Risk Analysis, Vol. , No. , 0000 DOI: 000

A Conjugate Class of Utility Functions for Sequential Decision Problems

Brett Houlding,[1] Frank P.A. Coolen[2], Donnacha Bolger[1].

The use of the conjugacy property for members of the exponential family of distributions is commonplace within Bayesian statistical analysis, allowing for tractable and simple solutions to problems of inference. However, despite a shared motivation, there has been little previous development of a similar property for using utility functions within a Bayesian decision analysis. As such, this paper explores a class of utility functions that appear to be reasonable for modeling the preferences of a decision maker in many real-life situations, but which also permit a tractable and simple analysis within sequential decision problems.

Keywords: Decision Analysis, Utility Theory, Risk Analysis, Preference Modelling, Matched Updating, Normative Choice Theory.

1. INTRODUCTION

We consider the problem a Decision Maker (DM) faces when seeking the optimal selection strategy within a Bayesian sequential decision problem. Such a situation is a form of dynamic programming problem, the solution to which involves the use of backward induction and Bellman's Equation (Bellman 1952, 1957). However, this technique requires that the DM evaluate a nested sequence of maximizations and integrations, with the latter not necessarily having closed form solutions. Sequential decision problems occur naturally and commonly (*e.g.*, Cox, 2012), and hence appropriate methodologies are applicable in a wide range of realistic problems, such as in the area of clinical trials within Medical Statistics (see Brockwell & Kadane, 2003), in trade-off problems between cost and performance within complicated systems (Muller, 1999 or Benjaafar *et. al,*, 1995), and in portfolio management in an economic framework (Guler, 2007).

In general, we can express the problem in the following manner: a DM is facing a sequential problem of length n in which a decision $d_i \in \mathcal{D}$ must be selected at each decision epoch $i = 1, \ldots, n$. The DM's objective is to select decisions so as to maximise the utility of the entire sequence $d_1, \ldots, d_n$, and this is assumed to depend on the value of an *a priori* uncertain parameter θ. In the following work we shall assume that this unknown θ will remain constant over time, *i.e.*, it is a static rather than a dynamic parameter. The DM has assigned or obtained a prior distribution $P(\theta)$ for θ and may learn about its true value by observing return variable r_i following decision selection d_i subject to the likelihood $P(r_i|d_i, \theta)$. As such, the DM may influence the information they receive concerning the parameter θ by making suitable decision selections $d_1, \ldots, d_{n-1}$ (the DM will also learn about θ following decision d_n, but as there will be no further decision to be made within the

[1]Discipline of Statistics, Trinity College Dublin, Ireland, houldinb@tcd.ie.
[2]Department of Mathematical Sciences, Durham University, UK.

 0272-4332//0100-0001$22.00/1 ©

problem's planned horizon, such information will not be of use and will thus not be considered).

An example of such a situation exists when each decision selection d_i results in a return $r_i \in \mathcal{R}$ with probability $P_{d_i}(r_i|\theta) \equiv P(r_i|d_i, \theta)$. In this case we may assume that the DM has a known utility function u for the return stream $r_1, \ldots, r_n$, which may then be converted to a utility function U over the decision sequence $d_1, \ldots, d_n$ via the expected utility representation, *i.e.*, $U(d_1, \ldots, d_n) = E[u(r_1, \ldots, r_n)]$, with this expectation taken with respect to the return stream $r_1, \ldots, r_n$. In the remainder of the material presented we will use the notation U to represent both the utility of a decision sequence for a given value of the parameter θ, and for the expected utility of a decision sequence with respect to beliefs over θ.

In this case the optimal decision for selection in epoch i, denoted π_i, will be a function of the history h_i of the previously made decisions and their observed returns, *i.e.*, $h_i = \{(d_1, r_1), \ldots, (d_{i-1}, r_{i-1})\}$. Implementation of Bellman's Equation then leads to the following optimal decision selection strategy (where for notational convenience we have set $U_i = U(d_1, \ldots, d_i, \pi_{i+1}(h_{i+1}), \ldots, \pi_n(h_n))$):

$$\begin{aligned} \pi_i(h_i) &= \arg\max_{d_i \in \mathcal{D}} E_{h_{i+1}|h_i, d_i} \\ &\quad [\cdots E_{h_n|h_{n-1}, \pi_{n-1}(h_{n-1})}[U_i]] \\ &\quad \text{for } i = 1, \ldots, n-1, \qquad (1) \\ \pi_n(h_n) &= \arg\max_{d_n \in \mathcal{D}} U_n. \qquad (2) \end{aligned}$$

Equation (1) contains a nested sequence of expectations, each of which requires the DM to consider a distribution of the form $P_{h_j|h_{j-1}, d_{j-1}}$. This conditioning argument implies that the only remaining uncertainty in h_j is with respect to the outcome r_{j-1} that will occur following selection of decision d_{j-1}. It should also be noted that Equations (1) and (2) need to be evaluated for all possible histories h_i for all possible i, hence explaining why sequential decision problems suffer from the so-called 'curse of dimensionality'.

Apart from the maximisation step, which can be readily solved if there are only a relatively small number of options within the set of possible decisions $\mathcal{D}$, problems in solving Equation (1) arise when the DM's beliefs are represented by non-discrete probability distributions. In such a situation Equation (1) requires that the DM evaluate a nested series of integrals, where in each case other than the innermost integral, the integrand is the product of the inner integral and the appropriate probability distribution. The innermost integrand, however, is the product of the appropriate probability distribution and the DM's utility function for this problem. As such the solution to such a nested series of integrals will not generally be available in closed form.

One approach to resolving this problem is to determine sufficiently accurate approximations by using techniques from numerical analysis such as discretisation, or possibly by using Monte Carlo methods akin to Brennan *et al.* (2007), Berry *et al.* (2010) and Muller (1999). However, despite continued advances in computational power, the problem concerning the solution of such a nested series of integrals still persists due to a need for fast and accurate solutions to high dimensional problems of interest. Attempts have been made to decrease this computation time, notably by Brockwell & Kadane (2003), who suggest gridding on sufficient statistics of exponential family members, and creating an algorithm which is linear in the number of sequential stages involved in the decision making procedure, diminishing the curse of dimensionality somewhat, and providing an approximate solution to the sequential decision problem considered. Yet despite this, there remains motivation to search for classes of utility functions that are both reasonable and flexible for representing the DM's true preferences in various situations, but which also allow closed form solutions to the nested series of integrals when the DM's beliefs are represented by appropriate probability distributions.

The remainder of this paper is as follows: in Section 2 we discuss the concept of conjugate utility as first considered by Lindley (1976), drawing comparison with the much more commonly used concept of conjugate probability distributions, and highlighting the problem with a straightforward sequential extension. In Section 3 we introduce a polynomial class of utility functions that is conjugate to certain probability distributions and which is also quite flexible for representing preferences. Section 4 demonstrates the effectiveness of this utility class in a couple of examples, whilst Section 5 concludes.

