Presenting the optimal decision, and the uncertainty associated with that decision, when assessing the cost-effectiveness of multiple options

Garry R. Barton,[1] Andrew H. Briggs,[2] & Elisabeth A. L. Fenwick.[2]

[1] School of Economics, University of Nottingham (mcxgrb@nottingham.ac.uk).

[2] Public Health & Health Policy, University of Glasgow (a.briggs@clinmed.gla.ac.uk; e.fenwick@clinmed.gla.ac.uk).

Objective: To demonstrate how the optimal decision, and level of uncertainty associated with that decision, can be presented when assessing the cost-effectiveness of multiple options. To explore and explain potentially counter-intuitive results that can arise when analysing multiple options.

Methods: A template was created, based on the assumption of multivariate normality, in order to replicate previous analyses which have compared the cost-effectiveness of multiple options. We used this template to explain the different shapes that the cost-effectiveness acceptability curve (CEAC), the cost-effectiveness acceptability frontier (CEAF), and the expected value of perfection information (EVPI) may take, with changing correlation structure and variance between the multiple options.

Results: We show that it is possible for i) an option, which is subject to extended dominance, to have the highest probability of being cost-effective for some willingness to pay (WTP) values, ii) the most cost-effective (optimal) option to never have the highest probability of being cost-effective, and iii) the EVPI to increase when the probability of making the wrong decision decreases. Changing the correlation structure between multiple options did not, however, change the presentation of results on the cost-effectiveness plane.

Conclusion: The cost-effectiveness plane has limited use in representing the uncertainty surrounding multiple options as it cannot represent correlation between the options. CEACs can represent decision uncertainty, but should not be used to determine the optimal decision. Instead, the CEAF shows the decision uncertainty surrounding the optimal choice, this can be augmented by the EVPI to show the potential gains to further research.

INTRODUCTION

The cost-effectiveness acceptability curve (CEAC) is often used to present the results of cost-effectiveness analysis [1-14]. However, it is possible that the CEAC may be misinterpreted [15, 16]. This possibility is highlighted by two potentially counter-intuitive results – Goeree et al. [2] have previously found i) that an option which was subject to extended dominance [17, 18] had the highest probability of being cost-effective, and ii) that an optimal option (as it provided the greatest benefit for a given cost) never had the highest probability of being cost-effective. The aim of this paper is to elucidate the appropriate presentation of uncertainty in cost-effectiveness studies in the case of multiple options, and to explore and explain why such potentially counter-intuitive results can arise when analysing multiple options.

When making decisions about the allocation of scarce health-care resources it has been argued that two questions are fundamental [13, 19, 20]. First, which option is estimated to be cost-effective, on the basis of existing evidence? Second, should further research be undertaken in order to reduce the level of uncertainty associated with that decision? This paper seeks to show how both of these questions can be addressed when assessing multiple options, and is structured as follows. First, we outline a template which can be used to represent uncertainty in the decision among a number of mutually exclusive treatment options for a patient group. The calculation of CEACs, the cost-effectiveness acceptability frontier (CEAF) and the expected value of perfect information (EVPI) are then described and demonstrated. The template is subsequently used to replicate two previously reported cost-effectiveness studies of multiple options [2, 3]. Finally, different assumptions about the correlation structure, and the variances of the outputs, are explored in order to understand how these can influence the shape of the resulting CEACs, CEAF, and EVPI.

METHODS

Probabilistic methods can be used to describe the uncertainty associated with input parameters in a model, and, in turn, estimate the uncertainty in the model outputs of cost, effect, and cost-effectiveness [21]. For the purposes of exploring potential scenarios relating to the costs and effects of multiple options, we developed a template (in Excel™) that approximates the joint distributions of cost and effects of multiple options under the assumption of multivariate normality. The template works with up to seven different options and requires the user to input the means and standard errors for all options. In addition, the user also specifies the correlation structure between the costs and effects both within and between options. From this information the variance-covariance matrix is generated and (through the use of the Cholesky decomposition technique [21]) the template draws 1000 simulations of both the cost and effect parameters of each option from the multivariate normal distribution. Finally, using the methods described below, the template can be used to estimate the CEACs, CEAF and EVPI.

Cost-effectiveness acceptability curve (CEAC)

The CEAC was originally introduced in order to graphically represent the level of uncertainty associated with the estimated cost-effectiveness of an option, providing an alternative to confidence intervals [22]. In the case of multiple options, a separate CEAC can be plotted for each option, where each CEAC represents the probability of (each option) being cost-effective at different levels of WTP [21, 23]. This probability was estimated in the following manner. In each of the 1000 iterations the total cost (C) and total effect (E) was estimated, for each option, and for a particular WTP (λ), the net monetary benefit (NMB) was estimated: NMB= λ*E - C. In each of the iterations the option with the highest net benefit was then identified. The probability of being cost-effective was then equivalent to the proportion of the 1000 iterations for which each option had the highest net benefit. CEACs were estimated by plotting these proportions (y-axis) for different WTP values (x-axis).

Cost-effectiveness acceptability frontier (CEAF)

Fenwick et al. [24] have, however, shown that the probability of being cost-effective can not be used to determine the optimal option, and that if the societal objective is to maximise health gain then decisions should be taken on the basis of expected net benefit, regardless of the uncertainty (probability) associated with the decision. Thus, the option with the highest probability of being cost-effective (highest CEAC) at any WTP value need not be the optimal option. The solution to the fact that CEACs can not identify the optimal option is to plot the CEAF [24]. The CEAF plots the probability that the option is cost-effective only for the optimal options (at different WTP values), where the optimal option (at each value of WTP for health gain) is identified as that option which has the highest expected net benefit (see Stinnett & Mullahy [25] for an explanation of net benefit). One minus the probability given by the CEAF at any WTP value is equivalent to the probability that a "wrong" decision will be made [20]. Additionally, since only optimal options are represented, the CEAF may not always correspond to the options with the highest probability of being cost-effective. The lower value of the range of WTP values for which each option was optimal denotes the ICER for that particular option and the upper value denotes the ICER for the next most costly option. However, as the net benefit is determined by the expectations of model (i.e. costs and effects) this range will vary when changes are made to input parameters e.g. the correlation structure between options [21].

Expected Value of Perfect Information (EVPI)

There is always some level of uncertainty surrounding the decision about which option is optimal. The EVPI estimates the value of undertaking further research in order to reduce that uncertainty, at different WTP values, and is dependent upon the probability of a wrong decision being made, and the consequences of that wrong decision [21, 26]. The EVPI (for an individual patient) is calculated, for a given WTP, by first calculating the maximum net benefit (across all options) for each iteration, and taking the average of these values in order to estimate the average maximum net benefit. The maximum average net benefit (i.e. the net benefit of the optimal decision averaged across all iterations) was then subtracted from the average maximum net benefit in order to estimate the potential net benefit forgone as a result of the uncertainty (the EVPI). The EVPI was calculated for differing WTP values and plotted on the same diagram as the CEAF, it shows the maximum amount that one would be WTP to eliminate uncertainty.

