UNCERTAINTY IN COST-EFFECTIVENESS ANALYSES: PROBABILISTIC UNCERTAINTY ANALYSIS AND STOCHASTIC LEAGUE TABLES

Rob M.P.M. Baltussen
Raymond C.W. Hutubessy
David .B. Evans
Christopher JL Murray

GPE Discussion Paper Series: No.34

EIP/GPE/EQC
World Health Organization

Abstract

As more prospective cost-effectiveness analyses (CEA) are undertaken, providing stochastic data on costs and effects, interest has grown in the application of standard statistical inferential techniques to the calculation of cost-effectiveness ratios (CER). However, it is important to temper the enthusiasm for stochastic methods by recognizing that it is very difficult to undertake prospective, controlled trials of many public health interventions which means that individual level data will not be available in many cases. We propose a method of taking uncertainty into account in the analysis involving probabilistic uncertainty analysis using Monte Carlo Simulations. This could be used in combination with non-parametric bootstrapping techniques where appropriate.

This paper also discusses how information on uncertainty should be communicated to policy makers. Current approaches to uncertainty analysis in CEA present study results in terms of some type of uncertainty interval, and little or no attention is paid to the question of how decision makers should interpret the results where uncertainty intervals overlap. We show how the choice between interventions can be facilitated with reference to a 'stochastic league table'. Such tables inform decision-makers about the probability that a specific intervention would be included in the optimal mix of interventions for various budget levels, given the uncertainty surrounding the intervention. This information helps decision makers decide on the relative attractiveness of different intervention mixes, and also on the implications for trading gains in efficiency for gains in other goals such as reducing health inequalities and increasing health system responsiveness.

Introduction

As more prospective cost-effectiveness analysis (CEA) studies are undertaken, providing stochastic data on costs and effects, interest has grown in the application of statistical techniques to the calculation of cost-effectiveness ratios (CER). Several methods have been developed including confidence planes (1), mathematical techniques (2), and the net health benefit approach (3).

However, it is important to recognize that many public health interventions do not lend themselves to the collection of sampled individual level data (by patient, health facility, region etc.), especially in a developing country context. For example, it is difficult to develop a feasible experimental design to identify the costs and effects of a national radio health education program, or a policy to subsidize the use of essential pharmaceutical products. Many economic evaluations require non-stochastic parameter estimates and modeling assumptions.

Typically, uncertainty stemming from the use of such non-sampled secondary data sources in CEA has been dealt with by sensitivity analysis[1] (4,5). These deterministic analyses draw inferences from point estimates of variables but interpretation is conditional upon a range of uncertainty that is assumed for critical variables. There are three major limitations to this approach. Firstly, the analyst has discretion as to which variables and what alternative values are included. Secondly, interpretation is essentially arbitrary because there are no comprehensive guidelines or standards as to what degree of variation in results is acceptable evidence that the analysis is robust. Thirdly, variation of uncertain parameters one at a time carries a risk that interactions between parameters may not be captured (6).

This paper examines the application of probabilistic uncertainty analysis with Monte Carlo simulations in this context. This builds on work already described in the literature (4,7-9) and requires that analysts assume some distributional form for costs and effects from which repeated samples are drawn to determine a distribution for the CER. The definition of an uncertainty range for CER is hampered by the instability of sample estimates of CERs, causing its mean to vary (13). The present paper applies the simple percentile method – usually applied to estimate uncertainty ranges for CER in non-parametric bootstrapping – to estimate uncertainty intervals for CERs involving probabilistic uncertainty analysis. This could be used in combination with non-parametric bootstrapping techniques where appropriate.

---

[1] There is confusion in the literature as to the definition of sensitivity analysis on the one hand and uncertainty analysis on the other hand. We argue that sensitivity analysis refers to uncertainty about social choices such as the discount rate or the inclusion of productivity costs. Uncertainty analysis refers to variation in the distribution of costs and effects (stemming from either non-sampled or sampled data). Following that definition, we prefer to use the term 'probabilistic uncertainty analysis' rather than 'probabilistic sensitivity analysis' to describe the process of drawing repeated samples from non-sampled data, i.e. from some *a priori* defined distributional form of costs and / or effects. We use the term 'non-parametric bootstrapping' only in relation to drawing samples from sampled data.