2. CONJUGATE UTILITY

The concept of conjugate utility originates from Lindley (1976), but subsequently appears to have received little further development, with the only notable instances being Novick and Lindley (1978) which considered use of utility functions as a mixture of k conjugate utility forms in a particular educational environment, and Islam (2011) which formally extends Lindley's conjugate functions to a bivariate setting. General research has been carried out in trying to find a suitable class of utility functions to ease the decision making process in different strands of decision theory, for instance by LiCalzi & Sorato (2006) who advocated the use of HARA (Hyperbolic Absolute Risk Aversion) utility functions, which are those functions $u(r)$ such that $-\frac{u''(r)}{u'(r)} = \frac{1}{a+br}$ for $a, b \in \mathcal{R}$, *i.e.*, the coefficient of absolute risk aversion is the reciprocal of a linear function of r. This approach is noteworthy in its applicability both to the normative decision making approach of expected utility theory (Von Neumann & Morgenstern, 1947), and also to the descriptive decision making methodology of Prospect Theory (Kahneman & Tversky, 1979).

Lindley's motivation was to explore a related idea to that of the conjugate prior in Bayesian statistical inference, but where instead a so-called 'conjugate' family of utility functions is sought with the property that they are both suitably 'matched' to a probability structure and realistic for application. This latter requirement is key, as unlike a probability distribution the only constraint placed upon a utility function is that it be a bounded function of its arguments. It is therefore quite easy to determine a utility form that satisfies any specified 'matching' criteria, but unless these functions represent a reasonable model of the actual subjective preferences of the DM, they will be unsuitable for inclusion in any meaningful decision analysis.

The idea of the conjugate prior for probability parameters in Bayesian statistical analysis was introduced by Raiffa & Schlaifer (1961), and is now a commonly used concept. In this setting a class of prior probability distributions $P(\theta)$ is said to be conjugate to a class of likelihood functions $P(r|\theta)$ if the resulting posterior distributions $P(\theta|r)$ are in the same family as $P(\theta)$, *i.e.*, if both $P(\theta|r)$ and $P(\theta)$ have the same algebraic form as a function of θ. A particular result of this definition is that all members of the exponential family of probability distributions have conjugate priors. The exponential family is a particularly important class of probability distributions that is commonly used in statistical modelling. In the Appendix we present a brief overview of conjugate updating in the case of exponential family members, as well as providing a brief overview of the work of Lindley (1976) in choosing an appropriate form of utility function to use in conjunction with a data generating mechanism from the exponential family.

Unfortunately, the techniques and ideas used within Lindley's utility class of do not readily generalise to allow exact expected utility calculations within a sequential decision problem. In order to provide a demonstration of this consider a sequential problem of length $n = 2$ and the following likelihood for decision return with normalising constant $G(\theta, d_i)$:

$$P(r_i|d_i, \theta) \quad = \quad G(\theta, d_i)H(r_i, d_i)e^{\theta r_i}. \tag{3}$$

Equation (3) is a simple generalisation of an exponential family member, as defined in the Appendix, which ensures that, regardless of decision d_i, the density itself will be a member of the exponential family. However, the actual density that Equation (3) represents is allowed to depend on the particular d_i selected. This is an important distinction from the work of Lindley that is necessary if the theory is to be applicable for sequential problems of interest.

Given suitable hyper-parameters n_0, $\boldsymbol{r}_0 = (r_{01}, \ldots, r_{0n_0})$, and $\boldsymbol{d}_0 = (d_{01}, \ldots, d_{0n_0})$, a natural conjugate prior for this likelihood is the following, where the notation $\boldsymbol{x}[j]$ will be used to denote the j-th element in $\boldsymbol{x}$, and $K(n_0, \boldsymbol{r}_0, \boldsymbol{d}_0)$ is the normalising constant:

$$P(\theta|n_0, \boldsymbol{r}_0, \boldsymbol{d}_0) = K(n_0, \boldsymbol{r}_0, \boldsymbol{d}_0)\Big(\prod_{i=1}^{n_0} G(\theta, \boldsymbol{d}_0[i])\Big) \times e^{\theta \sum_{i=1}^{n_0} \boldsymbol{r}_0[i]}. \tag{4}$$

Here $\boldsymbol{d}_0$ denotes a collection of n_0 hypothetical decisions, with $\boldsymbol{r}_0$ the corresponding collection of n_0 hypothetical returns resulting from these. This may be seen as analogous to, for instance, the choosing of hyperparameters for a Beta prior distribution, with these hyperparameters indicative of the number of hypothetical successes and failures witnessed in a collection of hypothetical trials, with an increased number of trials symptomatic of an increased confidence in the prior. After having selected decision d_1 the DM will observe return r_1 and update beliefs over θ to the following (where $\boldsymbol{r}_1 = (r_{01}, \ldots, r_{0n_0}, r_1)$ and $\boldsymbol{d}_1 = (d_{01}, \ldots, d_{0n_0}, d_1)$):

$$P(\theta|r_1,d_1,n_0,\boldsymbol{r}_0,\boldsymbol{d}_0) = K(n_0+1,\boldsymbol{r}_1,\boldsymbol{d}_1) \times \Big(\prod_{i=1}^{n_0+1} G(\theta,\boldsymbol{d}_1[i])\Big) \times e^{\theta\sum_{i=1}^{n_0+1}\boldsymbol{r}_1[i]}. \quad (5)$$

In generalising Lindley's utility form we make the requirement that a utility function for a decision stream does indeed depend on all the decisions within that stream. As such, one possibility is to consider a utility function that is a product of un-normalised densities of the form given in Lindley (1976), discussed in the Appendix (with the number of terms within the product being determined by the length of the decision sequence). In the two-period sequential problem this leads to the following, which is a sequential extension of that proposed by Lindley:

$$U(d_1,d_2,\theta) = F(d_1,d_2)G(\theta,d_1)^{n_1(d_1)}G(\theta,d_2)^{n_2(d_2)} \times e^{\theta(n_1(d_1)f_1(d_1)+n_2(d_2)f_2(d_2))}. \quad (6)$$

Hence, when preferences are as stipulated in Equation (6), the expected utility of a decision d will have a known 'closed' form when beliefs over the uncertain parameter follow a distribution of the exponential family. This can also be said for the utility family derived by Lindley, which is included as Equation (25) in our Appendix. We argue that the utility family of Equation (6) may be reasonable for representing preferences, and Lindley (1976) provides discussion of scenarios and special occasions in which its use may be appropriate. Given that the utility form of Equation (6) follows the formula of an un-normalised density, it will in general be most useful when corresponding to the un-normalised density of a uni-modal distribution with mode θ. In this case the utility of a decision will be measured by how accurately its value approximates the true value of the uncertain parameter θ. Note that this is not the only possibility for generalising Lindley's utility form for the case of a sequential problem; an alternative would be to use a sum of un-normalised densities rather than a product. Nevertheless, both choices suffer from the same problem and so we restrict attention to the product form of Equation (6). In what follows we assume that the returns $r_1,\ldots,r_n$ are conditionally independent given the parameter θ, which follows from our earlier comment about the static nature of θ.

