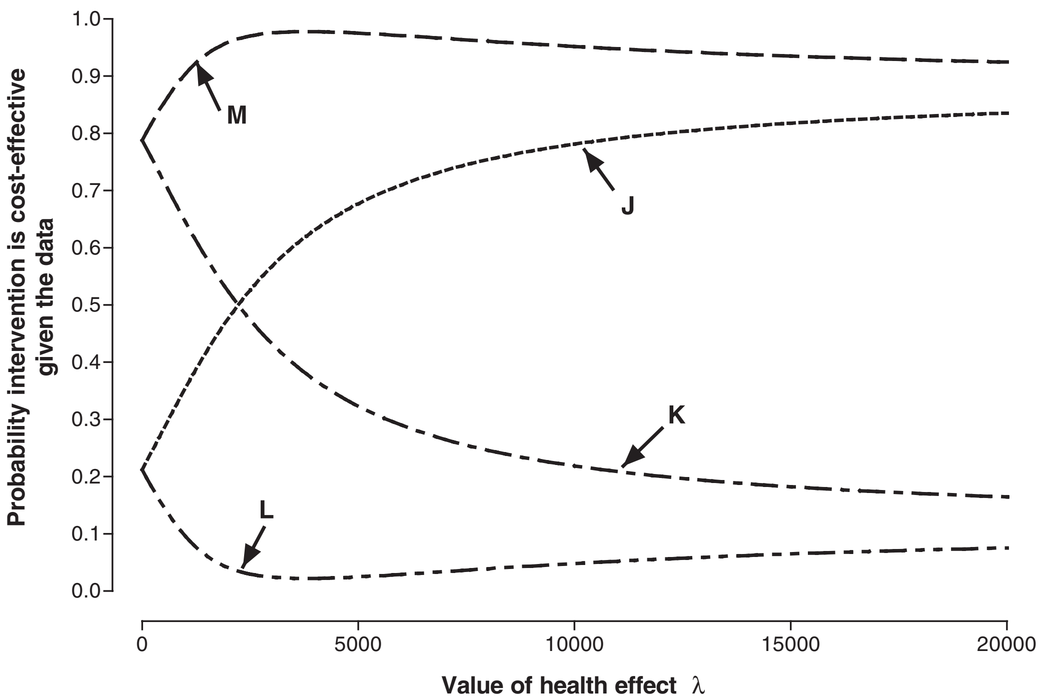

Figure 6. CEACs when the joint density covers all quadrants of the CE plane: scenarios J–M

result the CEAC forms a horizontal line through 0.5 (Figure 3). This 'bullseye' scenario is an extreme case that is unlikely to occur in practice but which provides a useful illustration of the principles of the CEAC.

(J) The ellipse falls within all of the quadrants, with the majority contained within the NE quadrant (more costly, more effective). The CEAC does not cut the y-axis at 0 because some of the density involves cost-savings (here 21%). The CEAC does not asymptote to 1 because not all of the density involves health gains (here only 88%).

(K) The ellipse falls within all of the quadrants, with the majority contained within the SW quadrant (less costly, less effective). The CEAC does not cut the y-axis at 0 because some of the density involves cost-savings (here 79%). The CEAC does not asymptote to 1 because not all of the density involves health gains (here only 12%). This scenario is the inverse of scenario J.

(L) The ellipse falls within all of the quadrants, with the majority contained within the NW quadrant (more costly, less effective). The CEAC does not cut the y-axis at 0 because some of the density involves cost-savings (here 21%). The CEAC does not asymptote to 1 because not all of the density involves health gains (here only 11%). However, the CEAC is not strictly a decreasing function of λ, due to the position of the joint density $(\Delta C, \Delta E)$ in the NE and SW quadrants. The trade-off between cost and effect implied by the joint density $(\Delta C, \Delta E)$ in the SW quadrant is lower than that implied by the joint density $(\Delta C, \Delta E)$ in the NE quadrant. Hence, as λ increases the joint density $(\Delta C, \Delta E)$ in the SW quadrant is excluded for no longer being cost-effective before the density in the NE quadrant is included as cost-effective, and the CEAC falls before rising.

