

Constructing Multiattribute Utility Functions for Decision Analysis

Ali E. Abbas

Department of Industrial and Enterprise Systems Engineering, College of Engineering,
University of Illinois at Urbana–Champaign, Urbana, Illinois 61820, aliabbas@uiuc.edu

Abstract A fundamental step in decision analysis is the accurate representation of the decision maker's preferences. When the decision situation is deterministic, each alternative leads to a single prospect (consequence). A prospect may be characterized by one or more attributes, such as health state and wealth. A value function that ranks the prospects is sufficient to rank order the decision alternatives in deterministic decision problems. When uncertainty is present, each alternative may result in a number of possible prospects, each characterized by a number of attributes. A von Neumann–Morgenstern utility function, defined over the domain of the attributes, is then required for each prospect we face. The best decision alternative is the one with the highest expected utility. This chapter surveys a variety of methods for constructing multiattribute utility functions. These methods include (i) using a deterministic value function and a one-dimensional utility function over value, (ii) using a general expansion theorem for multiattribute utility functions, (iii) using an independence assumption and a utility diagram to simplify the assessment of the conditional utility functions, (iv) using an attribute dominance condition to reduce the number of assessments, (v) using a utility copula to incorporate dependence using single-attribute assessments, (vi) deriving the functional form by asserting the number of preference switches that may occur for lotteries defined on a subset of the attributes, and (vii) characterizing preferences using functional equations to derive the functional form.

Keywords utility theory, multiple attributes, dependence, invariance, functional equations, graph-based methods

1. Introduction: Why Do We Need a von Neumann–Morgenstern Utility?

1.1. Review of Expected Utility Decision Making

In many decision problems, a decision maker faces a number of alternatives from which she needs to choose. When the decision situation is deterministic, each alternative leads to a single prospect. A prospect is a deterministic characterization of a scenario that may occur once the decision is made. When uncertainty is present, a decision alternative may result in a number of possible prospects, each occurring with a specified probability. It is the uncertainty about the actual prospect we will get that makes decision making difficult. To illustrate, consider a decision faced by a manager in an automobile company. He may need to choose between two alternatives: introducing a new innovative vehicle or making slight changes to an existing design. Each alternative may result in a continuum of prospects ranging from very high profit to the organization (if the demand is high) to major financial losses (if the demand is low or if there are recalls). The first alternative may result in a much larger profit for the company, but it also has a higher probability of a downside loss. The manager believes that the second alternative is safer (a proven design), but at the same time, it has a lower probability of the large upside profits that the first alternative may provide. Which of these two alternatives should that decision maker choose?

If the decision maker follows the axioms of expected utility theory (von Neumann and Morgenstern [34]), then the problem can be simplified because (i) any decision alternative he faces (with any number of prospects) can be reduced into an equally preferred alternative with only two prospects: the best and the worst. Moreover, (ii) the probability of getting the best prospect in this equivalent binary alternative is equal to the expected utility of that alternative, and (iii) the best decision alternative is the one with the highest expected utility. This result is illustrated graphically in Figure 1. Here, we use the axioms proposed by Howard [24]. These axioms (rules) are as follows:

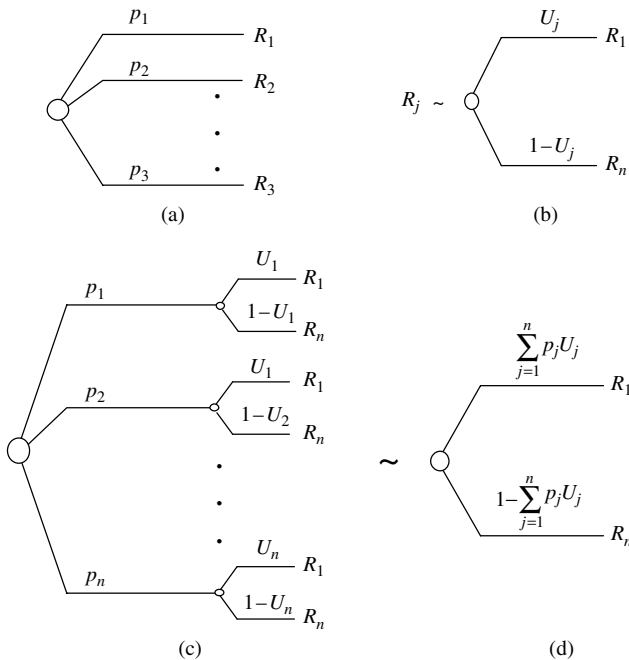
(1) **The Probability Rule.** This rule states that (i) I can characterize each decision alternative I face in terms of the set of possible prospects that may occur, and (ii) I can assign a probability of receiving each prospect for every alternative I choose. Put very generally, this rule states that we can draw the decision tree with all of its nodes and numbers.

(2) **The Order Rule.** This rule states that I can rank order the deterministic prospects that I have characterized in a list from the most preferred to the least preferred. Ties are allowed but preferences must be transitive; i.e., if I prefer prospect A to B, and if I prefer prospect B to C, then I must prefer prospect A to C.

(3) **The Equivalence Rule (von Neumann–Morgenstern utility rule).** This rule states that given three ordered prospects with strict preference, $A \succ B \succ C$, I can assign a preference probability, p_B , that would make me indifferent between receiving prospect B for certain, or a binary deal having a probability p_B of receiving A and $1 - p_B$ of receiving C.

(4) **The Substitution Rule.** This rule states that whenever I face a prospect B, for which I have stated a preference probability p_B of receiving A and $1 - p_B$ of receiving C in the equivalence rule, I would be indifferent between receiving this binary deal and receiving prospect B. This rule allows us to make various substitutions of a binary deal (with prospects A and C) and the deterministic prospect, B, whenever either occurs in the decision tree.

FIGURE 1. Interpretation of expected utility of a lottery as the probability of the best outcome of an equivalent lottery containing only two outcomes, the best and the worst.



(5) **The Choice Rule.** This rule states that if I face two binary decision alternatives, $L1$ and $L2$, both yielding the same prospects A and B (prospect A is preferred to B), and if $L1$ has a higher probability of getting A , then I should choose $L1$ over $L2$.

Figure 1(a) shows a set of prospects, R_1, \dots, R_n , that have been characterized by the probability rule and their corresponding probabilities, p_1, \dots, p_n , for a given decision alternative. The figure also shows the order rule ranking (assume R_1 is the best prospect and R_n is the worst). Figure 1(b) shows an example of a von Neumann and Morgenstern utility assessment, where a preference probability, U_j , is assigned to prospect R_j in terms of the best and worst prospects, R_1 and R_n . Figure 1(c) shows the substitutions made for the prospects in terms of their utilities (preference probabilities), and Figure 1(d) shows the equivalent lottery after multiplying the probabilities in Figure 1(a) by the von Neumann and Morgenstern utilities. Note that the probability of the best prospect in Figure 1(d) is equal to the expected utility of the original alternative, $\sum_{j=1}^n p_j U_j$.

A decision maker who follows the five rules mentioned above will be indifferent between receiving the lotteries of Figures 1(a) and 1(d). Thus, *any decision alternative can be reduced into an equivalent binary alternative using von Neumann and Morgenstern utility (preference probability) assessments.* Furthermore, *the expected utility of the original alternative is equal to the probability of getting the best prospect in the equivalent binary alternative.* Because any other alternative can also be reduced into a binary alternative with the same prospects, R_1 and R_n (but with different probabilities of achieving them), the choice rule determines the best decision alternative: we choose the decision alternative (lottery) with the highest expected utility.

We now observe the following:

1. Expected utility maximization is a direct result of the rules. Not any arbitrary numeric measure representing some form of preferences can be used in this formulation. The measure whose expected value is to be maximized over the set of alternatives must be the von Neumann–Morgenstern utility (that reduces the lottery into an equivalent binary lottery). If we choose to optimize any other arbitrary measure, such as minimizing the “value at risk” or minimizing the “maximum regret,” or any arbitrary “risk measure” or “score,” then we would be violating one of these rules. Which rule would you like to violate in your decision making?

2. Utility functions are not always needed. In some cases (such as *deterministic deals*), we have no uncertainty in the decision problem. In other cases, we may have uncertainty, but also have *deterministic dominance*, where the worst possible prospect that can be obtained by one alternative is better than the best possible prospect obtained by another. We may also have *first order stochastic dominance conditions*, where we always have a higher probability of exceeding any monetary amount with one alternative than with another. *In all these cases, we do not need utility values* because the preference between these lotteries can be determined by the order rule ranking alone (such as preferring more of an attribute to less). Our focus will be the more general case where such dominance conditions do not necessarily exist and where a utility function is indeed required.

3. The same rules apply to multiple attributes. In many cases, prospects can be characterized by multiple attributes, such as health state and wealth. The formulation we discussed above is not restricted to prospects characterized by a single attribute. It simply assigns the utility value to a prospect using the equivalence rule interpretation regardless of how it is characterized.

4. No major restrictions on the functional form of the utility function. This formulation does not pose restrictions on the shape or the functional form of the utility (or multiattribute utility) function, except as imposed by the equivalence rule and that the higher the utility value, the more preferred is the prospect.

1.2. von Neumann–Morgenstern Multiattribute Utility Functions

It is often the case that prospects of a decision can be characterized by more than a single attribute. A medical decision situation, for example, may involve health state and wealth, and the purchase of a new home may involve attributes such as comfort level, view, and price. The same rules we discussed above still apply: we still need to characterize the prospects in terms of the multiple attributes and assign a probability to them, we need to order the prospects, we need to assign a utility value to the prospects, we need to follow the substitution rule, and we need to follow the choice rule.

Now suppose that the prospects are characterized in terms of some level (x, y) of two attributes, X and Y , such as salary (US\$100 K) and vacation time (days). In principle, we can still order the prospects in accordance with the order rule without much difficulty. For example, if the decision maker prefers more vacation and more salary to less, then the prospect $(4, 6)$ is preferred to either $(2, 6)$ or $(4, 3)$. But if we face two prospects, one having a higher salary but less vacation, such as $(4, 6)$ and $(6, 5)$, then the order rule ranking requires a bit more thought. In these situations, it is helpful to use a value function to rank order the prospects.

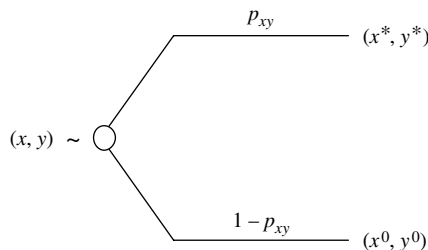
The value function captures the preference ordering of the prospects and maps each prospect (x, y) to some numeric value $V(x, y)$. The higher the magnitude of this value function, the more preferred is the prospect. Thus, the value function helps with the order rule ranking when multiple attributes are present. When there is no uncertainty, this value function determines the best decision alternative. The next section discusses value functions in more detail.

The equivalence rule is also straightforward when multiple attributes are present. For two attributes, X and Y , let (x^*, y^*) be the most preferred prospect we face and (x^0, y^0) be the least preferred. For each prospect (x, y) , the equivalence rule implies that we can state a preference probability, p_{xy} , of receiving the best prospect (x^*, y^*) and a probability $1 - p_{xy}$ of receiving the worst prospect (x^0, y^0) , which makes us indifferent to receiving (x, y) for certain (Figure 2).

If there is a relatively small number of prospects, then we can state a preference probability, p_{xy} , for each prospect individually by direct assessment, and the problem is solved. For salary and vacation time, for example, we would need to assign a preference probability of getting $(2, 3)$ or a binary deal giving $(4, 6)$ or $(0, 0)$. Once these utility values are assigned, the best decision alternative is the one with the highest expected utility. However, if the alternatives we face have numerous prospects (such as a continuum of prospects on the domain of the multiple attributes), then this utility elicitation task may be tedious. We therefore need some tools to help us assign utility values in these situations. The assessment of such multiattribute utility functions will be the focus of the next few sections.

The remainder of this chapter is organized as follows: §2 discusses value functions that rank order deterministic prospects (for either univariate or multiattribute problems). Section 3 shows how to construct a multiattribute utility function once the value function is specified. Section 4 presents a basic expansion theorem that expresses a multiattribute utility function in terms of conditional utility assessments (Abbas [3]). Sections 5 and 6 show

FIGURE 2. Equivalence rule is the same with multiple attributes.



how utility independence conditions (Keeney and Raiffa [26]) and attribute dominance conditions (Abbas and Howard [10]) can simplify the assessments needed for the basic expansion theorem. Section 7 presents bidirectional utility diagrams (graphical representations of the independence relations in multiattribute utility functions). Sections 8 and 9 present methods to assess multiattribute utility functions using single-attribute assessments using utility copula structures (Abbas [7]) and one-switch independence formulations (Abbas and Bell [9]). Section 10 shows how functional equations can also be used to determine the multiattribute utility function, and §11 presents further applications of value functions to multiattribute decision problems.

2. Deterministic Multiattribute Decision Problems (Value Functions)

When there is no uncertainty, each decision alternative leads to a single prospect. The problem is reduced to ranking the prospects in accordance with the order rule. As we discussed, we do not need utility values to determine the best decision alternative in this case because this task can be achieved using a more general value function. This is often referred to as the “deterministic phase of a decision analysis” (Matheson and Howard [29]). In principle, the actual measure that we use in the value function output is not important if we are merely interested in ranking the prospects (provided that the higher the magnitude of the value function, the more preferred is the prospect). Any monotonic transformation of the value function would also suffice. But if we are interested in meaningful increments between prospects to determine other properties, such as willingness to pay to exchange one prospect for another, then a measurable value function (Dyer and Sarin [18]) would be desirable. In particular, a monetary value measure, which allows us to place a dollar value on each prospect, would be convenient.

When the attributes are not all represented monetarily (such as the attribute health state), we need further clarity on the constructed scale that is used. In particular, health state on a numeric scale, for example, is often determined by time trade-off questions that convert health prospects into equivalent healthy life time.

To simplify the construction of the value function, and to eliminate double counting, we need to think carefully about the attributes we use in the characterization of the prospects. It is helpful to make a distinction between “direct and indirect value attributes” (Howard [24]). *Indirect value attributes will not contribute to a change in the order rule ranking of the prospects.*

To further clarify the distinction between direct and indirect values, consider the following example. In buying a house, we may place a direct value on the view, because it affects desirability, and on the total cost associated with it. However, we would only place an indirect value on the cost of maintenance because it is only a part of the total cost, an attribute that we have placed direct value on. In principle, we do not care about the breakdown of the total cost in terms of maintenance, insurance, or utilities, provided the total cost remains the same. If we faced two houses identical in every aspect but having different maintenance costs yet the same total costs, then we would be indifferent between them. Note, however, that if we were very proud to have a house with a low maintenance cost and it would give us great pleasure to tell our friends about it, then we would place a direct value on maintenance cost, and it would be another attribute of direct value. Once the direct value attributes are clarified, a value function is then used to rank the prospects. The best (deterministic) decision alternative is the one whose prospect ranks highest using this value measure.

An appropriate value function is one that captures the qualitative and quantitative features of the problem. The following example from Abbas and Howard [10] presents a value function of a “peanut butter and jelly sandwich.” The attributes involved are thickness of

the slices of bread, b , thickness of the peanut butter, p , thickness of the jelly, j , and the fraction of thickness of peanut butter to jelly, f ,

$$V(p, j, b, f) = V_{\max} \frac{pjbf}{(p^*j^*b^*f^*)^2} (2p^* - p)(2j^* - j)(2b^* - b)(2f^* - f), \quad (1)$$

where p^* , j^* , b^* , and f^* are the optimal values of p , b , j , and f , respectively, and V_{\max} is the dollar amount that the decision maker is willing to pay for the optimal sandwich. Note that the value function is not necessarily increasing with any of the attributes on the entire domain, and that there is dependence in preferences among the attributes. For example, the preference for the thickness of peanut butter will depend on the thickness of jelly.