Moreover, this paper discusses how information on uncertainty should be communicated to policy makers. The above mentioned techniques all present study results in terms of some type of uncertainty interval. However, little or no attention is paid to the question of how decision makers should interpret the results where uncertainty intervals overlap. This absence of guidance to decision makers is exacerbated in sectoral CEA based on the implicit or explicit use of cost-effectiveness league tables (10,11). Sectoral analysis requires that interventions are ranked on the basis of their cost-effectiveness ratios. In deterministic analysis, decision makers are assumed to work down the list, starting with the most cost-effective, and to stop funding interventions when the resources run out. The addition of uncertainty to this analysis is more realistic, but uncertainty intervals of many of the ratios may overlap and the decision maker is left with no guidance in the literature. It is simply assumed that no decision about which intervention is more efficient can be made. Yet, decision makers must and do make decisions about which interventions to encourage even when uncertainty is high (e.g. with overlapping confidence intervals).

We propose a new way for presenting the results in a way that will help decision makers allocate their scarce resources taking into account uncertainty, in the form of a 'stochastic league table'. This informs decision-makers about the probability that a specific intervention would be included in the optimal mix of interventions for various levels of resource availability, taking into account the uncertainty surrounding costs and effectiveness.

Probabilistic uncertainty analysis by Monte Carlo simulations

Probabilistic uncertainty analysis using Monte Carlo simulations has been well descibed elsewhere (4,7-9). Most applications assume a distributional form (e.g. normal, uniform, binominal) for each estimated (but non-sampled) variable. Repeated samples are then drawn from these distributions to determine an empirical distribution for some construct of the variables, such as CERs.

To illustrate the procedure, consider three interventions A, B and C, each with a number of mutually exclusive alternatives. Table 1 and Figure 1 present the hypothetical costs and effect data first presented in Murray et al (12). To reflect uncertainty, costs are assumed to be log normally distributed with standard deviation of 20 and effects are assumed to be normally distributed with standard deviation of 20. The covariance is assumed to be zero. Both interventions are defined at the population level, i.e. they refer to total costs and total effects of interventions in a population. As these distributions also refer to the population level rather than the individual level, this implies that draws taken from the samples have size n=1. The procedure to generate a sample distribution for the incremental CER from expanding the intervention, for example from C1 to C2, is then as follows:

1. Take one sample of costs and effects from the distribution of costs and effects from C1: C_{c1} and E_{c1}, and one sample of cost and effects from the distribution of costs and effects of c2: C_{c2} and E_{c2}

2. The sample estimate of the incremental CER is then given by C_{c2}-C_{c1} divided by E_{c2}-E_{c1}.
3. Repeating this three-stage process B times gives a vector of sample estimates which is the empirical sampling distribution of the incremental CER statistic.

There is little stability in these CER estimates where the distributions of costs or effectiveness overlap - some simulations will produce negative net effects and some positive net effects, for example. This can lead to positive or negative incremental CERs. Figure 2 shows the mean value of the sampled estimates of the incremental CERs for the alternatives of intervention C. For the assumed ranges of costs and effects, even after a large number of samples (B = 10,000) some means do not stabilize, and may even take on negative values(e.g. C2→C3). (The mean CER of 0→C1 is relatively stable because the origin is fixed and costs and effects of intervention C1 constitute its only source of uncertainty e.g. it is an average versus incremental ratio.)

The simple percentile method allows us to estimate confidence intervals in the presence of these unstable means. This approach takes the $100(\alpha/2)$ and the $100(1-(\alpha/2))$ percentile values of the bootstrap distribution as the upper and lower confidence limits for the CER. Table 2 shows the 90% confidence intervals for the alternatives of intervention C.

