CONSIDERING RISK ATTITUDE IN A VALUE OF INFORMATION PROBLEM

Hsiao, Chuck; Malak, Richard
Texas A&M University, United States of America

Abstract

In many decisions, one of the available alternatives is to gather more information about the situation at hand, which incurs a cost but leads to a more informed and thus improved decision. Thus, the decision problem is two-fold: first, whether or not to gather additional information, and second, which course of action or design to select based on the available information as a result of the first decision. Such problems are Value of Information problems, which seek to quantify the value of the potential information to guide the decision maker on whether or not it is worthwhile.
However, approaches to Value of Information problems typically implicitly assume that the decision maker is risk neutral, in the formulation of the problem. This paper considers how the inclusion of risk attitude affects a decision maker's decision about whether or not additional information is worthwhile. This leads to a more accurate model of decision problems typically facing decision makers. It discusses some of the mathematical complexities and illustrates the problem with an engineering example.

Keywords: Decision making, Value of Information, Risk attitude

Contact:
Chuck Hsiao
Texas A&M University
Mechanical Engineering
United States of America
chuckhsiao@tamu.edu

Please cite this paper as:
Surnames, Initials: *Title of paper*. In: Proceedings of the 20th International Conference on Engineering Design (ICED15), Vol. nn: Title of Volume, Milan, Italy, 27.-30.07.2015

1 INTRODUCTION

As society pursues more and more complex projects, it becomes more and more important to ensure that project resources are spent efficiently. Although much research has investigated decisions about designed artefacts or systems, relatively little study has been devoted to decisions about the design process itself. Of particular interest to designers are process decisions about whether to gather more information about an uncertainty in the design problem. This activity can reduce uncertainty and result in better-informed decisions about the design, but the costs and time required to gather this information are not necessarily worthwhile when one considers the benefits of doing so. The ability to model information gathering decisions more accurately can be critical for ensuring that limited design resources are well-spent.

When an engineer is making a decision under uncertainty, one of the process alternatives that are typically available is to expend resources to reduce that uncertainty, in order to make a more informed decision about a design. Examples include building prototypes, making sophisticated computer simulations, and conducting clinical trials or flight testing. These alternatives only indirectly affect the design decision that engineers are making about the project artefact, but their importance increases as we continue to pursue projects of greater sophistication.

The study of how additional information can affect a decision typically falls under the umbrella term of Value of Information (Howard 1966). The fundamental approach is to quantify the value of a quantity of information by comparing its benefits, in terms of the decision maker being able to make a more informed decision, with the cost of acquiring the information.

Typically, in Value of Information approaches such as Expected Value of Perfect Information, the decision maker is assumed to be risk neutral when making decisions under uncertainty. Although this simplifies the analysis, engineers and other decision makers in a project can instead be risk averse or even risk prone; some empirical studies show that managers tend to be risk averse instead of risk neutral (Tull and Hawkins 1976). Modelling decisions under risk neutrality, rather than taking the risk attitude of the decision maker into account in formulating a decision, can lead to sub-optimal decisions because the decision model – and thus the model's recommended optimal decision – does not accurately reflect the decision maker's preferences.

From a philosophical standpoint, both Value of Information and risk attitude are concerned with how the decision maker should act under uncertainty. Value of Information is concerned with quantifying how acquiring information can reduce uncertainty in the outcomes. On the other hand, risk attitude is concerned with how the decision maker prefers different outcomes in an uncertain situation. Both are relevant in a decision problem with uncertainty, and considering one without considering the other can lead to an incomplete understanding of the problem.

The goal of this paper is to explore some of the issues involved when risk attitude is considered in a Value of Information engineering decision problem. The paper will demonstrate how considering risk attitude may lead to a different value of information, and therefore different optimal process and subsequent design decisions, compared with under a risk-neutral risk attitude. Thus, risk attitude should be considered when formulating the decision process in engineering problems involving the gathering of information.

Section 2 discusses some of the literature that is relevant to this work. Section 3 investigates the mathematical formulation of a decision involving information from a decision analysis standpoint, and illustrates some of the difficulties involved when considering risk attitude. Section 4 uses an engineering example of a manufacturer deciding on the number of production and inspection machines in a TFT-LCD plant to demonstrate the effect of risk attitude on the optimal decision. Section 5 discusses some of the issues involved when considering risk attitude with Value of Information and points to further topics of research. Section 6 concludes this paper.

2 RELATED LITERATURE

Using expected value to quantify the effects of information has been studied since the 1960s. Howard's (1966) seminal work described the approach and gave some fundamental results. However, much of the literature has implicitly assumed risk neutrality. For example, Howard formulated the value of information as the difference between the expected profit as a result of clairvoyance about the outcomes and the expected profit without clairvoyance. As discussed in Section 3, this is only generally true under risk neutrality. Many other works use this same formulation to investigate issues with value of information (Bradley and Agogino 1994, Mehrez and Stulman 1982, Eckermann et al. 2010).