Replication of previous studies

By replicating the comparison of multiple options by Kamath et al. [3], and by Goeree et al. [2], we sought to depict both the optimal decision, and the level of uncertainty associated with that decision. Kamath et al. [3] estimated the cost-effectiveness of five different treatment options for knee osteoarthritis. Costs were estimated at year 2000 levels (US$) and effectiveness (pain efficacy) was measured in terms of the number of patients who achieved a minimum perceptible clinical improvement (MPCI) on the Western Ontario and McMaster Universities Osteoarthritis Index (WOMAC) pain subscale [27]. In order to replicate the results, for each of the five options, we extracted the mean cost (and associated standard error) and mean effect (and associated standard error) from the results of their probabilistic analysis (which were reported for a cohort of 1000 patients). As no information was provided about the correlation structure between the options we made the assumption that the costs and effects, both within and between options, were independent (r=0). We then used the previously described template to estimate the variance–covariance matrix, estimate a distribution for the costs and effects of each of the five options, plot the results on the cost-effectiveness plane, and estimate the CEACs, the CEAF, and the EVPI. In doing so, we demonstrate that the CEACs do not show the value at which each option becomes the optimal choice, as was claimed by Kamath et al. [3].

The same principles were used to replicate the results of Goeree et al. [2], who estimated the cost-effectiveness of multiple options for the management of gastroesophageal reflux disease. Costs were estimated at year 2000 levels (US$), and effects were measured in terms of QALYs (quality adjusted life years). However, as we had the actual simulations of the probabilistic analysis undertaken by Goeree et al. [2], we were able to construct the correlation matrix and hence the variance–covariance matrix for their analyses. The first analysis we undertook (base-case) was a replication of the analysis undertaken by Goeree et al. [2]. By altering the assumptions underlying this analysis we then ran a series of sensitivity analyses in order to explain why they found the aforementioned two potentially counter-intuitive results. In the second analysis we assumed that the costs and effects of different options were independent, but that the relationship between costs and effects within options remained the same. In the third analysis we used the original correlation matrix but made the assumption that the variance associated with the mean cost and mean effect of each option was one fifth of that in the base-case (which may have occurred if the original analysis had involved a larger sample size). Finally, the fourth analysis repeated that of the third, but with the assumption that the variance was five times greater than that in the base-case (a smaller sample size).

RESULTS

Replication of the analyses by Kamath et al.

The template was first used to replicate the probabilistic analyses of Kamath et al. [3], using the assumption that the costs and effects were independent, both within and between options. In line with their results, celecoxib and ibuprofen plus prophylaxis were both dominated, the ICER for ibuprofen was estimated to be $606 per additional MCPI response (relative to acetaminophen), and rofecoxib was estimated to cost $14,114 per additional MCPI response (relative to ibuprofen). The per patient cost and effect pairs for each of the 1000 iterations were plotted on the cost-effectiveness plane (Figure 1, panel I); the lines which join the estimated mean cost and mean effect for acetaminophen, ibuprofen and rofecoxib represent the ‘efficiency frontier’ [2, 21]. This demonstrates that acetaminophen was the optimal option when the WTP per patient achieving MCPI was <$606, followed by ibuprofen (WTP≥$606 to <$14,114) and rofecoxib (WTP≥$14,114).

<< Figure 1 >>

The corresponding CEACs are also plotted in Figure 1 (panel II). At very low levels of WTP acetaminophen was estimated to have a relatively high probability of being cost-effective (98.1% when WTP=$100), but as the WTP increased the probability of ibuprofen being cost-effective initially increased (to 99.6% when WTP=$3500), but was superseded by rofecoxib at high levels of WTP. In addition, despite the fact that ibuprofen plus prophylaxis was dominated, it still has a non-zero possibility of being cost-effective. Thus, to exclude dominated options from the CEAC calculations, may result in an underestimation of the probability of a wrong decision being made.

Figure 1 also includes the CEAF (panel III). The ICER for ibuprofen and for rofecoxib are plotted as vertical lines in order to demonstrate how the optimal option was identified, for different levels of WTP. Prior to the first vertical line (WTP<$606) aceptaminophen was optimal, between the two vertical lines (WTP≥$606 to <$14,114) ibuprofen was optimal and rofecoxib was optimal after the second vertical line (WTP≥$14,114). As the CEAF is constructed by plotting the probability of the optimal option being cost-effective over the range of WTP values for which each option was optimal it should be noted that the CEAF may be disjointed, as here at the second vertical line (which denotes the ICER for rofecoxib). At this value of WTP the probability of ibuprofen being cost-effective was 50.4%, whereas the corresponding probability for rofecoxib was 47.2%. Moreover, at this point, the optimal option did not have the highest probability of being cost-effective (by choosing rofecoxib there was a 52.8% of making the wrong decision). This arose because the optimal option was determined by the highest expected net benefit, whereas the CEAC simply represents the proportion of iterations over which each option had the highest net benefit, as mentioned above CEACs should not be used to determine the optimal option.

Finally, Figure 1 includes the EVPI for an individual patient (panel III). This was relatively low when there was a high probability that the optimal option was cost-effective, but there were local maximums when the optimal option changed (and the probability of a wrong decision was relatively high) e.g. at a WTP of $606 and $14,114.

Replications of the analyses by Goeree et al.

Base-case

The template was also used to replicate the analyses by Goeree et al. [2], the results of which were plotted on the cost-effectiveness plane (Figure 2, panel I). Options C, D and G were subject to extended dominance and the efficiency frontier was composed of options B, A (ICER=$7755, compared to B), F (ICER=$12,183, compared to A), and E (ICER=$110,845, compared to F). The CEACs (Figure 2, panel II) show that one of the optimal options (A) never had the highest probability of being cost-effective – initially B had the highest probability (WTP<$7970), followed by D (WTP ≥$7970 to <$12,540), C (WTP ≥$12,540 to <$14,850), F (WTP ≥$14,850 to <$112,120) and E (WTP≥$112,120). This again demonstrates that the CEACs should not and can not be used to identify the optimal option, particularly as an optimal option (A) never had the highest probability of being cost-effective. This is further highlighted by the fact that options which were subject to extended dominance (C and D) were associated with the highest probability of being cost-effective at certain levels of WTP (Figure 2, panel II).

<< Figure 2 >>

In order to simultaneously present the optimal option, and the level of uncertainty associated with that option, the CEAF was plotted (Figure 2, panel III). Again, we see that the CEAF is disjointed, and the optimal option does not always have the highest probability of being cost-effective e.g. at a WTP of $7755 (the ICER of A compared to B) the probability that B was cost-effective was 35.8%, compared to 25.3% for A. Indeed, at a WTP of $12,000, we would choose option A, as it had the highest expected net benefit, even though there was a >80% chance that this would be the wrong decision.