Once decision d_1 has been selected and r_1 returned, the DM will have updated beliefs to be as in Equation (5). Decision d_2 will then be selected so as to maximise the following (where $\boldsymbol{r'}_2$ is the vector consisting of $\boldsymbol{r}_1$ followed by $n_1(d_1)$ repetitions of $f_1(d_1)$ and by $n_2(d_2)$ repetitions of $f_2(d_2)$, whilst $\boldsymbol{d}_2$ is the vector consisting of $\boldsymbol{d}_1$ followed by $n_1(d_1)$ repetitions of d_1 and by $n_2(d_2)$ repetitions of d_2):

$$\begin{aligned} U(d_1,d_2|r_1) = & \int F(d_1,d_2)G(\theta,d_1)^{n_1(d_1)}G(\theta,d_2)^{n_2(d_2)} \\ & \times\ e^{\theta(n_1(d_1)f_1(d_1)+n_2(d_2)f_2(d_2))} \\ & \times\ K(n_0+1,\boldsymbol{r}_1,\boldsymbol{d}_1)\Big(\prod_{i=1}^{n_0+1} G(\theta,\boldsymbol{d}_1[i])\Big) \\ & \times\ e^{\theta\sum_{i=1}^{n_0+1}\boldsymbol{r}_1[i]}d\theta, \\ = & \frac{F(d_1,d_2)K(n_0+1,\boldsymbol{r}_1,\boldsymbol{d}_1)}{K(n_0+n_1(d_1)+n_2(d_2)+1,\boldsymbol{r'}_2,\boldsymbol{d}_2)}. \quad (7) \end{aligned}$$

Once the DM knows how they will select decision d_2 for any given values of d_1 and r_1 (with this d_2 found by application of Equation (2), namely that $\pi_2(r_1,d_1) = \arg\max_{d_2} U(d_1,d_2|r_1)$), attention will be focused on the selection of decision d_1. By Equation (1) this is found through the following:

$$\pi_1 = \arg\max_{d_1\in\mathcal{D}} E_{r_1|d_1}[E_{\theta|r_1,d_1}[U(d_1,\pi_2(r_1,d_1),\theta)]]. \quad (8)$$

The problem in generalising Lindley's utility form for sequential decision problems now occurs. Whilst the utility form of Equation (6) ensures that the term $E_\theta[U(d_1,\pi_2(r_1,d_1),\theta)]$ may be determined exactly, it does not guarantee that the expectation of this term with respect to the predictive distribution of r_1 has a closed form solution. The solution to the integral of the product of $E_\theta[U(d_1,\pi_2(r_1,d_1),\theta)]$ and $P(r_1|d_1)$ will in general depend on the specific functions included in both the exponential family expression and the selected sequential utility function, despite the fact that $P(r_1|d_1)$ can be expressed exactly (another result that follows from the use of exponential family distributions), in the form given below in Equation (9):

$$P(r_1|d_1) = \frac{H(r_1,d_1)K(n_0,\boldsymbol{r}_0,\boldsymbol{d}_0)}{K(n_0+1,\boldsymbol{r}_1,\boldsymbol{d}_1)}. \quad (9)$$

Using the notation where $\boldsymbol{r}'_1$ is the vector consisting of $\boldsymbol{r}_0$ followed by $n_1(d_1)$ repetitions of $f_1(d_1)$ and by $n_2(\pi_2(r_1,d_1))$ repetitions of $f_2(\pi_2(r_1,d_1))$, whilst $\boldsymbol{d}'_1$ is the vector consisting of $\boldsymbol{d}_0$ followed by $n_1(d_1)$ repetitions of d_1 and by $n_2(d_2)$ repetitions of

$\pi_2(r_1, d_1)$, the DM should select π_1 via the following:

$$\pi_1 = \arg\max_{d_1 \in \mathcal{D}} K(n_0, \boldsymbol{r}_0, \boldsymbol{d}_0) \int \frac{F(d_1, \pi_2(r_1, d_1))H(r_1, d_1)}{K(n_0 + n_1(d_1) + n_2(d_2) + 1, \boldsymbol{r}_1', \boldsymbol{d}_1')} dr_1. \quad (10)$$

The problem in solving Equation (10) arises because the integral it contains does not have a general closed form solution. Whilst the integral may be solved for certain and specific functions F, H and K, nothing in their definition ensures that this will always be the case.

3. THE POLYNOMIAL UTILITY CLASS

As an alternative to the un-normalised exponential family distribution form for utility that is suggested by Lindley, we now consider the use of a polynomial utility class. As will be demonstrated in the remainder, the proposed polynomial utility class allows closed form solutions to sequential decision problems when beliefs are represented by Normal distributions, whilst it is simultaneously flexible enough to be an adequate representation of beliefs in a variety of situations.

First assume that prior beliefs over θ are such that this parameter follows a Normal distribution with mean μ and variance σ^2. Furthermore, we assume that the distribution of return r_i also follows a Normal distribution (Markowitz (2012) discusses the use of this assumption, as well as that of quadratic utility functions, which we shall soon see are a subset of our polynomial utility class) with unknown mean $\mu_{d_i}(\theta) = \alpha_{d_i}\theta + \beta_{d_i}$ (α_{d_i} and β_{d_i} being known constants) and known variance $\sigma^2_{d_i}$.

In the case of an unknown mean but known variance, the Normal distribution is a member of the exponential family of distributions that is conjugate with itself. Hence returns $r_1, \ldots, r_n$ are observed following the selection of decisions $d_1, \ldots, d_n$, respectively, posterior beliefs for θ can be easily determined given the following respective likelihood and prior forms:

$$P_{d_1,..,d_n}(r_1,\ldots, r_n|\theta) = \prod_{i=1}^{n} P_{d_i}(r_i|\theta) \propto \exp\left\{-\sum_{i=1}^{n} \frac{(r_i - \mu_{d_i}(\theta))^2}{2\sigma^2_{d_i}}\right\}, \quad (11)$$

$$P(\theta) \propto \exp\left\{-\frac{(\theta-\mu)^2}{2\sigma^2}\right\}. \quad (12)$$

Using Normal-Normal conjugacy to combine these lead to posterior beliefs for θ which follow a Normal distribution with mean ν_n and variance η_n^2, where these parameters are specified by the following:

$$\nu_n = \frac{\mu + \sigma^2 \sum_{i=1}^{n} \frac{\alpha_{d_i}(r_i - \beta_{d_i})}{\sigma^2_{d_i}}}{1 + \sigma^2 \sum_{i=1}^{n} \frac{\alpha^2_{d_i}}{\sigma^2_{d_i}}}, \quad (13)$$

$$\eta_n^2 = \frac{\sigma^2}{1 + \sigma^2 \sum_{i=1}^{n} \frac{\alpha^2_{d_i}}{\sigma^2_{d_i}}}. \quad (14)$$