Contours of prospects providing the same value within the value function are known as “isopreference contours.” We should be indifferent between any two prospects that lie on the same isopreference contour (these prospects rank equally according to the order rule ranking).

The value function also determines the trade-offs between the levels of attributes. For example, we can tell the amount of increase in peanut butter that would make us just indifferent to a decrease in thickness of bread by observing two points on the same isopreference contour. In principle, the increase (or percentage increase) in one attribute that is needed to compensate for a decrease (or percentage decrease) in the other to achieve the same value varies across the domain of the attributes.

One type of value functions that is popular because of its simplicity is the additive value function. This value function implies that the increase in one attribute required to compensate for a decrease in the other is constant across the entire domain of the attributes. For two attributes, this value function has the form

$$V(x, y) = w_x x + w_y y, \quad (2)$$

where w_x, w_y are weight parameters. This value function, however, comes with the strong implication that preferences for levels of attribute X do not depend on the level of attribute Y ; i.e., for two points $(x_1, y_1), (x_2, y_2)$ on the same isopreference contour,

$$\frac{y_2 - y_1}{x_2 - x_1} = -\frac{w_x}{w_y},$$

which means that the increase in attribute X needed to compensate for a decrease in attribute Y does not depend on the levels of either X or Y . Here lies the problem with this value function. Think of two attributes, *wealth* and *amount of air we breathe*. It is obvious that the amount of *wealth* increase needed to compensate for a decrease in *air supply* will depend on the available air we have. The additive value function would not be suitable for this situation. However, it may be suitable for valuing deterministic cash flows for multiple time periods if the discount rate from one year to another is constant.

Another value function that is monotonically increasing with each of the attributes and may be of practical interest is one where the fractional increase in one attribute required to compensate for a fractional decrease in another is constant across the domain. This value function has the form

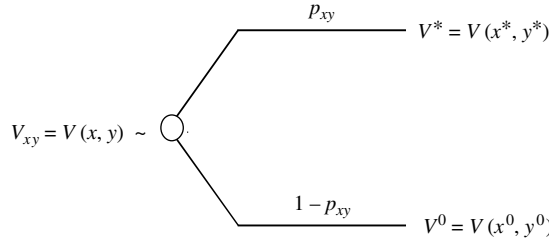
$$V(x, y) = xy^\eta, \quad (3)$$

where η is a trade-off parameter. It can be shown by direct derivatives across an isopreference contour that

$$\eta = -\frac{dx/x}{dy/y}.$$

This also means that for small percentage changes, the percentage increase in attribute X needed to compensate for a percentage decrease in attribute Y is constant.

FIGURE 3. Equivalence rule using a value function for multiple measures.



3. Constructing Multiattribute Utility Functions Using Value Functions

3.1. The Chain Rule for Risk Aversion (Dyer and Sarin [19], Matheson and Abbas [28])

Although the value function is sufficient to determine the best decision alternative in deterministic decision problems, we have seen that we need to assign utility values to the prospects when uncertainty is present to maximize the expected utility. This section illustrates how to assign a multiattribute utility function to the prospects once the value function is specified.

Let $V_{xy} = V(x, y)$ be the magnitude of the value function for any prospect (x, y) , and let $V^* = V(x^*, y^*)$ and $V^0 = V(x^0, y^0)$ be the magnitude of the value function for the most and least preferred prospects, respectively. The preference probability for any prospect (x, y) can now be determined by simply asking the decision maker for his indifference between the two deals in Figure 3. Note that the decision maker need not even be concerned about the actual values of X and Y . This assessment is simply a utility function (preference probability) assessment over a single measure, V , in terms of the best value, V^* , and the worst, V^0 . Note that *this measure need not be money*, but as is well known, money is simpler to assign a utility function to in comparison to other constructed scales. It also enables the calculation of the certain equivalent and the value of information on the uncertainties of interest.

Let U_V be the utility function we assign over the value measure. Given a prospect (x, y) whose value measure is $V(x, y)$, the utility of (x, y) , namely, $U(x, y)$, is equal to

$$U(x, y) = U_V(V(x, y)). \tag{4}$$

The problem is now reduced to a one-dimensional utility assessment over the value measure (Matheson and Abbas [28]).

Having a value function limits the individual utility functions that we can assign over the individual attributes. For example, if we have assigned a value function, $V(x, y)$, and a utility function over value, $U_V(V)$, then the conditional utility function of attribute X at the maximum value of attribute Y must be of the form

$$U(x, y^*) = U_V(V(x, y^*)).$$

Similarly, the conditional utility function at any other instantiation of Y can be obtained by direct substitution. We do not have another degree of freedom to assign individual utility functions over the attributes once the value function and the utility function over value have been specified.

The value function also relates the risk aversion functions (Arrow [14], Pratt [31]) among the different attributes. The risk aversion function over a single measure, V , is the negative ratio of the second to the first derivative of the utility function,

$$\gamma_V^U \triangleq - \frac{U_V''(V)}{U_V'(V)}. \tag{5}$$

We can use this relation to relate the risk aversion over value to the risk aversion over any attribute. Using the chain rule for partial derivatives, we take the first derivative of (5) with respect to x to get

$$\frac{\partial U(x, y)}{\partial x} = \frac{\partial U_V(V(x, y))}{\partial x} = \frac{\partial U_V(V(x, y))}{\partial V} \frac{\partial V(x, y)}{\partial x} = U'_V \frac{\partial V(x, y)}{\partial x}. \quad (6)$$

Now we use the chain rule again and evaluate the derivative of (6) with respect to x :

$$\begin{aligned} \frac{\partial^2 U(x, y)}{\partial x^2} &= \frac{\partial^2 U_V(V(x, y))}{\partial V^2} \left(\frac{\partial V(x, y)}{\partial x} \right)^2 + \frac{\partial U_V(V(x, y))}{\partial V} \frac{\partial^2 V(x, y)}{\partial x^2} \\ &= U''_V \left(\frac{\partial V(x, y)}{\partial x} \right)^2 + U'_V \frac{\partial^2 V(x, y)}{\partial x^2}. \end{aligned} \quad (7)$$

Substituting (6) and (7) into (5) gives the risk aversion function for attribute, X , in terms of the risk aversion function over the value measure,

$$\begin{aligned} \gamma_x^U &= -\frac{\partial^2 U(x, y)/\partial x^2}{\partial U(x, y)/\partial x} = -\frac{U''_V(\partial V(x, y)/\partial x)^2 + U'_V(\partial^2 V(x, y)/\partial x^2)}{U'_V(\partial V(x, y)/\partial x)} \\ &= -\gamma_V^{U_V} \frac{\partial V}{\partial x} - \frac{\partial^2 V(x, y)/\partial x^2}{\partial V(x, y)/\partial x}. \end{aligned} \quad (8)$$

For notational convenience, define γ_x^V as

$$\gamma_x^V \triangleq -\frac{\partial^2 V(x, y)/\partial x^2}{\partial V(x, y)/\partial x}. \quad (9)$$

Substituting Equation (9) into Equation (8) gives the “chain rule for risk aversion,”

$$\gamma_x^U = \gamma_V^U \frac{\partial V}{\partial x} + \gamma_x^V. \quad (10)$$

Equation (10) expresses the risk aversion function of an attribute X as a sum of two terms. The first term, $\gamma_V^U(\partial V/\partial x)$, is a product of the risk aversion function of utility toward value and the partial derivative of the value function with respect to the attribute of interest. The partial derivative, $\partial V/\partial x$, is in effect a unit conversion factor. The second term, γ_x^V , is the value function contribution to the risk aversion of the attribute of interest. We note here that Dyer and Sarin [19] used an equation similar to (10) for single attributes and defined the difference $\gamma_x^U - \gamma_x^V$ as the “relative risk version” of that attribute.

3.2. Utility Transversality (Matheson and Abbas [28])

We have assessed a utility function over the value measure and related it to the utility function over an attribute. If one of the attributes is expressed in units that are easy to relate to, such as money, but the value measure itself is not money, then we can assess the utility function over this monetary attribute to construct the utility function. We now relate the utility functions and risk aversion functions across the individual attributes.

If we assess a risk aversion function for one attribute, say, X , then we can rearrange Equation (10) for the risk attitude on value to obtain

$$\gamma_x^U - \gamma_x^V = \gamma_V^U \frac{\partial V}{\partial x}. \quad (11)$$

Similarly, for any other attribute Y , we get

$$\gamma_y^U - \gamma_y^V = \gamma_V^U \frac{\partial V}{\partial y}. \quad (12)$$

From (11) and (12),

$$\gamma_x^U = [\gamma_y^U - \gamma_y^V] \frac{\partial V / \partial x}{\partial V / \partial y} + \gamma_x^V. \tag{13}$$

Equation (13) may seem difficult to interpret at first, but we can simplify this equation by observing that the value function is constant along an “isopreference contour,”

$$V(x, y) = \text{constant along an isopreference contour.} \tag{14}$$

As a consequence, the total derivative of the value function across an isopreference contour is zero. Hence,

$$dV(x, y) \triangleq \frac{\partial V(x, y)}{\partial x} dx + \frac{\partial V(x, y)}{\partial y} dy = 0. \tag{15}$$

Rearranging (15) gives the trade-off between the two attributes x and y as

$$\left. \frac{\partial V(x, y) / \partial x}{\partial V(x, y) / \partial y} = - \frac{dy}{dx} \right|_{\text{isopreference contour}} \triangleq t(x, y), \tag{16}$$

where $t(x, y)$ defined above is the deterministic trade-off function between attributes y and x along an isopreference contour. Substituting (16) into (13) yields

$$\gamma_x^U(x, y) = [\gamma_y^U(x, y) - \gamma_y^V(x, y)]t(x, y) + \gamma_x^V(x, y). \tag{17}$$

To simplify the notation, we will drop the (x, y) terms from Equation (17) and rearrange to get

$$(\gamma_x^U - \gamma_x^V) = (\gamma_y^U - \gamma_y^V)t. \tag{18}$$

Equation (18) is a fundamental expression that relates the risk aversion functions across the different attributes to the trade-off function between them. Matheson and Abbas [28] introduce this relation and refer to it as “utility transversality.” The utility transversality relation asserts that the ratio of the relative risk aversion function of two attributes is equal to the trade-off function between them.

Example 1 (On Fates Comparable to Death). The following example is adapted from Howard [23] on making trade-offs concerning situations where a decision maker is exposed to fates comparable to death (such as outcomes of medical surgery). The decision maker 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 decision maker assigns a value function over consumption and health states as

$$V(x, y) = xy^\eta, \tag{19}$$

where x is expressed in dollars, y is the health state, and η determines the trade-off between x and y . The value function units are total well-being. One method to determine the multi-attribute utility function is to assess a utility function, $U_V(V)$, over the value function. For example, if $U_V(V)$ were exponential, then

$$U_V(V) = U_V(V(x, y)) = 1 - e^{-\gamma V(x, y)} = 1 - e^{-\gamma xy^\eta}, \tag{20}$$

where $\gamma = \gamma_V^U$ is the risk aversion coefficient toward value.

Using the chain rule for risk aversion of Equation (10), we can now deduce the risk aversion function toward each of the attributes of the decision problem. For example, in the case of attribute X , we have $\partial V/\partial x = y^\eta$ and $\gamma_x^V = 0$. Equation (10) reduces to

$$\gamma_x^U = \gamma_V^U \frac{\partial V}{\partial x} + \gamma_x^V = \gamma y^\eta. \tag{21}$$

For attribute Y , we have $\partial V/\partial y = \eta xy^{\eta-1}$ and $\gamma_y^V = -(\eta - 1)/y$,

$$\gamma_y^U = \gamma_V^U \frac{\partial V}{\partial y} + \gamma_y^V = \gamma \eta xy^{\eta-1} - \frac{\eta - 1}{y}. \tag{22}$$

Note that although the utility function over value exhibited constant absolute risk aversion, neither of the risk aversion expressions for an attribute in (21) or (22) is a constant. The following example from Matheson and Abbas [28] highlights some recurring misconceptions about discounting value versus discounting utility.

Example 2 (Discounting Value Not Utility). Consider a two-period cash flow and a value function with a discount factor, β , of the form

$$V(x, y) = x + \beta y,$$

where x is the amount received in the current year (no discounting), and y is the amount received in one year (with a discount factor). In the formulation of multiperiod cash flows, we are concerned with risk aversion, time preference, as well as the value function.

Suppose the decision maker has an exponential utility function over net present value. Then,

$$U(x, y) = -e^{-\gamma V(x, y)} = -e^{-\gamma(x + \beta y)} = -e^{-\gamma x} e^{-\gamma \beta y}.$$

This utility function is the product of two utility functions. The first has a risk aversion coefficient, γ , and the second has a risk aversion coefficient $\gamma\beta$. Thus, the risk aversion for the next year is multiplied by the discount factor. Note here we did not discount the utility function; we discounted the value of the cash flow a year from now (this message has been echoed by several authors; see, for example, Baucells and Sarin [15]). Moreover, the resulting two-attribute utility function does not need to have the product form (nothing in the equivalence rule dictates this). For example, if the decision maker had a logarithmic utility function over net present value, then the two-attribute utility function would be

$$U(x, y) = \log V(x, y) = \log(x + \beta y).$$

4. Basic Expansion Theorem for Multiattribute Utility Functions (Abbas [3, 6])

Some people may prefer to construct the utility function without the use of a value function by making direct conditional utility assessments over each attribute at certain instantiations of the remaining (complement) attributes. As we discussed, this degree of freedom can only occur if the value function is not specified. The remainder of this chapter is concerned with this assessment method as well as the derivation of analytical expressions that may be used as consistency checks with the value function approach. First, we observe that an analogous problem to this problem of expressing the multiattribute utility function in terms of conditional utility assessments exists in probability theory when we express a joint probability function in terms of marginal-conditional probability assessments. In probability theory, we have Bayes' expansion theorem that tells us the conditional probability assessments needed to construct a joint probability distribution in a given assessment order. This section derives an analogous expansion theorem for multiattribute utility functions. We follow Abbas [3, 6] to derive this basic expansion theorem and then illustrate its use in constructing multiattribute utility functions with conditional utility assessments.

4.1. Basic Definitions

We consider multiattribute utility functions with n attributes, X_1, X_2, \dots, X_n . We use the lowercase, x_i , to denote an instantiation of attribute X_i , and use the vector (x_1, \dots, x_n) to represent a prospect of the decision situation defined on a domain, \mathbb{D} , that forms a connected subset of \mathbb{R}^n . We use \bar{X}_i to denote the set of complement (remaining) attributes to X_i , and use \bar{x}_i to denote a certain instantiation of this complement. For simplicity, we also refer to a consequence, (x_1, \dots, x_n) , as (x_i, \bar{x}_i) , or (x_i, x_j, \bar{x}_{ij}) , where \bar{x}_{ij} is an instantiation of the complement attributes to X_i and X_j . We assume that the multiattribute utility function, $U(x_1, \dots, x_n)$ is (i) continuous and (ii) bounded, and (iii) for each attribute, X_i , there exist two reference values, x_i^0 and x_i^* , for which $U(x_i^*, \bar{x}_i) > U(x_i^0, \bar{x}_i)$ at all reference values of the complement attributes. We refer to this class of utility functions as class S . Because the utility function is bounded, we assume hereafter, with no further loss of generality, that it is normalized to range from zero to one. We also define a conditional utility function for attribute X as

$$U(x | \bar{x}) = \frac{U(x, \bar{x}) - U(x^0, \bar{x})}{U(x^*, \bar{x}) - U(x^0, \bar{x})}, \tag{23}$$

and note that $U(x^0 | \bar{x}) = 0$ and $U(x^* | \bar{x}) = 1$.

