Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Maximum Entropy Utility

Ali E. Abbas,

To cite this article:
Ali E. Abbas, (2006) Maximum Entropy Utility. Operations Research 54(2):277-290. http://dx.doi.org/10.1287/opre.1040.0204

Operations Research
Vol. 54, No. 2, March–April 2006, pp. 277–290
ISSN 0030-364X | EISSN 1526-5463 | 06 | 5402 | 0277

informs®

DOI 10.1287/opre.1040.0204

Maximum Entropy Utility

Ali E. Abbas
School of Engineering, University of Illinois at Urbana-Champaign, 104 S. Mathews Avenue, Urbana, Illinois 61801,
aliabbas@uiuc.edu

This paper presents a method to assign utility values when only partial information is available about the decision maker's preferences. We introduce the notion of a utility density function and a maximum entropy principle for utility assignment. 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.

Subject classifications: decision analysis: applications; utility/preference: estimation, multiattribute.
Area of review: Decision Analysis.
History: Received September 2002; revision received February 2004; accepted October 2004.

1. Introduction

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 et al. (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 kth 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 nonnegative 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.

We observe that the use of a normalized utility function has occurred in the literature with different interpretations and applications. For example, Berhold (1973) uses parametric forms of probability distributions to obtain convenient expressions for utility functions and certain equivalent calculations. In our work, we explore the analogy between probability and utility in both the discrete and continuous cases, for both numeric and nonnumeric variables, using parametric and nonparametric forms. 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 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. We have also found related work that maximizes the expected value of the entropy of the utility values of a given lottery (La Mura 1999). In comparison, we do not incorporate probability values into our utility entropy definition, and we use the analogy between probability and utility to incorporate the utility density function into the entropy expression.

The remainder of this paper is organized as follows. Section 2 presents several new definitions that highlight 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 §6 introduces minimum cross entropy utility, where additional knowledge about the shape of the utility function can be incorporated. Section 7 contains concluding remarks.

2. Utility-Probability Analogy

The analogy between utility and probability 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, in the presence of two additional prospects A and C such that $A \succ B \succ C$, 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 to von Neumann and Morgenstern type utility assessments, 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

$$U \triangleq (U_0, U_1, U_2, \ldots, U_{K-2}, U_{K-1}) = (0, U_1, U_2, \ldots, U_{K-2}, 1). \quad (1)$$

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 $0 \leqslant U_1 \leqslant U_2 \leqslant \cdots \leqslant U_{K-3} \leqslant U_{K-2} \leqslant 1$. This region, which we call the utility volume, has a volume equal to $1/(K-2)!$.

The second definition is a utility-increment vector, ΔU, 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

$$\Delta U \triangleq (U_1 - 0, U_2 - U_1, \ldots, 1 - U_{K-2}) = (\Delta u_1, \Delta u_2, \Delta u_3, \ldots, \Delta u_{K-1}). \quad (2)$$

Figure 1. A three-dimensional utility simplex.

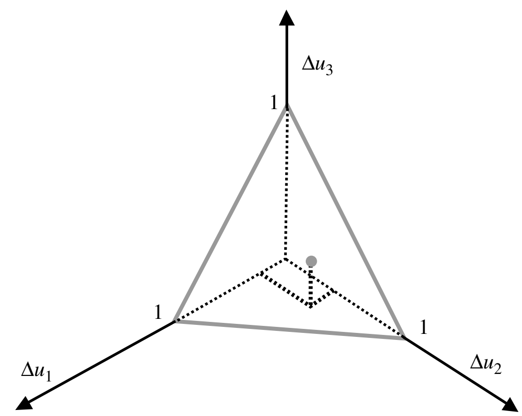

The coordinates of ΔU have two main properties: (1) they are all greater than or equal to zero and (2) sum to one. Therefore, any utility-increment vector can be represented as a point in a $(K-1)$-dimensional simplex $\{x: \sum_{i=1}^{K-1} x_i = 1, x_i \geqslant 0\}$. We will refer to this simplex as the utility simplex. A graphical illustration of a three-dimensional utility simplex is shown in Figure 1.

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

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 he assign? To answer this question, we observe that any point on 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 assert that the marginal probability density for each element of the utility increment vector is the Beta density, $\text{Beta}(1, K-2)$, while the marginal probability density for each element of the utility vector, U_j; $j = 1, \ldots, K-2$, is $\text{Beta}(j, K-j-1)$ (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). The analysis then 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 given the preference order, rather than marginal probability densities for the utility values 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.

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

$$\Delta u_i = \frac{1}{K-1}, \quad i = 1, \ldots, K-1. \tag{3}$$

The previous analysis 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, and results in a utility increment vector that lies in the center of the utility simplex. In §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 function, $U(x)$, over the given domain and is normalized to have values between zero and one. The normalized utility function has the same mathematical properties as a cumulative probability distribution as both are nondecreasing and range from zero to one. The normalization (and hence boundedness) of the utility function poses no 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 function (assuming the derivative exists) and we refer to it as a utility density function, $u(x)$ (Abbas 2002), i.e.,

$$u(x) \triangleq \frac{d}{dx} U(x). \tag{4}$$

If the utility function is normalized, then the utility density integrates to unity. The utility density function is nonnegative, due to the nondecreasing values of the utility function, and thus has the same mathematical properties as a probability density function: both are nonnegative 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.,

$$U(x) = \int_{x_{\min}}^{x} u(x)\, dx. \tag{5}$$

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

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 recall the famous 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 1812).

Laplace suggested that we assign *equal probabilities* to all outcomes unless there is *information* that suggests otherwise. When applied to Laplace's principle of insufficient reason, the utility-probability analogy suggests a method for assigning utility values that can be described as follows:

Table 1. Utility-probability analogy.

Probability		Utility	
Probability mass function	$P=(p_1,\ldots,p_K)$ p_i: $\sum_{i=1}^{K} p_i = 1,\ p_i \geqslant 0$	Utility increment vector	$\Delta U=(\Delta u_1,\ldots,\Delta u_{K-1})$ Δu_i: $\sum_{i=1}^{K-1}\Delta u_i = 1,\ \Delta u_i \geqslant 0$
Discrete cumulative probability	p_j: $\sum_{i=1}^{j} p_i,\ j=1,\ldots,K$ $P_j - P_{j-1} \geqslant 0$	Discrete utility value	U_j: $\sum_{i=1}^{j} u_i,\ j=1,\ldots,K-1$ $U_j - U_{j-1} \geqslant 0$
Probability density function	$f(x) \triangleq \frac{d}{dx}F(x)$ $\int_a^b f(x)\,dx = 1,\ f(x)\geqslant 0$	Utility density function	$u(x) \triangleq \frac{d}{dx}U(x)$ $\int_a^b u(x)\,dx = 1,\ u(x)\geqslant 0$
Cumulative distribution function	$F(x)=\int_{x_{\min}}^{x} f(x)\,dx$ $0\leqslant F(x)\leqslant 1,\ \frac{d}{dx}F(x)\geqslant 0$	Normalized utility function	$U(x)=\int_{x_{\min}}^{x} u(x)\,dx$ $0\leqslant U(x)\leqslant 1,\ \frac{d}{dx}U(x)\geqslant 0$

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. This utility assignment agrees with the mathematical results of the uniform Dirichlet distribution discussed earlier, and also 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 the next section, we present another application of the utility-probability analogy to measure the spread of a utility increment vector and utility density function.

