



AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search

<http://ageconsearch.umn.edu>

aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

No endorsement of AgEcon Search or its fundraising activities by the author(s) of the following work or their employer(s) is intended or implied.

Fiducial cost-benefit analysis of research: with an application to weather modification

Paper presented to the 2012 AARES Conference, February 7-9, Perth Australia

Stephen Beare¹ and Raymond Chambers²

Abstract

Environmental intervention is often seen as being high risk and high return. Traditional scientific hypothesis testing provides limited guidance to policy makers unless there is a high level of certainty in the supporting scientific evidence. Traditional cost-benefit analysis under uncertainty has shortcomings when considering high-risk investment, largely due to the choice of how to discount uncertainty outcomes. A corollary is that traditional cost-benefit analysis does not place a value on increased certainty, an important outcome of successful scientific research. A fiducial cost-benefit methodology is presented in this paper, which integrates hypothesis testing and traditional cost-benefit analysis. The fiducial approach is one way of objectively placing a value on changes in the level of uncertainty that does not depend on an assumption about a decision maker's attitudes towards variability in returns. This has two important implications. First, there is a level of uncertainty at which we would reject an investment with a positive expected net rate of return on the basis that the uncertainty associated with the outcome is too great. Second, it is possible to value a program of research that reduces the uncertainty about a critical decision parameter. An example based on data from a weather modification experiment conducted in South Australia is presented. The approach is the generalised using more traditional statistical methodology.

Key Words: Cost-benefit, risk, uncertainty, cost-benefit, fiducial inference

¹ Stephen Beare is the director of Analytecon and corresponding author, stephenbeare@analytecon.com .au.

² Raymond Chamber is a professor of statistics at the University of Wollongong.

1. Introduction

Economics and environmental science share a common basis in observational as opposed to controlled experimental data. Statistical models play a large role in systematically establishing the presence of changes and relationships. The effects in question are often small in relation to natural background variation and are often unevenly dispersed over a wide area. Furthermore, the source is often not linked to the effect with a clear measure of exposure and the system in which the effect is embedded is non-stationary and historical correlations are often not reliable benchmarks. The consequence is a level of uncertainty that is high and often greater than indicated by standard measures of statistical precision.

The detection and evaluation of intended or unintentional changes in natural systems due to human intervention is an area where this uncertainty is a dominant feature. In particular, this uncertainty is simply too great to influence or control through experimental design or to account for through an observable set of covariates. Hence, neither randomisation nor conditioning is able to isolate and measure the hypothesised effect with a level of accuracy that meets accepted scientific precedents within a practical experimental horizon.

At the same time, the effect in question can have large ramifications for public decision-making, as for example, should a government:

- Intervene on the basis of existing information;
- Delay interventions and fund further scientific investigation; or
- Accept the status quo.

The dilemma is then, given that there are substantive benefits to making the correct decision, there is, regardless of what decision is made, a substantive probability and cost of being incorrect. Cost-benefit analysis is a standard tool for decision-making by weighting uncertain benefits and costs, but it does not definitively answer the question of what is the best decision. In cases of very high levels of risk and return it may offer little if any effective guidance. The reasons for this lie with how the probabilities of alternative outcomes of a decision are assigned, and the rate at which uncertain outcomes are discounted.

Another way of thinking about discounting uncertainty is valuing additional information. There are a number of ways to approach this problem, one of which is through fiducial inference. A fiducial cost-benefit framework is developed in this paper to address the limitations of standard cost-benefit analysis. Fiducial inference is a type of inductive inference that was developed by R. A. Fisher in the 1940s. It subsequently lost favour with classical or frequentist statisticians because of its limited application and the fact that fiducial distributions did not always obey the laws of probability. However, recently there has been renewed interest in fiducial inference, particularly in the Bayesian community, largely attributable to its practical application (Hannig 2009). In particular, in certain cases (e.g. when the underlying distribution is Normal) it allows Bayesian credible intervals to be constructed without recourse to a prior.

In this paper we construct a framework for decision making by considering the opportunity costs of falsely rejecting the correct decision. This is done formally in the context of a hypothesis test and the opportunity costs of falsely rejecting either the null or the alternative hypothesis. For example, a simple one-tailed test might be of whether the effect in question exists (alternative hypothesis) or does not exist (null hypothesis). An appropriate probabilistic weighting of these outcomes is derived through fiducial inference. This weighting is expressed as a critical significance level above which the alternative hypothesis that an effect exists should not be rejected, since the cost of its rejection outweighs the costs of rejecting the null hypothesis. At a level below this critical level the null hypothesis should not be rejected. An important byproduct of this approach is that discounting of uncertainty is unnecessary and a reduction in uncertainty can be valued. A research evaluation focus is retained throughout the development of the approach.

