Calculating Partial Expected Value Of Perfect Information Via Monte-Carlo Sampling Algorithms.

Alan Brennan, MSc[(a)]

Samer Kharroubi, PhD[(b)]

Anthony O'Hagan, PhD[(c)]

Jim Chilcott, MSc[(a)]

(a) Health Economics and Decision Science, School of Health and Related Research, The University of Sheffield, Regent Court, Sheffield S1 4DA, England.

(b) Department of Mathematics, University of York, Heslington, York YO10 5DD, England.

(c) Department of Probability and Statistics, The University of Sheffield, Hounsfield Road, Sheffield S3 7RH, England.

Reprint requests to:

Alan Brennan, MSc

Director of Health Economics and Decision Science,

School of Health and Related Research,

The University of Sheffield,

Regent Court,

Sheffield S1 4DA,

England.

a.brennan@sheffield.ac.uk

ABSTRACT

Partial EVPI calculations can quantify the value of learning about particular subsets of uncertain parameters in decision models. Published case studies have used different computational approaches. This paper examines the computation of partial EVPI estimates via Monte-Carlo sampling algorithms. Our mathematical definition shows two nested expectations, which must be evaluated separately because of the need to compute a maximum between them. A generalised Monte-Carlo sampling algorithm uses nested simulation with an outer loop to sample parameters of interest and, conditional upon these, an inner loop to sample remaining uncertain parameters. Alternative computation methods and 'shortcut' algorithms are discussed and mathematical conditions for their use are considered. Maxima of Monte-Carlo estimates of expectations are biased upwards, and we demonstrate that using small samples results in biased EVPI estimates. Three case studies illustrate (i) the bias due to maximisation, and also the inaccuracy of shortcut algorithms (ii) when correlated variables are present and (iii) when there is non-linearity in net-benefit functions. If relatively small correlation or non-linearity is present, then the 'shortcut' algorithm can be substantially inaccurate. Empirical investigation of the numbers of Monte-Carlo samples suggest that fewer samples on the outer level and more on the inner level could be efficient and that relatively small numbers of samples can sometimes be used. Several remaining areas for methodological development are set out. Wider application of partial EVPI is recommended both for greater understanding of decision uncertainty and for analysing research priorities.

Acknowledgements:

The authors are members of ***CHEBS***: *The Centre for Bayesian Statistics in Health Economics,* University of Sheffield. Thanks also go to Karl Claxton and Tony Ades who were involved in our CHEBS "focus fortnight" event, to Gordon Hazen, Doug Coyle, Myriam Hunink and others for feedback on the poster at SMDM and to the UK National Coordinating Centre for Health Technology Assessment which originally commissioned two of the authors to review the role of modelling methods in the prioritisation of clinical trials (Grant: 96/50/02). Thanks also to Simon Eggington of ScHARR who programmed the new CHEBS add-in functions for sampling from the Dirichlet and Multi-Variate Normal distributions in EXCEL. Finally, many thanks to both of the referees whose substantial feedback through two revisions has helped to improve the scope, rigour and style of this publication.

INTRODUCTION

Quantifying expected value of perfect information (EVPI) is important for developers and users of decision models. Many guidelines for cost-effectiveness analysis now recommend probabilistic sensitivity analysis (PSA)[1,2] and EVPI is seen as a natural and coherent methodological extension[3,4]. Partial EVPI calculations are used to quantify uncertainty, identify key uncertain parameters, and inform the planning and prioritising of future research[5]. Many recent papers recommend partial EVPI, for sensitivity analysis rather than alternative 'importance' measures[6,7,8,9], or for valuing research studies in preference to 'payback' methods, but do not discuss computation methods in any detail. Some of the few published EVPI case studies have used slightly different computational approaches[10] and many analysts, who confidently undertake PSA to calculate cost-effectiveness acceptability curves, still do not use EVPI.

The concepts of EVPI are concerned with policy decisions under uncertainty. A decision maker's 'adoption decision' should be that policy which has the greatest *expected* pay-off given current information[11]. In healthcare, we use monetary valuation of health (λ) to calculate a single expected payoff e.g. expected net benefit $E(NB) = \lambda * E(QALYs) - E(Costs)$. Expected value of information (EVI) is a Bayesian[12] 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. The *value* of the additional information is the difference between the expected payoff that would be achieved under posterior knowledge and the expected payoff under current (prior) knowledge. 'Perfect' information means perfectly accurate knowledge i.e. absolute certainty about the values of parameters, and can be conceptualised as obtaining an infinite sample size, producing a posterior probability distribution that is a single point, or alternatively, as 'clairvoyance' – suddenly learning the true values of the parameters. For some values of the parameters the adoption decision would be revised, for others we would stick with our baseline adoption decision policy. By

investigating the pay-offs associated with different possible parameter values, and averaging these results, the '*expected'* value of perfect information is quantified. Obtaining perfect information on *all* the uncertain parameters gives 'overall EVPI', whereas 'Partial EVPI' is the expected value of learning the true value(s) of an individual or subset of parameters. Calculations are often done per patient, and then multiplied by the number of patients affected over the lifetime of the decision to quantify 'population EVPI'.

Reviews show that several methods have been used to compute EVPI. The earliest healthcare literature[13] used simple decision problems and simplifying assumptions, such as normally distributed net benefit, to calculate overall EVPI analytically via standard 'unit normal loss integral' statistical tables[14], but gave no analytic calculation method for partial EVPI. In 1998 and 2003[15], Felli and Hazen gave a fuller exposition of EVPI method, with a suggested general Monte-Carlo random sampling procedure for partial EVPI calculation and a 'shortcut' simulation procedure for use in certain defined circumstances. We review these procedures in detail in the next section. In the late 1990s, some UK case studies employed different algorithms to attempt to compute partial EVPI[16,17,18], but these algorithms actually computed "expected opportunity loss remaining" given perfect information on a subset of parameters, which is not the same as partial EVPI and can give substantially different results[10,19]. In 2002, a UK event helped to produce work resulting in a series of papers providing guidance on EVI method[10,19,20]. UK case studies since that time have used the two level Monte-Carlo sampling approach we examine in detail here[21,22]. Coyle at al. have used a similar approach[23], though sometimes using quadrature (taking samples at particular percentiles of the distribution) rather than random Monte-Carlo sampling to speed up the calculation of partial EVPI for a single parameter. Development of the approach to calculate expected value of sample information (EVSI) is also ongoing[20,24,25,26].

The EVPI literature is not confined to health economic policy analysis. A separate literature examines information gathering as the actual intervention e.g. a diagnostic or screening test that gathers

105 information to inform decisions on individual patients[27,28]. Risk analysis is the other most common
106 application area. Readers with a wider interest are directed to a recent review of risk analysis
107 applications[29], which showed, for example, Hammitt and Shlyakhter[30] building on previous authors'
108 work,[31,32,33,34] setting out similar mathematics to Felli and Hazen, and using elicitation techniques to
109 specify prior probability distributions when data are sparse.
110
111 The objective of this paper is to examine the computation of partial EVPI estimates via Monte-Carlo
112 sampling algorithms. In the next section, we define partial EVPI mathematically using expected value
113 notation. We then present a generally applicable nested 2 level Monte-Carlo sampling algorithm
114 followed by some variants which are valuable in certain circumstances. The impact of sampling error on
115 these estimates is covered including a bias caused by maximisation within nested loops. We lay out the
116 mathematical conditions when a 'short-cut' 1 level algorithm may be used. Three case studies are
117 presented to illustrate (i) the bias due to maximisation, (ii) the accuracy or otherwise of the shortcut
118 algorithm when correlated variables are present and (iii) the impact of increasingly non-linear net-
119 benefit functions. Finally, we present some empirical investigations of the required numbers of Monte-
120 Carlo samples and the implications for accuracy of estimates when relatively small numbers of samples
121 are used. We conclude with the implications of our work and some final remarks concerning
122 implementation.
123
124 **MATHEMATICAL FORMULATION**
125
126 **Overall EVPI**
127 We begin with some notation. Let,
128 θ be the vector of parameters in the model with joint probability distribution $p(\theta)$.
129 *d* denote an option out of the set of possible decisions; typically, d is the decision to adopt
130 or reimburse one treatment in preference to the others.

NB(d,θ) be the net benefit function for decision d for parameters values θ.

Overall EVPI is the value of finding out the true value of the currently uncertain θ. If we are not able to learn the value of θ, and must instead make a decision now, then we would evaluate each strategy in turn and choose the baseline adoption decision with the maximum expected net benefit, which we denote ENB0. ENB0, the expected net benefit given no additional information, is given by

$$\text{ENB0} = \max_d \left[E_\theta \{ \text{NB}(d, \theta) \} \right] \qquad (1)$$

E_θ denotes an expectation over the full joint distribution of θ, that is in integral notation:

$$E_\theta[f(\theta)] = \int_\theta f(\theta)\, p(\theta)\, d\theta$$

Now consider the situation where we might conduct some experiment or gain clairvoyance to learn the true values of the full vector of model parameters θ. Then, since we now know everything, we can choose with certainty the decision that maximises net benefit i.e. $\max_d \{ \text{NB}(d, \theta_{\text{true}}) \}$. This naturally depends on θ_{true}, which is unknown before the experiment, but we can consider the expectation of this net benefit by integrating over the uncertain θ.

$$\text{Expected net benefit given perfect information} = E_\theta \left(\max_d \left[\text{NB}(d, \theta) \right] \right) \qquad (2)$$

The overall EVPI is the difference between these two (2)-(1),

$$\text{EVPI} = E_\theta \left(\max_d \left[\text{NB}(d, \theta) \right] \right) - \max_d \left[E_\theta \{ \text{NB}(d, \theta) \} \right] \qquad (3)$$

It can be shown that this is always positive.

Partial EVPI

Now suppose that θ is divided into two subsets, θ^i and its complement θ^c, and we wish to know the expected value of perfect information about θ^i. If we have to make a decision now, then the expected

net benefit is ENB0 again, but now consider the situation where we have conducted some experiment to learn the true values of the components of $\theta^i = \theta^i_{true}$. Now θ^c is still uncertain, and that uncertainty is described by its conditional distribution, conditional on the value of θ^i_{true}. So we would now make the decision that maximises the expectation of net benefit over that distribution. This is therefore $ENB(\theta^i_{true}) = \max_d \left[E_{\theta^c|\theta^i true} \{NB(d,\theta)\} \right]$. Again, this depends on θ^i_{true}, which is unknown before the experiment, but we can consider the expectation of this net benefit by integrating over the uncertain θ^i.

Expected Net benefit given perfect info only on θ^i $= E_{\theta^i} \left(\max_d \left[E_{\theta^c|\theta^i} \{NB(d,\theta)\} \right] \right)$ (4).

Hence, the partial EVPI for θ^i is the difference between (4) and ENB0, i.e.

$$EVPI(\theta^i) = E_{\theta^i} \left(\max_d \left[E_{\theta^c|\theta^i} \{NB(d,\theta)\} \right] \right) - \max_d \left[E_\theta \{NB(d,\theta)\} \right] \quad (5)$$

This is necessarily positive and is also necessarily less than the overall EVPI.

Equation (5) clearly shows two expectations. The inner expectation evaluates the net benefit over the remaining uncertain parameters θ^c conditional on θ^i. The outer evaluates the net benefit over the parameters of interest θ^i. The conditioning on θ^i in the inner expectation is significant. In general, we expect that learning the true value of θ^i could also provide some information about θ^c. Hence the correct distribution to use for the inner expectation is the conditional distribution that represents the remaining uncertainty in θ^c after learning θ^i. The exception is when θ^i and θ^c are independent, allowing the unconditional (marginal) distribution of θ^c to be used in the inner expectation. The two nested expectations, one with respect to the distribution of θ^i and the other with respect to the distribution of θ^c given θ^i, may seem to involve simply taking an expectation over all the components of θ, but it is very important that the two expectations are evaluated separately because of the need to compute a maximum between them. It is this maximisation between the expectations that makes the computation of partial EVPI complex.

177

178 **COMPUTATION**

179

180 Three techniques are commonly used in statistics to evaluate expectations. The first is when there is an
181 analytic solution to the integral using mathematics. For instance, if X has a normal distribution with
182 mean μ and variance σ^2 then we can analytically evaluate the expectation of functions f(X) = X or X^2 or
183 of exp(X) i.e. $E[X] = \mu$; $E[X^2] = \mu^2 + \sigma^2$; $E[\exp(X)] = \exp(\mu + \sigma^2/2)$. This is the ideal but is all too often
184 not possible in practice. For instance, there is no analytical closed-form expression for $E[(1 + X^2)^{-1}]$.
185 The second common technique is quadrature, also known as numerical integration. There are many
186 alternative methods of quadrature which involve evaluating the value of the function to be integrated at a
187 number of points and computing a weighted average of the results[35]. A very simple example would
188 evaluate the net benefit function at particular percentiles of the distribution (e.g. at the 1^{st}, 3^{rd} ,5^{th} … 99^{th}
189 percentile) and average the results. Quadrature is particularly effective for low-dimensional integrals,
190 and therefore for computing expectations with respect to the distribution of a single or a small number of
191 uncertain variables. When larger numbers of variables exist, the computational load becomes
192 impractical. The third technique is Monte-Carlo sampling. This is a very popular method, because it is
193 very simple to implement in many situations. To evaluate the expectation of a function f(X) of an
194 uncertain quantity X, we randomly sample a large number, say N, of values from the probability
195 distribution of X. Denoting these by $X_1, X_2, \ldots ; X_N$, we then estimate $E\{f(X)\}$ by the sample mean
196 $\hat{E}\{f(X)\} = \frac{1}{N}\sum_{n=1}^{N} f(X_n)$. This estimate is unbiased and its accuracy improves with increasing N.
197 Hence, given a large enough sample we can suppose that $\hat{E}\{f(X)\}$ is an essentially exact computation
198 of $E\{f(X)\}$. It is the Monte-Carlo sampling approach which we now focus upon.

