

How to estimate cost-effectiveness acceptability curves, confidence ellipses and incremental net benefits alongside randomised controlled trials

Nixon RM¹, Wonderling D², Grieve R³.

1. (Address for correspondence)
MRC Biostatistics Unit
Institute of Public Health
University Forvie Site
Robinson Way
Cambridge, UK
CB2 2SR
richard.nixon@mrc-bsu.cam.ac.uk
Tel: 01223 330382
Fax 01223 330388
2. National Collaborating Centre for Acute Care, Royal College of Surgeons of England, UK
3. Department of Public Health and Policy, London School of Hygiene & Tropical Medicine, UK

Keywords: CEA, cost-effectiveness acceptability curve, confidence ellipse, incremental net benefit, central limit theorem.

Summary

Decision-makers require appropriate measures of the sampling uncertainty that surround the results of cost-effectiveness analysis. This paper uses the central limit theorem to derive 95% confidence intervals around the incremental net benefit, cost-effectiveness acceptability curves and confidence ellipses when data is collected from a single randomised controlled trial. The estimation of each of these measures is illustrated using a trial based cost-effectiveness analysis. The paper provides practical guidance for future researchers to follow and is designed to encourage a more widespread and informed use of these techniques.

1 Introduction

The last decade has seen a rapid advance in the statistical methods used in cost-effectiveness analyses (CEA) conducted alongside randomised controlled trials (RCTs) [1,2]. Problems with incremental cost-effectiveness ratios (ICERs), in particular the intractability of the associated standard error, have been recognised and alternative measures developed [3-7]. A general consensus has emerged in the methodological literature, that incremental net benefits (INB), cost-effectiveness acceptability curves (CEACs) and confidence ellipses are appropriate ways of presenting the results [1,8,9]. National policy-making agencies, such as the National Institute for Health and Clinical Excellence (NICE) now insist that these measures are used as part of the formal health technology appraisal process [10]. Health economists and health service researchers therefore need to know how to present results using these methods.

Methodological papers have defined CEACs [6,9,11,12], confidence ellipses [6,13], and standard errors (SE) for INB [2,3,6,7,8,14,15]. However, none of these papers explains, in a manner accessible to readers with only a basic or intermediate knowledge of statistics, exactly *how* to estimate each of the measures described.

This paper therefore aims to make explicit the estimation of confidence intervals (CIs) for the INB, CEACs and confidence ellipses. The paper focuses on using the central limit theorem (CLT) to estimate these measures in CEA alongside RCTs. An application of these techniques is illustrated using an example from a recently published CEA.

2 Methodology

Decision makers want information on the net benefits of interventions and the uncertainty around them. The focus in this paper is on the use of the CLT to summarise the uncertainty due to sampling error that surrounds the mean estimate of net benefit. Firstly the notation used is defined (Section 2.1), and the relevance of the CLT in this context is explained (2.2). The distribution of the estimates of the population mean differences in costs and effects is described (2.3). Each of the measures of interest - the confidence intervals around the INB, the CEAC and confidence ellipses, is defined algebraically and its estimation is illustrated using a case study (2.4-2.6). An Excel spreadsheet implementing all the methodology in this paper is available to download from [16].

2.1 Notation

A trial based CEA can provide the required information on the differences in mean costs and outcomes between trial arms (the incremental costs and effects). In this context, we denote by C_{ij} a random variable that is the total cost for individual $j = 1 \dots n_i$ who is given treatment $i = 1, 2$; n_i is the number of individuals given treatment i . This is a variable that has different values that follow a population probability distribution, but which have yet to be observed. We assume that the costs for those individuals given treatment i are independently drawn from the same distribution with mean μ_{ci} and variance σ_{ci}^2 , and that

the costs from the two treatment groups are independent. E_{ij} denotes the random variable for health outcome for individual j given treatment i ; μ_{ei} and σ_{ei}^2 are the mean and variance of the distribution from which this variable is drawn.