If x_i^0, x_i^* are the minimum and maximum values of X_i , respectively, and, furthermore, if $U(x_i, \bar{x}_i)$ lies between $U(x_i^0, \bar{x}_i)$ and $U(x_i^*, \bar{x}_i)$, then we maintain the von Neumann and Morgenstern normalized indifference probability interpretation for the conditional utility function across the entire domain of the argument, x_i . Although these two conditions are not essential in our formulation, they simplify the elicitation of the conditional utility function and provide a meaningful interpretation for its assessment. We therefore refer to the instantiations x_i^0 and x_i^* as the minimum and maximum reference values of X_i , respectively, while keeping in mind that they can be set arbitrarily. Finally, we define a normalized “conditional disutility function,” $\bar{U}(x | \bar{x})$, as

$$\bar{U}(x | \bar{x}) = 1 - U(x | \bar{x}). \tag{24}$$

4.2. Expansion of a Utility Function Around a Single Attribute

To express a multiattribute utility function in terms of the normalized conditional assessment, $U(x | \bar{x})$, of attribute X , we rearrange (23) and substitute into (24) to get

$$U(x, \bar{x}) = U(x^*, \bar{x})U(x | \bar{x}) + U(x^0, \bar{x})\bar{U}(x | \bar{x}). \tag{25}$$

We refer to this expression as “expanding the utility function around X .” The pattern associated with (25) is that when X appears as x^* in the term $U(x^*, \bar{x})$, it is followed by the conditional assessment of X on its complement. When it appears as x^0 in $U(x^0, \bar{x})$, it is followed by the conditional disutility of X on its complement.

4.3. Expansion of a Utility Function Around Multiple Attributes

Three basic terms result from the expansion (25): $U(x^*, \bar{x})$, $U(x^0, \bar{x})$, and $U(x | \bar{x})$. Suppose we further expand the two terms $U(x^*, \bar{x})$ and $U(x^0, \bar{x})$ with respect to a second attribute, Y (note that we do not yet expand the third term $U(x | \bar{x})$). This formulation would imply

$$U(x^*, \bar{x}) = U(x^*, y^*, \bar{x}y)U(y | x^*, \bar{x}y) + U(x^*, y^0, \bar{x}y)\bar{U}(y | x^*, \bar{x}y) \quad \text{and} \tag{26}$$

$$U(x^0, \bar{x}) = U(x^0, y^*, \bar{x}y)U(y | x^0, \bar{x}y) + U(x^0, y^0, \bar{x}y)\bar{U}(y | x^0, \bar{x}y). \tag{27}$$

Substituting from (26) and (27) into (25) gives

$$U(x, y, \bar{x}y) = U(x^*, y^*, \bar{x}y)U(x | \bar{x})U(y | x^*, \bar{x}y) + U(x^*, y^0, \bar{x}y)U(x | \bar{x})\bar{U}(y | x^*, \bar{x}y) \\ + U(x^0, y^*, \bar{x}y)\bar{U}(x | \bar{x})U(y | x^0, \bar{x}y) + U(x^0, y^0, \bar{x}y)\bar{U}(x | \bar{x})\bar{U}(y | x^0, \bar{x}y). \tag{28}$$

We have chosen to write (28) without any simplification to illustrate the pattern that is evolving, and to highlight some important points about this expansion:

(1) The terms $U(x^*, y^*, \bar{x}\bar{y})$, $U(x^*, y^0, \bar{x}\bar{y})$, $U(x^0, y^*, \bar{x}\bar{y})$, $U(x^0, y^0, \bar{x}\bar{y})$ in the right-hand side require a joint utility assessment of the attributes that have not been expanded, namely, $\bar{x}\bar{y}$, at all possible combinations of the minimum and maximum values of the attributes that have. We refer to these terms as the *unexpanded terms* and write them as $U(x_K^{*0}, \bar{x}_K)$.

(2) When the instantiation of X is x^* in the unexpanded terms, it is followed by a conditional utility assessment, $U(x | \bar{x})$, and when its instantiation is x^0 , it is followed by a conditional disutility assessment, $\bar{U}(x | \bar{x})$. The same applies to attribute Y and any other attribute we wish to expand.

(3) Equation (28) implies an expansion order: we first expand around X and then Y .

(4) The first attribute that is expanded is always conditioned on its complement. The second attribute is also conditioned on its complement, but the first term in the conditioning, X , has a value equal to its instantiation in the unexpanded terms.

(5) Equation (28) can also be written in a more compact form as

$$U(x, y, \bar{x}\bar{y}) = \sum_{(x, y)^{*0} \in \{X, Y\}^{*0}} U((x, y)^{*0}, \bar{x}\bar{y}) g_x(x | \bar{x}) g_y(y | x^{*0}, \bar{x}\bar{y}), \quad (29)$$

where the summation is carried out over all possible instantiations of the set $\{X, Y\}^{*0} \triangleq \{(x^*, y^*), (x^*, y^0), (x^0, y^*), (x^0, y^0)\}$, the term $(x, y)^{*0}$ represents an element of the set $\{X, Y\}^{*0}$, x^{*0} is the instantiation of x in $(x, y)^{*0}$, and $g_x(x | \bar{x})$, $g_y(y | x^{*0}, \bar{x}\bar{y})$ are either conditional utility functions or disutility functions depending on the instantiations of x , y in the terms, $U((x, y)^{*0}, \bar{x}\bar{y})$. In principle, we can expand around as many or as few attributes as we wish, and we can change the order of the expansion (because the derivation is not specific to a *particular* expansion order). For example, it must also be true that

$$U(x, y, \bar{x}\bar{y}) = \sum_{(x, y)^{*0} \in \{X, Y\}^{*0}} U((x, y)^{*0}, \bar{x}\bar{y}) g_y(y | \bar{y}) g_x(x | y^{*0}, \bar{x}\bar{y}), \quad (30)$$

which yields the exact same functional form although it is expressed in terms of different conditional assessments. To generalize these results, define X_K as the set of attributes that have been expanded, X_K^{*0} as the set of all possible instantiations of X_K when the attributes are at a minimum or a maximum value, and $U(x_K^{*0}, \bar{x}_K)$ as a joint utility assessment of \bar{X}_K at an instantiation of $x_K^{*0} \in X_K^{*0}$. Furthermore, define X_{iP} as the set of nodes expanded prior to X_i , x_{iP}^{*0} as the instantiation of x_{iP} in the term $U(x_K^{*0}, \bar{x}_K)$, and X_{iF} (the set of future nodes) as the set of complement attributes to $\{X_i\} \cup X_{iP}$. Finally, let

$$g(x_i | x_{iP}^{*0}, x_{iF}) = \begin{cases} U(x_i | x_{iP}^{*0}, x_{iF}), & \text{if } x_i = x_i^* \text{ in the unexpanded term } U(x_K^{*0}, \bar{x}_K); \\ \bar{U}(x_i | x_{iP}^{*0}, x_{iF}), & \text{if } x_i = x_i^0 \text{ in the unexpanded term } U(x_K^{*0}, \bar{x}_K). \end{cases} \quad (31)$$

By induction, we now state the following Theorem from Abbas [3, 6].

Theorem 1 (Basic Expansion Theorem for Multiattribute Utility). *Any multiattribute utility function $U(x) \in S$ can be expressed as*

$$U(x) = \sum_{x_K^{*0} \in X_K^{*0}} U(x_K^{*0}, \bar{x}_K) \prod_{X_i \in X_K} g(x_i | x_{iP}^{*0}, x_{iF}). \quad (32)$$

This basic expansion theorem does for utility functions what Bayes' expansion theorem does for probability functions, and expresses the multiattribute utility function in terms of conditional utility assessments. Of course, assessing the conditional utility functions in their current form may not be an easy task (cognitively), because a function like $U(x_n | x_1, \dots, x_{n-1})$, for example, is an n -variate assessment.

The next three sections discuss methods to simplify the assessments needed for the basic expansion theorem. One approach uses what is known as an *independence* assumption to reduce the order of the required assessments into single-attribute functions, thereby simplifying the assessment task (cognitively). Another approach uses what is known as an *attribute dominance* assumption to reduce the number of assessments needed. A third approach considers partial utility independence conditions and a *bidirectional utility diagram* that reduces the functional form even when full utility independence conditions among the attributes do not exist.

5. Simplifying Conditional Utility Assessments with Independence Assertions (Keeney and Raiffa [26])

One way to simplify the conditional utility assessments and the functional form of the basic expansion theorem is to condition the assessments on a smaller number of attributes. To do that, we can use what is known as a utility independence assumption (Keeney and Raiffa [26]). A set of attributes X_K is *utility independent of another set X_I given X_D* , written $(X_K \text{ UI } X_I | X_D)$, with $X_I \cup X_D = \bar{X}_K$, if preferences for lotteries over X_K do not depend on the instantiations of X_I . This definition is equivalent to saying that $(X_K \text{ UI } X_I | X_D)$ implies that the function that maps $x_K \rightarrow U(x_K, x_I, x_D)$ changes by only a linear transformation as x_I changes,

$$U(x_K, x_I, x_D) = l(x_I, x_D) + d(x_I, x_D)h(x_K, x_D), \tag{33}$$

for some functions l, d, h .

Note that if every attribute is utility independent of its complement, then we have a set of equations of the form (33) for each attribute. Moreover, from (23), this implies that the conditional utility assessment of any attribute satisfies

$$U(x | \bar{x}) = U(x | \bar{x}^0), \tag{34}$$

a single-attribute assessment at a fixed instantiation of the complement. Therefore, the basic expansion theorem reduces to a multilinear combination of single-attribute utility assessments. Keeney and Raiffa [26] show that this case leads to the multilinear form

$$U(x_1, \dots, x_n) = \sum_{i=1}^n k_i U_i(x_i) + \sum_{i=1}^n \sum_{j>i} k_{ij} U_i(x_i) U_j(x_j) + \dots + k_{123, \dots, n} U_1(x_1) U_2(x_2), \dots, U_n(x_n), \tag{35}$$

where the scaling constants satisfy

$$k_i = U(x_i^*, \bar{x}_i^0), \quad i = 1, \dots, n$$

$$k_{ij} = U(x_i^*, x_j^*, \bar{x}_{ij}^0) - k_{x_i} - k_{x_j}, \quad i = 1, \dots, n, j > i$$

$$k_{ijk} = U(x_i^*, x_j^*, x_l^*, \bar{x}_{ijl}^0) - k_{x_i x_j} - k_{x_i x_l} - k_{x_j x_l} - k_{x_i} - k_{x_j} - k_{x_l}, \quad i = 1, \dots, n, l > j > i$$

$$k_{ij, \dots, n} = U(x_i^*, x_j^*, \dots, x_n^*) - \sum_i k_{x_1, \dots, x_{i-1} x_{i+1}, \dots, n} - \dots - \sum_i \sum_{j>i} k_{x_i x_j} - \sum_i k_{x_i}.$$

A stronger independence condition also holds if every subset of the attributes is utility independent of its complement and the utility function has either the multiplicative form

$$1 + kU(x_1, \dots, x_n) = \prod_{i=1}^n [1 + k k_i U_i(x_i)] \tag{36}$$

or the additive form

$$U(x_1, \dots, x_n) = \sum_{i=1}^n k_i U_i(x_i). \quad (37)$$

The simplicity of constructing multiattribute utility functions using single-attribute utility assessments has popularized the use of the functional forms (35), (36), and (37) in many decision analysis applications, but they do come with strong implications. Note that the assumption of mutual utility independence for two attributes, *health* and *wealth*, as an example, would imply that our preferences for decision alternatives leading to different health states will not change with our wealth state, and that our preference for decision alternatives leading to lotteries over money will not change with our health state.

If utility independence conditions do not hold, then we are faced with the need to incorporate utility dependence among the attributes into our analysis as implied by the basic expansion theorem. Several authors have discussed this issue and have presented methods to incorporate utility dependence (see, for example, Fishburn [21, 22], Farquhar [20], Kirkwood [27], Bell [16], Keeney and Raiffa [26], Keeney [25], Abbas and Howard [10]). In particular, Farquhar [20] proposes a decomposition theorem for a class of multiattribute utility functions with certain utility independence structures known as fractional hypercubes. He shows that special cases of this formulation reduce to the multilinear form. Kirkwood [27] incorporates utility dependence by defining “parametrically dependent” conditional utility functions. Bell [16] incorporates utility dependence by assessing conditional utility functions at the boundaries and deriving the remaining utility values by interpolation. Keeney [25] takes an alternate approach and reformulates the attributes of a decision problem to derive a new set of attributes with a lesser degree of utility dependence.

6. Using an Attribute Dominance Condition (Abbas [2])

Another way to simplify the construction of the utility function using the basic expansion theorem is to reduce the number of conditional utility assessments it requires. This number can be reduced significantly when an “attribute dominance” condition exists. To illustrate this basic idea, suppose that the multiattribute utility function is minimum (zero) if any one of the attributes is at its minimum value. For example, in a decision involving health state and consumption, if either the health state or consumption is a minimum, then the utility function is minimum. Leg space, mileage, and comfort level may also be considered utility dominant attributes for some specified minimum levels. If all attributes are utility dominant, then all terms having an instantiation of an attribute at its minimum value are zero and the basic expansion theorem reduces to the product form

$$U(x) = U(x_K^*, \bar{x}_K) \prod_{X_i \in X_K} U(x_i | x_{iP}^*, x_{iF}),$$

which is similar in form to Bayes’ expansion theorem using products of conditional assessments. If some of the attributes are utility dominant and others are not, then we have a partial sum of products in the basic expansion theorem instead of this simple product.

Of particular interest to medical decision situations with two attributes, *health* and *wealth*, suppose that attribute dominance conditions exist. Define the marginal utility function over a single attribute, X , as the numerical value of the utility function when all other attributes are set at their maximum values. For two attributes we have (we use the superscript d to indicate a utility function corresponding to a utility dominant attribute)

$$U_x^d(x) \triangleq U_{xy}^d(x, y_{\max}). \quad (38)$$

The conditional utility function, $U_{y|x}^d(y|x)$, for attribute Y reduces to

$$U_{y|x}^d(y|x) \triangleq \frac{U_{xy}^d(x,y)}{U_{xy}^d(x,y_{\max})} = \frac{U_{xy}^d(x,y)}{U_x^d(x)}. \tag{39}$$

Rearranging (39) gives a simplified form of the basic expansion theorem as the product

$$U_{xy}^d(x,y) = U_x^d(x)U_{y|x}^d(y|x). \tag{40}$$

Equation (40) allows the construction of attribute dominance utility functions using marginal-conditional utility assessments analogous to the marginal-conditional approach of joint probability distributions. As we shall see, this decomposition enables many of the tools that were developed to incorporate probability dependence to be applied to utility dependence. In a similar manner, we can also define $U_{x|y}^d(x|y)$ as

$$U_{x|y}^d(x|y) \triangleq \frac{U_{xy}^d(x,y)}{U_y^d(y)}. \tag{41}$$

Combining Equations (39) and (41), we have

$$U_{x|y}^d(x|y) = \frac{U_{y|x}^d(y|x)U_x^d(x)}{U_y^d(y)}. \tag{42}$$

It is important to pause here and think about the interpretation and implications of Equation (42). In the analogous world of probability inference, we update our state of information about the outcome of an uncertain event by conditioning on the outcome of another event. This updating of information can be made even if there is no decision to be made, and often helps us think more clearly about our degree of belief on the outcomes of uncertain events. In a similar manner, Equation (42) provides a method to update our utility values over an attribute when we are guaranteed a fixed amount of another attribute. It is natural to extend our analogy and refer to Equation (42) as “Bayes’ rule for utility inference.” The following example demonstrates the elicitation of attribute dominance utility functions using the marginal-conditional approach and an application of Bayes’ rule for utility inference.

