

---

# Maximum Entropy Utility

---

Ali E. Abbas  
[aliabbas@stanford.edu](mailto:aliabbas@stanford.edu)

Department of Management Science and Engineering  
School of Engineering, Stanford University  
Terman Engineering Center 700-406  
Stanford, California 94305-4026

## Abstract

This paper presents a method to assign utility values when partial information is available about the decision maker's preferences. We introduce an analogy between probability and utility through the notion of a utility density function and illustrate the application of this analogy to the maximum entropy principle. The maximum entropy utility solution embeds a large family of utility functions that includes the most commonly used functional forms. We discuss the implications of "maximum entropy utility" on the preference behavior of the decision maker and present an application to competitive bidding situations where only previous decisions are observed by each party. We also present "minimum cross entropy utility" which incorporates additional knowledge about the shape of the utility function into the maximum entropy formulation, and work through several examples to illustrate the approach.

*Received* September 2002; revisions received April 2003, December 2003.

*(Subject Classifications: Decision Analysis: applications. Utility/Preference: estimation, multiattribute)*

# 1. INTRODUCTION

In this paper we present a method to assign utility values when only partial information is available about the decision maker's preferences. We assume in all of our analyses that the decision maker follows the axioms of normative utility theory (von Neumann and Morgenstern 1947) and in particular (1) can provide the complete preference order for the prospects (consequences) of a decision situation and (2) has transitive preferences. These requirements may seem difficult at first; however, we note that in many cases of decision analysis practice, such as when monetary prospects are involved, both transitivity and the complete order of the prospects are reasonable assumptions if the decision maker prefers more money to less.

We remind the reader that when a decision problem is deterministic, the order of the prospects is sufficient to determine the optimal decision alternative. However, when uncertainty is present, the von Neumann and Morgenstern utility values need to be assigned. Our approach starts with the ordinal preference of the prospects and ends with the assignment of cardinal utilities when partial preference information is available. By partial preference information we mean any information that does not include the utility values but does include the order of the prospects. Partial preference information includes knowing some utility values, observing previous decisions made by the decision maker, or even knowing bounds on the domain of the prospects. Partial preference information is often encountered in practice where (1) time or health constraints prevent complete elicitation of utility values; (2) the decision maker is unavailable or unwilling to assign utility values; (3) there is no single decision maker but rather a group, which may be able to reach consensus only on the preference order but not on the utility values (Kirkwood and Sarin 1985); and (4) in competitive bidding situations where a decision maker is trying to infer the utility values of others by observing their decisions.

In our search of the literature we have found some related work that ranks multicriteria decision alternatives using additive value functions when only the rank order of the attributes is available. Butler, Jia and Dyer (1997) use simulation and joint sensitivity analysis to select the optimal decision alternative when only the rank order is available; Rao and Sobel (1980) use the rank order to derive a marginal distribution for the  $k^{th}$  largest weight; Barron and Barret (1996) compare three approximate formulas to estimate the weights; and Jessop (1999) uses normalized attributes and a maximum entropy formulation to determine the weights. In comparison, our formulation applies to utility functions (not to value functions) and makes no assumptions about the structure of the utility function or the value function being additive.

The core idea of our approach uses a utility function that is normalized to range from zero to one. We define a utility density function as the derivative of a normalized utility function. Based on this definition, a utility density function has two main properties: it is non-negative and integrates to unity. These two properties form the basis of an analogy between probability and utility that transfers many tools from one domain into the other. In this paper, we build on this analogy to assign utility values with partial preference information.

The utility-probability analogy that we develop in this paper has not been seen in our search of the literature. Berhold (1973) rescales probability distributions to obtain convenient expressions for utility functions but he does not introduce this analogy. Castagnoli and LiCalzi (1996) interpret a normalized utility function as a probability distribution of an uncertain target that is independent of the lotteries faced by the decision maker. In contrast we do not interpret utility values as describing anything other than preferences and we interpret the normalized utility function as simply representing the preferences of the decision maker using the von Neumann and Morgenstern approach. Our work thus preserves the separation of beliefs about the likelihood of events from preferences over the results of those events.

The remainder of this paper is organized as follows. Section 2 presents several definitions on the analogy between probability and utility. Section 3 provides an interpretation for the entropy of a utility function in both the discrete and continuous cases, and proposes the maximum entropy utility principle to assign utility values based on partial preference information. Section 4 presents the maximum entropy utility solution that includes the most commonly used forms of utility functions. Section 5 presents several applications of maximum entropy utility, and Section 6 presents minimum cross entropy utility, where additional knowledge about the shape of the utility function can be incorporated.

## **2. UTILITY – PROBABILITY ANALOGY**

The analogy between probability and utility appears naturally in the probabilistic equivalence used in the von Neumann and Morgenstern utility assessments. Recall that when eliciting the *utility* value of a prospect,  $B$ , we have three ordered prospects,  $A \succ B \succ C$ , and we are indifferent between receiving  $B$  for sure and a binary gamble with a *probability*,  $U_B$ , of yielding  $A$ , and a *probability*  $1-U_B$  of yielding  $C$ . Howard (1992) observes this correspondence and suggests that the von Neumann and Morgenstern *utility* be called a “*preference probability*”.

In this paper, we present new definitions in both the discrete and continuous cases that highlight the analogy between probability and utility, and translate many tools from one domain into the other.

### **2.1. Discrete Case: Utility Vector and Utility-Increment Vector**

The first definition is a utility vector for a set of  $K$  ordered prospects. A utility vector contains the utility values of the prospects starting from lowest to highest. We assume that there is at least one prospect, which has strict preference to exclude the case of absolute indifference between the  $K$  prospects. With no loss of generality, we assign a utility value of zero to the least

preferred prospect,  $u_0$ , and a utility value of one to the most preferred prospect,  $u_{K-1}$ . The utility vector has  $K$  elements defined as

$$U \triangleq (u_0, u_1, u_2, \dots, u_{K-2}, u_{K-1}) = (0, u_1, u_2, \dots, u_{K-2}, 1). \quad (1)$$

Note that any utility vector of dimension  $K$  can be represented as a point in a  $(K-2)$ -dimensional space in the region defined by  $0 \leq u_1 \leq u_2 \leq \dots \leq u_{K-3} \leq u_{K-2} \leq 1$ . This region, which we call the utility volume, has a volume equal to  $\frac{1}{(K-2)!}$ .