3. The Entropy of a Utility Function

3.1. Entropy Measure for Discrete Prospects

Shannon (1948) introduced the term $H(p_1,\ldots,p_n) = -\sum_{i=1}^{n} p_i \log(p_i)$ 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

$$H(\Delta u_1,\Delta u_2,\Delta u_3,\ldots,\Delta u_{K-1}) = -\sum_{i=1}^{K-1}\Delta u_i \log(\Delta u_i). \qquad (6)$$

There are other measures that can be used for the spread, 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) *Monotonicity*: If the utility increments are equal, the measure of spread should be a monotonically increasing function of the number of prospects, K. 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) *Continuity*: 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) *Decomposition*: The measure of spread of the utility increment vector can be decomposed in terms of a weighted average of the spread of its subsets,

$$\begin{aligned} &H(\Delta u_1,\Delta u_2,\Delta u_3)\\ &\quad = H(\Delta u_1,(\Delta u_2+\Delta u_3))\\ &\qquad +(\Delta u_2+\Delta u_3)H\left(\frac{\Delta u_2}{\Delta u_2+\Delta u_3},\frac{\Delta u_3}{\Delta u_2+\Delta u_3}\right). \end{aligned} \qquad (7)$$

One implication of this axiom is that the order in which we calculate the measure of spread should not matter: if we calculate the spread of a utility increment vector directly using Equation (6), or if we calculate the spread of subsets of the utility increment vector separately then take a weighted average using (7), we should get the same result. For example, if we have a utility increment vector $\Delta U = (0.25, 0.5, 0.25)$, we can calculate its entropy directly as

$$H(0.25, 0.5, 0.25) \stackrel{\Delta}{=} -0.25\log(0.25) - 0.5\log(0.5) - 0.25\log(0.25) = 1.5\log(2).$$

If the last two elements (Δu_2 and Δu_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 $\Delta U_{23} = (2/3, 1/3)$. The original increment vector now reduces to only two coordinates, $\Delta U_R = (0.25, 0.75)$. The spread of $\Delta U_R = (0.25, 0.75)$ is therefore less than that of ΔU, and as a consequence, its entropy must also be less. However, when we add the weighted entropy due to the spread in ΔU_{23}, we have $H(0.25, 0.75) + 0.75H(2/3, 1/3) = 1.5\log(2)$. Both methods provide the same spread amount.

The rationale for having a weighted average in (7), rather than a simple sum, is that the increase in spread contributed by the elements of ΔU_{23} should depend on their utility mass. For example, consider another utility increment vector $\Delta U = (0.98, 0.01, 0.01)$. The last two elements contribute to the spread, but their contribution is not significant compared to that of the first element. As a result, the entropy of the normalized vector containing the last two elements, $\Delta U_{23} = (0.5, 0.5)$, should not contribute as much to the overall spread. This is why it is weighted by an amount 0.02 in the decomposition expression.

The entropy measure of Equation (6) is the only measure that satisfies the three axioms mentioned above. Now we discuss the implications of maximizing the entropy of the utility increment vector subject to given preference information.

Example 1 (Utility Increment Vector with Maximum Entropy). A decision maker faces a decision situation with K ordered prospects. We are interested in the utility increment vector, which maximizes the entropy measure subject to this information alone. The mathematical formulation for the problem is

$$\Delta U_{\text{maxent}} = \operatorname*{arg\,max}_{\Delta u_1, \Delta u_2, \ldots, \Delta u_{K-1}} -\sum_{i=1}^{K-1} \Delta u_i \log(\Delta u_i)$$

such that

$$\sum_{i=1}^{K-1} \Delta u_i = 1, \qquad (8)$$

$$\Delta u_i \geqslant 0, \quad i = 1, \ldots, K-1.$$

The solution to (6) yields a utility increment vector with equal increments

$$\Delta u_i = \frac{1}{K-1}, \quad i = 1, \ldots, K-1. \qquad (9)$$

The entropy of the vector described by (9) is $\log(K-1)$, which is the maximum value that can be attained for K prospects. Recall that this vector has the same utility increments as those described by the uniform Dirichlet distribution when only the preference order of the prospects was available. We have arrived at the same result by maximizing the entropy of the utility increment vector.

If additional information is available about the decision maker's preferences, this can also be incorporated into the formulation. For example, suppose that we know of preference ties between two prospects, U_{l-1} and U_l. In this case, we add an additional constraint in (8) that $\Delta u_l = 0$. This addition results in a utility increment vector of the form

$$\Delta U = (\Delta u_1, \ldots, \Delta u_{l-1}, \Delta u_l, \Delta u_{l+1}, \ldots, \Delta u_{K-1}) = \left(\frac{1}{K-2}, \ldots, \frac{1}{K-2}, 0, \frac{1}{K-2}, \ldots, \frac{1}{K-2}\right). \qquad (10)$$

Maximizing the entropy of the utility increment vector with given preference constraints therefore yields a vector that (1) satisfies the given constraints, and (2) produces (whenever the constraints allow) equal increments in utility values. Note that the utility increment vector of (10) has entropy of $\log(K-2)$, which is less than that of (9). The reduction in entropy is due to the additional preference information that $\Delta u_l = 0$. Let us now consider the minimum entropy solution to this problem:

$$\Delta U_{\text{minent}} = \operatorname*{arg\,min}_{\Delta u_1, \Delta u_2, \ldots, \Delta u_{K-1}} -\sum_{i=1}^{K-1} \Delta u_i \log(\Delta u_i)$$

such that

$$\sum_{i=1}^{K-1} \Delta u_i = 1, \qquad (11)$$

$$\Delta u_i \geqslant 0, \quad i = 1, \ldots, K-1.$$

The minimum entropy solution to (11) is not unique, and results in a utility increment vector with one nonzero element equal to one, and the remaining $(K-2)$ elements equal to zero. Different permutations in the position of this nonzero element yield different solutions to the minimum entropy problem that occur at the extreme (corner) points of the utility simplex. For example, if the nonzero element is Δu_l, we have

$$\Delta U = (\Delta u_1, \ldots, \Delta u_{l-1}, \Delta u_l, \Delta u_{l-1}, \ldots, \Delta u_{K-1}) = (0, \ldots, 0, 1, 0, \ldots, 0). \qquad (12)$$

The utility increment vector of (12) has zero entropy. This minimum entropy solution implies (1) preference ties for all prospects with utility values U_i, $i \leqslant l-1$; (2) a steep change of preferences for the prospect with utility U_l; and (3) preference ties for all prospects with utility values U_i, $i \geqslant l$. The utility increment vector of (12) therefore implies more information about a decision maker's preferences than either (9) or (10). As a consequence, it has less entropy.

