Health Economists' Study Group January 1999

Estimating uncertainty ranges for incremental costs by combining probabilistic sensitivity analysis with non-parametric bootstrapping

Joanne Lord

Health Care Evaluation Unit, Department of Public Health Sciences, St. George's Hospital Medical School, Cranmer Terrace, London SW17 ORE.

j.lord@sghms.ac.uk

ABSTRACT

In the past, economic evaluations were usually deterministic, with a single point estimate of incremental cost or cost-effectiveness. The impact of uncertainty over underlying data and assumptions was tested through various forms of sensitivity analysis.

Health economists are increasingly basing their cost estimates on individual resource utilisation data for a sample of patients. In such cases, statistical analysis can be used to obtain interval, as well as point, estimates of costs. With stochastic outcome data, interval estimates for cost-effectiveness ratios can also be obtained. This is not straightforward, because of the positive skew in resource utilisation data, and the difficulty in estimating confidence intervals for ratios. There is currently much interest in methods for overcoming these problems, including non-parametric bootstrapping. However, confidence intervals derived by such methods do not represent the totality of uncertainty. Uncertainty over the values ascribed to unit costs, the choice of discount rate, and other modelling assumptions are also important.

This paper uses cost data from a trial of patient education for osteo-arthritis of the knee to illustrate a method for combining uncertainty estimates for sampled and non-sampled data. Monte Carlo simulation with non-parametric bootstrapping and probabilistic sensitivity analysis (or parametric bootstrapping) is used to derive 95% "Uncertainty Ranges" for incremental costs. This method is easily extended to incremental cost-effectiveness ratios.

INTRODUCTION

In the past, economic evaluations of health care were usually deterministic.[1] Costs and cost-effectiveness ratios were presented as single point estimates representing average results for patients within the relevant population. The effects of uncertainty over the values of the models' input parameters were tested, if at all, through various forms of sensitivity analysis.[2,3] This is the only possible approach when data is not available at an individual patient level.

Increasingly, individual data on resource utilisation or costs is being collected for samples of patients.[4] Such data may be collected along with clinical outcome data as part of a controlled trial, or it may be collected separately to feed into a secondary economic analysis. The presence of sampled data provides an opportunity to move economic evaluation beyond deterministic analysis.[5] However, stochastic economic analysis is not simple. First, there is the problem of skewness in resource utilisation, and hence cost, data - a large proportion of costs are often due to a small proportion of patients.[1] This means that care has to be taken in calculating conventional confidence limits, or in hypothesis testing.[6] Second, sample sizes in clinical trials are usually based on clinical outcomes, and are often inadequate for economic analysis because of the large variation in cost data.[1] Third, the aim of economic evaluation is often to estimate a ratio of random variables, such as the incremental cost-effectiveness ratio. Here methods of traditional statistical inference can not be used. There is a growing literature on methods for overcoming these problems, including non-parametric bootstrapping.[7-16]

A further difficulty with stochastic economic evaluation arises because calculation of the desired result almost invariably requires the use of estimated parameters alongside sampled data. For instance, costs are not directly measurable. Instead they are estimated by mulitplying unit costs by resource quantities. The unit costs are usually treated as fixed parameters, though they might be seen as stochastic variables along with resource quantities and clinical outcomes.[5] Even when costs are estimated by adjusting individual patient charges for profits, unit costs are implicit in the institutional billing system(s). Other modelling parameters are often required to estimate missing resource utilisation data, or to extrapolate results over time or to broader patient populations. There is uncertainty surrounding each of these parameters. In addition, there is uncertainty around the choice of the basic analytical methods, which are embodied in the functional form of the model. Thus, estimates of uncertainty based on variation in sampled data do not capture the true extent of our uncertainty.

So, the introduction of statistical analysis into economic evaluations will not be sufficient, sensitivity analysis will still be required. This paper suggests a method through which statistical analysis of sampled data can be combined with probabilistic sensitivity analysis of model parameters. The aim is to estimate an uncertainty range that is closer to the true level of uncertainty. This is done through Monte Carlo simulation, using both non-parametric bootstrappings of sampled data and probabilistic sensitivity analysis (or parametric bootstrapping) of model parameters.

The approach is illustrated using cost data from a randomised controlled trial of patient education for osteo-arthritis of the knee.[16] No significant difference was found between the intervention and control groups in clinical outcomes at one year. The economic evaluation thus aimed to test whether there was any difference between the groups in net healthcare costs (a cost-minimisation study). An uncertainty range is estimated for the incremental cost using the combined parametric/non-parametric bootstrap method. Though the example is limited to cost data, the method can also be applied to an incremental cost-effectiveness ratio.

The combined parametric/non-parametric bootstrap algorithm

Let $\mathbf{x}_1$ and $\mathbf{x}_0$ be observed data for samples of patients from the intervention group and control group, respectively. Each $\mathbf{x}_i$ is a $(1 \times n_i)$ vector of data points, where n_i is the number of patients in group i. The data points may themselves be vectors rather than single numbers. For example, they might include quantities of various resources used by a given patient over a given time interval. Or they might include resource utilisation for given periods before and after randomisation. They might also include clinical outcome data.

Suppose that we are seeking to estimate some economic variable, ϑ. This might be the incremental cost, incremental cost-effectiveness, or net benefit. ϑ is estimated by some function (P) of the observed data and of some set of parameters (Γ).

$$\hat{\theta} = P(\mathbf{x}_0, \mathbf{x}_1, \Gamma)$$

Again, the Γ might include a wide range of parameters, such as unit costs, health state valuations and the discount rate. For the moment these parameters are set at their reference case levels, and treated as fixed.