The second definition is a utility-increment vector,  $\Delta U$ , whose elements are equal to the difference between the consecutive elements in the utility vector. The utility-increment vector has  $(K-1)$  elements defined as

$$\Delta U \triangleq (u_1 - 0, u_2 - u_1, \dots, 1 - u_{K-2}) = (\Delta u_1, \Delta u_2, \Delta u_3, \dots, \Delta u_{K-1}). \quad (2)$$

The coordinates of  $\Delta U$  have two main properties: they are all greater than or equal to zero and sum to one. Therefore, any utility-increment vector can be represented as a point in a  $(K-1)$ -dimensional simplex  $\left\{ x : \sum_{i=1}^{K-1} x_i = 1, x_i \geq 0 \right\}$ . We will refer to this simplex as the utility simplex. A

graphical illustration of a three-dimensional utility simplex is shown in Figure 1.

**Figure 1.** A 3-dimensional utility simplex.



The utility simplex provides the space of all possible utility values for the given preference order. The geometric representation of the utility simplex presented above and the two main properties of the utility increment vector form the basis of an analogy between probability and utility that is the basic premise of this paper.

## 2.2. Utility Assignment for Discrete Ordered Prospects

Now let us consider the following problem: a decision maker provides the preference order for a set of  $K$  prospects. If a decision analyst would like to assign utility values on behalf of the decision maker (or if a decision maker would like to infer another person's utility values) based on this preference order alone, what utility values should s/he assign? To answer this question, we observe that any point in the utility simplex satisfies the decision-maker's preference order of the prospects but assigns different utility values to them. In other words, knowledge of the preference order alone tells us nothing about the location of the utility increment vector over the utility simplex. If all we know about the prospects is their ordering, it is reasonable to assume, therefore, that the location of the utility increment vector is uniformly distributed over the utility simplex. This assumption gives equal likelihood to all utility values that satisfy the decision maker's preference order, and adds no further information about the location of the utility increment vector other than knowledge of the order of the prospects.

From a mathematical point of view, the assumption of a uniform distribution for the location of the utility increment vector over the utility simplex implies that its location is described by a Dirichlet distribution whose  $K-1$  parameters are all equal to one. Furthermore, properties of Dirichlet distributions suggest that the marginal probability density function for each element of the utility increment vector is the Beta density,  $Beta(1, K-2)$ , while the

marginal probability density for each element of the utility vector,  $U_j$ ;  $j=1,\dots,K-2$ , is  $Beta(j, K-j-1)$  (Degroot 1970).

The previous analysis treats a decision maker's unknown utility values as random variables from the decision analyst's viewpoint (except for the most preferred and least preferred prospects which have values of  $u_{K-1} = 1$  and  $u_0 = 0$  respectively). The analysis uses the preference order to derive a marginal probability density for each utility value. The problem that we seek to solve, however, is the assignment of utility values to all the prospects given the preference order. Fortunately, the remaining part of the problem, which assigns a utility value given its marginal probability density, is a relatively easy task and has had a large share of literature coverage. For example, Howard (1970) shows that the mean of a random variable is a natural assignment given its marginal distribution. We summarize this result for the utility increment vector below.

### **Utility Increment Assignment given the Preference Order**

When only the preference order is known, the marginal probability density for the increments in utility values of  $K$  ordered prospects is  $Beta(1, K-2)$ . The utility increment assignment is the mean of this distribution and is equal to  $\frac{1}{K-1}$ .

The previous result shows a method to assign utility values for a decision maker when only the preference order of the prospects is known. This assignment produces equal increments in utility values. In Section 3, we extend the analysis further and present a method to assign utility values given the preference order and any other information that may be available about the decision maker's preferences.

### 2.3. Continuous Case: Utility Functions and Utility Density Functions

Now we extend the previous definitions to the continuous case where the number of prospects,  $K$ , is infinite. We start with prospects of a decision situation, which have only one attribute,  $x$ , and discuss the case of multiple attributes in Section 5. A common example of one-attribute prospects in the continuous case is monetary outcomes over a continuous domain.

In the continuous case, the utility vector is a utility curve,  $U(x)$ , over the given domain and is normalized to have values between zero and one. The normalized utility curve has the same mathematical properties as a cumulative probability distribution as both are non-decreasing and range from zero to one. The normalization of the utility function poses no major limitations to von Neumann and Morgenstern utility values that are also bounded and range from zero to one.

The utility increment vector is now the derivative of the normalized utility curve (assuming the derivative exists) and we refer to it as a utility density function,  $u(x)$ , i.e.

$$u(x) \triangleq \frac{d}{dx}U(x) \quad (3)$$

If the utility curve is normalized, then the utility density integrates to unity. The utility density function is non-negative, due to the non-decreasing values of the utility curve, and thus has the same mathematical properties as a probability density function: both are non-negative and integrate to unity.

Note that the utility value,  $U(x)$ , of a given prospect,  $x$ , can be determined by integrating the utility density from the least preferred prospect,  $x_{\min}$ , (or the lower bound of the monetary prospects) up to that prospect,  $x$ . i.e.

$$U(x) = \int_{x_{\min}}^x u(x)dx. \quad (4)$$

We summarize the definitions and analogy between probability and utility in Table 1.

**Table 1.** Utility-probability analogy.

Probability		Utility	
Probability Mass Function	$P = (p_1, \dots, p_K)$ $p_i : \sum_{i=1}^K p_i = 1, p_i \geq 0$	Utility Increment Vector	$\Delta U = (\Delta u_1, \dots, \Delta u_{K-1})$ $\Delta u_i : \sum_{i=1}^{K-1} \Delta u_i = 1, \Delta u_i \geq 0$
Discrete Cumulative Probability	$P_j : \sum_{i=1}^j p_i, j = 1, \dots, K$ $P_j - P_{j-1} \geq 0$	Discrete Utility Value	$U_j : \sum_{i=1}^j \Delta u_i, j = 1, \dots, K-1$ $U_j - U_{j-1} \geq 0$
Probability Density Function	$f(x) \triangleq \frac{d}{dx} F(x)$ $\int_a^b f(x) dx = 1, f(x) \geq 0$	Utility Density Function	$u(x) \triangleq \frac{d}{dx} U(x)$ $\int_a^b u(x) dx = 1, u(x) \geq 0$
Cumulative Distribution Function	$F(x) = \int_{x_{\min}}^x f(x) dx$ $0 \leq F(x) \leq 1, \frac{d}{dx} F(x) \geq 0$	Utility Function	$U(x) = \int_{x_{\min}}^x u(x) dx$ $0 \leq U(x) \leq 1, \frac{d}{dx} U(x) \geq 0$

## 2.4. Utility Assignment by Analogy with Probability Assignment

The utility-probability analogy translates many tools from one domain into the other. To demonstrate one example, we refer to the problem of assigning a probability to the outcome of an uncertain event in the absence of perfect information. This problem dates back to Laplace’s “principle of insufficient reason”.