3.2. Entropy Measure for Continuous Prospects

In continuous form, the entropy expression is known as the differential entropy, $h(x)$. The differential entropy of a utility density function on a domain, $[a, b]$, is

$$h(u(x)) = -\int_a^b u(x)\ln(u(x))\,dx. \tag{13}$$

The maximum entropy solution subject to the normalization and nonnegativity constraints $\int_a^b u(x)dx = 1$, $u(x) \geqslant 0$, in the continuous case has the form (see the appendix)

$$u(x) = \frac{1}{(b-a)}, \quad a \leqslant x \leqslant b. \tag{14}$$

The utility density function of (14) integrates to a linear (risk-neutral) utility function

$$U(x) = \frac{x-a}{b-a}, \quad a \leqslant x \leqslant b. \tag{15}$$

The maximum entropy solution on a continuous bounded domain therefore exhibits risk-neutral behavior. To gain some further intuition about the implications of the maximum entropy and minimum entropy assignments in the continuous case, let us consider the following example.

Example 2 (Entropy of the CARA Utility). Consider the following constant absolute risk aversion (CARA) utility density that is normalized over the domain [0, 1]:

$$u(x) = \frac{\gamma e^{-\gamma x}}{1-e^{-\gamma}}, \quad 0 \leqslant x \leqslant 1, \tag{16}$$

where γ is the decision maker's risk-aversion coefficient.

By direct integration and the use of Equation (13), the differential entropy is

$$h(u(x)) = 1 + \frac{\gamma e^{\gamma}}{1-e^{\gamma}} - \ln\left(\frac{\gamma}{e^{\gamma}-1}\right). \tag{17}$$

Figure 2 plots the differential entropy versus the risk aversion, γ, from Equation (17).

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 $\gamma \to 0$,

$$u(x) = \frac{\gamma e^{-\gamma x}}{1-e^{-\gamma}} \to 1, \quad 0 \leqslant x \leqslant 1. \tag{18}$$

The utility density that maximizes the entropy expression is uniform (as expected) and integrates to a linear utility (risk-neutral) function, $U(x) = x$, $0 \leqslant x \leqslant 1$. Note that this maximum entropy solution is unique, and is both concave and convex. The maximum entropy utility assignment, therefore, favors no direction of risk attitude (as it occurs at the boundary of the two domains). From Figure 2, we also 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 of (17) has no lower bound because $h(u(x)) \to -\infty$ as $\gamma \to +\infty$ (the case of extreme

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

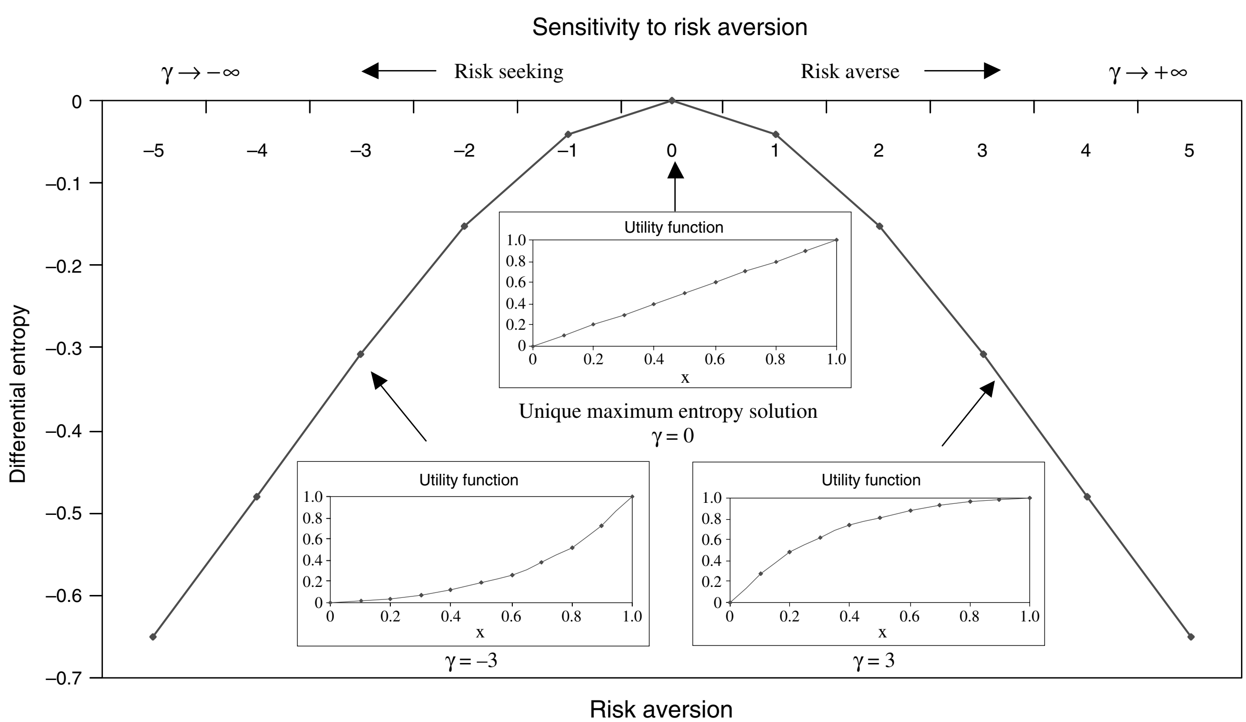

risk-averse behavior) and the utility density approaches an impulse density, $\delta(x)$,

$$\frac{\gamma e^{-\gamma x}}{1-e^{-\gamma}} \to \delta(x) \quad \text{as } \gamma \to +\infty. \tag{19}$$

The impulse utility density, $\delta(x)$, is a minimum entropy solution and integrates to a step (aspiration) utility function that jumps at the lower bound of the domain. The step utility function implies both (1) extreme risk averse behavior, and (2) a steep change in preferences at $x=0$.