By replicating the Goeree et al. [2] analysis it was also possible to estimate the EVPI for an individual patient (Figure 2, panel III). Again, there were local maximums in the EVPI when the optimal option changed (at a WTP of $12,183 and $110,845). However, this was not always the case, as when the optimal option changed from B to A there was a point of inflexion, but the EVPI continued to increase. This arose, despite there being a local peak in the error probability, because the falling probability of making an incorrect decision was outweighed by the increasing consequences of a wrong decision. Contrary to the previous example (see Figure 1, panel III), at high levels of WTP, the EVPI fell to virtually zero. This was due to virtually non-existent uncertainty at these levels of WTP, where option E had a high probability of being cost-effective due to strong evidence that it was the most effective option.

Sensitivity Analysis

When the analyses by Goeree et al. [2] were replicated, with the assumption that the costs and effects between options were independent, the cost-effectiveness plane was, visually, the same as in the base-case analyses (see Figure 2, panel I) as near identical cost and effect pairs (within options) were drawn in both of these probabilistic analyses. Thus, the ICERs were very similar to that in the base-case analyses. Conversely, the CEACs changed, with the curves tending to converge – those options which previously had the highest probability of being cost-effective, for a particular WTP, tended to have a lower probability of being cost-effective, and other options tended to have a higher probability of being cost-effective (Figure 3, panel I). Importantly, this meant that option A (which was the optimal option between a WTP of $7755 and $12,183) now had the highest probability of being cost-effective over a range of WTP values, whereas option C never did. Thus, the correlation structure can highly influence the probability of an option being estimated to be cost-effective, and explains why i) an option which is subject to extended dominance can have the highest probability of being cost-effective, and ii) an option which is optimal may never have the highest probability of being cost-effective.

<< Figure 3 >>

In the Goeree et al. [2] analysis these two results arose because all the options were positively correlated. Consequently, when, for example, option B had a low net benefit in a particular iteration it was likely that, in the same iteration, other options would also have a low net benefit, and at low levels of WTP, this meant that option B had the highest net benefit in a high proportion of iterations (and thereby a high CEAC). Conversely, when the options were independent, when option B had a low net benefit then there was a higher chance of another option having a higher net benefit, in the same iteration, and option B thereby had the highest net benefit in a lower proportion of iterations.

The assumption that the costs and effects between options were independent also changed the CEAF (Figure 3, panel II), where the optimal options tended to have a lower probability of being cost-effective, suggesting that there was more uncertainty. As a result, the EVPI also tended to be higher (Figure 3, panel II). This even occurred when the probability of the optimal option being cost-effective increased e.g. at a WTP of $12,000, in the base-case the

EVPI was 41.4 when option A had a 17.9% probability of being cost-effective, yet it rose to 107.4 when the options were assumed to be independent and the probability of option A being cost-effective was 19.8%. This change in the EVPI can be explained by the fact that when option B, for example, had a high net benefit in a particular iteration in the base-case (independent) analysis it was more (less) likely that other options would also have a high net benefit. Thus, in the base-case the consequences of the wrong decision (represented by the incremental net benefit) were estimated to be less than in the analysis where the options were assumed to independent.

The final two analyses had the same correlation structure as the base-case, but differing levels of variance. In both analyses the cost and effect pairs for each option were centred around similar means, and thereby ICERs, but the pairs were more condensed when the level of variation was reduced by a factor of five (Figure 4, panel I) and more spread out when the variation was inflated by a factor of five (Figure 4, panel II). Looking at the CEACs, reducing the level of variation tended to increase the probability of optimal options being cost-effective, at the expense of the probability associated with other options (Figure 5, panel I). Indeed, option A now had the highest probability of being cost-effective over a range of WTP values, and options C/D only briefly did so (from WTP ≥\$11,690 to <\$12,950). Thus, the level of variation in the costs and effects can highly influence the probability of an option being estimated to be cost-effective, and explains why i) an option which is subject to extended dominance can have the highest probability of being cost-effective, and ii) an option which is optimal may never have the highest probability of being cost-effective.

<< Figure 4 >>

In the Goeree et al. [2] analysis these two results arose because options C and D, which were subject to extended dominance, had relatively high levels of variation in net benefit, compared to option A. For example, at a WTP of \$12,000 the mean net benefits for options A, C, and D were \$9844, \$9841 and \$9840, and the variances were \$7426, \$10,967 and \$11,027. Thus, as the probability distributions surrounding the mean net benefit of options C/D were more spread out than for option A, within a particular iteration, there was a higher probability of a higher net benefit being drawn for options C/D, than for option A. When the level of variation was reduced, the distributions became more tightly centred around the mean, and within a particular iteration, option A had a higher probability of having a higher net benefit than options C/D (see CEACs in Figure 5, panel I). Increasing the level of variation had the opposite effect on the CEACs (Figure 6, panel I), as those options which had a relatively high level of variation (e.g. options C and D) had the highest net benefit in a greater proportion of iterations, and thereby a greater probability of being cost-effective, than in the base-case, at the expense of the those options which had a relatively low level of variation (e.g. options A and B).

<< Figure 5 >>

<< Figure 6 >>

As reducing the level of variation increased the proportion of iterations where optimal options had the highest net benefit, the CEAF became less disjointed (Figure 5, panel II) and, as optimal options tended to have a higher probability of being cost-effective, the probability of making a wrong decision was consistently lower. This, combined with the lower level of variation in net benefit, which reduced the consequences of making a wrong decision, meant that the EVPI was also lower when the variation was reduced (Figure 5, panel II). However, the converse occurred when the level of variation increased – the CEAF became more disjointed, the probability of making the wrong decision increased, as did the consequences of making a wrong decision, and in turn the EVPI also increased (Figure 6, panel II).

DISCUSSION

We have demonstrated that CEACs can sometimes show i) that options which are subject to extended dominance have the highest probability of being cost-effective, or ii) that optimal options never have the highest probability of being cost-effective. Thus, we have shown that CEACs should not be used to determine the optimal decision, as previously suggested by Kamath et al. [3]. In addition, we have explained that these two effects can arise because of the correlation structure, and differences in the level of variation, between multiple options. Furthermore, in order to present the optimal decision, and the level of uncertainty associated with that decision, we have demonstrated that the CEAF, along with the EVPI, should be calculated. We have similarly explained how these can be constructed and interpreted, and that it is, for example, possible for uncertainty to increase (as assessed by the EVPI) when the probability of the optimal option being cost-effective also increases.

Implications

In line with others [28, 29] we have shown that the correlation structure between the costs and effects of different options can greatly influence the level of uncertainty. However, we have shown that the presentation of results on the cost-effectiveness plane is insensitive to the correlation structure. Thus, in order to simultaneously present the optimal decision, along with the uncertainty associated with that decision, the CEAF and the EVPI should be plotted. Indeed, the only time it may be appropriate to present uncertainty on the cost-effectiveness plane is when two options are being compared. This arises, because, in contrast to the situation for multiple options, the incremental cost and incremental effects can be plotted thus directly revealing the correlation structure.