Most simply, if ϑ is the incremental cost, $\bar{\mathbf{x}}_0$ and $\bar{\mathbf{x}}_1$ are (1 x m) vectors of mean resource quantities used by the control and intervention groups, and Γ is the (m x 1) vector of fixed unit costs:

$$\hat{\vartheta} = \bar{\mathbf{x}}_1 * \Gamma - \bar{\mathbf{x}}_0 * \Gamma$$

The non-parametric bootstrap method[17] can be used to obtain an interval estimate for ϑ. Random samples of size n_0 and n_1 are taken from the data sets $\mathbf{x}_0$ and $\mathbf{x}_1$, respectively, with replacement. Data for some patients might be included in this bootstrap sample, ($\mathbf{x}_0^*$, $\mathbf{x}_1^*$), more than once, whilst data for other patients might not be included at all. A bootstrap replicate of ϑ is then estimated using this bootstrap sample:

$$\vartheta^* = P(\mathbf{x}_0^*, \mathbf{x}_1^*, \Gamma)$$

This process is then repeated a large number of times, B, to obtain a set of bootstrap replicates, ϑ_k^* (k=1,2,...,B). A simple interval estimate is obtained by taking the $100(\alpha/2)$th to $100((1-\alpha)/2)$th percentiles of the bootstrap replicates (where α is the significance level). However, the coverage of this simple percentile interval is poor. An improved bias-corrected and accelerated (BC_a) interval estimate is obtained by adjusting the upper and lower percentile cut-off points.[17]

Even the improved BC_a interval, however, does not allow for uncertainty over the parameter values, Γ. This can be introduced through parametric bootstrapping. First, prior probability distributions are assigned to each of the parameters. These distributions should represent current knowledge and beliefs about the parameter values. B sets of parameter values Γ° are obtained by random sampling (with replacement) from the prior distributions. Each of these *parametric* bootstrap samples is then combined with a *non-parametric* bootstrap sample of the data to provide a replicate of ϑ:

$$\vartheta^{*\circ} = P(\mathbf{x}_0^*, \mathbf{x}_1^*, \Gamma^\circ)$$

An interval estimate is then obtained by taking a BC_a percentile range of these revised replicates.

The question remains of how large B should be. The accuracy of bootstrap estimates depends on two factors: sampling variation, which falls as the sample sizes, n_1 and n_0, increase; and bootstrap re-sampling variation, which falls as B increases. Efron and Tibrishani recommend B=1,000 for BC_a percentile interval estimates, though fewer iterations are necessary for standard error estimates.[17] Bootstrap replicates of ratio variables, such as cost-effectiveness, can be highly volatile, leading to big jumps in standard error estimates as

B increases. However, the BC_a percentile method for interval estimation smoothes out these fluctuations, and 1,000 replications is probably still sufficient.[14] The effect of adding in variability from the prior parameter distributions not clear. One thousand replications might still be adequate, or additional replications might be required. A check can be performed by plotting the bootstrap interval estimates against increasing B. If they appear to stabilise by B=1,000, then this is probably large enough. If not, more bootstrap samples can be added.

Implementing this combined parametric/non-parametric bootstrap algorithm is computationally demanding, but feasible with modern personal computers and software. A spreadsheet can be used, though for more complex models a statistical package with programming capacity or a basic programming language would be more manageable. Care must be taken over the quality of the random number generator, as these do vary between proprietary software packages. The period of the generator must be less than the total number of random numbers required.

Example - The OAK Study

Background

The St. George's Osteo-Arthritis of the Knee (OAK) study, was a randomised controlled trial of patient education.[16] A number of studies have suggested that patient education can be beneficial in a range of chronic conditions[18,19], including arthritis.[20,21] The aim of such educational programmes is to enable patients to better manage their own conditions - to use medication and health services more appropriately, to use pain management and joint protection techniques, relaxation and exercise. However, the quality of evidence for the efficacy of patient education for OA is not good.

General practices were recruited from a South London health authority area. They were randomised to the intervention or control group (n=10 and 12 respectively). Patients with radiographic signs of knee OA were recruited in these practices between November 1995 and May 1997. In all, 105 patients were recruited from the intervention practices, and 65 from the control practices. The patients were assessed at recruitment by a baseline interview. Patients in the intervention practices were then invited to take part in four, one-hour group sessions, at weekly intervals. Up to six patients attended each group. The sessions were held at the GPs' surgeries and were led by a research nurse. All patients were followed up by postal questionnaire at one, three, six and twelve months. Health outcome instruments used included:

the Western Ontario and McMaster Universities Arthritis Index (WOMAC), the Arthritis Helplessness Index (AHI), the SF-36 and GHQ.

125 (74%) of the 170 patients with a baseline interview returned a full outcome questionnaire at one year, and a further 18 (11%) returned an abbreviated questionnaire. The control practices had significantly more partners and patients than the intervention practices, but there were no significant differences in follow-up rates or baseline socio-demographic characteristics, clinical indicators or outcome scales.

Data was analysed by practice, since that was how randomisation was conducted[22]. Analysis was conducted on an intention-to-treat basis.[6] An overall deterioration in knee OA and general health was observed for the patients in both groups over one year. The only significant clinical difference between the groups was for the WOMAC stiffness score, which rose for the intervention practices (by 2.4 percentage points) and fell for the control group (by 6.9 percentage points) ($p<0.05$). Similar changes were seen for the WOMAC pain score, and for the SF-36 physical dimension, although these differences were not significant. It was concluded that the OAK intervention did not improve clinical outcomes - though one dimension of one scale did show an improvement, no adjustment had been made for multiple significance testing.

Method of economic evaluation

An economic evaluation of the OAK study was independently funded. The OAK baseline interview included some questions about the use of health and social care resources, and about the impact of knee OA on paid work. However, additional information was required for the economic evaluation. This was collected from two sources: firstly, patients were interviewed again in their own homes after their final outcome assessment at one year; secondly, GP case notes were reviewed. Information was collected for each patient, for each "cost-generating event", over a two year period (from one year before baseline to one year after). Cost-generating events included:

- Attendance at the OAK educational sessions.
- The issue of prescriptions or purchase of over-the-counter or complementary medications.
- Consultations with primary care staff.
- Hospital inpatient stays, outpatient attendances, day case treatments and emergency visits.

- Consultations with paramedical and complementary therapists.

Each event was categorised as: due to knee OA; partly due to knee OA; not due to knee OA; or unknown. However, this classification was problematic.

126 patients (74%) were interviewed at one year, and GP case notes were reviewed for 137 patients (81%). Patients were not included in the economic analysis if their notes were not available for review - since the bulk of cost-generating events were identified from this source. However, patients were included if their notes had been reviewed, even if they had not been interviewed at one year. Missing data might be expected to reduce overall costs - particularly since two patients died within the study period, and their notes were not available for review. However, there was no difference in the follow-up rate between the groups.

