Calculating Partial Expected Value Of Perfect Information in Cost-Effectiveness Models via a Two Level Monte-Carlo Algorithm .

Alan Brennan, MSc[(a)] Samer Kharroubi, PhD[(b)] Anthony O'Hagan, PhD[(b)] Jim Chilcott, MSc[(a)]
(a) ScHARR, (b) Probability and Statistics, The University of Sheffield, England.

INTRODUCTION

Quantifying expected value of information is important for developers and users of decision models. It can be used in decision-making, and to guide planning and prioritising further data collection. Many guidelines for cost-effectiveness analysis recommend multi-way probabilistic sensitivity analysis[1,2]. Expected value of perfect information (EVPI) is a natural extension of probabilistic sensitivity analysis and has been recommended as the most coherent methodology[3,4]. In planning and prioritising future data collection, EVPI calculations can be used to show whether further data collection on uncertain parameters would provide benefit to decision makers by reducing decision uncertainty and to quantify the maximum potential value of further data collection on different subsets of model parameters.

The few published EVI case studies can leave analysts confused on the correct approach[5]. Analysts who confidently calculate cost-effectiveness acceptability curves, still do not use EVPI. Early literature on EVPI in healthcare[6] used stylised examples to calculate EVPI analytically. Felli and Hazen[7] gave a fuller exposition, setting out some mathematics using expected value notation, attempting a general Monte Carlo simulation procedure for partial EVPI calculation, and setting out a "shortcut" 1 level algorithm for use in some circumstances. In the late 1990s, a series of EVPI case studies in the UK used a different 1 level algorithm to calculate partial EVPI[8,9,10]. Other recent papers focus on partial EVPI as a sensitivity analysis measure, comparing with alternative 'importance' measure approaches[11,12,13,14].

Our aim is to clarify the calculation of partial EVPI and encourage its use health economic decision models. We assess alternative methods of calculating partial EVPI, clarify the correct approach via mathematical description, and investigate the mathematical conditions when alternative 'short-cut' approaches may be used appropriately. We use case studies to demonstrate the extent to which 'short-cut' algorithms approximations might be accurate. In particular, we investigate 3 cases:- firstly, where the cost-effectiveness model is made up of simple sum-products of its statistically independent parameters (such as simple decision tress); secondly, where the cost-effectiveness model has increasing levels correlation between model input parameters, and finally, where the cost-effectiveness model has increasing levels of non-linearity (such as many period Markov models). Because the correct algorithm is relatively computationally intensive, we also assess whether relatively small numbers of iterations are inherently biased and investigate the number of iterations, which must be undertaken to ensure accuracy.

METHODS

A systematic review of the literature on the use of modelling in planning and prioritising clinical trials was undertaken[15]. In contrast to systematic reviews of health technologies, this methodology review required both broad searching of publication databases and a method of "pearl growing", which uses the references and citations of core method papers to find the relevant literature. Methodology and applied studies were retrieved. Recent conference presentations on EVI methods were also reviewed. The analysis of the literature focussed on the algorithms for undertaking EVPI calculations, the number of simulations required and the mathematical description of the processes. Following the literature review, we developed mathematical descriptions for each of the various EVPI calculation methods found. A description using expected value notation allowed comparison of the mathematical validity and equivalence, or otherwise, of the algorithms. A description using summation notation allowed an understanding of their multi-level loop nature.

To investigate the effects of the different approaches, we developed three case studies. The first, and simplest case study is a cost-effectiveness model, characterised by sum-products of statistically independent parameters i.e. a simple decision tree. The model compares two drug treatments T0 and T1 (Figure 1). Costs for each strategy are for drugs and hospitalisations (e.g. central estimate costT0 = "cost of drug", plus "percentage admitted to hospital" x "days in hospital" x "cost per day" = £1,000 + 10% x 5.20 x £400 = £1,208). QALYs gained come from response and side-effect decrement (e.g. central estimate QALYT0 = "% responding" x "utility improvement for a specified duration", and "% side-effects" x "a utility decrement for a specified duration" = 70% responders x 0.3 x 3 years + 25% side effects x –0.1 x 0.5 years = 0.6175). The decision threshold cost per QALY (λ=£10,000) enables net benefits calculations (central estimate NetBenT0 = 10,000* 0.6175 – £1,208 = £4,967. The 19 uncertain model parameters are characterised with independent normal distributions with a mean (columns a, b) and a standard deviation (columns d, e). Monte Carlo simulations are undertaken using EXCEL formulae (=NORMINV(RAND(), mean, standard deviation) and loop macros to record the input and output values for each sample – similar to producing a cost-effectiveness plane[16] or a cost-effectiveness acceptability curve[17]. A single sample result is shown in columns f and g (cost difference =£629, QALY = -0.2446, and so net benefit in favour of T0 in this sample). Averaging 1,000 sample results (Figure 1b) shows that the baseline decision given current information is to adopt strategy T1. The cost effectiveness plane shows uncertainty to be larger in QALY differences than in costs, whilst the CEAC at threshold = £10,000 shows T1 to be cost-effective with a probability of 54.5%.

Case study 2 uses the same model, but investigates correlation. We examined 5 different levels of correlation (0, 0.1, 0.2, 0.3, 0.6) between 4 different sets of parameters. Firstly, positive correlations are

Figure 1: Case Study 1 Model Framework and Basic Results

Part a: Model Framework

Illustrative Model

	a	b	c	d	e
Number of Patients in the UK	1,000			Uncertainty in Parameter Means	
Threshold cost per QALY	£10,000			Standard Deviations	Sampled Values

Mean Model	T0	T1	Increment (T1 over T0)	T0	T1	T0	T1	Increment (T1 over T0)
Cost of drug	£ 1,000	£ 1,500	£ 500	1	1	£ 1,000	£ 1,499	£ 499
% admissions	10%	8%	-2%	2%	2%	10%	11%	0%
Days in Hospital	5.20	6.10	0.90	1.00	1.00	5.01	7.05	2.05
Cost per Day	£ 400	£ 400	£ -	200	200	£ 589	£ 589	£ -
% Responding	70%	80%	10%	10%	10%	88%	70%	-18%
Utility change if respond	0.3000	0.3000	-	0.1000	0.0500	0.3119	0.2120	- 0.0998
Duration of response (years)	3.0	3.0	-	0.5	1.0	3.1	4.1	1.0
% Side effects	25%	20%	-5%	10%	5%	31%	24%	-7%
Change in utility if side effect	-0.10	-0.10	0.00	0.02	0.02	-0.06	-0.14	-0.08
Duration of side effect (years)	0.50	0.50	-	0.20	0.20	0.70	0.65	- 0.05
Total Cost	£ 1,208	£ 1,695	£ 487			£ 1,308	£ 1,937	£ 629
Total QALY	0.6175	0.7100	0.0925			0.8394	0.5948	-0.2446
Cost per QALY	£ 1,956	£ 2,388	£5,267			£ 1,558	£ 3,256	-£2,570
Net Benefit of T1 versus T0	**£ 4,967**	**£ 5,405**	**£437.80**			**£7,086**	**£4,011**	**-£3,075**