In order to make inferences about the population mean differences in costs and in effects, ($\mu_{\Delta c} = \mu_{c2} - \mu_{c1}$ and $\mu_{\Delta e} = \mu_{e2} - \mu_{e1}$), we need firstly to estimate the parameters μ_{ci} and μ_{ei} , the population means in each treatment arm. The estimates of the mean parameters are the sample means

$$\begin{aligned}\hat{\mu}_{ci} &= \bar{C}_i = \frac{1}{n_i} \sum_{j=1}^{n_i} C_{ij} \\ \hat{\mu}_{ei} &= \bar{E}_i = \frac{1}{n_i} \sum_{j=1}^{n_i} E_{ij}\end{aligned}\tag{1}$$

2.2 Application of the central limit theorem (CLT) to CEA

The CLT states that whatever the shape of the population distributions of the costs C_{ij} and effects E_{ij} , the distributions of the sample means, \bar{C}_i and \bar{E}_i , will converge to normal distributions as n_i increases. The sampling distributions have the same means as the population distributions but each has a variance that is n_i times smaller [17]:

$$\begin{aligned}\hat{\mu}_{ci} = \bar{C}_i &\sim N\left(\mu_{ci}, \frac{\sigma_{ci}^2}{n_i}\right) \\ \hat{\mu}_{ei} = \bar{E}_i &\sim N\left(\mu_{ei}, \frac{\sigma_{ei}^2}{n_i}\right)\end{aligned}\tag{2}$$

We are interested in estimating the mean cost and effect differences between arms. Estimates of these are the differences between the sample means,

$$\begin{aligned}\hat{\mu}_{\Delta c} &= \hat{\mu}_{c2} - \hat{\mu}_{c1} = \bar{C}_2 - \bar{C}_1 = \Delta \bar{C} \\ \hat{\mu}_{\Delta e} &= \hat{\mu}_{e2} - \hat{\mu}_{e1} = \bar{E}_2 - \bar{E}_1 = \Delta \bar{E}\end{aligned}\tag{3}$$

2.3 Distribution of the estimates of the population mean differences in costs and in effects.

As the observations from the two treatment arms are assumed to be independent, and since the difference between two normal distributions is also a normal distribution, the distributions of $\hat{\mu}_{\Delta c}$ and $\hat{\mu}_{\Delta e}$ are approximately

$$\begin{aligned}\hat{\mu}_{\Delta c} &\sim N\left(\mu_{c2} - \mu_{c1}, \frac{\sigma_{c2}^2}{n_2} + \frac{\sigma_{c1}^2}{n_1}\right) = N(\mu_{\Delta c}, \sigma_{\Delta c}^2) \\ \hat{\mu}_{\Delta e} &\sim N\left(\mu_{e2} - \mu_{e1}, \frac{\sigma_{e2}^2}{n_2} + \frac{\sigma_{e1}^2}{n_1}\right) = N(\mu_{\Delta e}, \sigma_{\Delta e}^2)\end{aligned}\tag{4}$$

Here $\sigma_{\Delta c}^2$ and $\sigma_{\Delta e}^2$ are the variances of the estimated population mean cost and effect differences respectively. σ_{ci}^2 and σ_{ei}^2 $i = 1, 2$ are the variances of the distributions from which the cost and effects data from arm i are sampled. Estimates of $\sigma_{\Delta c}^2$ and $\sigma_{\Delta e}^2$ are

$$\begin{aligned}\hat{\sigma}_{\Delta c}^2 &= \frac{\hat{\sigma}_{c2}^2}{n_2} + \frac{\hat{\sigma}_{c1}^2}{n_1} \\ \hat{\sigma}_{\Delta e}^2 &= \frac{\hat{\sigma}_{e2}^2}{n_2} + \frac{\hat{\sigma}_{e1}^2}{n_1}\end{aligned}\quad (5)$$

where $\hat{\sigma}_{ci}^2$ and $\hat{\sigma}_{ei}^2$ $i = 1, 2$ are the estimates of the cost and effects distribution variance parameters in each trial arm. These are the sample variances in each trial arm:

$$\begin{aligned}\hat{\sigma}_{ci}^2 &= \frac{1}{n_i - 1} \sum_{j=1}^{n_i} (C_{ij} - \bar{C}_i)^2 \\ \hat{\sigma}_{ei}^2 &= \frac{1}{n_i - 1} \sum_{j=1}^{n_i} (E_{ij} - \bar{E}_i)^2\end{aligned}\quad (6)$$

Costs and effects will generally be correlated. The covariance between the individual costs and effects data in each treatment arm are denoted by σ_{ce1} and σ_{ce2} . These covariances are related to the correlations ρ_{ce1} and ρ_{ce2} of costs and effects data in each arm by

$$\begin{aligned}\sigma_{ce1} &= \rho_{ce1} \sigma_{c1} \sigma_{e1} \\ \sigma_{ce2} &= \rho_{ce2} \sigma_{c2} \sigma_{e2}\end{aligned}\quad (7)$$