(M) The ellipse falls within all of the quadrants, with the majority contained within the SE quadrant (less costly, less effective). The CEAC does not cut the y-axis at 0 because some of the density involves cost-savings (here 79%). The CEAC does not asymptote to 1 because not all of the density involves health gains (here only 89%). However, the

CEAC is not strictly an increasing function of λ, due to the position of the joint density $(\Delta C, \Delta E)$ in the NE and SW quadrants. The trade-off between cost and effect implied by the joint density $(\Delta C, \Delta E)$ in the SW quadrant is higher than that implied by the joint density $(\Delta C,\ \Delta E)$ in the NE quadrant. Hence, as λ increases the joint density in the NE quadrant is included as cost-effective before the joint density in the SW quadrant is excluded as no longer cost-effective, and the CEAC rises before falling. This scenario is the inverse of scenario L.

Some additional issues with CEACS

Bayesian or frequentist interpretation of CEACs?

Within this paper and more widely within the literature CEACs have been interpreted as representing the 'probability that the intervention is cost-effective' given the data [26–30]. It has been suggested that this interpretation is the most natural and useful as it directly addresses the issue of interest for decision-makers [26,31]. However, this interpretation is only valid in a Bayesian framework where parameters are taken to be random variables with associated probability distributions. Within the classical frequentist approach, parameters are taken as fixed, unvarying quantities. As such there is no place, within the frequentist approach, for attaching probabilities to possible values of the parameter. Instead, frequentist inference involves determining the probability (p-value) of attaining the data observed (or data more extreme than that observed) given a hypothesis regarding the parameter. It has been shown that the p-value curve, which plots the one-sided p-value on a test of positive incremental net benefit over a range of values of λ is the mirror image of the CEAC [19,26,32,33]. The p-value curve intersects the y-axis at one minus the one-sided p-value on a test of differences between costs, and asymptotes to the one-sided p-value on a test of differences between effectiveness [26,32]. However, the y-axis can only be interpreted as the 'probability that the intervention is cost-effective' when a Bayesian analysis incorporating a prior (vague or informative) has been undertaken [26].

Joint significance versus individual significance

Any ellipse that crosses the x- or y-axis is associated with a non-significant difference in costs (x-axis) and/or effects (y-axis). However, these scenarios may be associated with significant differences in cost-effectiveness for some range of λ values. For example, in scenarios E, F, G, H, L and M there are values of λ for which 95% of the joint density is considered cost-effective (scenarios E, F and M) or cost-ineffective (scenarios G, H and L) despite the lack of individual significance of cost and/or effect. However, note that this is not the case for scenarios J and K where there are no values of λ for which 95% of the joint density is cost-effective/cost-ineffective.

CEACs and dominance

The CEAC does not provide an illustration of the probability of dominance (determined by the presence and proportion of the joint density located within the SE quadrant) or the inverse, the probability of being dominated (determined by the presence and proportion of the joint density located within the NW quadrant). These probabilities are not presented on the curve as there is no value of λ for which these quadrants are isolated. However, this information may be of interest to the decision-maker and, hence, we recommend that this information be reported alongside the CEAC, either numerically or graphically as additional lines on the figure (as in Van Hout [18]).

The CEAC and value-of-information (VOI) analysis

The CEAC is a device that allows a visual impression of the joint uncertainty in the single dimension of the decision space. In particular, the CEAC overcomes the problems associated with estimating confidence intervals for cost-effectiveness ratios. A number of commentators have argued that the objective of the health system is to maximise health gain for limited resources and that, given this objective, it is the expected value of an intervention that is important for the immediate adoption decision [34,35]. As a consequence it has been argued that inference, including ranges of equivalence and benchmark error probabilities, is not useful or consistent with rational decision-making [34].