199

200 **Two-level Monte-Carlo computation of partial EVPI**

201

Box 1 displays a detailed description of a Monte- Carlo sampling algorithm to evaluate the expectations when estimating overall and partial EVPI. The process involves two nested simulation loops because the first term in (5) involves two nested expectations. The outer loop undertakes K samples of θ^i. In the inner loop it is important that many (J) values of θ^c are sampled from their conditional distribution, conditional on the value for θ^i that has been sampled in the outer loop. If θ^i and θ^c are independent we can sample from the unconditional distribution of θ^c. Note that, although the EVPI calculation depends on the societal value of health benefits λ, the whole algorithm does not need repeating for different λ thresholds. If the mean cost and mean effectiveness are recorded separately for each strategy at the end of each inner loop, then partial EVPI is quick to calculate for any λ. When evaluating overall EVPI, the inner loop is redundant because there are no remaining uncertain parameters and the process is similar to producing a cost-effectiveness plane[36] or a cost-effectiveness acceptability curve[37].

We can use summation notation to describe these Monte-Carlo estimates. We define the following:

θ_k^i is the k'th random Monte-Carlo sample of the vector of parameters of interest θ^i,

θ^c_{jk} is the *j*th sample taken from the conditional distribution of θ^c given that $\theta^i = \theta_k^i$.

θ_n is the vector of the n'th random Monte-Carlo samples of the full set of parameters θ, and

D is the number of decision policies.

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

$$\text{Estimated partial EVPI} = \frac{1}{K}\sum_{k=1}^{K}\left(\max_{d=1 to D}\left(\frac{1}{J}\sum_{j=1}^{J}\left[NB\left(d,\theta_k^i,\theta_{jk}^c\right)\right]\right)\right) - \max_{d=1 to D}\left[\frac{1}{L}\sum_{l=1}^{L}\left(\text{NB}(\text{d},\theta_l)\right)\right], \qquad (5s)$$

where, K is the number of different sampled values of parameters of interest θ^i; J, the number of different sampled values for the other parameters θ^c conditional upon each given θ_k^i; L, the number of different sampled values of all the parameters together when calculating the expected net benefit of the baseline adoption decision.

Felli and Hazen[4,15] gave a different Monte-Carlo procedure known as MC1 (see Appendix 1). When compared with Box 1, there are two important differences. The first is that MC1 appears as a single loop. Felli and Hazen assume that there is an algebraic expression for the expected payoff conditional on knowing θ^i, and thus the inner expectation in the first term of (5) can be evaluated analytically without using an inner Monte-Carlo sampling loop. This is not always possible and the inner loop in Box 1 provides a generalised method for any net benefit function. Note also that, although the procedure takes a concurrent random sample of the parameters of interest (θ^i) and the remaining parameters (θ^c), the assumption of an algebraic expression for the expected payoff is still made, and the sampling of θ^c is not used to evaluate the inner expectation. The second difference is that MC1 step 2ii recommends estimating the improvement obtained given the information, immediately as each sample of the parameters of interest is taken. Our 2 level algorithm can be amended to estimate the improvement given by the revised decision $d^*(\theta^i)$ over the baseline adoption decision d* at the end of each outer loop iteration (see Box 2).

The Box 2 algorithm is based on an alternative formula for partial EVPI, which combines the first and second terms of (5) into a single expectation.

$$\text{EVPI}(\theta^i) = E_{\theta^i}\left(\max_d\left[E_{\theta^c|\theta^i}\{\text{NB}(d,\theta)\}\right] - E_{\theta^c|\theta^i}\{\text{NB}(d^*,\theta)\}\right). \qquad (6)$$

The summation notation provides a mathematical description of the Box 2 estimate:

$$\text{EVPI}(\theta^i)\ \text{estimate} = \frac{1}{K}\sum_{k=1}^{K}\left(\max_{d=1 to D}\left[\frac{1}{J}\sum_{j=1}^{J}\left\{NB\left(d,\theta_k^i,\theta_{jk}^c\right)\right\}\right] - \frac{1}{J}\sum_{j=1}^{J}\left\{NB\left(d^*,\theta_k^i,\theta_{jk}^c\right)\right\}\right), \qquad (6s)$$

With large numbers of samples the estimates provided by the general algorithm (Box 1) and that computing improvement at each iteration (Box 2) will be equivalent. The difference between them concerns when to estimate the improvement. In Box 1 we estimate the second term of (5s) just once for the whole decision problem. In Box 2, we make K estimates of the improvement versus the baseline

adoption decision conditional on knowing the parameter of interest. If the same numbers of inner and outer samples are taken, then there is little difference in computation time because the same total number of samples and net benefit function evaluations are undertaken in both. The potential advantage of Box 2 is that the improvement is computed as exactly zero whenever the revised decision $d^*(\theta^i) = d^*$. Because of this, with small numbers of samples the Box 2 algorithm might have some marginal reduction in noise compared with Box 1. Furthermore, if the net benefit functions are positively correlated, then the Box 2 algorithm is less susceptible to noise and will provide marginally more accurate partial EVPI estimates for a given small number of samples. The number of Monte-Carlo samples required is our next consideration.

Monte-Carlo Sampling Error

Monte-Carlo sampling estimates of any expectations including those in (5) are subject to potential error. Consider a function *f* of parameters θ, for which the true mean $E_\theta[f(\theta)]$ is say μ. The estimator

$$\hat{\mu} = \frac{1}{N}\sum_{j=1}^{N}\left[\mathrm{f}(\theta_j)\right] \tag{7}$$

is an unbiased estimator of the true mean μ. The standard approach to ensuring that a Monte-Carlo expectation is estimated with sufficient accuracy is to increase the number of samples N, until the standard error of the estimator, $\mathrm{S.E.}(\hat{\mu})$, is less than some defined acceptable level. The Monte-Carlo sampling process provides us with an estimate of the variance of $f(\theta)$,

$$\hat{\sigma}^2 = \frac{1}{N-1}\sum_{j=1}^{N}\left(\mathrm{f}(\theta_j) - \hat{\mu}\right)^2 \tag{8}$$

and the estimated standard error of the Monte-Carlo estimator is defined by

$$\hat{\mathrm{s}} = \mathrm{S.E.}\left(\hat{\mu}\right) = \frac{\hat{\sigma}}{\sqrt{N}} \tag{9}$$

The standard error in the Monte-Carlo estimate of an expectation S.E.($\hat{\mu}$) reduces in proportion to the square root of the number of random Monte-Carlo samples taken.

Applying this approach to estimating the net benefits given current information is straightforward. For each decision option we can consider f(θ)=NB(d,θ) and denote the estimators of expected net benefit E_θ[NB(d,θ)] as $\hat{\mu}_d$, with associated variance estimators $\hat{\sigma}_d$ and standard errors $\hat{s}_d$. Running a probabilistic sensitivity analysis (as in steps 1 to 3 of Box 1), we can establish the mean and variance estimators and choose a sample size N to achieve a chosen acceptable level of standard error.

However, estimating the potential Monte-Carlo error in partial EVPI computation is more complex because we have a nested loop when we are repeatedly estimating expectations. In computing partial EVPI, we have K outer loops, and for each sampled θ^i_k we estimate the *conditional* expected net benefit using J samples of $\theta^c|\theta^i_k$ in the inner loop. We can denote the Monte-Carlo estimator of the expected net benefit for decision option *d* conditional on a particular value of the parameters of interest θ^i_k, as

$$\hat{\mu}_{dk} = \frac{1}{J}\sum_{j=1}^{J}\left[NB\left(d,\theta_k^i,\theta_{jk}^c\right)\right] \quad (10)$$

Denoting $\hat{\sigma}_{dk}$ as the estimator of the variance in the net benefit conditional on the k'th sample $\theta^i_{k,}$ then the standard error of this Carlo estimate is therefore estimated by:

$$\hat{S}_{dk} = \text{S.E.}\left(\hat{\mu}_{dk}\right) = \frac{\hat{\sigma}_{dk}}{\sqrt{J}} = \sqrt{\frac{1}{J}\frac{1}{(J-1)}\sum_{j=1}^{N}\left(NB\left(d,\theta_k^i,\theta_{jk}^c\right) - \hat{\mu}_{dk}\right)^2} \quad (11)$$

We might expect that the standard error of the estimated conditional expected net benefit $\hat{s}_{dk}$ will be lower than the overall standard error $\hat{s}_d$, because we have learned the value of sample θ^i_k and hence reduced uncertainty. If it is, then the number of inner loop samples required to reach a specified

tolerance level could reduce. However, this will not necessarily always be the case and we give an example in the case study section when knowing θ^i_k. is at a particular value can actually increase the variance in net benefit and the standard error. In general it is worth checking how stable these standard errors are for different sampled values of the parameters of interest early in the process of partial EVPI computation.

Having estimated the conditional expected net benefit for each of the D options, we take the maximum. The partial EVPI estimate is therefore made up of K*D Monte-Carlo expectations, each estimated with error, within which K maximisations take place. With the maximisation taking place between the inner and the outer expectations there is no analytic form for describing the standard error in the partial estimate. Oakley et al. have recently developed a first suggestion for an algorithmic process for this estimation based on small numbers of runs[38]. This process of taking the maximum of Monte-Carlo estimates has one further important effect.

Bias when taking maxima of Monte-Carlo expectations

Although the Monte-Carlo estimate of an expectation is unbiased, it turns out that the estimate of the maximum of these expectations is biased, and biased upwards. To see this, consider 2 treatments with net benefit functions NB1(θ) and NB2(θ) with true but unknown expectations μ_1 and μ_2 respectively . If μ_1 and μ_2 are quite different from each other then any error in the Monte-Carlo estimators $\hat{\mu}_1 = \frac{1}{N}\sum_{j=1}^{N}\left[\text{NB1}(\theta_j)\right]$ and $\hat{\mu}_2 = \frac{1}{N}\sum_{j=1}^{N}\left[\text{NB2}(\theta_j)\right]$ is unlikely to affect which treatment is estimated to have the highest expected net benefit. However, if μ_1 and μ_2 are close, then the Monte-Carlo sampling error can cause us to mistakenly believe that the other treatment has the higher expectation, and this will tend to cause us to over-estimate the maximum. Mathematically, we have that

$$\text{E}[\max\{\hat{\mu}_1, \hat{\mu}_2)] \geq \max\{\text{E}[\hat{\mu}_1], \text{E}[\hat{\mu}_2]\} = \max\{\text{E[NB1]},\text{E[NB2]}\} = \max\{\mu_1, \mu_2\} \quad (12)$$

Thus, the process of taking the maximum of the expectations (when they are estimated via a small number of Monte-Carlo samples) creates a *bias* i.e. an *expected error* due to Monte-Carlo sampling.

The bias affects partial EVPI estimates because we evaluate maxima of expectations in both the first and second terms of (5s). For the first term, the process of estimating the maximum of Monte-Carlo expectations is undertaken for each different sample of the parameters of interest (θ_k^i). Each of the K evaluations is biased upwards and therefore the first term in (5s) is biased upwards. The larger the number of samples J in the inner loop, the more accurate and less biased the estimator $\hat{\mu}_{dk}$ given each θik. The larger the number of samples K in the outer loop the more accurate the average of the maximum expected net benefits i.e. $\hat{\mu}(\theta^i) = \frac{1}{K}\sum_{k=1}^{K} \max_d \{\hat{\mu}_{dk}\}$. If J is small and K is very large then we will get a very accurate estimate of the wrong i.e. biased partial EVPI. If $\hat{\mu}_d(\theta^i)$ is the Monte-Carlo estimator of expected net benefit for decision option d given parameters θi , and $\mu_d(\theta^i)$ is the true expected net benefit for decision option d given parameters θi, then the size of the expected bias in the first term of (5s) is given by the formula:

$$\text{Expected Bias in first term of (5s)} = E_{\theta^i}\left(E_{\theta^c|\theta^i}\left(\max_d\left[\hat{\mu}_d(\theta^i)\right]\right) - \max_d\left[\mu_d(\theta^i)\right]\right) \quad (13)$$

The magnitude of the bias is directly linked to the degree of separation between the true expected net benefits. When the expected net benefits for competing treatments are close, and hence parameters have an appreciable partial EVPI, then the bias is higher.

Because the second term in (5s) is also upwards biased, the overall bias in partial EVPI estimates can be either upwards or downwards. The size and direction of the bias will depend on the net benefit functions, the characterised uncertainty and the numbers of samples used. Increasing the sample size J

reduces the bias of the first term. Increasing the sample size L reduces the bias of the second term. If we compute the baseline adoption decision's net benefit with very large L, but compute the first term with very small number of inner loops J, then such partial EVPI computations will be upward biased. It is important also to note that the size K of the outer sample in the 2-level calculation does not affect bias. For overall EVPI, the first term in (3s) is unbiased but the second (negative) term is biased upwards and hence, the Monte-Carlo estimate of overall EVPI is biased downwards. As with Monte-Carlo error in partial EVPI estimates, the size of the expected bias cannot generally be calculated analytically. The investigation of methods to develop an algorithm for this bias estimation is continuing.