So if there is no correlation between costs and effects then the covariance is also zero.. The CLT leads us to an approximate distribution for $\hat{\mu}_{\Delta c}$ and $\hat{\mu}_{\Delta e}$ for large values of n_i

$$\begin{pmatrix} \hat{\mu}_{\Delta c} \\ \hat{\mu}_{\Delta e} \end{pmatrix} \sim \text{BVN} \left(\begin{pmatrix} \mu_{\Delta c} \\ \mu_{\Delta e} \end{pmatrix}, \begin{pmatrix} \sigma_{\Delta c}^2 & \sigma_{\Delta ce} \\ \sigma_{\Delta ce} & \sigma_{\Delta e}^2 \end{pmatrix} \right)\quad (8)$$

where $\sigma_{\Delta ce}$ is the covariance between $\hat{\mu}_{\Delta c}$ and $\hat{\mu}_{\Delta e}$ and BVN denotes a bivariate normal distribution [18]. We need to estimate $\sigma_{\Delta ce}$, the covariance between the estimates of the population mean cost and effects difference. In Appendix A we show that

$$\sigma_{\Delta ce} = \text{Cov}(\hat{\mu}_{\Delta c}, \hat{\mu}_{\Delta e}) = \frac{\sigma_{ce2}}{n_2} + \frac{\sigma_{ce1}}{n_1}\quad (9)$$

That is, the covariance between the estimates of the population mean cost and effect difference is the sum of the covariances between the cost and effects data in each arm divided by the respective sample sizes. An estimate of this covariance is therefore

$$\hat{\sigma}_{\Delta ce} = \frac{\hat{\sigma}_{ce2}}{n_2} + \frac{\hat{\sigma}_{ce1}}{n_1} \quad (10)$$

Estimates of the covariances σ_{ce1} and σ_{ce2} are the sample covariances of the individual data

$$\hat{\sigma}_{cei} = \frac{1}{n_i - 1} \sum_{j=1}^{n_i} (C_{ij} - \bar{C}_i)(E_{ij} - \bar{E}_i) \quad (11)$$

The correlation between the estimate of the population mean costs and effects difference is defined as

$$\rho_{\Delta ce} = \frac{\sigma_{\Delta ce}}{\sigma_{\Delta c} \sigma_{\Delta e}} \quad (12)$$

and is estimated by

$$\hat{\rho}_{\Delta ce} = \frac{\hat{\sigma}_{\Delta ce}}{\hat{\sigma}_{\Delta c} \hat{\sigma}_{\Delta e}} \quad (13)$$

The next sections describe how the CLT can be applied to estimate each of the measures of interest in CEA, and is illustrated using a CEA based on an RCT. The trial was designed to evaluate the cost-effectiveness of acupuncture in the management of chronic headache compared with usual care [19]. The key statistics are presented in Table 1.

2.4 Incremental net benefit (INB)

To summarise the results of a CEA the incremental net (monetary) benefit (INB) of one treatment compared to another may be reported, together with a 95% CI. The INB is defined as

$$\text{INB}(K) = K\mu_{\Delta e} - \mu_{\Delta c} \quad (14)$$

where K represents the decision-makers willingness to pay for a one unit gain in health outcome. Thus the new treatment is cost-effective if and only if $\text{INB}(K) > 0$. The value of K is generally unknown, so it is usual to plot the estimated value of $\text{INB}(K)$, for various values of K . $\mu_{\Delta c}$ and $\mu_{\Delta e}$ are estimated by $\hat{\mu}_{\Delta c}$ and $\hat{\mu}_{\Delta e}$. The expected value and variance of $\hat{\text{INB}}(K) = K\hat{\mu}_{\Delta e} - \hat{\mu}_{\Delta c}$ are given by

$$\begin{aligned} E[\hat{\text{INB}}(K)] &= KE[\hat{\mu}_{\Delta e}] - E[\hat{\mu}_{\Delta c}] \\ &= K\mu_{\Delta e} - \mu_{\Delta c} \\ \text{Var}[\hat{\text{INB}}(K)] &= K^2 \text{Var}[\hat{\mu}_{\Delta e}] + \text{Var}[\hat{\mu}_{\Delta c}] - 2K \text{Cov}[\hat{\mu}_{\Delta e}, \hat{\mu}_{\Delta c}] \\ &= K^2 \sigma_{\Delta e}^2 + \sigma_{\Delta c}^2 - 2K \rho_{\Delta ce} \sigma_{\Delta e} \sigma_{\Delta c} \end{aligned} \quad (15)$$

The means, variances and correlation of the population mean differences can then be estimated using the formulae given in Section 2.3 above.