A second implication is that, as previously suggested [22, 24, 30], instead of using the CEAC to estimate which option is optimal, the sole purpose of presenting the CEAC is to demonstrate the level of uncertainty associated with an option. Moreover, we have also shown that, if our objective is to maximise the expected net benefit regardless of the level of uncertainty, as argued previously by Claxton [31], then this can result in choosing options where we know there is a high probability of making a wrong decision. For example, in the study by Goeree et al. [2], at a WTP of $12,000, option A was optimal despite a probability that this was the wrong decision that exceeded 80% (see Figure 3, panel II).

Strengths and weaknesses

A crucial assumption within this paper is that optimal decisions will be made when options with the highest expected net benefit are implemented (i.e. when benefits are maximised for a given cost). Coast [15] has, however, argued that such an objective does not comply with how society would wish to allocate scarce healthcare resources, as, for example, it ignores the notion of equity. Palmer & Smith [32] have also shown that this objective ignores the extent to which the decision can be deferred or reversed. Both these points are valid however, neither change the fact that it is *possible* for an optimal option (unless optimal is defined in terms of probability) not to have the highest probability of being cost-effective, and thus the CEAC should not be used to determine the optimal allocation of resources.

We are only aware of one previous possible explanation as to why an optimal option may not have the highest probability of being cost-effective. Namely, that the incremental net benefit was positively skewed [24]. However, as this was not the case in the Goeree et al. [2] analysis, positive skewness could not account for why it had previously been shown that an option which was subject to extended dominance had, at certain WTP values, the highest probability of being cost-effective, nor could it account for why an option which was cost-effective never had the highest probability of being cost-effective [2]. Thus, the main strength of this paper is to demonstrate that these instances can be brought about by particular correlation structures, or by differences in the relative levels of variance, between multiple options. By constructing a template we have also demonstrated that the shapes of the CEACs, the CEAF, and the EVPI are highly dependent upon assumptions made about these variables. The template is based on the assumption of multivariate normality, thus, it may not be appropriate to use this template to replicate studies where there is substantial departure from this distribution. It is however possible to use the Cholesky decomposition technique to incorporate correlations between other types of distributions, see Briggs et al. [21] for further explanation.

CONCLUSION

The cost-effectiveness plane has limited use in representing uncertainty in the multiple option case as it cannot represent correlation between the options, which can play a very important role. CEACs can represent decision uncertainty, but should not be used to determine the optimal decision. Instead, the CEAF shows the optimal choice and this can be augmented using the EVPI plot to show the potential gains to further research to reduce uncertainty.

Acknowledgements: We thank the authors of a previous publication (Goeree R, O'Brien B, Blackhouse G, Marshall J, Briggs A, & Lad R) for providing us with the results of their analyses, which were used to inform some of the analyses in this paper.

REFERENCES

1. Konski A, Watkins-Bruner D, Feigenberg S et al. Using decision analysis to determine the cost-effectiveness of intensity-modulated radiation therapy in the treatment of intermediate risk prostate cancer. Int J Radiat Oncol Biol Phys 2006; 66: 408-415.
2. Goeree R, O'Brien BJ, Blackhouse G et al. Cost-effectiveness and cost-utility of long-term management strategies for heartburn. Value Health 2002; 5: 312-328.
3. Kamath CC, Kremers HM, Vanness DJ et al. The cost-effectiveness of acetaminophen, NSAIDs, and selective COX-2 inhibitors in the treatment of symptomatic knee osteoarthritis. Value Health 2003; 6: 144-157.
4. Fenwick E, Marshall DA, Levy AR et al. Using and interpreting cost-effectiveness acceptability curves: an example using data from a trial of management strategies for atrial fibrillation. BMC Health Serv Res 2006; 6: 52.
5. Noyes K, Dick AW, Holloway RG et al. Pramipexole and levodopa in early Parkinson's disease: dynamic changes in cost effectiveness. Pharmacoeconomics 2005; 23: 1257-1270.

6. Hoch JS, Rockx MA, Krahn AD. Using the net benefit regression framework to construct cost-effectiveness acceptability curves: an example using data from a trial of external loop recorders versus Holter monitoring for ambulatory monitoring of "community acquired" syncope. BMC Health Serv Res 2006; 6: 68.
7. 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: 80-90.
8. Muntz R, Hutchings J, Edwards R-T et al. Economic evaluation of treatments for children with severe behavioural problems. J Ment Health Policy Econ 2004; 7: 177-189.
9. Ramsey SD, Berry K, Etzioni R et al. Cost effectiveness of lung-volume-reduction surgery for patients with severe emphysema. N Engl J Med 2003; 348: 2092-2102.
10. O'Reilly J, Lowson K, Young J et al. A cost effectiveness analysis within a randomised controlled trial of post-acute care of older people in a community hospital. BMJ 2006; 333: 228-232.
11. Weintraub WS, Barnett P, Chen S et al. Economics methods in the Clinical Outcomes Utilizing percutaneous coronary Revascularization and Aggressive Guideline-driven drug Evaluation (COURAGE) trial. Am Heart J 2006; 151: 1180-1185.
12. Iglesias C, Nixon J, Cranny G et al. Pressure relieving support surfaces (PRESSURE) trial: cost effectiveness analysis. BMJ 2006; 332: 1416-1420.
13. Fenwick E, Palmer S, Claxton K et al. An iterative Bayesian approach to health technology assessment: Application to a policy of pre-operative optimization for patients undergoing major elective surgery. 2006; 26: 480-496.
14. Legood R, Gray A, Wolstenholme J et al. Lifetime effects, costs, and cost effectiveness of testing for human papillomavirus to manage low grade cytological abnormalities: results of the NHS pilot studies. BMJ 2006; 332: 79-85.
15. Coast J. Is economic evaluation in touch with society's health values? BMJ 2004; 329: 1233-1236.
16. Groot Koerkamp B, Hunink M, Stijnen T et al. Limitations of acceptability curves for presenting uncertainty in cost-effectiveness analysis. Med Descis Making in press;
17. Weinstein MC. Principles of cost-effective resource allocation in health care organizations. Int J Technol Assess Health Care 1990; 6: 93-103.
18. Cantor SB. Cost-effectiveness analysis, extended dominance, and ethics: a quantitative assessment. Med Decis Making 1994; 14: 259-265.