Unit costs were applied to each cost-generating event, and mean costs per patient calculated for each practice, for the two years. Net costs were calculated for the practices, by subtracting costs for the year before baseline from costs for the year after. The incremental cost was then estimated as the mean difference in practice net costs for the intervention group compared to the control group. No attempt was made to estimate cost-effectiveness, because of the poor evidence of any health improvement resulting from the intervention.

All costs were estimated in 1996/7 pounds sterling (adjusting for inflation if necessary). Costs were estimated from three perspectives: that of the NHS; that of private individuals; and that of society as a whole. Although, in this paper, only societal costs are reported. Both direct and indirect costs were estimated, the latter by applying the median national wage rate to time reported to have been taken off paid work. Patient transport costs were estimated by applying a standard mileage rate to estimated distances between the patients' homes and the health care facilities.

As far as possible, unit cost estimates for the reference case analysis were based on published national data (Table 1). Because of the very large number of individual unit cost estimates - 735 resources were included in the database, including 525 medications - it was not possible to assign separate prior distributions to each estimate. So, resources were grouped into seven main categories: the OAK intervention, medications, primary care services, hospital services, community services, patient time and transport. Prior probability distributions were then assigned to these seven resource groups (Table 2). Where data was available on the level of variation between providers or patients, this was used. Otherwise, estimates of variation were made by the author. The cost of medications was assumed to be fixed , because of national price agreements, although variation in pharmacy dispensing costs (including on-cost allowance, fees, container allowance and discounts were included).

The prior distributions were used, together with non-parametric bootstrapping of the resource utilisation data to estimate a 95% BC_a uncertainty range for the incremental costs.

Results

Figure 1 shows the monthly social cost per patient over the two year study period for the reference case analysis. Much of the observed volatility is due to a few inpatient episodes for individual patients. For example, the control group peak at month 18 was due to a single patient who had a fall, possibly caused by her knee OA, and had to spend over a month in hospital. It can be argued that inpatient costs should be excluded, since it is highly unlikely that the intervention would influence them. The monthly costs excluding costs relating to episodes of inpatient (or day-patient) care are shown in Figure 2. The remaining peak around month twelve for the intervention patients is due to the cost of the OAK educational sessions. Results are presented below both including and excluding inpatient costs. This illustrates the effect of differing data variability on combined parametric/non-parametric bootstrap uncertainty ranges. It also emphasises the point that bootstrapping does not remove the need for traditional simple sensitivity analysis to investigate the effect of methodological uncertainty.

Figure 3 and Figure 4 show net social costs per patient, by practice. Again, they demonstrate that removing inpatient costs greatly reduces variation. However, quite large differences between the individual practices remain. The means of these net social costs per patient for the intervention and control group practices are shown in Table 3. If we consider all health care costs, the estimated incremental cost is -£147, suggesting that the increase in costs for the intervention group is less than that for the control group. The 95% confidence interval for this estimate, based on the traditional statistical method, is wide -£539 to £245. If we remove inpatient and day-patient costs, the situation is changed, with an incremental cost of +£135 (£27 to £244) - there is a significant increase in costs for the intervention practices, but no change for the control practices.

The bootstrap replicates for the analyses including inpatient costs are shown in Figure 5. The BC_a interval based on non-parametric bootstrapping is -£650 to +£146. The difference between this and the conventional parametric estimate given above suggests that the cost data is not normally distributed, and that the sample sizes are not sufficient to appeal to the central limit theorem. Parametric bootstrapping (probabilistic sensitivity analysis) gives a much narrower interval estimate of -£133 to -£33. Combining the two methods gives a total uncertainty range of -£713 to +£167. Here the variability in the resource utilisation

data is much greater than that produced by the prior distributions. Figure 6 shows the equivalent bootstrap replicates if we exclude inpatient and day-case episodes. Here the BC_a interval is £39 to £253, which is close to the conventional parametric confidence interval. The parametric bootstrap interval is £62 to £213, and the combined interval is £20 to £291. Figure 7 and Figure 8 show the various BC_a interval estimates against increasing numbers of replications. In this case, one thousand replications appears to be sufficient, as all the estimates are relatively stable by B=1,000.

DISCUSSION

This example shows how interval estimates can be derived for incremental costs to reflect uncertainty due to sample variation and some of the modelling assumptions. This combined interval will be wider than estimates based on either source of variation alone. In the above example, the addition of parametric bootstrapping would not have changed the conclusion reached by non-parametric bootstrapping or conventional statistical analysis alone. However, in other situations it could well do so. The breadth of the interval estimate provides information about the power of the study. The very wide intervals for the OAK study, if all costs are included, suggest that the sample size is not sufficient for this analysis, because of the cluster randomisation.

It is not possible to build an interval that covers all possible sources of uncertainty. In particular, uncertainty due to basic analytical methods is not suitable for probabilistic analysis.[2] So, for example, the choice of whether to include all costs, or to restrict the analysis to certain types of cost, is outside the scope of the stochastic analysis. In the OAK study this is a crucial decision, since excluding inpatient costs changes the results from a non-significant negative incremental cost to a significant positive one. The appropriate way to investigate such methodological uncertainty is through thinking back to basic economic principles, an understanding of the health care context and simple sensitivity analysis.

The method described in this paper entails certain assumptions. First, it assumes that all the parameters are independent of one another. For the OAK analysis this suggests that, for example, the unit cost of primary care is not related to the cost of community or hospital services. In reality costs for different types of services within a given geographical area might well be inter-related. Second, the method assumes that the parameters are independent of the sampled data. Thus the quantity of resources used must not be related to the unit cost of those resources. However, if the price of one resource is high, healthcare professionals or

patients might well substitute some other resource.[5] Third, the method applies one set of parameter estimates to all patients in each bootstrap sample. This implies that the same parameter values apply to all patients. In practice, parameters, such as unit costs, might well vary between patients or between health care institutions. All of these assumptions could be relaxed by building a more complicated model, with joint probability distributions for example. It is not clear how much difference a more sophisticated approach would make in practice.

CONCLUSION

A recent article in the British Medical Journal criticised the quantity and quality of statistical analysis in published economic evaluations: "*Economic outcomes should be evaluated with the same statistical standards that are now expected for clinical outcomes*".[23] There are aspects of economic evaluation that make statistical analysis particularly difficult: cost data is often skewed; sample sizes are often insufficient for the large variations in costs; analysis of complex variables such as incremental cost-effectiveness ratios can not be conducted with conventional statistical methods; and uncertainty over the values of modelling parameters and assumptions must be considered. Progress has been made in the economic literature on overcoming the first three of these obstacles. This paper presents a method for tackling the fourth problem. By using a mixture of parametric and non-parametric bootstrapping, an interval estimate can be obtained that reflects uncertainty over both sampled and non-sampled data. Further work is required to investigate the behaviour of these combined uncertainty ranges with different data sets. The effect of some of the simplifying assumptions employed in the method described here also needs to be tested.