Does this mean that uncertainty is not important for decision-making and, by implication, that CEACs are redundant? The answer is of course no. Decision uncertainty affects the expected value of the decision (through the possibility that the incorrect decision is made) and hence further information, which reduces this uncertainty, has value. Formal value-of-information (VOI) techniques assess the expected costs of uncertainty surrounding a decision made on the basis of current information, in order to determine whether further research should be conducted and how this should be designed. The CEAC (or more precisely, its complement) provides the probability that the intervention of interest is not cost-effective, i.e. the probability that a decision to choose that intervention is incorrect. It is this probability that is combined with the loss function associated with incorrect decision-making to give the VOI. Further discussion of these issues, together with an introduction to the concept of the cost-effectiveness acceptability frontier (CEAF) which shows the portion of competing CEACs that have the maximum expected value, are given in Fenwick *et al.* [27]. In conclusion, the accurate assessment of uncertainty is necessary, although not sufficient, for a VOI analysis, suggesting a continuing role for the CEAC as a step on the route determining if there is value in conducting further research.

Conclusions

Through a series of cost-effectiveness scenarios we have illustrated some facts and fallacies about CEACs. We have shown that:

(i) CEACs do not always cut the y-axis at 0.
(ii) CEACs do not always asymptote to 1.
(iii) CEACs are not always monotonically increasing in λ and as a consequence.
(iv) CEACs are not cumulative distribution functions (see Appendix A).

When constructing and interpreting CEACs the simple rules to remember are:

(i) *Cost-savings and CEAC origins*. The value at which the CEAC crosses the y-axis is determined by the presence and proportion of the joint density ($\Delta C, \Delta E$) within the southerly quadrants (SW and/or SE).
(ii) *Effect gains and CEAC asymptotes*. The value to which the CEAC asymptotes is determined by the presence and proportion of the joint density (ΔC, ΔE) within the easterly quadrants (NE and/or SE).
(iii) *Position and shape*. The shape of the CEAC is determined by the location and proportion of the joint density ($\Delta C, \Delta E$) within each of the quadrants.

Acknowledgements

This paper is dedicated to our late friend, colleague and co-author Bernie O'Brien. The paper was prepared whilst Elisabeth Fenwick was a postdoctoral fellow supervised by Bernie O'Brien at the Centre for Evaluations of Medicines (CEM), funded through a fellowship from Roche, Canada. Andrew Briggs is the recipient of a Public Health Career Scientist Award from the UK Department of Health. At the time that the paper was written, Bernie O'Brien was supported by a Senior Investigator Award from the Canadian Institutes of Health Research.

Notes

a. If the position of the intervention and control were reversed the joint density would fall within the NE quadrant.

Appendix A

A cumulative distribution function (cdf) reflects the probability that a random variable X is below a certain value x. It is the integral of the probability density function (pdf) below x:

$$F(x) = P(X < x) = \int_{-\infty}^{\infty} xf(x)\,\mathrm{d}x$$

A cdf is non-decreasing (monotonically increasing) in x, that is where

$$x_1 \leq x_2 \Rightarrow F(x_1) \leqslant F(x_2)$$

In addition

$$\lim_{x\to-\infty} F(x) = 0$$

and

$$\lim_{x\to\infty} F(x) = 1$$

Here, the random variable X represents the joint density $(\Delta C, \Delta E)$ and x represents the value of λ. As described above, the construction of CEACs involves including the joint density $(\Delta C, \Delta E)$ that falls in the NE quadrant ($\Delta E > 0$, $\Delta C > 0$) as the value of λ increases, whilst simultaneously excluding the joint density $(\Delta C, \Delta E)$ that falls in the SW quadrant (where $\Delta E < 0$ and $\Delta C < 0$). Hence, the CEAC is not non-decreasing in x – as λ increases, the probability that the joint density is below λ may rise or fall. In addition, it is not true that the limit of the CEAC is 1 as λ tends to infinity. As stated above this depends upon the location and proportion of the joint density that falls within the easterly quadrants (NE and/or SE). Hence, the CEAC does not represent a cumulative density function.