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

$$\frac{\gamma e^{-\gamma x}}{1-e^{-\gamma}} \to \delta(x-1) \quad \text{as } \gamma \to -\infty \tag{20}$$

From Equations (19) and (20), 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, $\delta(x-x_0)$, $0 \leqslant x_0 \leqslant 1$, 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 utility assignment (1) produces equal increments in utility values for the case discrete prospects; (2) makes no assumptions about the direction of the risk attitude in the continuous case; and (3) makes no assumptions about steep changes in the decision maker's preferences (unless additional preference information is incorporated into the constraints as we shall see in §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 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 nonuniform 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 function (*or utility vector*) *whose utility density* (*or utility increment vector*) *has maximum entropy subject to whatever preferences are known.*

This method of utility assignment provides an analogy with Jaynes' maximum entropy principle for probability inference. 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 formulation with constraints, $\int_a^b h_i(x)u(x)\,dx = \mu_i$, $i=1,\ldots,n$, as well as the normalization and nonnegativity constraints, $\int_a^b u(x)\,dx = 1$ and $u(x) \geqslant 0$, is

$$\begin{aligned} &\underset{u(x)}{\arg\max} - \int_a^b u(x)\ln(u(x))\,dx \\ &\text{subject to} \\ &\int_a^b h_i(x)u(x)\,dx = \mu_i, \quad i=1,\ldots,n, \\ &\int_a^b u(x)\,dx = 1, \quad \text{and} \quad u(x) \geqslant 0. \end{aligned} \tag{21}$$

Formulation (21) yields a utility density of the form (see the appendix)

$$u_{\text{maxent}}(x) = e^{-\alpha_0 - 1 - \alpha_1 h_1(x) - \alpha_2 h_2(x) - \cdots - \alpha_n h_n(x)}, \tag{22}$$

where $u_{\text{maxent}}(x)$ is the maximum entropy utility solution, $[a, b]$ is the domain of prospects, $h_i(x)$ is a given preference constraint, μ_is are given constants, and α_is are the Lagrange multipliers for the given constraints.

As we shall see, expression (22) is a general expression that includes the most commonly used functional forms of utility functions. For example, the risk-neutral utility function, which has a uniform utility density, arises when the constraints $h_i(x)=0$ in (21) and (22). When $n=1$ and $h_1(x)=x$, the maximum entropy utility is a CARA utility on the positive domain. When $n=2$, $h_1(x)=x$ and $h_2(x)=x^2$, the maximum entropy utility is a Gaussian utility density, which integrates to an S-shaped utility function on the real domain, similar in form to the value function in prospect theory (Kahneman and Tversky 1979).

The maximum entropy utility solution also embeds the hyperbolic absolute risk-averse utility function (HARA) on a bounded domain,

$$U(x) = \frac{1}{1-\gamma}\left(\gamma\left(\beta + \frac{\alpha}{\gamma}x\right)^{1-\gamma} - 1\right), \tag{23}$$

with a utility density of the form

$$u(x) = \alpha\left(\beta + \frac{\alpha}{\gamma}x\right)^{-\gamma} = e^{\ln(\alpha) - \gamma\ln(\beta + (\alpha/\gamma)x)}, \tag{24}$$

where α, β, and γ are given constants. The HARA utility function reduces to a risk-neutral utility function when $\gamma = 0$; to a CARA utility function when $\gamma \to \pm\infty$; and to a constant relative risk-averse utility function (CRRA) when $\beta = 0$ and $\gamma > 0$. Comparing Equations (22) and (24) shows that HARA utility functions can be expressed by the maximum entropy utility solution when $n = 1$ and $h_1(x) = \ln(\beta + (\alpha/\gamma)x)$.

Using Equation (22), and Arrow-Pratt's definition of local risk aversion (Pratt 1964, Arrow 1965), the maximum entropy utility function with constraints $h_i(x)$, $i = 0, 1, \ldots, n$, has a risk aversion, $\gamma_{\text{maxent}}(x)$, of

$$\gamma_{\text{maxent}}(x) = -\frac{d}{dx}\ln(u_{\text{maxent}}(x)) = \alpha_1 h_1'(x) + \alpha_2 h_2'(x) + \cdots + \alpha_n h_n'(x), \qquad (25)$$

where $h_i'(x) = (d/dx)h_i(x)$. Equation (25) shows the linear effect contributed by the derivative of each preference constraint on the risk-aversion function, and the role of each constraint, $h_i(x)$, in determining the overall risk aversion. For example, Equation (25) explains how the constraints $n = 2$, $h_1(x) = x$, and $h_2(x) = x^2$ lead to an S-shaped utility function on the real domain, because $h_1'(x)$ is a constant and $h_2'(x) = 2x$ changes sign from the positive domain to the negative domain leading to a change in the risk-aversion function from risk aversion to risk seeking. Equation (25) also shows the wide range of risk-aversion expressions that can be modeled by the maximum entropy utility family.

4.1. Interpreting the Preference Constraints

Now we provide an interpretation for the equality constraint $\int_a^b h_i(x)u(x)\,dx = \mu_i$ for the case where $h_i(x)$ is a bounded nondecreasing function on the domain of prospects. Expected utility theory asserts that for any given lottery with cumulative distribution $F(x)$ on the domain $[a, b]$, there exists an equivalent two-outcome lottery, with a probability, p, of getting b and $(1-p)$ of getting a such that a decision maker is indifferent between the two lotteries (Figure 3). If the utility function is normalized, then the indifference probability, p, is equal to the expected utility of both lotteries. Suppose that we know the value of the indifference probability that makes a decision maker indifferent between a given lottery and its equivalent two-outcome lottery.

Equating the expected utility of both sides of Figure 3 for a normalized utility function gives

$$\int_a^b U(x)dF(x) = p. \qquad (26)$$

Using the rule of integration by parts,

$$U(x)F(x)|_a^b - \int_a^b F(x)u(x)\,dx = p$$
$$\Rightarrow 1 - 0 - \int_a^b F(x)u(x)dx = p, \qquad (27)$$

Figure 3. Indifference between a lottery $F(x)$ and an equivalent two-outcome lottery.