REFERENCES

1 Briggs A, Gray A. The distribution of health care costs and their statistical analysis for economic evaluation. *Journal of Health Service Research and Policy* 1998;**3**:233-245.

2 Briggs A, Sculpher MJ, Buxton MJ. Uncertainty in the economic evaluation of health care technologies: the role of sensitivity analysis. *Health Economics* 1994;**3**:95-104.

3 Briggs A, Sculpher MJ. Sensitivity analysis in economic evaluation: a review of published studies. *Health Economics* 1995;**4**:355-371.

4 Adams ME, McCall NT, Gray DT, Orza MJ, Chalmers TC. Economic analysis in randomized control trials. *Medical Care* 1992;**30**:231-243.

5 Drummond MF, O'Brien BJ. Clinical importance, statistical significance and the assessment of economic and quality-of-life outcomes. *Health Economics* 1993;**2**:205-212.

6 Bland JM. *An introduction to medical statistics*. Oxford: Oxford University Press, 1995.

7 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. *Medical Care* 1994;**32**:150-163.

8 van Hout BA, Maiwen JA, Gordon GS, Rutten FFH. Costs, effects and c/e-ratios alongside a clinical trial. *Health Economics* 1994;**3**:309-319.

9 Rutten-van Molken MP, van Doorslaer EK, van Vliet RC. Statistical analysis of cost outcomes in a randomized controlled clinical trial. *Health Economics* 1994;**3**:333-345.

10 Waaker P, Klaassen MP. Confidence intervals for cost/effectiveness ratios. *Health Economics* 1995;**4**:373-381.

11 Chaudhary MA, Stearns SC. Estimating confidence intervals for cost-effectiveness ratios: an example from a randomized trial. *Statistics in Medicine* 1996;**15**:1447-1458.

12 Willan AR, O'Brien BJ. Confidence intervals for cost-effectiveness ratios: an application of Fieller's theorem. *Health Economics* 1996;5:297-305.

13 Polsky D, Glick HA, Willke R, Schulman K. Confidence intervals for cost-effectiveness ratios: a comparison of four methods. *Health Economics* 1997;**6**:243-252.

14 Briggs A, Wonderling DE, Mooney CZ. Pulling cost-effectiveness analysis up by its bootstraps: a non-parametric approach to confidence interval estimation. *Health Economics* 1997;**6**:327-340.

15 Tambour M, Zethraeus N. Bootstrap confidence intervals for cost-effectiveness ratios: some simulation results. *Health Economics* 1998;**7**:143-147.

16 Lord J, Victor C, Littlejohns P, Ross F, Axford J. Economic evaluation of a primary care based education programme for patients with osteoarthritis of the knee. Final report submitted to the NHS R&D National Coordinating Centre for Health Technology Assessment.

17 Efron B, Tibshirani R. *An introduction to the bootstrap*. New York: Chapman and Hall, 1993.

18 Mazzuca SA. Does patient education in chronic disease have therapeutic value? *Journal of Chronic Diseases* 1982;**35**:521-529.

19 Gibson PG, Coughlan J, Wilson AJ, et al. *The effects of limited (information only) patient education programs on the health outcomes of adults with asthma*. Cochrane Database of Systematic Reviews, 1997.

20 Mullen PD, Laville EA, Biddle AK, Lorig K. Efficacy of psychoeducational interventions on pain, depression, and disability in people with arthritis: a meta-analysis. *Journal of Rheumatology* 1987;**14 Suppl 15**:33-39.

21 Hirano PC, Laurent DD, Lorig K. Arthritis patient education studies, 1987-1991: a review of the literature. *Patient Education & Counseling* 1994;**24**:9-54.

22 Kerry SM, Bland JM. Analysis of a trial randomised in clusters. *BMJ* 1998;**316**:54.

23 Barber JA, Thompson SG. Analysis and interpretation of cost data in randomised controlled trials: review of published studies. *BMJ* 1998;**317**:1195-1200.

24 British Medical Association, Royal Pharmaceutical Society of Great Britain. *British National Formulary*. London: British Medical Association, 1996.

25 Prescription Pricing Authority. *Annual report*. Newcastle upon Tyne: PPA, 1996.

26 Bellamy N, Buchanan W, Goldsmith C, Campbell J, Stitt L. Validation study of WOMAC: a health status instrument for measuring clinically important patient relevant outcomes to anti-rheumatic drug therapy in patients with osteoarthritis of the hip or knee. *Journal of Rheumatology* 1988;**15**:1833-1840.

27 Netten A, Dennett J. *Unit costs of health and social care 1997*. Canterbury: Personal Social Services Research Unit, University of Kent, 1997.

28 Office for National Statistics. *New Earnings Survey 1996*. London: HMSO, 1996.

Table 1. Sources of societal unit cost estimates for the reference case analysis

Resource category	Source
OAK sessions	Estimate based on trial experience. Including nurse time (for development, training, organisation and running the group sessions), the use of GP surgery rooms, travel and consumables. Costs were allocated over 174 patients, 34 groups and 20 practices - the numbers that could be recruited within a single district within one year.
Prescription medications	BNF net price[24] plus average dispensing costs[25]
Over-the-counter medicines and complementary medicines	Same as prescribed medication (if in BNF) otherwise retail price less VAT
Hospital outpatient and inpatient services	HFMA/CIPFA estimates[26]
Hospital day case treatment	PSSRU estimates[27]
Primary and community health services, social services, and aids and appliances	PSSRU estimates[27]
Complementary therapists	Reported patient expenditure
Patient time	National median wage rates[28]
Transport	Mileage rate for NHS staff reported in Netten & Dennett[27]