There are two separate effects of using Monte-Carlo sampling to estimate the first term in (5) – the random error if J and K are small and the bias if J is small. The bias will decrease with increasing inner loop sample sizes, but for a chosen acceptable accuracy we typically need much larger sample sizes when computing EVPI than when computing a single expectation. We investigate some of the stability of partial EVPI estimates for different inner and outer sample numbers in the case studies. We also examine a very simple 2 treatment decision problem, in which it is possible to compute the bias in formula (13) analytically.

The 'Short-Cut' 1 Level Algorithm

In some simple models, it is possible to evaluate expectations of net benefit analytically, particularly if parameters are independent. Suppose $NB(\theta)=\lambda*\theta_1 - \theta_2* \theta_3$, and the parameters θ_2 and θ_3 are independent, so that the expected net benefit can be calculated analytically simply by running the model with the parameters set equal to their mean values, $E_\theta\{\mathrm{NB}(d,\theta)\} = \lambda * \overline{\theta_1} - \overline{\theta_2} * \overline{\theta_3}$. Although simple, there are economic models in practice, particularly decision tree models, which are of this form.

In such circumstances, the 2 level partial EVPI algorithm can be simplified to a 1 level process (Box 3). This performs a one level Monte-Carlo sampling process, allowing parameters of interest to vary, keeping remaining uncertain parameters constant at their prior means. It is much more efficient than the two- level Monte-Carlo method, since we replace the many model runs by a single run in each of the expectations that can be evaluated without Monte Carlo. Mathematically, we compute analytic solutions for the inner expectations in the 1st term of (5) and all of the expectations in the 2nd term of (5). Note that the expectations of maxima cannot be evaluated in this way. Thus, the expectation in the first term of (3) and the outer expectation in the first term of (5) are still evaluated by Monte-Carlo in Box 3. Felli and Hazen give a similar procedure, which they term a 'shortcut' (MC2) and is identical to MC1 described earlier but with those parameters not of interest set to their prior means i.e. $\theta^c = \overline{\theta^c}$. Note that a misunderstanding of the Felli and Hazen 'short cut' method previously led some analysts to use a quite inappropriate algorithm, which focussed on reduction in opportunity loss[16,17]. The level of inaccuracy in estimating partial EVPI which resulted from this incorrect algorithm is discussed elsewhere.

The 1 level algorithm is correct under the following conditions. Mathematically, the outer level expectation over the parameter set of interest θ^i is as per equation (5), but the inner expectation is replaced with net benefit calculated given the remaining uncertain parameters θ^c set at their prior mean.

1 level partial EVPI for $\theta i = E_{\theta^i}\left\{\max_d\left[\mathrm{NB}(d,\theta^i,\overline{\theta^c})\right]\right\} - \max_d\left\{E_\theta \mathrm{NB}(d,\theta)\right\}$ (14)

Note that we now have just one expectation, and that the 1-level approach is equivalent to the 2 level algorithm if (5) ≡ (14), i.e. if

$$E_{\theta^i}\left(\max_d\left[E_{\theta^c|\theta^i}\left\{\mathrm{NB}(d,\theta)\right\}\right]\right) \equiv E_{\theta^i}\left\{\max_d\left[\mathrm{NB}(d,\theta^i,\overline{\theta^c})\right]\right\} \qquad (15)$$

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

Felli and Hazen comment that the 1 level procedure can apply successfully "when all parameters are assumed probabilistically independent and the pay-off function is multi-linear i.e. linear in each individual parameter", in other words condition (15) will hold if:

A1. For each d the function NB(d, θ) can be expressed as a sum of products of components of θ

A2. All of the components of θ are mutually probabilistically independent of each other.

Condition (15) will also hold in a second circumstance. It is not necessary for *all* of the parameters to be independent of each other provided that the net benefit functions are linear. In fact, the 1 level procedure can apply successfully for *any* chosen partition of the parameter vector θ into parameters of interest θ^i, and their complement θ^c if the conditions below are satisfied:

B1. For each d, the function $NB(d, \theta) = NB(d, \theta^i, \theta^c)$ is a linear function of the components of θ^c, whose coefficients may depend on d and θ^i. If θ^c has m components, this linear structure takes the form $NB(d, \theta^i, \theta^c) = A1(d, \theta^i)\times\theta^c(1) + A2(d, \theta^i)\times\theta^c(2) + \ldots + Am(d, \theta^i) \times\theta^c(m) + b(d, \theta^i)$.

B2. The parameters θ^c are probabilistically independent of the parameters θ^i.

Thus, provided the net benefit function takes the form in sufficient condition (B1), then the one-level algorithm will be correct in the cases where there are (a) no correlations at all, (b) correlations only within θ^i, (c) correlations only within θ^c, or (d) correlations within θ^i and within θ^c but no correlations between θ^i and θ^c. If the net benefits are linear functions of the parameters, it is only when the correlations are between members of θ^c and θ^i that the 1 level algorithm will be incorrect.

The specifications of the sufficient conditions in (A1,A2) and (B1,B2) above are actually slightly stronger than the necessary condition expressed mathematically in (15) but it is unlikely in practice that the one-level algorithm would correctly compute partial EVPI in any economic model for which one or other of the two circumstances described did not hold. In the next section we consider how accurate the shortcut 1-level estimate might be as the parameters move from independent to being more highly correlated, and as the net benefit functions move from linear to greater non-linearity.

CASE STUDIES

Case Study Model 1: Analytically tractable model to illustrate effects of bias

Case study 1 has 2 treatments with a very simple pair of net benefit functions, NB1 = 20,000*θ1, NB2 = 19,500*θ2, where θ1 and θ2 are statistically independent uncertain parameters each with a normal distribution N(1,1). Analytically, we can evaluate max{E(NB1), E(NB2)} as max{20000,19500} = 20,000. We compare the analytic results with repeatedly using very small numbers of Monte-Carlo samples to evaluate the expectations of NB1 and NB2, and illustrate the scale of the bias due to taking maxima of two Monte-Carlo estimated expectations. In this very simple example with statistically independent, normally distributed net benefit functions, it is also possible to derive analytically, both the partial EVPI's and the expected bias due to taking maxima of Monte-Carlo estimated expectations.

Case Study 1 Results - Bias

In all of the case study results, the partial EVPI estimates are presented not in absolute financial value terms but rather relative to the overall EVPI for the decision problem. Thus, if we have an overall EVPI of say £1400, which we 'index' to 100, then a partial EVPI of £350 would be reported as 'indexed partial EVPI' = 25.

The effect of Monte-Carlo error induced bias in partial EVPI estimates depends upon the numbers of inner samples J used in the first term (5s) and the number of samples L used to estimated the expected net benefit of the baseline adoption decision in the second term of (5s). In this very simple example with

439 statistically independent, normally distributed net benefit functions, it is actually possible to derive
440 analytically, both the partial EVPIs and the bias due taking maxima of Monte-Carlo estimated
441 expectations (See Appendix 2). Table 3 shows the resulting bias for a range of J and L sample sizes.
442 When L is small, the second term in (5s) is over-estimated due to the bias. In this case study the effect is
443 strong enough, for example at L=1000, that the partial EVPI estimate is actually downwards biased for
444 any value of J over 100. As L is increased the second term converges to its true value. When J is small
445 and L is large, we can expect the first term in (5s) to be over-estimated and the resulting partial EVPI
446 estimate to be upwards biased. The bias when J=100 is 0.49% of the true EVPI, and this decreases to
447 0.1% at J=500 and 0.05% at J=1,000. Note that the actual error in a Monte-Carlo estimated EVPI can be
448 considerably greater than this on any one run if small numbers of outer samples are used because over
449 and above this bias we have the usual Monte-Carlo sampling error also in play.
450
451 **Case Study Model 2: Accuracy of 1 level estimate in a decision tree model with correlations**
452
453 The second case study is a decision tree model comparing two drug treatments T0 and T1 (Table 1).
454 Costs and benefits for each strategy depend upon 19 uncertain parameters characterised with
455 multivariate normal distributions. We examine 5 different levels of correlation (0, 0.1, 0.2, 0.3, 0.6)
456 between 6 different parameters. Zero correlation of course implies independence between all of the
457 parameters. Correlations are anticipated between the parameters concerning the two drugs' mean
458 response rates and the mean durations of response i.e. θ5, θ7, θ14 and θ16 all are correlated with each
459 other. Secondly, correlations are anticipated between the two drugs' expected utility improvements, θ6
460 and θ15. To implement this model we randomly sample the multi-variate normal correlated values
461 using [R] statistical software[39]. We also implemented an extension of Cholesky decomposition in
462 EXCEL Visual Basic to create a new EXCEL function =MultiVariateNormalInv (see CHEBS
463 website)[40].

464

465 **Case Study 2 Results – Effects of Correlation on Accuracy of 1 Level Algorithm**

466

467 In the circumstance where correlation is zero, Figure 1 shows 1 level and 2 level partial EVPI estimates
468 for a range of parameter(s) of interest. The estimates are almost equivalent, with the 2 level estimates
469 just slightly higher than the 1 level estimates for each of the parameter(s) of interest examined. The
470 largest difference is just 3% of the overall EVPI. This reflects the mathematical results that (a) the 1
471 level and 2 level EVPI should be equivalent, because the cost-effectiveness model has net benefit
472 functions that are sum-products of statistically independent parameters, and (b) the 2 level estimates are
473 upwardly biased due to the maximisation of Monte-Carlo estimate in the inner loop. Note also that
474 partial EVPI for groups of parameters is lower than the sum of the EVPIs of individual parameters e.g.
475 utility parameters combined ($\theta 6$ and $\theta 15$) = 57%, compared with individual utility parameters =
476 46%+24% = 70%.

477

478 If correlations are present between the parameters, then the 1 level EVPI results sometimes substantially
479 under estimate the true EVPI. The 1 level and 2 level EVPI estimates are broadly the same when small
480 correlations are introduced between the important parameters. For example, with correlations of 0.1, the
481 2 level result for the utility parameters combined ($\theta 6$ and $\theta 15$) is 58%, 6 percentage points higher than
482 the 1 level estimate. However, if larger correlations exist, then the 1 level EVPI 'short-cut' estimates
483 can be very wrong. With correlations of 0.6, the 2 level result for the utility parameters combined ($\theta 6$
484 and $\theta 15$) is 18 percentage points higher than the 1 level estimate, whilst for the response rate parameters
485 combined ($\theta 5$ and $\theta 14$) shows the maximum disparity seen, at 36 percentage points. As correlation is
486 increased the disparity between 2 level and 1 level estimates increases substantially. The results
487 demonstrate that having linear or sum-product net benefit functions is not a sufficient condition for the 1

488 level EVPI estimates to be accurate and that the second mathematical condition, i.e. that parameters are
489 statistically independent, is just as important as the first.
490
491 The 1 level EVPI results should be the same no matter what level of correlation is involved, because the
492 1 level algorithm sets the remaining parameters θ^c at their prior mean values no matter what values are
493 sampled for the parameters of interest. The small differences shown in Fig 1 between different 1 level
494 estimates are due to random chance of different samples of θ^i. The 2 level algorithm correctly accounts
495 for correlation, by sampling the remaining parameters from their *conditional* probability distributions
496 within the inner loop. It could be sensible to put the conditional mean for θ^c given θ^i into the 1 level
497 algorithm rather than the prior mean, but only in the very restricted circumstance when the elements of
498 θ^c are conditionally independent given θ^i and the net benefit function is multi-linear. In case study 2,
499 such a method would not apply for any of the subgroups of parameters examined, because the elements
500 of the vector of remaining parameters θ^c are correlated with each other.
501
502 **Case Study Model 3: Accuracy of 1 level estimate in an increasingly non-linear Markov model**
503
504 Case study 3 extends the Case study 2 model incorporating a Markov model for the natural history of
505 continued response. Table 2 shows that the parameters for mean duration of response (θ7 and θ16) are
506 replaced with 2 Markov models of natural history of response to each drug with health states
507 "responding", "not responding" and "died" (θ20 to θ31). The mean duration of response to each drug is
508 now a function of multiple powers of Markov transition matrices. To investigate the effects of
509 increasingly non-linear models, we have analysed time horizons of Ptotal = 3, 5, 10, 15 and 20 periods
510 in a Dirichlet distribution. To implement the models we sampled from the Dirichlet distribution in the
511 statistical software R[41], and also extended the method of Briggs[42] to create a new EXCEL Visual Basic
512 function = DirichletInv. We have characterised the level of uncertainty in these probabilities by

assuming that each is based on evidence from a small sample of just 10 transitions. We use a Bayesian framework with a uniform prior of Dirchlet(1,1,1), and thus the posterior transition rates used in sampling for those "responding" to the health states "responding", "not responding" and "died" are Dirichlet(7,4,2) and the equivalent transition rates for non-responders are Dirichlet (1,10,2).. We have assumed statistical independence between the transition probabilities for those still responding and those no longer responding and also between the transition probabilities for T1 and T0.