19. Claxton K, Sculpher M, Drummond M. A rational framework for decision making by the National Institute For Clinical Excellence (NICE). Lancet 2002; 360: 711-715.
20. Sculpher M, Claxton K. Establishing the cost-effectiveness of new pharmaceuticals under conditions of uncertainty--when is there sufficient evidence? Value Health 2005; 8: 433-446.
21. Briggs AH, Sculpher MJ, Claxton K. Decision modelling for health economic evaluation. New York: Oxford University Press, 2006.
22. van Hout BA, Al MJ, Gordon GS et al. Costs, effects and C/E-ratios alongside a clinical trial. Health Econ 1994; 3: 309-319.
23. Briggs AH, Goeree R, Blackhouse G et al. Probabilistic analysis of cost-effectiveness models: choosing between treatment strategies for gastroesophageal reflux disease. Med Decis Making 2002; 22: 290-308.
24. Fenwick E, Claxton K, Sculpher MJ. Representing uncertainty: the role of cost-effectiveness acceptability curves. Health Econ 2001; 10: 779-787.
25. Stinnett AA, Mullahy J. Net health benefits: a new framework for the analysis of uncertainty in cost-effectiveness analysis. Med Decis Making 1998; 18: S68-S80.
26. Claxton K. Bayesian approaches to the value of information: implications for the regulation of new pharmaceuticals. Health Econ 1999; 8: 269-274.
27. Bellamy N, Buchanan WW, Goldsmith CH et al. Validation study of WOMAC: a health status instrument for measuring clinically important patient relevant outcomes to antirheumatic drug therapy in patients with osteoarthritis of the hip or knee. J Rheumatol 1988; 15: 1833-1840.
28. Ades AE, Lu G. Correlations between parameters in risk models: estimation and propagation of uncertainty by Markov Chain Monte Carlo. Risk Anal 2003; 23: 1165-1172.
29. Chessa AG, Dekker R, van Vliet B et al. Correlations in uncertainty analysis for medical decision making: an application to heart-valve replacement. Med Decis Making 1999; 19: 276-286.
30. Fenwick E, O'Brien BJ, Briggs AH. Cost-effectiveness acceptability curves--facts, fallacies and frequently asked questions. Health Econ 2004; 13: 405-415.
31. 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.
32. Palmer S, Smith PC. Incorporating option values into the economic evaluation of health care technologies. J Health Econ 2000; 19: 755-766.

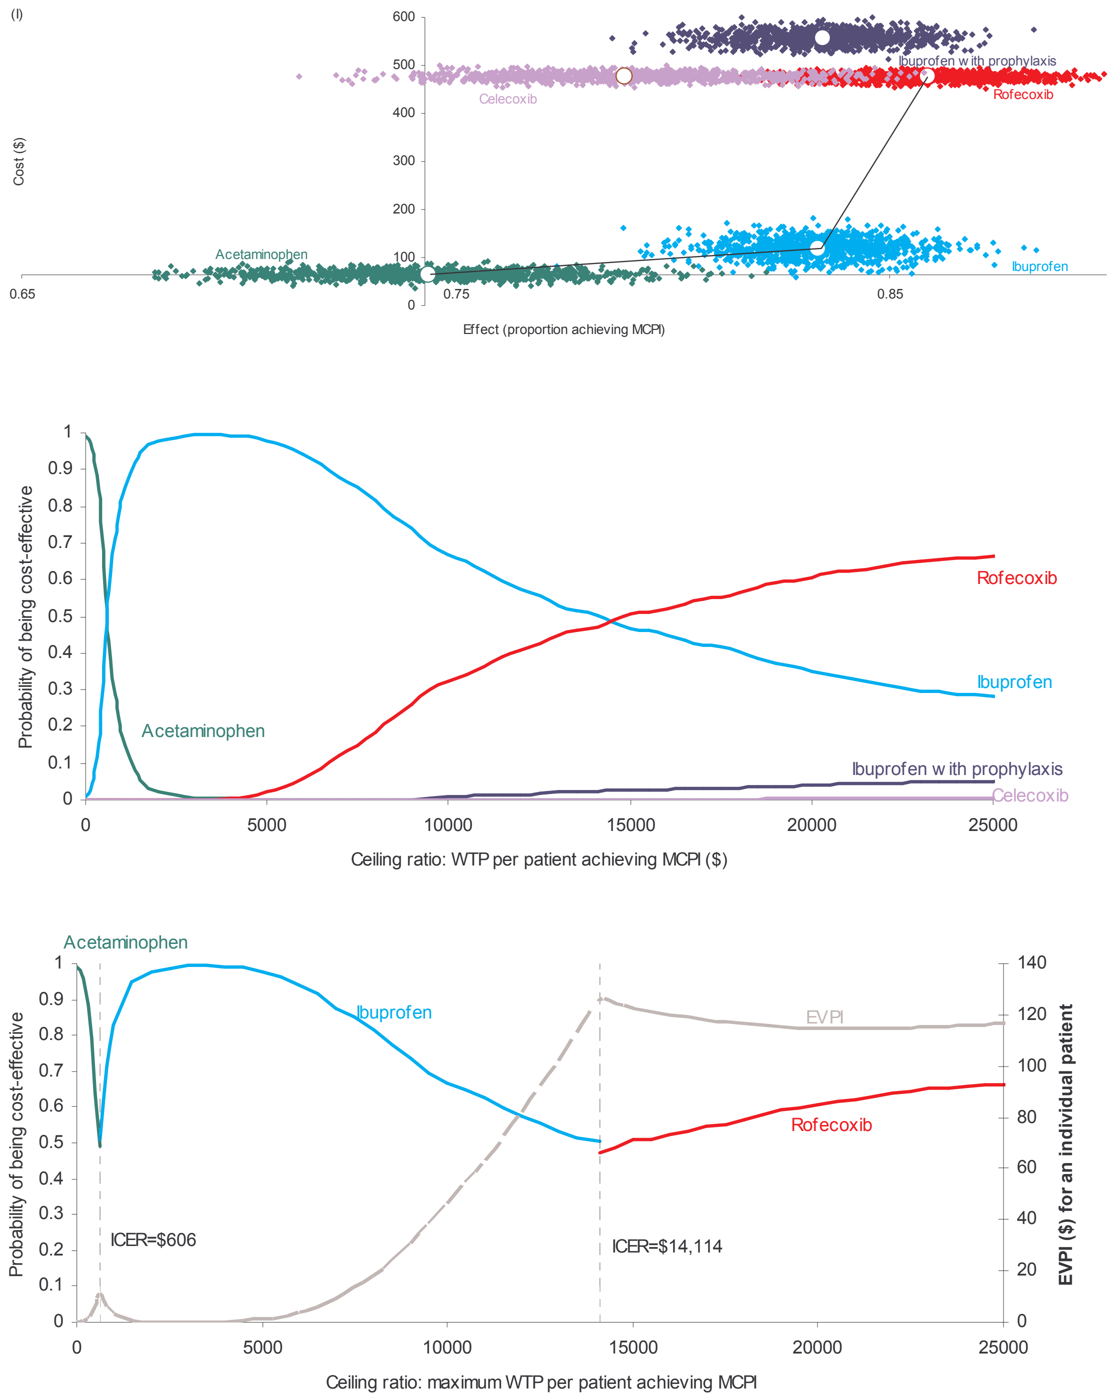

Figure 1. Replication of the analysis by Kamath et al. [3] presented in terms of the cost-effectiveness plane (panel I), cost-effectiveness acceptability curves (CEACs) (panel II), and the cost-effectiveness acceptability frontier (CEAF) and the expected value of perfect information (EVPI) (panel III).