Table 2. Prior probability distributions

Parameter	Cumulative probability distribution	Source
OAK sessions (£ per patient)	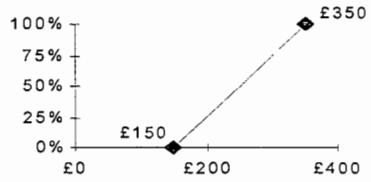	Estimated by the author (reference case ± £100)
Dispensing cost (% of BNF net price)	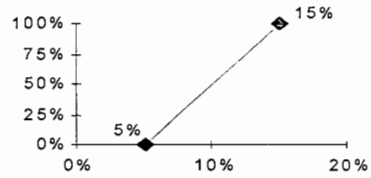	Estimated by the author (reference case ± 5%)
Cost of hospital care (% of reference costs)	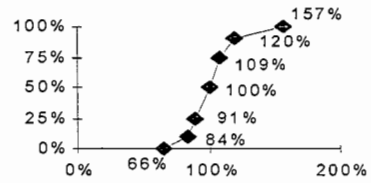	Distribution (quartiles and upper and lower deciles) of an index of day-case and inpatient Trust prices (not adjusted for market forces). National Reference Cost Index, NHS Executive 1998.
Cost of community services (% of reference costs)	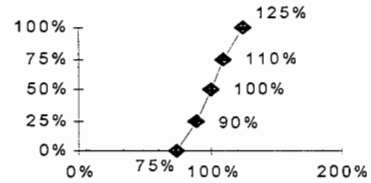	Distribution of unit labour costs for community staff. Harper J. in Netten & Dennett 1997.
Cost of primary care (% of reference costs)	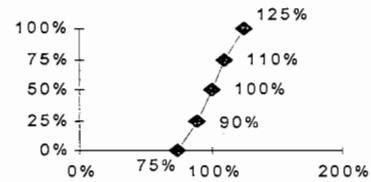	Assumed to be the same as community health services, in the absence of any other information.
Patient wage (£ per hour)	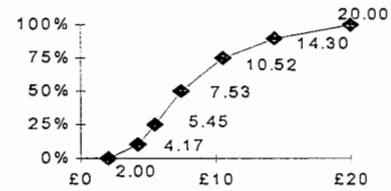	New Earnings Survey 1996.
Patient transport (£ per mile)	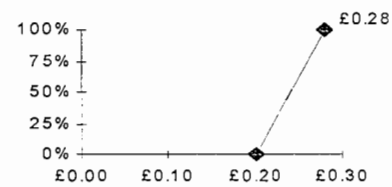	Estimated by author (reference case ± 4p per mile)

Table 3. Net social cost per patient for the reference case analysis

	Mean for control practices (n=11)			Mean for intervention practices (n=9)			Mean difference		
	Year before	Year after	Change	Year before	Year after	Change	Year before	Year after	Change
OAK sessions	£0	£0	£0	£11	£186	£175	£11	£186	£175
Medications	£149	£144	-£5	£147	£143	-£4	-£2	-£1	£1
IP and DP care	£25	£225	£199	£76	£90	£14	£50	-£135	-£185
OP care	£103	£88	-£15	£81	£96	£15	-£22	£8	£30
Community services	£20	£58	£38	£64	£35	-£29	£44	-£23	-£67
GP care	£78	£63	-£15	£87	£85	-£2	£9	£22	£13
Other	£36	£21	-£15	£47	£14	-£33	£11	-£6	-£17
Transport	£9	£8	-£1	£8	£7	-£1	£0	£0	£0
Indirect	£18	£111	£93	£36	£32	-£4	£18	-£79	-£97
Including costs for all health care categories									
TOTAL	£438	£717	£279	£558	£689	£132	£120	-£27	-£147
SD	£264	£583	£621	£286	£262	£212			
Lower CL	£282	£372	-£88	£371	£518	-£7	-£123	-£412	-£539
Upper CL	£594	£1,061	£646	£745	£860	£270	£363	£357	£245
Excluding inpatient and day-case costs									
TOTAL	£395	£381	-£13	£446	£568	£122	£51	£187	£135
SD	£223	£198	£141	£168	£142	£105			
Lower CL	£263	£264	-£97	£336	£475	£53	-£120	£37	£27
Upper CL	£526	£498	£70	£555	£660	£191	£223	£336	£244

Figure 1. Social cost per patient by month - including all costs

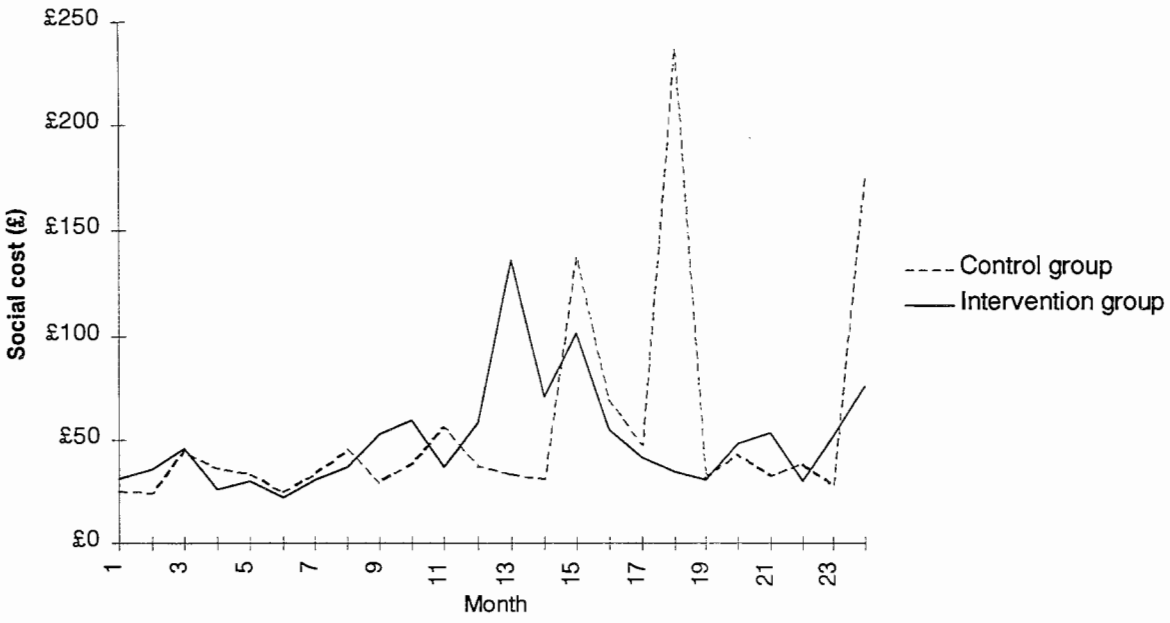

Figure 2. Social cost per patient by month - excluding inpatient and day-case costs