More recently, the issue of incorporating risk attitude with Value of Information approaches has been recognized and studied. Mehrez (1985) gave bounds on the value of information for a risk-averse decision maker compared with a risk-neutral decision maker. Nadiminti et al (1996) considered the value of information based on various forms of payment in a risk-averse context. Eeckhoudt and Godfroid (2000) examined mathematically why higher risk aversion does not always lead to a higher value of information. Bickel (2008) investigated the relationship between perfect and imperfect information in a two-alternative problem when risk attitude is considered. Though they gave important insights on risk attitude and Value of Information, these works all considered a two-alternative decision problem, where the objective is to consider whether or not to gather information in a go/no-go, accept/reject, invest/don't invest etc. type of problem, rather than the more general problem given here, where there are a large number of possible alternatives. Additionally, this paper applies the analysis to a manufacturing context, which is a more sophisticated model of the design decision than is typically studied.

For more general considerations of the Value of Information problem, Gould (1974) found that, contrary to intuition, an increase of the uncertainty in a decision maker's prior beliefs do not necessarily entail an increase in the value of information, nor does an increase in the number of uncertain parameters. Hilton (1981) showed that increasing the number of alternatives available to the decision maker does not necessarily monotonically increase the value of information. Miller (1975) showed that in a sequential decision, the value of gaining information about individual uncertain parameters tend to increase. Samson et al. (1989) showed that the value of information about independent sources of uncertainty is generally not additive, i.e. the value of information about one source of uncertainty and another source of uncertainty may be greater than, less than, or equal to the value of information about both sources together. Although these works do not deal specifically with risk attitude, they demonstrate some important considerations when formulating a problem involving the effects of information, in the context of the relationship between uncertainty and the value of information.

3 MATHEMATICAL FORMULATION

Howard (1966) described (perfect) value of information as "if a perfect clairvoyant appeared and offered to eliminate one [or more] of the uncertainties in the problem, we would be willing to offer him a financial consideration. The question is how large should this financial consideration be." The value of the information is tied to how much a decision maker would be willing to pay such that it increases his expected value over not acquiring the information. However, Howard then gave the mathematical form as the difference between the expected profit if the information were obtained and the expected profit if the information were not obtained. Mathematically, this is denoted as:

$$\mathrm{EVI}_\theta = \mathrm{E}_\theta\left[\max_i\{\mathrm{E}_y[u(x^{(i)}, y)|\theta]\}\right] - \max_i\{\mathrm{E}_y[u(x^{(i)}, y)]\} \tag{1}$$

where EVI_θ is the Expected Value of Information

θ is the information received from some source of information

i is an index of the set of possible design parameters $x^{(i)}$, to delineate different possible sets

$x^{(i)}$ is the ith set of design parameters that can be chosen

y is the set of uncertain parameters

$u(\cdot)$ is a value or utility function expressing the decision maker's preferences

$E_y[g(y)|\theta]$ is the expected value of $g(y)$ taken over y conditioned on θ, that is, given that information θ is known

In Equation (1), for the first term, $\mathrm{E}_y[u(x^{(i)}, y)|\theta]$ is the expected profit given information θ, with the expectation taken over the possible uncertain parameter y values. Then, $\max_i\{\cdot\}$ is to maximize this expected value by varying design parameters, that is, across possible sets of design parameter values (indexed by i). Finally, $\mathrm{E}_\theta[\cdot]$ is to take the expectation across all possible values of information θ that may be received. Thus, the first term expresses mathematically the expected profit if the information were obtained. The second term expresses the current decision without gathering information, namely, finding the set of design parameters $x^{(i)}$ that maximizes the expected value across the uncertain y.

For a decision maker using EVI to help make a decision, it is implied that if the actual cost of information is greater than this value, then the information should not be acquired, since it is not worth the expected benefit from being able to make a better (more informed) decision. Similarly, if the actual cost is less than this value, then the information should be acquired. Additionally, it is implied that after he makes this information decision, he then selects the ith set of design parameters that maximizes his expected value.

This formulation is suitable when the decision maker is risk neutral. However, if the decision maker is either risk averse or risk prone, then this formulation does not suffice. A decision maker is risk averse if he prefers the expected consequence of a lottery to that lottery, and risk prone if he prefers a lottery to its expected consequence (Keeney and Raiffa 1993). Simply put, in an uncertain situation, a risk-averse decision maker prefers to avoid the potential downsides of an alternative with uncertain outcomes, while a risk-prone decision maker prefers the potential upsides.