Part b: Summary Results

Results Statistics	Mean of 1,000 Samples	Central Estimate Using Prior Means
Mean Cost T0	£1,210	£1,208
Mean Cost T1	£1,692	£1,695
Mean QALY T0	0.6216	0.6175
Mean QALY T1	0.7044	0.7100
Mean Net Benefit T0	£5,006	
Mean Net Benefit T1	£5,352	
Mean Incremental Cost (T1 v T0)	£483	£487
Mean Incremental QALY (T1 v T0)	0.08280	0.0925
Mean Incremental Cost Per QALY (T1 v T0)	£5,833	£5,267
Mean Incremental Net Benefit (T1 v T0)	£ 345	£437
Percentage of Monte Carlo Samples When T1 is The Correct Decision	54.5%	-
Expected Net Benefit Achievable if Perfect Information Were Available	£6,704	-
Value of Perfect Information	£1,352	-

Part c: Cost Effectiveness Plane

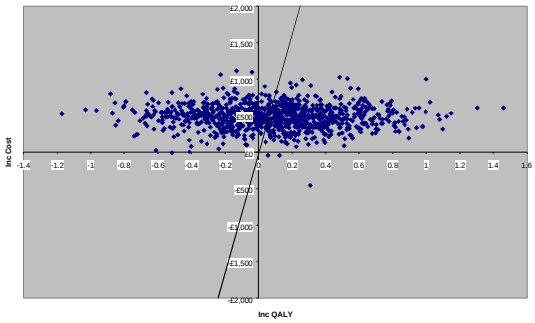

anticipated between the two drugs' response rates i.e. if T0 response rate were high, then we would expect to see a high response rate for T1. Secondly, positive correlations are anticipated between the two drugs' utility improvements. Finally, there are positive correlations between mean response rate and the mean duration of that response i.e. response rate for T0 is positively correlated with duration of response for T0, and the same for T1's response rate and with duration of response. For Monte-Carlo sampling, we convert this correlation structure to a variance-covariance matrix for the correlated normally distributed variables. We used R software, to randomly sample correlated values from a multi-variate normal distribution. We also implemented an extension of Cholesky decomposition in EXCEL Visual Basic to create a new EXCEL function =MultiVariateNormalInv (CHEBS website)[18].

Case study 3 uses an adaptation of the Case study 1 model incorporating Markov transitions to investigate the effects of including non-linear functions of the variables. The parameters for mean duration of response on T0 and T1 are replaced with 2 models of the history of response to each drug over several years based on Markov transition matrices with health states "responding", "not responding" and "died". The central estimate transition probabilities per period for both drugs are 0.6, 0.3 and 0.1 from responding to responding, not responding and died respectively. The transition probabilities for patients already in the not responding state are 0.0, 0.9, 0.1. Death is an absorbing state. Thus, mean duration of response is a function of multiples of Markov transition matrices. To investigate the effects of increasingly non-linear models, we have analysed time horizons of 3, 5, 10, 15 and 20 periods. To undertake Monte Carlo sampling, we have characterised the level of uncertainty in these transition rates assuming that they are each based on a small sample of just 10 transitions i.e. 6 responders, 3 non responders and 1 death. We have assumed statistical independence between the transition probabilities for responders and non-responders. We also assumed statistical independence between the transition probabilities for T1 and T0. We sampled from the Dirichlet distribution in R, and also extended the method of Briggs[19] to create a new EXCEL Visual Basic function =DirichletInv[18].

For each case study, we analysed EVPI via different methods for 19 individual parameters and for groups. The groups represent different types of proposed data collection exercises – a trial to obtain data on response rate parameters alone, a utility study to obtain data on mean utility gain of responders, and a long term follow-up observational study could obtain data on duration of response to both drugs.

RESULTS

Literature Review

The literature reveals several methods to calculate EVPI. Early healthcare[6,20] literature used stylised decision problems and simplifying assumptions, such as uncertainty in net benefit is normally distributed, and showed overall (but not partial) EVPI calculated analytically via standard 'unit normal loss integral' statistical tables[4]. (via standard 'unit normal loss integral' statistical tables[21]), but gave no analytic calculation method for partial EVPI. In 1998, Felli and Hazen[7] gave a fuller exposition of EVPI method, setting out some mathematics using expected value notation, with a general Monte Carlo simulation procedure for partial EVPI calculation. This "MC1" procedure appeared to suggest that only the parameters of interest (ξ_I) need to be sampled but, following discussions with the authors of this paper, this was recently corrected[22] (both ξ_I and ξ_I^C sampled), to show mathematical notation with nested expectations. Felli and Hazen also provided a "shortcut" simulation procedure, for use when all parameters are assumed probabilistically independent and the payoff function is multi-linear. In the late 1990s, a series of EVPI case studies in the UK used a different algorithm to calculate partial EVPI[8,9,10].

This used a one level simulation algorithm to analyse the "expected opportunity loss remaining" if perfect information were obtained on a subset of parameters. Some work has been undertaken on the number of simulations required, and Coyle uses quadrature (taking samples at particular percentiles of the distribution) rather than random Monte Carlo sampling to speed up the calculation of EVPI for a single parameter[12]. Recent papers[11,12,13,14], compare partial EVPI to alternative 'importance' measures for sensitivity analysis, and others[15] compare EVPI to 'payback' methods for prioritising research, concluding that partial EVPI is the most logical and coherent approach without discussing in detail the exact EVPI calculation methods required. A further strand of literature is concerned with the value of gathering clinical information on a specific patient or group[23,24]. Here, the information gathering is actually an intervention (e.g. a diagnostic test or a screening strategy that gathers information in order to improve decision making) and the "value of the information" is the net benefit of the intervention. This is slightly different from EVPI on uncertain parameters in a policy decision model, and the unwary reader can become confused. EVPI methods in risk analysis were recently reviewed[25]. In 1999, building upon previous work by Gould[26], Hilton[27], Howard[28] and Hammitt[29], a case study by Hammitt and Shlyakhter[30] set out similar mathematics to Felli and Hazen and examined the use of elicitation methods to quantify prior probability distributions if data were sparse.

Concepts

EVPI concerns decisions on policy options under uncertainty. Decision theory shows that the 'adoption decision' should be the policy with the greatest *expected* pay-off given current information[31]. In healthcare, we use monetary valuation of health (λ) to calculate a single payoff synthesising health and cost consequences e.g. expected net benefit $E(NB) = \lambda * E(QALYs) - E(Costs)$. EVPI is a Bayesian[32] approach that works by taking current knowledge (a prior probability distribution), adding in proposed information to be collected (data) and producing a posterior (synthesised probability distribution) based on all available information. In defining EVPI, 'information' refers to the additional data we might collect to further inform the adoption decision. 'Perfect' information means perfectly accurate knowledge, or absolute certainty, about the value of a parameter, and can be thought of as an infinite sample size, producing a posterior probability distribution that is a single point. Perfect information on all parameters implies no uncertainty about the optimal adoption decision. The 'value' of additional information is quantified as the additional pay-off if a revised adoption decision is made. If existing uncertainty is small, then additional information is unlikely to revise the decision, but when existing uncertainty is large, then acquiring the true value of a parameter could indicate an alternative policy with higher expected net benefit. EVPI calculation is concerned with the value of *proposed* data collection and we assess the 'expected' impact of different possible data from across the full range of current uncertainty. For some data the adoption decision would be revised, for other data we would stick with

our baseline adoption decision policy. By investigating the pay-offs associated with different possible collected data, and averaging these results, the *'expected'* value of perfect information is quantified. 'Partial EVPI' examines the value of individual (or subsets of) parameters. For example, we might collect efficacy data whilst other parameters, such as costs, remain uncertain. Perfect information on all parameters gives 'overall EVPI'. Calculations are often done per patient, and we multiply by the number of patients affected over the lifetime of the decision to quantify 'population EVPI'. The concept then is clear, we simulate obtaining additional information, consider whether this data would result in a revised adoption decision, and quantify the net benefit of our original baseline decision as compared with the extra net benefit provided by the revised decision given new data.