Now consider the following polynomial utility form, which is independent of θ given the return stream:

$$u(r_1, \ldots, r_n) = \sum_{k_1=0}^{m_1} \sum_{k_2=0}^{m_2} \cdots \sum_{k_n=0}^{m_n} \Big(a_{k_1,k_2,\ldots,k_n} \times r_1^{k_1} r_2^{k_2} \cdots r_n^{k_n}\Big). \quad (15)$$

When beliefs over decision return follow a Normal distribution, this polynomial utility for return streams leads to the following utility for decision streams, given our aforementioned assumption that returns are conditionally independent of each other given θ. Note that here once again the expectation is taken with respect to the return stream $r_1, \ldots, r_n$, yielding the following:

$$\begin{aligned} U(d_1, \ldots, d_n|\theta) &= E_{r_1,\ldots,r_n|d_1,\ldots,d_n,\theta}[u(r_1, \ldots, r_n|\theta)] \\ &= E_{r_1,\ldots,r_n|d_1,\ldots,d_n,\theta}[u(r_1, \ldots, r_n)], \\ &= \sum_{k_1=0}^{m_1} \cdots \sum_{k_n=0}^{m_n} a_{k_1,\ldots,k_n} \times \prod_{i=1}^{n} \int r_i^{k_i} P(r_i|d_i, \theta) dr_i. \end{aligned} \quad (16)$$

The solution to the integrals on the right hand side of Equation (16) are the raw moments of the Normal distribution (see Table I below for the first five), which can be readily expressed in closed form by making an appropriate substitution and expressing them in terms of Gaussian Integrals, (see, *e.g.*, Papoulis 1991). In particular for the Normal distribution, the raw moments are of the following polynomial form (for suitable constants a_{ij} and $k \in \mathbb{N}$):

$$X|\mu, \sigma \sim \mathcal{N}(\mu, \sigma^2) \Rightarrow E[X^k] = \int x^k P(x|\mu, \sigma^2) dx = \sum_{i=0}^{k} b_{ki} \mu^i \sigma^{k-i}. \quad (17)$$

Table I . Raw Moments of the Normal Distribution

Order	Raw Moments
1	μ
2	$\mu^2 + \sigma^2$
3	$\mu^3 + 3\mu\sigma^2$
4	$\mu^4 + 6\mu^2\sigma^2 + 3\sigma^4$
5	$\mu^5 + 10\mu^3\sigma^2 + 15\mu\sigma^4$

The polynomial utility class of Equation (15) is very flexible, allowing for a reasonable model of preferences in many real-life situations. For example, the utility function $u(r_1, \dots, r_n) = \sum_{i=1}^{n} \phi^{i-1} r_i$ is included as a special case, and is suitable for representing preferences in situations where future decision returns are subject to discounting at rate $\phi \in [0, 1]$ (this is referred to as the Exponential Discounting Model). A further possibility is for when a trade-off exists between returns received in differing periods, *i.e.*, when $u(r_1, \dots, r_n) = \sum_{i=1}^{n} \phi_i u_i(r_i)$, where $\phi_i \geq 0$ is a trade-off weight satisfying the constraint $\sum_{i=1}^{n} \phi_i = 1$, and where u_i is a known polynomial function of return r_i. This may be an appropriate model for the situation in which each return represented a single attribute within a larger multi-attributed return $(r_1, \dots, r_n)$.

An alternative possibility for when $(r_1, \dots, r_n)$ is considered a multi-attribute return (see for instance, Huang *et. al* (2013), Chang (2011) and Musal *et. al* (2012)), and one that does not make such strong independence assumptions regarding preferences over differing attributes levels, is to take $u(r_1, \dots, r_n) = \left[\Pi_{i=1}^{n} [\phi^* \phi_i u_i(r_i) + 1] - 1\right] / \phi^*$, where $\phi_i \in (0, 1)$ and $\phi^* > -1$ represent non-zero scaling constants (see, *e.g.*, Keeney 1972). Again, provided u_i is a polynomial function of r_i, this utility function is also a member of the polynomial class of Equation (15).

Even when the appropriate utility function is known to be of a specific algebraic form that is not of the class described by Equation (15), *e.g.*, exponential or logarithmic utility, the polynomial utility class can still be used to provide an approximation through the use of Taylor Polynomials (see, for example, Hlawitschka 1994, Diamond & Gelles 1999). For an infinitely differentiable function f of a single variable x, the Taylor Series is defined on an open interval around a as $T(x) = \sum_{n=0}^{\infty} \frac{d^n f(x)}{dx^n}|_{x=a}(x - a)^n/n!$. The function f can then be approximated to a specified degree of accuracy by taking a partial sum of this series, and each such partial sum will be of the form of a polynomial in x. This result is also generalizable for approximating multivariate functions, hence allowing a greater class of utility functions to be approximated by the polynomial class of Equation (15).

We briefly comment on some of the shortcomings of the use of polynomial utility functions. Frequently polynomials are not monotonically increasing functions of their arguments. This may be problematic as a DM is likely to want to assign utility values to returns that are strictly increasing as the returns grow larger. If a DM is not careful then they may end up associating a greater utility value to a lower return then to a slightly higher one, *i.e.*, $u(r + \epsilon) < u(r)$ for some $\epsilon > 0$. Evidently care should be taken to ensure a DM avoids potentially illogical situations like this, as otherwise they may end up making irrational decisions in the event of extreme returns occurring. Polynomials are also unbounded and tend towards positive or negative infinity in extreme cases. This may cloud the decision making process of an individual, by placing a utility of an unreasonable magnitude on a particular event. In many realistic settings lower and upper bounds can be placed on the potential outcomes which may occur, and hence the domain of the utility function can be made compact. In section 4.2 we shall consider the use of a utility function which is unbounded, but it is assumed that values which will cause the function to misbehave have a negligibly small probability of occurrence, and hence will not have adverse consequences to a DM.

Finally we note that in the following examples the problem at hand is of a discrete decision nature. At each epoch an individual must choose one decision from a finite collection of possible alternatives, with the optimal decision path being that which maximises expected utility over the uncertain parameter θ. Another potential setting is one in which there is a continuum of possible decisions open to an individual, *i.e.*, they make their choice from an infinite set of alternatives. In one of the following examples we consider a decision problem where a DM must choose whether to buy a fixed amount of stock A, stock B or neither. She has a clearly finite amount of options. An illustration of how this problem could be mapped into a continuous decision domain would be where she wishes to decide how large a quantity of a particular stock to buy, meaning she now faces an infinite collection of possible choices. The work presented above is generalisable to a continuous decision setting, but incorporates some further complications in the maximisation aspect of calculations. This was discussed in Section 1, specifically how this maximisation is straightforward when dealing with a relatively small collection of options, but complexity increases with the number of options, which in the continuous case is an infinite amount.