To see why the above formulation is not suitable if a decision maker were not risk neutral, consider the actual goal of the decision maker. He is deciding between whether or not he should gather the information, based on whether it will increase his expected value. That is, he is selecting:

$$\text{Maximize}\left(\mathrm{E}_\theta\left[\max_i\{\mathrm{E}_y[u(x^{(i)},y,c_\theta)|\theta]\}\right],\ \max_i\{\mathrm{E}_y[u(x^{(i)},y)]\}\right) \tag{2}$$

where $c_\theta < 0$ is the cost of that information. The decision maker should choose to gather the information only if:

$$\mathrm{E}_\theta\left[\max_i\{\mathrm{E}_y[u(x^{(i)},y,c_\theta)|\theta]\}\right] - \max_i\{\mathrm{E}_y[u(x^{(i)},y)]\} > 0 \tag{3}$$

Note that this is in the same as Equation (1) except for the addition of the cost c_θ to the reward function $u(\cdot)$, and the addition of "> 0" at the end. In this case, the cost of gathering the information is explicitly considered. Oftentimes, the cost of information c_θ is incurred separate from the outcome attributes $f(x^{(i)},y)$, so this can be written as:

$$\mathrm{E}_\theta\left[\max_i\{\mathrm{E}_y[u(f(x^{(i)},y)+c_\theta)|\theta]\}\right] - \max_i\left\{\mathrm{E}_y\left[u\left(f(x^{(i)},y)\right)\right]\right\} > 0 \tag{4}$$

Now, if the decision maker were risk neutral, then:

$$\begin{aligned}\mathrm{E}_\theta\left[\max_i\{\mathrm{E}_y[u(f(x^{(i)},y)+c_\theta)|\theta]\}\right] &= \mathrm{E}_\theta\left[\max_i\left\{\mathrm{E}_y\left[u\left(f(x^{(i)},y)\right)+u(c_\theta)\middle|\theta\right]\right\}\right] \\ &= \mathrm{E}_\theta\left[\max_i\left\{\mathrm{E}_y\left[u\left(f(x^{(i)},y)\right)\middle|\theta\right]\right\}\right] + u(c_\theta)\end{aligned} \tag{5}$$

and therefore, the decision maker would only gather information if:

$$\begin{aligned}&\mathrm{E}_\theta\left[\max_i\left\{\mathrm{E}_y\left[u\left(f(x^{(i)},y)\right)\middle|\theta\right]\right\}\right] + u(c_\theta) - \max_i\left\{\mathrm{E}_y\left[u\left(f(x^{(i)},y)\right)\right]\right\} > 0 \\ &\rightarrow \mathrm{E}_\theta\left[\max_i\left\{\mathrm{E}_y\left[u\left(f(x^{(i)},y)\right)\middle|\theta\right]\right\}\right] - \max_i\left\{\mathrm{E}_y\left[u\left(f(x^{(i)},y)\right)\right]\right\} > -u(c_\theta)\end{aligned} \tag{6}$$

This is strategically equivalent to Equation (1), and thus Equation (1) is simply a reduction of this more general formulation in Equation (4) when risk neutrality is assumed. However, if the decision maker were not risk neutral, then in general the utility function $u(f(x^{(i)},y)+c_\theta)$ *cannot* be separated in this way. Rather, the cost of information must be considered explicitly as part of the problem; it cannot be separated and placed on the other side of the inequality. The value of information then is c_θ where c_θ is such that Equation (4) is an equality rather than an inequality.

Under some assumptions about risk attitude, simplifications can be made in calculating the value of information. For example, assume that the risk attitude follows constant absolute risk aversion (CARA). Then it can be written in the form (Keeney and Raiffa 1993):

$$\begin{aligned} u(x)\sim -e^{-cx} &\leftrightarrow r(x)\equiv c>0 \quad (\mathit{constant\ risk\ aversion}) \\ u(x)\sim x &\leftrightarrow r(x)\equiv 0 \quad (\mathit{risk\ neutrality}) \\ u(x)\sim e^{-cx} &\leftrightarrow r(x)\equiv c<0 \quad (\mathit{constant\ risk\ proneness})\end{aligned} \tag{7}$$

where $r(x)$ is the local risk aversion at x and c is a given constant, the *risk coefficient*. In this case, because of the exponential nature of the utility function and the logarithm nature of the certainty equivalent, the inputs to the utility function *can* be separated when considering the values as certainty equivalents of the expected utility, as shown below for the continuous risk averse case:

$$\begin{aligned}
&\mathrm{E}_\theta\left[\max_i\{\mathrm{E}_y[u(f(x^{(i)},y)+c_\theta)|\theta]\}\right] \\
&= \int_\theta\left[\max_i\left\{\int_y[-\exp(-cf(x^{(i)},y)-cc_\theta)]\,p(y|\theta)dy\right\}\right]p(\theta)d\theta \qquad (8)\\
&= \int_\theta\left[\max_i\left\{\int_y\left[-\exp\left(-cf(x^{(i)},y)\right)\exp(-cc_\theta)\right]p(y|\theta)dy\right\}\right]p(\theta)d\theta \\
&= -\exp(-cc_\theta)\int_\theta\left[\max_i\left\{\int_y\left[\exp\left(-cf(x^{(i)},y)\right)\right]p(y|\theta)dy\right\}\right]p(\theta)d\theta
\end{aligned}$$

The certainty equivalent of this is:

$$\begin{aligned}
&u^{-1}(-\exp(-cc_\theta))+u^{-1}\left(\int_\theta\left[\max_i\left\{\int_y -\exp\left(-cf(x^{(i)},y)\right)p(y|\theta)dy\right\}\right]p(\theta)d\theta\right) \qquad (9)\\
&= c_\theta+u^{-1}\left(\mathrm{E}_\theta\left[\max_i\left\{\mathrm{E}_y\left[u\left(f(x^{(i)},y)\right)\middle|\theta\right]\right\}\right]\right)
\end{aligned}$$

Thus, when considering the problem in terms of certainty equivalents under a CARA risk attitude, the value of information is the certainty equivalent of deciding with the additional information minus the certainty equivalent of deciding without the additional information:

$$c_\theta = u^{-1}\left(\mathrm{E}_\theta\left[\max_i\left\{\mathrm{E}_y\left[u\left(f(x^{(i)},y)\right)\middle|\theta\right]\right\}\right]\right)-u^{-1}\left(\max_i\left\{\mathrm{E}_y\left[u\left(f(x^{(i)},y)\right)\right]\right\}\right) \qquad (10)$$

where $u^{-1}(\cdot)$ is the certainty equivalent for an expected utility, found via the inverse of the utility function. The importance of this is that from a practical standpoint, the certainty equivalents for the case of deciding with information and the case of deciding without information are substantially easier to calculate than the iterative procedure of finding c_θ such that it satisfies Equation (4) as an equality by varying the value of c_θ.

4 ENGINEERING EXAMPLE

4.1 Problem Description

This example is a highly simplified model meant to illustrate the concepts discussed in this paper and how they may be used on a practical engineering problem.

An LCD manufacturer is planning the layout for a new Thin-Film Transistor (TFT) manufacturing plant. This plant produces substrates with the semi-transparent controller electronics that form part of every LCD panel. Each TFT substrate produced by the plant is then mated with a colour filter (CF) substrate, and then liquid crystal (LC) is added between the two substrates, to form an LCD panel which can then be sold for profit.

In this model, the plant consists of production machines and inspection machines. Production machines perform each processing step needed to convert a bare substrate into a TFT substrate. However, production machines can also introduce defects, which render the LCD unusable. If the defective substrates are not found during production, then they will finally be discovered when the LCD panel is activated. At that point, the entire LCD panel must be scrapped, which includes the CF substrate and LC.

The manufacturer thus uses inspection machines to detect defects in substrates during production. A defective substrate, when found, is scrapped immediately. In this way, the plant does not expend additional production resources on substrates that will eventually be thrown away.

Each inspection machine takes the place of a production machine, however, so each inspection machine reduces the potential output of the TFT manufacturing plant. The total number of machines in the plant is fixed. Thus, the LCD manufacturer is looking to find the number of production machines, and therefore the number of inspection machines, that maximizes the plant's overall profit including the costs of the CF substrate and LC. Too few production machines leads to a low production output, decreasing profit. Too many production machines leads to too few inspection machines, resulting in excessive CF and LC scrap. The manufacturer must find the right balance that will maximize the overall profit of the plant.

For simplicity, the TFT production process consists of five layers. Each layer requires five processing steps. At the end of each layer, the substrate can be examined by an inspection machine for defects. If there are not enough inspection machines to cover every layer for every substrate, the substrates are sampled instead, where each substrate is randomly selected to be inspected, depending on inspection capacity. It is assumed that the sampling rate for each layer is the same. The inspection machine can detect defects from previous layers as well as the current layer, but multiple defects on a single substrate do not increase the detection rate. The outcome of an inspection is simply that a defect was detected, resulting in a scrap of the substrate, or that no defect was detected, and the substrate continues with production. There is also a chance of detecting a defect when none exists, resulting in the scrapping of a good TFT substrate.