Algorithms and Mathematics

Whilst the literature suggests several approaches, the two level algorithm below is applicable to any decision model (Box 1). Note the EVPI calculation depends on λ, but does not need repeating for different thresholds. If mean cost and mean effectiveness are recorded separately for each strategy at the end of each inner loop (the end of step 5), then partial EVPI is quick to calculate for any λ. Overall EVPI is just partial EVPI for the whole parameter set, so the inner loop is redundant because there are no remaining uncertain parameters. Box 1 is a fuller description of Felli and Hazen's MC1 approach. The revised MC1 procedure[22] seems to suggest concurrently generating random samples of the parameters of interest (ξ_I) and the remaining uncertain parameters (ξ_I^C). In fact, the nested nature of the sampling is implicit in the mathematics of MC1 step 2. Our algorithm shows the nested loops transparently – first sample parameters of interest and only then sample the remaining uncertain parameters, *given* the sampled parameters of interest. The MC1 procedure assumes there is an algebraic expression for the expected net benefit of the revised adoption decision given new data (step 2i). For simple decision models, algebraic integration of net benefit functions can be tractable, but our inner loop provides a generalised method for any model. Finally, MC1 step 2ii suggests calculating the improvement (i.e. net benefit of the revised minus the baseline decision) within an inner loop, which is correct, but not necessary as mathematics will show.

Overall EVPI Mathematics

θ	be the parameters for the model, with defined prior probability distributions.
d	be the set of possible decisions or strategies.
$NB(d,q)$	be the function of net benefit for decision d, and parameters θ
$E[f(x)]$	denote the expectation of a function f(x)

Box 1: General 2 level Monte Carlo Algorithm for Calculation of Partial EVPI on a Parameter Subset of Interest

Preliminary Steps

0) Set up a decision model comparing different strategies and set up a decision rule e.g. Cost per QALY is < λ

1) Characterise uncertain parameters with probability distributions
e.g. normal(μ,σ^2), beta(a,b), gamma (a,b), triangular(a,b,c) … etc

2) Simulate (say 10,000) sample sets of uncertain parameter values (Monte Carlo).

3) Work out the baseline adoption decision given current information
i.e. the strategy giving (on average over the 10,000 simulations) the highest expected net benefit.

Partial EVPI for a parameter subset of interest

The algorithm has 2 nested loops

4) Simulate a perfect data collection exercise for your parameter subset of interest by:
sampling each parameter of interest once from its prior uncertain range **(outer level simulation)**

5) calculate the best strategy given this new knowledge on the parameter of interest by
- fixing the parameters of interest at their sampled values
- simulating the other remaining uncertain parameters (say 10,000 times) allowing them to vary according to their prior uncertainty **(inner level simulation)**
- calculating the mean net benefit of each strategy
- choosing the revised adoption decision to be the strategy which has the highest expected net benefit given the new data on the parameters of interest

6) Loop back to step 4 and repeat steps 4 and 5 (say 10,000 times) and then calculate the average net benefit of the revised adoption decisions given perfect information on parameters of interest

7) The EVPI for the parameter of interest =
average net benefit of revised adoption decisions given perfect information on parameters (6)
minus
average net benefit given current information i.e. of the baseline adoption decision (3)

Overall EVPI

The algorithm for overall EVPI requires only 1 loop

8) For each of the 10,000 sampled sets of parameters from step (3) in turn,
- work out the optimal strategy given that particular sampled sets of parameters,
- record the net benefit of the optimal strategy

9) With "perfect" information (i.e. no uncertainty in the values of each parameter) we would always choose the optimal strategy.
Overall EVPI =
average net benefit of optimal adoption decisions given perfect information on all parameters (8)
minus
average net benefit given current information i.e. of the baseline adoption decision (3)

First, consider the situation with only the current (prior) information. We evaluate each strategy in turn and choose the baseline adoption decision with the maximum expected net benefit.

Expected net benefit | current information = $\max_d \{E_q \, \mathrm{NB}(d,q)\}$ (1)

Second, consider obtaining perfect information on all model parameters. We would choose the strategy with known highest net benefit. At this stage we do not have the data, but we can calculate the *expectation* of the net benefits we will receive by integrating over the uncertain ranges of the parameters.

Expected net benefit | perfect information = $E_q\left\{\max_d \text{NB}(\text{d},q)\right\}$ (2)

The expected additional value of collecting perfect information on all uncertain model parameters is therefore given by equation (2) minus equation (1), as shown below.

Overall EVPI = $E_q\left\{\max_d \text{NB}(\text{d},q)\right\} - \max_d\left\{E_q \text{NB}(\text{d},q)\right\}$ (3)

Both (1) and (2) have one expectation and so (3) can be computed using concurrent 1 level simulation.

Partial EVPI Mathematics

θ_i be the parameters of interest for partial EVPI (i.e. those for which data collection is proposed).

θ_{-i} be the other remaining uncertain parameters i.e. those <u>not</u> of interest

If we obtain a particular value perfect information on θ_i , say θ_i^p, but have remaining uncertainty on θ_{-i} we should choose the strategy providing maximum expected net benefit i.e. $\max_d\left(E_{q_{-i}} \text{NB}(\text{d},q) \mid q_i^p\right)$

At this stage we do not have the perfect data on θ_i. So, we need to calculate the *expectation* of the net benefits we will receive by also integrating over the uncertain ranges of the θ_i parameters. The expected net benefit given the proposed collection of perfect information on the set of parameters θ_i is therefore:

Expected Net benefit, given perfect info only on θ_i = $E_{q_i}\left\{\max_d\left(E_{q_{-i}} \text{NB}(\text{d},q) \mid q_i\right)\right\}$ (4)

The conditioning on θ_i in the inner expectation is significant. In general, we expect that learning the true value of θ_i would also provide some information about θ_{-i}. Hence the correct distribution to use for the inner expectation is the conditional distribution that represents the remaining uncertainty in θ_{-i} after learning θ_i. The exception is when θ_i and θ_{-i} are independent, allowing the unconditional (marginal) distribution of θ_{-i} to be used in the inner expectation. Although such independence is often assumed in economic model parameters (as we do in Case Study 1), the assumption is rarely fully justified.

The additional value of collecting perfect information on a subset θ_i of uncertain model parameters is therefore given by equation (4) minus equation (1).

Partial EVPI obtaining data on θ_i = $E_{q_i}\left\{\max_d\left(E_{q_{-i}} \text{NB}(\text{d},q) \mid q_i\right)\right\} - \max_d\left\{E_q \text{NB}(\text{d},q)\right\}$ (5)

(5) clearly shows two expectations. The inner expectation evaluates the net benefit over the remaining uncertain parameters θ_{-i}. The outer evaluates the net benefit over the parameters of interest θ_i. This is the mathematical expression of the need for 2 level Monte Carlo simulation algorithm.

'Short-Cut' Algorithms

The literature revealed two 1 level algorithms. Felli and Hazen[7] give a 'shortcut procedure' (MC2). This performs one level simulation, allowing parameters of interest to vary, keeping remaining uncertain parameters constant at their prior means. **Box 2** describes our equivalent full transparent algorithm.