Case Study 3 Results – Effects of Non-Linearity on Accuracy of 1 Level Algorithm

We investigated the extent of non-linearity for each Markov model by expressing the net benefits as functions of the individual parameters using simple linear regression and noting the resulting adjusted R^2 for each. Increasing the number of periods in Markov model (e.g. 3, 5, 10, 15, 20) results in greater non-linearity (i.e decreasing adjusted R^2 = 0.97, 0.95, 0.90, 0.87, 0.83 respectively). Figure 2 shows the effects on partial EVPI estimates. The 1 level estimates are substantially lower than the 2 level for the trial (θ5, θ14) and utility parameters (θ6, θ15) and for their combination. Indeed, the 1 level partial EVPI estimates are actually negative for the trial parameters (θ5, θ14) 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 $E_{\theta^i}\left\{\max_d\left[\mathrm{NB}(\mathrm{d},\theta \mid \theta^c = \overline{\theta^c})\right]\right\}$ is actually lower than the second term, $\max_d\{E_\theta \mathrm{NB}(\mathrm{d},\theta)\}$. Thus, when we set the parameters we are not interested in (θ^c) to their prior means in term 1, the net benefits obtained are lower than in term 2 when we allow all parameters to vary. Estimated partial EVPI for the Markov transition probabilities for duration of disease (θ^i = θ20 to θ31) show a high degree of alignment between the 1 level and 2 level methods. This is because, after conditioning on θ^i the net benefit functions are now linear in the remaining statistically independent parameters. It is very important to note that even quite high adjusted R^2 does not imply that 1 level and 2 level estimates will be equal or even of the same order of magnitude. For example for trial parameters (θ5, θ14) when correlation is set

538 at 0.1, the adjusted R^2 is 0.973 but the 2 level EVPI estimate is 30 compared with a 1 level of 19. This
539 suggests that the 2 level EVPI algorithm may be necessary, even in non-linear Markov models very well
540 approximated by linear regression.
541
542 **Results On Numbers of Inner and Outer Samples Required**
543
544 We can use the Monte-Carlo sampling process to quantify the standard errors in expected net benefits
545 for a given number of samples quite easily. For example, 1000 samples in case study 2 with zero
546 correlation provided an estimator for the mean[NB(T0)] $\hat{\mu}_{T0}$ = £5,006, with an estimator for the sample
547 standard deviation [NB(T0)] $\hat{\sigma}_{T0}$ = £2510, giving a standard error of $\left(\hat{\sigma}_{T0}/\sqrt{1000}\right)$ = £2.51. The
548 equivalent figures for T1 are mean estimator £5351, sample standard deviation estimator £2864 and
549 standard error £2.87. This shows clearly that the 95% confidence intervals for the expected net benefits
550 (£5006±5 and £5351±6) do not overlap and we can see that 1000 samples is enough to indicate that the
551 expected net benefit of T1 given current information is higher than that for T0.
552
553 As discussed earlier, it is likely that, conditioning on knowing the value of θ^i_k, will give estimators of the
554 variance in net benefits $\hat{\sigma}_{dk}$ which will be lower than the prior variance $\hat{\sigma}_d$ because knowing θ^i_k means
555 we are generally less uncertain about net benefits. However, this is not necessarily always the case, and
556 it is possible that posterior variance can be greater. When estimating EVPI(θ_7) in case study 2 with zero
557 correlation, we found for example that our k=4th sampled value (θ^i_4 = 4.4 years) in the outer loop
558 combined with J=1000 inner samples provided a higher standard error $\left(\hat{\sigma}_{T0}/\sqrt{1000}\right)$ = £3.25 as
559 compared with £2.51.
560

We further examined the number of Monte-Carlo samples required for accurate unbiased estimates of partial EVPI using case study 2, assuming zero correlation, and focusing only on the partial EVPI for parameters (θ5 and θ14). Figure 3 illustrates how the estimate converges as increasing numbers of inner and outer samples are used. With very small numbers of inner and outer level samples the partial EVPI estimate can be wrong by an order of magnitude. For example, with J=10 and K=10, we estimated the indexed EVPI(θ5,θ14) at 44 compared to a converged estimate of 25.using J=10,000 and K= 1,000. However, even with these quite small numbers of samples the fact that the current uncertainty in variables θ5 and θ14 is important in the decision between treatments is revealed. As the numbers of inner and outer samples used are extended cumulatively in Figure 3, the partial EVPI result begins to converge. The order of magnitude of the EVPI(θ5,θ14) estimates is stable to within 2 indexed percentage points once we have extended the sample beyond K=100 outer and J=500 inner samples. The number of samples needed for full convergence is not symmetrical for J and K. For example, over K=500 the EVPI(θ5,θ14) estimate converges to within 1 percentage point, but for the inner level, where there is a 4 point difference between J=750 and J=1000 samples, and it requires samples of J=5,000 to 10,000 to converge to within 1 percentage point. 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.

Of course, the acceptable level of error when calculating partial EVPI depends upon their use. If analysts want to clarify broad rankings of sensitivity or information value for model parameters then knowing whether the indexed partial EVPI is 62, 70 or 78 is probably irrelevant and a standard deviation of 4 may well be acceptable. If the exact value needs to be established within 1 indexed percentage point then higher numbers of samples will be necessary.

Having seen that K=100, J=500 produced relatively stable results for one parameter set in Case study 2, we decided to investigate the stability of partial EVPI estimates using relatively small numbers of

586 samples in four different parameter groups using the 5 models in case study 3 i.e. 20 parameter sets in
587 total. By repeatedly estimating the partial EVPI, we were able to produce a distribution of results and
588 hence estimate the standard deviation in the partial EVPI estimates. Figure 4 shows the standard
589 deviations obtained for different numbers of inner and outer samples. The results show that when we
590 increase the number of outer samples from (K=100 to K=300, with J set at 500), the standard deviations
591 fall substantially, on average by a factor of 0.62. This is in line with a reduction in proportion to the
592 square root of the number of outer samples i.e. reduction in standard deviation $\propto (\sqrt{100})/(\sqrt{300})=0.58$.
593 In contrast, the reductions in standard deviation due to increases in the number of inner samples are not
594 so marked. When we increase the number of inner samples from (J=100 to J=500, with K set at 100),
595 the standard deviations fall on average by a factor of just 0.89, which is a much smaller reduction than if
596 reductions were in proportion to the square root of the number of inner samples ($\sqrt{100}/\sqrt{500}$)= 0.45.
597 This demonstrates that improving the accuracy of partial EVPI estimates requires proportionately greater
598 effort on the inner level than the outer. It is also clear that the higher the true partial EVPI, the greater
599 the level of noise that might be expected. Figure 5 shows 'confidence intervals', (± 1.96 * s.d.) for the
600 partial EVPI estimates with relatively small numbers of samples. Parameters with low EVPI are
601 estimated with low EVPI even with as small a number of samples as K=100, J=100. Parameters with
602 much higher EVPI's are estimated with relatively high EVPI but also have a larger confidence interval
603 around them.
604
605 Finally, we used case study 3 to compare the algorithm that computes improvement after each iteration
606 (Box 2) with the general algorithm (Box 1), to assess whether estimates might exhibit less noise. We
607 undertook 30 runs using both Box 1 and Box 2 algorithms with K=100 outer and J=100 inner samples.
608 Figure 6a shows the results for the four different parameter sets and five different time period models.
609 The results show that standard deviations in the indexed partial EVPI results are almost equivalent for
610 the Box 2 algorithm compared with the Box 1 algorithm. Over all of the 20 parameters examined, the

611 average reduction in standard deviation in estimates is just 1%. This is because the net benefit functions
612 in case study 3 are almost uncorrelated (the only linked variable is θ4). We then repeated this process,
613 but this time assumed that the natural history of response using the Markov model was the same for both
614 treatments. That is, parameter θ26=θ20, θ27=θ21, … θ31=θ25. Because these parameters are now
615 linked, the net benefit functions for the two treatments are now correlated.(correlation = 0.33, 0.44, 0.59,
616 0.66 and 0.71 for the models with 3, 5, 10, 15 and 20 total periods respectively). Figure 6b shows that
617 the standard deviations of the Box 2 algorithm EVPI estimates are now lower than those for Box 1, with
618 an average reduction in standard deviation in estimates of 9%. The reduction in standard deviation
619 observed was higher for the models with higher correlations in net benefit (estimated reduction in
620 standard deviation in partial EVPI estimates = 1%, 6%, 15%, 11%, and 13% respectively). The standard
621 deviation in partial EVPI estimates is reduced by approximately 2% for every 0.1 increase in the
622 correlation between the net-benefits. Using a square root of n, rule of thumb, this suggests that using the
623 Box 2 algorithm might require roughly 4% fewer samples for every 0.1 increase in correlation between
624 the net-benefits to achieve the same level of accuracy in partial EVPI as the Box 1 algorithm.
625
626 **DISCUSSION**
627
628 This paper describes the calculation of partial EVPI, with the evaluation of two expectations, an outer
629 expectation over the parameter set of interest and an inner expectation over the remaining parameters. A
630 generalised algorithm of nested outer and inner loops can be used to compute Monte-Carlo estimates of
631 the expectations and the maxima required for each outer loop. In specific circumstances, a 'short-cut' 1
632 level algorithm is equivalent to the 2 level algorithm and can be recommended for use in simple models
633 with linear and independent parameters. If net benefits are non-linear functions of parameters, or where
634 model parameters are correlated, the 1 level algorithm can be substantially inaccurate. The scale of
635 inaccuracy increases with non-linearity and correlation, but not always predictably so in scale. Case
636 studies here show the 1 level algorithm under-estimating partial EVPI but elsewhere we have shown a

637 case study where over-estimates are also possible. In practice, the 1 level 'short-cut' algorithm could be
638 useful to screen for parameters which do not require further analysis. If parameters do not affect the
639 decision, our case studies show that their partial EVPI will be very close to zero using both the 2 level
640 and the 1 level algorithm. Thus, the 1 level algorithm might be used with a relatively small number of
641 iterations (e.g. 100) to screen for groups of parameters in very large models. The 2 level Monte-Carlo
642 algorithm is applicable in any model, provided there is computing resource to run a large enough
643 number of samples.
644
645 The number of inner and outer level simulations required depends upon the number of parameters, their
646 importance to the decision, and the model's net benefit functions. The standard error of each Monte-
647 Carlo estimated expectation in the algorithm reduces in proportion to the square root of samples used but
648 when this accumulates over many inner and outer loops and the maxima taken, the standard error of
649 partial EVPI estimates is not generally able to be computed analytically. We recommend analysing the
650 convergence of estimates to ensure a threshold accuracy of partial EVPI estimates fit for the specific
651 purpose of the analysis. Our empirical approach, in a series of alternative models, suggests that the
652 number of inner and outer samples should not in general be equal. In these case studies, 500 inner loops
653 for each of the 100 outer loop iterations (i.e. 50,000 iterations in total) proved capable of estimating the
654 order of magnitude of partial EVPI reasonably well in our examples, although it is likely that higher
655 numbers may be needed in some situations. For very accurate calculation or in computationally
656 intensive models, one might use adaptive processes to test for convergence in the partial EVPI results,
657 within a pre-defined threshold.
658
659 A further consequence of Monte-Carlo sampling error is the existence of an over-estimating bias in
660 evaluating maximum expected net benefit across decision options when using small numbers of samples.
661 This can result in over or under-estimating the partial EVPI depending on the number of iterations used
662 to evaluate the first and second terms. Previous authors have investigated mathematical description of

663 Monte-Carlo bias outside the EVPI context[43]. Again, analytical computation of this bias is generally not
664 possible and analysis of the convergence of estimates as the number of inner samples increases is
665 recommended. In our case studies the bias appeared as no more than 1or 2 percentage points of the
666 overall EVPI when using 1000 inner samples. Further theoretical investigation of Monte-Carlo bias in
667 the context of partial EVPI would be useful and work is ongoing on a theoretical description of the
668 Monte-Carlo bias in partial EVPI calculation, and on using this theory to develop algorithms to quantify
669 the inner level sample size required for a particular threshold of accuracy[38,44].
670
671 The differences between EVPI results using the general algorithm (Box 1) and that computing
672 improvement at each iteration (Box 2) were relatively small in case study 3 when net benefit functions
673 had low correlation. If $EVPI(\theta^i)$ is small, then even small numbers of samples provide good estimates
674 using either algorithm. If $EVPI(\theta^i)$ is large, then on a high proportion of occasions a different decision
675 option would be taken i.e. $d^*(\theta^i_k) \neq d^*$. Box 1 provides K estimates of $E_{\theta c}[NB(d^*(\theta^i_k), \theta)| \theta^i_k] -$
676 $E_\theta[NB(d^*,\theta)]$. In contrast, Box 2 provides K estimates of $E_{\theta c}[\{NB(d^*(\theta^i_k), \theta) - NB(d^*, \theta)\}| \theta^i_k]$. If the
677 net benefit functions are highly positively correlated, then the Box 2 algorithm is less susceptible to
678 noise and provides marginally more accurate partial EVPI estimates for a given number of samples. It is
679 important also to note that if the net benefit functions are negatively correlated then Box 2 estimates
680 would display higher variance than Box 1 estimates. From a computation time perspective, a further
681 refinement to the Box 2 algorithm could also be useful in the circumstance when there are very many
682 strategies and evaluating the net benefit functions takes appreciable computation time. This refinement
683 would use as small a number of inner loop iterations as possible to identify with reasonable certainty
684 which of the many strategies is $d^*(\theta^i_k)$. If $d^*(\theta^i_k) = d^*$, then there is zero improvement and we need no
685 further calculation. If $d^*(\theta^i_k) \neq d^*$, then we can use a larger number of inner loop samples just to
686 estimate the improvement in expected net benefit between the 2 relevant strategies $d^*(\theta^i_k)$ and d^*. Such
687 an adaptive approach can be useful when undertaking large numbers of Monte-Carlo samples becomes
688 too time-consuming.