The approach is applied to data from a weather modification experiment conducted in South Australia, and addresses the question of whether future trials should be conducted. Weather modification technologies, if effective, can generate very large net economic benefits. However, over 50 years of scientific investigation has not established conclusive evidence that cloud seeding results in increased precipitation (NRCNAS, 2003).

Interest in weather modification goes beyond rainfall enhancement to the unintended effects on climate of particulate emission from power plants and urbanisation and other forms of land use change. Geo-engineering options to mitigate climate change are being proposed, such as using unmanned ships to spray seawater into the atmosphere to increase cloud albedo on a global scale to reflect solar energy back out to space (Satler et al. 2008). Furthermore, the issues extend to how land, marine and water resources are managed more generally, with many aspects of theory only

tentatively supported by empirical evidence. In this context, the questions of interest go beyond what are the best management options to the value of research aimed at gaining a better understanding of the effectiveness of different management options.

2. Cost-Benefit Analysis with Uncertainty

The terms risk and uncertainty are sometimes used interchangeably. There is an economic distinction that uncertainty pertains to outcomes with known or estimated probabilities and risk is where no meaningful probabilities can be ascribed. (Knight, 1921). Here we focus on uncertainty. However, it can be useful to think of uncertainty in terms of possible states of nature and the consideration of risk arises when there are competing strategies to manage across those states of nature. This perspective provides context to the question of how uncertainty is mapped into risk.

In the first and most common application of the cost-benefit framework, uncertainty is considered through the use of expected values. Outcomes of an uncertain investment decision are discounted at a rate equal to their probability of occurrence or in a continuum by their probability density function. Probability weighted outcomes of an uncertain process are used to calculate the expected net present value of a flow of costs and benefits arising from an investment decision. In the absence of any costs or benefits bearing additional or less risk, the expected value can be equated with certainty. A decision maker is faced with a range of potential but uncertain investment returns can rank investment options on the basis of expected rates of return. Effectively, there is no association between uncertainty and risk and the decision maker is said to be risk neutral.

Where the bearing of risk is not costless the comparison of investment options requires a measure of risk about different rates of return in order to rank them in a consistent way. Mapping uncertainty into risk raises the question of what are we uncertain about? In cost-benefit analysis uncertainty is about the size of the sampling error. That is, the difference between an unknown true value and its corresponding estimate.

Expected utility and portfolio theory provide different means of extending cost-benefit analysis to explicitly consider risk. Within each approach uncertainty is represented by the sampling distribution of an estimate, in this case the estimated rate of return. The most common measure of uncertainty is the standard error of an estimate. If the sampling distribution of the estimation error is normal, the distribution is symmetric and fully defined by its expected value and variance. The

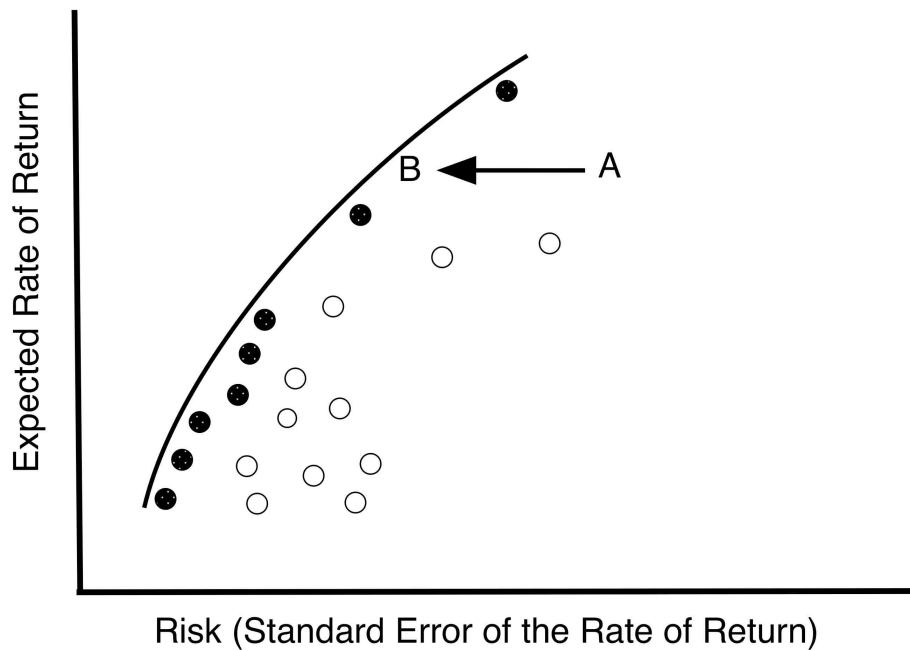
skewness and other moments of an error distribution can substantively change the characteristics of inferences drawn from that estimate. Here normality is assumed, as it does not distract attention from the underlying question as to how the sampling distribution of an estimate can effectively characterise risk.