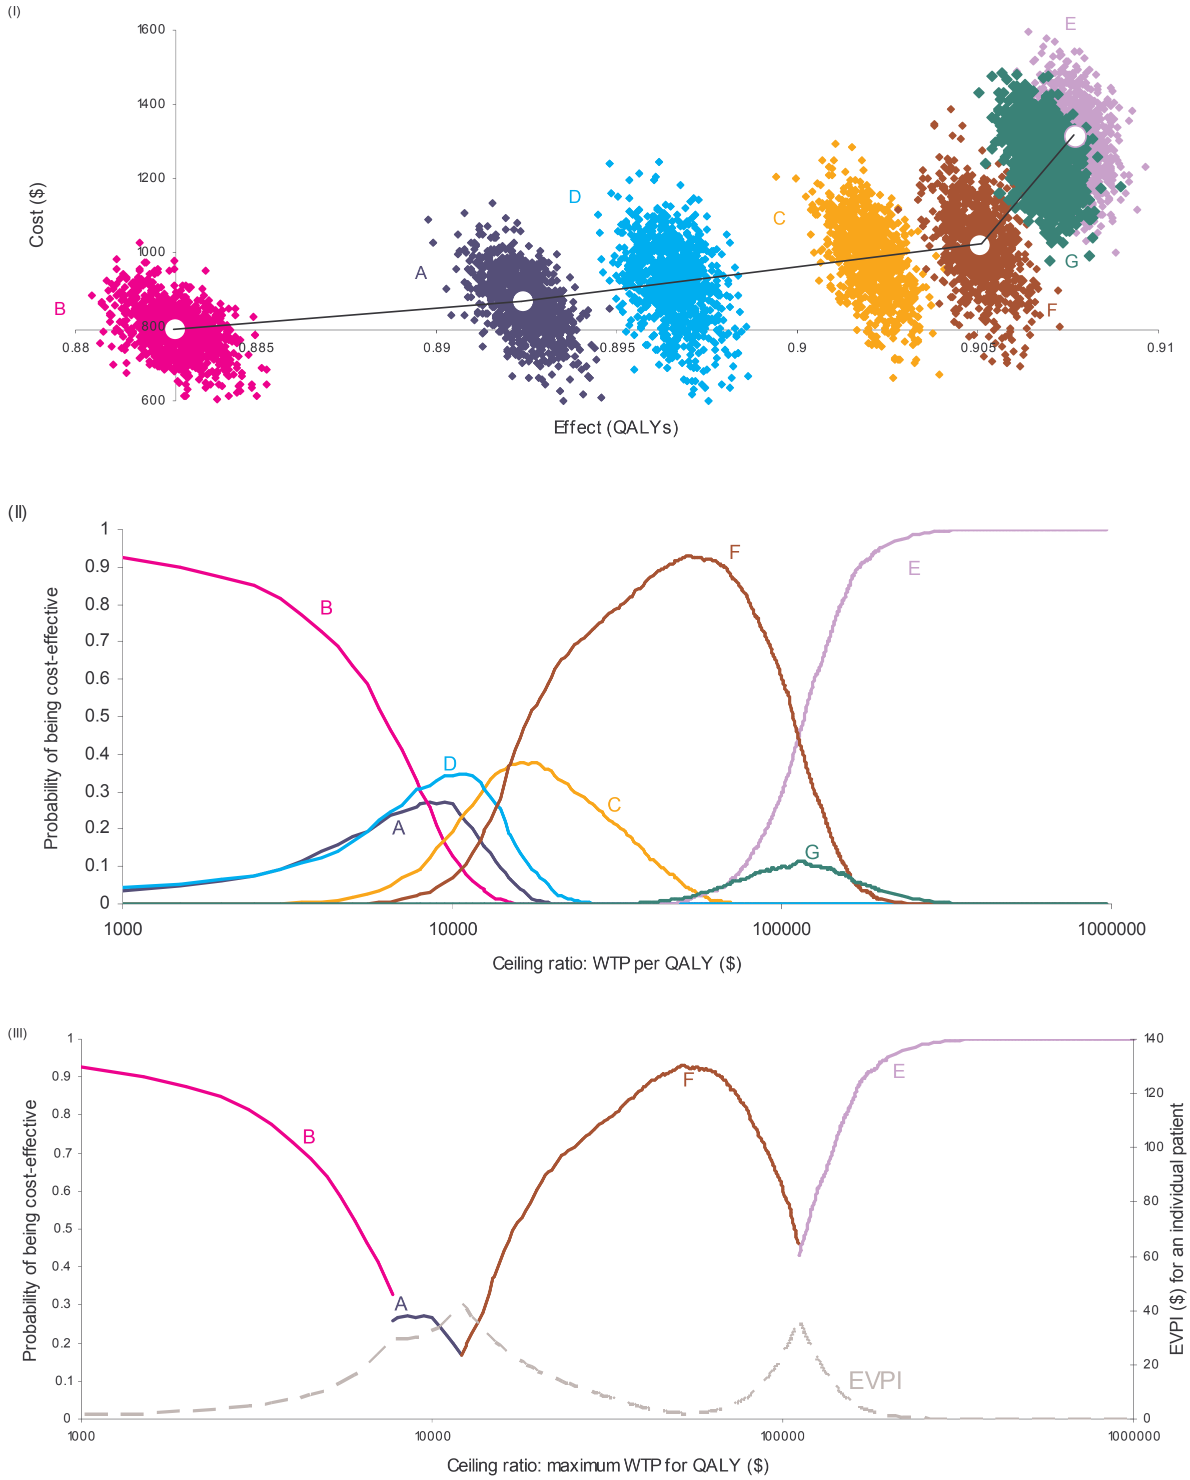

Figure 2. Replication of the analysis by Goeree et al. [2] presented in terms of the cost-effectiveness plane (panel I), cost-effectiveness acceptability curves (CEACs) (panel II), and the cost-effectiveness acceptability frontier (CEAF) and the expected value of perfect information (EVPI) (panel III).