689

690 There are non-Monte-Carlo methods that can be used to compute partial EVPI. Quadrature often has
691 limited use, because there is often a large number of uncertain parameters in economic models.
692 However, if the number of parameters in either θ^i or θ^c is small, then quadrature can be used for the
693 relevant computations in partial EVPI, and where θ^i is a single parameter, this can cut the number of
694 values of θ^i required from around 1000 (which is what would typically be needed for Monte Carlo) to 10
695 or fewer. A quite different approach is set out in by Oakley et al.[45], who use a Bayesian approach based
696 on the idea of emulating the model with a Gaussian process. Although this method is technically much
697 more complex than Monte Carlo, it can dramatically reduce the number of model runs required and the
698 authors recommend its application if many EVPI calculations are required in a model which has
699 individual runs taking more than a few seconds.

700

701 There remain some areas where further methodological research would be useful. Computing
702 population EVPI demands estimated patient numbers involved in the policy decision. Incidence and
703 prevalence are important, as are the likely lifetime of the technology and potential changes in competitor
704 strategies. There are arguments over the validity of analysing phased adoption of the intervention over
705 time explicitly versus full adoption implied by the decision rule. When trading off against the costs of
706 data collection, timing of data collection is important too. Some parameters may be collectable quickly
707 (e.g. utility for particular health states), others take longer (e.g. long term side-effects), and still others
708 may be inherently unknowable (e.g. the efficacy of an influenza vaccine prior to the arrival of next years
709 strain of influenza).

710

711 EVPI is important, both in decision-making, and in planning and prioritising future data collection.
712 Policy makers assessing interventions are keen to understand the level of uncertainty, and many
713 guidelines recommend probabilistic sensitivity analysis[20]. The common representations of uncertainty,

714 the cost-effectiveness plane and the cost-effectiveness acceptability curve[46] show the relative importance
715 of uncertainty in costs and effectiveness. Partial EVPI extends these by giving the breakdown by
716 parameter, so that decision makers see clearly the source and scale of uncertainty. This paper seeks to
717 encourage analysts to extend the approach to calculation of overall and partial EVPI. The theory and
718 algorithms required are now in place. The case study models have shown the feasibility and performance
719 of the method, indicating the numbers of samples needed for stable results. Wider application will bring
720 greater understanding of decision uncertainty and research priority analysis.

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 $< \lambda$

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 L (say L=10,000) sample sets of uncertain parameter values (Monte Carlo).

3) Work out the baseline adoption decision d* given current information
i.e. the strategy giving (on average over L=10,000 simulations) the highest estimated 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 the parameter subset of interest once from its joint prior distribution **(outer level simulation)**

5) estimate the net benefit of the best strategy given this new knowledge on the parameters of interest by
- fixing the parameters of interest at their sampled values θ^i_k
- simulating the other remaining uncertain parameters $\theta^c_{j\,k}$ (say J=10,000 times) allowing them to vary according to their conditional probability distribution (conditional upon the parameter subset of interest at its sampled value θ^i_k) **(inner level simulation)**
- calculating the conditional expected net benefit of each strategy $E(\theta|\ \theta^i_k)[NB(d,\theta)]$ given θ^i_k by evaluating the net benefit at each $(\theta^c_{j\,k},\ \theta^i_k)$ and averaging
- choosing the revised adoption decision $d^*(\ \theta^i_k\ )$ to be the strategy which has the highest estimated expected net benefit given the sampled value for the parameters of interest

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

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

Overall EVPI

The algorithm for overall EVPI requires only 1 loop (which can be done at the same time as steps (2,3)

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

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

Box 2: 2 level Monte-Carlo Algorithm for Calculation of Partial EVPI using Improvement In each Iteration

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) estimate the net benefit of the best strategy given this new knowledge on the parameters of interest by
- fixing the parameters of interest at their sampled values θ^i_k
- simulating the other remaining uncertain parameters $\theta^c_{j\,k}$ (say J=10,000 times) allowing them to vary according to their conditional probability distribution (conditional upon the parameter subset of interest at its sampled value θ^i_k) **(inner level simulation)**
- calculating the conditional expected net benefit of each strategy $E_{(\theta|\,\theta ik)}[NB(d,\theta)]$ given θ^i_k by evaluating the net benefit at each $(\theta^c_{j\,k},\ \theta^i_k)$ and averaging
- choosing the revised adoption decision $d^*(\theta^i_k)$ to be the strategy which has the highest estimated expected net benefit given the sampled value for the parameters of interest
- **compute improvement in conditional mean net benefit as the difference between the revised decision given θ^i_k and the baseline adoption decision d* given θ^i_k i.e. $E_{(\theta|\,\theta ik)}[NB(d^*(\theta^i_k),\ \theta)] - E_{(\theta|\,\theta ik)}[NB(d^*,\ \theta)]$**

6) Loop back to step 4 and repeat steps 4 and 5 (say K=10,000 times)

7) The EVPI for the parameter of interest = **average of the improvements recorded in step 5**

Box 3: 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 the parameter subset of interest once from its prior distribution **(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 sampled value for the parameters of interest

6) Loop back to step 4 and repeat steps 4 and 5 (say K=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)

773 **Table 1: Case Study 2 Model**

	Treatment T0			Treatment T1		
	Param No.	Prior Mean	Std Dev	Param No.	Prior Mean	Std Dev
Cost of drug	θ 1	£1,000	£1	θ 11	£1,500	£1
% admissions	θ 2	10%	2%	θ 12	8%	2%
Days in Hospital	θ 3	5.20	1.00	θ 13	6.10	1.00
Cost per Day	θ 4	£400	£200	θ 4	£400	£200
% Responding	θ 5	70%	10%	θ 14	80%	10%
Utility change if respond	θ 6	0.3000	0.1000	θ 15	0.3000	0.0500
Duration of response (years)	θ 7	3.0	0.5	θ 16	3.0	1.0
% Side effects	θ 8	25%	10%	θ 17	20%	5%
Change in utility if side effect	θ 9	-0.1000	0.0200	θ 18	-0.1000	0.0200
Duration of side effect (years)	θ 10	0.50	0.20	θ 19	0.50	0.20

774

775 $\lambda = £10{,}000$

776 $NBT0 = \lambda*(\theta5*\theta6*\theta7 + \theta8*\theta9*\theta10) - (\theta1 + \theta2*\theta3*\theta4)$

777 $NBT1 = \lambda*(\theta14*\theta15*\theta16 + \theta17*\theta18*\theta19) - (\theta11 + \theta12*\theta13*\theta4)$

778

Table 2: Case Study 3 Model

	Treatment T0			Treatment T1		
	Param No.	Prior Mean	Std Dev	Param No.	Prior Mean	Std Dev
Cost of drug	θ 1	£1,000	£1	θ 11	£1,500	£1
% admissions	θ 2	10%	2%	θ 12	8%	2%
Days in Hospital	θ 3	5.20	1.00	θ 13	6.10	1.00
Cost per Day	θ 4	£400	£200	θ 4	£400	£200
% Achieving Initial Response	θ 5	70%	10%	θ 14	80%	10%
Utility change if respond	θ 6	0.3000	0.1000	θ 15	0.3000	0.0500
% Side effects	θ 8	25%	10%	θ 17	20%	5%
Change in utility if side effect	θ 9	-0.1000	0.0200	θ 18	-0.1000	0.0200
Natural History Model for Duration of Continued Response if Initial Response is Achieved						
Markov Transition Probabilities						
p(Responding --> Responding)	θ 20	60%		θ 26	60%	
p(Responding --> Not Responding)	θ 21	30%	Dirichlet (7,4,2)	θ 27	30%	Dirichlet (7,4,2)
p(Responding --> Die)	θ 22	10%		θ 28	10%	
p(Not Responding --> Responding)	θ 23	0%		θ 29	0%	
p(Not Responding --> Not Responding)	θ 24	90%	Dirichlet (1,10,2)	θ 30	90%	Dirichlet (1,10,2)
p(Not Responding --> Die)	θ 25	10%		θ 31	10%	
p(Die --> Die)		100%				

θ22 = 1-θ20-θ21. θ25 = 1-θ23-θ24. . θ28 = 1-θ26-θ27 θ31 = 1-θ29-θ30.

$S0 = (\theta 5, 1-\theta 5, 0)^T$, $U0 = (\theta 6, 0, 0)^T$, $S1 = (\theta 14, 1-\theta 14, 0)^T$, $U1 = (\theta 15, 0, 0)^T$.

$$M0 = \begin{pmatrix} \theta 20 & \theta 21 & \theta 22 \\ \theta 23 & \theta 24 & \theta 25 \\ 0 & 0 & 1 \end{pmatrix}, \; M1 = \begin{pmatrix} \theta 26 & \theta 27 & \theta 28 \\ \theta 29 & \theta 30 & \theta 31 \\ 0 & 0 & 1 \end{pmatrix}$$

Net Benefit functions depend upon the number of Markov periods used (Ptotal = 3, 5, 10, 15, 20)

$$NBT0 = \lambda * \left[\left\{ \sum_{p=1}^{Ptotal} \left((S0)^T * (M0)^p * U0 \right) \right\} + \theta 8 + \theta 9 + \vartheta 10 \right] - \left(\theta 1 + \theta 2 * \vartheta 3 * \vartheta 4 \right)$$

$$NBT1 = \lambda * \left[\left\{ \sum_{p=1}^{Ptotal} \left((S1)^T * (M1)^p * U1 \right) \right\} + \theta 17 + \theta 18 + \vartheta 19 \right] - \left(\theta 11 + \theta 12 * \vartheta 13 * \vartheta 4 \right)$$

791

792 **Table 3: Bias in Monte-Carlo Estimates of EVPI Dependent on Number of Samples**

793 **(Bias in partial EVPI for parameter θ^1 in Case Study 1 as a % of its true EVPI)**

794

795

	Number of Samples in 2nd Term of (5s)				
	L= 1,000	3,000	10,000	100,000	1,000,000
Number of Samples in 1st Term of (5s)					
J = 100 .	-1.55%	-0.08%	0.44%	0.49%	0.49%
300	-1.87%	-0.41%	0.11%	0.16%	0.16%
500	-1.94%	-0.47%	0.05%	0.10%	0.10%
1,000	-1.99%	-0.52%	0.00%	0.05%	0.05%
10,000	-2.03%	-0.57%	-0.05%	0.00%	0.00%
100,000	-2.03%	-0.57%	-0.05%	0.00%	0.00%

796

797

798 **Figure 1: Impact of Increasing Correlation on Inaccuracy of 1 level method to calculate partial EVPI**
799