From Section 2.3, we know that $\hat{\mu}_{\Delta_c}$ and $\hat{\mu}_{\Delta_e}$ approximately have a bivariate normal distribution. Furthermore, a linear combination of normal variables is also normal, so

$$\hat{\text{INB}}(K) \sim N\left(E[\hat{\text{INB}}(K)], \text{Var}[\hat{\text{INB}}(K)]\right) \quad (16)$$

and therefore CIs can be constructed in the usual way:

$$100(1-\alpha)\% \text{ CI} = K\hat{\mu}_{\Delta_e} - \hat{\mu}_{\Delta_c} \pm z_{1-\alpha/2} \sqrt{K^2 \hat{\sigma}_{\Delta_e}^2 + \hat{\sigma}_{\Delta_c}^2 - 2K\hat{\rho}_{\Delta_{ce}} \hat{\sigma}_{\Delta_e} \hat{\sigma}_{\Delta_c}} \quad (17)$$

where $z_{1-\alpha/2}$ is the value of the standard normal distribution such that 100(1- α)% of the area falls within $\pm z_{1-\alpha/2}$. If $\alpha = 0.05$, then $z_{0.975} = 1.96$ and a 95% CI is constructed.

Applying the above to the case study, we can see how in this example the estimated incremental net benefit, $\hat{\text{INB}}(K)$, increases as the threshold value of a QALY, K , increases (Table 2).

The INB is greater the more value we place on each QALY gained as $\hat{\mu}_{\Delta_e}$ is positive in this case. The confidence intervals for the INB become wider as K increases – this is because the uncertainty around the value of the effect is magnified by the value of K placed on the effect.

2.5 Cost-effectiveness acceptability curve (CEAC)

Ultimately, we are interested in assessing whether one treatment is cost-effective compared to the other, that is, assessing whether the INB is positive. As the population mean costs and effects are not known exactly, but are estimated from a sample, we can never be certain that the INB is positive for a particular value of K . The CEAC is the estimated probability that $\hat{\text{INB}}(K) > 0$ plotted against K (see Figure 1) and is defined as:

$$\begin{aligned} Q(K) &= P(\hat{\text{INB}}(K) > 0) \\ &= P\left(\frac{\hat{\text{INB}}(K) - E[\hat{\text{INB}}(K)]}{\sqrt{\text{Var}[\hat{\text{INB}}(K)]}} > \frac{-E[\hat{\text{INB}}(K)]}{\sqrt{\text{Var}[\hat{\text{INB}}(K)]}}\right) \end{aligned} \quad (18)$$

Under the CLT, the left hand side of this inequality has a standard normal $N(0,1)$ distribution so

$$\begin{aligned} Q(K) &= 1 - \Phi\left(\frac{-E[\hat{\text{INB}}(K)]}{\sqrt{\text{Var}[\hat{\text{INB}}(K)]}}\right) \\ &= \Phi\left(\frac{K\mu_{\Delta_e} - \mu_{\Delta_c}}{\sqrt{K^2 \sigma_{\Delta_e}^2 + \sigma_{\Delta_c}^2 - 2K\rho_{\Delta_{ce}} \sigma_{\Delta_e} \sigma_{\Delta_c}}}\right) \end{aligned} \quad (19)$$

where $\Phi(z)$ is the cumulative density function of a standard normal distribution, that is the probability that a $N(0,1)$ variables is less than z . Again the means, variances and

correlation of the population mean differences are estimated as in 2.3 above. The calculation is repeated for various values of K. From a frequentist perspective, the CEAC is one minus the p-value of a one-sided test $H_1: \text{INB}(K) > 0$, plotted against K. The probability that the $\text{INB}(K)$ is positive is a Bayesian concept - it is the posterior probability that the $\text{INB}(K)$ is positive given the data. From a Bayesian perspective the population parameters $\mu_{\Delta_c}, \mu_{\Delta_e}, \sigma_{\Delta_c}^2, \sigma_{\Delta_e}^2$ and ρ have distributions.