Box 2: One level Monte Carlo Algorithm for Calculation of Partial EVPI on a Parameter Subset of Interest

Preliminary Steps …. As in Box 1

One level Partial EVPI for a parameter subset of interest

The algorithm has 1 loop

4) Simulate a perfect data collection exercise for your parameter subset of interest by:
sampling each parameter of interest once from its prior uncertain range **(one level simulation)**

5) calculate the best strategy given this new knowledge on the parameter of interest by
- fixing the parameters of interest at their sampled values
- fixing the remaining uncertain parameters of interest at their prior mean value
- calculating the mean net benefit of each strategy given these parameter values
- choosing the revised adoption decision to be the strategy which has the highest net benefit given the new data on the parameters of interest

6) Loop back to step 4 and repeat steps 4 and 5 (say 10,000 times) and then calculate the average net benefit of the revised adoption decisions given perfect information on parameters of interest

7) The EVPI for the parameter of interest =
average net benefit of revised adoption decisions given perfect information on parameters (6)
minus
average net benefit given current information i.e. of the baseline adoption decision (3)

This algorithm does not apply in general (it ignores uncertainty in parameters not of interest). But, it can apply in certain situations. Felli and Hazen suggest "when all parameters are assumed probabilistically independent and the pay-off function is multi-linear i.e. linear in each individual parameter".

The mathematics shows the outer level expectation as per equation (5) but the inner expectation is replaced with net benefit calculated given the remaining uncertain parameters set at their prior mean.

1 level partial EVPI for θi = $$E_{q_i}\left\{\max_d \text{NB}(d, q_i \mid q_{-i} = \overline{q_{-i}})\right\} - \max_d \{E_q \text{NB}(d,q)\} \qquad (6)$$

Conditions for 'Short-cut' Accuracy

The 1-level approach is equivalent to the correct 2 level algorithm if (5) ≡ (6), i.e.

if $$E_{q_i}\left\{\max_d \left(E_{q_{-i}} \text{NB}(d,q) \mid q_i\right)\right\} \equiv E_{q_i}\left\{\max_d \text{NB}(d, q_i \mid q_{-i} = \overline{q_{-i}})\right\} \qquad (7)$$

This is true if the left hand side inner bracket (expectation of net benefit, integrating over θ_{-i}) is equal to the net benefit obtained when θ_{-i} are fixed at their prior means (i.e. $q_{-i} = \overline{q_{-i}}$) in the right hand side.

The expected value of a vector function Fn(θ_{-i}) is equal to Fn($\overline{q_{-i}}$) if the function is a simple linear function i.e. Fn(θ_{-i}) = $A_1 * \theta_{-i(1)} + A_2 * \theta_{-i(2)} + A_3 * \theta_{-i(3)} + \ldots +$ constant. This is a sufficient condition but not necessary. Fn(θ_{-i}) can equal Fn($\overline{q_{-i}}$) without this linearity condition holding (see Case Study 1).

Furthermore, even with linearity, we can only have equality in (7) if, for all possible values of θ_i, the conditional expectation of θ_{-i} given θ_i equals the unconditional expectation of $\overline{q_{-i}}$. This therefore imposes a further condition, which is satisfied if θ_i and θ_{-i} are independent, but not generally otherwise. Thus (7) is satisfied if (a) the net benefit functions are linear functions of the θ_{-i} for all of the decisions *d* and all of the possible values of the parameter set of interest θ_i, and (b) if θ_i and θ_{-i} are independent.

Alternative 1 level Algorithm

The other 1 level algorithm used to evaluate partial EVPI in some recent case studies involves considering the reduction in opportunity loss given new data. The overall opportunity loss inherent in a decision problem is given by the overall EVPI from equation (3). In order to calculate this for a subset of parameters, this alternative 1 level algorithm simulates by keeping the parameters of interest constant at their prior mean values $\overline{q_i}$, and allowing the remaining unknown parameters to vary according to prior uncertainty. The algorithm is almost equivalent to that in Box 2 but the parameters kept constant are reversed. The expected opportunity loss given $q_i = \overline{q_i}$ is calculated. Mathematically,

$$\text{New opportunity loss} = E_{q_{-i}}\left\{\max_d \text{NB}(\text{d}, q_{-i} \mid q_i = \overline{q_i})\right\} - \max_d \left\{E_{q_{-i}} \text{NB}(\text{d}, q_{-i} \mid q_i = \overline{q_i})\right\} \qquad (3b)$$

Thus the alternative 1 level EVPI i.e. reduction in opportunity loss given θ_i is given by (3) - (3b)

$$= E_q\left\{\max_d \text{NB}(\text{d}, q)\right\} - \max_d \left\{E_q \text{NB}(\text{d}, q)\right\} -$$

$$\left(E_{q_{-i}}\left\{\max_d \text{NB}(\text{d}, q_{-i} \mid q_i = \overline{q_i})\right\} - \max_d \left\{E_{q_{-i}} \text{NB}(\text{d}, q_{-i} \mid q_i = \overline{q_i})\right\}\right) \qquad (8)$$

This expression is actually a measure of how much remaining uncertainty there would be in the decision problem if we had perfect knowledge that the parameters of interest are at their prior mean values. It is not actually measuring partial EVPI but rather an approximation of residual EVPI.

Implementing the Two Level Expectation Calculation via Monte Carlo Sampling

Summation notation can better describe the Monte Carlo sample mean calculation.

$$\text{Overall EVPI} = \frac{1}{N}\sum_{n=1}^{N}\left[\max_{d=1 to D}\left(\text{NB}(\text{d}, q_n)\right)\right] - \max_{d=1 to D}\left[\frac{1}{L}\sum_{l=1}^{L}\left(\text{NB}(\text{d}, q_l)\right)\right] \qquad (3s)$$

$$\text{Partial EVPI} = \frac{1}{K}\sum_{k=1}^{K}\left(\max_{d=1 to D}\left(\frac{1}{J}\sum_{j=1}^{J}\left[NB\left(d, q_j^{-i} \mid q^i = q_k^i\right)\right]\right)\right) - \max_{d=1 to D}\left[\frac{1}{L}\sum_{l=1}^{L}\left(\text{NB}(\text{d}, q_l)\right)\right] \qquad (5s)$$

For partial EVPI, we denote by K, the number of different sampled values of parameters of interest θ_i, and J, the number of sets of values for the other parameters θ_{-i}. D is the number of decision policies.

L is the number of sampled sets of values used for all the parameters together when calculating the expected net benefit of the baseline adoption decision i.e. the second term. This can be done once for the

whole decision problem rather than each time when evaluating partial EVPI for different parameter subsets. In the first term of (5s), the inner expectation is estimated by taking a mean net benefit over J Monte Carlo samples of θ_{-i} given the particular k^{th} sample of θ_i. The outer expectation is estimated by taking a mean over K Monte Carlo samples of the parameters of interest θ_i. An equal number for inner and outer sampling is not necessary and may not be efficient. In considering the number of samples required, we need to consider the precision of the resulting estimate in EVPI. The larger the number of Monte Carlo samples, the smaller the likely error (variance) in the Monte Carlo estimate of an expectation. This is true of all 3 summations in (5s). The accuracy with which they estimate the corresponding expectations in equation (5) will increase as we increase J, K and L.