References

1. Multicentre Aneurysm Screening Group. Multicentre aneurysm screening study (MASS): cost effectiveness analysis of screening for abdominal aortic aneurysms based on four year results from a randomised controlled trial. *Br Med J* 2002; **325**: 1135–1138.
2. Bower P, Byford S, Barber J *et al.* Meta-analysis of data on costs from trials of counseling in primary care: using individual patient data to overcome sample size limitations in economic analyses. *Br Med J* 2003; **326**(7401): 1247–1250.
3. Chilcott J, McCabe C, Tappenden P *et al.* Modelling the cost effectiveness of interferon beta and glatiramer acetate in the management of multiple sclerosis. Commentary: evaluating disease modifying treatments in multiple sclerosis. *Br Med J* 2003; **326**(7388): 522–524.
4. Gray A, Raikou M, McGuire A *et al.* Cost effectiveness of an intensive blood glucose control policy in patients with type 2 diabetes: economic analysis alongside randomised controlled trial (UKPDS 41). *Br Med J* 2000; **320**(7246): 1373–1378.
5. UK Prospective Diabetes Study Group. Cost effectiveness analysis of improved blood pressure control in hypertensive patients with type 2 diabetes: UKPDS 40. *Br Med J* 1998; **317**(7160): 720–726.
6. Manns BJ, Lee H, Doig CJ, Johnson D, Donaldson C. An economic evaluation of sctivated protein C treatment for severe sepsis. *New Engl J Med* 2002; **347**(13): 993–1000.
7. Ramsey SD, Berry K, Etzioni R *et al.* Cost effectiveness of lung-volume-reduction surgery for patients with severe emphysema. *New Engl J Med* 2003; **348**(21): 2092–2102.
8. Delaney BC, Wilson S, Roalfe A *et al.* Cost effectiveness of initial endoscopy for dyspepsia in patients over age 50 years: a randomised controlled trial in primary care. *Lancet* 2000; **356**(9246): 1965–1969.
9. Sheldon R, O'Brien BJ, Blackhouse G *et al.* Effect of clinical risk stratification on cost-effectiveness of the implantable cardioverter-defibrillator: the Canadian implantable defibrillator study. *Circulation* 2001; **104**(14): 1622–1626.
10. Drummond MF, Becker DL, Hux M *et al.* An economic evaluation of sequential i.v./po moxifloxacin therapy compared to i.v./po co-amoxiclav with or without clarithromycin in the treatment of community-acquired pneumonia. *Chest* 2003; **124**(2): 526–535.
11. Grieve R, Thompson S, Normand C, Suri R, Bush A, Wallis C. A cost-effectiveness analysis of rhDNase in children with cystic fibrosis. *Int J Technol Assess Health Care* 2003; **19**(1): 71–79.
12. Miller P, Chilvers C, Dewey M *et al.* Counseling versus antidepressant therapy for the treatment of mild to moderate depression in primary care: economic analysis. *Int J Technol Assess Health Care* 2003; **19**(1): 80–90.
13. Kamath CC, Kremers HM, Vanness DJ, O'Fallon WM, Cabanela RL, Gabriel SE. The cost-effectiveness of acetaminophen, NSAIDs, and selective COX-2 inhibitors in the treatment of symptomatic knee osteoarthritis. *Value Health* 2003; **6**(2): 144–157.
14. Goeree R, O'Brien BJ, Blackhouse G, Marshall J, Briggs A, Lad R. Cost-effectiveness and cost-utility of long-term management strategies for heartburn. *Value Health* 2002; **5**(4): 312–328.
15. Price MJ, Briggs AH. Development of an economic model to assess the cost effectiveness of asthma management strategies. *Pharmacoeconomics* 2002; **20**(3): 183–194.
16. Griffin AD, Perry AS, Fleming DM. Cost-effectiveness analysis of inhaled zanamivir in the treatment of influenza A and B in high-risk patients. *Pharmacoeconomics* 2001; **19**(3): 293–301.
17. Byford S, Knapp M, Greenshields J *et al.* Cost-effectiveness of brief cognitive behaviour therapy versus treatment as usual in recurrent deliberate self-harm: a decision-making approach. *Psychol Med* 2003; **33**(6): 977–986.
18. Van Hout BA, Al MJ, Gordon GS, Rutten FFH. Costs, effects and c/e-ratios alongside a clinical trial. *Health Econ* 1994; **3**: 309–319.
19. Lothgren M, Zethraeus N. Definition, interpretation and calculation of cost-effectiveness acceptability curves. *Health Econ* 2000; **9**: 623–630.
20. Lothgren M, Zethraeus N. *On the interpretation of cost-effectiveness acceptability curves.* Stockholm School of Economics Working Paper Series in Economics and Finance, 1999; (323).