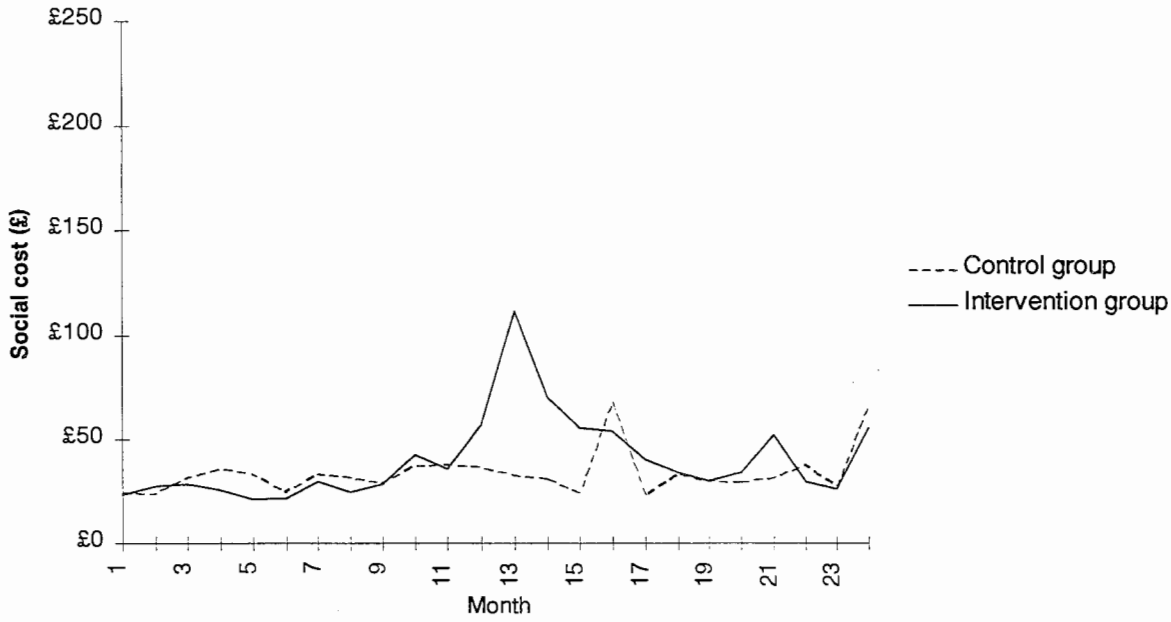

Figure 3. Net Social cost per patient by practice - including all costs

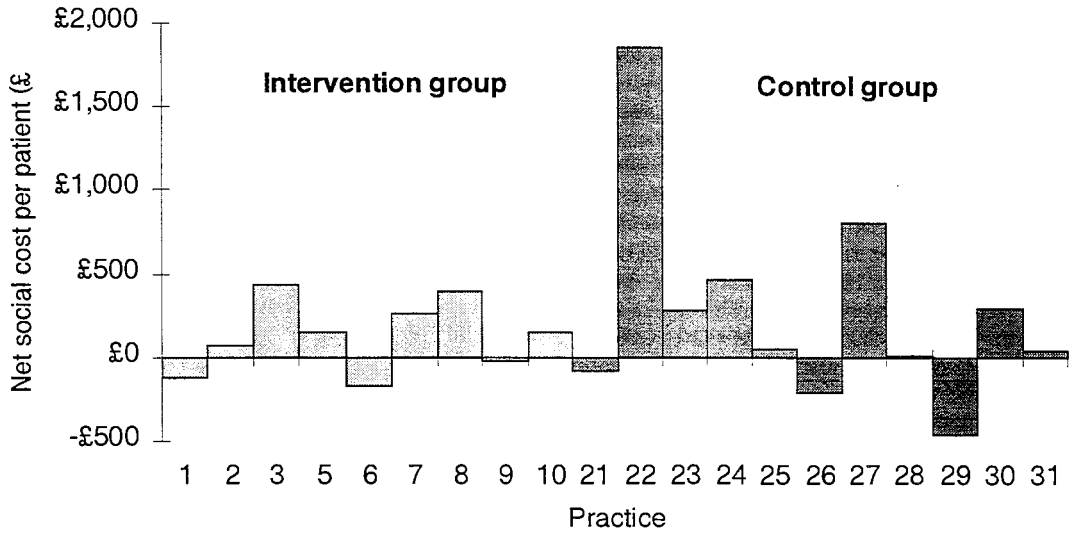

Figure 4 Net Social cost per patient by practice - excluding inpatient and day-case costs