4. EXAMPLES

4.1 Example 1

First consider the case where prior beliefs about unknown θ are such that $\theta \sim \mathcal{N}(0,1)$. At each of two epochs a decision must be made - either d_A or d_B in both cases. The returns associated with these decisions have Normal distributions $R_i|d_A,\theta \sim \mathcal{N}(\theta,1)$, and $R_i|d_B,\theta \sim \mathcal{N}(-\theta,2)$. A decision tree for the above sequential problem is given below in Fig. 1 (note that the shading indiciates the continuous nature of the potential returns):

Fig. 1. Decision Tree for Example 1

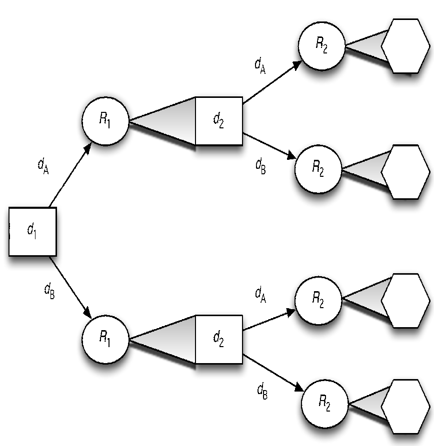

For a utility function consider $u(r_1,r_2) = r_1^2 + r_2$, which belongs to the polynomial utility class (15) as required. Then

$$\begin{aligned} u(r_1,r_2,\theta) &= \int (r_1^2 + r_2) \prod_{i=1}^{2} P(r_i|d_i,\theta) dr_1 dr_2, \\ &= \int r_1^2 \prod_{i=1}^{2} P(r_i|d_i,\theta) dr_1 dr_2 \\ &+ \int r_2 \prod_{i=1}^{2} P(r_i|d_i,\theta) dr_1 dr_2, \\ &= \int r_1^2 P(r_1|d_1,\theta) dr_1 \\ &+ \int r_2 P(r_2|d_2,\theta) dr_2. \end{aligned} \tag{18}$$

Now using (17), and that for the Normal Distribution the first raw moment is μ and the second is $\mu^2+\sigma^2$, we obtain Table II below containing utility outcomes for potential decision streams.

Table II . Utility for given decisions

d_1	d_2	$U(d_1,d_2,\theta)$
d_A	d_A	$\theta^2+\theta+1$
d_A	d_B	$\theta^2-\theta+1$
d_B	d_A	$\theta^2+\theta+2$
d_B	d_B	$\theta^2-\theta+2$

Using (13) and (14) updated beliefs about θ given the history of decisions taken and returns observed can be obtained after one and two decision epochs. The results are given below in Table III.

Table III . Posterior beliefs about θ

d_1	d_2	R_1	R_2	$P(\theta\|$ Rest $)$
d_A	-	r_1	-	$\mathcal{N}(\frac{r_1}{2},\frac{1}{2})$
d_B	-	r_1	-	$\mathcal{N}(-\frac{r_1}{3},\frac{2}{3})$
d_A	d_A	r_1	r_2	$\mathcal{N}(\frac{r_1+r_2}{3},\frac{1}{3})$
d_A	d_B	r_1	r_2	$\mathcal{N}(\frac{2r_1-r_2}{5},\frac{2}{5})$
d_B	d_A	r_1	r_2	$\mathcal{N}(\frac{2r_2-r_1}{5},\frac{2}{5})$
d_B	d_B	r_1	r_2	$\mathcal{N}(\frac{-r_1-r_2}{4},\frac{1}{2})$

Now consider the expected utility values at the far right hand side of the decision tree. For instance in the case of choosing d_A at both epochs we have the following (where we denote by h_2 the history of decisions made and returns observed up to this point):

$$\begin{aligned}
&E_{\theta|h_2,r_2,d_2=d_A}[\theta^2+\theta+1] \\
&= \int (\theta^2+\theta+1)P(\theta|h_2,d_2=d_A,r_2)d\theta, \\
&= \int \theta^2 P(\theta|h_2,d_2=d_A,r_2)d\theta \\
&+ \int \theta P(\theta|h_2,d_2=d_A,r_2)d\theta + 1, \\
&= \frac{r_1^2+2r_1r_2+r_2^2+3r_1+3r_2+12}{9}.
\end{aligned}$$

Note the use of the raw moments of the Normal distribution from (17) to compute the integrals above. Now to continue to "roll back" along the tree, moving from right to left, requires the predictive distribution of R_2 given the history h_2 of decisions and rewards observed up to that point. However $P(R_2|h_2,d_2) = \int P(R_2|d_2,\theta)p(\theta|r_1,d_1)d\theta$, which gives, for example, $E[R_2|r_1,d_1=d_A,d_2=d_A]=\frac{r_1}{2}$. The predictive expected values of R_2 are then used in the expected utility equations, *i.e.*, by replacing the R_2 terms by their expected value in terms of expressions in r_1. For instance, in the case of choosing d_A at both epochs we now have expected utility $\frac{r_1^2}{4}+\frac{r_1}{2}+\frac{4}{3}$.

At the decision nodes labeled d_2 the maximum utility path is chosen from those available as detailed above. In the case where $d_1=d_A$ this is $\max\Big\{(\frac{r_1^2}{4}+\frac{r_1}{2}+\frac{4}{3}),(\frac{r_1^2}{4}-\frac{r_1}{2}+1.4)\Big\}$ and in the case where $d_1=d_B$ this is $\max\Big\{(\frac{r_1^2}{9}-\frac{r_1}{3}+2.4),(\frac{r_1^2}{9}+\frac{r_1}{3}+2.5)\Big\}$. Obviously, determining the maximum is dependent upon the unknown value of R_1, however, it transpires that, when $d_1=d_A$ then $\pi_2=d_A$ if $r_1>0.667$, and $\pi_2=d_B$ otherwise. Similarly, when $d_1=d_B$ then $\pi_2=d_B$ if $r_1>-0.15$ and $\pi_2=d_A$ otherwise. Finally, to work out the expected values of the maximum, we evaluate

$$\begin{aligned}
&E_{R_1|d_1=d_A}\Big[\max\{f_1(r_1),f_2(r_2)\}\Big] \\
&= \int_{-\infty}^{0.0667} f_1(r_1)P(r_1|d_A)dr_1 \\
&+ \int_{0.0667}^{\infty} f_2(r_1)P(r_1|d_A)dr_1,
\end{aligned}$$

where

- $f_1(r_1)=\frac{r_1^2}{4}-\frac{r_1}{2}+1.4$,
- $f_2(r_1)=\frac{r_1^2}{4}+\frac{r_1}{2}+\frac{4}{3}$,
- $R_1|d_A\sim\mathcal{N}(0,1)$.