If the acceptance of risk is not costless then when presented with two options with equal expected rates of return, the one with the lower level of risk is preferred. This leads to the idea of a frontier defined by an efficient set of possible investments that offer the lowest risk for a given rate of return. A stylised example is given in Figure 1. The investment alternatives that are shown in black form the frontier, dominating other alternatives, shown as open circles, that generate similar rates of return at higher levels of risk.

Investment options along the expected returns frontier can't be ranked without the introduction of a decision metric that allows returns with a higher level of uncertainty to be discounted. In the context of portfolio theory the metric is based on the minimisation of diversifiable risk in a portfolio that includes a risk free rate of return. Here discounting of risk takes the form of risk premiums or hurdle rates over what is assumed to be a risk free rate of return. In expected utility theory the preferences of the decision maker for risk and return are maximised. Here the discount rates are a measure of individual's aversion or preference for risk.

Extrapolating such empirically derived discount rates into the realm of private and public investment options with very high risk and return profiles is open to question, largely due to subjective choice of what is being optimised. While imposing the assertion of risk neutrality may help maintain a greater degree of consistency and objectivity, it may not be particularly prudent. Further, the decision to delay a policy intervention to allow time for further research, or even a program of ongoing adaptive management, is difficult to justify. It would be highly desirable to be able to attach an objective value to a reduction in uncertainty, shown in Figure 1 as a movement from A to B, which is often the primary outcome of research.

Figure 1 A stylised risk and return frontier



2.1 An alternative view of uncertainty

Uncertainty can be framed in the context of a decision rule as opposed to an estimate. That is, the uncertainty is associated with the probability of making a correct or incorrect decision. The decision rule considered here is to minimise the opportunity cost of making an incorrect decision. This underpins the fiducial approach to cost benefit analysis. However, a simple coin-tossing example can be used to illustrate how an alternative view of uncertainty can be used to value information

Consider the following game. There are two coins, one fair and the other biased, with a probability of throwing heads of 60 per cent. These probabilities are known. The cost of playing the game is \$20 and player selects one coin at random and flips it 10 times receiving \$20 for each head and losing \$20 dollars for each tail. It is straightforward to determine that the expected payout is zero.

Now suppose the payer is offered the opportunity to twice toss the coin that they have randomly chosen before they commit the \$20 play to play the game. The question is how much is this opportunity to trial the coin worth? The problem can be approached in different ways. Here, the opportunity cost of the incorrect decision to decline or play the game is considered:

- The opportunity cost of incorrectly rejecting the hypothesis that the coin is fair and deciding to play which is clearly a loss of \$20; versus
- The opportunity cost of incorrectly rejecting the hypothesis that coin is not fair and declining to play which again is clearly a loss of \$20.

Set out this way the problem revolves around the costs and probabilities of false rejection. The probability of false rejection is simply one minus the probability of selecting either coin. The contingency table for the game prior to trialling the selected coin is shown below.

Coin	HH	HT	TT	Row Total
Fair	0.125	0.250	0.125	0.500
Biased	0.180	0.240	0.080	0.500
Col Total	0.305	0.490	0.205	1.000

Prior to the trial the probability of selecting either coin is 0.50. This gives equal weight to the opportunity costs of false rejection and is in line with the zero expected payoff, and consequent indifference between the decisions to play or to decline the game.

Given the three possible results from trialling the selected coin, the posterior probabilities associated with the selected coin are shown below.