800

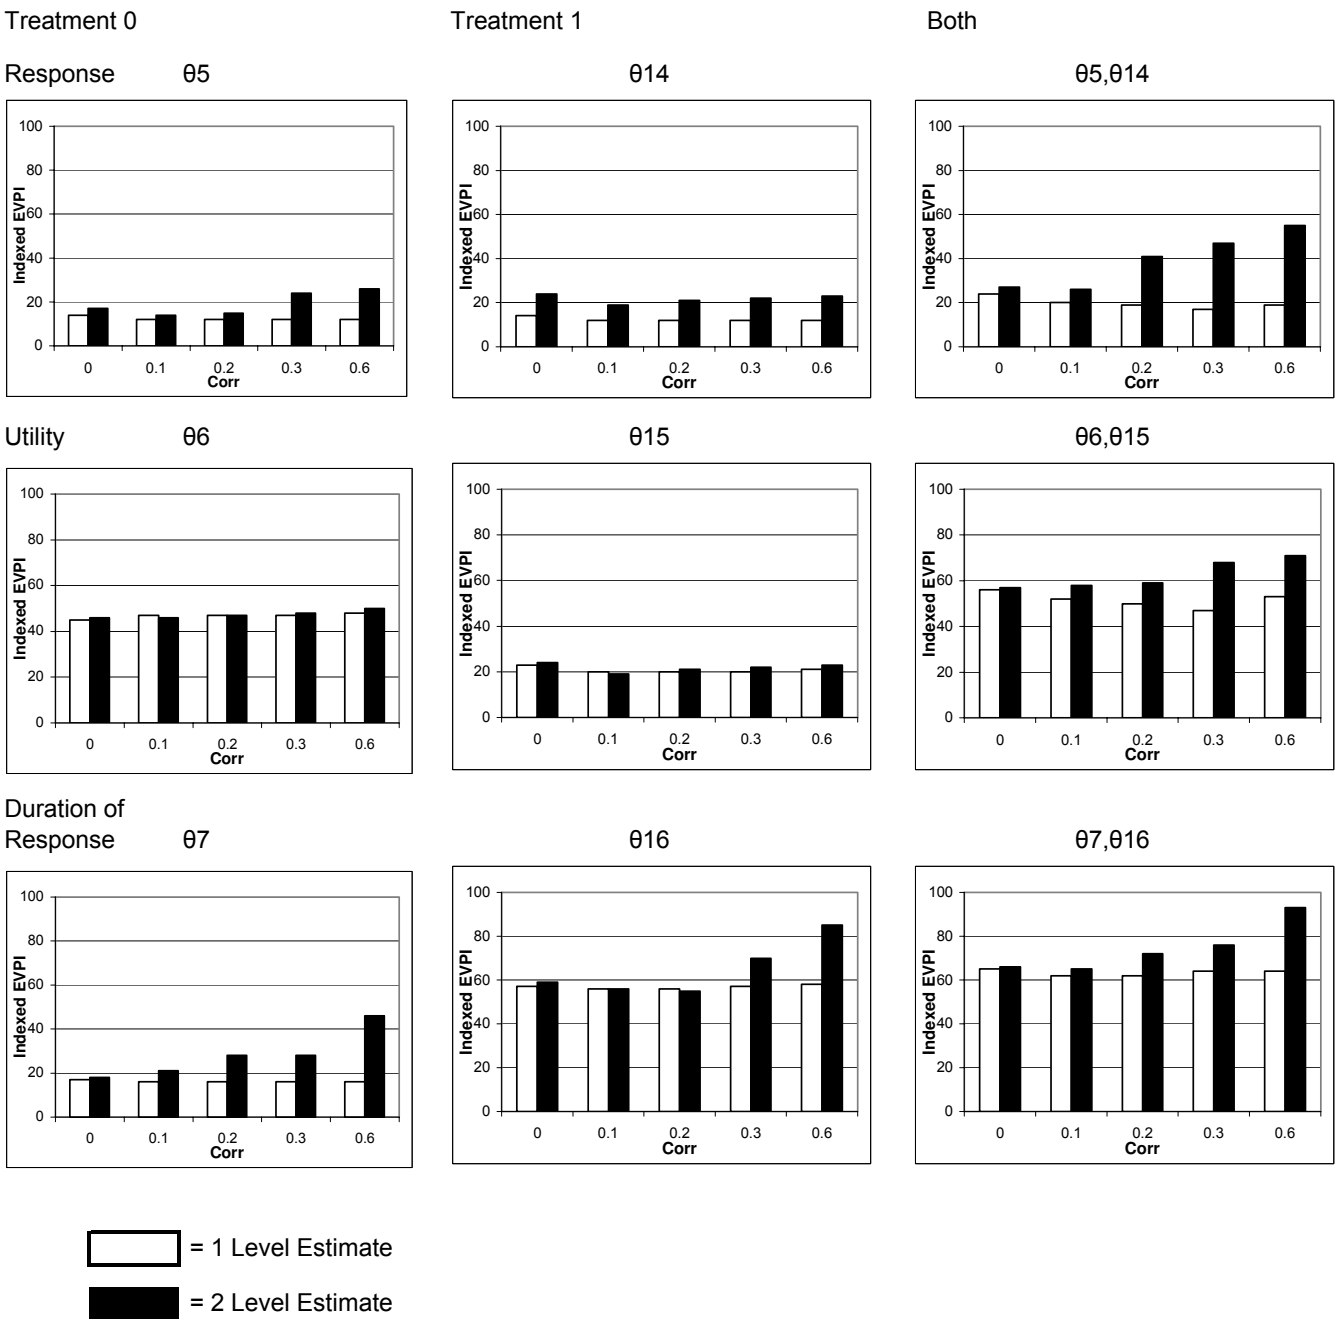

801

802 K=1000 outer and J=1000 inner samples

803 Indexed to overall EVPI per patient = 100

Figure 2: Impact of Increasing Non-Linearity on Inaccuracy of 1 level method to calculate partial EVPI

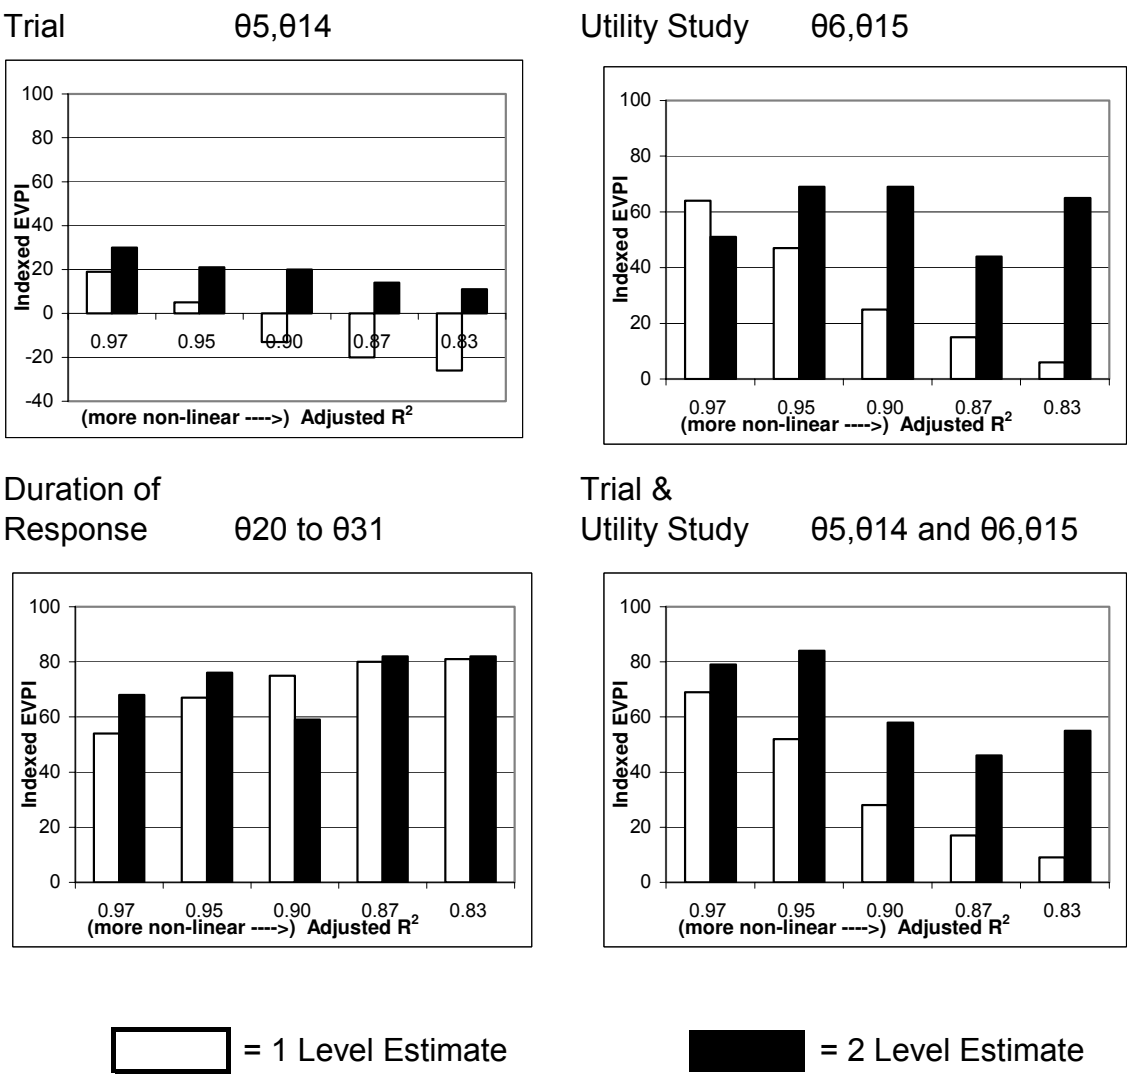

809 **Figure 3: Illustration of stability of Monte-Carlo EVPI Estimates as the inner and outer samples (J and K) are**
810 **extended (Parameters θ5, θ14 Case Study 2 correlation = 0)**
811

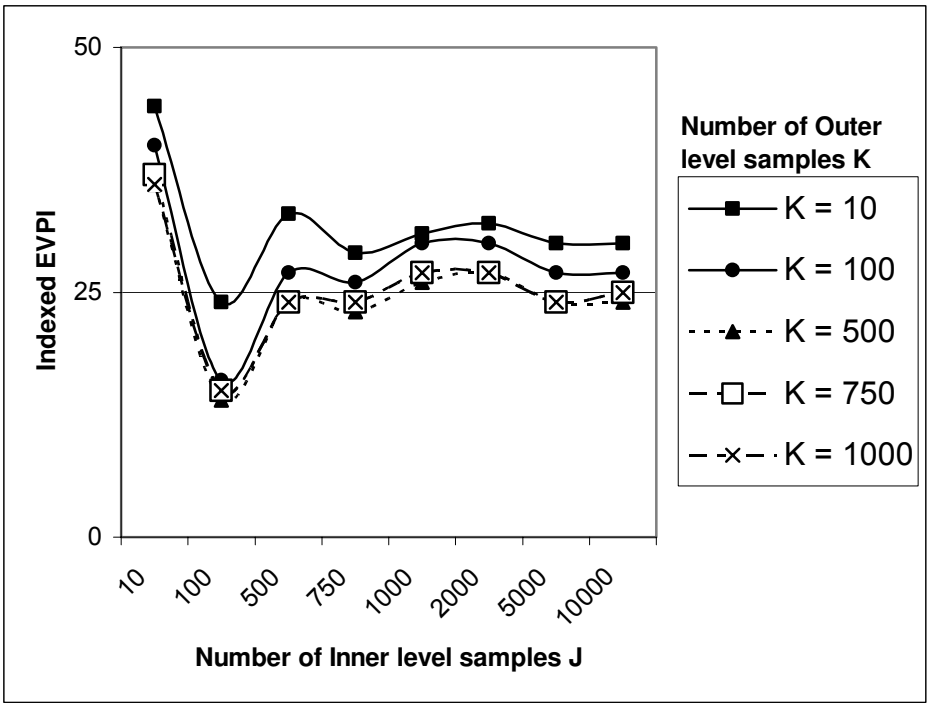

812
813

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

814

Figure 4:

(a) Stability of partial EVPI estimates using relatively small numbers of samples (Case Study 3)