Example 3 (An Attribute Dominance Utility Function for a Medical Decision).

A patient undergoing a cancer treatment is deciding whether to have chemotherapy or radiotherapy. The two attributes involved are health state, Y (measured by the quality of life on a normalized scale from zero to one), and consumption levels, X (measured in millions of dollars on a normalized scale from zero to one). The patient’s preferences are such that when any of these attributes has a minimum value, the prospect is considered a least preferred prospect.

To assess the multiattribute utility function, we start by assessing the marginal utility function for consumption. Let us assume that the decision maker is risk neutral over the range of wealth levels considered, when the quality of life is at a maximum.

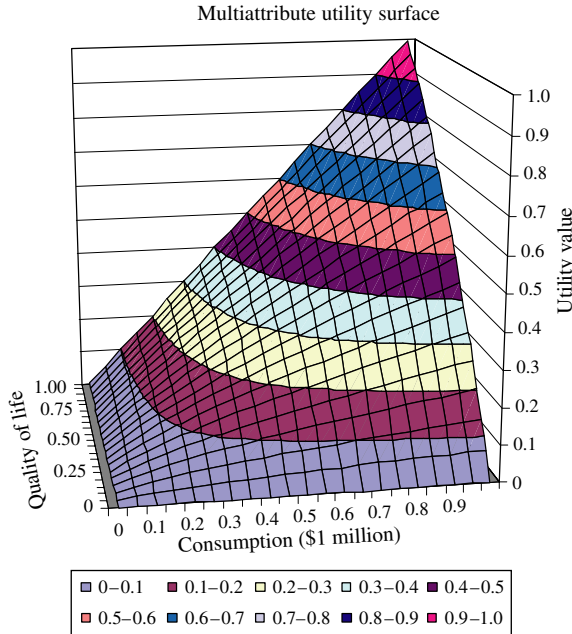
$$U_x^d(x) = x, \quad 0 \leq x \leq 1, \tag{43}$$

where x is expressed in millions of dollars.

Now we assess the conditional utility function for quality of life given consumption. The decision maker states he is risk averse over the quality of life, with a risk aversion function that depends on the value of the wealth attribute. The decision maker assigns an exponential conditional utility function for quality of life given wealth as

$$U_{y|x}^d(y|x) = \frac{1 - e^{-\gamma(x)y}}{1 - e^{-\gamma(x)}} \quad 0 \leq y \leq 1, 0 \leq x \leq 1, \tag{44}$$

FIGURE 4. Attribute dominance utility surface for two attributes.



where y is the quality of life, $\gamma(x) = 1/(\rho_0 + x)$ is the risk aversion function, and ρ_0 is a constant chosen to be US\$0.3 million in this example.

The conditional utility function of Equation (44) is normalized to range from zero to one and can thus be assessed directly without the need for any further scaling of the utility function or assessing its boundary values. The multiattribute utility function obtained using Equation (40) is

$$U_{xy}^d(x, y) = \frac{x(1 - e^{-(\gamma(x))y})}{1 - e^{-(\gamma(x))}}, \quad 0 \leq y \leq 1, 0 \leq x \leq 1. \quad (45)$$

Figure 4 shows the multiattribute utility function for the decision maker's prospects. Now we determine the marginal utility function for quality of life as

$$U_y^d(y) \triangleq U_{xy}^d(x_{\max}, y) = \frac{1 - e^{-y/1.3}}{1 - e^{-1/1.3}}, \quad 0 \leq y \leq 1, \quad (46)$$

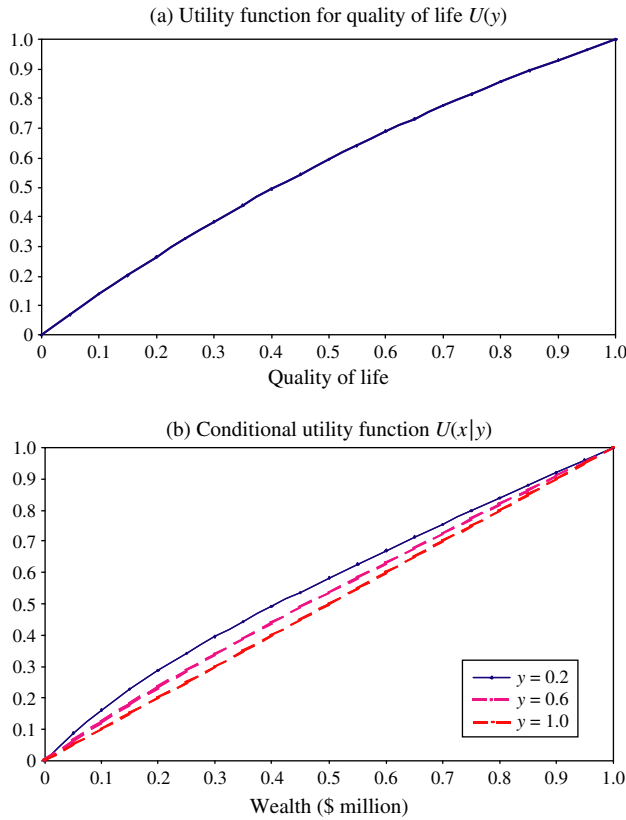
where $\gamma(x_{\max}) = 1/(0.3 + 1) = 1/1.3$. Using Bayes' rule for utility inference, we can also determine the conditional utility function for wealth given health state as

$$U_{x|y}^d(x | y) = \frac{U_{xy}^d(x, y)}{U_y^d(y)} = \frac{x(1 - e^{-1/1.3})(1 - e^{-(\gamma(x))y})}{(1 - e^{-y/1.3})(1 - e^{-(\gamma(x))})}, \quad 0 \leq y \leq 1, 0 \leq x \leq 1. \quad (47)$$

Figure 5(a) shows the marginal utility function for quality of life obtained by utility inference, and Figure 5(b) shows the conditional utility functions for wealth given different values of y using Bayes' rule for utility inference. Note that the decision maker is risk neutral for wealth only when the quality of life, Y , is at a maximum. For any other value of Y , the decision maker is risk averse for wealth.

If we use the product form of the marginal utility functions in our analysis, then the utility function for wealth would not update with different values of quality of life. The previous analysis thus provides the flexibility to model more representative utility functions that have the same marginal utility function but incorporate utility dependence between the attributes.

FIGURE 5. (a) Marginal utility function, $U_y^d(y)$, and (b) conditional utility function, $U_{x|y}^d(x|y)$.



6.1. Utility Diagrams for Attribute Dominance Utility Functions

Abbas [2] shows that utility independence is symmetric for attribute dominance utility functions (as is the case with probabilistic independence for joint probability functions). Therefore, we can construct “utility diagrams” to model the independence relations in attribute dominance utility functions that are analogous in form to relevance diagrams for joint probability functions. Now let us start with the simplest utility diagram and assume that we have a decision situation with two attributes, A and B , having utility independence. We represent this situation as a diagram with two nodes and no arrows connecting them (Figure 6). The absence of the arrow asserts that the two attributes are utility independent given our current state of preference.

Similar to probability independence, the presence of utility independence can greatly simplify the elicitation of attribute dominance utility functions. The multiattribute utility

FIGURE 6. Two attributes with utility independence.

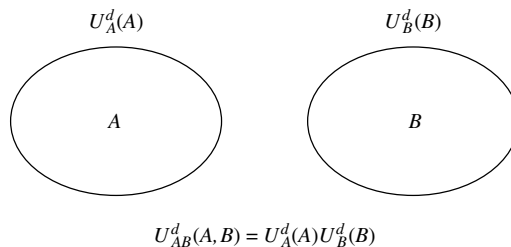
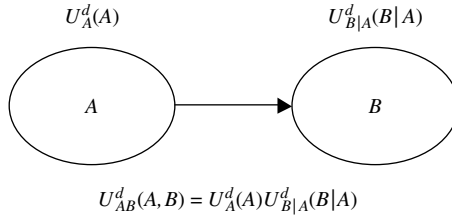


FIGURE 7. Two attributes with possible utility dependence.



function in this case is the product of the marginal utility functions of the attributes:

$$U_{(x_1, \dots, x_n)}^d = \prod_{i=1}^n U_{x_i}^d(x_i). \tag{48}$$

When we have not asserted utility independence between two attributes, we add an arrow between their nodes in the utility diagram. The arrow implies the *possibility* of utility dependence between the attributes given our current state of preferences (Figure 7).

This convention is often helpful in the early stages of the multiattribute utility assessment. For example, if we have not confirmed the utility independence assumptions between the attributes, we add an arrow between the two nodes and verify it later during the utility assessments. As we have seen, when the utility dominant attributes have utility dependence, the attribute dominance utility function takes the form

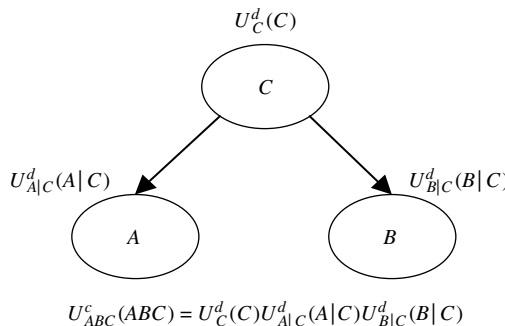
$$U_{x_1, \dots, x_n}^d = U_{x_1}^d(x_1)U_{x_2|x_1}^d(x_2|x_1) \cdots U_{x_n|x_1, \dots, x_{n-1}}^d(x_n|x_1, \dots, x_{n-1}). \tag{49}$$

If two attributes, A and B , have conditional utility independence given a third attribute, C , then we can represent this relationship using the utility diagram shown in Figure 8. The attribute dominance utility function takes the form

$$U_{ABC}^d(A, B, C) = U_C^d(C)U_{B|C}^d(B|C)U_{A|C}^d(A|C). \tag{50}$$

Similar to the arrows in a relevance diagram, the arrows in the utility diagram represent a given utility assessment order. It may sometimes be useful to change this assessment order into one that is more comfortable to the decision maker to simplify the utility assessments or to verify the utility independence assumptions that have been made. To do this, we need a set of rules for inference and arrow reversals in utility diagrams. Analogous to relevance diagrams, the rules of arrow reversals in utility diagrams are as follows: (i) you can add an arrow between any two nodes of the utility diagram provided you do not create a cycle, and (ii) you can flip an arrow between two nodes provided they are conditioned on the same attributes. Furthermore, the arbitrary removal of an arrow from the utility diagram is not permitted because it asserts the existence of some utility independence conditions between the attributes that may not be available.

FIGURE 8. Three attributes with conditional utility independence.



7. Bidirectional Utility Diagrams (Abbas [3, 6])

We have seen how utility independence and attribute dominance conditions simplify the construction of the multiattribute utility function using the basic expansion theorem. This section considers situations where not all the attributes are utility independent of their complement and the attributes need not all be utility dominant. Because we relax the attribute dominance conditions, utility independence assertions will not necessarily be symmetric. Therefore, the utility diagrams we discussed in the previous section would need to model this asymmetric independence relationship.

Note that $(X_K \text{ UI } X_I | X_D)$ implies that the more general conditional utility function

$$U(x_K | x_I, x_D) \triangleq \frac{U(x_K, x_I, x_D) - U(x_K^0, x_I, x_D)}{U(x_K^*, x_I, x_D) - U(x_K^0, x_I, x_D)} \quad (51)$$

does not depend on the instantiation of x_I , where x_K^0 and x_K^* are the instantiations of the attributes in X_K at their minimum and maximum values, respectively. It will therefore be convenient to express the independence relation $(X_K \text{ UI } X_I | X_D)$ by setting x_I arbitrarily to x_I^0 in the normalized conditional utility function to get the following “equivalent definition of utility independence”:

$$(X_K \text{ UI } X_I | X_D) \Leftrightarrow U(x_k | x_I, x_D) = U(x_k | x_I^0, x_D), \quad \forall x_k, x_I, x_D. \quad (52)$$

We can now envision a wide variety of possible utility independence assertions among the attributes. To facilitate the representation of these independence relations, we provide a graphical representation of utility independence below.

Definition 1 (Bidirectional Utility Diagrams) (Abbas [3]). A bidirectional utility diagram as a graph whose nodes represent the attributes X_1, X_2, \dots, X_n , and whose arrows represent the *possibility of* utility dependence between them. The absence of an arrow between two nodes X_i and X_j asserts the utility independence relations $(X_i \text{ UI } X_j | \bar{X}_{ij})$ and $(X_j \text{ UI } X_i | \bar{X}_{ij})$; a unidirectional arrow from X_i to X_j , written as $X_i \rightarrow X_j$, asserts the relation $(X_i \text{ UI } X_j | \bar{X}_{ij})$; and a bidirectional arrow between X_i and X_j , written as $X_i \leftrightarrow X_j$, does not assert any independence relations.

The arrow notation used here is similar to that of influence diagrams where the absence of an arrow asserts probability independence and an arrow merely allows for the *possibility of* dependence. If it turns out later during the actual assessment of the utility functions that the two attributes are utility independent, then the functional form would still be consistent with the independence assertions already provided.

Figure 9 shows some bidirectional diagrams and utility dependence matrices for three attributes.

Proposition 1. *The following are equivalent:*

- (i) $(X_K \text{ UI } X_I | X_D)$ and
- (ii) the set of pairwise utility independence relations $(X_K \text{ UI } X_j | \bar{X}_{jK}), \forall X_j \in X_I$.

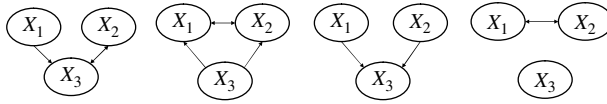
Proposition 1 shows that we can represent the independence relation $(X_K \text{ UI } X_I | X_D)$ with a set of pairwise relations between the set X_K and each of the individual attributes in X_I . This result forms the basic premise of the graph-based representations that we discuss. But how do we derive the functional form of bidirectional diagrams? To derive this functional form, we first write the set of constraints (functional equations) that it needs to satisfy. For example, Figure 9(d) implies the system of equations

$$U(x_1, x_2, x_3) = l_1(x_2, x_3) + d_1(x_2, x_3)h_1(x_1, x_2), \quad (53)$$

$$U(x_1, x_2, x_3) = l_2(x_1, x_3) + d_2(x_1, x_3)h_2(x_1, x_2), \quad (54)$$

$$U(x_1, x_2, x_3) = l_3(x_1, x_2) + d_3(x_1, x_2)h_3(x_3). \quad (55)$$

FIGURE 9. Examples of three-attribute bidirectional utility diagrams (Abbas [3]).