Laplace suggested that we assign *equal probabilities* to all outcomes unless there is *information* that suggests otherwise. If we apply the utility-probability analogy to Laplace’s principle of insufficient reason, we have a method for assigning utility values that can be expressed as follows: when only the preference order of the prospects is available, we assign *equal increments in utility values* unless there is *preference information* that suggests otherwise.

The utility assignment suggested by this result agrees with the intuitive assignment a decision analyst would make when only the order of the prospects is known. The rationale is that if we know only the order of the prospects, there should be no reason for one increment in utility

values to be larger than the other unless there is preference information that suggests otherwise. Assigning unequal increments implies additional information about the decision maker's preferences that is not included in the preference order alone.

In this example, the application of the utility-probability analogy to a well-known problem in probability resulted in a methodology for assigning utility values that agrees with our intuition and with the mathematical results of the uniform Dirichlet distribution over the utility simplex. In the next section, we present another application of the utility-probability analogy to measure the spread of the utility increment vector and the utility density function.

### **3. THE ENTROPY OF A UTILITY FUNCTION**

#### **3.1. Entropy Measure for Discrete Prospects**

Shannon (1948) introduced the term  $H(p_1, \dots, p_n) = -\sum_{i=1}^n p_i \log(p_i)$  as a measure of uncertainty about a discrete random variable having a probability mass function,  $p$ . He called this term the entropy. Shannon's entropy term is also a measure of the spread of a probability distribution that achieves its maximum value when the distribution assigns equal probabilities to all outcomes. Building on this idea, Jaynes (1957) proposed the use of a prior probability distribution that maximizes Shannon's entropy measure (has maximum spread) and satisfies the partial information constraints when no further information is available. Jaynes' proposition is considered to be an extension of Laplace's principle of insufficient reason, as it incorporates additional information, and is known as the maximum entropy principle. It has found wide use in the assignment of prior probabilities using partial information.

It is natural to extend our analogy by considering Shannon's entropy definition as a measure of spread for the coordinates of the utility increment vector

$$H(\Delta u_1, \Delta u_2, \Delta u_3, \dots, \Delta u_{K-1}) = -\sum_{i=1}^{K-1} \Delta u_i \log(\Delta u_i). \quad (5)$$

If we take the first partial derivative of equation (5) with respect to  $\Delta u_i$  and equate it to zero, we find that this measure achieves its maximum value when the utility increments are all equal. In other words, the utility increment vector that maximizes this entropy measure has the same utility increments as those described by the uniform Dirichlet distribution over the utility simplex. Maximizing the entropy of the utility increment vector with certain preference constraints yields a utility vector that satisfies the given constraints and produces (whenever the constraints allow) equal increments in utility values.

There are other measures that can be used for the spread in the utility increment vector but the entropy measure uniquely satisfies three essential axioms that were proposed by Shannon. We discuss these axioms as they relate to a measure of spread for the utility increment vector below.

(1) The measure of spread of the utility increment vector is a monotonically increasing function of the number of prospects,  $K$ , when the utility increments are equal. The rationale for this axiom is that the larger the number of prospects with equal utility increments the wider is the spread of the utility increment vector.

(2) The measure of spread of a utility increment vector should be a continuous function of the increments. If one of the utility increments changes slightly, the measure of spread should not change abruptly but should change in accordance with the corresponding change in spread.

(3) The order in which we calculate the measure of spread should not matter. For example, if we calculate the spread of a utility increment vector directly using equation (5), or if we calculate the spread of subsets of the utility increment vector separately then take a weighted average, we should get the same result,

$$H(\Delta u_1, \Delta u_2, \Delta u_3) = H(\Delta u_1, (\Delta u_2 + \Delta u_3)) + (\Delta u_2 + \Delta u_3)H\left(\frac{\Delta u_2}{\Delta u_2 + \Delta u_3}, \frac{\Delta u_3}{\Delta u_2 + \Delta u_3}\right). \quad (6)$$

For example, if we have a utility increment vector  $\Delta U = (0.25, 0.5, 0.25)$ , we can calculate its entropy directly,  $H(0.25, 0.5, 0.25) \triangleq -0.25 \log(0.25) - 0.5 \log(0.5) - 0.25 \log(0.25) = 1.5 \log(2)$ . If the last two elements ( $\Delta u_2$  and  $\Delta u_3$ ) are combined, they have a weight of  $(\Delta u_2 + \Delta u_3) = 0.75$  and together they form a utility increment vector that is normalized as  $\Delta U_{23} = (\frac{2}{3}, \frac{1}{3})$ . The original increment vector now reduces to two co-ordinates,  $\Delta U_R = (0.25, 0.75)$ . The entropy of  $\Delta U_R$  is therefore less than that of  $\Delta U$ , but when we add the weighted entropy due to the spread in  $\Delta U_{23}$ , we have  $H(0.25, 0.75) + 0.75H(\frac{2}{3}, \frac{1}{3}) = 1.5 \log(2)$ . Both methods thus provide the same spread. The entropy measure of equation (5) is the only measure that satisfies these three axioms.

### 3.2. Entropy Measure for Continuous Prospects

Now we discuss the differential form of the entropy expression,  $h(x)$ , when applied to a utility density function on a domain,  $[a, b]$

$$h(u(x)) = -\int_a^b u(x) \ln(u(x)) dx. \quad (7)$$

As shown in the Appendix, if we take the derivative of equation (7) with respect to  $u(x)$  and equate it to zero, the corresponding utility density is uniform over the bounded domain,  $u(x) = \frac{1}{(b-a)}$ , and the differential entropy has a maximum value of  $\ln(b-a)$ . The uniform utility density integrates to a linear (risk neutral) utility function. Any other utility density has less spread with this entropy measure. To gain some further intuition about the implications of the maximum entropy and minimum entropy assignments, let us consider the following example.

### Example 1: Entropy of the CARA Utility

Consider the following constant absolute risk aversion (CARA) utility density over the domain  $[0, 1]$

$$u(x) = \frac{\gamma e^{-\gamma x}}{1 - e^{-\gamma}}, \quad 0 \leq x \leq 1, \quad (8)$$


where  $\gamma$  is the decision maker's risk aversion coefficient.

By direct integration and the use of equation (7), the differential entropy is

$$h(u(x)) = 1 + \frac{\gamma e^{\gamma}}{1 - e^{\gamma}} - \ln\left(\frac{\gamma}{e^{\gamma} - 1}\right). \quad (9)$$

Figure 2 plots the differential entropy vs. the risk aversion,  $\gamma$ , from equation (9).