The CEAC for the case study is presented in Table 2 and Figure 1.

2.6 Confidence ellipses

To understand the uncertainty around the estimates of the population mean differences in costs and in effects it is useful to plot the point estimates on the cost-effectiveness plane, together with confidence ellipses. The frequentist definition of these is that if one were to repeat the experiment a large number of times, then the true incremental cost and incremental effect would lie in the ellipses generated this way 50%, 75% or 95% of the time (see Figure 2). In the Bayesian paradigm the credible ellipses are regions with a 50%, 75% or 95% probability of containing the mean incremental cost and incremental effect.

A convenient way of defining an ellipse is with the following two equations.

$$\begin{aligned}\Delta_c &= \sqrt{-2 \log(1 - \alpha)} \sigma_{\Delta_c} \cos\left(\theta - \frac{\arccos(\rho_{\Delta_{ce}})}{2}\right) + \mu_{\Delta_c} \\ \Delta_e &= \sqrt{-2 \log(1 - \alpha)} \sigma_{\Delta_e} \cos\left(\theta + \frac{\arccos(\rho_{\Delta_{ce}})}{2}\right) + \mu_{\Delta_e}\end{aligned}\tag{20}$$

As before, $\alpha = 0.05$ corresponds to a 95% confidence ellipse. As the angle θ is varied, taking on values between 0 and 2π radians, the ellipse will be traced out by Δ_c and Δ_e . Explanation of how these equations are derived is given in Appendix B. In practice we do not know these parameters, so the means, variances and correlation of the population mean differences are estimated by the estimates as in 2.3 above.

To calculate confidence ellipses, co-ordinates can be estimated using the equations above for a number of values of θ between 0 and 2π radians (Table 3). When these points are plotted on the cost-effectiveness plane they form ellipses. The estimated correlation between incremental cost and QALYs gained determines the orientation of the ellipses; in this case there is a modest inverse correlation between the mean incremental cost and effect.

3 Discussion

This paper uses the CLT to derive and define 95% CI around the INB, CEACs, and confidence ellipses. Applying the CLT provides a robust, practical way of estimating the sampling uncertainty surrounding the results of trial-based CEA. The estimation of each measure has been illustrated using a trial-based CEA with a link provided to the spreadsheet used [16]. The outcome of interest in a CEA is the mean incremental cost-effectiveness for the population concerned. Cost data in particular may be drawn from

populations with highly skewed distributions and an advantage of the CLT is that it avoids assuming that the data are sampled from a normal distribution, and will give an asymptotically unbiased estimate of the population mean. However, this raises the question of what sample size is required to invoke the CLT? As Cochran [20] points out that there are no hard and fast rules for determining how large the sample size should be before assuming the CLT is reasonable. The author suggests though that, when the main deviation from normality of the population distribution is due to positive skewness, the sample size should be larger than $25\eta^2$, where η is the empirical skewness coefficient. One approach would therefore be to extend sample size calculations for CEA alongside RCTs to allow for the skewed nature of cost data, thus making it more reasonable to apply the CLT. An alternative would be to use the non-parametric bootstrap [21]. However, as O'Hagan and Stevens [22] explain the non-parametric bootstrap also relies on asymptotic assumptions.

The methods presented apply to circumstances where the baseline characteristics of patients in each intervention group are similar. However in other contexts, particularly CEA alongside observational studies or non-randomised trials, the methods described may need to be extended to adjust for baseline characteristics. Different approaches have been described that allow the statistics described in this paper to be estimated, adjusting for baseline differences. The approach developed by Hoch et al [23] estimates individual net-benefits for each patient in a trial for a given level of K and then uses ordinary least squares (OLS) regression to estimate the INB and associated measures of uncertainty. This has the disadvantage of using a single set of covariates to adjust for baseline differences. Instead, the mean cost and effectiveness differences can be estimated separately using a system of seemingly unrelated regression equations [24]. This method enables different covariates to be used when estimating costs and effects. Unlike the methods presented in this paper both approaches assume that costs and effects are drawn from normal distributions. A more flexible method for adjusting for baseline differences is presented by Nixon and Thompson [25] whose methods consider costs and effects jointly while allowing for a range of different distributions (e.g. gamma or log normal).