An alternate way to derive the functional form is to use the basic expansion theorem. Let X_K be the set of nodes that provide at least one independence assertion in a diagram, let X_{iD} be the set of nodes that have arrows into X_i , and let X_{iI} be the set of nodes with no arrows into X_i . Any utility function that is consistent with the independence assertions of the diagram must satisfy

$$\forall X_i \in X_K, (X_i \text{ UI } X_{iI} | X_{iD}) \text{ in the diagram} \\
 \Rightarrow U(x_i | x_{iD}, x_{iI}) = U(x_i | x_{iD}, x_{iI}^0), \quad \forall x_i, x_{iD}, x_{iI}, i \in K \text{ in the functional form.} \quad (56)$$

Further divide X_{iD} into X_{iPD}, X_{iFD} , where X_{iPD} is the set of attributes in X_{iD} expanded prior to X_i , and X_{iFD} is the set of attributes in X_{iD} to be expanded after X_i . Similarly, divide X_{iI} into X_{iPI}, X_{iFI} . Therefore, $\bar{X}_i = X_{iI} \cup X_{iD}$, $X_{iI} = X_{iPI} \cup X_{iFI}$, $X_{iD} = X_{iPD} \cup X_{iFD}$. Substituting into the basic expansion theorem gives

$$U(x_1, \dots, x_n) = \sum_{x_K^0 \in X_K^0} U(x_K^0, \bar{x}_K) \prod_{i \in K} g(x_i | x_{iPD}^0, x_{iFD}, x_{iPI}^0, x_{iFI}). \quad (57)$$

Any utility independence relation in the diagram must be satisfied on the whole domain of the attributes, and therefore, it must be satisfied at any specific reference values, i.e., it is necessary that

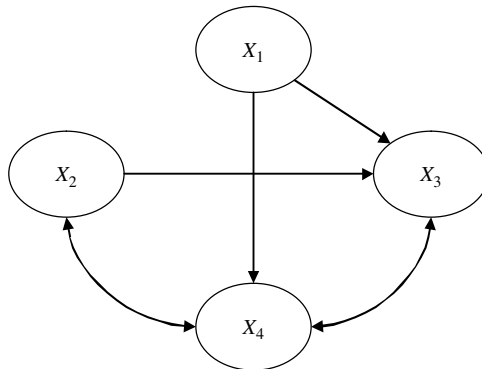
$$(X_i \text{ UI } X_{iI} | X_{iD}) \Rightarrow U(x_i | x_{iPD}^0, x_{iFD}, x_{iPI}^0, x_{iFI}) = U(x_i | x_{iPD}^0, x_{iFD}, x_{iI}^0), \\
 \forall x_i, x_{iFD}, x_{iFI}. \quad (58)$$

This necessary condition allows us to replace every term $U(x_i | x_{iPD}^0, x_{iFD}, x_{iFI}^0, x_{iFI})$ in the right-hand side of (57) with $U(x_i | x_{iPD}^0, x_{iFD}, x_{iI}^0)$ to get the decomposed form

$$U(x_1, \dots, x_n) = \sum_{x_K^0 \in X_K^0} U(x_K^0, \bar{x}_K) \prod_{i \in K} g(x_i | x_{iPD}^0, x_{iFD}, x_{iI}^0). \quad (59)$$

Example 4. Consider the bidirectional utility diagram of Figure 10. In this figure, only two rows provide independence equations, namely,

FIGURE 10. Example of a bidirectional utility diagram for four attributes.



Node X_1 : Asserts that $U(x_1 | x_2, x_3, x_4) = U(x_1 | x_2^0, x_3^0, x_4^0)$.

Node X_2 : Asserts that $U(x_2 | x_1, x_3, x_4) = U(x_2 | x_1^0, x_3^0, x_4^0)$.

To derive the functional form, we expand the utility function around attributes X_1, X_2 and substitute for the independence relations to get the decomposed form

$$\begin{aligned}
 U(x_1, x_2, x_3, x_4) &= U(x_1^*, x_2^*, x_3, x_4)U(x_1 | x_2^0, x_3^0, x_4^0)U(x_2 | x_1^0, x_3^0, x_4^0) \\
 &\quad + U(x_1^*, x_2^0, x_3, x_4)U(x_1 | x_2^0, x_3^0, x_4^0)\bar{U}(x_2 | x_1^0, x_3^0, x_4^0) \\
 &\quad + U(x_1^0, x_2^*, x_3, x_4)\bar{U}(x_1 | x_2^0, x_3^0, x_4^0)U(x_2 | x_1^0, x_3^0, x_4^0) \\
 &\quad + U(x_1^0, x_2^0, x_3, x_4)\bar{U}(x_1 | x_2^0, x_3^0, x_4^0)\bar{U}(x_2 | x_1^0, x_3^0, x_4^0).
 \end{aligned} \tag{60}$$

The functional form (60) requires six different assessments: $U(x_1 | x_2^0, x_3^0, x_4^0)$, $U(x_2 | x_1^0, x_3^0, x_4^0)$, $U(x_1^*, x_2^*, x_3, x_4)$, $U(x_1^*, x_2^0, x_3, x_4)$, $U(x_1^0, x_2^*, x_3, x_4)$, and $U(x_1^0, x_2^0, x_3, x_4)$. To verify that (60) satisfies the independence equations, we calculate the normalized conditional utility functions it implies. By direct substitution, $U(x_1 | x_2, x_3, x_4) = U(x_1 | x_2^0, x_3^0, x_4^0)$ and $U(x_2 | x_1, x_3, x_4) = U(x_2 | x_1^0, x_3^0, x_4^0)$. The conditional utility functions are thus equal to the lower order conditional assessments provided by the decision maker at the boundaries, namely, $U(x_1 | x_2^0, x_3^0, x_4^0)$, $U(x_2 | x_1^0, x_3^0, x_4^0)$. There are no constraints on the actual functional form of the conditional utility function assessments required in this example, and independence will always be satisfied. This result is always true for *canonical diagrams*, which we define below.

Let X_K be the set of nodes having at least one independence assertion and $|X_K|$ be the cardinality of the set X_K .

Definition 2. A bidirectional diagram with $|X_K| = 0$ or $|X_K| = 1$ is canonical. A diagram with $|X_K| \geq 2$ is canonical if and only if there are no unidirectional or bidirectional arrows connecting any two different nodes of X_K with each other.

Example 5 (Reduction to the Canonical Multilinear Form). If every attribute is utility independent of its complement (the utility function has a multilinear form), then the set X_K contains all the attributes, and there are no arrows in the entire diagram. This multilinear representation is therefore canonical. Moreover, because every node provides an independence equation, the terms $U(x_K, \bar{x}_K)$ go through a full expansion of the n attributes and become the constants $U(x_n^{*0})$. Substituting into (59) gives

$$U(x_1, \dots, x_n) = \sum_{x_n^{*0} \in X_n^{*0}} U(x_n^{*0}) \prod_{X_i \in X_n} g(x_i | \bar{x}_i^0), \tag{61}$$

which is another (compact) way of expressing the multilinear form.

If a diagram is not canonical, we still use (59) to derive the functional form, but now we need an extra step to verify that the assessed conditional utility functions match the utility independence assertions of the diagram when substituted into (59).

Extensions to higher order independence relations where a subset of attributes is utility independent of another subset are also possible. Abbas [3] shows that such higher order independence relations yield functional forms that are special cases of the more general functional form implied by a bidirectional diagram because they require additional constraints on the boundary values. For example, the functional form implied by the independence assertion $(X_1 X_2 UI X_3)$ is a special case of the functional form implied by the two independence relations $(X_1 UI X_3 | X_2)$ and $(X_2 UI X_3 | X_1)$. As a further illustration, the condition of mutual utility independence (which implies that every subset of the attributes is utility independent of its complement) is a special case of the more general multilinear form (which implies that every attribute is utility independent of its complement) with additional constraints on the normalizing constants.

But how do we assess the portion of a bidirectional diagram that contains bidirectional arrows? This assessment could still be cognitively difficult. For example, the assessments

$U(x_2 \mid x_1^0, x_3^0, x_4)$, $U(x_1^*, x_2^*, x_3, x_4)$, $U(x_1^*, x_2^0, x_3, x_4)$, $U(x_1^0, x_2^*, x_3, x_4)$, and $U(x_1^0, x_2^0, x_3, x_4)$ of the previous example each require a two-dimensional utility assessment over the domain of two attributes. The next three sections present tools to conduct these types of assessments. One approach uses a *utility copula* that constructs these assessments using single-attribute functions at the boundary values of the domain. Another approach uses a condition known as “one-switch utility independence,” and a third approach uses *functional equations* to derive their functional form.

8. Multiattribute Utility Copulas (Abbas [5])

The decompositions that result from the basic expansion theorem and the independence relations of the bidirectional diagram may still involve assessments for nodes connected with bidirectional arrows. These joint utility assessments may be difficult to assess. This section presents a method to capture this type of utility dependence in terms of single-attribute utility functions. The idea is to assess single-attribute utility functions for each attribute at a boundary value (or an instantiation of the complement attributes) and construct a utility surface that matches these single-attribute assessments. Abbas [2] explores this approach defining utility copulas for attribute dominance utility functions, and Abbas [5] extends this idea to more general utility functions. This type of analysis does for utility functions what copula structures do for joint probability functions (Sklar [33]).

Definition 3 (Multiattribute Utility Copula). A multiattribute utility copula, $C_\lambda(v_1, \dots, v_n)$, is a multivariate function of n variables that satisfies the following conditions:

(1) *Normalized Range and Domain.* The function, C_λ , is a continuous mapping from the n -dimensional hypercube $[0, 1]^n$ to the interval $[0, 1]$ and is normalized such that

$$C_\lambda(0, \dots, 0) = 0 \quad \text{and} \quad C_\lambda(1, \dots, 1) = 1. \tag{62}$$

(2) *Positive linear transformation at reference values.* For each argument v_i , there exist some reference values, $\lambda_{i,j}, i \neq j$, for which the utility copula satisfies

$$C_\lambda(\lambda_{i,1}, \dots, \lambda_{i,i-1}, v_i, \lambda_{i,i+1}, \dots, \lambda_{i,n}) = a_i v_i + b_i, \quad i = 1, \dots, n, \tag{63}$$

where $0 \leq \lambda_{i,j} \leq 1$, $0 < a_i \leq 1$, and $0 \leq b_i < 1$.

Proposition 2. Any multiattribute utility function that is (i) continuous, (ii) bounded, and (iii) strictly increasing with each argument for at least one reference value of the complement attributes can be expressed in terms of normalized conditional utility functions, $U_i(x_i \mid \bar{x}_i^{\lambda_i})$, $i = 1, \dots, n$, and a multiattribute utility copula, C_λ , as

$$U(x_1, \dots, x_n) = C_\lambda(U_1(x_1 \mid \bar{x}_1^{\lambda_1}), \dots, U_n(x_n \mid \bar{x}_n^{\lambda_n})). \tag{64}$$

With the regularity conditions on the utility function, we are guaranteed the existence of at least one instantiation $\bar{x}_i^{\lambda_i}$ for which $U_i(x_i \mid \bar{x}_i^{\lambda_i})$ is strictly increasing. If the utility function is strictly increasing with each argument across the entire domain, then we can define a conditional utility function at any instantiation of the complement attributes. We can also define 2^{n-1} conditional utility functions for each attribute, X_i , at all possible combinations of the boundary (minimum and maximum) values of its complement. In particular, we define the maximum marginal utility function as

$$U_i(x_i \mid \bar{x}_i^*) = \frac{U(x_i, \bar{x}_i^*) - U(x_i^0, \bar{x}_i^*)}{1 - U(x_i^0, \bar{x}_i^*)}, \quad U(x_i^0, \bar{x}_i^*) \neq 1, \tag{65}$$

and the minimum marginal utility function as

$$U_i(x_i \mid \bar{x}_i^0) = \frac{U(x_i, \bar{x}_i^0)}{U(x_i^*, \bar{x}_i^0)}, \quad U(x_i^*, \bar{x}_i^0) \neq 0. \tag{66}$$

INFORMS holds copyright to this article and distributed this copy as a courtesy to the author(s). Additional information, including rights and permission policies, is available at http://journals.informs.org/.

Abbas [5] introduces two classes of utility copulas: class 1 utility copulas that satisfy the condition

$$C_1(1, \dots, 1, v_i, 1, \dots, 1) = a_i v_i + b_i, \quad i = 1, \dots, n, \tag{67}$$

and match the normalized conditional utility functions at the maximum boundaries of the complement, as well as class 0 utility copulas that satisfy

$$C_1(0, \dots, 0, v_i, 0, \dots, 0) = a_i v_i + b_i, \quad i = 1, \dots, n, \tag{68}$$

and match the minimum marginal utility functions.

Our focus in this chapter will be utility functions that are nondecreasing with each of their arguments. An example of such utility copulas that match the utility functions at the maximum boundaries of the domain of the attributes is the extended Archimedean form, which we discuss in more detail below.

Definition 4 (Archimedean Functional Form). We define the extended Archimedean functional form, $E(v_1, \dots, v_n)$, as

$$E(v_1, \dots, v_n) = a\psi^{-1} \left[\prod_{i=1}^n \psi(l_i + (1 - l_i)v_i) \right] + b, \tag{69}$$

where $a = 1/(1 - \psi^{-1}[\prod_{i=1}^n \psi(l_i)])$; $b = 1 - a$; the function ψ satisfies the conditions that: (i) $\psi(v)$ is continuous on the domain $v \in [0, 1]$, (ii) $\psi(v)$ is strictly increasing on the domain $v \in [0, 1]$, and (iii) $\psi(0) = 0$ and $\psi(1) = 1$; and the parameters $l_i, i = 1, \dots, n$ must be greater than or equal to zero and less than 1.

The function ψ has the same mathematical properties as either a strictly increasing cumulative probability distribution or a strictly increasing normalized utility function on the domain $[0, 1]$. With this observation, all well-known functional forms of cumulative distributions, such as beta distributions, can be used as the function ψ in (69). The extended Archimedean utility copula matches the single-attribute conditional utility functions at the maximum value of the complement attributes. The following Corollary illustrates how to determine the parameters l_i .

Corollary 1 (Assessing the parameters of the Extended Archimedean Utility Copula). *The parameters $l_i, i = 1, \dots, n$ of the multiattribute utility function obtained using (69) satisfy*

$$a(1 - l_i) = 1 - U(x_i^0, \bar{x}_i^*), \quad i = 1, \dots, n. \tag{70}$$

It is straightforward to see that if at least one attribute, X_j , is a utility dominant attribute ($U(x_j^0, \bar{x}_j) = 0, \forall \bar{x}_j$ or equivalently, $l_j = 0$), then $a = 1$ and $l_i = U(x_i^0, \bar{x}_i^*)$ for all attributes. This situation provides a simple mechanism for determining the parameters of the functional form. On the other hand, if none of the attributes are utility dominant, then $l_i \neq 0, i = 1, \dots, n, a > 1$, and therefore, $U(x_i^0, \bar{x}_i^*) < l_i$. Because l_i is bounded from above by one, we will not be able to model all values of $U(x_i^0, \bar{x}_i^*)$ as it gets closer to unity because of the strict inequality. However, from a cognitive view, we probably would not want to use the maximum margin utility functions for large values of $U(x_i^0, \bar{x}_i^*)$, because the difference between the utility values at the maximum and minimum values of X_i at this instantiation of the complement would be too small. In this case, as we shall see, a class 0 utility copula would be simpler to reason about. For two attributes, extended Archimedean utility copulas span the boundary ranges of $U(x_1^0, x_2^*) + U(x_1^*, x_2^0) < 1$, and for more than two attributes they span the range $\sum_{i=1}^n U(x_i^0, \bar{x}_i^*) < n - 1$, which shows that this constraint becomes less restrictive with larger numbers of attributes. As shown in Abbas [5], we can also model more flexible utility surfaces using linear and composite combinations of utility copulas.

The following example presents an application of extended Archimedean utility copulas.