**Figure 2.** Sensitivity analysis of the differential entropy to the risk aversion coefficient.



From the results of Figure 2, we can see that the differential entropy has a unique maximum that occurs when  $\gamma = 0$ . Using L'Hopital's formula, we can show that when  $\gamma \rightarrow 0$

$$u(x) = \frac{\gamma e^{-\gamma x}}{1 - e^{-\gamma}} \rightarrow 1, \quad 0 \leq x \leq 1. \quad (10)$$

The maximum entropy utility density of equation (10) integrates to a linear utility function,  $U(x) = x$ ,  $0 \leq x \leq 1$ , that exhibits risk neutral behavior. The unique maximum entropy utility solution is both concave and convex, and favors no direction of risk attitude (as it occurs at the boundary of the two domains). From Figure 2, we note that the entropy is symmetric around  $\gamma = 0$ . Therefore the entropy of a risk averse utility function with  $\gamma = 3$  (for example) is the same as that of a risk seeking utility function with  $\gamma = -3$ .

The differential entropy has no lower bound since  $h(u(x)) \rightarrow -\infty$  as  $\gamma \rightarrow +\infty$  (the case of extreme risk averse behavior) and the utility density approaches an impulse density,  $\delta(x)$

$$\frac{\gamma e^{-\gamma x}}{1 - e^{-\gamma}} \rightarrow \delta(x) \text{ as } \gamma \rightarrow +\infty. \quad (11)$$

The impulse utility density integrates to a step (aspiration) utility function that jumps at the lower bound of the domain. The step utility function implies both extreme risk averse behavior and a steep change in preferences at  $x = 0$ .

As  $\gamma \rightarrow -\infty$  (extreme risk seeking behavior),  $h(u(x)) \rightarrow -\infty$ , (again) and the utility density approaches an impulse density,  $\delta(x-1)$ , at  $x=1$ .

$$\frac{\gamma e^{-\gamma x}}{1 - e^{-\gamma}} \rightarrow \delta(x-1) \text{ as } \gamma \rightarrow -\infty. \quad (12)$$

From equations (11) and (12) we observe that both cases of extreme risk averse and risk seeking behavior correspond to minimum entropy solutions. These solutions imply more about the decision maker's preferences than only the order of the prospects as they also favor one direction of risk attitude over the other. Furthermore any impulse function,  $\delta(x-x_0)$ ,  $0 \leq x_0 \leq 1$ , is also a minimum entropy solution that implies a steep change of preferences at the prospect,  $x_0$ . The maximum entropy solution, on the other hand, makes no assumptions about steep changes in the decision maker's preferences.

To summarize the results of this section, the maximum entropy assignment produces equal increments in utility values for the discrete case, and makes no assumptions about the direction of the risk attitude or about steep changes in preferences for the continuous case (unless this information is explicitly incorporated into the constraints as we shall see in Section 5).

### **3.3. The Maximum Entropy Utility Principle**

Based on the previous results, we are now ready to answer the following question. “Given the partial preference information we know about the decision maker, there may be several utility values that satisfy the given preference information constraints. What is the unbiased assignment of utility values that we should make?” By “unbiased” utility values, we mean those that do not lead to arbitrary assumptions of preference information that is not available. For example, the assignment of either risk averse or risk seeking behavior to a decision maker is a biased assignment unless there is preference information to support it, and the assignment of a non-uniform distribution over the utility simplex is a biased assignment when only the order of the prospects is available as it gives a set of utility values more likelihood than others.

To answer the utility assignment question, we propose the following maximum entropy utility principle:

“In making inferences on the basis of partial preference information, we use the utility curve (or utility vector) whose utility density function (or utility increment vector) has maximum entropy subject to whatever preferences are known”.

This method of assigning utility values provides an analogy with Jaynes’ maximum entropy principle for probability inference. It can be applied to both the continuous and the discrete utility forms. We call the utility values obtained from this principle the maximum entropy utility. In the next section, we discuss the maximum entropy utility solution given preference

information constraints and present some common forms of utility functions that this solution provides.

#### 4. THE MAXIMUM ENTROPY UTILITY FAMILY

The maximum entropy utility solution for constraints,  $\int_a^b h_i(x)u(x)dx = \mu_i$ ,  $i = 1, \dots, n$ ; with

$\int_a^b u(x)dx = 1$  and  $u(x) \geq 0$ , is shown in the Appendix as

$$u_{\maxent}(x) = e^{-\alpha_0 - 1 - \alpha_1 h_1(x) - \alpha_2 h_2(x) - \dots - \alpha_n h_n(x)}, \quad (13)$$

where  $u_{\maxent}(x)$  is the maximum entropy utility solution,  $[a, b]$  is the domain of prospects,  $h_i(x)$  is a given preference constraint,  $\mu_i$ 's are a given sequence of utility values or moments of the utility function, and  $\alpha_i$  is the Lagrange multiplier for each utility value or moment constraint.

The first application of the maximum entropy formulation is that it provides us with a general expression for utility functions that includes the most commonly used functional forms. For example, the risk neutral utility function, which has a uniform utility density, is a special case of equation (13) where the constraints,  $h_i(x)$ , are equal to zero. When  $h_1(x) = x$  and the remaining constraints are zero, the maximum entropy utility is a CARA utility on the positive domain. When  $h_1(x) = x$  and  $h_2(x) = x^2$ , the maximum entropy utility is a Gaussian utility density, which integrates to an S-shaped prospect theory utility function on the real domain (Kahneman and Tversky, 1979).

The maximum entropy utility solution also embeds the hyperbolic absolute risk averse utility function (HARA), which has the form

$$U(x) = \frac{1}{1-\gamma} (\gamma(\beta + \frac{\alpha}{\gamma}x)^{1-\gamma} - 1), \quad (14)$$

and a utility density of the form

$$u(x) = \alpha \left( \beta + \frac{\alpha}{\gamma} x \right)^{-\gamma} = e^{\frac{\ln(\alpha) - \gamma \ln(\beta + \frac{\alpha}{\gamma} x)}{\gamma}}, \quad (15)$$

where  $\alpha, \beta$ , and  $\gamma$  are given constants. The HARA utility function reduces to a risk neutral utility function when  $\gamma = 0$ ; to a CARA utility function when  $\gamma \rightarrow \pm \infty$ ; and to a constant relative risk averse utility function (CRRA) when  $\beta = 0$  and  $\gamma > 0$ . Comparing equations (13) and (15), shows that HARA utility functions can be expressed by the maximum entropy utility solution when  $h_1(x) = \ln(\beta + \frac{\alpha}{\gamma} x)$  and the remaining constraints are zero.

#### 4.1. Maximum Entropy Risk Aversion

Using equation (13), and Arrow - Pratt's definition of local risk aversion (Pratt 1964 and Arrow 1965), the maximum entropy utility function with constraints  $h_i(x)$ ,  $i = 0, 1, \dots, n$  has a risk aversion,  $\gamma_{\maxent}(x)$ , of

$$\gamma_{\maxent}(x) = -\frac{d}{dx} \ln(u_{\maxent}(x)) = \alpha_1 h_1'(x) + \alpha_2 h_2'(x) + \dots + \alpha_n h_n'(x), \quad (16)$$

where  $h_i'(x) = \frac{d}{dx} h_i(x)$ .