Of special interest are interventions that are weakly dominated on the basis of the point estimate of their CER but have a wide confidence interval. In such cases, some simulations might show them to be no longer dominated. In figure 1, see for example intervention B1. Because its mean is located north-west of the line between the origin and intervention B2, it appears to be weakly dominated. However, the uncertainty range of the CER of the implementation of B1 ranges from $5.27 to $12.63 per health effect. The $5.27 is lower than some of the values in the uncertainty range of the CER of the implementation of B2. That is, we cannot be sure that B1 should be excluded of the set of alternatives under consideration.

Combining probabilistic uncertainty analysis with non-parametric bootstrap procedures

In the situation in which individual level data are available for some component of costs or effects, one feasible approach is to combine probabilistic uncertainty analysis with non-parametric bootstrapping to estimate a total ‘uncertainty range’ for CERs (9). The use of non-parametric bootstrapping has been advocated by many authors (13-16) and has been extensively applied to empirical data (2,7,13,14,16-23). Unlike probabilistic uncertainty analyses, the bootstrap approach is a non-parametric method that makes no distributional assumptions concerning the statistic in question. Instead it employs the original data in a resampling exercise in order to give an empirical estimate of the sampling distribution of that estimate.

The basic concept behind non-parametric bootstrapping is to treat the study sample as if it were the population, the premise being that it is better to draw inferences from the sample in hand rather than make potentially unrealistic assumptions about the underlying population. Using the non-parametric bootstrap approach, successive random draws are taken with replacement from the study sample data As such, the fact that an observation has been selected does not preclude it from being selected again for the same resample, which leads to the construction of different bootstrap resamples. The statistic of interest and its distribution is calculated from these resamples. The number of bootstrap resamples, B, should at least be 1,000 to construct confidence intervals, in order to ensure that the tails of the empirical distributions are filled (13). An important advantage of the non-parametric bootstrap approach is that it is of no consequence whether the original sample is a well-behaved distribution because it forms its own probability density function.

To illustrate the combination of probabilistic uncertainty analysis and non-parametric bootstrapping, consider a CEA with costs being the product of vectors of unit prices and resource utilization. By defining a probability distribution of unit prices[2], and with resource utilization and effectiveness data stemming from sampled data, a total uncertainty range[3] can be estimated by combining probabilistic uncertainty analysis with non-parametric bootstrapping. To start with, a large number (B) of samples of size n_p of sets of unit prices are obtained by random sampling from the prior distributions, and the mean price is calculated for each of the B samples. Similarly, B bootstrap samples of size n_q are taken from the resource utilization and effectiveness data, and the mean resource utilization and effectiveness is estimated for each of the B bootstrap samples. Then B replicates of the CER can be obtained by combining B bootstraps of both resource utilisation and effectiveness data with the B sets of prices sampled from the prior distributions. These are then used to calculate a percentile interval.

Uncertainty at the allocation level

Traditionally, the above results, as reported in Table 2, are placed in a single league table to inform decision-makers about the relative value of a set of (mutual exclusive) interventions. One problem with the league table approach is that the rank ordering which takes place is usually made on the basis of point estimates of CE alone (11). For example, Goodman et al. incorporated the estimated uncertainty interval for their estimates of the CE of interventions against malaria, but because the intervals overlapped, they were unwilling to suggest which ones should be given preference in the event of a shortage of resources (24). If policy makers are to be fully informed when making decisions based on the results of economic evaluations, it is important that analysts assist

[2] Statistical analysis of costs are generally performed conditional on a set of unit prices, thereby suppressing any uncertainty associated with those unit price estimates although varying unit price estimates can have a non-negligible effect on results (27).

[3] Instead of calling this a confidence interval, the term *uncertainty range* could be used as such an interval incorporates uncertainty in sampled and non-sampled data.

them more than is done by presenting either point estimates of CERs or uncertainty intervals which overlap.