In conclusion, this paper derives recommended measures for presenting sampling uncertainty in CEA by using the CLT. The approach described and illustrated provides practical guidance for researchers to follow and can help ensure a more widespread and informed use of these techniques.

Appendix

A Derivation of the covariance between the estimates of population mean cost and effects difference

$$\begin{aligned}
 \sigma_{\Delta ce} &= \text{Cov}(\hat{\mu}_{\Delta c}, \hat{\mu}_{\Delta e}) \\
 &= \text{Cov}(\Delta \bar{C}, \Delta \bar{E}) \\
 &= \text{Cov}(\bar{C}_2 - \bar{C}_1, \bar{E}_2 - \bar{E}_1)
 \end{aligned} \tag{21}$$

The two treatment arms are still independent, even though costs and effects within a treatment arm are correlated, so

$$\begin{aligned}
 \sigma_{\Delta ce} &= \text{Cov}(\bar{C}_2, \bar{E}_2) + \text{Cov}(\bar{C}_1, \bar{E}_1) \\
 &= \frac{1}{n_2} \sum_{j=1}^{n_2} \text{Cov}(C_{2j}, E_{2j}) + \frac{1}{n_1} \sum_{j=1}^{n_1} \text{Cov}(C_{1j}, E_{1j})
 \end{aligned} \tag{22}$$

as $\text{Cov}(C_{2j}, E_{2j}) = \sigma_{ce2}$ for all j and $\text{Cov}(C_{1j}, E_{1j}) = \sigma_{ce1}$ for all j then

$$\sigma_{\Delta ce} = \frac{\sigma_{ce2}}{n_2} + \frac{\sigma_{ce1}}{n_1} \tag{23}$$

B Derivation of confidence ellipses

If $\hat{\mu}_{\Delta c}$ and $\hat{\mu}_{\Delta e}$ have a bivariate normal distribution

$$\begin{pmatrix} \hat{\mu}_{\Delta c} \\ \hat{\mu}_{\Delta e} \end{pmatrix} \sim \text{BVN} \left(\begin{pmatrix} \mu_{\Delta c} \\ \mu_{\Delta e} \end{pmatrix}, \begin{pmatrix} \sigma_{\Delta c}^2 & \rho_{\Delta ce} \sigma_{\Delta c} \sigma_{\Delta e} \\ \rho_{\Delta ce} \sigma_{\Delta c} \sigma_{\Delta e} & \sigma_{\Delta e}^2 \end{pmatrix} \right) \tag{24}$$

then

$$\frac{1}{1 - \rho_{\Delta ce}^2} \left(\left(\frac{\Delta \bar{C} - \mu_{\Delta c}}{\sigma_{\Delta c}} \right)^2 - 2\rho \left(\frac{\Delta \bar{C} - \mu_{\Delta c}}{\sigma_{\Delta c}} \right) \left(\frac{\Delta \bar{E} - \mu_{\Delta e}}{\sigma_{\Delta e}} \right) + \left(\frac{\Delta \bar{E} - \mu_{\Delta e}}{\sigma_{\Delta e}} \right)^2 \right) \tag{25}$$

has a chi squared distribution with 2 degrees of freedom [26]. A $100\alpha\%$ confidence area is contained by the ellipse formed by setting this expression to equal k , the critical value

of a chi squared distribution corresponding to $100\alpha\%$. As $P(\chi_2^2 < k) = 1 - e^{-\frac{k}{2}}$ then $k = -2\log(1 - \alpha)$. So the equation for the ellipse becomes

$$\left(\frac{\Delta_c - \mu_{\Delta c}}{\sigma_{\Delta c}} \right)^2 - 2\rho \left(\frac{\Delta_c - \mu_{\Delta c}}{\sigma_{\Delta c}} \right) \left(\frac{\Delta_e - \mu_{\Delta e}}{\sigma_{\Delta e}} \right) + \left(\frac{\Delta_e - \mu_{\Delta e}}{\sigma_{\Delta e}} \right)^2 = -2(1 - \rho_{\Delta ce}^2) \log(1 - \alpha) \tag{26}$$

Practically, this equation will be difficult to use as part of a computer program for plotting the confidence ellipses, so we reparameterise it in terms of a variable $\theta \in [0, 2\pi)$

$$\begin{aligned}
\Delta_c &= \sqrt{-2 \log(1-\alpha)} \sigma_{\Delta c} \cos\left(\theta - \frac{\arccos(\rho_{\Delta ce})}{2}\right) + \mu_{\Delta c} \\
\Delta_e &= \sqrt{-2 \log(1-\alpha)} \sigma_{\Delta e} \cos\left(\theta + \frac{\arccos(\rho_{\Delta ce})}{2}\right) + \mu_{\Delta e}
\end{aligned}
\tag{27}$$