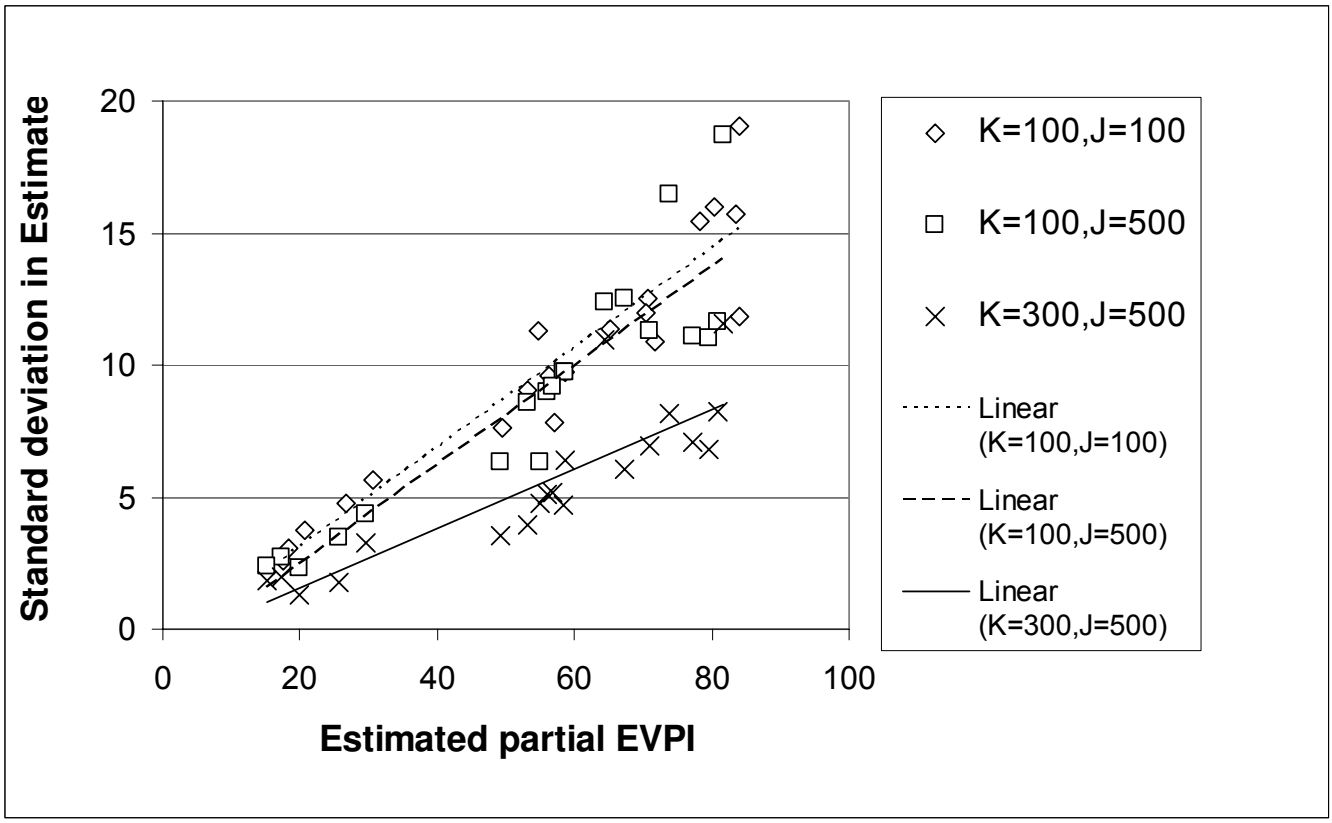

(b) Stability of partial EVPI estimates using Box 1 versus Box 2 Algorithms

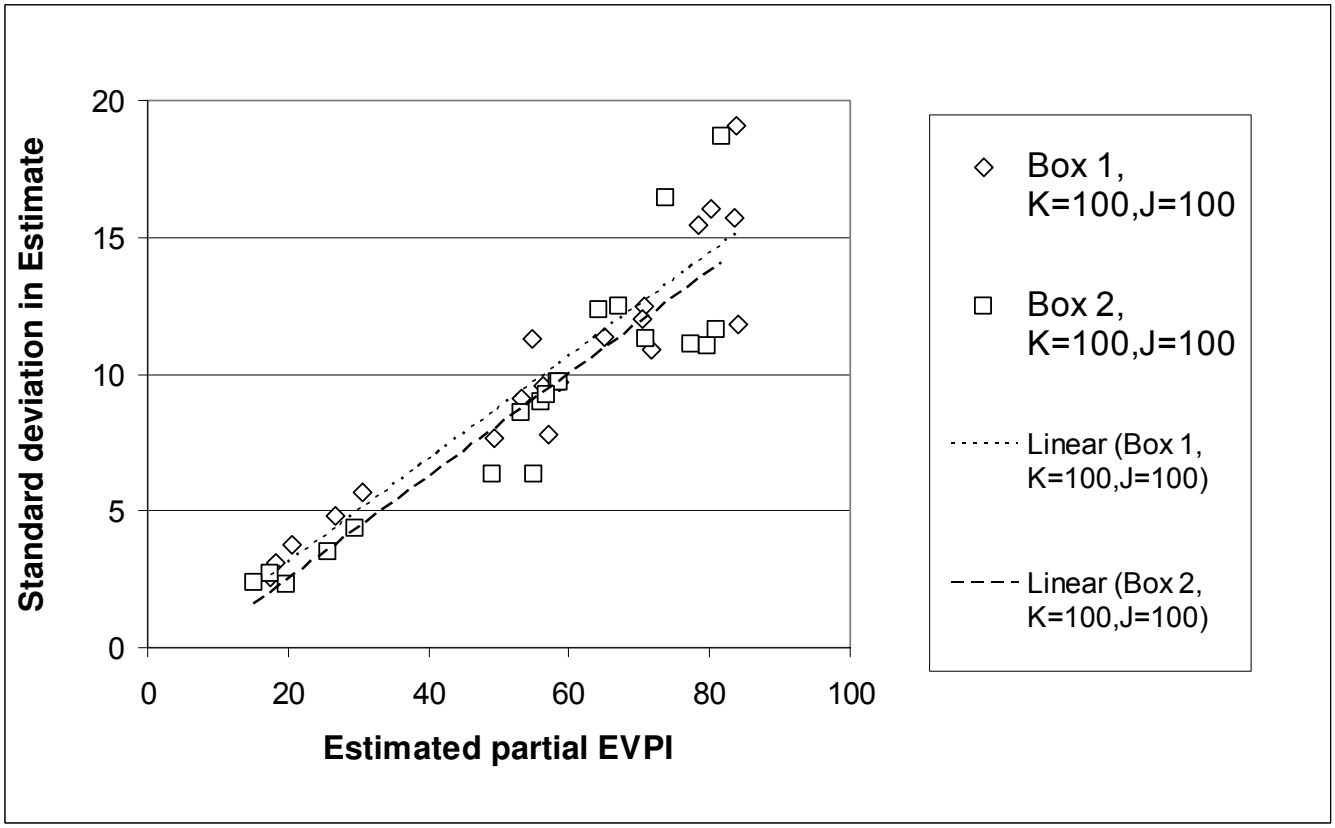

822
823 **Figure 5 'Confidence intervals' for partial EVPI estimates in Case Study 3**

K=100, J=100

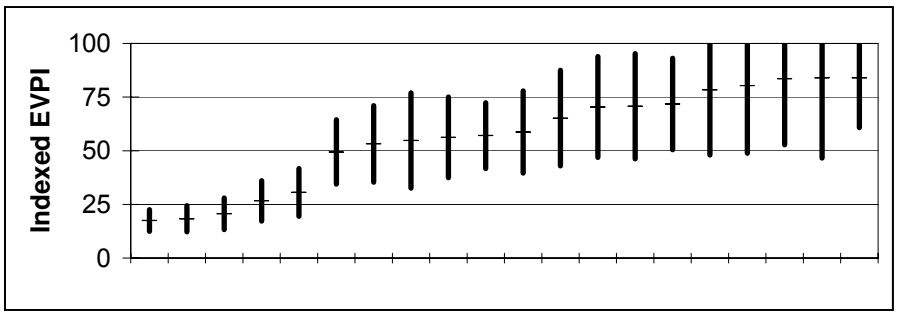

K=100, J= 500

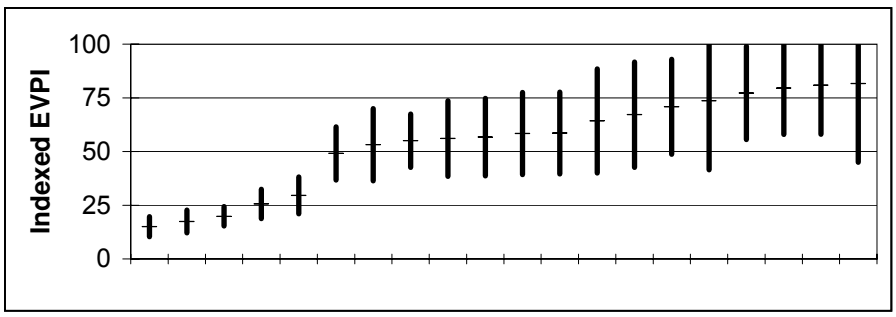

K=300, J=100

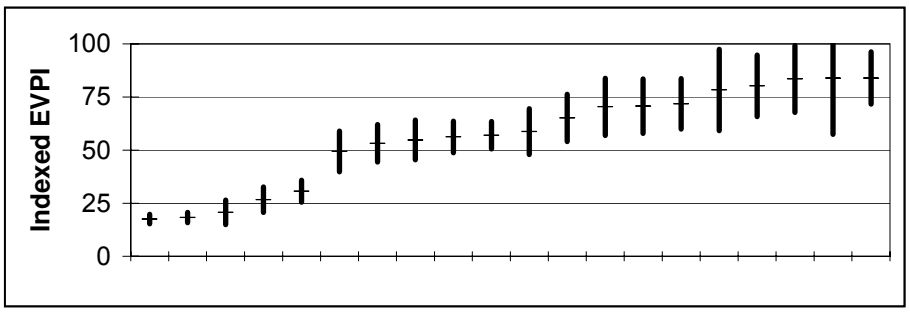

K=300, J=500

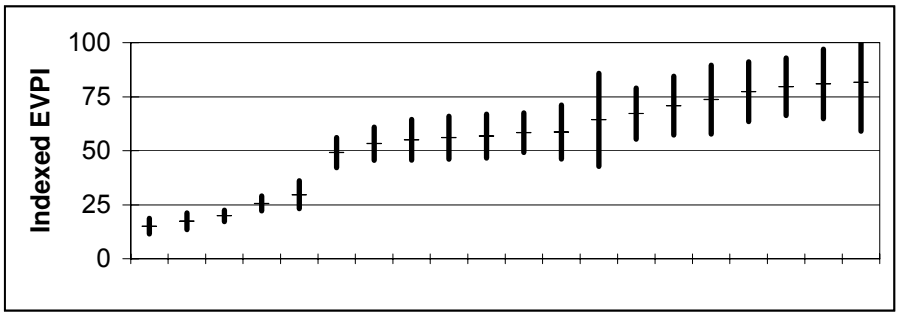

824
825 Standard deviations based on 30 runs.

826 **Figure 6: Comparison of Box 1 and Box 2 Algorithm Noise**
827 **(a) Stability of EVPI estimates using Box 1 versus Box 2– Net benefit Functions with Very Low Correlation**
828

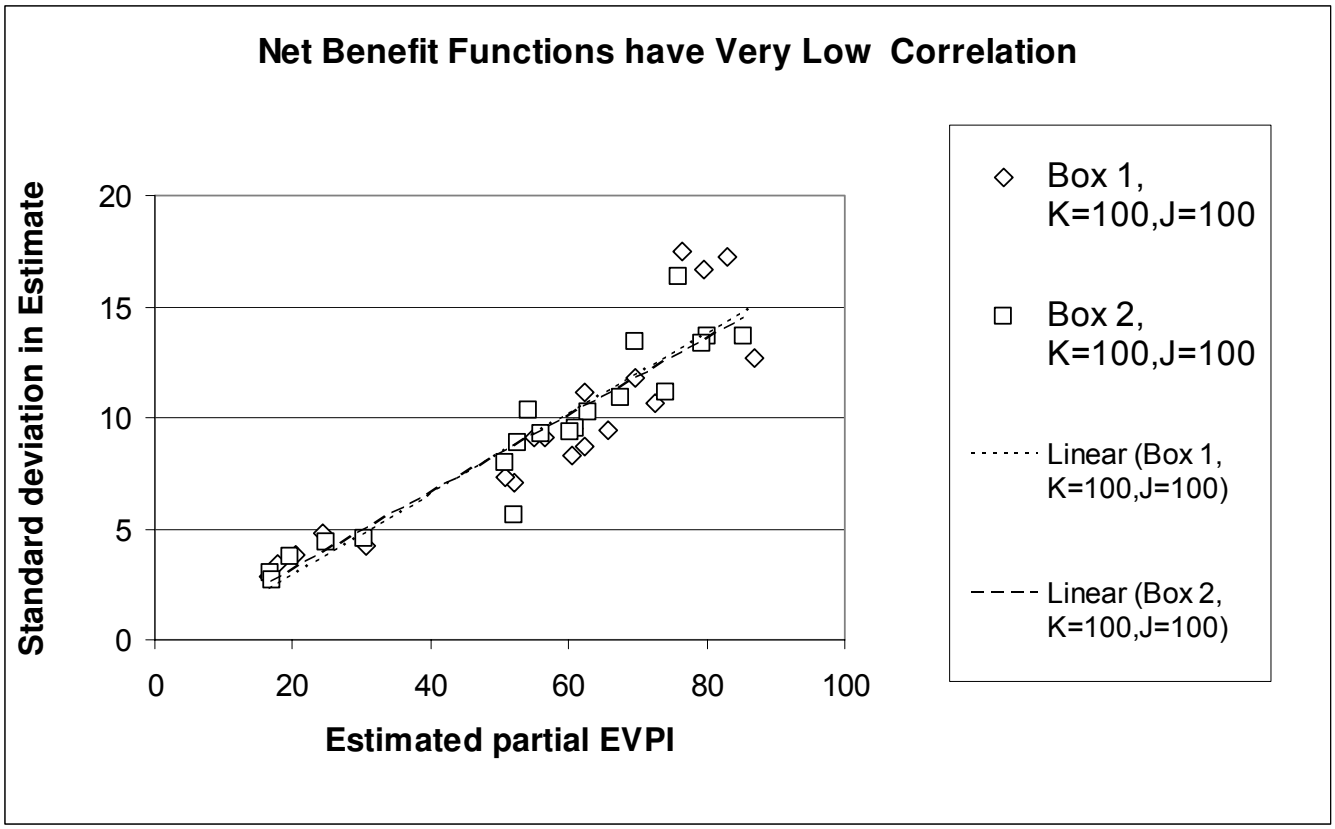

829

830 **(b) Stability of EVPI estimates using Box 1 versus Box 2– Net benefit Functions with High Correlation**
831

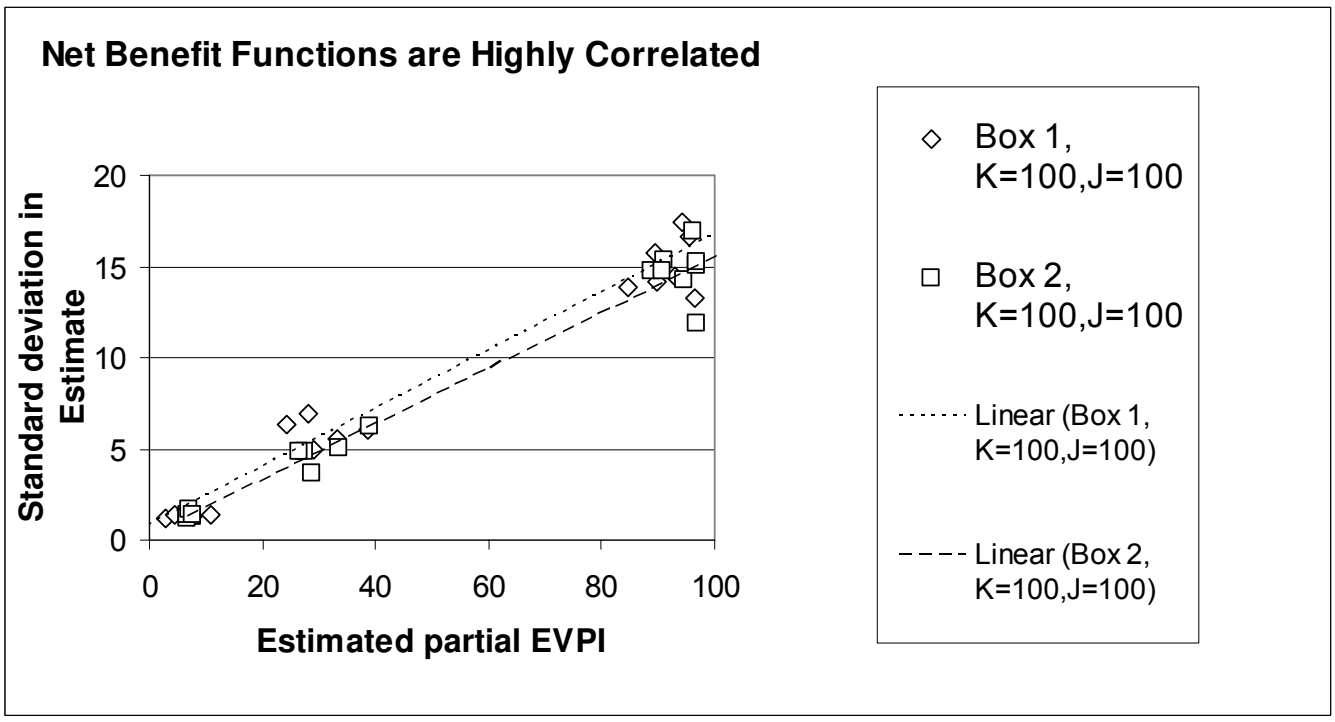

832

Appendix 1: Felli and Hazen MC1 Monte Carlo Procedure for Partial EVPI [4,15]