We propose a different approach building on the above work; i.e. the use of stochastic league tables. The approach provides the likelihood of inclusion of a specific intervention in the optimal mix of interventions given the uncertainty surrounding the intervention. The construction of stochastic league tables requires four steps[4]. Firstly, using Monte Carlo simulations, random draws are taken from distributions of costs and effects. This step is identical to the first step in probabilistic uncertainty analysis as described above, but random draws are taken from distributions for all interventions.

The second step is to determine the optimal mix of interventions for given levels of resource availability following the procedure for choosing between mutually exclusive and independent interventions outlined in Murray et al. (12). The most efficient intervention in the set of mutually exclusive interventions is evaluated according to its average CE ratio (versus doing nothing), while the CE of others in the mutually exclusive set are evaluated incremental to the most efficient intervention.

Thirdly, this process is repeated a large number of times (here 10,000) to provide 10,000 estimates of the optimal mix of interventions. If P equals the number of times that an intervention is included in the optimal mix, P/10,000 is the probability that the intervention is included. In our example, for resources equal to 50, C1 is included 4,323 times, a 43% probability of being included (Table 3). P for C2 equals 1,406, a probability of inclusion of 14%. In the remaining cases (43% of all random draws), costs of each possible option overrun the available resources and no intervention can be funded fully. This explains why the probabilities do not add up to 100%.

The fourth step involves repeating this procedure for various levels of resource availability to reveal the 'resource expansion path', showing the probability that each intervention will be included at different levels of resource availability (Table 3). Decision makers can use this information to prioritize interventions should more resources become available for health care. The probability that a more expensive alternative will be included increases with the level of resource availability. For example, the probability C2 is included increases from 14% to 50% when resources increase from 50 to 100. In our example, no intervention is included in the optimal mix with certainty – even at high levels of resource availability – because of the relative large standard deviations assumed for costs and effects.[5]

The degree of uncertainty in costs and effects of an intervention can have a large impact on its probability of inclusion in the optimal mix. If we change the

---

[4]We are in the process of developing a software program, MCLeague, to carry out this process.

[5] At these high budget levels, the sum of the probabilities of inclusion of mutual exclusive interventions equals 100%, as there will always be resources to finance one of the interventions.

standard deviation of the cost distribution for intervention A2 from 20 to 70, its probability of inclusion at a level of resource availability of 300 increases from 5% to 22% (Table 4). This is because intervention costs now are sometimes very low thereby rendering the intervention relatively cost-effective (with resources ≥ 600, its probability of inclusion decreases because it now has to compete with the more cost-effective interventions A3 and A4 which can be afforded). The general conclusion is that the higher the uncertainty in costs and effects, the more equal the probabilities of inclusion of interventions will be, other things equal. This is true both within the same mutual exclusive set as well as between independent sets of interventions.

In Table 3, the numbers in bold represent interventions that would be selected in a traditional league table based on the cost-effectiveness ratios calculated in Table 1. These interventions would also be chosen by the stochastic league table because of their higher probabilities of inclusion. However, the stochastic league table provides additional information to the decision maker. With resources of 200, a traditional league table would choose intervention C2 whereas our stochastic league table shows almost identical probabilities of inclusion of C1 and C2 in the optimal mix of interventions. This information provides decision makers with more information than simply presenting the confidence intervals for all CERs. For example, it allows decision makers to better evaluate the impact of trading off the efficiency goal against other objectives such as reducing health inequalities in their selection of interventions (25). In general, the more interventions (belonging to the same mutually exclusive set) differ regarding their probabilities of inclusion in the optimal mix, the more efficiency decision makers give up if they choose to over-ride the results in favor of other goals in their choice of interventions - the stochastic league table in our example informs decision makers that they are not likely to lose much in terms of efficiency if they decide to select C1 rather than C2 for equity reasons. This important information is not revealed in deterministic league tables or in the traditional approach to uncertainty in CEA.