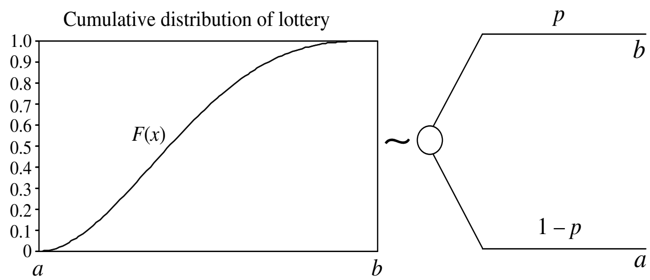

where $U(b)F(b) = 1$ and $U(a)F(a) = 0$ for a normalized utility function.

Re-arranging gives

$$\int_a^b F(x)u(x)\,dx = 1 - p. \qquad (28)$$

Comparing (28) to the preference constraint, $\int_a^b h_i(x) \cdot u(x)\,dx = \mu_i$, shows that the preference constraint can be interpreted as knowledge of the decision maker's indifference probability for a given lottery. With this interpretation, $h_i(x)$ can be interpreted as a cumulative distribution, $F(x)$, and μ_i can be interpreted as the decision maker's indifference probability for its least preferred outcome. In particular, if we know the indifference probability for a uniform lottery, the maximum entropy utility assignment yields an exponential (CARA) utility function.

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 function 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. We illustrate this application through the following example.

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

Example 3 (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.

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

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

such that

$$\int_0^{20} u(x)\,dx = 0.32, \quad \int_0^{40} u(x)dx = 0.57, \tag{29}$$

$$\int_0^{50} u(x)dx = 0.67, \quad \int_0^{90} u(x)dx = 0.95,$$

$$\int_0^{100} u(x)\,dx = 1, \quad u(x) \geqslant 0.$$

If we compare the preference constraints, $h_i(x)$, of Equation (29) to those of (22), we find that they are in effect indicator functions over certain intervals. For example, the constraint

$$\int_0^{20} u(x)\,dx = \int_0^{100} I_{20}(x)u(x)\,dx = \int_0^{100} h_1(x)u(x)\,dx,$$

where $I_{20}(x)$ is an indicator function for $x \in [0, 20]$. From (22), we can see that the solution to this problem has the form

$$u_{\text{maxent}}(x) = e^{-\alpha_0 - 1 - \alpha_1 I_{20}(x) - \alpha_2 I_{40}(x) - \alpha_3 I_{50}(x) - \alpha_4 I_{90}(x)},$$
$$0 \leqslant x \leqslant 100. \tag{30}$$

Equation (30) is the staircase utility density shown in Figure 4(a). Figure 4(a) also shows Kim's CARA utility density (not known to the decision analyst). In Figure 4(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.

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 functional form that is chosen for the fit. The maximum entropy approach, however, provides a unique utility function that makes no prior assumptions about the structure of the utility function: the functional form is determined by the available preference information constraints, $h_i(x)$. We note that Equation (29) does not incorporate any information about Kim's risk attitude over the sub-intervals. In §6, we will refer back to this example and incorporate risk aversion into Kim's formulation.

Figure 4. Comparison of both Kim's utility function and the maximum entropy utility function.

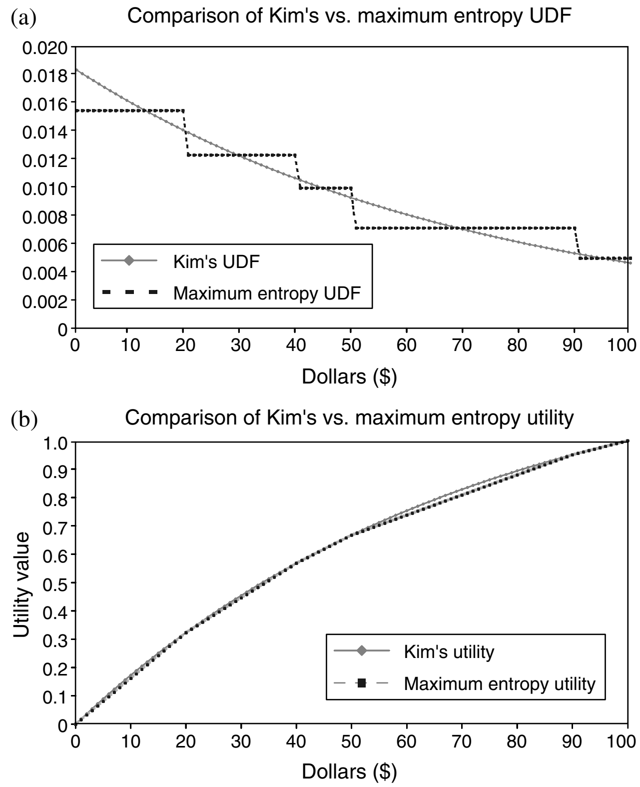

5.2. Inferring Utility Values by Observing Decisions

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

$$\int_a^b U(x)\,dF(x) \geqslant \int_a^b U(x)\,dG(x)$$

$$\Rightarrow U(x)F(x)\big|_a^b - \int_a^b F(x)u(x)\,dx$$

$$\geqslant U(x)G(x)\big|_a^b - \int_a^b G(x)u(x)\,dx$$

$$\Rightarrow \textstyle\int_a^b [G(x)-F(x)]u(x)\,dx \geqslant 0, \qquad (31)$$

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

The maximum entropy utility formulation becomes

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

such that

$$\int_a^b u(x)\,dx = 1, \quad u(x) \geqslant 0, \qquad (32)$$

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

If stochastic dominance exists between the two distributions ($F(x) \leqslant G(x)\ \forall x$), then any decision maker (even the risk seeking) would prefer $F(x)$ to $G(x)$. This implies that we should not learn anything new about the decision maker's utility function by observing this decision. This situation is reflected in the maximum entropy formulation by noting that any utility density function satisfies condition (31) (as it is nonnegative), and so condition (31) is never binding.

If stochastic dominance does not exist, then we may learn about the decision maker's utility function by observing his decisions. Formulation (32) can be solved using the Karush-Kuhn-Tucker optimality conditions or by discretization. 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 feasible set (for a reference on convex optimization, see Boyd and Vandenberghe 2004). We demonstrate an application of (32) through the following example.

EXAMPLE 4 (INFERRING UTILITY VALUES). 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 5(a) shows the maximum entropy utility function and the decision maker's utility function. If the decision maker faces the two lotteries of Figure 5(b) and prefers lottery 2 to lottery 1, an additional inequality constraint can be added into the maximum entropy formulation as described above. Figure 5(b) shows the effect of observing this decision on the maximum entropy utility function. Figures 6(a) and 6(b) show how the maximum entropy utility function is updated after observing more decisions made by the decision maker.