Equation (16) shows the linear effect contributed by the derivative of each preference constraint on the overall risk aversion function. Equation (16) also shows the wide range of risk aversion expressions that can be modeled by the maximum entropy utility family.

## 5. APPLICATIONS OF MAXIMUM ENTROPY UTILITY

In this section, we discuss several applications of the maximum entropy utility principle to infer utility values in practice using partial preference information.

## 5.1. Knowledge Of Some Utility Values

When eliciting a utility curve in practice, we often start by eliciting the utility values for some of the prospects. If we know only some utility values and would like to assign a utility function over a continuous domain, we solve for the maximum entropy utility density function subject to the given utility values. We illustrate this application through the following example.

### Example 2: The Party Problem

The party problem, introduced by Ronald Howard at Stanford University, can be summarized as follows: Kim is interested in having a party. She has three alternatives: Indoors, Outdoors, and on the Porch. However, she is uncertain about the weather situation, which can be sunny or rainy. She orders the prospects from best to worst, and assigns utility values and dollar equivalents to the prospects she is facing. These values are shown in Table 2.

**Table 2.** Kim's utility values for some of the monetary prospects.

Prospect	Dollar Value (\$)	Utility Values
Outdoors, Sunny	100	1
Porch, Sunny	90	0.95
Indoors, Rainy	50	0.67
Indoors, Sunny	40	0.57
Porch, Rainy	20	0.32
Outdoors, Rainy	0	0

Kim has a CARA utility; but let us assume that this information is not provided to the decision analyst. Now we would like to determine her continuous maximum entropy utility function over the domain of monetary prospects she is facing. The maximum entropy formulation for the utility density is

$$u_{\maxent}(x) = \arg \max \left( - \int_0^{100} u(x) \ln(u(x)) dx \right)$$

such that

$$\int_0^{20} u(x) dx = 0.32, \int_0^{40} u(x) dx = 0.57, \int_0^{50} u(x) dx = 0.67, \tag{17}$$

$$\int_0^{90} u(x) dx = 0.95, \int_0^{100} u(x) dx = 1, u(x) \geq 0.$$

If we compare the preference constraints,  $h_i(x)$ , of equation (17) to those of (13), we find they are in effect indicator functions over certain intervals. For example, the constraint

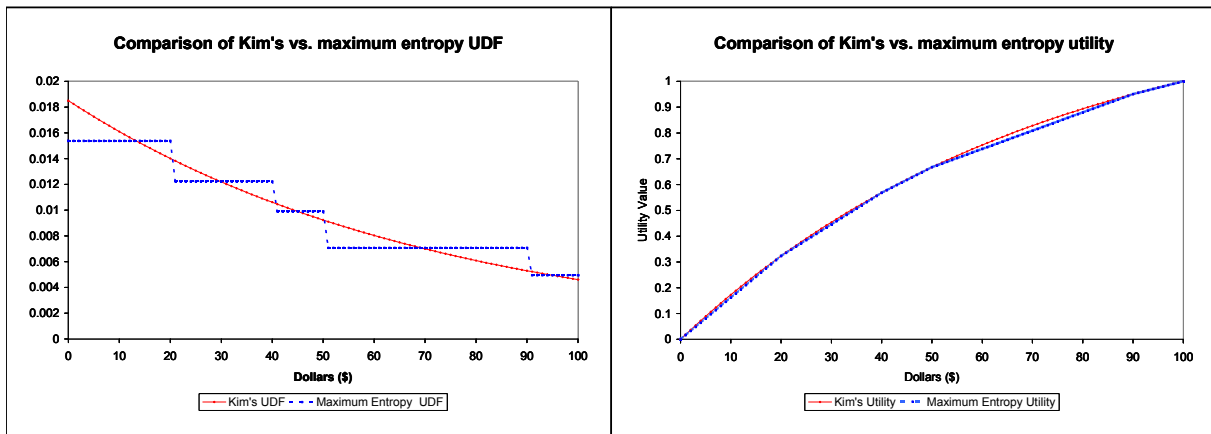
$$\int_0^{20} u(x) dx = \int_0^{100} I_{20}(x) u(x) dx = \int_0^{100} h_1(x) u(x) dx, \text{ where } I_{20}(x) \text{ is an indicator function for } x \in [0, 20].$$

From (13), we can see that the solution to this problem has the form

$$u_{\maxent}(x) = e^{-\alpha_0 - 1 - \alpha_1 I_{20}(x) - \alpha_2 I_{40}(x) - \alpha_3 I_{50}(x) - \alpha_4 I_{90}(x)}, \quad 0 \leq x \leq 100 \tag{18}$$

Equation (18) is the staircase utility density shown in Figure 3(a) together with Kim's CARA utility density (not known). In Figure 3(b), we compare the corresponding maximum entropy utility function to Kim's CARA utility. The maximum entropy utility function is piecewise linear connecting the given utility values. It exhibits risk neutral behavior over certain ranges and risk averse behavior over the whole domain.

**Figure 3.** Comparison of both Kim's and the maximum entropy utility function.



One question that may arise in practice here is the rationale for using the maximum entropy utility function rather than finding the best curve fit for the utility assessments provided. The basic motivation for the maximum entropy approach is that curve-fitting methods assume a certain structure (e.g. the utility function used or the order of the splicing polynomial). The fitted utility function will depend on the structure that is chosen for the fit. The maximum entropy approach, however, provides a unique utility function that makes no assumptions about the structure unless there is preference information to support it. We note that equation (17) does not incorporate any information about Kim's risk attitude over the sub-intervals. In Section 6 we will refer back to this example and incorporate risk aversion into Kim's formulation.

## 5.2. Inferring Utility Values by Observing Decisions

We now apply the maximum entropy utility principle to infer a decision maker's utility function by observing previous decisions. We assume that the decision maker maximized her expected utility in making these decisions and that the lotteries she was facing are known. If the decision maker prefers a lottery with cumulative distribution  $F(x)$  to a lottery  $G(x)$  we add an additional inequality constraint that the expected utility of  $F(x)$  is greater than or equal to that of  $G(x)$ . Using the rule of integration by parts we can show that

$$\begin{aligned}
& \int_a^b U(x)dF(x) \geq \int_a^b U(x)dG(x) \\
\Rightarrow & U(x)F(x) \Big|_a^b - \int_a^b F(x)u(x)dx \geq U(x)G(x) \Big|_a^b - \int_a^b G(x)u(x)dx \quad (19) \\
\Rightarrow & \int_a^b [G(x) - F(x)]u(x)dx \geq 0,
\end{aligned}$$

where  $U(x)F(x) \Big|_a^b = U(x)G(x) \Big|_a^b = 1$  due to the use of a normalized utility function.

