

Decision-theoretic Sensitivity Analysis using Value of Information

Jeremy Oakley

Department of Probability and Statistics, University of Sheffield, The Hicks Building,
Hounsfield Road, Sheffield, S3 7RH, UK;
E-mail: j.oakley@sheffield.ac.uk

1. INTRODUCTION

When a computer model is to be used to guide a decision, it is important for the decision-maker to acknowledge and investigate the uncertainty in the model. Typically, there will be uncertainty surrounding the true values of the input parameters in the model that should be used for the decision problem in question, and this then induces uncertainty in the output of the model. If the decision-maker considers their probability distribution for each unknown input in the model, they can then derive their probability distribution for the model output. The combination of their output distribution and an appropriate utility/loss function can then guide their decision.

In some cases, it may be possible to learn more about some or all of the uncertain input parameters before a final decision is made. In this case, it is then desirable to assess the importance of each uncertain input parameter in the model. Quantifying parameter importance is known as global or probabilistic sensitivity analysis. A measure of parameter importance that has been advocated previously is the variance-based measure (see 1). Variance-based measures consider the contribution of each uncertain input parameter to the variance of the model output. However, uncertainty about the model output as characterised by its variance is not necessarily equivalent to uncertainty about the optimum decision. Consequently, using variance-based measures to establish parameter importance in decision problems can in some cases produce misleading results, even as far as ranking the parameters in the wrong order of importance.

An alternative measure of parameter importance can be derived within the framework of utility theory. The idea is to determine whether different values of a particular input parameter lead to different optimum decisions, and if so, how much the expected utility/loss under alternative optimum decisions varies. Specifically, the expected utility of learning the true numerical value of an uncertain input parameter before the decision is made can be calculated. This quantity is known as the partial expected value of perfect information (partial EVPI), and precisely quantifies the importance of an uncertain input variable. When the specific purpose of the model is to guide a decision within a clearly defined utility/loss structure, we advocate the partial EVPI as the single correct measure of an uncertain parameter's importance.

Further author information:

J.O.: Telephone: +44 (0)114 222 3853

In most practical situations, the decision-maker will not be able to learn the true value of an uncertain input parameter precisely, even if they desire to do so. The more likely possibility is that they may have the option of collecting more data to reduce their uncertainty about the unknown parameter. The expected value of perfect information framework can be extended to consider the expected value of collecting this data before making the decision; this is known as the expected value of sample information (EVSI). EVSI measures can then be used for deriving optimal sample sizes.

Both partial EVPIs and EVSIs can be computed using Monte Carlo methods. Unfortunately, to obtain these measures accurately, very large numbers of model evaluations are needed, potentially millions. For computationally expensive computer models, evaluating these measures may then require prohibitively lengthy computing times. However, in many cases it will be possible to exploit a feature of the computer model to dramatically speed up the computation; the function mapping inputs to output is often a smooth function. If the model is run at a particular set of input values and the output is observed, we will then also have information about the likely output at neighbouring sets of input parameter values.

When the time needed for a single run of the model is non-trivial, it can be highly advantageous to construct an *emulator*, a statistical approximation to the original computer model based on a fairly small number of different runs of that model. The emulator can then be used to give a fast approximation to the computer model regardless of the complexity of the model. An emulator is a regression model, and any regression technique can be employed. Our preferred option is the Gaussian process model. The Gaussian process emulator is a non-parametric approach that with the exception of continuity, makes no other assumptions about the functional form of the computer model. Gaussian processes have been used successfully before for efficient computation in other areas of sensitivity and uncertainty analysis. It will be demonstrated that the Gaussian process approach is of the order of 1000 times more efficient than Monte Carlo methods in terms of numbers of model runs, for computing partial EVPIs and EVSIs.

An application is given in the field of health economics. Economic models are used to estimate the cost-effectiveness of new treatments under consideration. A decision-maker will use the output of the model to help decide whether or not to approve the new treatment. There is always uncertainty regarding the values of the input parameters needed for the model; for example, it will not be known exactly how effective the treatment is, or what financial resources the patients on the treatment will use. There will be particular interest in conducting a probabilistic sensitivity analysis when using the model. It will often be possible to obtain more data regarding some of the model parameters, and hence reduce input uncertainty. Additionally, a certain class of models, known as patient simulation models, require an extensive simulation to produce the output for a single choice of input parameters. These models can be very computationally expensive, requiring in some cases in excess of an hour per run. In this scenario, emulator methods are essential for computation of EVPIs and EVSIs.