This yields $E_{R_1|d_1=d_A}\Big[\max\{f_1(r_1),f_2(r_2)\}\Big]=2.0165$ and a similar calculation for the bottom branch yields an expected utility of 3.050. Hence the decision sequence which maximises expected utility is $\pi_1=d_B$, followed by $\pi_2=d_B$ if $r_1>-0.15$, and $\pi_2=d_A$ otherwise. Note, however, that the given utility function $u(r_1,r_2)=r_1^2+r_2$ is slightly more risk seeking for positive returns than it is risk averse for negative returns. This is consistent with the optimal strategy that requires the DM to initially select the decision with the greatest variance of distribution of return, given that both decisions lead to an expected return of 0.

4.2 **Example 2**

As a second, slightly more practical example, consider the following scenario. When a DM enters into a long futures contract for a stock, she makes a commitment to purchase that stock at a fixed price (known as the strike price) at a fixed time in the future. When that fixed time is reached and the purchase is made, the DM has made a profit if the current market price exceeds the strike price, and a loss if not. Here a DM has been given the chance to enter into a long futures contract on stock A, and also the chance to enter into a long futures contract on stock B. There are two decision epochs, and she may only enter into (at most) one long futures contact in each epoch. We denote by d_A the decision to enter into a long futures contract on stock A, and by d_B the decision to enter into a long futures contract on stock B. She may also choose to enter into neither, a decision we denote by d_N. At the first epoch, a decision d_1 must be made. The time t, after which the DM must purchase the stock at the strike price, is assumed to be the length of an epoch, *e.g.*, if a DM chooses d_A at the first epoch, then she purchases stock A and makes either a profit or loss precisely before she must make her next decision. Hence after an epoch learning occurs about the unknown parameter of interest θ, which is related to stock performance and hence to the profit or loss made.

Suppose $\theta\sim\mathcal{N}(1,3)$, *i.e.*, the DM expects stock prices to be more than the strike price. We also have

- $R_i|d_A,\theta\sim\mathcal{N}(2\theta,2)$,
- $R_i|d_B,\theta\sim\mathcal{N}(\theta+1,3)$,

- $P(R_i = 0|d_N) = 1$.

The utility function of the DM over returns r_1 and r_2 resulting from decisions d_1 and d_2 respectively are given by the sum of exponential utility functions in Equation (19), the use of which is common in financial decision situations (see for instance Gerber & Pafumi, (1998) or Tsanakas & Desli, (2005)). The 0.9 multiplying the second term in the function pertains to a discounting rate, by which it is meant that returns received in the future are regarded as inferior to identical returns received in the present. The discounting rate is a measure of the degree to which the DM prefers returns now to those at some future time, with the discounting becoming more extreme over time. It is also noted that this form of utility function will yield negative utility values, but this is of no concern as utility functions are invariant to affine linear transformations, so the values may be translated to any desired region.

$$u(r_1, r_2) = -e^{-0.1r_1} - 0.9e^{-0.1r_2}. \tag{19}$$

This can be approximated by a multivariate Taylor expansion to ensure that it belongs to the polynomial utility class. This expansion is taken about the point $(2, 2)$, corresponding to the most likely values of r_1 and r_2 resulting from the prior distribution. This yields the approximation

$$\begin{aligned} u(r_1, r_2) \quad \approx \quad & -1.9005 + 0.0978r_1 + 0.0917r_2 \\ & -0.004r_1^2 - 0.0045r_2^2. \end{aligned} \tag{20}$$

The decision tree is given below (Fig. 2).

The solution to this problem is found in the same manner as in Example 1. In particular we have

$$\begin{aligned} u(r_1, r_2, \theta) = & \int \Big(-1.9005 + 0.0978r_1 + 0.0917r_2 \\ & -0.004r_1^2 - 0.0045r_2^2\Big) \prod_{i=1}^{2} P(r_i|d_i, \theta) dr_1 dr_2 \\ = & \ -1.9005 + 0.0978 \int r_1 P(r_1|d_1, \theta) dr_1 \\ + & \ 0.0917 \int r_2 P(r_2|d_2, \theta) dr_2 \\ - & \ 0.004 \int r_1^2 P(r_1|d_1, \theta) dr_1 \\ - & \ 0.0045 \int r_2^2 P(r_2|d_2, \theta) dr_2. \end{aligned}$$

Using the raw moments of the Normal distribution as before allows the calculation of the expected utility for possible decision paths, *e.g.*, $u(d_A, d_B, \theta) = -1.8348 + 0.2783\theta - 0.0205\theta^2$, while $u(d_A, d_N, \theta) = -1.9085 + 0.1956\theta - 0.016\theta^2$, *etc.* Again we use (13) and (14) to update beliefs about θ. For instance there are three conditional values at the first epoch, for example $\theta|d_1 = d_A, r_1 \sim \mathcal{N}(\frac{3r_1+1}{7}, \frac{3}{7})$, and seven at the second epoch, for example $\theta|d_1 = d_A, d_2 = d_B, r_1, r_2 \sim \mathcal{N}(\frac{3r_1+r_2}{8}, \frac{3}{8})$. Expected utility values are then found as in Example 1, but for illustrative purposes consider the expected utility having first made decision d_A, followed by d_B, and observing r_1 and r_2 respectively, then:

Fig. 2. Decision Tree for Example 2

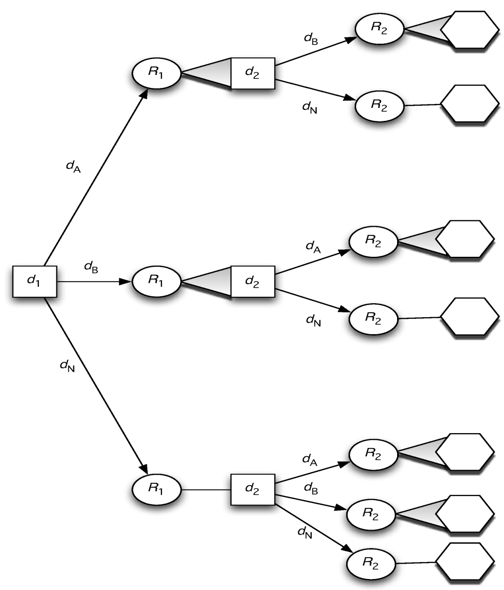

$$\begin{aligned} & E_{\theta|h_2, r_2, d_2 = d_B}[u(d_A, d_B, \theta)] \\ = & \int u(d_A, d_B, \theta) P(\theta|h_2, d_2 = d_B, r_2) d\theta \\ = & \int \big(-1.8348 + 0.2783\theta - 0.0205\theta^2\big) \\ \times & \quad P(\theta|h_2, d_2 = d_B, r_2) d\theta, \\ = & -1.8348 + 0.2783 \int \theta P(\theta|h_2, d_2 = d_B, r_2) d\theta \\ - & \ 0.0205 \int \theta P(\theta|h_2, d_2 = d_B, r_2) d\theta \\ = & -1.8438 + 0.1044r_1 + 0.034r_2 - 0.0028r_1^2 \\ - & \ 0.0019r_1 r_2 - 0.00032r_2^2 \end{aligned}$$