Model parameters are ξ, net benefit function V, and the decision options as A. Let $E[V|\ \xi, A]$ be the decision maker's expected payoff as a function of ξ and A. The baseline adoption decision is denoted A*. ξ_I is a collection of parameters whose EVPI we wish to calculate and let ξ_I^C be the set of remaining parameters in the problem $\xi = (\xi_I, \xi_I^C)$,. The decision which maximises expected value conditional upon particular values for the parameters of interest ξ_I is denoted $A^*(\xi_I)$. The procedure then is:

MC1: General Monte Carlo Simulation Procedure

1. Repeatedly generate random parameter values $\xi = (\xi_I, \xi_I^C)$
2. For each generated $\xi = (\xi_I, \xi_I^C)$,
    i. Determine $A^*(\xi_I)$ as the decision option A maximizing $E[V|\ \xi_I, A]$.
    ii. Calculate the improvement achieved by using $A^*(\xi_I)$

$$\text{Improvement} = E[V|\ \xi_I, \xi_I^C, A^*(\xi_I)] - E[V|\ \xi_I, \xi_I^C, A^*]$$

End For

3. Estimate EVPI (ξ_I) as the average of the calculated improvement values.

Here it is assumed in Step 2i of the procedure that there is an algebraic expression for the quantity $E[V|\ \xi_I, A] = E_{\xi_I^c}[E[V|\ \xi_I, \xi_I^C, A|\ \xi_I]$.

Appendix 2: Means of maxima of two independent normally distributed variables

Suppose X_1,X_2 are independent normal random variables with parameters μ_1,σ_1 and μ_2,σ_2, respectively. Suppose $\mu_1 > \mu_2$. Then

$$E[\max\{X_1,X_2\}] = E[X_1 + \max\{0,X_2-X_1\}] = \mu_1 + E[\max\{0,X_2-X_1\}] = \mu_1 + E[\max\{0,Y\}]$$

where $Y = X_2 - X_1 \sim \text{normal}(\mu_Y,\sigma)$, with $\mu_Y = \mu_2-\mu_1$, $\sigma^2 = \sigma_1^2 + \sigma_2^2$. We have

$$E[\max\{0,Y\}] = E[\max\{0,\mu + \sigma Z\}] = \sigma E[\max\{0,Z - c\}]$$

where Z is a standard normal variable and $c = -\mu_Y/\sigma$. Then with $\varphi(z)$ the standard normal density and $\Phi(c) = \int_{-\infty}^{c} \phi(z)dz$, the cumulative standard normal distribution function, we have

$$E[\max\{0,Z-c\}] = \int_{-\infty}^{\infty} \max\{0,z-c\}\phi(z)dz = \int_{-\infty}^{c} 0.\phi(z)dz + \int_{c}^{\infty}(z-c)\phi(z)dz$$

$$= \int_{c}^{\infty} z\phi(z)dz - c\int_{c}^{\infty}\phi(z)dz$$

$$= \int_{c}^{\infty} z\phi(z)dz - c[1-\Phi(c)] = \int_{c}^{\infty} z\phi(z)dz - c[1-\Phi(c)]$$

$$= \int_{c}^{\infty} z\frac{1}{\sqrt{2\pi}}e^{-z^2/2}dz - c[1-\Phi(c)]$$

$$= \frac{1}{\sqrt{2\pi}}e^{-c^2/2} - c[1-\Phi(c)]$$

Hence, $E[\max\{X_1,X_2\}] = \mu_1 + \sigma.\left(\frac{1}{\sqrt{2\pi}}e^{-c^2/2} - c[1-\Phi(c)]\right)$

Application to Case Study 1

When NB1, NB2 are independent normally distributed with parameters (μ_1,σ_1) and (μ_2,σ_2), then

$$E[\max\{NB1,NB2\}] = EMAX(\mu_1,\sigma_1,\mu_2,\sigma_2) = \mu + \sigma.\left(\frac{1}{\sqrt{2\pi}}e^{-c^2/2} - c[1-\Phi(c)]\right) \quad (10)$$

where $\Phi(\cdot)$ is the standard normal distribution function, and

$$\mu = \max\{\mu_1,\mu_2\} \qquad \sigma = (\sigma_1^2 + \sigma_2^2)^{1/2} \qquad c = |\mu_2-\mu_1|/\sigma$$

Because θ1 and θ2 are independent normal(1,1) random variables, the Monte-Carlo estimate of EVPI(θ^i) when using K outer, J inner and L samples for each component is given by

$$\text{MCEVPI}_{J,K,L}(\theta^1) = \frac{1}{K}\sum_k \left[\max_d \frac{1}{J}\sum_j NB(d,\theta_{1k},\theta_{2j}) \right] - \max_d \frac{1}{L}\sum_l NB(d,\theta_l)$$

$$= \frac{1}{K}\sum_k \left[\max\left\{ \frac{1}{J}\sum_j 20\theta_{1k}, \frac{1}{J}\sum_j 19.5\theta_{2j} \right\} \right] - \max\left\{ \frac{1}{L}\sum_l 20\theta_{1l}, \frac{1}{L}\sum_l 19.5\theta_{2l} \right\}$$

$$= \frac{1}{K}\sum_k \left[\max\left\{ 20\theta_{1k}, 19.5\bar{\theta}_{2J} \right\} \right] - \max\left\{ 20\bar{\theta}_{1L}, 19.5\bar{\theta}_{2L} \right\}$$

where, $\bar{\theta}_{2J} = \frac{1}{J}\sum_j \theta_{2j} \sim \text{normal}(1, 1/\sqrt{J})$.

Therefore, we can calculate the expected value of a Monte-Carlo estimate as,

$$\text{E}\left[\text{MCEVPI}_{J,K,L}\ \text{EVPI}(\theta^1)\right] = E\left[\max\left\{ 20\theta_{1k}, 19.5\bar{\theta}_{2J} \right\} \right] - E\left[\max\left\{ 20\bar{\theta}_{1L}, 19.5\bar{\theta}_{2L} \right\} \right]$$

$$= \text{EMAX}(20,\ 20,\ 19.5,\ 19.5/\sqrt{\text{J}}) - \text{EMAX}(20,\ 20/\sqrt{\text{L}},\ 19.5,\ 19.5/\sqrt{\text{L}})$$

The true expected value of perfect information on θ^1 is given by

$$\text{EVPI}(\theta^1) = E_{\theta_1}\left[\max_d E\left[NB(d,\theta) \middle| \theta_1 \right] \right] - \max_d E[NB(d,\theta)]$$

$$= E_{\theta_1}\left[\max\left\{ 20\theta_1, 19.5 \right\} \right] - \max\left\{ 20, 19.5 \right\}$$

$$= \text{EMAX}(20,\ 20,\ 19.5,\ 0) - 20.$$

Then Bias(J,L) = $\text{E}[\text{MCEVPI}_{J,K,L}(\theta^1)] - \text{EVPI}(\theta^1)$

$$= \{\text{EMAX}(20,\ 20,\ 19.5,\ 19.5/\sqrt{\text{J}})$$

$$- \text{EMAX}(20,\ 20/\sqrt{\text{L}},\ 19.5,\ 19.5/\sqrt{\text{L}})\}$$

$$- \{\text{EMAX}(20,\ 20,\ 19.5,\ 0) - 20\}.$$

886 **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 D, Buxton MJ, O'Brien BJ. Measures of importance for economic analysis based on decision modeling. Journal of Clinical Epidemiology 2003:56;989–997.

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

5 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).

6 Felli JC, Hazen GB. A Bayesian approach to sensitivity analysis. Health Economics 1999;8:263-68.

7 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.

8 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.

9 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.

10. Brennan, A., Chilcott, J., Kharroubi, S. A. and O'Hagan, A.. A two level Monte Carlo approach to calculating expected value of perfect information:- Resolution of the uncertainty in methods. Poster at the 24th Annual Meeting of the Society for Medical Decision Making, October 23rd, 2002, Baltimore. 2002. Available at http://www.shef.ac.uk/content/1/c6/03/85/60/EVPI.ppt

11 Raiffa H. Decision analysis: Introductory lectures on choice under uncertainty. 1968. Addison-Wesley, Reading, Mass.

12 Spiegelhalter DJ, Myles JP, Jones DR, Abrams KR. Bayesian Methods in Health Technology Assessment: A Review. *Health Technology Assessment* 2000;**4**.

13 Claxton K,.Posnett J. An economic approach to clinical trial design and research priority-setting. Health Economics 1996;5:513-24.

14 Lapin, L. L. Quantitative methods for business decisions. 1994. Orlando, Florida, Dryden Press.

15 Felli JC, Hazen GB. Correction: Sensitivity Analysis and the Expected Value of Perfect Information, Medical decision Making, 2003;23:97

16 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.

17 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.

18 Payne N, Chilcott JB, McGoogan E. Liquid-based cytology in cervical screening. Health Technology Assessment 2000.

19 Brennan, A., Chilcott, J., Kharroubi, S. A. and O'Hagan, A. (2002). A two level Monte Carlo approach to calculating expected value of perfect information:- Resolution of the uncertainty in methods. Available at http://www.shef.ac.uk/content/1/c6/02/96/05/Brennan.doc

20 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.

21 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.

22 Tappenden P, Chilcott J., Eggington, S., Oakley, J., McCabe C. Methods for expected value of information analysis in complex health economic models: developments on the health economics of interferon-β and glatiramer acetate for multiple sclerosis. Health Technol Assess 2004;8(27).

23 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.

24 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. Available at http://www.shef.ac.uk/content/1/c6/03/85/60/EVSI.ppt

25 Brennan, A., Chilcott, J., Kharroubi, S. A. and O'Hagan, A. A two level Monte Carlo approach to calculating expected value of sample information:- How to value a research design. Discussion paper, 2002. . Available at http://www.shef.ac.uk/content/1/c6/02/96/29/EVSI.doc

26 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.

27 Phelps CE, Mushlin AI. Focusing technology-assessment using medical decision-theory. Med Decis Making. 1988;8:279-289.

28 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.

29 Yokota F, Thompson K. Value of Information Literature Analysis: A Review of Applications in Health Risk Management, Medical Decision Making, 2004;24:287-98.

30 Hammitt JK, Shlyakhter AI. The Expected Value of Information and the Probability of Surprise, Risk Analysis, 1999;19:135-152.

31 Gould J. Risk,stochastic preference and the value of information, Journal of Economic Theory, 1974;8: 64–84.

32 Hilton R. Determinants of information value, Management Science, 1981;27: 57–64.

33 Howard RA. Decision Analysis: Practice and Promise, Management Science, 1988;34:679–695.

34 Hammitt JK. Can More Information Increase Uncertainty? Chance, 1995;8:15–17.

35 Eric W. Weisstein. "Quadrature." From MathWorld--A Wolfram Web Resource. http://mathworld.wolfram.com/Quadrature.html

36 Briggs AH, Gray A. Handling uncertainty when performing economic evaluation of healthcare interventions. Health Technology Assessment 1999;3.

37 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.

38 Oakley J, Tappenden P., Chilcott J., Brennan A. Sample Sizes for Monte Carlo Partial EVPI Calculations. Research Report Department of Probability and Statistics, University of Sheffield, 2006.

39 Venables W. N., Smith D. M. and the R Development Core Team. An Introduction to R Notes on R: A Programming Environment for Data Analysis and Graphics Version 2.2.1 (2005-12-20). Available at http://www.r-project.org/

40 http://www.shef.ac.uk/chebs/news/news1.html

41 An Introduction to R version 2.2.1 (2005-12-20). ISBN 3-900051-12-7 . http://cran.r-project.org/doc/manuals/R-intro.html

42 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

43 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.

44 Tappenden P, Chilcott J., Eggington, S., Oakley, J., McCabe C. Methods for expected value of information analysis in complex health economic models: developments on the health economics of interferon-β and glatiramer acetate for multiple sclerosis. Health Technol Assess 2004;8(27).

45 Oakley J. Value of information for complex cost-effectiveness models. Research Report No. 533/02 Department of Probability and Statistics, University of Sheffield, 2002.

46 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.