Another advantage concerns the information provided in the expansion path, illustrated in Table 3. With resources of 200, there is little to choose between B2 and B3 but preference would be given to B2. However, if the decision maker felt that additional resources would become available in the near future, and that the costs of switching from B2 to B3 might be substantial, it would be sensible for them to choose B3. Again, this type of information is not provided in the standard approach to uncertainty.

Stochastic league tables may also show that interventions that would otherwise have been ruled out by dominance in traditional league tables might well be included in some draws. As noted above, in our example, intervention B1 will never be eligible for selection in a deterministic league table because it is (weakly) dominated by B2. However, taking into account uncertainty the stochastic league table (Table 3) shows that B1 has a low but non-zero probability of being included in the optimal mix. Whether decision makers will actually select such interventions depends on the probability of inclusion compared to other mutually exclusive alternatives, and the trade-off between efficiency and other objectives of health systems.

Figure 1 depicts an alternative way of visualizing the information of Table 3. The vertical axis shows the probability of being chosen at the level of resource availability on the horizontal axis. The logic is the same as that described for the interpretation of the tables.

Discussion

This paper presents an extension and generalization of previously described methods examining uncertainty in cost-effectiveness studies. Whereas previous studies applied the concept of bootstrapping and Monte Carlo simulations in the contexts of clinical trials, here we apply them to decision models analyzing cost-effectiveness based on any combination of primary and secondary data. Given the prevailing scarcity of sampled data on costs and effects of many public health interventions in developing countries, we propose the use of probabilistic uncertainty analysis using Monte Carlo simulations, in combination with non-parametric bootstrapping techniques where appropriate. The simple percentile method can be applied to estimate uncertainty ranges.

However, the reporting of some type of uncertainty range of CER in individual studies ignores the question of how policy makers should interpret the results where uncertainty results overlap. The stochastic league table developed in this paper is a new way of presenting uncertainty around costs and effects to decision makers. It provides additional information beyond that offered by the traditional treatment of uncertainty in CEA, presenting the probability that each intervention is included in the optimal mix for given levels of resource availability. The most likely optimal mix will be the one that contributes the most to maximizing population health for that level of resources. Decision makers can then decide the extent to which they should trade off gains in efficiency for gains in other goals of the health system.

Stochastic league tables are conceptually different from the recently suggested portfolio approach, borrowed from financial economics and characterizing health care resources allocation as a risky investment problem (26). This approach provides the optimal intervention mix given decision-makers' explicit preferences concerning risk and return. Our stochastic league table provides the probability of an intervention being chosen in the optimal mix, given uncertainty. Risk-neutral decision makers would choose the most likely combination of interventions.

A drawback to our framework (and to the portfolio approach for that matter) is that distributions of costs and effects are assumed to be independent e.g. no joint distributions are defined. Moreover, the definition of the prior distributions is left to the analyst, who may have very little information about the actual distribution, but whose choice is likely to have a large effects on the results. It is technically possible to include covariance between costs and outcomes in the analysis, but this requires more information about covariances than is usually available. The development of stochastic league tables is an important step forward in the interpretation of uncertainty at the decision making level.

Table 1. Costs and effects for three independent sets of mutually exclusive alternatives

Interventions	Costs		Health Benefits		Cost-effectiveness*
	Mean	S.D.	Mean	S.D.	
A1	120	20	1	2	-
A2	140	20	5.5	2	25.4
A3	170	20	3	2	-
A4	190	20	7	2	33.3
B1	100	20	12	2	-
B2	120	20	17	2	7.1
B3	150	20	20	2	10.0
C1	50	20	22	2	2.3
C2	70	20	24.5	2	8.0
C3	120	20	29	2	11.1
C4	170	20	31	2	25.0

* cost-effectiveness ratios after exclusion of dominated interventions

Table 2. Bootstrap confidence interval for incremental CERs of four mutually exclusive alternatives of intervention C using the 90% simple percentile method

Intervention	Incremental CER	
	Mean	90% Confidence Interval
0-C1	2.29	0.77 - 3.88