Coin	HH	HT	TT
Fair	0.410	0.510	0.610
Biased	0.590	0.490	0.390

If one or less heads are tossed, the opportunity cost of incorrectly rejecting the hypothesis that the chosen coin is fair exceeds the alternative and the optimal decision is to decline the game and incur no cost. Given that two heads are tossed, the decision to reject the hypothesis that the chosen coin is fair and to play has an opportunity cost that is \$3.60 lower than the decision to decline the game. This is the ex post expected value of the decision to play. To calculate the ex ante value of the trial,

the expected return to each decision must be weighted by the probability of the trial outcomes, $(0 \times .205 + 0 \times .490 + 3.6 \times .305)$, which comes out to be approximately \$1.10.

The example does not touch on how the trial affects the uncertainty associated with the variability in the pay-off, which is trivial, given there are only two coin tosses. The value of the trial is derived from a change in expectations as opposed to a change in variance. Consideration of the latter is reason for introducing fiducial inference.

3. Fiducial Cost-Benefit Analysis

Hypothesis testing, which underpins the scientific method, provides another perspective on the idea of risk and return. Consider the following standard one-tailed hypothesis test:

$$H_0: \text{Effect} = 0$$

$$H_a: \text{Effect} > 0$$

As with the previous example, risk and return are couched in terms of the direct and opportunity costs of false rejection. That is, the cost of falsely rejecting the null hypothesis H_0 (type-I error) versus the cost of false rejection of the alternative hypothesis H_a (type-II error). Clearly, if the cost of falsely rejecting H_0 is far greater than the cost of falsely rejecting H_a it would be rational to try to limit the probability of a type-I error and set a significance level S for the test which is quite high. Conversely, if the cost of falsely rejecting H_0 is less than falsely rejecting H_a , it would be rational to try to limit the probability of a type-II error and hence set a significance level S for the test which is relatively low. The point is that the costs of false rejection can be seen as a form of weighting the uncertainty about whether to reject either H_0 or H_a . The objective is to formalize this weighting in an objective way.

3.1 The fiducial framework

The issue addressed is whether to reject either H_0 or H_a on the basis of the relative costs of making an incorrect judgment, given the experimental information. What is of interest is what value of S would leave us indifferent between the costs of falsely rejecting either H_0 or H_a . Put another way,

what critical value of the relevant test statistic would equate the costs of falsely rejecting either H_0 or H_a .

In conducting the test a critical value is chosen for the test statistic T , which is assumed to have a known sampling distribution under H_0 . The critical value is a mapping of the significance level S into the percentiles of the sampling distribution of T under H_0 . If the value of the T falls below the critical value, H_a is rejected in favour of H_0 (the rejection of H_a being the de-facto consequence of failing to reject H_0). Alternatively, T can be mapped through its sampling distribution under H_0 to an observed significance level (or p-value) defined by the order of the percentile of this distribution that equates to the observed value of T . If this p-value is greater than S , H_a is rejected in favour of H_0 .

The significance level S is a probability bound on the realised value of T under H_0 . That is, all observed values of T less than or equal to the critical value $C(S)$ are 'consistent' with H_0 being true, and hence one can consider the set $T \leq C(S)$ as defining a confidence region under H_0 with confidence level equal to $1 - S$. How this confidence region is interpreted distinguishes our approach to inference. Within the frequentist approach it is a statement about the sampling distribution of T under H_0 , i.e. $\Pr\{T \leq C(S)\} = 1 - S$ under H_0 . Here T is a random variable calculated from the data about the fixed parameter value given by H_0 . Under this approach, and with $S = 0.1$ significance level, one would expect that if 100 samples were drawn, then, if H_0 is true, the 90 per cent confidence interval $(0, C(S = 0.1))$ would be expected to contain the observed value of T 90 per cent of the time.

When T is a pivotal statistic, i.e. its sampling distribution is the same irrespective of whether H_0 or H_a is true, this confidence region can be mapped to an interval, say $E(T, S)$, of potential values for the effect, all of which cannot be distinguished from the null value given the chosen significance level S and the value of T . The essence of the fiducial approach is that it interprets $E(T, S)$ as a statement that the probability that the unknown effect takes values within this interval given the observed value of T is $1 - S$. That is, given a one-sided test with significance level S leading to a $100(1 - S)$ per cent confidence interval $(0, T = C(S))$, the fiducial approach asserts that the probability that the unknown effect lies in $E(T, S)$ is $1 - S$. Since this holds for any value of S , one can derive a distribution for the unknown effect given the observed value of T . Given this distribution, the key to reformulating the cost benefit analysis is then to use it to estimate of the probabilities of falsely rejecting H_0 or H_a from the sample data.