The inspection machine vendor is introducing a new model, and thus there is uncertainty about the quality of the newer inspection machines. Specifically, there is uncertainty about the newer inspection machines' throughput (number of substrates processed per day) and defect detection rate (probability of detecting a defect). To help make the decision about the number of production machines to purchase, and thus also the number of inspection machines, the manufacturer can gather information about the newer inspection machine model to better understand its capabilities.

The manufacturer would like to maximize the plant's profit at the end of two years from the start of construction. As a simplification, it is assumed that during the first year the plant is involved with moving in the different machines and setting them up, while the second year is when the machines are in full operation. There are costs to purchasing each production and inspection machine, as well as costs of production such as for each bare substrate and for each processing step.

The profit over the two years for different inspection machine throughputs and defect detection rates, as a function of the number of production machines purchased, is shown in Figure (1).

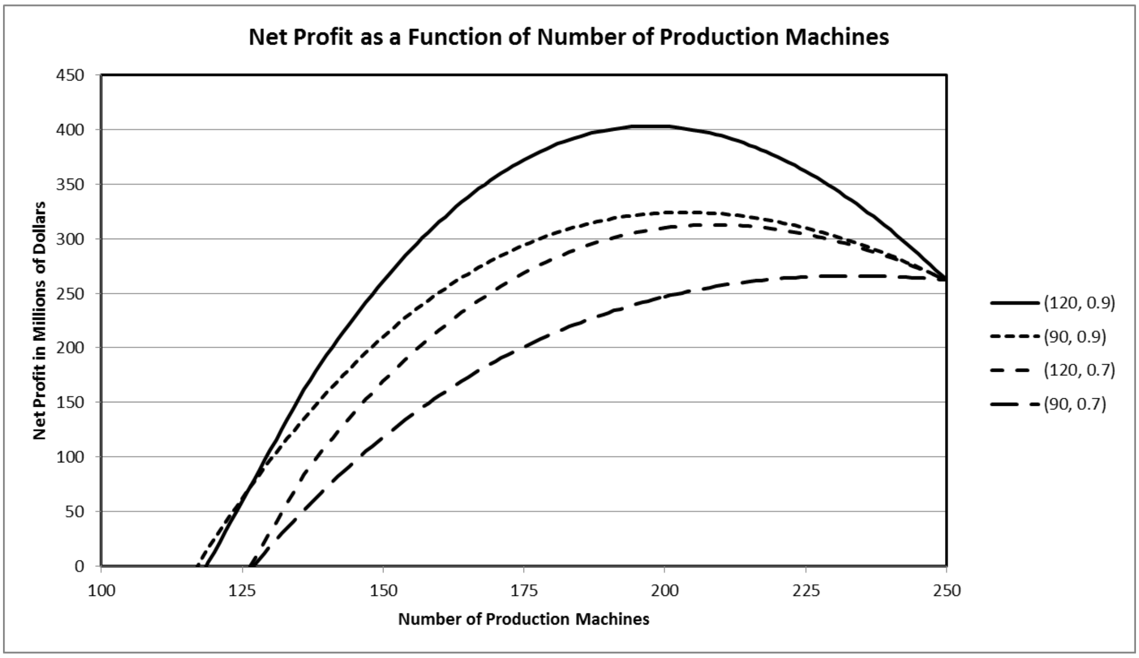

Figure 1. Net Profit Based on Number of Production Machines Purchased.

Depending on the inspection machine throughput (first value) and the defect detection rate (second value), there is a rapid decrease in profit if the manufacturer selects less than about 180 production machines. However, the optimal number of production machines can vary from 197 to 233 depending on the inspection machine parameters.

For this paper, the manufacturer's prior beliefs are a uniform distribution for the inspection throughput between 90 and 120, and a uniform distribution for the defect detection rate between 0.7 and 0.9. If the manufacturer gathers information, the information is perfect; the actual values for both of these parameters will be revealed perfectly.

4.2 Problem Formulation

The goal is to formulate the manufacturer's decision process mathematically, including both the information-gathering aspect as well as the production decision (i.e. number of production machines to purchase) aspect. For the latter, the manufacturer's production decision is how many production machines to purchase, with the inspection machine throughput and defect detection rate as uncertain parameters. These make up the design parameter $x^{(i)}$ and uncertain parameters y respectively. Since there is only one design parameter in this example, the superscript (i) will be dropped. The model described in Section 4.1 and shown in Figure (1) represents $f(x, y)$. It represents the relationship between the manufacturer's production decision and uncertain parameters with the outcome under interest in terms of the plant's net profit.