C1-C2	8.05	-47.34	-	63.09
C2-C3	-73.56	-8.08	-	59.49
C3-C4	-14.51	-112.61	-	116.97

Table 3. Stochastic league table presenting the probability of inclusion (%) of three independent sets of mutual exclusive interventions in the optimal mix of interventions at different levels of resource availability*

Intervention	Resource availability							
	50	100	150	200	300	400	600	800
A1	0	0	0	0	1	2	0	0
A2	0	0	1	1	5	25	28	29
A3	0	0	0	0	0	1	4	4
A4	0	0	0	0	0	11	67	67
B1	0	3	14	11	2	0	0	0
B2	0	1	18	35	26	21	15	15
B3	0	0	4	36	72	78	85	85
C1	43	40	30	36	8	3	0	0
C2	14	47	25	37	17	10	0	0
C3	0	13	38	16	45	333	24	23
C4	0	0	8	11	31	54	76	76

* numbers in bold represent interventions that would be listed in a traditional league table.

Table 4. As Table 3, with standard deviation for costs of intervention A2 increased from 20 to 70

Intervention	Resource availability							
	50	100	150	200	300	400	600	800
A1	0	0	0	0	1	2	1	1
A2	7	9	14	12	22	37	28	28
A3	0	0	0	0	0	1	5	4
A4	0	0	0	0	1	9	67	67
B1	0	3	15	10	1	1	0	0
B2	0	1	19	35	28	18	15	14
B3	0	0	3	37	71	81	85	86
C1	43	40	30	39	11	3	0	0
C2	15	47	30	33	17	7	1	1
C3	0	13	34	17	45	30	23	24
C4	0	0	6	11	28	60	77	75

* numbers in bold represent interventions that would be listed in a traditional league table.