References

1. Briggs AH (Ed). *Statistical methods for cost-effectiveness research: a guide to current issues and future developments*. OHE: London, 2003.
2. O'Hagan A, Stevens JW. A framework for cost-effectiveness analysis from clinical trial data. *Health Econ* 2001; **10**: 300-315.
3. Willan AR and Lin DY. Incremental net benefit in randomised clinical trials, *Stat Med* 2001; **20**: 1563-1574.
4. Chaudhary MA, Stearns SC. Confidence intervals for cost-effectiveness ratios: an example from a randomised trial. *Stat Med* 1996; **15**: 1447-1458.
5. Briggs AH, Mooney CZ, Wonderling DE. Constructing confidence intervals for cost-effectiveness ratios: an evaluation of parametric and non-parametric techniques using monte carlo simulation. *Stat Med* 1999; **18**: 3245-3262.
6. Van Hout BA, Al MJ, Gordon GS, Rutten FFH. Costs, effects and C/E ratios. *Health Econ* 1995; **4**: 373-381.
7. Stinnett AA and Mullahy J. Net health benefits: a new framework for the analysis of uncertainty in cost-effectiveness analysis. *Med Decis Making* 1998; **18** suppl: S68-80.
8. Willan AR, Chen EB, Cook RJ, Lin DY. Incremental net benefit in randomised clinical trials with quality-adjusted survival. *Stat Med* 2003; **22**: 353-362.
9. Fenwick E, O'Brien BJ, Briggs A. Cost-effectiveness acceptability curves - facts, fallacies and frequently asked questions *Health Econ* 2004; **13**: 405-415.
10. National Institute for Clinical Excellence (NICE). *Guide to the methods of Technology Appraisal*. NICE: London, 2004.
11. Lothren M, Zethraues N. Definition, interpretation and calculation of cost-effectiveness acceptability curves. *Health Econ* 2000; **9**: 623-630.
12. Fenwick E, Claxton K, Sculpher M. Representing uncertainty: the role of cost-effectiveness acceptability curves. *Health Econ* 2001; **10**: 779-787.
13. Briggs A, Fenn P. Confidence intervals or surfaces? Uncertainty on the cost-effectiveness plane. *Health Econ* 1998; **7**: 723-740.
14. Tambour M, Zethraeus N, Johannesson M. A note on confidence intervals in cost-effectiveness analysis. *Int J Technol Assess Health Care* **14**: 467-471.
15. Laska EM, Meisner M, Siegel C, Wanderling J. Statistical cost-effectiveness analysis of two treatments based on net health benefits. *Stat Med* 2001; **20**: 1279-1302.
16. <http://www.mrc-bsu/personal/richard/downloads.html>
17. Feller W. *An introduction to probability theory and its applications*, Volume 1, New York ; Chichester : Wiley, 1968.

18. DeGroot MH. *Probability and Statistics*. 2nd Ed. Reading, Mass. ; Wokingham : Addison-Wesley, 1989]
19. Wonderling D, Vickers AJ, Grieve R, and McCarney R. Cost effectiveness analysis of a randomised trial of acupuncture for chronic headache in primary care *BMJ* 2004; **328**: 747–749
20. Cochran WG. *Sampling techniques*. 3rd ed. New York ; London : Wiley 1977.
21. Thompson SG, Barber JA. How should cost data in pragmatic randomised trials be analysed? *BMJ* 2000; **320**: 1197-1200.
22. O'Hagan A, Stevens JW. Assessing and comparing costs: how robust are the bootstrap and methods based on asymptomatic normality? *Health Econ* 2003; **12**: 33-49.
23. Hoch JS, Briggs AH, Willan AR. Something old, something new, something borrowed, something blue: a framework for the marriage of health econometrics and cost-effectiveness analysis. *Health Econ* 2002; **11**: 415-430.
24. Willan AR, Briggs AH, Hoch JS. Regression methods for covariate adjustment and subgroup analysis for non-censored cost-effectiveness data. *Health Econ* 2004 **13**: 461-475.
25. Nixon RM, Thompson SG. Incorporating covariate adjustment, subgroup analysis and between-centre differences into cost-effectiveness evaluations. *Health Econ* 2005 (in press).
26. Kotz S, Balakrishnan, N, Johnson NL. *Continuous multivariate distributions Vol. 1, Models and applications*. 2nd ed. New York ; Chichester : Wiley, 2000.