5.3. Maximum Entropy Multiattribute Utility

When the decision situation has multiple attributes, a value function is constructed to rank order the prospects, and

Figure 5. (a) Exponential vs. maximum entropy utility function. (b) Two lotteries faced by the decision maker and updated utility function.

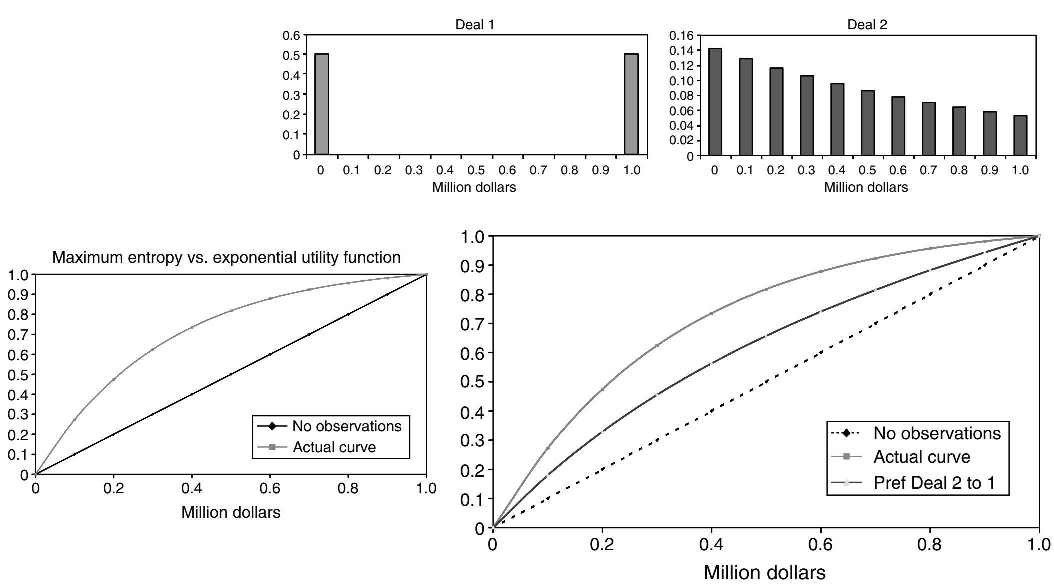

Figure 6. (a) Second observation (deal 4 is preferred to deal 3). (b) Third observation (deal 6 is preferred to deal 5).

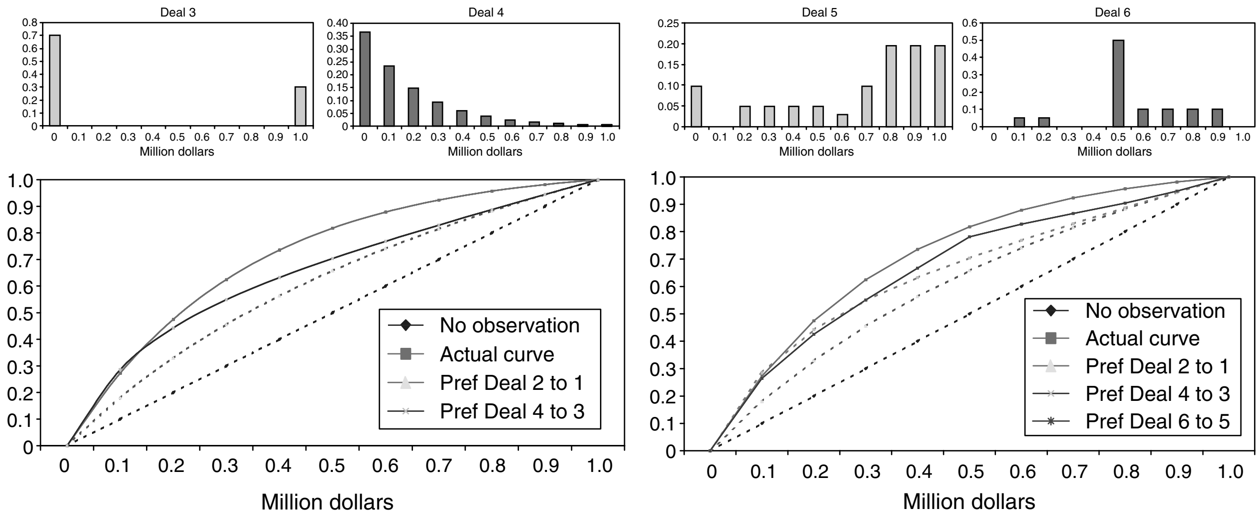

a utility function is assigned over the value function to represent the decision maker's risk attitude towards value (Matheson and Howard 1968; Dyer and Sarin 1979, 1982; Keeney and Raiffa 1976). Using this approach, a maximum entropy multiattribute utility function can be constructed with partial preference 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 5 (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 zero (instant painless death) to one (current health with no disability). The value model over consumption and health states is a Cobb-Douglas function given as

$$V(x,y)=yx^{\eta},\quad 0\leqslant x\leqslant 1,\ 0\leqslant y\leqslant 1, \tag{33}$$

where x is in millions of dollars, y is the health state, and η 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,

$$\begin{aligned}U_{\text{maxent}}(V(x,y))&=V(x,y)\\&=yx^{\eta},\quad 0\leqslant x\leqslant 1,\ 0\leqslant y\leqslant 1.\end{aligned} \tag{34}$$

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

$$U(x)=x^{\eta},\ 0\leqslant x\leqslant 1\quad\text{and}\quad U(y)=y,\ 0\leqslant y\leqslant 1. \tag{35}$$

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, η: the decision maker is risk averse for consumption if $\eta<1$ and risk seeking if $\eta>1$. Figure 7 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.

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. In the continuous case, minimum cross-entropy formulations for a utility density, $u(x)$, and a target density, $q(x)$, take the form

$$\begin{gathered}u_{\text{minXent}}(x)=\arg\min_{u(x)}\left(\int_a^b u(x)\ln\left(\frac{u(x)}{q(x)}\right)dx\right)\\ \text{such that}\\ \int_a^b h_i(x)u(x)\,dx=\mu_i,\quad i=1,\ldots,n,\\ \int_a^b u(x)\,dx=1\quad\text{and}\quad u(x)\geqslant 0.\end{gathered} \tag{36}$$

Using the method of Lagrange multipliers, we have

$$\begin{aligned}L(u(x))={}&\int_a^b u(x)\ln\left(\frac{u(x)}{q(x)}\right)dx+\alpha_0\left\{\int_a^b u(x)\,dx-1\right\}\\&+\sum_{i=1}^n\alpha_i\left\{\int_a^b h_i(x)u(x)\,dx-\mu_i\right\}.\end{aligned} \tag{37}$$

Figure 7. (a) Utility contour plot. (b) Maximum entropy utility surface for $\eta = 0.7$.