Figure 1. Costs and effects of three sets of mutually exclusive interventions

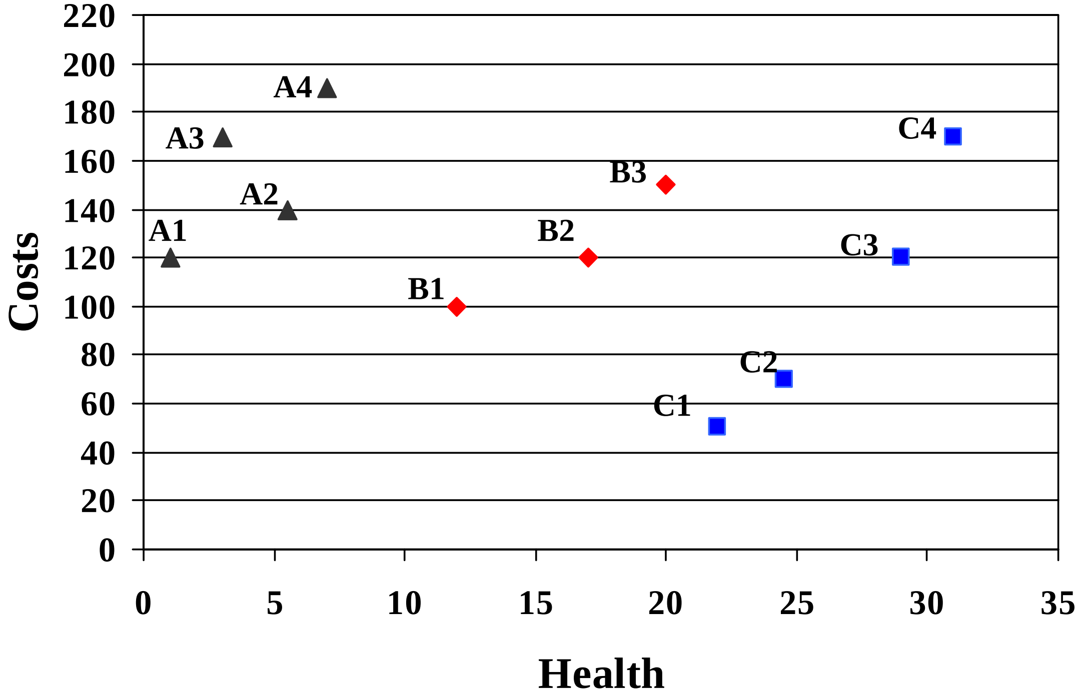

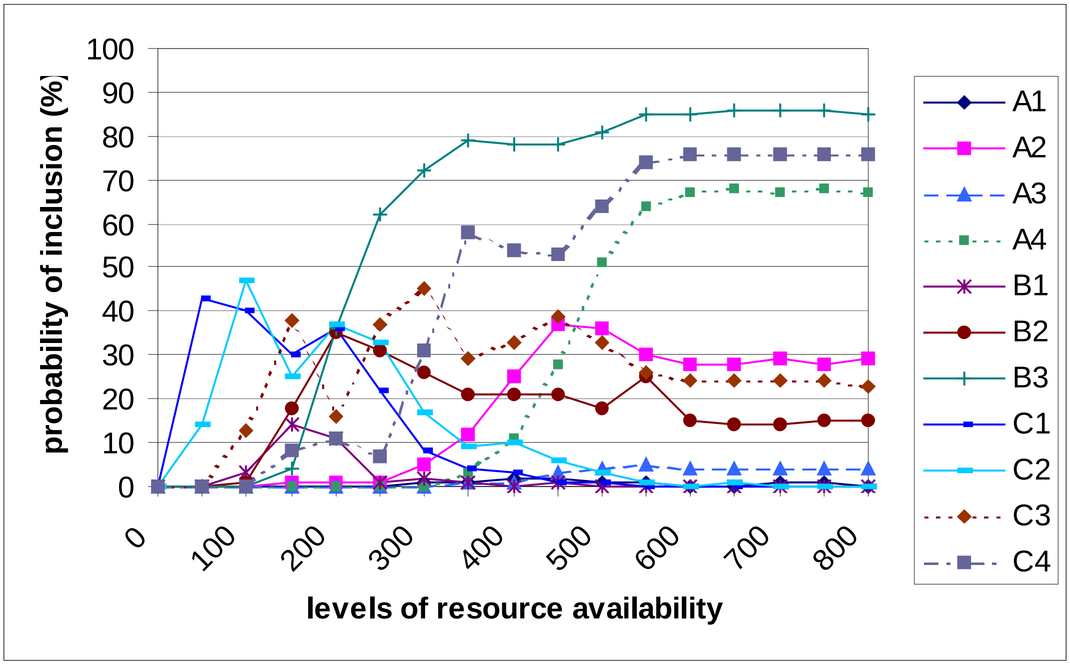

Figure 3. Probability of inclusion (%) of three independent sets of mutually exclusive interventions

Reference List

1. O'Brien BJ, Drummond MF, Labelle RJ, Willan A. In search of power and significance: issues in the design and analysis of stochastic cost-effectiveness studies in health care. Med.Care 1994;32(2):150-63.

2. Polsky D, Glick HA, Willke R, Schulman K. Confidence intervals for cost-effectiveness ratios: a comparison of four methods. Health Econ. 1997;6(3):243-52.

3. Briggs AH. A Bayesian approach to stochastic cost-effectiveness analysis. Health Econ 1999;8(3):257-61.

4. Briggs A, Sculpher M, Buxton M. Uncertainty in the economic evaluation of health care technologies: the role of sensitivity analysis. Health Econ. 1994;3(2):95-104.

5. Briggs A, Sculpher M. Sensitivity analysis in economic evaluation: a review of published studies. Health Econ. 1995;4(5):355-71.

6. O'Brien BJ, Drummond MF. Statistical versus quantitative significance in the socioeconomic evaluation of medicines. PharmacoEconomics 1994;5(5):389-98.

7. Critchfield GC, Willard KE, Connelly DP. Probabilistic sensitivity analysis methods for general decision models. Comput.Biomed.Res. 1986;19(3):254-65.