Example 6 (Constructing an Extended Archimedean Utility Copula). Suppose that a decision maker faces two attributes, X and Y , defined on the normalized domain $[0, 1] \times [0, 1]$. He states that his conditional utility functions $U(x | y^*)$ and $U(y | x^*)$ are exponential with risk aversion coefficients $\gamma_x = 3$ and $\gamma_y = 2$, respectively. He also provides the utility values $U(x^0, y^*) = 0.4$ and $U(x^*, y^0) = 0.2$ using indifference probability assessments. For example, the utility value $U(x^0, y^*) = 0.4$ asserts that he is indifferent between receiving (x^0, y^*) for certain or receiving a binary deal that provides (x^*, y^*) with probability 0.4 and (x^0, y^0) with probability 0.6. We now consider the following proper class 1 generating function:

$$\psi(v) = \frac{1 - e^{-\delta v}}{1 - e^{-\delta}}, \quad 0 \leq v \leq 1, \delta \in \mathbb{R} \setminus \{0\}. \tag{71}$$

Substituting from (71) into (69) gives the class 1 utility copula

$$C_\lambda(v_x, v_y) = a \frac{1}{\delta} \ln \left(1 - \frac{(1 - e^{-\delta(l_x + (1-l_x)v_x)})(1 - e^{-\delta(l_y + (1-l_y)v_y)})}{(1 - e^{-\delta})} \right) + b. \tag{72}$$

By definition, $a = 1/(1 + 1/\delta) \ln(1 - (1 - e^{-\delta l_x})(1 - e^{-\delta l_y})/(1 - e^{-\delta}))$, and from Corollary 1, the parameters, l_x and l_y must satisfy

$$a(1 - l_x) = 0.6, \quad a(1 - l_y) = 0.8. \tag{73}$$

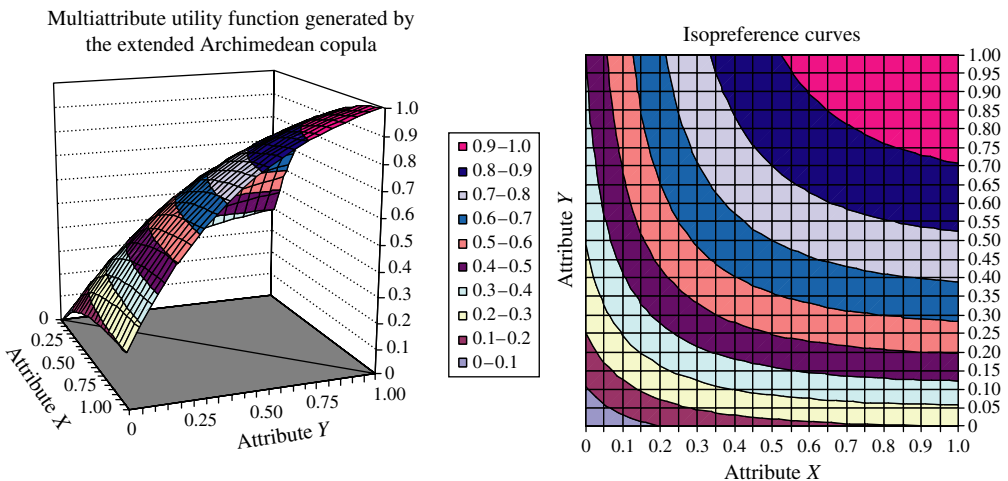
If we use $\delta = 1$ in the generating function, then the solution to Equation (73) gives $l_x = 0.54$, $l_y = 0.39$, with $a = 1.32$ and $b = -0.32$. The multiattribute utility function and the isopreference curves obtained by this Archimedean functional form are shown in Figure 11.

Changing the value of the parameter δ yields different multiattribute utility surfaces having the same normalized conditional utility functions at the maximum margins but provides different trade-off functions among the attributes. The value of δ can be determined by changing its value to match the decision maker's trade-offs. Alternatively, we can assess several points on the surface of the utility function and substitute into (72) to derive the value of δ .

Note that the utility function obtained with (69) using a generating function $\psi(v) = v$, $0 \leq v \leq 1$ is the multiplicative form of mutual utility independence,

$$U(x_1, \dots, x_n) = \frac{\prod_{i=1}^n (l_i + (1 - l_i)U_i(x_i | \bar{x}_i)) - \prod_{i=1}^n l_i}{1 - \prod_{i=1}^n l_i}, \tag{74}$$

FIGURE 11. Multiattribute utility function generated by the extended Archimedean copula.



where the normalized conditional utility functions, $U_i(x_i | \bar{x}_i), i = 1, \dots, n$, are assessed at any instantiations of the complement attributes. For the case of two attributes, we can expand this form as

$$\begin{aligned}
 IU(x_1, x_2) &= \frac{(l_1 + (1 - l_1)U_1(x_1 | \bar{x}_1))(l_2 + (1 - l_2)U_2(x_2 | \bar{x}_2)) - l_1 l_2}{1 - l_1 l_2} \\
 &= a_1 U_1(x_1 | \bar{x}_1) + a_2 U_2(x_2 | \bar{x}_2) + (1 - a_1 - a_2)U_1(x_1 | \bar{x}_1)U_2(x_2 | \bar{x}_2), \quad (75)
 \end{aligned}$$

from which we can see immediately that with $0 \leq l_i < 1$, we cannot construct utility surfaces having mutual utility independence and negative values of $(1 - a_1 - a_2)$ using the extended functional form. This does not mean that the Archimedean form cannot model utility functions with a negative mixed partial derivative $\partial^2 U(x_1, x_2) / \partial x_1 \partial x_2$. In fact, we can obtain a negative or positive mixed partial derivative by appropriate choice of the generating function. Abbas [5] shows that if the generating function is log concave, then the two-attribute mixed partial derivative is negative, and if it is log convex, then it is positive. Moreover, linear combinations of utility copulas can also add more general utility surfaces with varying mixed partial derivatives.

The result of reducing the Archimedean form to the multiplicative form also enables us to conduct *sensitivity to utility independence* for a given copula. Using a sensitivity analysis, Abbas [5] shows that about 25% of the decisions we make could be incorrect (75% correct decisions) if we use a multiplicative form of utility independence when in fact the attributes have utility dependence using the extended Archimedean form. "Compare this to a coin toss that would give 50% correct decisions."

Abbas [5] also defines the scaled Archimedean functional form as a class 0 utility copula,

$$S(v_1, \dots, v_n) = a\eta^{-1}\left(\prod_{i=1}^n \eta(m_i v_i)\right), \quad 0 \leq v_i \leq 1, \quad (76)$$

where $0 < m_i \leq 1$, $a = 1/(\eta^{-1}(\prod_{i=1}^n \eta(m_i)))$, and the function $\eta(v)$ is (i) continuous and (ii) strictly decreasing on the domain $v \in [0, 1]$ with (iii) $\eta(0) = 1$ and $\eta(1) = 0$. Abbas [5] shows that the scaled Archimedean functional form (76) is a class 0 utility copula with the parameters m_i satisfying

$$am_i = U(x_i^*, \bar{x}_i^0), \quad i = 1, \dots, n. \quad (77)$$

If $a = 1$, which implies that there exists at least one attribute with $m_i = 1 = U(x_i^*, \bar{x}_i^0)$, then $m_i = U(x_i^*, \bar{x}_i^0), i = 1, \dots, n$. On the other hand, if $a > 1$, then $m_i < U(x_i^*, \bar{x}_i^0)$. Because m_i is bounded from below, this implies that there is a lower bound on the boundary values $U(x_i^*, \bar{x}_i^0)$ that can be modeled with a scaled Archimedean form. In fact, this lower bound is $\sum_{i=1}^n U(x_i^*, \bar{x}_i^0) > 1$, a constraint that is less restrictive as the number of attributes increases. Moreover, when $\sum_{i=1}^n U(x_i^*, \bar{x}_i^0) < 1$, we probably would not want to use a class 0 utility copulas, because it would be easier (cognitively) to use a class 1 utility copula and match the maximum marginal values instead.

Although we have limited the discussion of utility copulas to utility functions that are strictly increasing with each of their arguments at a reference value of the complement attributes, we can extend this condition to utility functions that are strictly monotonic with each attribute at a reference value of the complement. With this extension, we also generalize the marginal condition, (63), of a utility copula to be $C_\lambda(\lambda_{i,1}, \dots, \lambda_{i,i-1}, v_i, \lambda_{i,i+1}, \dots, \lambda_{i,n}) = a_i v_i + b_i, i = 1, \dots, n$, where $0 \leq \lambda_{i,j} \leq 1$ and $a_i \neq 0$. Archimedean forms can also model these new utility copulas. For example, $\psi^{-1}[\psi(u)\psi(1 - v)]$ is a hybrid utility copula (class 1 for u and class 0 for v), and is in fact decreasing with the argument v . Other directions for future work on utility copulas also include deriving copula functions that match more than one boundary utility assessment, as well as exploring the interpretation of notions of tail dependence and microcorrelations of probability copulas into utility copulas.

9. One-Switch Utility Independence (Abbas and Bell [9])

We now present another method that may be useful in capturing the utility dependence between two nodes connected with bidirectional arrows, if its conditions are satisfied. First, let us refer back to the implications of utility independence between two attributes: *preferences for lotteries over one attribute will not change as we change the value of the other*. Now consider what happens when preferences do change but change only once. Consider a decision with two attributes that may be relevant in the decision to purchase a home, *wealth* and *home quality*. It seems plausible that our preferences for lotteries over quality of the home will change as we get richer. Indeed, there are a lot of uncertainties in the quality of the new home at the purchase stage. We may like house A better than house B but settle for B because of financial constraints. But it is quite likely that our preferences may change as we become wealthier, but will change only once no matter how wealthy we get. If preferences for lotteries over attribute X can change, but can change only once, as we vary another attribute, Y , we say that X exhibits one-switch independence of Y . This idea extends the notion of one-switch utility functions introduced by Bell [17], where preferences between monetary lotteries change only once as our wealth state changes, to multiattribute utility functions. Here preferences between lotteries over X change only once as we vary another attribute (or parameter), Y . The following proposition characterizes the one-switch independent utility functions.

Proposition 3 (Abbas and Bell [9]). *X is one-switch independent of Y if and only if*

$$u(x, y) = g_0(y) + g_1(y)[f_1(x) + f_2(x)\phi(y)], \quad (78)$$

where

$$g_0(y) = u(x^0, y), g_1(y) = [u(x^*, y) - u(x^0, y)], f_1(x) = u(x | y^0), f_2(x) = u(x | y^*) - u(x | y^0),$$

and

$$\phi(y) \triangleq \frac{u(x | y) - u(x | y^0)}{u(x | y^*) - u(x | y^0)} \quad (79)$$

is a monotonic function of y and independent of x .

Note that this condition captures utility dependence between two attributes and provides a nice interpretation for the resulting functional form. Of course we can envision many extensions of this result, such as mutual one-switch independence, or extensions to multiple attributes where every attribute is one-switch independent of its complement or every subset of the attributes is one-switch independent of its complement. Further generalizations could also include deriving one-switch value functions for ordinal preferences.

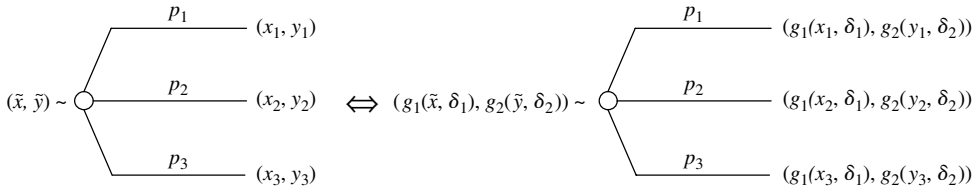
10. Using Functional Equations and Invariance Arguments

Functional equations provide an excellent method for deriving the functional form of a utility function if we can characterize our preferences in such forms (see, for example, Pfanzagl [30], Aczél [13], Abbas [4, 7], Abbas and Aczél [8]). This section gives some examples of constructing multiattribute utility functions using such functional equations.

To illustrate, suppose you face a lottery whose prospects are characterized by multiple attributes, such as an uncertain cash flow over n successive years (for simplicity, let us just consider two years). Of course the certainty equivalent of this situation is in fact a contour of points having constant value (or constant expected utility),

$$\{(x, y): U(x, y) = \text{expected utility}\}.$$

FIGURE 12. An example of two-attribute transformation invariance.



Now suppose that each of the lottery outcomes is modified by a multivariate transformation of the form $(x, y) \rightarrow (g_1(x, \delta_1), g_2(y, \delta_2))$, where g_1, g_2 are monotonic transformations, and δ_1, δ_2 are transformation parameters. This formulation is very general. For example, the transformations can represent shift transformations on the lottery outcomes or scale transformations representing a fraction of the lottery at each year.

Suppose further that the certain equivalent contour of this new lottery is a mapping that results when any point on the certain equivalent contour of the previous lottery is modified by the same transformation $(x, y) \rightarrow (g_1(x, \delta_1), g_2(y, \delta_2))$, and that this applies for any lottery. What does this tell us about the multiattribute utility function? Abbas [7] defines this situation as multiattribute transformation invariance, an example of which is shown in Figure 12.

To illustrate how this formulation derives the functional form, we first define a multiattribute utility function, $U(x_1, x_2, \dots, x_n)$, as invariant to the transformation $(g_1(x_1, \delta_1), \dots, g_n(x_n, \delta_n))$ over sets $\Delta_i, \delta_i \in \Delta_i$ if

$$U(g_1(x_1, \delta_1), \dots, g_n(x_n, \delta_n)) = k(\delta_1, \dots, \delta_n)U(x_1, x_2, \dots, x_n) + d(\delta_1, \dots, \delta_n), \quad \forall \delta_i \in \Delta_i, \quad (80)$$

for some functions $k(\delta_1, \dots, \delta_n)$ and $d(\delta_1, \dots, \delta_n)$.

The following proposition from Abbas [7] derives the functional form of a utility function satisfying multiattribute transformation invariance.

Proposition 4. *A decision maker with a continuous multiattribute utility function, $U(x_1, x_2, \dots, x_n)$, that is strictly monotonic with each of its arguments satisfies multiattribute transformation invariance with the multivariate transformation $(g_1(x_1, \delta_1), \dots, g_n(x_n, \delta_n))$ for all $\delta_i \in \Delta_i$ if and only if he has an invariant multiattribute utility function to this given transformation for all $\delta_i \in \Delta_i$.*

Example 7 (Invariance to Multivariate Shift Transformations). Suppose that a decision maker faces a multivariate lottery over any n attributes such as (i) profit earned in years 1 through n , (ii) a portfolio of n stocks, where each stock represents an attribute, or (iii) n attributes in the design of an automobile such as mileage, fuel efficiency, or maximum odometer speed. Suppose he is indifferent between receiving this uncertain lottery or receiving $(\tilde{x}_1, \dots, \tilde{x}_n)$ for certain. If the outcomes of the lottery are modified by shift amounts $g_i(x_i, \delta_i) = x_i + \delta_i, i = 1, \dots, n$, and if the decision maker is indifferent between receiving this modified lottery and receiving $(\tilde{x}_1 + \delta_1, \dots, \tilde{x}_n + \delta_n)$ for certain, then his multiattribute utility function must satisfy the functional equation

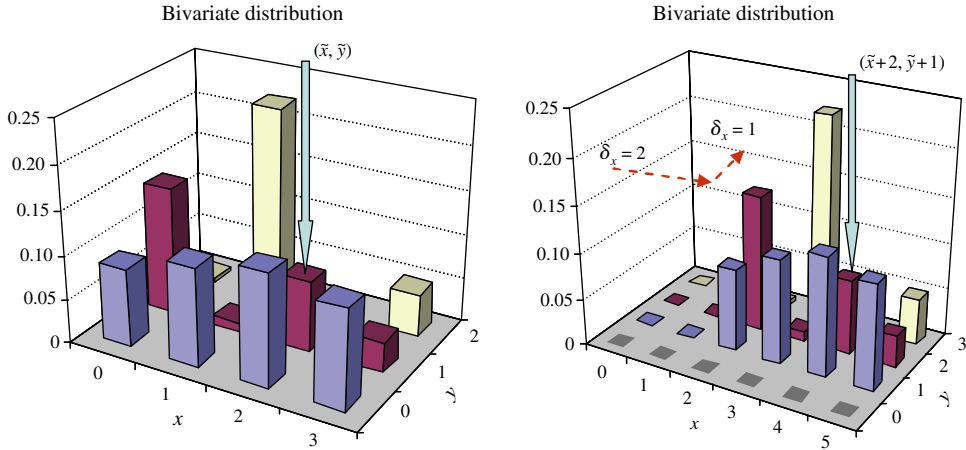
$$U(x_1 + \delta_1, \dots, x_n + \delta_n) = k(\delta_1, \dots, \delta_n)U(x_1, x_2, \dots, x_n) + d(\delta_1, \dots, \delta_n), \quad \forall \delta_i \in \Delta. \quad (81)$$

An example of a bivariate lottery that is shifted by $\delta_x = 2, \delta_y = 1$ is shown in Figure 13.