An important issue of bias in these estimates emerges from examining the mathematics. Monte Carlo estimates of maximums of expectations are biased upwards i.e. tend to be overestimated if small samples are used. If X and Y are random quantities, then in general max(E[X],E[Y]) is upward biased if E[X], and E[Y] are estimated with error. The argument is : (a) in the inner loop of equation 5 we need to evaluate max(E[NB1], E[NB2]), (b) if we use simulation to do this there is simulation error in E[NB1] and E[NB2], (c) our estimate of max(E[NB1], E[NB2]), is increasingly biased as Var(E[NB1]) and Var(E[NB2]) increase. Thus, in the first term of (5s), the smaller the number of samples J, the larger the bias, and hence the greater tendency to overestimate the maximum of the expected net benefits. Taking maximums of Monte Carlo expectations occurs in both terms of the computation of partial EVPI. Both terms will be subject to upward bias. Whether the overall bias in equation (5s) is upwards or downwards will depend on the net benefit functions, the characterised uncertainty and the relative sizes of J and L. Increasing the sample size J reduces that bias of the first term. Increasing the sample size L reduces that bias of the second term. The size K of the outer sample in the 2-level calculation is less crucial because it does not induce bias. For overall EVPI, only the second term of the computation is biased. The first term in (3s) is unbiased. Hence, the Monte Carlo estimate of overall EVPI with small samples is biased downwards.

Case Study Results

Table 1 shows partial EVPI results for each parameter subset for Case study 1. Using 1,000 simulations the overall EVPI per patient is £1,352. The population overall EVPI (maximum benefit to society) is £1,352,000, usually interpreted as meaning that a study costing more could not possibly be efficient. The 2 level and 1 level results are very similar. When expressed as a percentage of the overall EVPI, the largest absolute difference between 2 level and 1 level results is 3%. The 2 level algorithm results are slightly higher than the 1 level results in every case. This reflects the mathematical results shown earlier. Firstly, the 1 level and 2 level EVPI calculations are mathematically equivalent if the cost-

effectiveness model has net benefit functions that are sum-products of statistically independent parameters. And secondly, the 2 level estimates in this case are upwardly biased due to the maximisation of Monte-Carlo estimate in the inner loop. Six of the 19 individual parameters are important i.e. response rates, utility parameters and durations of response. The remaining thirteen (not shown) had zero partial EVPI. Note that the partial EVPI for groups of parameters is always lower than the sum of the EVPIs of the individual parameters (e.g. utility parameters combined = 57%, compared with individual utility parameters = 46%+24% = 70%). We also compared the one level opportunity loss reduction method (Equation 8). The results were substantially lower for every individual parameter and group, sometimes less than half that of the 2 level algorithm (e.g. utility parameters combined estimate = 27%), and confirming that this incorrect method could substantially under-estimate partial EVPI.

Table 1: Case Study Results Comparing the Algorithms

		ß---------- Partial EVPI (Indexed - Overall EVPI = 100) ----------à										Overall EVPI
Parameters		% Resp T0	Utility Change if Respond T0	Duration of Response T0	% Resp T1	Utility Change if Respond T1	Duration of Response T1	Trial	Utility Only	Trial + Utility	Durations	All
Parameter No.		5	6	7	14	15	16	5,14	6,15	5,6,14,15	7,16	
Case Study 1	**Independent Linear Cross-product model**											
	2 level	17	46	18	16	24	59	27	57	68	66	£ 1,352
	1 level	14	45	17	14	23	57	24	56	67	65	£ 1,352
Opportunity Loss	1 Level	10	24	10	12	16	32	6	27	35	33	£ 1,352
Case Study 2	**Correlated Parameters Linear Cross-product model**											
(a) 0.1	2 level	14	46	21	14	19	56	26	58	60	65	£ 1,306
(b) 0.2	2 level	15	47	28	23	21	55	41	59	74	72	£ 1,295
(c) 0.3	2 level	24	48	28	27	22	70	47	68	81	76	£ 1,275
(d) 0.6	2 level	26	50	46	44	23	85	55	71	91	93	£ 1,255
(a) 0.1	1 level	12	47	16	12	20	56	20	52	62	62	£ 1,306
(b) 0.2	1 level	12	47	16	12	20	56	19	50	58	62	£ 1,295
(c) 0.3	1 level	12	47	16	12	20	57	17	47	55	64	£ 1,275
(d) 0.6	1 level	12	48	16	12	21	58	19	53	62	64	£ 1,255
Case Study 3	**Non Linear Markov Models**											
(a) N=3	2 level	-	-	-	-	-	-	30	51	79	68	£ 903
(b) N=5	2 level	-	-	-	-	-	-	21	69	84	76	£ 1,154
(c) N=10	2 level	-	-	-	-	-	-	20	69	58	59	£ 1,616
(d) N=15	2 level	-	-	-	-	-	-	14	44	46	82	£ 1,898
(e) N=20	2 level	-	-	-	-	-	-	11	65	55	82	£ 2,119
(a) N=3	1 level	-	-	-	-	-	-	19	64	69	54	£ 903
(b) N=5	1 level	-	-	-	-	-	-	5	47	52	67	£ 1,154
(c) N=10	1 level	-	-	-	-	-	-	-13	25	28	75	£ 1,616
(d) N=15	1 level	-	-	-	-	-	-	-20	15	17	80	£ 1,898
(e) N=20	1 level	-	-	-	-	-	-	-26	6	9	81	£ 2,119

The case study 2 results show that if correlations are present between the parameters, then the 1 level EVPI results can substantially underestimate the true EVPI. The 1 level and 2 level EVPI estimates are broadly the same when small correlations are introduced between the important parameters. For example, with correlations of 0.1, the 2 level result for the utility parameters combined is 58%, 6 points higher than the 1 level estimate. However, if larger correlations exist, then the 1 level EVPI 'short-cut'

results can be very wrong. With correlations of 0.6, the 2 level result for the utility parameters combined is 71%, now 18 percentage points higher than the 1 level estimate. The maximum disparity is seen for the response rate parameters combined (36 points higher). As correlation is increased, Figure 2 shows that the average disparity over the parameters examined increases substantially.

Figure 2: Impact of Increasing Correlation on Inaccuracy of 1 level method to calculate partial EVPI

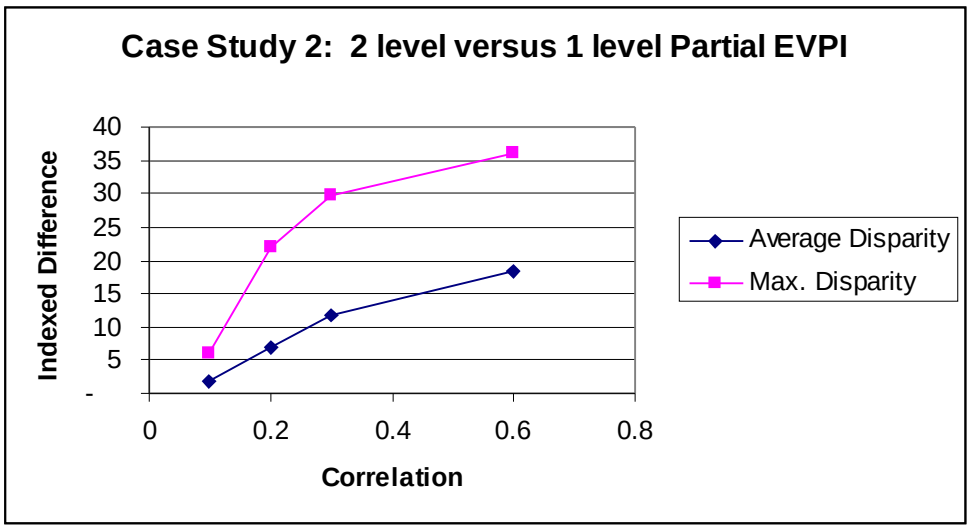

For Case study 2, the 1 level EVPI results are the same no matter what level of correlation is involved. This is because the 1 level algorithm sets the remaining parameters (those not of interest) at their prior mean values no matter what values are sampled for the parameters of interest. The 2 level algorithm correctly accounts for correlation, by sampling the remaining parameters from their *conditional* probability distributions within the inner loop i.e. *given* the values for the parameters of interest sampled in the outer loop. Note that these results further demonstrate that having linear or sum-product net benefit functions is not a sufficient condition for the 1 level EVPI estimates to be accurate. The second mathematical condition i.e. that parameters are statistically independent is just as important as the first.