The manufacturer can choose to gather information, incurring a cost of c_θ. For this paper, perfect information will be assumed, and thus:

$$f(x, y)|\theta = f(x, \theta) \tag{11}$$

The decision problem confronting the manufacturer is thus:

$$\text{Maximize}\left(\mathrm{E}_\theta\left[\max_x\{u(f(x,\theta) + c_\theta)\}\right],\ \max_x\{\mathrm{E}_y[u(f(x,y))]\}\right) \tag{12}$$

The outcome of solving Equation (12) gives the highest expected utility. In the course of solving the problem, the arguments to those maximizations, namely whether or not to gather information and the number of production machines to purchase, will also be found, which inform the manufacturer of what alternatives to choose. Although the first alternative involves an expectation over θ while the second alternative involves an expectation over y, the manufacturer's beliefs about θ prior to receiving the information is the same as his beliefs over y itself; the two distributions are identical. Note that the differences between the two alternatives are the flipped order of expectation and maximization operations, and the cost of gathering information. The latter alternative involves one maximization over uncertain parameters y, while the former alternative involves many maximizations, one for each possible value of θ. Thus, one involves a single "larger" optimization problem, while the other involves many simpler optimization problems. A decision tree of the decision is shown in Figure (2).

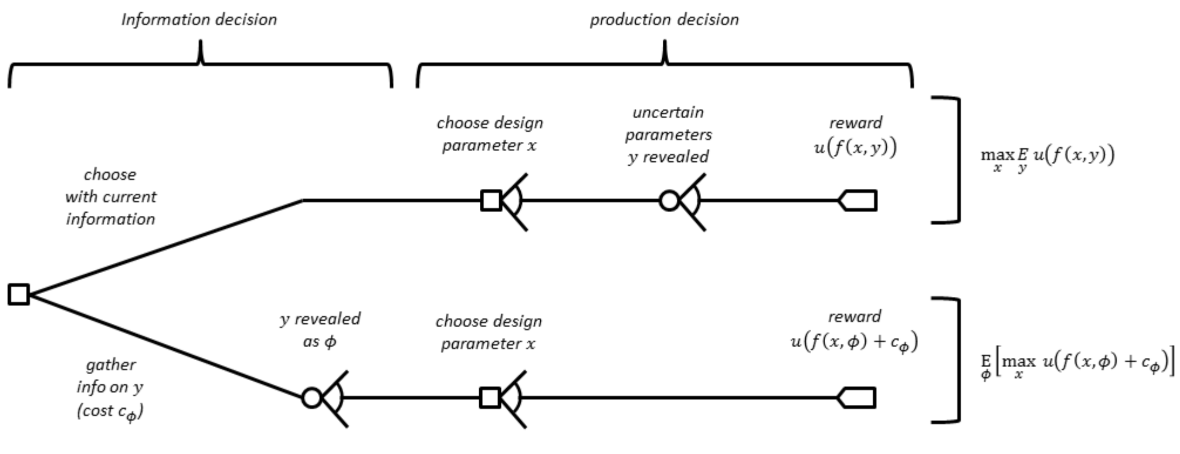

Figure 2. Decision Tree for Engineering Example.

In this paper, the manufacturer will be assumed to have a CARA risk attitude.

4.3 Results

The above engineering example was simulated via quadrature of the uncertain parameters of inspection machine throughput and inspection machine defect detection rate. The expected value of perfect information based on risk attitude, in terms of the certainty equivalent of the expected utility, is shown in Figure (3).

If the risk coefficient is less than zero, indicating a more risk-prone risk attitude, then the expected value of perfect information is also lower compared with under risk neutrality. Intuitively, this means that if the manufacturer is more willing to accept the downsides, then the value of receiving perfect information about the uncertain parameters is lower; the manufacturer does not value the information as highly. Similarly, if the risk coefficient is greater than zero, indicating a more risk-averse risk attitude, then the expected value of perfect information is higher. This means that the manufacturer is more willing to have additional information about the potential outcomes in order to avoid the possible downsides of the design decision.

As an example, say the manufacturer's risk coefficient were 6 per $1 million, and the actual cost of gaining perfect information were $1.40 million. If the manufacturer analysed the decision without considering risk attitude, i.e. under risk neutrality, the calculated value of perfect information would be $1.37 million. Under risk neutrality, then, the recommended solution would be to not gather information, and to make the purchase decision under the currently available information. However, with the manufacturer's risk attitude taken into account, the calculated value of perfect information would be $1.48 million. Thus, the manufacturer *should* expend resources to gather more information, or study this decision further in the event that the information is not perfect. If the actual cost of information is within the band in between $1.37 million for risk neutrality and the actual values of perfect information for different risk coefficients as shown in Figure (3), then the manufacturer will make a non-optimal decision if risk attitude is not taken into account. Although this difference of about a million dollars may not seem significant in light of the possible profit of hundreds of millions of dollars, the decision itself is still worth a significant amount of money, regardless of the total expenditure for this construction project.