We have now converted knowledge of multiattribute shift invariance into the problem of solving the multivariate functional Equation (80), which is a generalized Pexider equation, whose general continuous solutions for shift amounts $\delta_1, \dots, \delta_n$ on intervals of positive length are (i) the additive utility function

$$U(x_1, \dots, x_n) = \sum_{j=1}^n a_j x_j + B, \quad k = 1, \quad d(\delta_1, \dots, \delta_n) = \sum_{j=1}^n a_j \delta_j \quad (82)$$

FIGURE 13. Shift invariance with a bivariate lottery.



Note. The certain equivalent prospect (\bar{x}, \bar{y}) follows the same shift transformation; its relative position with respect to the lottery remains the same.

and (ii) the multiplicative utility function

$$U(x_1, \dots, x_n) = D \prod_{i=1}^n e^{-a_i x_i} + C, \quad k(\delta_1, \dots, \delta_n) = \prod_{i=1}^n e^{-a_i \delta_i},$$

$$d(\delta_1, \dots, \delta_n) = C \left(1 - \prod_{i=1}^n e^{-a_i \delta_i} \right). \tag{83}$$

There may be situations in practice where we can assert transformation invariance with other multivariate transformations (such as scale transformations). This information can also be used to derive the functional form of the utility function by building on the results of shift invariance. To illustrate, consider the class of transformations, $g_i(x_i, \delta_i)$, that can be converted into a shift transformation with a monotonic transformation of variables, $g_{\text{shift}}(x)$, such that

$$g_{\text{shift}}(g(x, \delta)) = g_{\text{shift}}(x) + \beta(\delta), \tag{84}$$

and $\beta(\delta)$ is a function of δ but not of x . This class applies to a wide variety of transformations. For example, the scale transformation, $g(x, \delta) = \delta x$, belongs to this class because we can take the logarithm of both sides to get $\ln(g(x, \delta)) = \ln(\delta) + \ln(x)$, with $g_{\text{shift}}(x) = \beta(x) = \ln(x)$. The power transformation, $g(x, \delta) = x^\delta$, also belongs to this class as it can be converted into a shift transformation by taking the double logarithm of both sides to get $\ln(\ln(g(x, \delta))) = \ln(\delta) + \ln(\ln(x))$. In this case, $g_{\text{shift}}(x) = \ln(\ln(x))$ and $\beta(\delta) = \ln(\delta)$. With this change of variables, a multiattribute utility function that satisfies transformation invariance with a transformation $g_i(x_i, \delta_i)$, $i = 1, \dots, n$ satisfying (84), for x_i, δ_i on an interval of positive length, has either an additive functional form,

$$U(x_1, \dots, x_n) = \sum_{i=1}^n a_i g_{i\text{-shift}}(x_i) + a_0, \tag{85}$$

or a multiplicative functional form,

$$U(x_1, \dots, x_n) = a \prod_{i=1}^n e^{a_i g_{i\text{-shift}}(x_i)} + b, \tag{86}$$

where $a, a_i, i = 1, \dots, n$ are nonzero, but otherwise $a, a_i, i = 1, \dots, n$, and a_0, b are arbitrary constants.

The functional forms (85) and (86) define a large family of multiattribute utility functions that satisfy mutual utility independence. For example, we can see that a decision maker who satisfies multiattribute scale invariance for all values $\delta_1, \dots, \delta_n$ over intervals of positive length must have either (i) an additive utility function (when $k(\delta_1, \dots, \delta_n) = 1$),

$$U(x_1, x_2, \dots, x_n) = \sum_{i=1}^n a_i \ln(x_i) + a_0, \tag{87}$$

or (ii) a multiplicative utility function (when $k(\delta_1, \dots, \delta_n) = \prod_{i=1}^n \delta_i^{a_i}$),

$$U(x_1, x_2, \dots, x_n) = a \prod_{i=1}^n x_i^{a_i} + b. \tag{88}$$

In a similar manner, we can derive many other functional forms of multiattribute utility functions satisfying (85) and (86) using their invariant transformation. For example, the multiattribute utility function that is invariant to scale transformations $g_i(x_i, \delta_i) = \delta_i x_i$ for attributes x_1, \dots, x_k , shift transformations $g_i(x_i, \delta_i) = \delta_i + x_i$ for attributes x_{k+1}, \dots, x_l , and power transformations for attributes x_l, \dots, x_n for all $\delta_i, i = 1, \dots, n$ is either (i) a multiplicative utility function,

$$U(x_1, x_2, \dots, x_n) = a \prod_{i=1}^k x_i^{a_i} \prod_{i=k+1}^l e^{a_i x_i} \prod_{i=l+1}^n (\ln x_i)^{a_i} + b, \tag{89}$$

or (ii) an additive utility function,

$$U(x_1, x_2, \dots, x_n) = \sum_{i=1}^k a_i \ln(x_i) + \sum_{i=k+1}^l a_i x_i + \sum_{i=l+1}^n a_i \ln(\ln x_i) + a_0. \tag{90}$$

Multiattribute transformation invariance thus provides an interpretation for a wide variety of functional forms of multiattribute utility functions. Whereas the examples we discussed so far lead to assertions of mutual utility independence among the attributes, this notion also extends to situations that capture utility dependence among the attributes. The following example illustrates the generality of this approach.

Example 8 (Invariance with Identical Values of the Transformation Parameters). Suppose that the decision maker satisfies multiattribute transformation invariance when tested with values of the transformation parameters that are (i) identical for each attribute, i.e., $\delta_1 = \delta_2 = \dots = \delta_n = \Delta$, and (ii) where Δ is satisfied on an interval of positive length. From the previous results, the multiattribute utility function must satisfy the functional equation

$$U(x_1 + \Delta, x_2 + \Delta, \dots, x_n + \Delta) = k(\Delta)U(x_1, x_2, \dots, x_n) + d(\Delta). \tag{91}$$

Abbas et al. [12] derive the continuous solutions of (91) as either

$$U(x_1, \dots, x_n) = f(x_2 - x_1, \dots, x_n - x_1) + Ax_1, \quad k(z) = 1, \quad d(z) = Az, \tag{92}$$

or

$$U(x_1, \dots, x_n) = e^{Cx_1} f(x_2 - x_1, \dots, x_n - x_1) + E, \quad k(z) = e^{Cz}, \quad d(z) = E(1 - e^{Cz}), \tag{93}$$

where f is an arbitrary continuous nonconstant function, and where $C \neq 0, E, A$ are arbitrary constants.

We have again derived the functional form of a multiattribute utility function by knowledge of transformation invariance. Equations (92) and (93) show that invariance with identical shift amounts on all the attributes does indeed provide some functional decomposition of the utility function, but allows for more general functional forms than those that are invariant for all values of $\delta_1, \dots, \delta_n$. Here utility independence need not be enforced. In fact, this formulation also allows for partial utility independence conditions. To illustrate, for two attributes, we have

$$U(x, y) = \begin{cases} e^{Cy} f(x - y) + E, \\ f(x - y) + Ay. \end{cases} \quad (94)$$

If $f(t) = a_1 e^{(a+b)t} + a_3 e^{at} + a_2$, $C = (a + b)$, and $E = a_4$ in the multiplicative solution, we get

$$U(x, y) = a_1 e^{(a+b)x} + a_2 e^{(a+b)y} + a_3 e^{(ax+by)} + a_4, \quad k \neq 1. \quad (95)$$

The functional form (95) does not assert mutual utility independence or constant absolute risk aversion unless either (i) $U(x, y) = a_1 e^{(a+b)x} + a_2 e^{(a+b)y} + a_4$, when $a_3 = 0$, or (ii) $U(x, y) = a_3 e^{(ax+by)} + a_4$, when $a_1, a_2 = 0$. A third case of mutual utility independence also occurs when $f(t) = a_1 t$ in the additive solution $f(x - y) + Ay$ to get

$$U(x, y) = ax + by + c. \quad (96)$$

Invariance with identical values of transformation parameters can also lead to solutions with conditional utility independence. For example, when $a_1 = 0$, in (95), $U(x, y) = a_2 e^{(a+b)y} + a_3 e^{(ax+by)} + a_4$, where attribute X is conditionally utility independent of attribute Y . From the previous results, we observe that we need to test for transformation invariance over independent sets of the shift parameters for all the attributes before asserting some forms of utility independence.

In some follow-up work, Abbas and Chudziak [?] derive the functional forms of multiattribute utility functions that lead to a maximum of one-switch change in preferences between uncertain cash flows as the initial wealth increases in the form of constant annuity payments.

10.1. Verifying the Functional Form of the Utility Function with Invariance Arguments

We have shown that invariance arguments can be used to derive the functional form of a utility function. We now show how they can also be used to verify the functional form of the multiattribute utility function. For example, suppose that a decision maker states that he satisfies the conditions of mutual utility independence and he provides individual utility functions for each of the attributes. Invariance arguments can be used to verify that mutual utility independence is indeed satisfied.

Abbas [7] shows that any multiattribute utility function with mutual utility independence must satisfy transformation invariance with a multivariate transformation of the form $(x_1, x_2, \dots, x_n) \rightarrow (g_1(x_1), g_2(x_2), \dots, g_n(x_n))$. An additive utility function $U(x_1, \dots, x_n) = a \sum_{i=1}^n \mu_i U_i(x_i) + b$ must satisfy transformation invariance with

$$g_i(x_i, \delta_i) = U_i^{-1}(cU_i(x) + \delta_i), \quad c \neq 0, \quad (97)$$

and a product utility function $U(x_1, \dots, x_n) = a \prod_{i=1}^n U_i(x_i) + b$ must satisfy invariance with

$$g_i(x_i, \delta_i) = U_i^{-1}(\delta_i U_i(x_i)), \quad \delta_i \neq 0, \quad (98)$$

where $a \neq 0$ and $cU_i(x) + \delta_i$ and $\delta_i U_i(x_i)$ lie on the range of the utility function, but otherwise $a, b, c, \delta_i, i = 1, \dots, n$ are arbitrary.

These results apply directly to the more general multiplicative form $U(x_1, \dots, x_n) = a \prod_{i=1}^n [1 + \mu\mu_i U_i(x_i)] + b$, because we can define $U_i^M(x_i) = [1 + \mu\mu_i U_i(x_i)]$ to get the product form $U(x_1, \dots, x_n) = a \prod_{i=1}^n U_i^M(x_i) + b$, and use (98) to derive its invariant transformation as

$$g_i(x_i, \delta_i) = U_i^{-1} \left(\frac{\delta_i(1 + \mu\mu_i U_i(x_i)) - 1}{\mu\mu_i} \right), \quad \delta_i, \mu_i, \mu \neq 0. \quad (99)$$

These results provide consistency checks that must be verified when assuming the widely used forms of utility independence. We can also envision future work that derives transformation invariance relations for other multivariate transformations of the form $(x, y) \rightarrow (g_1(x, y), g_2(x, y))$ that derive the functional form of more general utility functions.

11. Additional Benefits of the Value Function

My purpose in this section is to show how using the value function can simplify many other features that relate to multiattribute decision problems. We have already seen that using a value function reduces the whole multiattribute utility assessment into a one-dimensional assessment. Therefore, we can simplify the problem without having to think about utility independence or dependence among the attributes.

But in addition, when reasoning about the individual attributes instead of the single attribute of value, we can sometimes arrive at inconsistent formulations. The notion of a *multiattribute target* is one such example that illustrates how thinking in terms of individual attributes may lead to inconsistencies in decision making.

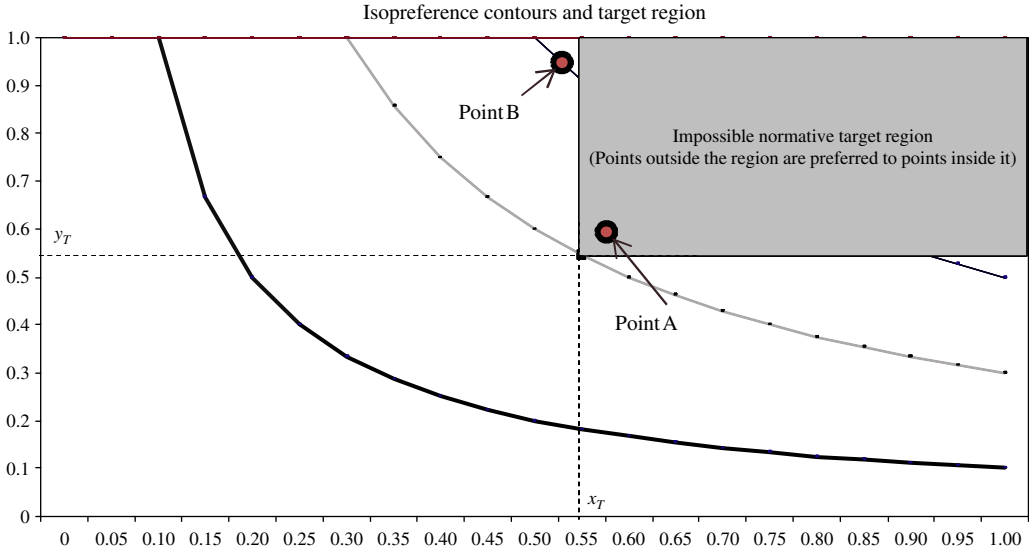
To illustrate, consider two attributes, X and Y , where we prefer more of each attribute to less. Figure 14 shows the contours of the value function of a firm having these two attributes. Now suppose that the firm also sets an incentive structure (requirement) to exceed levels x_T and y_T of attributes X and Y , respectively. This requirement forms the shaded region in Figure 14 and is quite popular in practice with managerial settings such as management by objectives and balanced score cards. A simple look at the contours of the value function in relation to the shaded region is sufficient to highlight the inconsistency with this type of target-setting behavior.

We can see from Figure 14 that this incentive structure (although popular in practice and in some of the academic literature) will reward performance point A versus point B (because point A satisfies the target and point B does not), even though point B has a higher value for the organization as determined by the value function. *This target violates the value consistency requirements* of the firm and is induced by thinking about the individual attributes instead of the value function.

The notion of bivariate (or even *multivariate*) *stochastic dominance* is another example where the value function can significantly reduce the analysis of the problem. Multivariate stochastic dominance derives conditions on the multiattribute utility function based on properties of joint distributions, extending the notion of stochastic dominance from the univariate case to the multiattribute case.

We observe here that *the multivariate stochastic dominance problem is also a univariate problem because the value function can reduce the multivariate lottery into a one-dimensional lottery over the value measure*. This works for both discrete and continuous lotteries. To illustrate, consider a bivariate lottery $\langle p_1, (x_1, y_1), \dots, p_n, (x_n, y_n) \rangle$. If we define a value function, $V(x, y)$, over the attributes then this lottery becomes the univariate lottery $\langle p_1, v_1, \dots, p_n, v_n \rangle$ over value. Properties of first order or any higher order stochastic dominance that were derived for the univariate case will work just as well for utility functions over value (and for the multiple attributes). In fact they will even work for situations where the utility function

FIGURE 14. Value inconsistency when trying to exceed multiple thresholds.



is not necessarily increasing with each of its arguments, such as the peanut butter and jelly sandwich.