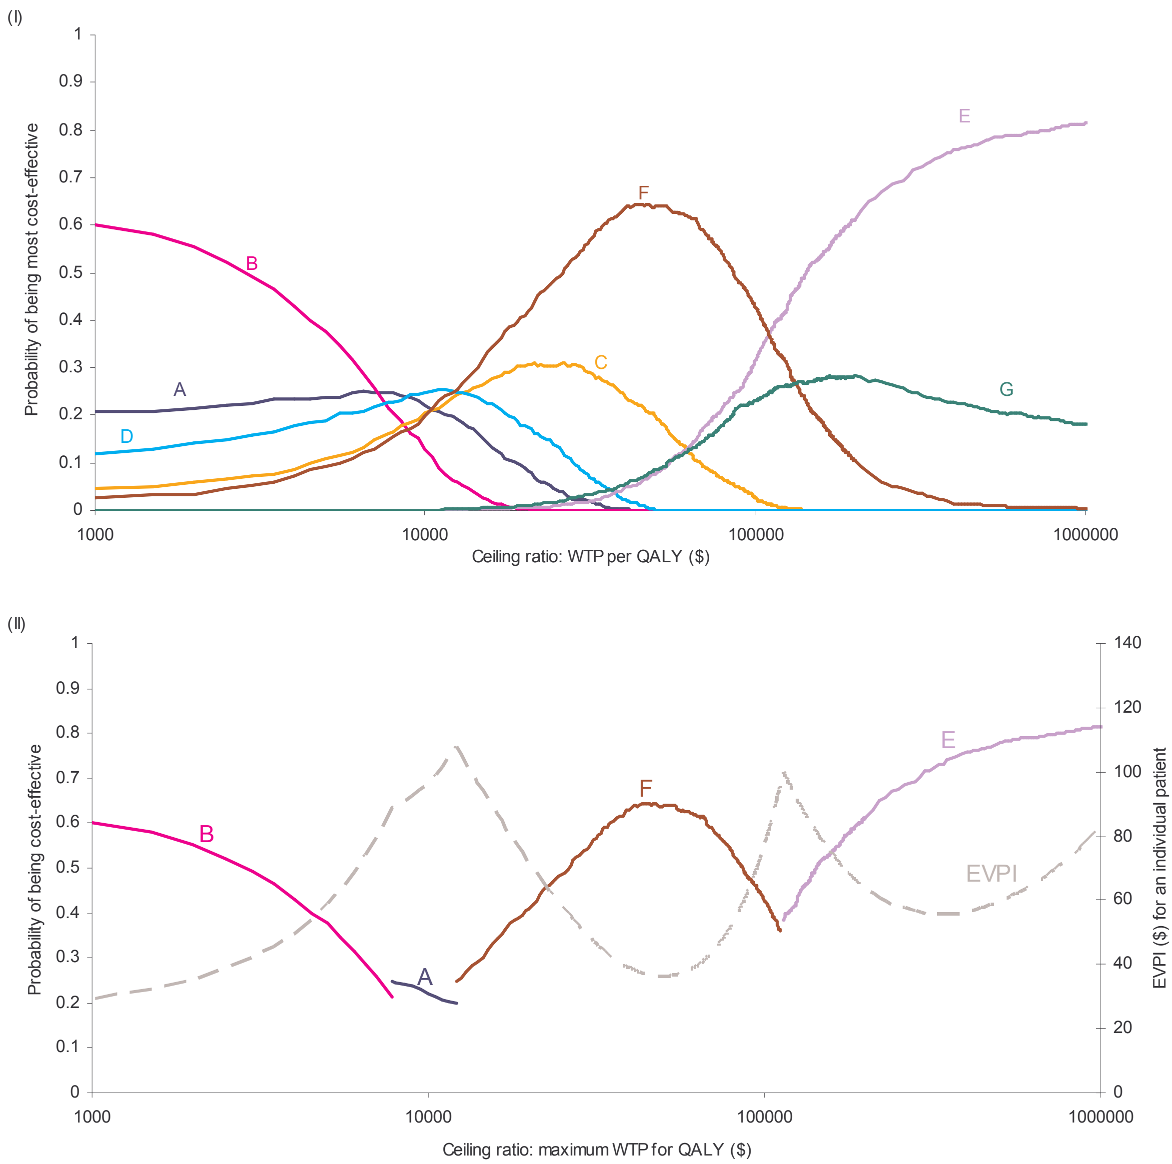

Figure 3. Replication of the analysis by Goeree et al. [2], with the assumption that the costs and effects between options were independent, presented in terms of cost-effectiveness acceptability curves (CEACs) (panel I), and the cost-effectiveness acceptability frontier (CEAF) and the expected value of perfect information (EVPI) (panel II).

Figure 4. Replications of the analysis by Goeree et al. [2] presented on the cost-effectiveness plane, with the assumptions that the level of variation was one fifth of that (panel I), and five times that (panel II), in the base-case analysis.