21. Briggs AH, Fenn P. Confidence intervals or surfaces? Uncertainty on the cost-effectiveness plane. *Health Econ* 1998; **7**: 723–740.
22. Sculpher M, Poole L, Cleland J, Drummond M. Low doses versus high doses of angiotensin converting enzyme inhibitor lisinopril in chronic heart failure: a cost-effectiveness analysis based on the Assessment of Treatment with Lisinopril and Survival Analysis (ATLAS) study. *Eur J Heart Failure* 2000; **2**: 447–454.
23. Fenwick E, Claxton K, Sculpher M, Briggs A. *Improving the efficiency and relevance of health technology assessment*: *the role of iterative decision analytic modelling*. CHE Discussion Paper, 2000; 179.
24. Black WC. The CE plane: a graphic representation of cost-effectiveness. *Med Decision Making* 1990; **10**: 212–214.
25. Torrance GW, Siegel JE, Luce BR. Framing and designing the cost-effectiveness analysis. In *Cost-effectiveness in Health and Medicine*, Gold MR, Seigel JE, Russell LB, Weinstein MD (eds). Oxford University Press: New York, 1997.
26. Briggs AH. A Bayesian approach to stochastic cost-effectiveness analysis. *Health Econ* 1999; **8**(3): 257–261.
27. Fenwick E, Claxton K, Sculpher M. Representing uncertainty: the role of cost-effectiveness acceptability curves. *Health Econ Lett* 2001; **10**: 779–787.
28. Luce BR, Claxton K. Redefining the analytical approach to pharmacoeconomics. *Health Econo* 1999; **8**(3): 187–189.
29. O'Hagan A, Stevens JW. The probability of cost-effectiveness. *BMC Med Res Methodol* 2002; **2**(1): 5–10.
30. Heitjan DF, Moskowitz AJ, Whang W. Bayesian estimation of cost-effectiveness ratios from clinical trials. *Health Econ* 1999; **8**(3): 191–201.
31. Briggs AH, Gray AM. Handling uncertainty when performing economic evaluation of healthcare interventions. *Health Technol Assess* 1999; **3**(2).
32. O'Hagan A, Stevens JW, Montmartin J. Inference for the C/E acceptability curve and C/E ratio. *Pharmacoeconomics* 2000; **17**(4): 339–349.
33. O'Brien BJ, Briggs AH. Analysis of uncertainty in health care cost-effectiveness studies: an introduction to statistical issues and methods. *Stat Methods Med Res* 2002; **11**: 455–468.
34. Claxton K, Sculpher M, Drummond MF. A rational framework for decision making by the National Institute For Clinical Excellence (NICE). *Lancet* 2002; **360**(9334): 711–716.
35. Claxton K. The irrelevance of inference: a decision making approach to the stochastic evaluation of health care technologies. *J Health Econ* 1999; **18**: 341–364.