The Case study 3 results show the effects of increasingly non-linear Markov models (Table 1). The 1 level estimate is substantially lower than the 2 level for the trial and utility parameters. Indeed, it actually negative for the trial parameters for the 3 most non-linear case studies. This is because the net benefit function is so non-linear that the first term in the 1 level EVPI equation $E_{q_i}\left\{\max_d \text{NB}(\text{d},q_i \mid q_{-i}=\overline{q_{-i}})\right\}$ is actually lower than the second term $\max_d\{E_q \text{NB}(\text{d},q)\}$. Thus, when we set the parameters we are not interested in (θ_{-i}) to their prior means in term 1, the net benefits obtained are lower than in term 2 when we allow all parameters to vary. Figure 3 shows that more non-linear net-benefit functions often produce less accurate 1 level EVPI estimates. However, this is not always the case. Partial EVPIs for the Markov transition probabilities for duration of disease show a high degree of alignment between the 1 level and 2 level methods. It is very important to note that even quite high adjusted R^2 (e.g. 0.973 in 3a as compared to 0.829 in 3e) does not imply that 1 level and 2 level

estimates will be the equal or even of the same order of magnitude. For example in 3a, the 2 level EVPI for trial parameters is 30 compared with a 1 level result of 19. The 2 level EVPI algorithm is necessary, even in non-linear Markov models which might be very well approximated by linear regression.

Figure 3: Impact of Increasing Non-Linearity on Inaccuracy of 1 level method to calculate partial EVPI

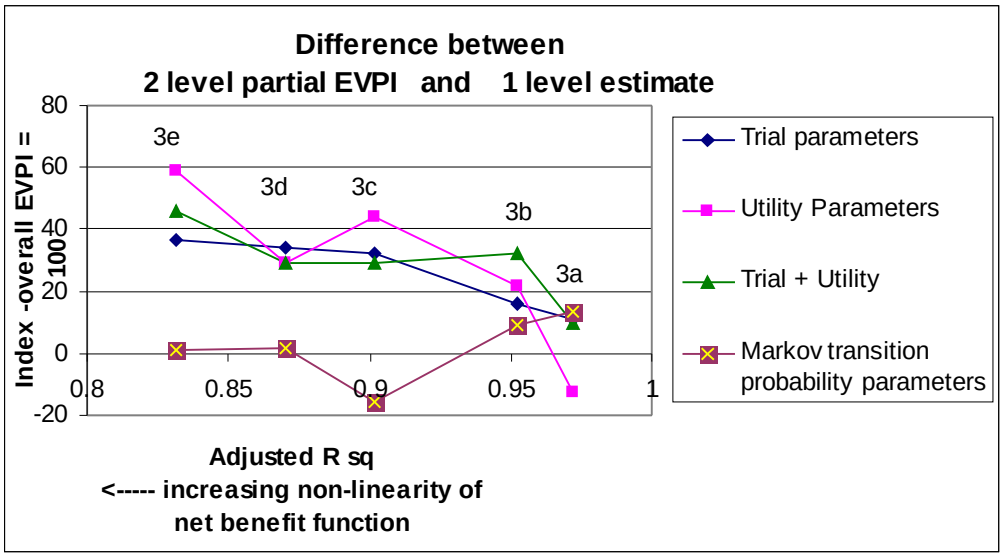

Numbers of Samples Required

The number of Monte-Carlo samples required for accurate unbiased estimates of partial EVPI has been examined in case studies 1 and 3. Case study 1, (satisfying the multi-linear and independence criteria for equivalence) showed 2 level results with 1000 inner samples that were always slightly higher than the 1 level due to the upward bias. We examined the response parameters combined (5,14) with larger samples. Figure 4 shows how the 2 level result changes depending on the number of inner level samples.

Figure 4: Number of Monte Carlo Iterations and the Effect on Accuracy of EVPI estimate for Response Parameters 5,14. in an Independent Linear Cross-product model

J (Inner Level) \ K (Outer level)	10	100	500	750	1000
10	44	40	36	37	36
100	24	16	14	15	15
500	33	27	24	24	24
750	29	26	23	24	24
1000	31	30	26	27	27
2000	32	30	27	27	27
5000	30	27	24	24	24
10000	30	27	24	25	25

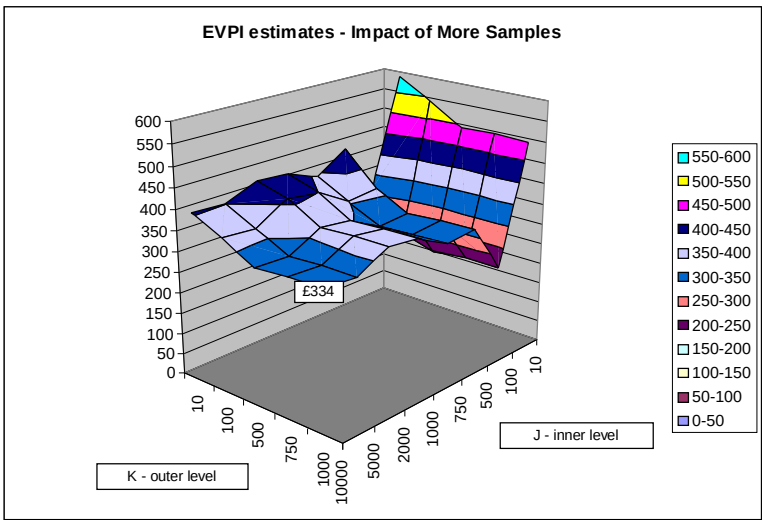

The correct indexed result is 25. With very small numbers of inner samples the 2 level results can be wrong by an order of magnitude (e.g. 10 inner samples and 1000 outer samples gives an indexed result

of 36. With 1,000 inner level samples the indexed result is 27 (£359) and with 10,000 inner samples the 2 level result is 25 (£334). Once over 500 outer level samples (K), the partial EVPI estimate converges to within 1 percentage point. This contrasts with the inner level (J), where the partial EVPI shows a 4 point difference between 750 and 1000 samples, and where it required samples of between 5,000 and 10,000 for the partial EVPI estimates to converge to within 1 percentage point. The number of samples needed for convergence is not symmetrical for J and K, indeed the results suggest that fewer samples on the outer level and larger numbers of samples on the inner level could be the most efficient approach to gaining accurate EVPI estimates. However, case study 1 does suggest that the order of magnitude of the partial EVPI estimates might be stable over 100 outer and 500 inner samples.

The stability of partial EVPI estimates using 100 outer and 500 inner samples has been tested on four different parameter groups using case study 3. Each of the five case study 3 models was re-run 3 times.