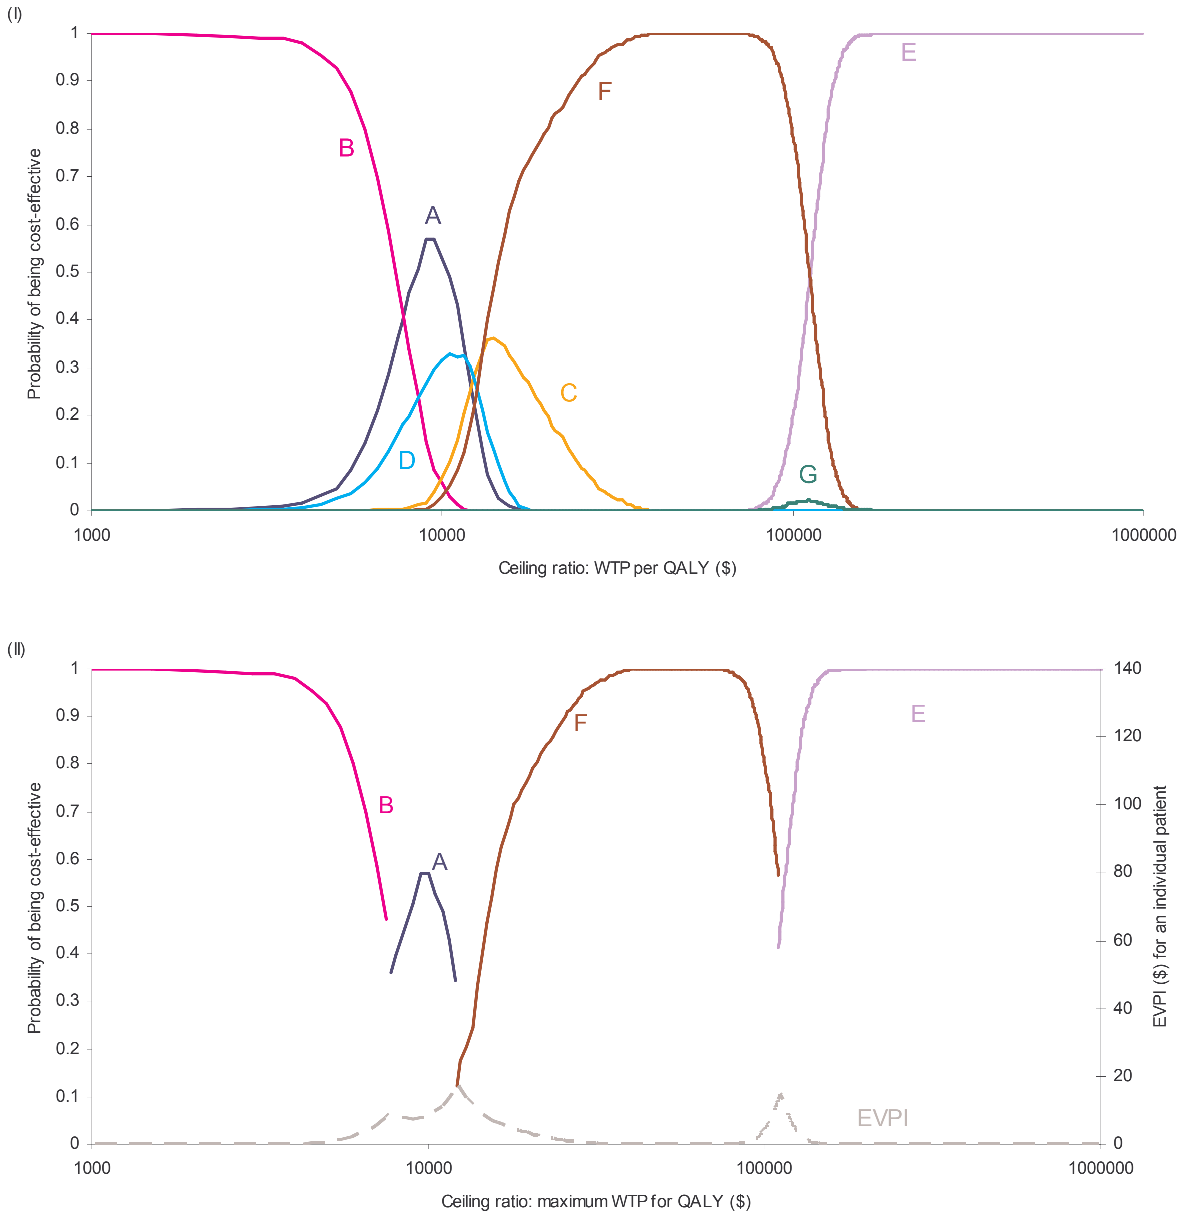

Figure 5. Replications of the analysis by Goeree et al. [2] the assumption that the variance associated with the mean cost and mean effect of each option was one fifth of that in the base-case analysis, results are presented in terms of cost-effectiveness acceptability curves (CEACs) (panel I), and the cost-effectiveness acceptability frontier (CEAF) and the expected value of perfect information (EVPI) (panel II).

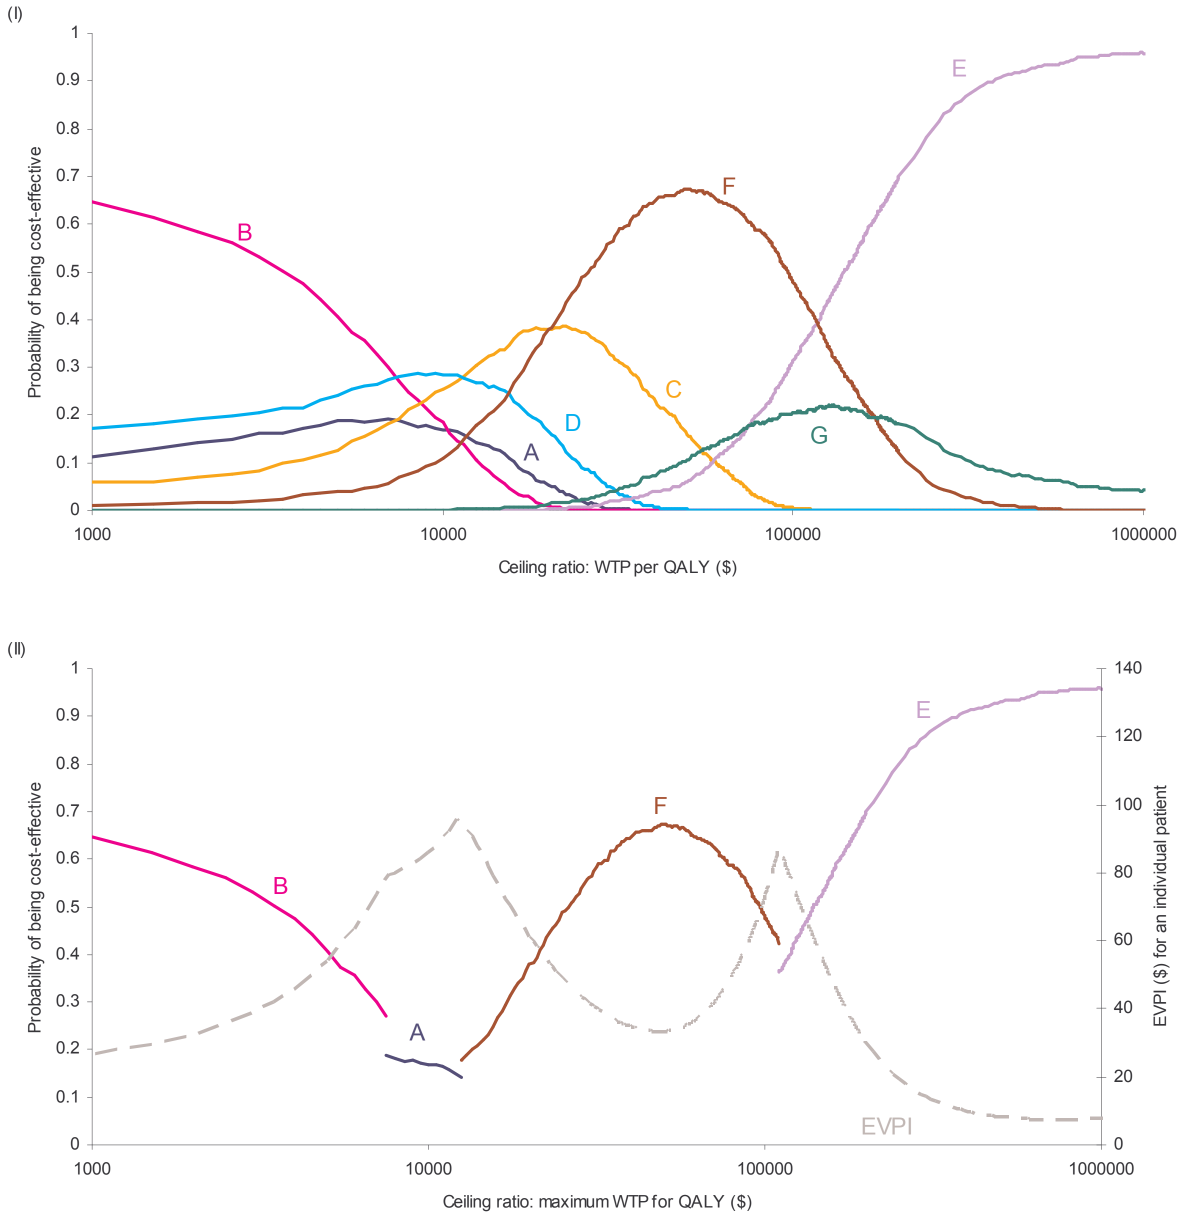

Figure 6. Replications of the analysis by Goeree et al. [2] the assumption that the variance associated with the mean cost and mean effect of each option was five times that in the base-case analysis, results are presented in terms of cost-effectiveness acceptability curves (CEACs) (panel I), and the cost-effectiveness acceptability frontier (CEAF) and the expected value of perfect information (EVPI) (panel II).