As a side note, there are several kinks in the graph in Figure (3). This is because the manufacturer must choose an integer number of production machines to purchase, particularly in the case of not gathering more information, even though the optimal number may include fractions of a machine. For example, for risk coefficient c from -10 to -5.4, if information is not gathered, the optimal number of production machines to purchase is 204. For $-5.3 \leq c \leq 0.2$, the optimal number of production machines to purchase is 205, and so forth. This discrete number of allowable machines to purchase leads to different segments with noticeable kinks at their boundaries for the certainty equivalents.

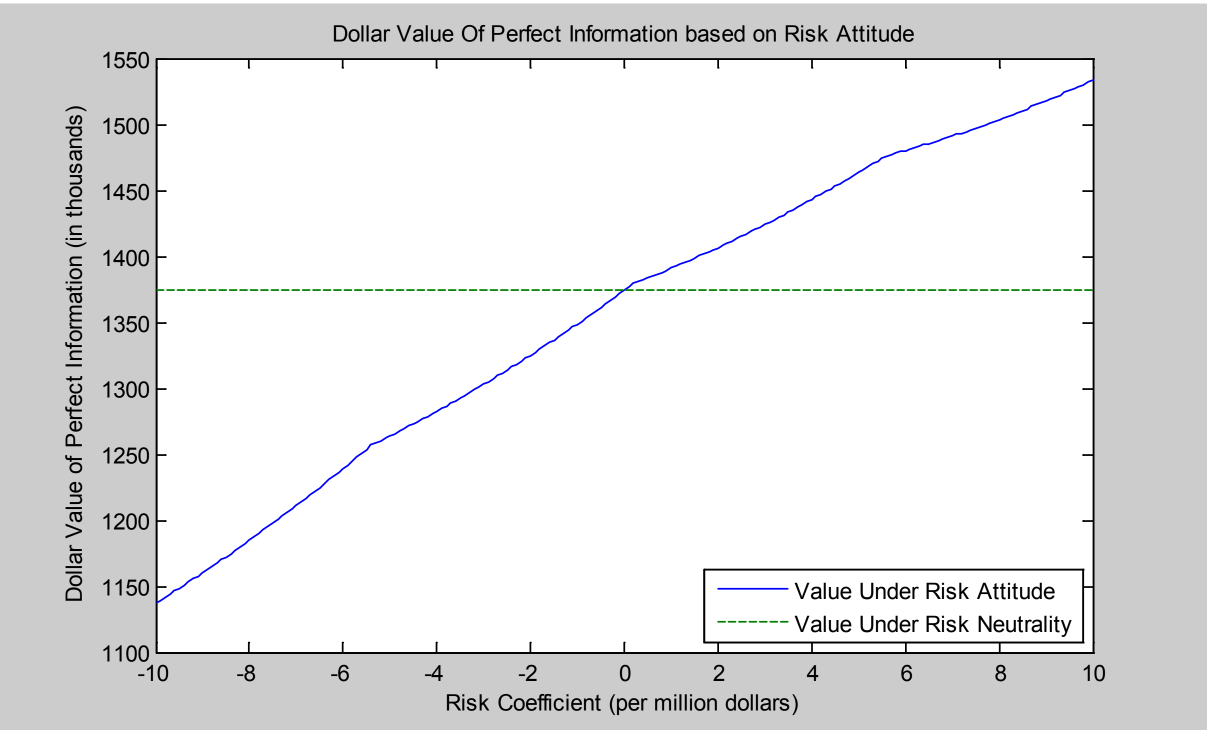

Figure 3. Expected Value of Perfect Information Based on Risk Attitude.

5 DISCUSSION AND FURTHER WORK

In the engineering example, it was found that greater risk aversion did lead to a higher value of perfect information. Although Eeckhoudt and Godfroid (2000) have previously investigated a counterexample of this, further insight into the conditions under which greater risk aversion does lead to a higher value of perfect information, and indeed of the relationship between them, would be beneficial to decision makers in an engineering context.

Additionally, risk aversion and the value of information of subsets of uncertain parameters can be studied. Although this paper only investigates the case of perfect information versus not gathering information, the engineering example purposely includes multiple uncertain parameters for future study of this type of problem. This allows for modelling the gathering of information in stages (i.e. sequentially), which better represents the typical decision problem confronting engineers in projects. In such situations, the upstream decisions will affect the downstream decisions, and thus that effect should be considered in making the upstream decisions. Such decisions can be very difficult to computationally solve because of the large number of nested expectations and maximizations involved. Thus, many practicing decision makers tend to view the situation more heuristically, rather than using a formal decision theoretic procedure. A better understanding of this effect in a Value of Information context, combined with the decision maker's risk attitude, will give better guidance for decision makers when considering the effect of information in a decision problem.