Table 2: Stability of partial EVPI estimates using 100 outer and 500 inner samples

	Trial	Utility Only	Trial + Utility	Durations
	5,14	6,15	5,6,14,15	7,16
	ß----------Partial EVPI		(Indexed - Overall EVPI = 100) ----------à	
1st Run (a)	33	73	89	50
(b)	26	60	74	83
(c)	18	52	58	54
(d)	14	70	55	81
(e)	13	49	53	89
2nd Run (a)	34	81	89	65
(b)	22	56	76	85
(c)	19	52	64	67
(d)	19	63	40	90
(e)	15	54	54	80
3rd Run (a)	26	56	87	47
(b)	22	66	68	81
(c)	20	54	54	76
(d)	16	49	60	78
(e)	13	44	48	70
Average (Standard deviation) of 3 estimates				
Average (a)	31	70	88	54
(b)	23	61	72	83
(c)	19	53	59	65
(d)	17	60	52	83
(e)	14	49	52	80
Std Dev (a)	4	13	1	10
(b)	2	5	4	2
(c)	1	1	5	11
(d)	2	11	10	7
(e)	2	5	3	9

The results do suggest some level of stability. For example, the partial EVPI for trial parameters in case study 3a shows 3 separate results of 33, 34 and 26. However, for other parameters the results do exhibit more noise, for example, the EVPI estimates for utility parameters in case study 3a are 73, 81, 56. We compared the standard deviation in the 3 estimates against the adjusted R^2 to see whether the level of non-linearity in the model has an impact on the "noise". The results showed that the noise is independent of the level of non-linearity in the net benefit functions. However Figure 5 does show that the higher the partial EVPI, the greater the level of noise that might be expected using the 100 outer and 500 inner samples.

Figure 5: Stability of partial EVPI estimates using 100 outer and 500 inner samples

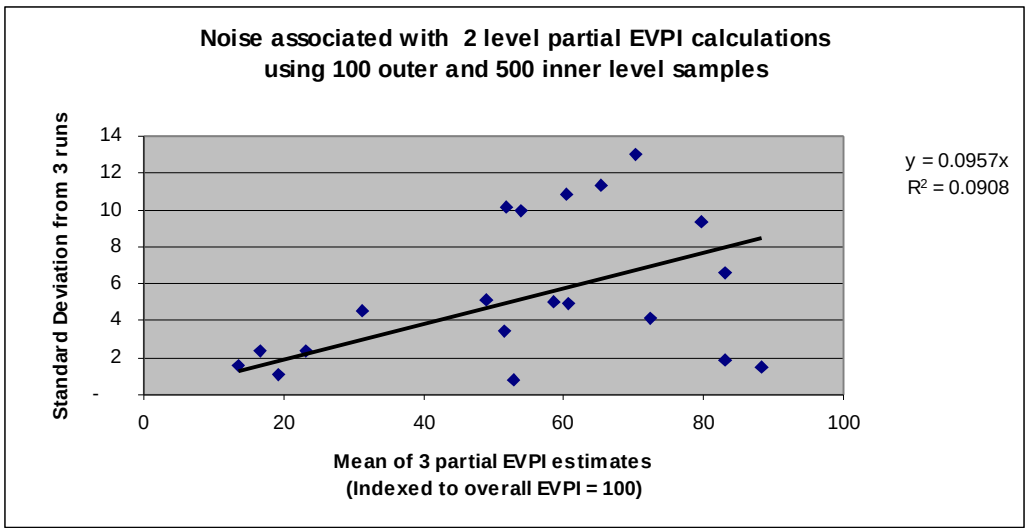

DISCUSSION

This paper mathematically describes the correct way to calculate partial EVPI, with the evaluation of two expectations, an outer expectation over the parameter set of interest and an inner expectation over the remaining parameters. A 'short-cut' 1 level algorithm is equivalent to the 2 level algorithm and can be recommended for use in models, which are simple enough for certain conditions to hold. If net benefits are non-linear functions of parameters or where model parameters are correlated, the 1 level algorithm can be substantially inaccurate. The scale of inaccuracy often increases with non-linearity and correlation, but can sometimes be unpredictable in scale and direction. The formerly used alternative 1 level algorithm does not measure partial EVPI (but rather residual EVPI) and so is not recommended. The 2 level Monte Carlo algorithm is recommended for general use.

The number of inner and outer level simulations required depends upon the number of parameters, their importance to the decision, and the model's net benefit functions. Our empirical approach, in a series of alternative models, suggests they do not need to be equal. 500 inner loops for each of the 100 outer loop iterations (i.e. 50,000 iterations in total) proved capable of estimating the order of magnitude of partial EVPI reasonably well. For very accurate calculation or in computationally intensive models, one might use an adaptive process to test for convergence in the partial EVPI results, within a pre-defined threshold.

Our mathematics has also shown the existence of an over-estimating bias in evaluating maximum expected net benefit across strategies using small numbers of Monte Carlo samples. This can result in over or under-estimating the partial EVPI depending on the number of iterations used to evaluate the first and second terms. Previous authors have investigated mathematical description of Monte Carlo bias outside the EVPI context[33]. Further mathematical definition of the relationship between iterations and accuracy in partial EVPI, will be of benefit[34].

The 1 level 'short-cut' algorithm could be useful to screen for parameters, which do not require further analysis. If parameters do not affect the decision, then their partial EVPI will be very close to zero using both the 2 level and the 1 level algorithm. Thus, the 1 level algorithm might be used with a relatively small number of iterations (e.g. 100) to screen for groups of parameters in very large models. The use of such an approach in non-linear or correlated models is a trade-off.

There is a relationship between the common representations of uncertainty, the cost-effectiveness plane and the cost-effectiveness acceptability curve[35], and the overall EVPI results. The cost-effectiveness

plane shows the relative importance of uncertainty in costs and effectiveness. Partial EVPI shows the breakdown by parameter, so decision makers see clearly the source and scale of uncertainty.

Since first presenting our mathematics and algorithm [5,36] a small number of case studies have been developed. For the UK National Institute for Clinical Excellence and NCCHTA, Claxton et al. present six such case studies[37]. In Canada, Coyle at al. have used a similar approach for the treatment of severe sepsis[38]. Authors continue to recommend EVPI [14,25]. Development of the approach to calculate expected value of sample information (EVSI) is also ongoing. The 2 level algorithm requires adaptation[39,40], undertaking Bayesian updating to move from a prior distribution for the uncertain parameter, via the simulated data collection, to a pre-posterior distribution. For some probability density functions this can be analytic, but others require numerical approaches or the use of WinBUGS simulation. Recent case studies include analysis of pharmaco-genetic tests in rheumatoid arthritis[41] and a mathematical description is now available for a number of scenarios[42].

Methodological research would be useful for population EVPI, which demands estimated patient numbers involved in the policy decision. Incidence and prevalence are important, as are the likely lifetime of the technology and potential changes in competitor strategies. There are arguments over the validity of analysing phased adoption of the intervention over time explicitly versus full adoption implied by the decision rule. Timing of data collection is important too. Some parameters may be collectable quickly (e.g. utility for particular health states), others take longer (e.g. long term side-effects), and still others may be inherently unknowable (e.g. the efficacy of an influenza vaccine prior to the arrival of next years strain of influenza). There are trade-offs on the investment in different data collection, quicker and cheaper versus fuller reduction in uncertainty.

EVPI is important, both in decision-making, and in planning and prioritising future data collection. Policy makers assessing interventions are keen to understand the level of uncertainty, and many guidelines recommend probabilistic sensitivity analysis[20]. This paper seeks to encourage analysts to extend the approach to calculation of partial EVPI. The theory and algorithms required are now in place. The case study models have shown the feasibility and performance of the method, indicating the numbers of samples needed for stable results. Wider application will bring greater understanding of decision uncertainty and research priority.