8. Doubilet P, Begg CB, Weinstein MC, Braun P, McNeil BJ. Probabilistic sensitivity analysis using Monte Carlo simulation. A practical approach. Med.Decis.Making. 1985;5(2):157-77.

9. Lord J, Asante MA. Estimating uncertainty ranges for costs by the bootstrap procedure combined with probabilistic sensitivity analysis. Health Econ. 1999;8(4):323-33.

10. Birch S, Gafni A. Cost-effectiveness ratios: in a league of their own [see comments]. Health Policy. 1994;28(2):133-41.

11. Drummond M, Torrance G, Mason J. Cost-effectiveness league tables: more harm than good? Soc.Sci.Med. 1993;37(1):33-40.

12. Murray CJ, Evans DB, Acharya A, Baltussen RM. Development of WHO guidelines on generalized cost-effectiveness analysis. Health Econ 2000;9(3):235-51.

13. Briggs AH, Wonderling DE, Mooney CZ. Pulling cost-effectiveness analysis up by its bootstraps: a non- parametric approach to confidence interval estimation. Health Econ. 1997;6(4):327-40.

14. Chaudhary MA, Stearns SC. Estimating confidence intervals for cost-effectiveness ratios: an example from a randomized trial [see comments]. Stat.Med 1996;15(13):1447-58.

15. Manning WG, Jr., Fryback DG, Weinstein MC. Gold MR, Siegel JE, Russel LB, Weinstein MC, editors.Cost-effectiveness in health and medicine. 1 ed. New York: Oxford University Press; 1996; 8, Reflecting uncertainty in cost-effectiveness analysis. p. 247-175.

16. Briggs AH, Gray AM. Handling uncertainty when performing economic evaluation of healthcare interventions. Health Technol.Assess. 1999;3(2):1-134.

17. Briggs A, Fenn P. Confidence intervals or surfaces? Uncertainty on the cost-effectiveness plane. Health Econ. 1998;7(8):723-40.

18. Campbell MK, Torgerson DJ. Bootstrapping: estimating confidence intervals for cost-effectiveness ratios. QJM. 1999;92(3):177-82.

19. Hunink MG, Bult JR, de Vries J, Weinstein MC. Uncertainty in decision models analyzing cost-effectiveness: the joint distribution of incremental costs and effectiveness evaluated with a nonparametric bootstrap method. Med Decis.Making 1998;18(3):337-46.

20. Mennemeyer ST, Cyr LP. A bootstrap approach to medical decision analysis. J Health Econ. 1997;16(6):741-7.

21. Obenchain RL, Melfi CA, Crompton DW, Bueshing DP. Bootstrap analyses of cost-effectiveness in antidepressant pharmacotherapy. PharmacoEconomics 1997;11:464-72.

22. Severens JL, De Boo TM, Konst EM. Uncertainty of incremental cost-effectiveness ratios. A comparison of Fieller and bootstrap confidence intervals. Int J Technol.Assess.Health Care 1999;15(3):608-14.

23. Tambour M, Zethraeus N. Bootstrap confidence intervals for cost-effectiveness ratios: some simulation results [comment]. Health Econ. 1998;7(2):143-7.

24. Goodman CA, Mills AJ. The evidence base on the cost-effectiveness of malaria control measures in Africa. Health Policy Plan 1999;14(4):301-12.

25. Murray CJ, Frenk J. A framework for assessing the performance of health systems. Bull World Health Organ 2000;78(6):717-31.

26. O'Brien BJ, Sculpher MJ. Building uncertainty into cost-effectiveness rankings: portfolio risk- return tradeoffs and implications for decision rules. Med Care 2000;38(5):460-8.

27. Rittenhouse BE, Dulisse B, Stinnett AA. At what price significance? The effect of price estimates on statistical inference in economic evaluation. Health Econ. 1999;8(3):213-9.

Figure 2. Mean incremental CERs of four mutually exclusive alternatives of intervention C as a function of number of samples (in total 10,000).