Table 1 – Calculation of statistics for the case study:

	Treatment Group (Acupuncture)		Control Group (Usual Care)		Difference	
	Eq	Estimate	Eq	Estimate	Eq	Estimate
sample size		n_2 136		n_1 119		n_1+n_2 255
<u>Effect (QALYs)</u>						
mean	1	$\hat{\mu}_{e2}$ 0.727	1	$\hat{\mu}_{e1}$ 0.708	3	$\hat{\mu}_{\Delta e}$ 0.018
variance of mean		$\frac{\hat{\sigma}_{e2}^2}{n_2}$ 0.00010		$\frac{\hat{\sigma}_{e1}^2}{n_1}$ 0.00011	5	$\hat{\sigma}_{\Delta e}^2$ 0.00021
<u>Cost (£)</u>						
mean	1	$\hat{\mu}_{c2}$ 403	1	$\hat{\mu}_{c1}$ 217	3	$\hat{\mu}_{\Delta c}$ 186
variance of mean		$\frac{\hat{\sigma}_{c2}^2}{n_2}$ 936		$\frac{\hat{\sigma}_{c1}^2}{n_1}$ 1,985	5	$\hat{\sigma}_{\Delta c}^2$ 2,920
<u>cost and effect</u>						
covariance	7, 11	$\hat{\sigma}_{ce2}$ -9.250	7, 11	$\hat{\sigma}_{ce1}$ -12.578	10	$\hat{\sigma}_{\Delta ce}$ -0.174
correlation		$\hat{\rho}_{ce2}$ -0.218		$\hat{\rho}_{ce1}$ -0.231	13	$\hat{\rho}_{\Delta ce}$ -0.222
Incremental cost-effectiveness ratio (£ per QALY)						$\frac{\hat{\mu}_{\Delta c}}{\hat{\mu}_{\Delta e}}$ 10,088

Table 2 - Case study: Incremental net benefit of Acupuncture vs Usual care

K	E(INB)	Var(INB)	95% CI for INB		Q
Eq	15	15	17	17	19
£0	-£186	£2,920	-£292	-£80	0.000
£5,000	-£94	£9,888	-£289	£101	0.172
£10,000	-£2	£27,315	-£326	£322	0.496
£15,000	£91	£55,204	-£370	£551	0.650
£20,000	£183	£93,552	-£417	£782	0.725
£25,000	£275	£142,361	-£464	£1,015	0.767
£30,000	£368	£201,631	-£513	£1,248	0.793
£35,000	£460	£271,361	-£561	£1,481	0.811
£40,000	£552	£351,551	-£610	£1,714	0.824
£45,000	£644	£442,202	-£659	£1,948	0.834
£50,000	£737	£543,313	-£708	£2,181	0.841

Table 3 Case study: Confidence ellipses for Acupuncture vs Usual care*

Radians	95%		75%		50%	
	Incremental Effect	Incremental Cost	Incremental Effect	Incremental Cost	Incremental Effect	Incremental Cost
	Δ_e	Δ_c	Δ_e	Δ_e	Δ_e	Δ_e
0	0.041	£269	0.033	£242	0.029	£226
0.2 π	0.053	£192	0.042	£190	0.035	£189
0.4 π	0.052	£113	0.041	£137	0.034	£151
0.6 π	0.038	£62	0.032	£102	0.028	£127
0.8 π	0.017	£59	0.017	£99	0.018	£125
1.0 π	-0.004	£104	0.003	£130	0.008	£147
1.2 π	-0.016	£180	-0.005	£182	0.002	£183
1.4 π	-0.015	£259	-0.004	£236	0.003	£221
1.6 π	-0.001	£310	0.005	£270	0.009	£246
1.8 π	0.020	£314	0.020	£273	0.019	£248
2.0 π	0.041	£269	0.033	£242	0.029	£226

* See equation 20

Figure 1 - Case Study: CEAC for Acupuncture vs Usual Care




Figure 2 - Case study: Confidence ellipses for Acupuncture vs Usual care