3.2 Using the fiducial framework to price a change in uncertainty

To apply the fiducial framework two economic values need to be calculated. For example, in the context of adopting or commercialising a research finding:

- (1) The cost of falsely rejecting H_0 which is taken to be the costs of adoption or commercialisation; and
- (2) The cost of falsely rejecting H_a , which is taken to be the benefits foregone if a successful commercialisation is not attempted.

The actual calculation of these values would present a wide range of issues. However, to facilitate the development of the fiducial framework an example is constructed in the context of a field trial that generates a positive but uncertain treatment effect. The benefits of adopting are assumed to be proportional to the size of the treatment effect.

Given the data obtained in the trial, it is assumed that it is possible to generate a fiducial distribution for the unknown treatment effect, x , with mean μ and variance σ^2 . Note that both μ and σ^2 will be functions of the trial data. For convenience, it is also assumed that this distribution is normal, i.e. the treatment effect, x , is assumed to be drawn from a normal distribution with mean μ and variance σ^2 . To calculate the cost of falsely rejecting H_a we need to calculate the probability that H_a is true. This can be done in different ways. Here it is done in two steps. First, the expected net benefit, given that H_a is true, is computed as:

$$NB_{x>0} = E(x|x > 0)V - \text{Adoption Costs}$$

where NB is the net benefit, V is the value of the treatment effect and E denotes expectation with respect to the fiducial distribution of x . Given the assumed normality of this distribution,

$$E(x|x > 0) = \mu + \sigma \frac{\phi(-\sigma^{-1}\mu)}{1 - \Phi(-\sigma^{-1}\mu)}$$

where ϕ and Φ are the standard normal probability and cumulative density functions respectively.

In the second step x is replaced by $\tilde{x} = kx$, where $k = E(x|x > 0)\mu^{-1}$, so $E(\tilde{x}) = \tilde{\mu} = E(x|x > 0)$ and $.sd(\tilde{x}) = \tilde{\sigma} = E(x|x > 0)\mu^{-1}\sigma$. For a test that generates a confidence interval with a specified confidence level of 100λ per cent, the probability of falsely rejecting H_0 is by design $1 - \lambda$. The probability of falsely rejecting H_a at this confidence level depends on the distribution of

$$\frac{\tilde{x} - \tilde{\mu}}{\tilde{\sigma}}$$

which is taken to be standard normal (as an approximation). The probability of falsely rejecting H_a is therefore a power calculation based on the value of λ :

$$F_a(\lambda) = \frac{1}{1 - \lambda_0} \int_0^b \phi\left(\frac{z - \tilde{\mu}}{\tilde{\sigma}}\right) dz$$

where $\lambda_0 = \Phi\left(\frac{-\tilde{\mu}}{\tilde{\sigma}}\right)$ and $b = \tilde{\mu} + \tilde{\sigma}\Phi^{-1}(\lambda)$, where $\Phi\{\Phi^{-1}(p)\} = p$.

The total net benefit is therefore

$$NB(\lambda) = (1 - \lambda) \times \text{Program Costs} + F_a(\lambda)NB_{x>0} \quad (1)$$

and the break-even confidence level is the value of λ that equates this total net benefit to zero:

$$\lambda = \dots$$

This relationship effectively values the effect of an increase or decrease in the level of precision with which the treatment effect is measured.

3.3 An example

A rainfall enhancement trial using ground-based ion generating arrays was conducted in South Australia in 2009 (Beare *et al.* 2011). A primary aim of the trial was to test the hypothesis that operation of the system in the assessment region leads to increased rainfall downwind of the generators. The trial used two sites operated using a randomised crossover design. The sites had similar orography and a high level of correlation in average rainfall.

The analysis of the trial data used gauge-level data and a statistical model that incorporated fixed meteorological, orographic and operating effects as well as random effects. The random effects were added to sweep out latent or unmeasured influences and included an effect for each gauge and downwind blocks that would capture the correlation in rainfall between successive downwind rain gauges as rain-bearing clouds move through the area.

The attribution or enhancement effect was defined as the ratio of observed rainfall to observed rainfall less predicted change in rainfall given the status of a generator was switched from off to on. This is a complex statistic in the sense that its sampling distribution is unknown and there are potentially complex correlations in the gauge-level residual errors, as the location of the downwind target gauges varied with wind direction. Random effect spatio-temporal block bootstrap simulations were used to assess the variability of the estimated change in rainfall (Chambers and Chandra, 2012). The blocking was done to isolate downwind and cross wind variation in the correlation between gauge residuals.