The maximum entropy utility formulation becomes

$$u_{\maxent}(x) = \arg \max \left( - \int_a^b u(x) \ln(u(x)) dx \right)$$

*such that*

$$\int_a^b u(x) dx = 1, u(x) \geq 0 \tag{20}$$


$$\int_a^b [G(x) - F(x)] u(x) dx \geq 0.$$

This problem can be solved using the Karush-Kuhn-Tucker optimality conditions or by discretization and a numerical optimization package. The use of an equality constraint in equation (19) provides at least one feasible solution to this problem in the form of equation (13), where  $h_1(x) = [G(x) - F(x)]$ . The feasible region is convex due to the linear inequality constraints, and the concavity of the entropy expression provides a unique maximum entropy solution over the set of feasible solutions. To demonstrate an application of this formulation we consider the following example.

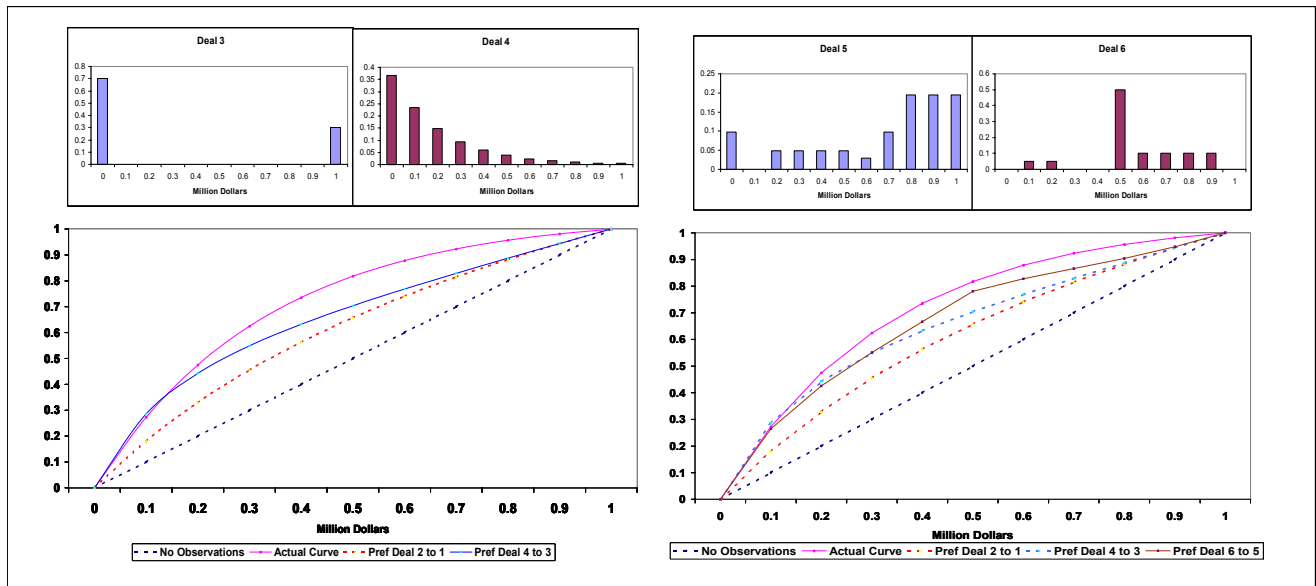
**Example 3:** A decision maker with an exponential CARA utility function and a risk tolerance of \$300,000 faces a deal whose prospects range from \$0 to \$1 Million.

Let us assume that an observer is trying to infer the decision maker's utility function, and that the only information that is available to him is the domain of monetary prospects that the decision maker is facing. As explained above, the maximum entropy utility solution is risk neutral over this domain. Figure 4 (a) shows the maximum entropy utility function and the decision maker's utility function. If the decision maker faces the two lotteries of Figure 4 (b) and prefers lottery 2 to lottery 1, an additional inequality constraint can be added into the maximum entropy formulation as described above. Figure 4(b) shows the effect of observing this decision on the maximum entropy utility function. Figures 5(a) and 5(b) show how the maximum entropy utility function is updated after observing more decisions made by the decision maker.

**Figure 4.** (a) Exponential vs. Maximum Entropy Utility Curve. (b) Two lotteries faced by decision maker and updated utility function.



**Figure 5.** (a) Second Observation (deal 4 is preferred to deal 3). (b) Third Observation (deal 6 is preferred to deal 5)



### 5.3. Maximum Entropy Multiattribute Utility

When the decision situation has multiple attributes, a value function is constructed to rank order the prospects, and a utility function is assigned over the value function to represent the decision maker's risk attitude towards value (for more details on this method see Matheson and

Howard 1968, Dyer and Sarin 1979 and 1982, and Keeney and Raiffa 1976). Using this approach, a maximum entropy multiattribute utility function can be constructed with partial preference information using a utility assessment over the value function in the maximum entropy formulation. The following example, adapted from (Howard, 1980) illustrates this approach.

#### **Example 4: Utility Function for Health State and Consumption**

A decision maker facing prospects of medical surgery provides a value function over two attributes: consumption,  $x$ , and health state,  $y$ . The health state is a disability level normalized from 0 (instant painless death) to 1 (current health with no disability). The value model over consumption and health states is a Cobb-Douglas function given as

$$V(x, y) = yx^\eta, \quad 0 \leq x \leq 1, \quad 0 \leq y \leq 1, \quad (21)$$

where  $x$  is in millions of dollars,  $y$  is the health state, and  $\eta$  is the trade-off coefficient.

Now we assign a utility function over the value function. If all we know about the prospects is the domain of their attributes, the maximum entropy assignment produces a uniform utility density over value and a corresponding linear (risk neutral) utility function over the value model,

$$U_{\text{maxent}}(V(x, y)) = V(x, y) = yx^\eta, \quad 0 \leq x \leq 1, \quad 0 \leq y \leq 1. \quad (22)$$


The marginal utility functions for the individual attributes that correspond to this maximum entropy assignment are

$$U(x) = x^\eta, \quad 0 \leq x \leq 1 \quad \text{and} \quad U(y) = y, \quad 0 \leq y \leq 1. \quad (23)$$

In this example, the maximum entropy utility function is risk neutral over value, so the risk attitude towards each attribute is determined by the value function. The decision maker is risk neutral over health states but his utility function for consumption depends on the trade-off coefficient,  $\eta$ : the decision maker is risk averse for consumption if  $\eta < 1$  and risk seeking if

$\eta > 1$ . Figure 6 shows the isopreference contours and the maximum entropy utility surface when only the bounds on the domain of the attributes are available. If additional information is available (such as utility assessments) it can also be incorporated into the formulation as described above.