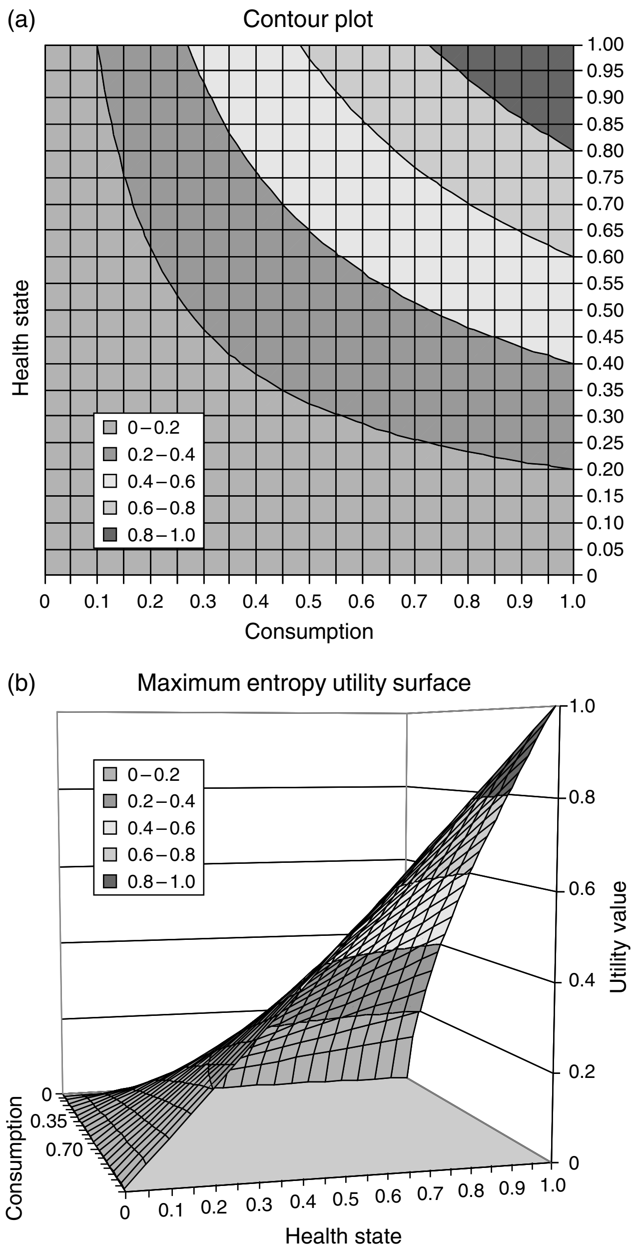

Taking the first partial derivative of Equation (36) with respect to $u(x)$ and equating it to zero gives

$$\frac{\partial L(u(x))}{\partial u(x)} = \ln\left(\frac{u(x)}{q(x)}\right) + 1 + \alpha_0 + \sum_{i=1}^{n} \alpha_i h_i(x) = 0. \quad (38)$$

Re-arranging (38) shows that the minimum cross entropy solution has the form

$$u_{\text{minXent}}(x) = q(x)e^{-\alpha_0 - 1 - \alpha_1 h_1(x) - \alpha_2 h_2(x) - \cdots - \alpha_n h_n(x)}, \quad (39)$$

where α_i is the Lagrange multiplier for each constraint and $u_{\text{minXent}}(x)$ is the minimum cross entropy utility density. From Equation (39), 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). As is the case with probability theory, minimum cross-entropy formulations guarantee the independence of the solution on the choice of the coordinate system (Jaynes 1968).

6.1. Minimum Cross Entropy Risk Aversion

If we take the logarithm of both sides of Equation (39), we have

$$\ln(u_{\text{minXent}}(x)) = \ln(q(x)) - \alpha_0 - 1 - \alpha_1 h_1(x) - \alpha_2 h_2(x) - \cdots - \alpha_n h_n(x). \quad (40)$$

From Equation (40) we can see that the risk-aversion function, $\gamma_{\text{minXent}}(x)$, for the minimum cross entropy utility solution is equal to the sum

$$\gamma_{\text{minXent}}(x) = -\frac{d}{dx}\ln(u_{\text{minXent}}(x)) = \gamma_{\text{target}}(x) + \alpha_1 h_1'(x) + \alpha_2 h_2'(x) + \cdots + \alpha_n h_n'(x), \quad (41)$$

where $\gamma_{\text{target}}(x) = -(d/dx)\ln(q(x))$ 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 hyperbolic utility density function, $q(x) = 1/(x+\alpha)$, 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 6 (The Party Problem Revisited). To demonstrate an application of minimum cross entropy utility we refer back to the party problem of Example 3 and assume that we have some additional knowledge about Kim's preferences. For example, we may know that she is risk averse, and, furthermore, we may have an estimate of her initial wealth being \$40. In this case, we use a target density of the form

$$q(x) = \frac{1}{\log(140/40)(x+40)}, \quad (42)$$

which gives a normalized logarithmic utility function over the domain [\$0, \$100].

Figure 8 shows the minimum cross entropy density, which is a piecewise hyperbolic 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 4 where no target density was available.

Figure 8. (a) Comparison of maximum entropy and minimum cross entropy utility densities. (b) Minimum cross entropy utility function vs. Kim's utility function.

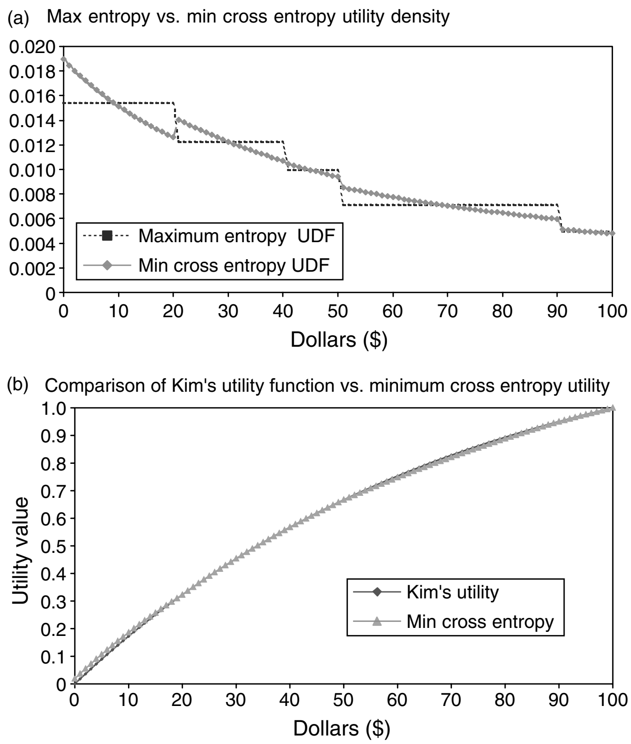

7. Conclusions

In this paper, we presented an analogy between utility and probability through the notion of a utility density function, and 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, because 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 that we presented in this paper 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.

