

STOCHASTIC LEAGUE TABLES: COMMUNICATING COST-EFFECTIVENESS RESULTS TO DECISION-MAKERS

RAYMOND C.W. HUTUBESSY^{a,*}, ROB M.P.M. BALTUSSEN^a, DAVID B. EVANS^a,
JAN J. BARENDREGT^b AND CHRISTOPHER J.L. MURRAY^a

^a *Global Programme on Evidence for Health Policy (GPE), World Health Organization (WHO),
Geneva, Switzerland*

^b *Department of Public Health, Erasmus University, Rotterdam, Netherlands*

SUMMARY

The presentation of the results of uncertainty analysis in cost-effectiveness analysis (CEA) in the literature has been relatively academic with little attention paid to the question of how decision-makers should interpret the information particularly when confidence intervals overlap. This question is especially relevant to sectoral CEA providing information on the costs and effects of a wide range of interventions.

This paper introduces stochastic league tables to inform 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. 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. Copyright © 2001 John Wiley & Sons, Ltd.

KEY WORDS — cost-effectiveness analysis; decision-making analysis; uncertainty analysis

INTRODUCTION

Uncertainty in cost-effectiveness analysis (CEA) has received a lot of attention in recent years, leading to the development of a range of approaches, such as non-parametric bootstrapping [1], the construction of confidence planes [2], mathematical techniques [3], probabilistic sensitivity analyses using Monte Carlo simulations [4], and the net health benefit approach [5]. These 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 [6,7]. 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

* Correspondence to: Global Programme on Evidence for Health Policy (GPE), World Health Organization (WHO), 20 Avenue Appia, CH-1211 Geneva, Switzerland. Tel.: +41 22 7913253; fax: +41 22 7914328; e-mail: hutubessyr@who.ch

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 approach to presenting decision-makers with the results of CEA including uncertainty through the construction 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 or resource availability, taking into account the uncertainty surrounding its total costs and effectiveness. Each intervention should be thought of as a national programme or policy, which can only be purchased at one point [8]. Although the argument is presented with reference to the generalized method for sectoral CEA, which we recently proposed, allowing decision-makers to assess the efficiency of the current mix of interventions as well as the relative attractiveness of changes to this mix should new resources become available [8], it is applicable to any form of sectoral analysis.

THE ANALYTICAL FRAMEWORK

The construction of stochastic league tables requires four steps (a software program, MCLLeague, is being developed to carry out this process). Firstly, using Monte Carlo simulations,

random draws are taken from estimated distributions of total costs and effects for the interventions under study. These distributions are *a priori* defined by the analyst and may take different forms, for example normal, log-normal, and uniform distributions [4]. Table 1 presents the hypothetical costs and effect data first presented in Murray *et al.* [8]. To reflect uncertainty, costs are here assumed to be log-normally distributed, with an S.D. of 20, and effects are assumed to be normally distributed with an S.D. of 2. The covariance is assumed to be zero. The conclusions are not dependent on these assumptions. Random draws are taken from these 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.* [8]. The most efficient intervention in the set of mutually exclusive interventions is evaluated according to its average cost-effectiveness ratio (versus doing nothing), while the cost-effectiveness 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 10000) to provide 10000 estimates of the optimal mix of interventions. If P equals the number of times that an intervention is included in the optimal mix, $P/10000$ is the probability that the intervention is included. Hence, P is the proportion of samples from the estimated

Table 1. Costs, effects and cost-effectiveness of three independent sets of mutually exclusive alternatives

Interventions	Total costs		Total effects		Cost-effectiveness ^a
	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

^a Cost-effectiveness ratios after exclusion of dominated interventions.

distribution for which the intervention is estimated to be optimal based on the sample average and incremental cost-effectiveness ratios. In our example, for resources equal to 50, C1 is included 4323 times, a 43% probability of being included (Table 2). P for C2 equals 1406, 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 2). 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 47% 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.

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 standard deviation of the cost distribu-

tion 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 3). 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 2, 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

Table 2. 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^a

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	33	24	23
C4	0	0	8	11	31	54	76	76

^a Numbers in bold represent interventions that would be listed in a traditional league table.

Table 3. As Table 2, with standard deviation for costs of intervention A2 increased from 20 to 70^a

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

^a Numbers in bold represent interventions that would be listed in a traditional league table.

against other objectives such as reducing health inequalities in their selection of interventions [9]. 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 favour 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 2. 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. 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 2) 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 2. 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.




Figure 1. Probability of inclusion (%) of three independent sets of mutually exclusive interventions

DISCUSSION

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 [10]. 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 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. Alternatively, where empirical data on patient' costs and effects are available, our framework could employ the technique of non-parametric bootstrapping in which samples are drawn with replacement from the original data. This approach has the advantage that it does not rely on parametric assump-

tions concerning the underlying distribution and that covariances between costs and effects can be easily incorporated [1]. The development of stochastic league tables is an important step forward in the interpretation of uncertainty at the decision-making level.

ACKNOWLEDGEMENTS

The views expressed are those of the authors and not necessarily those of the Organization they represent.

REFERENCES

1. 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**: 327–340.
2. Briggs A, Fenn P. Confidence intervals or surfaces? Uncertainty on the cost-effectiveness plane. *Health Econ* 1998; **7**: 723–740.
3. Willan AR, O'Brien BJ. Confidence intervals for cost-effectiveness ratios: an application of Fieller's theorem. *Health Econ* 1996; **5**: 297–305.
4. Lord J, Asante MA. Estimating uncertainty ranges for costs by the bootstrap procedure combined with probabilistic sensitivity analysis. *Health Econ* 1999; **8**: 323–333.
5. Stinnett AA, Mullahy J. Net health benefits: a new framework for the analysis of uncertainty in cost-effectiveness analysis. *Med Decis Making* 1998; **18**(Suppl): S68–S80.
6. Birch S, Gafni A. Cost-effectiveness ratios: in a league of their own. *Health Policy* 1994; **28**: 133–141.
7. Drummond M, Torrance G, Mason J. Cost-effectiveness league tables: more harm than good? *Soc Sc Med* 1993; **37**: 33–40.
8. Murray CJ, Evans DB, Acharya A, *et al.* Development of WHO guidelines on generalized cost-effectiveness analysis. *Health Econ* 2000; **9**: 235–251.
9. Murray CJ, Frenk J. A framework for assessing the performance of health systems. *Bull World Health Organ* 2000; **78**: 717–731.
10. O'Brien BJ, Sculpher MJ. Building uncertainty into cost-effectiveness rankings: portfolio risk-return tradeoffs and implications for decision rules. *Med Care* 2000; **38**: 460–468.