Although this paper assumed that the information source gave perfect information about the outcomes, the work can be extended to imperfect (sample) information. This would more accurately reflect the decision situation facing many decision makers in an engineering project. Since some imperfect information sources are repeatable, i.e. will give different results for each test (such as testing a material's yield strength, which will give different results for different samples), modelling imperfect information as well as sequential information can lead to greater insights.

The more general problem formulation given in Equation (3), compared with the typical formulation given in Equation (1), can be applied to a wider variety of situations. For example, it can be used to study the case where the cost of gathering information is itself uncertain. In a risk neutral situation, Equation (1) can be applied to this type of problem, by substituting expected cost for the cost. However, when risk attitude is considered, Equation (3) should be applied.

Finally, although this work assumed constant absolute risk aversion, other risk attitude models can be applied. Because non-neutral risk attitudes lead to greater mathematical complexity in the solution of information problems, one avenue of research is finding more efficient methods for computing the value of information under different utility functions. A related question of how much deviation from risk neutrality leads to how different of a solution can also be investigated. This would give insight to decision makers about whether or not the additional mathematical complexity of incorporating risk attitude is worth the greater understanding of the problem and more accurate solution compared with simply assuming risk neutrality, an information problem of its own.

6 CONCLUSION

This paper investigated some of the challenges involved when the Value of Information problem formulation is combined with risk attitude in an engineering context. It showed that the typical problem formulation implicitly assumes risk neutrality, and does not generally apply for a non-risk neutral risk attitude. The paper also showed the formulation for a decision maker with a constant absolute risk aversion (CARA) risk attitude. It demonstrated how the optimal decision changes when risk attitude changes in a manufacturing example.

Although decisions involving the gathering of information has been studied for many years under various Value of Information approaches, it is only fairly recently that incorporating risk attitude in modelling these decisions has been recognized as a topic of research. By incorporating a decision maker's risk attitude, the problem formulation is more complete and can more accurately reflect the considerations of the decision maker. As society embarks on more and more sophisticated projects, the ability to more accurately model decision problems and provide more accurate guidance for decision makers in complex engineering projects will lead to improved outcomes and better use of our planet's finite resources for a more sustainable future.

REFERENCES

Bickel, J. E. (2008) 'The Relationship Between Perfect and Imperfect Information in a Two-Action Risk-Sensitive Problem', *Decision Analysis,* 5(3), 116-128.

Bradley, S. and Agogino, A. M. (1994) 'Intelligent Real Time Design Methodology for Component Selection: An Approach to Managing Uncertainty', *Journal of Mechanical Design,* 116(4), 980-988.

Eckermann, S., Karnon, J. and Willan, A. R. (2010) 'The Value of Value of Information', *Pharmacoeconomics,* 28(9), 699-709.
Eeckhoudt, L. and Godfroid, P. (2000) 'Risk Aversion and the Value of Information', *The Journal of Economic Education,* 31(4), 382-388.
Gould, J. P. (1974) 'Risk, Stochastic Preference, and the Value of Information', *Journal of Economic Theory,* 8, 64-84.
Hilton, R. W. (1981) 'The Determinants of Information Value: Synthesizing Some General Results', *Management Science,* 27(1), 57-64.
Howard, R. A. (1966) 'Information Value Theory', *IEEE Transactions on Systems Science and Cybernetics,* 2(1), 22-26.
Keeney, R. L. and Raiffa, H. (1993) *Decisions with Multiple Objectives,* 2 ed., Cambridge, UK: Cambridge University Press.
Mehrez, A. (1985) 'The Effect of Risk Aversion on the Expected Value of Perfect Information', *Operations Research,* 33(2), 455-458.
Mehrez, A. and Stulman, A. (1982) 'Some Aspects of the Distributional Properties of the Expected Value of Perfect Information (EVPI)', *The Journal of the Operational Research Society,* 33, 827-836.
Miller, A. C. (1975) 'The Value of Sequential Information', *Management Science,* 22(1).
Nadiminti, R., Mukhopadhyay, T. and Kriebel, C. H. (1996) 'Risk Aversion and the Value of Information', *Decision Support Systems,* 16, 241-254.
Samson, D., Wirth, A. and Rickard, J. (1989) 'The Value of Information from Multiple Sources of Uncertainty in Decision Analysis', *European Journal of Operational Research,* 39, 254-260.
Tull, D. S. and Hawkins, D. I. (1976) *Marketing Research: Meaning, Measurement, and Method: A Text with Cases,* New York: MacMillan.