The predictive distribution of R_2 given the history up to that point is found by integration with respect to updated beliefs about θ, *i.e.*, $P(R_2|r_1, d_1, d_2) = \int P(R_2|d_2, \theta) P(\theta|r_1, d_1) d\theta$, allowing predicted expected values, *e.g.*, $E[R_2|r_1, d_1 = d_A, d_2 = d_B] = \frac{3r_1+1}{7}$, which in turn permit the expected utility equations in terms of r_1 only. This allows calculation of the expected utility, *e.g.*, when $d_1 = d_A$ and $d_2 = d_B$ then the expected utility is given by $-1.838 + 0.119r_1 - 0.0037r_1^2$.

Now considering the branch where d_A is chosen first, it is clear that d_B is preferred over d_N if and only if

$$\begin{aligned} & -1.838 + 0.119r_1 - 0.0037r_1^2 > \\ & -1.888 + 0.082r_1 - 0.0029r_1^2 \\ \iff & 0.050 + 0.037r_1 - 0.0008^2 > 0 \\ \iff & -1.31 < r_1 < 47.56 \end{aligned}$$

Hence:

$$\begin{aligned} & E_{R_1|d_1=d_A}\Big[\max\{f_1(r_1), f_2(r_2)\}\Big] \\ = & \int_{-\infty}^{-1.31} f_1(r_1)P(r_1|d_A)dr_1 \\ + & \int_{-1.31}^{47.56} f_2(r_1)P(r_1|d_A)dr_1 \\ + & \int_{47.56}^{\infty} f_1(r_1)P(r_1|d_A)dr_1, \end{aligned}$$

where

- $f_1(r_1) = -1.888 + 0.082r_1 - 0.0029r_1^2$,
- $f_2(r_1) = -1.838 + 0.119r_1 - 0.0037r_1^2$,
- $R_1|d_A \sim \mathcal{N}(2, 2)$.

This results in $E_{R_1|d_1=d_A}\Big[\max\{f_1(r_1), f_2(r_2)\}\Big] \approx -1.62$. A similar procedure for when $d_1 = d_B$ gives a value approximately equal to -1.66, and when $d_1 = d_N$ the expected utility is -1.66. Hence, performing the remainder of the calculations results in an optimal decision sequence of $\pi_1 = d_A$, followed by $\pi_2 = d_B$ if $-1.31 < r_1 < 47.56$. The DM has a slightly risk-averse utility function in this example, and hence it is unsurprising that she chooses the return with the smaller associated variance, given that both returns have (prior) equal means. Of course it should be noted that the upper bound for an observation r_1 before the decision $d_2 = d_A$ is no longer optimal can be explained by the fact that these values have negligibly small probability of occurring. Note that the derivative of $f_2(r_1)$ is negative for $r_1 > 16.1$, implying incoherent utility preferences, but the probability of seeing a value exceeding this threshold is less than 0.0001, *i.e.*, it is negligibly small.

In addition to the above computations we also conducted the analysis in this question using the original exponential utility function of Equation (19), rather than its polynomial approximation using the Taylor Series, as demonstrated in Equation (20). Use of this non-polynomial utility function meant we were unable to use the techniques derived in Section 3. Nevertheless it was possible to calculate a result (using numerical integration methods) which was that the optimal sequence was $\pi_1 = d_A$, followed by $\pi_2 = d_B$ if $r_1 > -1.308$. We see that it is essentially identical to the outcome determined above, with the expected utility also being -1.62 for buying stock A first. There is no upperbound on r_1 in this exponential case, while for its Taylor series approximation the upperbound was 47.56. The probability of witnessing a value above this bound is suitably smalll, yet significant enough to cause a minor difference in the finalised figures for expected utility in the two cases. Overall we see that both approaches garner the same outcome, but with the former, using polynomial utility, being significantly more tractable, and not requiring the use of numerical integration in the computations. Note that where we to continue to extend Equation (20) to a higher order Taylor expansion then the results in both cases would be indistinguishable.

5. DISCUSSION

As discussed in Section 2 Lindley's method of conjugate utility permits a closed form solution to a one-off decision problem. However, when extended to a sequential decision problem the solution is generally no longer available in closed form. As an alternative we have proposed the use of a polynomial utility class, which allows for tractable solutions in sequential decision problems when Normal distributions reasonably represent prior beliefs over the unknown decision parameter θ and decision returns R_i. Normal-Normal conjugacy ensures that subsequent utility values are in a closed form and are easily interpretable. Note that while on the surface the requirement for beliefs to follow a Normal distribution may seem like a somewhat restrictive facet of this methodology, the Normal distribution is a very prevalent one which occurs naturally in many realistic frameworks. Also furthering this claim is the Central Limit Theorem, which states that the sum of a suitably large amount of independent identically distributed random variables, having finite mean and variance, will converge to the Normal distribution, making it a commonly used limiting distribution in cases where the elements of interest are not themselves Normally distributed. Note that this methodology can naturally be extended to a multivariate Normal setting, as similar conjugacy to the univariate case is applicable here also.

It is noteworthy that the method outlined in this paper could potentially be extended to include the idea of adaptive utility, as seen in Cyert & DeGroot (1975) and Houlding & Coolen (2011). In adaptive utility there is uncertainty of a DM over her preferences, and she may learn about them over time. In sequential decision problems within this framework, the additional uncertainty further increases the computational complexity of a solution, which suffer greatly from the curse of dimensionality. The implementation of a method such as the polynomial utility class would be especially effective in this case, and the tractability of solutions would be all the more valuable, as it would greatly decrease computational cost.

An area for possible further exploration is in the use of sets or classes of utility functions to model imprecise utilities, discussed in Houlding & Coolen (2012). In cases of imprecise utility in sequential decision problems the tractability of computation is an even bigger issue then for single utility functions due to the need to track both a lower expected utility bound and an upper expected utility bound rather than a single expected utility value.

Also of future interest is whether this method could be extended beyond the Normal distribution to other members of the Exponential family or Stable family of distributions, or maybe to consider discretization, *e.g.*, using the Multinomial model with conjugate Dirichlet priors. Finally, we also mentioned that commonly used functions such as the exponential and logarithmic functions are not contained in the polynomial utility class, but may be approximated by Taylor series. Further study could be conducted to determine the level of accuracy of this approximation depending on the number of terms considered in the partial sum (as touched on at the end of Example 2) so as to allow a robust analysis and potential bounds on resulting errors.