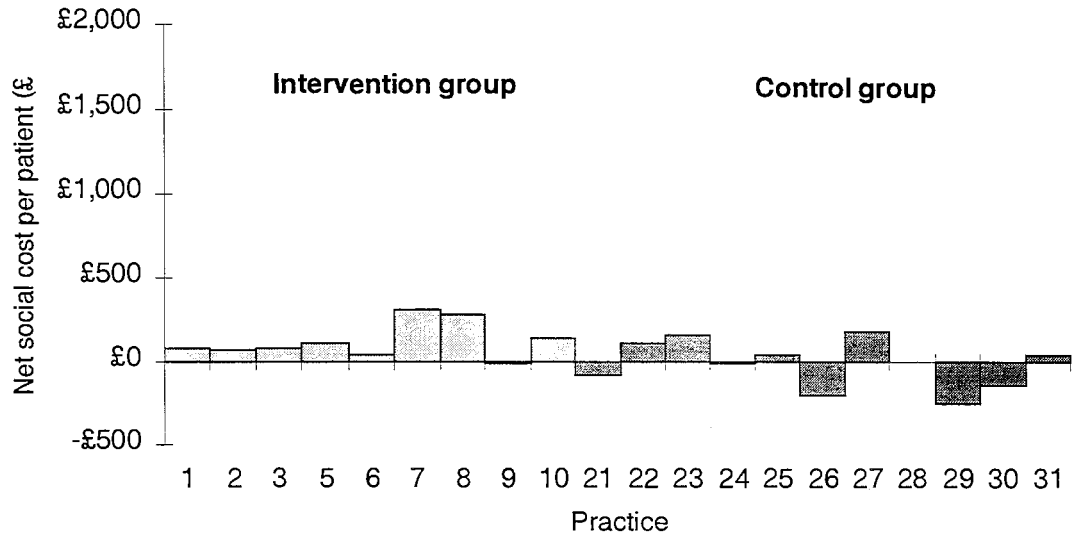

Figure 5. Distribution of bootstrap replicates - including all costs

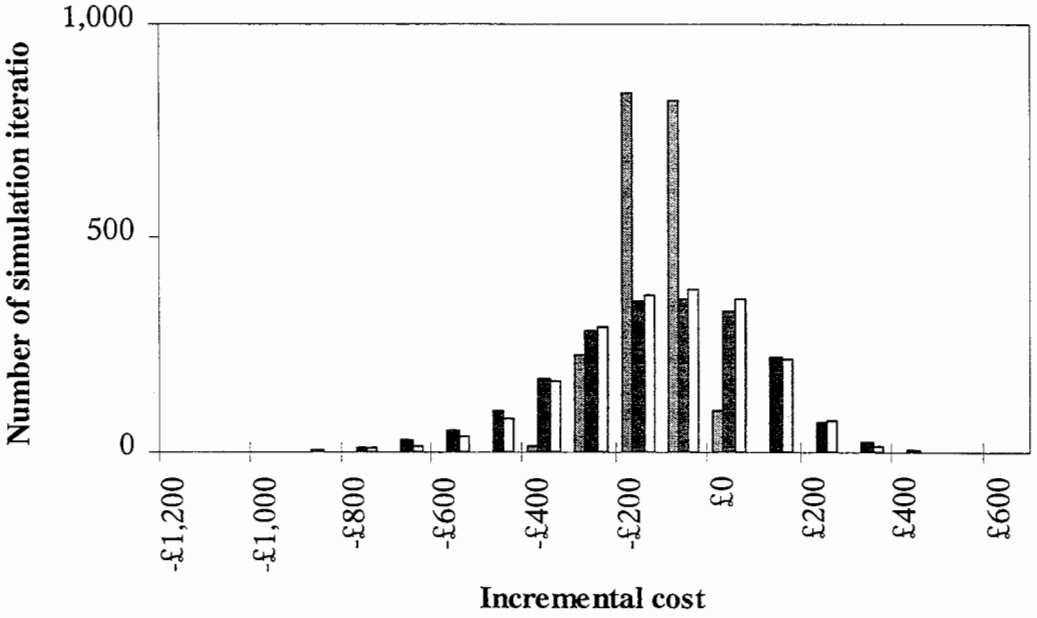

Figure 6. Distribution of bootstrap replicates - excluding inpatient and day-case costs