**Figure 6.** (a) Utility Contour Plot. (b) Maximum Entropy Utility Surface for  $\eta = 0.7$ .



## 6. MINIMUM CROSS ENTROPY UTILITY

In many situations we may have additional knowledge about the shape of the utility function (concave or convex) or its relation to a certain family of utility functions. In this case we can use the analogy with probability theory to minimize the cross entropy measure (Kullback and Leibler 1951) to a known utility density function. Minimum cross entropy formulations for a utility density,  $u(x)$ , and a target density,  $q(x)$ , take the form

$$u_{\min Xent}(x) = \arg \min_{u(x)} \left( \int_a^b u(x) \ln \left( \frac{u(x)}{q(x)} \right) dx \right)$$

such that

$$\int_a^b h_i(x) u(x) dx = \mu_i \quad i = 1, \dots, n; \tag{24}$$

$$\int_a^b u(x) dx = 1 \text{ and } u(x) \geq 0.$$

Using the method of Lagrange multipliers, we have

$$L(u(x)) = \int_a^b u(x) \ln\left(\frac{u(x)}{q(x)}\right) dx + \alpha_0 \left\{ \int_a^b u(x) dx - 1 \right\} + \sum_{i=1}^n \alpha_i \left\{ \int_a^b h_i(x) u(x) dx - \mu_i \right\} \quad (25)$$

Taking the first partial derivative of equation (24) with respect to  $u(x)$  and equating it to zero gives

$$\frac{\partial L(u(x))}{\partial u(x)} = \ln\left(\frac{u(x)}{q(x)}\right) + 1 + \alpha_0 + \sum_{i=1}^n \alpha_i h_i(x) = 0 \quad (26)$$

From equation (26), we can see that the minimum cross entropy solution takes the form

$$u_{\min Xent}(x) = q(x) e^{-\alpha_0 - 1 - \alpha_1 h_1(x) - \alpha_2 h_2(x) - \dots - \alpha_n h_n(x)}, \quad (27)$$

where  $\alpha_i$  is the Lagrange multiplier for each constraint and  $u_{\min Xent}(x)$  is the minimum cross entropy utility density. From equation (27), we can see that maximizing the entropy of  $u(x)$  is, therefore, a special case of minimizing the cross entropy when the target density,  $q(x)$ , is uniform (risk neutral utility function).

## 6.1. Minimum Cross Entropy Risk Aversion

If we take the logarithm of both sides of equation (27), we have

$$\ln(u_{\min Xent}(x)) = \ln(q(x)) - \alpha_0 - 1 - \alpha_1 h_1(x) - \alpha_2 h_2(x) - \dots - \alpha_n h_n(x). \quad (28)$$

From equation (28) we can see that the risk aversion function,  $\gamma_{\min Xent}(x)$ , for the minimum cross entropy utility solution is equal to the sum

$$\gamma_{\min Xent}(x) = -\frac{d}{dx} \ln(u_{\min Xent}(x)) = \gamma_{\text{target}}(x) + \alpha_1 h_1'(x) + \alpha_2 h_2'(x) + \dots + \alpha_n h_n'(x), \quad (29)$$

where  $\gamma_{\text{target}}(x) = -\frac{d}{dx} \ln(q(x))$  is the risk aversion function of the target density.

Target densities that produce common utility functions are the exponential utility density function,  $q(x) = \gamma e^{-\gamma x}$ , and the inverse utility density function,  $q(x) = \frac{1}{x + \alpha}$ , which lead to

exponential and logarithmic utility functions respectively. The target density is monotonically decreasing if the decision maker is risk averse and monotonically increasing if he is risk seeking.


**Example 5: The Party Problem Revisited**

To demonstrate an application of minimum cross entropy utility we refer back to the party problem of Example 2 and assume we have some knowledge of risk aversion for Kim. We use a

target density,  $q(x) = \frac{1}{\log(\frac{140}{40})(x + 40)}$ , which gives a normalized logarithmic utility function

over the domain  $[\$0, \$100]$ . Figure 7 shows the minimum cross entropy density, which is a piecewise inverse function that integrates to a piecewise logarithmic utility function and satisfies the given utility values. Incorporating knowledge of risk aversion through the target density,  $q(x)$ , contributes to the concavity of the utility function over the sub-intervals. The solution can be compared to the results of Figure 3 where no target density was available.

**Figure 7.** (a) Comparison of maximum entropy and minimum cross entropy utility densities. (b) Minimum cross entropy utility function vs. Kim’s utility function.



## 7. CONCLUSIONS

In this paper, we introduced an analogy between utility and probability through the notion of a utility density function and presented a maximum entropy utility principle to assign utility values with partial preference information.

Maximum entropy utility satisfies the main axioms of von Neumann and Morgenstern's normative utility theory. For example, since both transitivity and complete ordinal preferences were required for the prospects of the decision situation, the assigned maximum entropy utility values in turn satisfy transitivity and assign utility values for the complete set of prospects. The maximum entropy utility formulation assigns a unique utility value to each prospect due to the concavity of the entropy expression, and provides a continuous utility function over the domain of continuous attributes.

Jaynes (1968) proposed a basic desideratum for probability assignment, suggesting that in two problems where we have the same information, we should assign the same probabilities. In a similar fashion, the maximum entropy utility principle satisfies the analogous desideratum that in two problems where we have the same preference information, we should assign the same utility values. The maximum entropy utility formulation assigns the same utility values in different problems if the same preference information is incorporated into the constraints.

Maximum entropy utility also satisfies an essential desideratum of utility and probability independence that stems from the foundations of normative utility theory: the utility value of a prospect should not depend on the probability of getting that prospect due to the normative separation of beliefs from preferences (Samuelson 1952). The utility values assigned by maximum entropy utility do not depend on the lottery that the decision maker is facing.

The utility- probability analogy leads to further research on joint utility density functions for multiple attributes, utility inference mechanisms analogous to Bayes' rule for probability

inference, graphical representations of multiattribute utility functions, and duals to expected utility formulations with the roles of probability and utility reversed.

## ACKNOWLEDGEMENT

The author thanks Ronald Howard, Jim Dyer, the associate editor and three anonymous reviewers for many helpful suggestions on this paper.