2. THE EXPECTED VALUE OF PERFECT INFORMATION

We now give a decision-theoretic measure of importance of an uncertain input variable. This measure is based on a standard result from decision theory (see for example 2), and was advocated by (3) and (4).

Suppose a decision-maker has to chose one decision d from a set of possible decisions \mathcal{D} . The decision-maker has a computer model to aid their decision, denoted by $y = f(\mathbf{x})$, where y is the model output and \mathbf{x} are the model inputs. In the decision problem at hand, we suppose that there are ‘true’, uncertain values of the inputs that should be used in the model, and these uncertain values are denoted by \mathbf{X} with distribution G . We then suppose that the utility of a decision d will be some function of the true output of the model, $f(\mathbf{X})$, and is denoted by $U\{d, f(\mathbf{X})\}$.

The decision maker then chooses the decision d to maximise their expected utility $E_{\mathbf{X}}[U\{d, f(\mathbf{X})\}]$. We can now define the expected utility of the optimum decision to be U^* , where

$$U^* = \max_d E_{\mathbf{X}}\{U\{d, f(\mathbf{X})\}\}. \quad (1)$$

Now suppose that the decision maker decides that they will learn the value of \mathbf{X} before making their decision. Once they have learnt \mathbf{X} , their utility is then

$$\max_d U\{d, f(\mathbf{X})\}, \quad (2)$$

and so their expected utility of learning \mathbf{X} (i.e., before they find out what \mathbf{X} actually is) is

$$E_{\mathbf{X}}\{\max_d U\{d, f(\mathbf{X})\}\}. \quad (3)$$

The expected value of perfect information (EVPI) is then defined as the expected gain in utility:

$$E_{\mathbf{X}}\{\max_d U\{d, f(\mathbf{X})\}\} - \max_d E_{\mathbf{X}}\{U\{d, f(\mathbf{X})\}\}. \quad (4)$$

Now denote one of the uncertain input variables to be X_i . The same argument can be applied to derive the expected value of learning X_i before making the decision. Given X_i , we are still uncertain about the remaining input variables, \mathbf{X}_{-i} , and so we would choose the decision to maximise $E_{\mathbf{X}_{-i}|X_i}\{U\{d, f(\mathbf{X})\}\}$. The expected utility of learning X_i is then

$$E_{X_i} \left[\max_d E_{\mathbf{X}_{-i}|X_i}\{U\{d, f(\mathbf{X})\}\} \right], \quad (5)$$

and so the expected gain in utility, the partial EVPI of X_i is

$$E_{X_i} \left[\max_d E_{\mathbf{X}_{-i}|X_i}\{U\{d, f(\mathbf{X})\}\} \right] - \max_d E_{\mathbf{X}}\{U\{d, f(\mathbf{X})\}\}. \quad (6)$$

Here, we advocate the partial EVPI of X_i as a measure of importance of that variable in the model.

2.1. Computation

Partial EVPIs can be computed by Monte Carlo methods, but this can be computationally intensive and in some cases infeasible when a single run of the model takes a non-trivial amount of computing time. When the model is computationally expensive, a common approach is to use an *emulator*, a fast statistical approximation to the computer model based on regression (see 5; 6). This can be considerably more efficient than Monte Carlo when the output of the computer model is a smooth function of inputs. Full computational details for partial EVPI estimates using (Gaussian process) emulators are given in (7)

3. EXAMPLE: HEALTH ECONOMIC MODELLING AND THE GERD MODEL

One application area in which partial EVPIs are currently used is health economics. In health economics, the interest is in allocating health care resources as effectively as possible. The decision problem is to choose which out of set of competing treatments for an illness is the most cost-effective. Cost-effectiveness of a treatment is described with a single (financial) measure known as the *net benefit* of the treatment, and net benefits are often predicted using computer models. (This is because clinical trials typically only record information on clinical effectiveness). The models invariably require specification of parameters that are uncertain, and so there is interest in investigating the consequences of this input uncertainty. The utility of choosing a particular treatment is then given by the net-benefit of the treatment, and so it is possible to measure the importance of each input using partial EVPIs.