The attribution could be calculated as a simple average of all rain gauges. Gauges can also be weighted according to the size of the area they represent. A Voronoi area is the area defined as being geographically closest to a particular gauge location. A Voronoi area-weighted average will be more reflective of rainfall volumes over the trial area, however the precision of the estimate is likely to decline as greater weights are given to gauges that are more widely dispersed. While the simple gauge average may be preferable from a statistical perspective in that it can be more accurately measured, it does not provide the best estimate of the volume of rainfall falling in the target area and any of that attributable as rainfall enhancement. The weighted average is likely to be better aligned with the net economic benefits of any additional rainfall. The bootstrap distributions for the two enhancement estimates are shown in Figures 2 and 3 respectively.

The fiducial cost-benefit framework is used to look the uncertainty around each of these estimates, summarised as:

- Simple average attribution: A mean enhancement of 9.4 per cent and a bootstrap relative standard error of 72.3 per cent (significant at approximately the 90 per cent confidence level); and
- Voronoi area weighted attribution: A mean enhancement of 5.6 per cent and a bootstrap relative standard error of 114.3 per cent (significant at approximately the 80 per cent confidence level)

Figure 2 The random effect spatio-temporal block bootstrap distribution of the estimated change in rainfall using unweighted gauge-level averages





Figure 3 The random effect spatio-temporal block bootstrap distribution of the estimated change in rainfall using Voronoi area weighted gauge-level averages



However to avoid the complexity of looking at changes in the level and precision of the estimated enhancement effect, the mean enhancement effect for the simple average is scaled back to 5.6 per cent while the relative standard error is retained.

The opportunity cost of falsely rejecting the trial results was taken to be the research program costs including capital and operating costs as well as the costs of evaluation. The cost of falsely rejecting the alternative hypothesis taken to be the value of the useful water attributed to the operation of the Atlant generators. It was calculated from the estimated levels of attribution, either a simple gauge average or the Voronoi area weighted gauge average. The estimated change in rainfall was converted to effective run-off into urban and rural catchment areas and valued at the cost of bulk urban charges and traded prices for irrigation water, respectively. On average the per percentage point benefits were of the order of \$1.3 million.

The results are summarised in Figures 4 and 5 in which the total net benefit, defined by equation (1), is plotted against the confidence level λ . The curves are generated numerically and then smoothed. The smoothing results in the curves meeting prior to reaching the 50 and 100 per cent confidence levels where they do in fact meet.

The reduced precision of the simple average attribution shifts the breakeven confidence limit to the right as it increases the probability of falsely rejecting the alternative hypothesis at any given confidence limit as shown in Figure 4. While both results are above the breakeven point due to the high return attributed to a 5.6 per cent increase in rainfall.

For any given confidence level for the estimated increase in rainfall effect (shown at 80 and 90 per cent in Figure 5), the simple average measurement generates a greater net benefit than the weighted average unless there is certainty in which case they become equivalent. Conversely, a decision maker weighting the costs against the benefit of increased rainfall over the region would calculate a lower net benefit for any given level of confidence in the results.

It is possible to fall below the break-even point even if the project has a positive expected benefit using standard cost-benefit analysis. This occurs for the Voronoi area weighted estimate where the per percentage point return falls below about \$1 million. This is a fundamental difference between standard and fiducial cost-benefit; the level of uncertainty can be great enough to reject a project with a positive expected rate of return in its own right. No comparison need be drawn with other investment options.

Figure 4. Breakeven confidence limits for the area weighted and simple gauge average attributions





Figure 5. Graphical calculation of the benefits of reducing the level of uncertainty



In Figure 5, the benefit of being able to increase the precision of the Voronoi weighted average estimate from 80 to 90 per cent confidence level is shown. The additional precision therefore, if it could be achieved, would be valued at approximately \$1.1 million in the context of the decision to adopt the technology. Alternatively this might be seen as the cost of using area-based weights to gain more accurate assessment of volumetric rainfall.