86 **REFERENCES**

[1] National Institute for Clinical Excellence. Guidance for manufacturers and sponsors: Technology Appraisals Process Series No. 5. 2001. London, National Institute for Clinical Excellence.

[2] Manning WG, Fryback DG, Weinstein MC. Reflecting uncertainty in cost-effectiveness analysis. In Gold MR, Siegel JE, Russell LB, Weinstein MC, eds. Cost-Effectiveness in Health and Medicine. New York: Oxford university Press, 1996. p.247-75.

[3] Coyle ref

[4] Felli C,.Hazen GB. Sensitivity analysis and the expected value of perfect information. Medical Decision Making 1998;18:95-109.

[5]. Brennan A, Chilcott J, Kharroubi SA, O'Hagan A. A two level Monte Carlo approach to calculating expected value of perfect information:- Resolution of the uncertainty in methods. Centre for Health Economics and Bayesian Statistics, Discussion Paper, 2002 – Poster at SMDM

[6] Claxton K,.Posnett J. An economic approach to clinical trial design and research priority-setting. Health Economics 1996;5:513-24.
[7] Felli C,.Hazen GB. Sensitivity analysis and the expected value of perfect information. Medical Decision Making 1998;18:95-109.
[8] Fenwick E, Claxton K, Sculpher M, Briggs A. Improving the efficiency and relevance of health technology assessment: The role of iterative decision analytic modelling. CHE Discussion Paper 179, Centre for Health Economics 2000.
[9] Claxton K, Neuman PJ, Araki SS, Weinstein M. The value of information: an application to a policy model of Alzheimer's disease. International Journal of Health Technology Assessment 2001;17:38-55.
[10] Payne N, Chilcott JB, McGoogan E. Liquid-based cytology in cervical screening. Health Technology Assessment 2000.
[11] Felli JC, Hazen GB. A Bayesian approach to sensitivity analysis. Health Economics 1999;8:263-68.
[12] Coyle D, Grunfeld E, Wells G. The assessment of the economic return from controlled clinical trials. A framework applied to clinical trials of colorectal cancer follow-up. Eur J Health Econom 2003;4:6-11.
[13] Karnon J, Brennan A, Chilcott J. Commentary on Coyle et al. The Assessment of the economic return from controlled trials. Eur J Health Econom 2003;4:239-40.
[14] Meltzer D. Addressing uncertainty in medical cost-effectiveness analysis—implications of expected utility maximization for methods to perform sensitivity analysis and the use of cost-effectiveness analysis to set priorities for medical research. J Health Econ. 2001;20:109-129.
[15] Chilcott J, Brennan A, Booth A, Karnon J, Tappenden P. The role of modelling in prioritising and planning clinical trials. Health Technol Assess 2003;7(23).
[16] Briggs AH, Gray A. Handling uncertainty when performing economic evaluation of healthcare interventions. Health Technology Assessment 1999;3.
[17] Van Hout BA, Al MJ, Gordon GS, Rutten FF . Costs, effects and C/E-ratios alongside a clinical trial. Health Economics 1994;3:309-19.
[18] http://www.shef.ac.uk/chebs/
[19] Briggs AH, Ades AE, Price MJ. Probabilistic Sensitivity Analysis for Decision Trees with Multiple Branches: Use of the Dirichlet Distribution in a Bayesian Framework, Medical Decision Making 2003; 23:4:341-350
[20] Thompson KM, Evans JS. The value of improved national exposure information for perchloroethylene (Perc): a case study for dry cleaners, Risk Analysis 1997;17:253-71.
[21] Lapin, L. L. Quantitative methods for business decisions. 1994. Orlando, Florida, Dryden Press.
[22] Felli JC, Hazen GB. Correction: Sensitivity Analysis and the Expected Value of Perfect Information, Medical decision Making, 2003;23:97
[23] Phelps CE, Mushlin AI. Focusing technology-assessment using medical decision-theory. Med Decis Making. 1988;8:279-289.
[24] Bartell SM, Ponce RA, Takaro TK, Zerbe RO, Omenn GS, Faustman EM. Risk estimation and value-of-information analysis for three proposed genetic screening programs for chronic beryllium disease prevention. Risk Anal. 2000;20:87-99.
[25] Yokota F, Thompson K. Value of Information Literature Analysis: A Review of Applications in Health Risk Management, Medical Decision Making, 2004;24:287-98.
[26] Gould J. Risk,stochastic preference and the value of information, Journal of Economic Theory, 1974;8: 64–84.
[27] Hilton R. Determinants of information value, Management Science, 1981;27: 57–64.
[28] Howard RA. Decision Analysis: Practice and Promise, Management Science, 1988;34:679–695.
[29] Hammitt JK. Can More Information Increase Uncertainty? Chance, 1995;8:15–17.
[30] Hammitt JK, Shlyakhter AI. The Expected Value of Information and the Probability of Surprise, Risk Analysis, 1999;19:135-152.

[31] Raiffa H. Decision analysis: Introductory lectures on choice under uncertainty. 1968. Addison-Wesley, Reading, Mass.
[32] Spiegelhalter DJ, Myles JP, Jones DR, Abrams KR. Bayesian Methods in Health Technology Assessment: A Review. *Health Technology Assessment* 2000;**4**.
[33] Linville CD, Hobbs BF, Venkatesh BN. Estimation of error and bias in Bayesian Monte Carlo decision analysis using the bootstrap, Risk Analysis, 2001;21:63-74.
[34] Oakley J. Value of information for complex cost-effectiveness models. Research Report No. 533/02 Department of Probability and Statistics, University of Sheffield, 2002.
[35] Van Hout BA, Al MJ, Gordon GS, Rutten FF . Costs, effects and C/E-ratios alongside a clinical trial. Health Economics 1994;3:309-19.

[36] Brennan A, Chilcott J, Kharroubi SA, O'Hagan A. A two level Monte Carlo approach to calculating expected value of perfect information:- Resolution of the uncertainty in methods. Centre for Health Economics and Bayesian Statistics, Discussion Paper, 2002
[37] Claxton K, Eggington S, Ginnelly L, Griffin S, McCabe C, et al. A pilot study of value of information analysis to support research recommendations for the National Institute for Clinical Excellence, Society for Medical Decision Making Annual Conference, (abstract), 2004.
[38] Coyle, D., McIntyre, L., Martin, C., and Herbert, P. Economic analysis of the use of drotrecogin alfa (activated) in the treatment of severe sepsis in Canada. Poster #9 presented at the 24th Annual Meeting of the Society for Medical Decision Making, October 20th, 2002, Washington. 2002.
[39] Brennan, A. B., Chilcott, J. B., Kharroubi, S., and O'Hagan, A. A two level monte carlo approach to calculation expected value of sample information: how to value a research design. Presented at the 24th Annual Meeting of the Society for Medical Decision Making, October 23rd, 2002, Washington. 2002.
[40] Web address for my EVSI paper
[41] Brennan, A. B., Bansback, N., Martha, K., Peacock, M., and Huttner, K. Methods to analyse the economic benefits of a pharmacogenetic test to predict response to biological therapy in rheumatoid arthritis and to prioritise further research. Presented at the 24th Annual Meeting of the Society for Medical Decision Making, October 22nd, 2002, Washington. 2002.
[42] Ades AE, Lu G, Claxton K**.** Expected value of sample information calculations in medical decision modelling. Medical Decision Making. 2004 Mar-Apr;24(2):207-27.