The same type of one-dimensional analysis also applies to the continuous case. Abbas and Matheson [11] introduce the notion of a *value-based cumulative probability* function, $F(V_0)$, as the probability that the outcome of a joint probability distribution (or multivariate lottery) has a value less than or equal to V_0 as determined by the value function. For an n -variable probability density function, $f(x_1, x_2, \dots, x_n)$, define the value-based cumulative probability as

$$F(V_0) \triangleq \iint_D f(x_1, x_2, \dots, x_n) dx_1 dx_2, \dots, dx_n,$$

where $D = \{(x_1, x_2, \dots, x_n): V(x_1, x_2, \dots, x_n) \leq V_0\}$.

The multivariate lottery problem is now reduced to a single dimension of value. In comparing two lotteries, we can compare their value-based cumulative probability distributions and derive dominance conditions on the single-attribute utility functions over value. It would be quite difficult to think about this notion of multivariate stochastic dominance for multiple attributes when monotonicity is not present without the use of a value function. In fact, the value function presents the most general way of reasoning about this extension particularly for utility functions that are not necessarily monotonic with their arguments.

Another notion that might be of interest when constructing the multiattribute utility function is that of multivariate risk aversion (Richard [32]), which may often be considered a desirable property. A decision maker with a utility function, $U(x, y)$, facing prospects $x_0 < x_1, y_0 < y_1$, is said to be multivariate risk averse if prefers a lottery that provides a 50–50 chance of (x_0, y_1) and (x_1, y_0) to another that provides a 50–50 chance of (x_0, y_0) and (x_1, y_1) . This is equivalent to saying that

$$U(x_1, y_1) - U(x_0, y_1) - U(x_1, y_0) + U(x_0, y_0) < 0.$$

When the utility function is twice continuously differentiable, this condition also implies that the mixed partial derivative, $\partial^2 U(x, y) / \partial x \partial y$, is negative. This property is interesting to think about and may indeed be desirable in many situations (such as consumption levels (or calorie intake) in two successive years if (x_0, y_0) falls below survival level), but it need not apply to all. The peanut butter and jelly sandwich is a simple example where it need

INFORMS holds copyright to this article and distributed this copy as a courtesy to the author(s). Additional information, including rights and permission policies, is available at http://journals.informs.org/.

not be (and is not) satisfied. Moreover, this concept becomes difficult to reason about when more than two attributes are present. If we assign a value function and a utility function over value, then we do not need to think about this notion of multivariate risk aversion among the attributes while constructing the utility function. However, it may be useful as a consistency check after constructing the final functional form.

Another (equivalent) notion is that of *correlation aversion*. If a decision maker prefers a lottery providing a 50–50 chance of (x_0, y_1) and (x_1, y_0) to another that provides a 50–50 chance of (x_0, y_0) and (x_1, y_1) , then he prefers lotteries where the variables are negatively correlated to those that are positively correlated. As we shall see, we can also reduce this problem into a one-dimensional problem using a value function. First we illustrate this notion of correlation aversion in more detail, using the following example from Abbas [7]. This example relates correlation aversion to the simple notion of *variance aversion*.

A decision maker faces a bivariate Gaussian lottery over two attributes, X and Y , such that

$$f(x, y) = \frac{1}{2\pi\sigma_1\sigma_2\sqrt{1-\rho^2}} e^{-z/(2(1-\rho^2))}, \tag{100}$$

where $z = (x - \mu_1)^2/\sigma_1^2 + 2\rho(x - \mu_1)(y - \mu_2)/(\sigma_1\sigma_2) + (y - \mu_2)^2/\sigma_2^2$, $\rho = \text{corr}(x, y) = \sigma_{12}/(\sigma_1\sigma_2)$.

Suppose that the decision maker has an additive value function (which as we have seen need not apply to all situations but may be suitable for situations like constant discounting), and suppose that the utility function over value is exponential, i.e.,

$$V(x, y) = a_1x + a_2y \quad \text{and} \quad U_V(V) = -e^{-\gamma V}.$$

His expected utility is equal to the two-dimensional characteristic function of the bivariate Gaussian, which is

$$EU = L(a_1, a_2) = e^{-[a_1\mu_1 + a_2\mu_2 - (1/2)(\sigma_1^2 a_1^2 + 2\sigma_1\sigma_2\rho a_1 a_2 + \sigma_2^2 a_2^2)]}. \tag{101}$$

Because the utility function over value is exponential, then the value certain equivalent for this lottery is

$$\tilde{V} = \frac{-1}{\gamma} \ln(L(a_1, a_2)) = \frac{1}{\gamma} \left[a_1\mu_1 + a_2\mu_2 - \frac{1}{2}(\sigma_1^2 a_1^2 + 2\sigma_1\sigma_2\rho a_1 a_2 + \sigma_2^2 a_2^2) \right]. \tag{102}$$

Equation (102) is an exact expression for the certain equivalent of a bivariate Gaussian lottery for a decision maker who has a bivariate exponential utility function. This result shows that the value certain equivalent decreases when the correlation coefficient ρ is positive (but it is simply a result of saying that the variance increases when the correlation coefficient is positive and it decreases when it is negative, and that the certain equivalent decreases as the variance increases). To illustrate, consider the special case where $a_1 = a_2 = \gamma$ (such as equal risk aversion coefficients for the payoffs of two different stocks at the same time period); we have

$$\tilde{V} = \mu_1 + \mu_2 - \frac{\gamma}{2}(\sigma_1^2 + 2\sigma_1\sigma_2\rho + \sigma_2^2) = \mu_T - \frac{\gamma}{2}\sigma_T^2, \tag{103}$$

where μ_T and σ_T^2 are the mean and variance of the sum of the two variables, $X + Y$, a one-dimensional problem over a single attribute, money.

The notion of correlation aversion (which stems from variance aversion) makes perfect sense in a stock portfolio problem with Gaussian structure, where the dependence is captured entirely by the correlation coefficients, where we prefer less variance to more (assuming we are risk averse over the monetary prospects considered), and with all else held constant.

It is also straightforward to extend this example to a multivariate Gaussian distribution where the certain equivalent will be equal to the mean less half the risk aversion coefficient multiplied by the variance of the portfolio. Because less variance is preferred to more in the case of risk aversion toward money, we can derive the effects that the pairwise correlation coefficients will have on the overall variance and the certain equivalent of the portfolio.

But would you want to have this correlation aversion property for any attributes and any decision you face? If you were risk neutral over money in the previous example, $\gamma = 0$, then you would not be correlation averse, because you would not care about the variance in the first place. When the attributes become more general, such as one of the attributes representing the “left shoe” and the other representing the “right shoe,” most people would prefer a gamble providing a “pair of shoes” or nothing to a lottery that provides a single shoe (left or right). Correlation aversion would be difficult to justify in this example. In general, this property of correlation aversion will depend on the value function and the utility function over value as well as the types of attributes in the problem. As the value function becomes more complex (losing monotonicity) or if it is super additive, it will be more difficult to reason about why you would want such a property to be satisfied within the utility function. Thinking in terms of the value function and a one-dimensional utility function over value, you would not need to consider this property while constructing the utility function. Similar to the notion of multivariate risk aversion, its purpose could be simply to test this property in the final functional form.

12. Conclusions and the Road Ahead

Let us recap the main points. (i) Maximizing arbitrary measures such as value at risk or minimizing maximum regret may violate the choice obtained using the axioms of expected utility. If you use arbitrary measures, then you should realize this fact and you should think about that particular axiom of utility that you would like to violate in your decision making. (ii) If there is no uncertainty in the decision problem or if there is first order stochastic dominance, then we do not need to assess a utility function. A value function is sufficient. (iii) When uncertainty is present, the axioms of decision analysis tell us that we need to maximize the expected von Neumann–Morgenstern utility in our choice among decision alternatives. (iv) When the prospects have multiple direct value attributes, we assign a multiattribute utility function by a variety of methods that can in general be classified into two fundamental approaches: The first uses a value function and then assigns a one-dimensional utility function over value. The second derives the multiattribute utility function directly without the use of the value function.

In situations where the utility function is not necessarily increasing with each of the attributes (such as the peanut butter and jelly sandwich example), the construction of a utility function by direct utility assessments over the individual attributes may be much more complicated compared to the value function approach.

Utility independence conditions can significantly simplify the construction of the multiattribute utility function by direct utility assessments over the attributes (both cognitively and computationally) and can reduce the number and order of assessments needed in the basic expansion theorem. But if those conditions are not present in a given situation, then we should not enforce them. If we cannot formulate the attributes to yield utility independence, then we need to make a “declaration of utility dependence,” and incorporate this property into our formulation. Attribute dominance conditions, when present, can also reduce the number of assessments needed for the basic expansion theorem, and enable the use of a variety of tools for updating preference in a simple way such as Bayes’ rule for utility inference.

Bidirectional utility diagrams provide a convenient method for graphically capturing the partial utility independence relations, as well as the asymmetric nature of utility independence. When attribute dominance conditions exist, they reduce to the utility diagrams with

symmetric independence properties. The nodes in the bidirectional utility diagram that are connected with bidirectional arrows can be elicited in a variety of ways including a utility copula, a one-switch independence formulation, or a functional equation. We also showed how functional equations themselves are powerful tools for deriving the functional form of the whole utility function.

Which method should we use when constructing the utility function? The ultimate test is having a utility function that captures the decision maker's preferences and an approach that the decision maker is comfortable providing responses to. When formulating decision problems, the analyst can start by defining the direct value attributes to construct a value function. We have seen that when this task is achieved, the problem reduces to a single-attribute utility assessment over the value measure or any of the attributes considered. It is also particularly useful when the function is not increasing with each of the attributes. If, the decision maker is not comfortable providing a value function, then we can formulate the attributes to yield the multilinear form of utility independence or any of its special cases (such as the additive or multiplicative forms). If this formulation cannot be done, then he can formulate the attributes to yield bidirectional diagrams with a canonical form. We showed how this formulation enables a tractable functional form that provides a natural extension to the multilinear form and allows for flexibility in the utility assessments provided by the decision maker.

A hybrid combination of methods can also be used as a consistency check. For example, if we construct the utility function using the value function approach and observe that one of the attributes is one-switch independent of its complement, then we have a constraint on the utility function assessment over value or on the value function itself. Alternatively, if we know that the utility function satisfies transformation invariance with a given transformation, then we also have a consistency check and a condition that must be satisfied in the final functional form.

My interest in utility theory began during my PhD work at Stanford University (Abbas [1, 2]) and emerged from analogies between machinery developed for joint probability functions and my desire to build similar machinery for multiattribute (and single-attribute) utility functions. This area, I believe, is still very raw, and I am convinced that the future of multiattribute utility theory will bring a variety of new tools, invariance formulations, switching conditions, diagrams, curve-fitting methods, new utility copula formulations, and many other normative techniques that better capture preferences. I also believe that the value function approach for constructing multiattribute utility functions has had much less literature coverage than the direct utility assessment approach, and I see the need for more work on deriving more representative value functions that better capture the trade-offs among the different attributes and reduce the construction of the multiattribute utility function into a one-dimensional utility assessment.

Acknowledgments

This work was partially supported by National Science Foundation Career Award NSF-DRMS 0846417.

References

- [1] A. E. Abbas. An entropy approach for utility assignment in decision analysis. C. J. Williams, ed. *Proceedings of the 22nd International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering*. American Institute of Physics Conference Proceedings, Vol. 659. American Institute of Physics, College Park, MD, 659:328–338, 2002.
- [2] A. E. Abbas. *Entropy Methods in Decision Analysis*. Ph.D. thesis, Stanford University, Palo Alto, CA, 2004.
- [3] A. E. Abbas. Bidirectional utility diagrams. Technical Report UILU-ENG-2005-3001, University of Illinois at Urbana-Champaign, Urbana, 2005.

- [4] A. E. Abbas. Invariant utility functions and certain equivalent transformations. *Decision Analysis* 4(1):17–31, 2007.
- [5] A. E. Abbas. Multiattribute utility copulas. *Operations Research* 57(6):1367–1383, 2009.
- [6] A. E. Abbas. From Bayes' nets to utility nets. *Proceedings of the 29th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering..* American Institute of Physics Conference Proceedings, Vol. 1193, American Institute of Physics, College Park, MD, 3–12, 2009.
- [7] A. E. Abbas. Invariant multiattribute utility functions. *Theory and Decision* 68(1):69–99, 2010.
- [8] A. E. Abbas and J. Aczél. The role of some functional equations in decision analysis. *Decision Analysis* 7(2):215–228, 2010.
- [9] A. E. Abbas and D. E. Bell. One-switch independence for multiattribute utility functions. *Operations Research*. Forthcoming.
- [10] A. E. Abbas and R. A. Howard. Attribute dominance utility. *Decision Analysis* 2(4):185–206, 2005.
- [11] A. E. Abbas and J. E. Matheson. Normative decision making with multiattribute performance targets. *Journal of Multi-Criteria Decision Analysis* 16(3–4):67–78, 2009.
- [12] A. E. Abbas, J. Aczél, and J. Chudziak. Invariance of multiattribute utility functions under shift transformations. *Results in Mathematics* 54(1–2):1–13, 2009.
- [13] J. Aczél. *Lectures on Functional Equations and Their Applications*. Academic Press, New York, 1966.
- [14] K. J. Arrow. The theory of risk aversion. Lecture 2 in *Aspects of the Theory of Risk-Bearing*. Yrjö Jahnssonin Saatio, Helsinki, Finland, 1965.
- [15] M. Baucells and R. K. Sarin. Evaluating time streams of income: Discounting what? *Theory and Decision* 63(2):95–120, 2007.
- [16] D. E. Bell. Multiattribute utility functions: Decompositions using interpolation. *Management Science* 25(8):744–753, 1979.
- [17] D. E. Bell. One-switch utility functions and a measure of risk. *Management Science* 34(12):1416–1424, 1988.
- [18] J. Dyer and R. Sarin. Measurable value functions. *Operations Research* 27(4):810–822, 1979.
- [19] J. S. Dyer and R. K. Sarin. Relative risk aversion. *Management Science* 28(8):875–886, 1982.
- [20] P. H. Farquhar. A fractional hypercube decomposition theorem for multiattribute utility functions. *Operations Research* 23(5):941–967, 1975.
- [21] P. C. Fishburn. von Neumann–Morgenstern utility functions on two attributes. *Operations Research* 22(1):35–45, 1974.
- [22] P. C. Fishburn. Approximations of two-attribute utility functions. *Mathematics of Operations Research* 2(1):30–44, 1977.
- [23] R. A. Howard. On fates comparable to death. *Management Science* 30(4):407–422, 1984.
- [24] R. A. Howard. In praise of the old time religion. W. Edwards, ed. *Utility Theories: Measurements and Applications*. Kluwer, Boston, 27–55, 1992.
- [25] R. Keeney. Analysis of preference dependencies among objectives. *Operations Research* 29(6):1105–1120, 1981.
- [26] R. L. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Trade-Offs*. Wiley, New York, 1976.
- [27] C. Kirkwood. Parametrically dependent preferences for multiattributed consequences. *Operations Research* 24(1):92–103, 1976.
- [28] J. E. Matheson and A. E. Abbas. Utility transversality: A value-based approach. *Journal of Multi-Criteria Decision Analysis* 13(5–6):229–238, 2005.
- [29] J. E. Matheson and R. A. Howard. 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, 1968.
- [30] J. Pfanzagl. A general theory of measurement applications to utility. *Naval Research Logistics Quarterly* 6(4):283–294, 1959.

- [31] J. W. Pratt. Risk aversion in the small and in the large. *Econometrica* 32(1–2):122–136, 1964.
- [32] S. Richard. Multivariate risk aversion, utility independence and separable utility functions. *Management Science* 22(1):12–21, 1975.
- [33] A. Sklar. Fonctions de répartition à n dimensions et leurs marges. *Publications de l'Institut de Statistique de L'Université de Paris* 8:229–231, 1959.
- [34] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ, 1947.