Appendix. Maximum Entropy Solution

The maximum entropy formulation for a utility density function with given constraints is

$$\arg\max_{u(x)} -\int_a^b u(x)\ln(u(x))\,dx$$

subject to

$$\int_a^b h_i(x)u(x)\,dx=\mu_i,\quad i=1,\ldots,n, \tag{A1}$$

$$\int_a^b u(x)\,dx=1 \quad\text{and}\quad u(x)\geqslant 0,$$

where $[a,b]$ is the domain of the prospects, $h_i(x)$ is a function that describes some preference information, and μ_is are given constants.

Using the method of Lagrange multipliers, we have

$$L(u(x))=-\int_a^b u(x)\ln(u(x))\,dx-\alpha_0\left\{\int_a^b u(x)\,dx-1\right\}$$
$$-\sum_{i=1}^n \alpha_i\left\{\int_a^b h_i(x)u(x)\,dx-\mu_i\right\}, \tag{A2}$$

where α_i is the Lagrange multiplier for each preference constraint.

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

$$\frac{\partial L(u(x))}{\partial u(x)}=-\ln(u(x))-1-\alpha_0-\sum_{i=1}^n \alpha_i h_i(x)=0. \tag{A3}$$

Re-arranging Equation (A3) gives

$$u(x)=e^{-\alpha_0-1-\alpha_1 h_1(x)-\alpha_2 h_2(x)-\cdots-\alpha_n h_n(x)}. \tag{A4}$$

When only the normalization and nonnegativity constraints are available, the maximum entropy solution is uniform over a bounded domain,

$$u(x)=e^{-\alpha_0-1}=\frac{1}{b-a},\quad a\leqslant x\leqslant b. \tag{A5}$$

Furthermore, if a utility density function is of the form of (A4), then the constraint set needed for its assignment is

$$\int_a^b h_i(x)u(x)\,dx = \mu_i, \quad i = 0, \ldots, n. \tag{A6}$$

By writing any density function in the form of (A4), 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 function in the form

$$\begin{aligned} u(x) &= \frac{1}{\text{Beta}(m,n)} x^{m-1}(1-x)^{n-1} \\ &= e^{-\ln(\text{Beta}(m,n)) - (m-1)\ln x - (n-1)\ln(1-x)}, \quad 0 \leqslant x \leqslant 1. \end{aligned} \tag{A7}$$

Comparing (A4) and (A7), we can see that the constraint set needed to assign a Beta density function is

$$\int_0^1 \ln(x)u(x)\,dx = \mu_1, \quad \int_0^1 \ln(1-x)u(x)\,dx = \mu_2, \tag{A8}$$

where μ_1 and μ_2 are given constants.

Acknowledgments

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

References

Abbas, A. E. 2002. An entropy approach for utility assignment in decision analysis. C. Williams, ed. *Proc. 22nd Internat. Workshop on Bayesian Inference and Maximum Entropy Methods in Sci. Engrg.* Moscow, ID.

Arrow, K. J. 1965. The theory of risk aversion. *Aspects of the Theory of Risk-Bearing*, Lecture 2. Yrjo Jahnssonin Saatio, Helsinki, Finland.

Barron, F. H., B. E. Barret. 1996. Decision quality using ranked attribute weights. *Management Sci.* **42** 1515–1523.

Berhold, M. H. 1973. The use of distribution functions to represent utility functions. *Management Sci.* **19** 825–829.

Boyd, S., L. Vandenberghe. 2004. *Convex Optimization*. Cambridge University Press, Cambridge, UK.

Butler, J. C., J. Jia, J. S. Dyer. 1997. Simulation techniques for the sensitivity analysis of multi-criteria decision models. *Eur. J. Oper. Res.* **103** 531–545.

Castagnoli, E., M. LiCalzi. 1996. Expected utility without utility. *Theory Decision* **41** 281–301.

DeGroot, M. H. 1970. *Optimal Statistical Decisions*. McGraw Hill, New York.

Dyer, J. S., R. Sarin. 1979. Group preference aggregate rules based on strength of preference. *Management Sci.* **24**(9) 822–832.

Dyer, J. S., R. Sarin. 1982. Relative risk aversion. *Management Sci.* **28** 875–886.

Howard, R. A. 1970. Decision analysis: Perspectives on inference, decision, and experimentation. *Proc. IEEE* **58** 632–643.

Howard, R. A. 1980. On making life and death decisions. R. A. Howard, J. E. Matheson, eds. *The Principles and Applications of Decision Analysis*, Vol. II. Strategic Decisions Group, Menlo Park, CA.

Howard, R. A. 1992. In praise of the old time religion. W. Edwards, ed. *Utility Theories: Measurements and Applications*. Kluwer Academic Publishers, Boston, MA, 27–55.

Jaynes, E. T. 1957. Information theory and statistical mechanics. *Physical Rev.* **106** 620–630.

Jaynes, E. T. 1968. Prior probabilities. *IEEE Trans. Systems, Sci. Cybernetics* **4** 227–241.

Jessop, A. 1999. Entropy in multiattribute problems. *J. Multicriteria Decision Anal.* **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, New York.

Kirkwood, C. W., R. Sarin. 1985. Ranking with partial information: A method and an application. *Oper. Res.* **33** 38–48.

Kullback, S., R. A. Leibler. 1951. On information and sufficiency. *Ann. Math. Statist.* **2** 79–86.

La Mura, P. 1999. Foundations of multi-agent systems. Ph.D. dissertation, Chapter 6. Stanford University, Stanford, CA.

Laplace, P. S. 1812. *Theorie Analytique des Probabilities*.

Matheson, J. E., R. A. Howard. 1968. An introduction to decision analysis. R. A. Howard, J. E. Matheson, eds. *The Principles and Applications of Decision Analysis*, Vol. I. Strategic Decisions Group, Menlo Park, CA.

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 Anal.* **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 Systems Tech. J.* **27** 379–423, 623–656.

von Neumann, J., O. Morgenstern. 1947. *Theory of Games and Economic Behavior*, 2nd ed. Princeton University Press, Princeton, NJ.