We give the example used in (7) that also illustrate the efficiency of the emulator approach in the computation of partial EVPIs. The model compares treatment strategies for gastroesophageal reflux disease. In this example, we suppose that a decision has to be made regarding the adoption of one of three treatment strategies:

1. Acute treatment with proton pump inhibitors (PPIs) for 8 weeks, then continuous maintenance treatment with PPIs at the same dose.
2. Acute treatment with PPIs for 8 weeks, then continuous maintenance treatment with hydrogen receptor antagonists (H2RAs).
3. Acute treatment with proton pump inhibitors PPIs for 8 weeks, then continuous maintenance treatment with PPIs at the a lower dose.

The model was presented in (8). In the scenario that we are considering, there are twenty-three uncertain inputs, relating to quantities such as probabilities of healing and recurrence of the symptoms with each treatment, and resources used by patients such as number of visits to a general practitioner. Distributions for all the uncertain inputs are described in (9). The output of the model can be converted into a utility for each treatment.

Using 600 runs of the model, we estimate the partial EVPI of each patient. The GERD model is computationally cheap, so we can determine the true partial EVPIs based on

massive Monte Carlo samples (several hundred million in this case). Although there is some inaccuracy in the estimates, we have identified all the influential inputs in the model, to within what we believe would be an acceptable order of magnitude. For comparison, we also estimate the partial EVPIs using a combination of Simpson’s rule and Monte Carlo as described in section 2.

We give the actual values of the estimates and true values of the partial EVPIs for the six most important variables in table 1.

uncertain input parameter	true partial EVPI	Gaussian process estimate	Simpson/MC estimate estimates
hazard for healing on PPIs	1.286	1.194	3.465
no. of symptom weeks after surgery	2.271	2.500	4.229
Recurrence probability on PPIs (6-12) months	4.905	4.579	5.507
Recurrence probability on H2RAs (0-6) months	21.221	20.908	23.417
Recurrence probability on H2RAs (6-12) months	2.652	2.666	2.958
Recurrence probability on low dose PPIs (6-12) months	3.473	3.378	3.846

Table 1. True values, Gaussian process estimates and Simpson/Monte Carlo estimates of the partial EVPIs of the six most influential input variables. The Gaussian process estimates are based on 600 model runs, and the Simpson/Monte Carlo estimates are based on 410200 model runs.

These partial EVPIs can then be interpreted as (financial) values of learning the value of the corresponding parameter before choosing which treatment to use for the patient population. The figure represents dollars per patient, and so needs to be multiplied by the size of the patient population to give a final value.

References

- [1] A. Saltelli, K. Chan, and M. Scott, editors. *Sensitivity Analysis*. Wiley, New York, 2000.
- [2] A. O’Hagan. *Kendall’s Advanced Theory of Statistics, Volume 2B, Bayesian Inference*. Edward Arnold, London, 1994.
- [3] J. C. Felli and G. B. Hazen. Sensitivity analysis and the expected value of perfect information. *Medical Decision Making*, 18:95–109, 1998.
- [4] K. Claxton. Bayesian approaches to the value of information: implications for the regulation of new health care technologies. *Health Economics*, 8:269–274, 1999.

- [5] J. Sacks, W. J. Welch, T. J. Mitchell, and H. P. Wynn. Design and analysis of computer experiments. *Statist. Sci.*, 4:409–435, 1989.
- [6] A. O’Hagan, M.C. Kennedy, and J. E. Oakley. Uncertainty analysis and other inference tools for complex computer codes (with discussion). In J. M. Bernardo, J. O. Berger, A. P. Dawid, and A. F. M. Smith, editors, *Bayesian Statistics 6*, pages 503–524. Oxford: University Press, 1999.
- [7] J. E. Oakley. Sensitivity analysis for computationally expensive economic models. Technical report, Department of Probability and Statistics, University of Sheffield, 2003.
- [8] R. Goeree, B. O’Brien, R. Hunt, G. Blackhouse, A. Willan, and J. Watson. Economic evaluation of long term management strategies for erosive oesophagitis. *PharmacoEcon.*, 16:679–697, 1999.
- [9] A. H. Briggs, R. Goeree, G. Blackhouse, and B. J. O’Brien. Probabilistic analysis of cost-effectiveness models: Choosing between treatment strategies for gastroesophageal reflux disease. *Med Decis Making*, 22:290–308, 2002.