3.4 An alternative Empirical Bayes approach to pricing uncertainty

The problem of pricing a change in uncertainty can be recast as a problem in Empirical Bayes inference. This is more in line with the analysis of the coin-tossing game presented in Section 2. A standard treatment and effect approach is used. Note that *NID* stands for "Normal, independently distributed". Let θ = actual treatment effect and $\hat{\theta}$ = estimated treatment effect, with a measurement model:

$$\hat{\theta} = \theta + \epsilon$$

where ϵ is measurement error, and a process model:

$$\theta \sim N(\mu, \sigma^2)$$

Given these models, $\hat{\theta} | \theta \sim N(\theta, \sigma^2)$. So

$$\theta | \hat{\theta} \sim N\left(\frac{\sigma^2 \mu + \hat{\theta} \sigma^2}{\sigma^2 + \sigma^2}, \frac{\sigma^2 \sigma^2}{\sigma^2 + \sigma^2}\right)$$

and hence

$$E[\theta | \hat{\theta}] = \frac{\sigma^2 \mu + \hat{\theta} \sigma^2}{\sigma^2 + \sigma^2}$$

It immediately follows that

Now

so

Similarly

The Expected Net Benefit Given Decision to Proceed (expected gains less adoption costs) is then

and the Expected Net Benefit Given Decision Not to Proceed (foregone gains but no adoption costs)

is

We need to appropriately weight these two Expected Net Benefit expressions, taking into account the asymmetric nature of the Expected Net Benefits and the outcomes:

1. Reject when it is true -

2. Reject \square when it is true (Do not reject \square when it is not true) - \square .

One weighting is

Error! Not a valid embedded object.

where \square is a density on \square . However, it is not obvious how this density should be chosen.

Given \square , an alternative is to estimate \square by:

A lower bound on this estimate is obtained by setting:

4. Concluding Comments

The fiducial cost-benefit decision calculus based on formal hypothesis testing that has been put forward in this paper provides a more robust approach to decision making under substantial levels of uncertainty. It eliminates a key subjective aspect of standard cost-benefit analysis, as it does not require the choice of rate at which to discount uncertain outcomes.

The fiducial approach allows a value to be placed on changes in the level of uncertainty. This has two important implications. First, there is a level of uncertainty at which we would reject an investment with a positive expected net rate of return on the basis that the uncertainty associated with the outcome is too great. Second, it is possible to value a program of research that is able to reduce the uncertainty about a critical decision parameter.

However, these values are couched in terms of foregone opportunity costs associated with a mutually exclusive decision to, for example, continue or abandon a research program. In many

instances this may not be appropriate as the problem does allow the clear delineation of the opportunity costs of a hypothesis test. Further, it is open to any bias that might be introduced into the costing of the two alternatives. For example, assuming that a decision has irreversible consequences when in fact it might just represent the time value of a delay. The converse is another example.

This approach offers a means for evaluating investment decisions with respect to:

- Innovative approaches and technologies where the returns are large but highly uncertain.
- Research designed to reduce the uncertainty about key parameters that are seen to be critical determinants of policy intervention into natural systems.

In these instances there often isn't the capacity to conduct experiments under controlled conditions and there is a high reliance on observational data. Fiducial cost-benefit analysis then provides a framework for objectively comparing options where uncertainty is high but the cost of foreclosing on options or extended delay is substantial.

The restriction to normality that underpins the fiducial cost-benefit model allows the probabilities of false rejection to be calculated from the data. An alternative set out in Section 3.4 is to use an Empirical Bayes approach. The advantages and disadvantages of the two approaches remain to be tested. It may be that while it comes at a cost of an additional unknown parameter, the Empirical Bayes approach may be more robust in terms of application.

References

Beare S, Chambers R, Peak S, Ring J (2011). Accounting for Spatiotemporal Variation of Rainfall Measurements when Evaluating Ground-Based Methods of Weather Modification. *Journal of Weather Modification*. 43, 44-63.

Chambers R and Chandra H (2012). A Random Effect Block Bootstrap for Clustered Data. *Journal of Computational and Graphical Statistics* To appear.

Hannig, J (2009). On Generalized Fiducial Inference, *Statistica Sinica* 19, 491-544.

Knight FH (1921). Risk, Uncertainty, and Profit. Boston, MA: Hart, Schaffner & Marx; Houghton Mifflin Company.

National Research Council (NRCNAS) (2003). Critical Issues in Weather Modification Research. National Academy Press, 123 pp.

Salter S, Sortino G and Latham J (2008). Sea-going Hardware for the Cloud Albedo Method of Reversing Global Warming. *Phil.Trans.Roy. Soc. A* 366, 2989-4006.