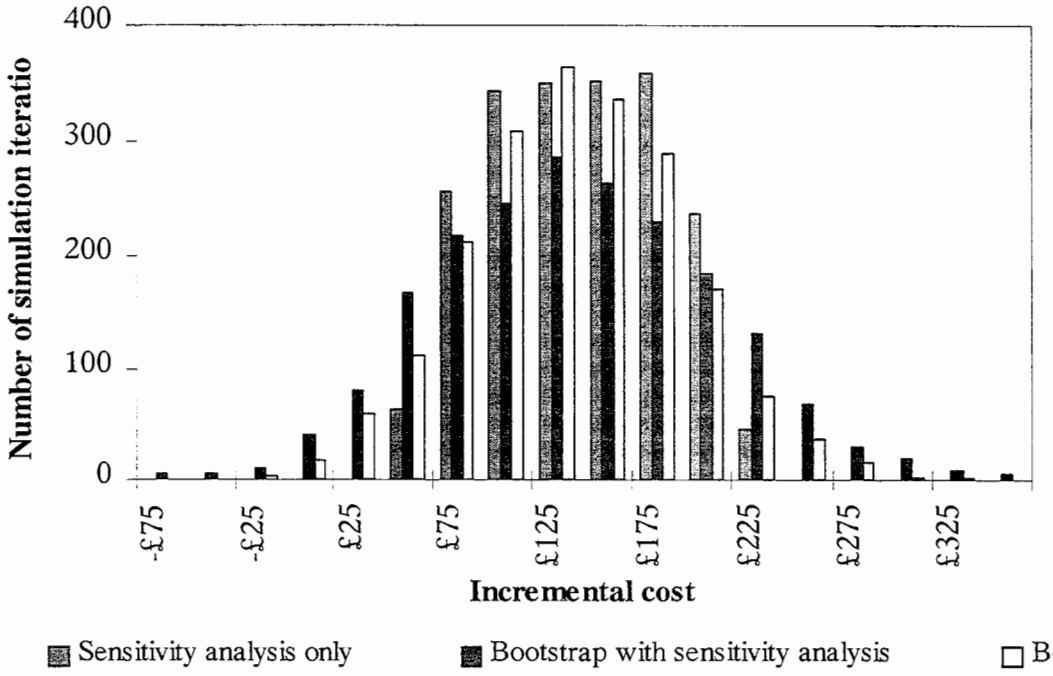

Figure 7. BC_a interval estimates by number of replications - all costs

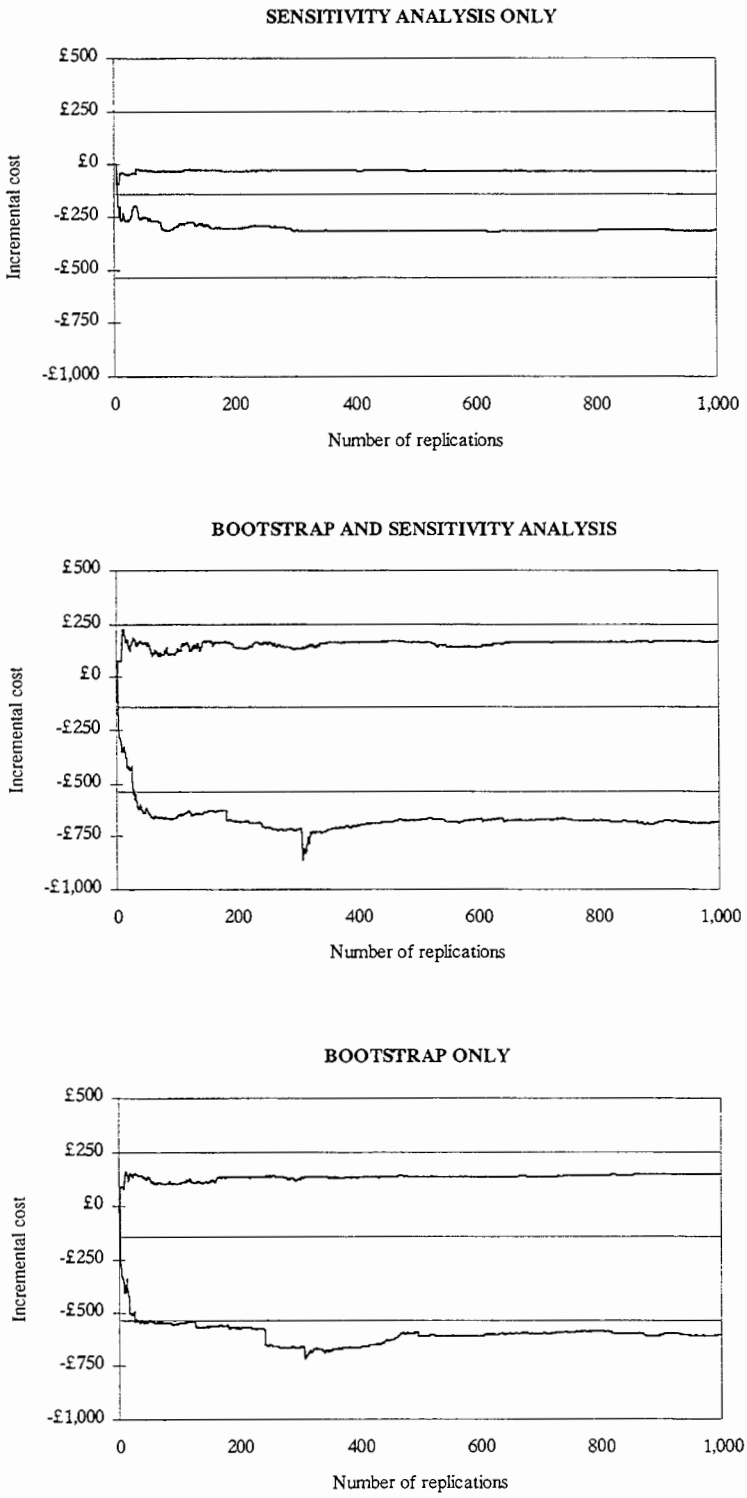

Faint lines show the reference case mean and 95% confidence limits estimated by the conventional parametric statistical method.

Figure 8. BC_a percentile interval estimates by the number of replications - excluding inpatient costs

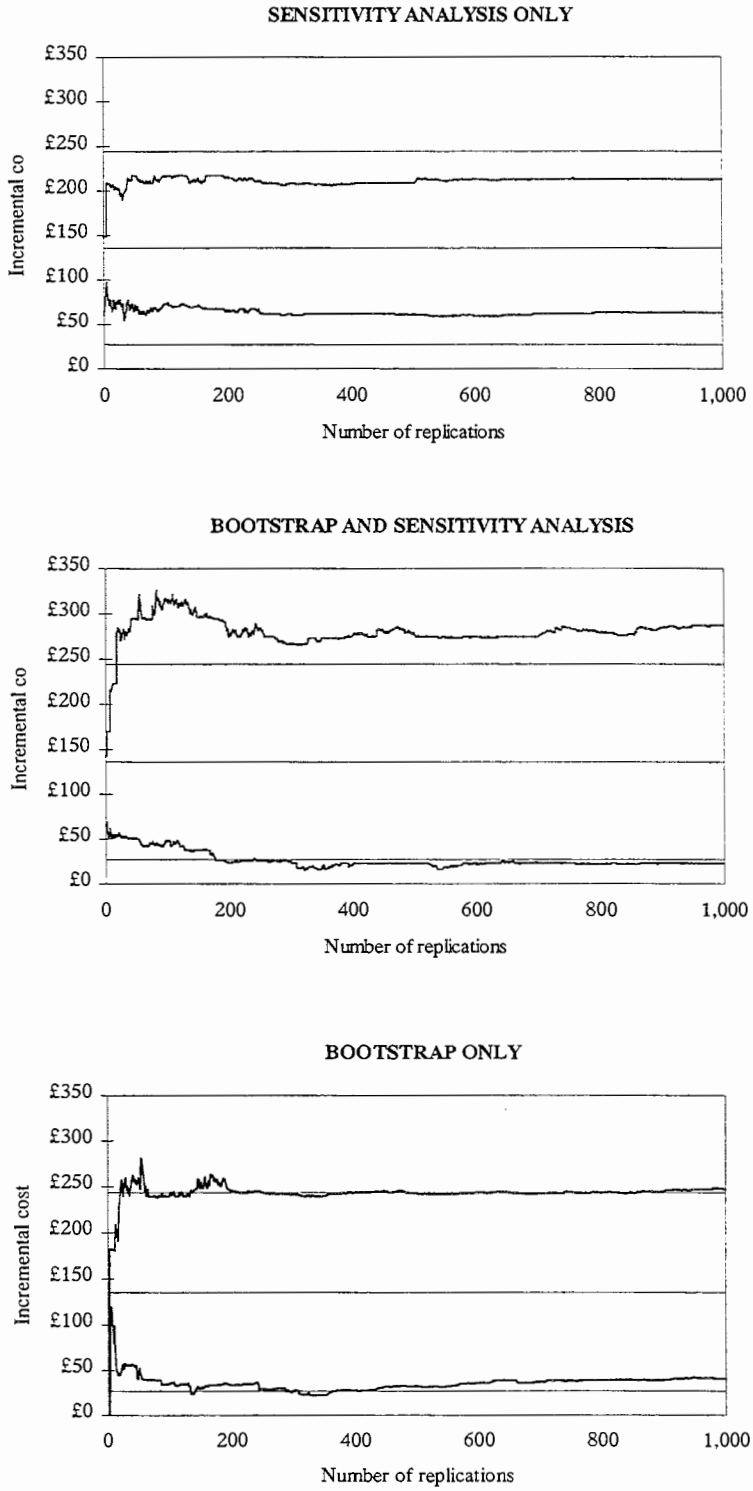

Faint lines show the reference case mean and 95% confidence limits estimated by the conventional parametric statistical method.