## APPENDIX: MAXIMUM ENTROPY SOLUTION

The maximum entropy formulation for moments and/ or fractile constraints is

$$\begin{aligned} \max_{f(x)} \quad & - \int_a^b f(x) \ln(f(x)) dx \\ \text{subject to} \quad & \int_a^b h_i(x) f(x) dx = \mu_i \quad i = 1, \dots, n; \\ & \int_a^b f(x) dx = 1 \text{ and } f(x) \geq 0, \end{aligned} \quad (\text{A.I})$$

where  $[a, b]$  is the support of the maximum entropy distribution,  $h_i(x)$  is either an indicator function over an interval for fractile constraints, or  $x$  raised to a certain power, for moment constraints, and  $\mu_i$ 's are a given sequence of fractiles or moments.

Using the method of Lagrange multipliers, we have

$$L(f) = - \int_a^b f(x) \ln(f(x)) dx - \alpha_0 \left\{ \int_a^b f(x) dx - 1 \right\} - \sum_{i=1}^n \alpha_i \left\{ \int_a^b h_i(x) f(x) dx - \mu_i \right\}, \quad (\text{A.II})$$

where  $\alpha_i$  is the Lagrange multiplier for each fractile or moment constraint.

Taking the partial derivative with respect to  $f(x)$  and equating it to zero gives

$$\frac{\partial L(f)}{\partial f(x)} = -\ln(f(x)) - 1 - \alpha_0 - \sum_{i=1}^n \alpha_i h_i(x) = 0. \quad (\text{A.III})$$

Re-arranging equation (A.III) gives

$$f(x) = e^{-\alpha_0 - 1 - \alpha_1 h_1(x) - \alpha_2 h_2(x) - \dots - \alpha_n h_n(x)}. \quad (\text{A.IV})$$

For example, when no constraints are available, except that the density function is normalized and non-negative, the maximum entropy solution is uniform over a bounded domain.

$$f(x) = e^{-\alpha_0 - 1} = \frac{1}{b-a}, \quad a \leq x \leq b. \quad (\text{A.V})$$

Conversely, if the density function is of the form of equation (A.IV), then the constraint set needed for its assignment is

$$\int_a^b h_i(x) f(x) dx = \mu_i, \quad i = 0, \dots, n. \quad (\text{A.VI})$$

By writing any density function in the form of (A.IV), we can determine the constraints in the maximum entropy formulation that lead to its assignment. This is known as the inverse maximum entropy problem. For example, we can rewrite a Beta density in the form

$$f(x) = \frac{1}{\text{Beta}(m, n)} x^{m-1} (1-x)^{n-1} = e^{-\ln(\text{Beta}(m, n)) - (m-1)\ln x - (n-1)\ln(1-x)}, \quad 0 \leq x \leq 1 \quad (\text{A.VII})$$

Comparing (A.IV) and (A.VII), we can see that the constraint set needed to produce a beta density function is

$$\int_0^1 \ln(x) f(x) dx = \mu_1, \quad \int_0^1 \ln(1-x) f(x) dx = \mu_2, \quad (\text{A.VIII})$$

where  $\mu_1$  and  $\mu_2$  are given constants.

## REFERENCES

Abbas, A.E. 2002. An Entropy Approach for Utility Assignment in Decision Analysis. Proceedings of the 22<sup>nd</sup> *International workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering*. Moscow, Idaho. C. Williams (Ed).

- Arrow, K.J. 1965. The Theory of Risk Aversion. Lecture 2 in *Aspects of the Theory of Risk-Bearing*. Yrjo Jahnssonin Saatio, Helsinki.
- Barron, F.H., B.E. Barret. 1996. Decision Quality using Ranked Attribute Weights. *Management Science*. 42, 1515-1523.
- Berhold, M.H. 1973. The use of distribution functions to represent utility functions. *Management Science*. 19, 825-829.
- Butler, J.C., J. Jia, J.S. Dyer. 1997. Simulation Techniques for the Sensitivity Analysis of Multi-criteria Decision Models. *European Journal of Operational Research*. 103, 531-545.
- Castagnoli, E., M. LiCalzi. 1996. Expected utility without utility. *Theory and Decision*. 41, 281-301.
- DeGroot, M.H. 1970. *Optimal Statistical Decisions*. McGraw Hill Book Company.
- Dyer, J.S., R. Sarin. 1979. Group Preference Aggregate Rules Based on Strength of Preference. *Management Science*, 24, 9, 822-832.
- Dyer, J.S., R. Sarin. 1982. Relative Risk Aversion, *Management Science*, 28, 875-886.
- Howard, R.A. 1970. Decision Analysis: Perspectives on Inference, Decision, and Experimentation. *Proceedings of the IEEE*. 58, 632-643.
- Howard, R.A. 1980. On Making Life and Death Decisions, in *The Principles and Applications of Decision Analysis*, Vol. II, R. A. Howard and J. E. Matheson (eds.). Strategic Decisions Group, Menlo Park, California, USA.
- Howard, R.A. 1992. In Praise of the Old Time Religion. *Utility Theories: Measurements and Applications*, pp. 27-55, Ed. W. Edwards, Kluwer Academic Publishers.
- Jaynes, E.T. 1957. Information theory and statistical mechanics. *Physical Review*. 106, 620-630.

Jaynes, E.T. 1968. Prior Probabilities. *IEEE Transactions on Systems, Science and Cybernetics*. 4, 227-241.

Jessop, A. 1999. Entropy in Multiattribute Problems, *Journal of Multicriteria Decision Analysis*, 8, 61-70.

Kahneman, D., A. Tversky. 1979. Prospect Theory: An analysis of decision under risk. *Econometrica*. 47, 263-291.

Keeney, R.L., H. Raiffa. 1976. Decisions with Multiple Objectives: Preferences and Value Tradeoffs. John Wiley and Sons, Inc.

Kirkwood, C.W., R. Sarin. 1985. Ranking with Partial Information: A Method and an Application. *Operations Research*. 33, 38-48.

Kullback, S., R.A. Leibler. 1951. On Information and Sufficiency. *Annals of Mathematical Statistics*. 2, 79-86.

Matheson, J.E., R.A. Howard. 1968. An Introduction to Decision Analysis in *The Principles and Applications of Decision Analysis*, Vol. I, R. A. Howard and J. E. Matheson (eds.). Strategic Decisions Group, Menlo Park, California, USA.

Pratt, J. 1964. Risk aversion in the small and in the large. *Econometrica*. 32, 122-136.

Rao, J.S., M. Sobel. 1980. Incomplete Dirichlet Integrals with Applications to Ordered Uniform Spacings. *J. Multivariate Analysis*. 10, 603-610.

Samuelson, P.A. 1952. Probability, Utility, and the Independence Axiom. *Econometrica*. 20, 670 - 678.

Shannon, C.E. 1948. A mathematical theory of communication. *Bell sys. Tech. Journal*, 27, 379-423, 623-656.

von Neumann, J., O. Morgenstern. 1947. *Theory of games and economic behavior*. Princeton, N.J. Princeton University. 2<sup>nd</sup> edition.