ACKNOWLEDGEMENTS

The authors would like to thank the anonymous reviewers whose suggestions and additional references have helped in our presentation, and to the participants of GDRR 2013: Third Symposium on Games and Decisions in Reliability and Risk, for their feedback on an early presentation of this work.

REFERENCES

Bellman, R. (1952). On the Theory of Dynamic Programming. *Proceedings of the National Academy of Sciences*, 38(8), 716-719.

Bellman, R. (1957). *Dynamic Programming.* Princeton University Press.

Benjaafar, S., Morinb, T.L. & Talavageb J.J. (1995). The Strategic Value of Flexibility in Sequential Decision Making, *European Journal of Operational Research*, 82(3), 438-457.

Berry, S.M., Carlin, B.P., Lee, J.J. & Muller, P. (2010). *Bayesian Adaptive Methods for Clinical Trials*. Chapman and Hall.

Brennan, A., S. Kharroubi, A. O'Hagan & J. Chilcott (2007). Calculating Partial Expected Value of Perfect Information via Monte Carlo Sampling Algorithms. *Medical Decision Making*, 27, 448-470.

Brockwell, A. E. & Kadane, J.B. (2003). A Gridding Method for Solving Bayesian Sequential Decision Problems. *Journal of Statistical Computation and Graphics*, 12(3), 566-584.

Chang, C. T. (2011). Multi-choice Goal Programming with Utility Functions. *European Journal of Operational Research.* 215(2), 439-445.

Cox, L. A. Jnr. (2012). Confronting Deep Uncertainties in Risk Analysis. *Risk Analysis*, 32(10), 1607-29.

Cyert, R.M. & DeGroot, M.H. (1975). Adaptive Utility. *Adaptive Economic Models*, 223-246, eds. R.H. Day & T. Groves, Academy Press, New York.

Diamond, H. & Gelles, G. (1999). Gaussian Appromximation of Expected Utility. *Economic Letters*, 64, 301-307.

Gerber, H. U. & Pafumi, G. (1998). Utility Functions: From Risk Theory to Finance. *North American Actuarial Journal*, 2(3), 74-91.

Guler, I. (2007). Throwing Good Money after Bad? Political and Institutional Influences on Sequential Decision Making in the Venture Capital Industry. *Administrative Science Quarterly*, 52(2), 248-285.

Haung, Y., Chang, W., Li, W. & Lin, Z. (2013). Aggregation of Utility-based Individual Preferences for Group Decision-Making. *European Journal of Operational Research*

Hlawitschka, W. (1994). The Emperical Nature of Taylor Series Approximations to Expected Utility. *The American Economics Review*, 84(3), 713-719.

Houlding, B. & Coolen, F.P.A. (2011). Adaptive Utility and Trial Aversion. *Journal of Statistical Planning and Inference*, 141(2), 734-747.

Houlding, B. & Coolen, F.P.A. (2012). Nonparametric Predictive Utility Inference. *European Journal of Operational Research*, 221(1), 222-230.

Islam, A. F. M. S. (2011). Loss Functions, Utility Functions and Bayesian Sample Size Determination. Doctoral Thesis, Queen Mary, University of London.

Kahneman, D & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, Vol. 47(2), 263-291.

Keeney, R.L. (1972). Multiplicative Utility Functions. *Operations Research*, 22, 22-34.

LiCalzi, M. & Sorato, A. (2006). The Pearson system of utility functions, *European Journal of Operational Research*, 172(2), 560573.

Lindley, D.V. (1976). A Class of Utility Functions. *The Annals of Statistics*, 4(1), 1-10.

Markowitz, H. (2012). Mean-variance Approximations to Expected Utility. *European Journal of Operational Research.*

Muller, P. (1999). Simulation-Based Optimal Design. *Bayesian Statistics 6*, 459-474 (Valencia Volume, edited by Bernardo *et. al.*).

Musal, R. M., Soyer, R., McCabe, C. & Kharroubi, S. A. (2012). Estimating the Population Utility Function: A Parametric Bayesian Approach. *European Journal of Operational Research*, 218(2), 438-547.

Novick, M. R. & Lindley, D. V. (1978). The Use of More Realistic Utility Functions in Educational Applications. *Journal of Educational Measurement*, 15, 181-191.

Papoulis, A. (1991). *Random Variables and Stochastic Processes.* McGraw-Hill.

Raiffa, H. & Schlaifer, R. (1961). *Applied Statistical Decision Theory.* Division of Research, Graduate School of Business Administration, Boston.

Tsanakas, A. & Desli, E. (2005). Measurement and Pricing of Risk in Insurance Markets. *Risk Analysis*, 25(6), 1653-1668.

Von Neumann, J. & Morgenstern, O. (1947). Theory of Games and Economic Behaviour. Second Edition, Princeton University Press.

APPENDIX

In the univariate case, a probability distribution is said to be a member of the exponential family if, following a suitable parameterization, it can be expressed in the following form, where the function $H(r)$ is non-negative, and $G(\theta)$ is the normalising constant:

$$P(r|\theta) \quad = \quad G(\theta)H(r)e^{\theta r}. \tag{21}$$

The natural conjugate prior for such densities is then, for suitable hyper-parameters n_0 and r_0, of the following form, with normalising constant $K(n_0, r_0)$:

$$P(\theta|n_0, r_0) \quad = \quad K(n_0, r_0)G(\theta)^{n_0}e^{\theta r_0}. \tag{22}$$

Following a sample of independent and identically distributed values $(r_1, r_2, \ldots, r_n) = \boldsymbol{r}$, with $P(r_i|\theta)$ as in Equation (3), posterior beliefs over θ will be as follows:

$$\begin{aligned} P(\theta|\boldsymbol{r}, r_0, n_0) \quad &= \quad K\Big(n + n_0, \sum_{i=0}^{n} r_i\Big) \\ &\times \quad G(\theta)^{n+n_0}e^{\theta \sum_{i=0}^{n} r_i}. \end{aligned} \tag{23}$$

Lindley (1976) shows that if the utility of a decision d depends on the value of some uncertain parameter θ that has posterior distribution according to Equation (23), then there is a natural conjugate 'matched' utility function that takes the following form, with $F(d)$ a positive function:

$$U(d, r, \theta) \quad = \quad F(d)G(\theta)^{n(d)}e^{\theta r(d)}. \tag{24}$$

When determining the expected utility of decision d with respect to posterior beliefs over θ, Equation (24) leads to the following (with $r = \sum_{i=0}^{n} r_i$ and $N = n + n_0$):

$$\begin{aligned} U(d, r) &= \int F(d)K(N, r)G(\theta)^{N+n(d)}e^{\theta(r+r(d))}d\theta, \\ &= \quad \frac{F(d)K(N, r)}{K(N + n(d), r + r(d))}. \end{aligned} \tag{25}$$